

# Chapter 8

## Neutrino Detectors



Leslie Camilleri

After a brief introduction describing the many sources of neutrinos, this article will describe the various detector techniques that are being used to observe neutrinos of energies ranging from a few MeV to hundred's of GeV.

### 8.1 Historical Introduction

In 1930 in order to explain the continuous energy spectrum of electrons emitted in beta decay, Pauli postulated [1] that these electrons were emitted together with a light neutral particle. This particle was subsequently named the neutrino. Their actual observation had to wait until 1953 when Reines and Cowan recorded [2] interactions of anti(electron)neutrinos emitted by a reactor in a cadmium doped liquid scintillator detector. Since then, in addition to the  $\nu_e$ , two other flavours of neutrinos were observed, the  $\nu_\mu$  and  $\nu_\tau$ . The  $\nu_\mu$ , which is produced in  $\pi \rightarrow \mu$  decay, was proved to be different [3] from the  $\nu_e$  in an experiment at Brookhaven using thick-plate optical spark chambers. The  $\nu_\tau$ , companion of the  $\tau$  lepton, was observed [4] at Fermilab in an emulsion cloud chamber detector consisting of iron plates interleaved with sheets of photographic emulsions. Although until recently neutrinos were thought to be massless and were described as such in the Standard Model, in the past decade they have been found to be massive [5, 6]. Furthermore each of the three flavour states mentioned above consists of a superposition of three mass states of unequal masses leading to oscillations of one flavour into another under the appropriate conditions. The characteristics of these oscillations depend on three mixing angles  $\theta_{13}$ ,  $\theta_{12}$  and  $\theta_{23}$  as well as on the difference of the square of the 3

---

L. Camilleri (✉)

Nevis Labs, Columbia University, Irvington-on-Hudson, NY, USA

e-mail: [camil@nevis.columbia.edu](mailto:camil@nevis.columbia.edu)

© The Author(s) 2020

C. W. Fabjan, H. Schopper (eds.), *Particle Physics Reference Library*,

[https://doi.org/10.1007/978-3-030-35318-6\\_8](https://doi.org/10.1007/978-3-030-35318-6_8)

337

neutrino masses,  $\Delta m_{12}^2$  referred to as the solar mass difference as it is of importance in oscillations of solar neutrinos,  $\Delta m_{13}^2 \sim \Delta m_{23}^2$  referred to as the atmospheric mass difference as it drives oscillations of neutrinos produced in the atmosphere through the decay of mesons produced in cosmic ray interactions. The flavour of interacting neutrinos can only be determined if the interaction is via a charged current. In these, the  $\nu_e$ ,  $\nu_\mu$  and  $\nu_\tau$  respectively produce a negative electron, muon or  $\tau$  lepton in the final state. Antineutrinos produce the corresponding positive charged lepton.

## 8.2 Sources of Neutrinos and Their Characteristics

Naturally occurring neutrinos and man-made neutrinos are produced through several different processes. Nature provides us with solar neutrinos emitted by the sun, atmospheric neutrinos produced by the interaction of cosmic rays in the atmosphere, cosmological neutrinos produced by a variety of deep space violent events, geological neutrinos produced by nuclear decays in the earth core as well as neutrinos produced in beta decay. Man made neutrinos are produced by nuclear reactors or by specially designed beams at accelerators or by highly radioactive sources. These processes are briefly described below.

### 8.2.1 Solar Neutrinos

They are emitted in nuclear reactions occurring in the sun [7]. The three main reactions are  $p + p \rightarrow d + e^+ + \nu_e$ , emitting a continuous spectrum of neutrinos with an end point at 0.4 MeV,  $e + {}^7\text{Be} \rightarrow {}^7\text{Li} + \nu_e$  with a monochromatic spectrum at 0.862 MeV and  ${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e$  also with a continuous spectrum with an end point at 15 MeV. Their total flux on earth is  $6.4 \times 10^{10} \text{ cm}^{-2}\text{s}^{-1}$ .

### 8.2.2 Atmospheric Neutrinos

Atmospheric neutrinos [8] are produced in the decays of  $\pi$  and K mesons produced in the interactions of cosmic rays in the upper atmosphere. Their energy ranges over several orders of magnitude up to hundreds of GeV. They are observed either coming from above or from below and in the latter case they will have traversed the earth. This allows us to observe them from a few kilometers to about 13,000 km from their production point, thus providing us with very different baselines over which to study oscillations. These predominantly  $\nu_e$  and  $\nu_\mu$  neutrinos are usually observed through their charged current interactions respectively producing electrons or muons.

### 8.2.3 Cosmological Neutrinos

The study of cosmological neutrinos [9] at the TeV scale is in its infancy. Their very low rate necessitates extremely large detectors. This has led to the use of naturally occurring detection media such as lake or sea water and Antarctic ice. The Cerenkov light or radio waves emitted by charged particles produced in their interactions in the medium are recorded, respectively, in strings of photomultiplier tubes or antennas.

### 8.2.4 Reactor Neutrinos

Nuclear reactors are an abundant source of antineutrinos, 6  $\bar{\nu}_e$  per nuclear fission on average, resulting in a flux of  $1.8 \times 10^{20}$  per GW thermal energy, emitted isotropically. The standard method to study them [10] is to observe the Inverse Beta Decay (IBD) reaction  $\bar{\nu}_e + p \rightarrow e^+ + n$  in a hydrogen-rich liquid scintillator detector. In addition to observing photons emitted as a result of the positron annihilation, the neutron can be detected by recording photons emitted by the neutron capture in the scintillator.

### 8.2.5 Accelerator Neutrinos

Accelerator neutrinos are produced [11] by the decay of  $\pi$  and K mesons themselves produced by the interaction of a proton beam on a target as illustrated in Fig. 8.1. The target must be thick enough along the beam to maximize the proton interaction probability and yet thin enough to minimize the reinteraction probability and multiple scattering of the produced mesons such as to produce as high an energy and as focussed a beam as possible. The usual target geometry consists of a series of thin rods of low Z material such as carbon or beryllium separated by a few cms but in line with the proton beam. The mesons are then focussed by a system of toroidal magnets. These, referred to as horns [12], consist of two concentric current sheets, parabolically shaped that provide a toroidal magnetic field. Its strength is inversally proportional to the radial displacement from the beam axis and the integral is such as to bend more the particles that are further away from the beam thus providing a near parallel beam. A second horn is usually provided such as to compensate for





Fig. 8.1 The principle of an accelerator produced neutrino beam




**Fig. 8.2** The principle of horn focusing of mesons in a neutrino beam

over-focused and under-focused particles as shown in Fig. 8.2. The particles then enter a long evacuated decay tunnel in which  $\pi \rightarrow \mu \nu_\mu$ ,  $K \rightarrow \mu \nu_\mu$  and  $K \rightarrow \pi e \nu_e$  decays occur producing a predominantly  $\nu_\mu$  beam with an admixture of  $\sim 1\%$  of  $\nu_e$ . Focussing positive mesons produces a neutrino beam whereas focussing negative mesons, achieved by a reversal of the polarity of the horns, produces an antineutrino beam.

An alternative to the horns is a system of bending magnets and quadrupoles. Such a technique [13] has been used in the Sign Selected Quadrupole Train, SSQT, neutrino beam at Fermilab. Its performance is described in [14]. An advantage of this technique is that the neutrino beam not being along the axis of the proton beam,  $\nu_e$  from  $K_L^0$  decays will not enter the detector since their parents will not be deflected. This is a distinct advantage in oscillation experiments looking for  $\nu_e$  appearance in a  $\nu_\mu$  beam in which the intrinsic  $\nu_e$  background is irreducible.

The above techniques produce a beam with a broad energy spectrum, referred to as a broad band beam. A narrower range of neutrino energies is sometimes desirable. Such narrow band beams are obtained by first momentum-selecting the parent pions and kaons before they decay using standard beam optics methods, thus reducing the range of neutrino energies. Furthermore, the neutrino energy can be deduced on an event by event basis as it is related to the neutrino production angle and this can be computed from the radial position of the event within the detector. The uncertainty on the energy depends on the momentum and angular spread of the meson beam and on the length of the decay channel. It is typically 5–20%. The intensity of these narrow band beams is necessarily lower than that of broad band beams.

Another way to expose the detector to neutrinos with a given narrow energy spectrum is to place the detector at an off-axis angle to the beam [15]. The kinematics of pion decay, shown in Fig. 8.3 are such that neutrinos observed at a non-zero angle to the proton beam have an approximately unique momentum irrespective of the momentum of their parent meson. Furthermore the value of this unique momentum depends on the off-axis angle, thus allowing a detector to be exposed to the neutrino momentum required by the physics under investigation by placing it at the appropriate angle.



**Fig. 8.3** The correlation between the pion momentum and its decay neutrino momentum, plotted for several neutrino directions relative to the proton beam axis

An accelerator neutrino beam can also, potentially, contain  $\nu_\tau$ 's. These are produced through the production of  $D_s$  mesons in the initial proton interactions and their subsequent decays  $D_s \rightarrow \tau \nu_\tau$  followed by  $\tau \rightarrow \nu_\tau + \dots$ . However in most accelerator beams the  $\nu_\tau$  content is negligible since the  $D_s$  production cross section is small at existing energies. One notable exception will be discussed in Sect. 8.3.4.

The semi-leptonic decays of charmed particles have also been used to produce neutrinos. In this case, because of the very short lifetime of charm, a decay tunnel is unnecessary. The beam is produced in a so-called beam dump [16], in which the incident proton beam and secondary pions and kaons are absorbed before they can decay. At CERN, the beam dump [17] was made of copper disks that could be separated thus altering its density between  $3$  and  $9 \text{ g} \cdot \text{cm}^{-3}$ . The normal neutrino flux from  $\pi$  and K decays was reduced by about 3 orders of magnitude. Since charm is produced in pairs in proton interactions, the beam contains an approximately equal amount of neutrinos and antineutrinos, the only difference being due to  $\pi$  and K mesons decays occurring before these mesons are absorbed. Furthermore, because of the equal  $e\nu_e$  and  $\mu\nu_\mu$  decay probabilities of charm an equal number of  $\nu_e$  and  $\nu_\mu$  are present in the beam. In this context, it is interesting to note [18] that hadronic colliders, and in particular LHC, produce a large quantity of charm and beauty particles in the forward directions, resulting in two well collimated neutrino beams emerging from each interaction region.

## 8.3 Detection Techniques


Because of the small interaction cross section of neutrinos, neutrino detectors must be massive. The exception is detectors addressing coherent neutrino interactions for which the cross section is orders of magnitude larger than for other neutrino interactions and which will be addressed in Sect. 8.3.1. The nature of these massive detectors depends on the physics being addressed. It usually involves observing the resulting hadronic part of an interaction and, if a charged current interaction, the observation of a charged lepton. If the physics merely requires the measurement of the total neutrino energy, a calorimetric detector suffices. If individual particles must be measured, then a more sophisticated tracking device is needed. The measurement of a final state muon is usually accomplished in a fairly straightforward way with magnetized iron because of the muon penetrating nature. A final state electron is more difficult to measure, especially its charge, because of bremsstrahlung and showering as it propagates through material. Neutrino detectors must then necessarily be of several types. Techniques must be suitable to detect neutrinos of energies ranging from a few MeV to about a PeV. They must be fine-grained enough to measure electrons, identify individual particles and observe secondary vertices of  $\tau$ 's or charmed particles or heavy enough to produce large number of interactions using calorimetric techniques. It is evident that neutrino detectors use most of the detecting techniques used in particle physics. They will be outlined in the following sections.

### 8.3.1 *Totally Active Scintillator Detectors*

Scintillator detectors can either use liquid or solid scintillator. If liquid is used the detector consists of either a single large tank or of tubes filled with liquid. Solid scintillator detectors usually consist of strips. The first neutrino detector [2] used by Reines and Cowan was intended to observe the interaction of reactor antineutrinos of a few MeV. The observation was made, as in subsequent reactor experiments, using the IBD reaction  $\bar{\nu}_e + p \rightarrow e^+ + n$  and using a detector consisting of liquid scintillator viewed by photomultipliers. In addition to observing light emitted from the positron annihilation, the neutron can be detected by also observing photons emitted by the neutron capture in the hydrogen of the scintillator. In order to enhance the neutron capture cross section, they added cadmium to the scintillator. They observed an excess of events when the reactor was in operation leading to the first detection of (anti)neutrino interactions and a subsequent Nobel prize. This technique is still being applied to this day [10], albeit with some refinements. Several recent experiments which will be described below, used it to search for  $\bar{\nu}_e$  oscillations to another flavour in the domain of the atmospheric  $\Delta m^2$ ,  $2.5 \times 10^{-3} \text{ eV}^2$ . Because of the low energy of reactor  $\bar{\nu}_e$ 's,  $\bar{\nu}_\mu$ 's or  $\bar{\nu}_\tau$ 's that they potentially oscillate to cannot be observed through their charged current

interactions since it is energetically impossible to produce  $\mu$ 's or  $\tau$ 's. Oscillations can then only be observed through the disappearance technique resulting in a reduction and distortion of the expected  $\bar{\nu}_e$  spectrum. Given the energy of reactor  $\bar{\nu}_e$ 's (a few MeV) and the value of the atmospheric  $\Delta m^2$ , CHOOZ [19] was located 1000 m from a reactor complex in order to be near oscillation maximum. It used a single large tank of liquid scintillator and was subjected to a cosmic muon rate of  $0.4 \text{ m}^{-2} \text{ s}^{-1}$ . One of the major backgrounds in this type of experiment is the background generated by cosmic ray muons. The first line of defense is to place the detector underground, at a depth of 300 m water equivalent (m.w.e) in the case of CHOOZ. A muon traversing the detector does not, in itself, simulate a signal event because the large amount of energy deposited can be well identified. However neutrons produced by muons traversing dead areas of the detector or the surrounding rock can elastically scatter on a proton, causing the proton and the subsequent neutron capture to simulate the signature of a reactor event. This background can be eliminated by vetoing on the passage of a nearby muon. In addition cosmic muons can produce long lived isotopes such as  $^6\text{He}$  and  $^9\text{Li}$  which subsequently can beta decay producing an electron and a neutron, thus simulating an antineutrino event. This background cannot be eliminated by vetoing on the passage of a muon because of the long lifetime of these decays (178 ms in the case of  $^9\text{Li}$ ) which would introduce an inordinate dead time. It must be estimated and subtracted. Palo Verde [20] was located at a shallower depth of 32 m.w.e. and chose to use acrylic cells filled with liquid scintillator. This extra segmentation was needed to reduce the larger muon induced neutron background caused by the larger cosmic muon flux of  $22 \text{ m}^{-2} \text{ s}^{-1}$  at this depth. Instead of cadmium, these experiments have been using a 0.1% admixture of gadolinium with a large neutron absorption cross section leading to an 84% capture fraction. Absorption in gadolinium leads to the emission of gamma rays with a total energy of 8 MeV, within  $\sim 30 \mu\text{s}$  and  $\sim 6 \text{ cm}$  of the positron annihilation, thus providing a well recognizable delayed coincidence. The CHOOZ target scintillator consisted of 50% by volume Norpar-15 [21] and IPB + hexanol (also 50% by volume). The wave-length shifters were p-PTP and bis-MSB (1 g/l). The gadolinium was introduced as a solution of  $\text{Gd}(\text{NO}_3)_3$  in hexanol. Because of oxygenation of the nitrate the 4 m light attenuation length in the scintillator decreased with time at a rate of  $(4.2 \pm 0.4) \cdot 10^{-3}$  per day. This required a careful monitoring of the scintillator transparency using calibration sources. The light yield was 5300 photons/MeV.

The observation of a modification of the expected  $\bar{\nu}_e$  spectrum necessitates a very precise knowledge of the antineutrino flux emitted by the reactor as well as of the antineutrino interaction cross section. They failed to observe a disappearance of antineutrinos and the limit set on this oscillation was governed by these sources of systematics uncertainty. In order to overcome these limitations more recent reactor oscillation experiments use a second identical detector located close to the reactor in order to measure the expected interaction rate before the neutrinos have a chance to oscillate. The detector used by one such experiment, Double Chooz [22], located at the same location as CHOOZ but using, in addition, a near detector placed at 410 m from the reactors, will be described as an example. The scintillator, amounting to



**Fig. 8.4** The Double Chooz detector

10 tons, is housed in a central tank, Fig. 8.4, consisting of a clear acrylic. It is surrounded by a gamma catcher consisting of undoped liquid scintillator housed in a second acrylic shell that provides additional gamma detection probability for interactions occurring near the boundary of the central tank. A third envelope consisting of stainless steel holds the photomultipliers that view the two scintillator volumes through the acrylic shells. A buffer consisting of mineral oil fills the space between the stainless steel shell and the acrylic shell of the gamma catcher. It serves the purpose of absorbing any radioactivity emitted by the photomultipliers. In order to veto on cosmic muons the photomultiplier shell is itself surrounded by yet another scintillator volume housed in a final outer shell that also holds a second set of photomultipliers viewing this inner veto layer. Lastly, planes of scintillator counters cover the ceiling of the detector cavern, to identify more muons that traverse the surrounding material and are a potential source of neutrons. The scintillator chosen for the target is a 20/80 mixture of phenyl-xylethane (PXE)/dodecane with 0.1% gadolinium doping introduced as a dipivaloymethane molecule,  $\text{Gd(dpm)}_3$ . This has demonstrated long term stability. With PPO and Bis-MSB as fluors the mixture has an attenuation length greater than 5 m at 450 nm and a light yield of 7000 photons/MeV resulting in 200 detected photoelectrons/MeV. The positron detection



threshold is less than 700 keV, well below the threshold of 1.022 MeV of the inverse beta decay reaction.

Calibration of the detector is required to determine, in both detectors, the efficiency for observing the inverse beta decay reaction, the energy scales for positrons, neutrons and gamma, the timing of the photomultipliers and the light transport properties. To this end gamma sources, neutron sources and laser light flashers are used and deployed throughout the detector volumes in order to map out the relevant parameters. In the target this is done with an articulated arm at the end of which is mounted the calibrating device, the position of which is determined by the length and azimuthal position of the arm. In the gamma catcher a guide tube into which a source can be inserted at the end of a wire is used. The geometry of the tube and the wire length determine the position of the source.

Whereas Double Chooz was the first to report a hint for a non-zero  $\theta_{13}$ , two experiments have since produced the best measurements of this angle. They use the same concept as Double Chooz but have used either a larger neutrino flux (RENO [23]) or more detectors and more flux (Daya Bay [24]). RENO, in South Korea, is exposed to the flux of 6 reactors in a row totalling  $16.4 \text{ GW}_{th}$ . Its far detector is 168 m underground and 1380 m from the central reactor whereas its near detector is 46 m underground and 290 m from the reactor line. Its inner target weighs 15.4 tons and is viewed by 340 photomultipliers. Daya Bay, in China, uses eight identical detectors and is located near three reactor complexes Daya Bay, Ling Ao I and II, a total of  $17.4 \text{ GW}_{th}$ . Its far detector hall is 324 m underground, 1540 m from Ling Ao and 1910 m from Daya Bay and houses four detectors. One near detector hall 363 m from the Daya Bay complex and another one about 500 m from the Ling Ao complex each house 2 detectors. Each detector includes a 20 ton neutrino target viewed by 192 8" photomultipliers. Their inner vetos are tanks of water in which Cerenkov light is viewed by photomultipliers. The Daya Bay energy resolution is  $\sigma_E/E = 7.5\%/\sqrt{E}$ . Using a variety of radioactive sources they are able to determine the absolute neutrino energy scale to 1% and the relative energy scale between detectors to  $<0.2\%$ . The relative detection efficiency uncertainty was 0.13% and was substantiated by comparing rates of detectors in the same hall. Table 8.1 compares the systematic uncertainties achieved in Daya Bay to the ones in the CHOOZ single detector experiment demonstrating the effectiveness of a multiple detector and multiple location experiment. It should be noted that all three experiments have observed a structure in the positron energy distribution between 4 and 6 MeV when compared to Monte Carlo predictions based on the present understanding of a reactor neutrino flux. This structure is also seen in their near detectors (see for instance [25]) and its amplitude is proportional to the reactor flux. It is therefore believed to be due to our lack of complete understanding of the complex origin of a reactor neutrino flux.

KamLAND [26] is also an experiment observing reactor antineutrinos but studies oscillations in the domain of the solar  $\Delta m^2$ ,  $7.5 \times 10^{-5} \text{ eV}^2$ , and is therefore situated at an average distance of about 180 km from 53 Japanese power reactors to compensate for the much smaller  $\Delta m^2$ . The same reaction and technique as described above are used. However to observe enough events at this distance


**Table 8.1** The Daya Bay systematic uncertainties compared to the ones in CHOOZ

Variable	CHOOZ [%]	Daya Bay [%]
$\nu$ flux and cross section	1.9	–
Reactor power	0.7	–
Energy per fission	0.6	–
Number of protons	0.3	0.03
H/C ratio and Gd concentration	1.2	0.1
Spatial effects	1.0	0.02
Live time	–	0.01
Analysis cuts	1.5	0.082
Total	2.7	<0.13

a detector consisting of 1 kiloton of liquid scintillator had to be used. The scintillator is housed in a 13 m diameter transparent nylon balloon suspended in non-scintillating oil acting as a buffer and viewed by 1879 photomultipliers. This inner detector is surrounded by a 3.2 kiloton water Cerenkov counter which has the dual purpose of reducing  $\gamma$  rays and neutrons from the surrounding rock and of detecting cosmic ray muons. As well as measuring the oscillation pattern as a function of L/E of reactor neutrinos within their detector, KamLAND also made a measurement of geological neutrinos [27].

The KamLAND detector would also be used in the IsoDAR project [28] searching for sterile neutrinos through  $\bar{\nu}_e$  disappearance at a  $\Delta m^2$  of  $\sim 1 \text{ eV}^2$ . A 60 MeV cyclotron would be placed a few meters from the surface of KamLAND and 16.5 m from its centre, Fig. 8.5. The cyclotron would accelerate  $\text{H}_2^+$  ions (a hydrogen molecule with one electron removed) as the single charge for two protons of  $\text{H}_2^+$  reduces the repulsive force within a bunch and hence minimizes the effect of space charge blow up of the beam which in turn keeps beam loss down. A high current source, currently under commissioning, would produce the  $\text{H}_2^+$  ions which would be bunched with a radio-frequency quadrupole placed vertically above the centre of the cyclotron. The bunched ions would be bent electrostatically into the plane of the cyclotron. After 96 turns they would be extracted with a thin septum, stripped and transported 50 m, resulting in 10 mA of protons impinging on a beryllium target placed near the KamLAND detector. Neutrons produced in this target stream through a sleeve consisting of small beryllium spheres surrounded by highly enriched (99.995%)  $^7\text{Li}$ . A graphite reflector surrounds the target and sleeve. The neutrons captured by the lithium produce  $^8\text{Li}$  which subsequently decays producing  $\bar{\nu}_e$ 's entering the KamLAND detector in which they can be detected via IBD. The  $(12 \text{ cm}/\sqrt{E_{\text{MeV}}})$  spatial resolution and  $(6.4\%/\sqrt{E_{\text{MeV}}})$  energy resolution of KamLAND coupled with the detector size allows the observation of event rate oscillations as a function of L/E within KamLAND in addition to an overall disappearance of  $\bar{\nu}_e$ 's.

Reactor complexes are next planned to be used as sources of antineutrinos to illuminate two larger versions ( $\sim 10$  kilotons up from  $\sim 10$  tons) of the current




**Fig. 8.5** The layout of the IsoDAR 60 MeV cyclotron and its input stage:  $H_2^+$  source, RFQ and inflector. The same layout is to be used as an input stage for the DAE $\delta$ ALUS project (see Sect. 8.4.3)

reactor detectors. These detectors RENO50 [29] and JUNO [32] would be located  $\sim 50$  km from the reactors. At this distance the disappearance of reactor  $\bar{\nu}_e$ 's is dominated by the solar  $\Delta m^2$ . However the additional small effect arising from the atmospheric  $\Delta m^2$  is mass hierarchy dependent. A precise energy measurement of the IBD positron therefore allows the determination of the mass hierarchy, as well as more accurate measurements of  $\theta_{12}$  and  $\Delta m_{12}^2$ .

The possibility to deploy a 10 kiloton liquid scintillator immersed off shore in the vicinity of a nuclear reactor complex in order to perform oscillation physics has been investigated [33]. This Hawaii Anti-Neutrino Observatory, Hanohano, could alternatively be deployed far from a reactor in order to observe geological neutrinos.

Several scintillator detectors are also planned for deployment very close to reactors for neutrino oscillation into a sterile neutrino, accurate flux measurements and reactor monitoring, all using the IBD reaction. For this purpose they need to be compact. They also need to be segmented to mitigate the background from reactor



**Fig. 8.6** An exploded view of a single PVT cube used in the SoLid detector prototype

neutrons. STEREO will use the same detector technique as described above for the  $\theta_{13}$  experiments. Prospect [30] will run in 2 phases near a reactor at ORNL. Phase I will use a 3 ton single volume  $^6\text{Li}$  loaded liquid scintillator detector movable between 7 and 12 m. The scintillator is EJ-309 from Eljen Technology to which  $^6\text{Li}$ , PPO fluors and bis-MSB wavelength shifters have been added resulting in 6500 detected photons/MeV and a 4 m attenuation length. The volume is segmented by low mass optical separators into 120 segments  $14.4 \times 14.4$  cm in cross-sectional area and 120 cm long, read by a pmt at each end. The positron deposits its energy in the liquid scintillator and the neutron is observed via its capture in hydrogen or  $^6\text{Li}$ . The phase II detector will have a larger 10 ton mass while maintaining the same segmentation geometry and cover baselines between 15 and 19 m. Another example is the 1.6 ton SoLid built after prototyping a 288 kg version [31] deployed near the Belgian BR2 reactor. The final detector is built out of 12,000  $5 \times 5 \times 5$  cm<sup>3</sup> ELJEN Technology EJ-200 polyvinyl toluene (PVT) cubes, Fig. 8.6. Sheets of  $^6\text{Li}:\text{ZnS}(\text{Ag})$ , 225  $\mu\text{m}$  thick allow the detection of the IBD neutrons through break up of the lithium to an alpha and  $^3\text{H}$  with a Q-value of 4.78 MeV. The signals of each cube are read through two wave length shifting fibres connected to Hamamatsu S12572-050P multi-pixel photon counters.


Borexino [34] is a detector installed in Italy at the Laboratorio Nazionale del Gran Sasso (LNGS) for the measurement of solar neutrinos and in particular the  $^7\text{Be}$  862 keV monochromatic line using the purely leptonic reaction  $\nu + e \rightarrow \nu + e$ . This reaction results in an electron spectrum with a sharp edge at 665 keV. Borexino consists of 300 tons of liquid organic scintillator (pseudocumene and  $1.5 \text{ g} \cdot \text{l}^{-1}$  PPO as fluor) housed in a nylon vessel itself suspended in a stainless

steel sphere on which are mounted 2200 photomultipliers. The sphere is filled with a pseudocumene solvent with a quencher acting as a shield for radioactivity coming mainly from the tubes and is itself immersed in a water Cerenkov tank viewed by an additional 200 photomultipliers to identify cosmic ray muons. The light yield of the detector is 500 photoelectrons/MeV actually recorded. Timing information from the photomultipliers allow the spatial reconstruction of the event and hence the determination that it occurred within the fiducial volume. The  $\alpha$  and  $\beta^+$  components of natural radioactivity can be reduced by pulse shape discrimination whereas the  $\beta^-$  and  $\gamma$  components are indistinguishable from the signal. A reduction and thorough understanding of the background has allowed them to observe solar neutrinos with an energy as low as 150 keV and, hence, make the first direct observation of pp fusion solar neutrinos as well as measure the solar beryllium line and geoneutrinos [35]. After a year during which the background has been reduced through six cycles of water extraction the radiopurity levels are now  $2.7 \times 10^{-18}$  for  $^{14}\text{C}/^{12}\text{C}$  and, at 95% CL,  $<9.7 \times 10^{-19} \text{ g} \cdot \text{g}^{-1}$  for uranium and  $<1.2 \times 10^{-18} \text{ g} \cdot \text{g}^{-1}$  for thorium. This will allow improved measurements of solar and geoneutrinos as well as a new very short baseline neutrino oscillation project, SOX [36].


SOX is intended to search for oscillations of decay  $\nu_e$  or  $\bar{\nu}_e$  from a radioactive source into sterile neutrinos at the level of  $\Delta m^2$  of  $1 \text{ eV}^2$ . The sources being considered are  $^{51}\text{Cr}$  and  $^{144}\text{Ce}$ , with the latter already approved. The  $^{144}\text{Ce}$  would be placed in a pit under the Borexino detector. The small size of the 3.7–5.0 PBq source (about 1 L) coupled with the large 7 m size of the detector and its good spatial resolution of 12 cm and energy resolution of 3.5% would allow the observation of oscillation waves as a function of L/E within the detector as well as an overall measurement of  $\bar{\nu}_e$  disappearance. Given an existing detector, the most taxing task is the source. It would be produced in a Russian laboratory from the reprocessing of nuclear fuel and must then be extensively shielded and transported to the Gran Sasso by a circuitous route for safety reasons.

Totally active liquid scintillator detectors have also been used in accelerator experiments producing higher energy neutrinos. MiniBooNE [37], looking for  $\nu_\mu \rightarrow \nu_e$  oscillations in the Fermilab Booster neutrino beam in order to investigate the LSND signal [38], is exposed to neutrinos of about 1 GeV. The detector consists of 800 tons of mineral oil ( $\text{CH}_2$ ) held in a spherical tank. The density of the oil is  $0.86 \text{ g} \cdot \text{cm}^{-3}$  and has an index of refraction of 1.47. The light attenuation in this medium varies from a few cm at 280 nm to 20 m at 400 nm. The inner region (575 cm radius) is viewed by 1280 8-inch photomultipliers held on an optical barrier that separates it from a 35 cm thick outer region. This outer region, itself viewed by 240 tubes is used to veto events caused by charged particles entering the detector and to tag events that include particles exiting the detector in order to identify contained events. Cosmic ray events are greatly reduced by restricting the triggers to those occurring within a  $19.2 \mu\text{s}$  window starting  $4.4 \mu\text{s}$  before the  $1.6 \mu\text{s}$  long beam spill. The energy of an event is related to the total amount of light observed.  $\nu_\mu$  and  $\nu_e$  events are identified by the flavour (muon or electron) of the lepton in the final state CC interaction. Muons are distinguished from electrons using the light pattern of their Cerenkov rings as shown in Figs. 8.7 and 8.8. Muons give a sharp ring filled on

**Fig. 8.7** The Cerenkov light pattern characteristic of a muon in the MiniBooNE detector



**Fig. 8.8** The Cerenkov light pattern characteristic of an electron in the MiniBooNE detector



its interior as the muon approaches the tubes. Electrons give a fuzzy ring because of the many electrons and positrons each moving in a slightly different direction within the showers. In addition to the intrinsic  $\nu_e$  component of the beam caused by  $\mu$ , K and  $\pi$  decays the background to the  $\nu_e$  appearance search comes from  $\pi^0$  decays to two photons. This background can be greatly reduced by the ability of the detectors to observe separately the two electron-like rings produced by the two photons. The  $\pi^0$ 's can be reconstructed with a mass resolution of  $20 \text{ MeV}/c^2$ . The event vertex, direction and energy resolutions with which  $\nu_e$  events are reconstructed are 22 cm,  $2.8^\circ$  and 11% respectively. The experiment observed an unexplained excess of electromagnetic low energy events, but was not able to determine whether they were due to single photons or electrons due to the similarity of the rings produced by them.

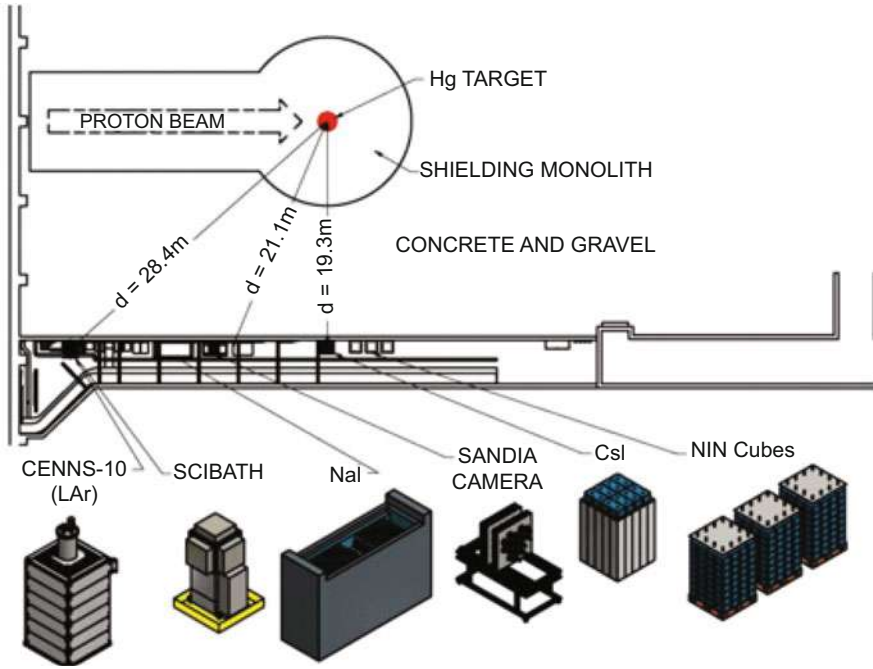
Liquid scintillator detectors can also be of a tracking kind, in which the scintillator is confined in tubes and read by wave length shifting (WLS) fibres. NOvA [39], an experiment that runs in the Fermilab 2 GeV off-axis NuMI beam at a distance of 810 km from the lab is such an example. It consists of planes of extruded PVC tubes alternating in the horizontal and vertical direction. Each tube is 3.87 cm by 6 cm in cross-sectional area, 15.6 m long and is filled with mineral

oil with 5% pseudocumene. Since one of its physics goals is  $\nu_e$  appearance it must be fine-grained enough to identify electrons and distinguish them from the showers produced from the decay photons of  $\pi^0$  mesons. To this end, each plane corresponds to a sampling frequency of only 17% of a radiation length. The WLS fibres are in the shape of a loop and are read at the end opposite the loop. Avalanche photodiodes with a quantum efficiency of about 80% are used and detect 40 photoelectrons for a minimum ionizing particle crossing a tube at the far end. They are produced in arrays of 16 diodes each with a cross-section of  $1.8 \times 1.0 \text{ mm}^2$  and must be run at a temperature of  $-15^\circ\text{C}$ . Each diode observes both ends of a fibre. The detector has an overall mass of 14 kilotons, consisting 70% of scintillator and the remainder of PVC. Its overall length is 67 m, with a cross section of  $15.7 \times 15.7 \text{ m}^2$ . The detector is located on the surface but the impact of cosmic rays is mitigated by the short beam spill of  $10 \mu\text{s}$  and the speed of the photodiodes. Nonetheless, to reduce the electromagnetic component of cosmic rays, the detector is covered by a 3 m overburden of concrete and barite.

Totally active tracking detectors can also be made of extruded solid scintillator bars usually read with WLS fibres embedded in a hole or a groove made in the scintillator. An example of such a detector is SciBar [40], a 15 ton detector consisting of 14,336 strips each of dimensions  $1.5 \times 2.5 \times 300 \text{ cm}^3$  and using 64-pixel multianode photomultipliers. It was first used in Japan on the KEK neutrino beam line and then moved to the NuMI beam line at Fermilab in the US.

Coherent Elastic Neutrino-Nucleus Scattering,  $\text{CE}\nu\text{NS}$ , is a process in which the neutrino interacts with the whole nucleus rather than with individual nucleons [41], leaving the nucleus whole and carrying very little energy since the momentum transfer must be small. The recoiling nucleus subsequently produces secondary recoils and scintillation light. The  $\text{CE}\nu\text{NS}$  cross section is several orders of magnitude larger than, for instance, the IBD cross section as it depends on the square of the number of neutrons in the nucleus. However, the smallness of the energy release made the process impossible to measure until recently. The COHERENT experiment [42] overcame this difficulty by using a 14.6 kg sodium-doped CsI crystal 34 cm long. The heavy cesium and iodine nuclei, provide the large cross sections and the large scintillation light yield necessary to detect low energy recoil nuclei down to a few keV. The crystal was read by a super bialkali low background Hamamatsu R877-100. The source of neutrinos was the decay of pions and muons produced by the Spallation Neutron Source at Oak Ridge National Laboratory. Protons on target (POTs) were delivered in  $1 \mu\text{s}$  long spills at a rate of 60 Hz resulting in  $4 \times 10^{18}$  isotropically emitted neutrinos per day. The detector was placed in a basement corridor at a location, Fig. 8.9, that provided 12 m of neutron-moderating concrete and gravel in the direct line of sight to the SNS target, thus reducing neutron-induced recoil nuclei background (NIN) to an acceptable level. Cosmic rays were also reduced with an 8 m water overburden. The detector was enclosed in high-density polyethylene, to reduce NIN, as well as in both low activity and in standard lead. Muon vetos and water tanks containing a neutron moderator completed the shielding of the detector. The photomultiplier signals were amplified and digitized at 500 MSamples/s over  $70 \mu\text{s}$  intervals starting  $55 \mu\text{s}$  before the POT






**Fig. 8.9** The COHERENT experiment layout, showing the proton beam, target, shielding and experimental area

signal. Two  $12\mu\text{s}$  windows, one preceding and one following the POT trigger, allowed the comparison of data respectively unrelated and related to the beam. The  $40\mu\text{s}$  interval preceding these two windows was used to veto events due to previous energy depositions. The window following the POT signal showed a distribution of events consistent both in energy with coherent scattering and in time distribution with pion and muon decays. These signals were absent in the window preceding the POT, allowing the experimenters to announce a first observation of  $\text{CE}\nu\text{NS}$  at the 6.7 standard deviation confidence level. Detectors with different technologies such as liquid argon and NaI[Tl] crystals are currently installed in the same location with further expansions being considered.

### 8.3.2 Water Cerenkovs

These are large volumes of ultra-pure water in which charged particles produced in neutrino interactions are detected through the Cerenkov light they emit. The measurement and separation of electrons, muons and  $\pi^0$ 's is as was described in the context of MiniBooNE and illustrated in Figs. 8.7 and 8.8. Several experiments [43, 44], originally conceived to search for proton decay have pioneered this technique





**Fig. 8.10** The electron direction relative to the sun position in solar neutrino candidate events observed in SuperKamiokande

for the study of atmospheric and solar neutrinos, and, as a byproduct, have made the first recording of neutrinos emitted by a supernova, namely SN1987A. These detectors have been followed by the most productive one, SuperKamiokande(SK) [45], a 50 kiloton detector placed in the Kamioka mine in Japan at a depth of 2700 m water equivalent. It consists of two concentric cylindrical detectors. The inner one (ID) is viewed by 11,146 photomultipliers of 20 inch diameter providing a 40% coverage while the outer one(OD) is viewed by 1885 8 inch tubes. The absence of signal in the OD distinguishes fully contained events from partially contained ones. SK has successfully observed neutrinos ranging in energy from a few MeV(solar neutrinos) to several tens of GeV (atmospheric and accelerator neutrinos). The electron energy and direction produced in the elastic scattering reaction used to observe solar neutrinos are related to the incident neutrino energy and direction. The pointing accuracy is such that the origin of these neutrinos can be clearly associated to the sun, Fig. 8.10. In order to observe as much of the solar neutrino spectrum as possible, it has minimized background such as to be able to lower their detection threshold to  $\sim 3.5$  MeV. They have observed neutrino interactions coming from above and from below and have observed the reduction of  $\nu_\mu$  interactions from below (long baseline) due to  $\nu_\mu \rightarrow \nu_\tau$  oscillations. Using a neural network approach they were also able to identify the resulting  $\nu_\tau$  CC component. SK is currently the heart of the T2K long baseline experiment [46] in which it is exposed to a beam of neutrinos or antineutrinos from the JPARC accelerator laboratory 295 km away. The beam is produced starting with 30 GeV protons and is an off-axis beam with a narrow neutrino energy spectrum peaked at 600 MeV. The experiment includes a near detector which will be described in Sect. 8.3.5. Their excellent electron and muon identification have allowed them to observe  $\nu_\mu$  disappearance as well as  $\nu_e$  appearance, and measurements of  $\sin^2\theta_{23}$ ,  $\Delta m_{23}^2$  as well as  $\theta_{13}$ . In long baseline  $\nu_e$  appearance experiments such as T2K or NOvA, the measured value of  $\theta_{13}$  is correlated to the yet unknown CP violation phase, leading these experiments to

quote their measurement of  $\theta_{13}$  as a range of values driven by the allowed CP phase values. Combining this range of  $\theta_{13}$  with the precise reactor experiments measurement of  $\theta_{13}$ , allows to limit the possible values of the CP violation phase.

Several long baseline projects based on the water Cerenkov technique have been proposed to measure the mass hierarchy and the CP phase via  $\nu_\mu \rightarrow \nu_e$  oscillations. However because of the low signal event rate expected, detectors of the order of the megaton are needed. This would enable these detectors to continue the very successful non-beam physics programme of SK, namely atmospheric and solar neutrino physics, proton decay searches and supernovae watches. The MEMPHYS [47] project was planned in the context of a potential neutrino beam [48] from CERN to a new laboratory in the Frejus tunnel. It consists of 3 cylindrical water Cerenkov counters placed in contiguous caverns for a total of 0.5 megatons. Hyper-Kamiokande(HK), described in their Letter of Intent [49], is a natural extension of SK and would use the same beam as SK (600 MeV off-axis at  $2.5^\circ$ ) but with its power upgraded from 0.75 MW to 1.35 MW, mostly by increasing the JPARC main ring repetition rate to 0.86 Hz. In their latest design the beam would impinge on a 0.52 Mton (0.38 Mton fiducial) water Cerenkov detector consisting of two 74 m diameter and 60 m high cylinders located 295 km from J-Parc in a cavern 650 m underground and 8 km south of SK. The Cerenkov light would be observed by 80,000 50 cm diameter Hamamatsu R12860 photomultipliers of a new Box and Line design providing 40% coverage with single photons detected with a 24% efficiency and a 1 ns timing resolution. The photomultipliers have survived extensive pressure and implosion tests. Alternative sensors are also being investigated such as Hybrid Photo Detectors and the multi-photomultipliers concept developed for KM3Net [72]. The extrapolation of the water Cerenkov technique to a megaton-sized detector is driven to a large extent by the cost of the optical sensors and their production schedule, making the spacing and size of the sensors of prime importance. Similar detectors were considered for installation at SURF DUSEL [50] in order to observe neutrinos from a new beam from Fermilab. However, as will be discussed below, the liquid argon technique has been adopted instead.

A test [52] in the Super-Kamiokande detector demonstrated that neutrons from the IBD reaction  $\bar{\nu}_e + p \rightarrow e^+ + n$  could be detected with the addition of gadolinium in the water. A 2.4 L acrylic vessel was filled with a 0.2%  $\text{GdCl}_3$  mixture. A BGO crystal containing an Am/Be radioactive source was placed in its middle. The  $\alpha$  particles emitted by the ameritium interacted in the beryllium via  $\alpha + {}^9\text{Be} \rightarrow {}^{12}\text{C}^* + n$ . By immersing the vessel in the SuperKamiokande detector and using the 4.43 MeV carbon deexcitation photon as a trigger it was demonstrated that the neutron could be detected via its absorption in the gadolinium, as described earlier in Sect. 8.3.1, with an efficiency of 66.7% with a 3 MeV threshold for delayed events. It was estimated that a background reduction of  $2 \times 10^{-4}$  could be achieved at a 10 MeV prompt event analysis threshold for  $\bar{\nu}_e$ . This opens the way for the use of the water Cerenkov technique for the detection of  $\bar{\nu}_e$ 's of geological or reactor origin.

Detectors that measured a deficit in the solar neutrino spectrum were all sensitive to  $\nu_e$  only. In order to definitely prove that the deficit was due to a flavour

transformation rather than a disappearance it remained to prove that the overall flux of neutrinos, including all three flavours, was as predicted by the standard solar model. This was achieved by SNO [51] by measuring neutral current reactions which can occur for all three flavours since they do not have the energy constraint of charged currents, namely the mass of the appropriate produced charged lepton. It's heavy water ( $D_2O$ ) target made it sensitive to three neutrino reactions, including neutral current reactions:

- Elastic scattering on electrons:  $\nu_{e,\mu,\tau} + e^- \rightarrow \nu_{e,\mu,\tau} + e^-$
- Charged current absorption of neutrinos by deuterons:  $\nu_e + d \rightarrow e^- + p + p$
- Neutral current disintegration of the deuteron with a threshold of 2.2 MeV:  $\nu_{e,\mu,\tau} + d \rightarrow \nu_{e,\mu,\tau} + n + p$ . This reaction could not be observed in light water because of the binding energy of oxygen being larger than the solar neutrino energies. The neutron was detected by absorption on deuterium or on  $^{35}Cl$  in added  $MgCl_2$ . At a later stage of the experiment an array of  $^3He$  filled proportional tubes was added providing a highly efficient neutron detection through the reaction  $^3He + n \rightarrow p + ^3H + 764 \text{ keV}$ .

Because of the overall neutron production of only a few tens per day, care had to be given to select radiopure materials. The detector is located in a mine at a depth of 6000 m.w.e, thus reducing the cosmic ray background to 70/day. It consists of an acrylic sphere containing 1000 tons of heavy water viewed by 9438 photomultipliers. It is immersed in a structure containing light water for shielding and support. The proportional counters were placed in the heavy water in a lattice with 1 m spacing. The counters were 5.08 cm in diameter and filled with 85%  $^3He$  and 15%  $CF_4$  at a pressure of 2.5 atm. Electrons were detected by the Cerenkov light they emitted. These included those produced in the primary interaction as well as those produced through Compton scattering on electrons of photons emitted through neutron absorption.

Cerenkov detectors are also the technique of choice for cosmological neutrinos. The scarcity of these very high energy neutrinos requires the use of large naturally occurring target and detection media such as a lake [53] or sea water [54–57] or Antarctic ice [58, 59] which can be instrumented with photomultipliers at the scale of  $1 \text{ km}^3$ . The photomultipliers are connected into vertical strings and lowered in the water or, in the case of ice, into holes melted using hot water. The strings have to be located at great depths to shield the detector from downgoing cosmic muons. This necessitates the inclusion of the photomultipliers in pressure vessels. They must also be in regions of high light transmission in order to maximize the spacing of photomultipliers and reduce the cost. The most advanced of these detectors is ICECUBE [59] in the Antarctic. It consists of 86 strings positioned in a 125 m hexagonal grid at a depth between 1450 and 2450 m. Each string includes 60 digital optical modules (DOM). Each DOM is a 35 cm pressure vessel containing a 25 cm diameter pmt, a wave form digitizer, a fast ADC and electronics self-triggering at the level of 1/4 of a photoelectron. Digital information is sent to the surface. It is complemented by a  $1 \text{ km}^2$  surface array consisting of 160 ice-filled tanks. The average absorption and scattering lengths of the ice at the detector depth are 110 m

and 20 m respectively at 400 nm. Its energy threshold is 100 GeV. A subarray, DeepCore, consisting of 8 strings closely spaced at 40–70 m and with a DOM separation of 7 m instead of 17 m allows the observation of neutrinos with energies as low as 10 GeV. ICECUBE can search for point sources with an angular resolution of  $1.5^\circ$ , based on the signal arrival time at the photomultipliers, which is also used to reject downgoing cosmic ray muons. ICECUBE made the first observation of cosmological neutrinos between 20 and 2000 TeV, at energies high enough that they could not be attributed to atmospheric neutrinos. Several extensions of ICECUBE are being considered. ICECUBE-Gen2 [60] consists of an additional 120 strings to augment the coverage by about an order of magnitude, coupled with new more directional sensitive detectors as well as smaller ones to reduce the hole diameter and hence the fuel cost. Another is PINGU [61], a proposal to study neutrino oscillations parameters using a sample of about 60,000 atmospheric neutrinos with a threshold energy of a few GeV obtained by instrumenting a 6 Mton clear ice volume at the bottom of ICECUBE with 26 closely spaced strings each carrying 192 optical modules.

In addition to optical detection of the Cerenkov light, Antarctic ice has also been used to detect the coherent radio signals emitted by the cascade resulting when a neutrino interacts in a dielectric medium, the ice. This kind of radiation was predicted by Askaryan [62] and is caused by propagating showers acquiring a negative charge excess through Compton scattering and the annihilation of positrons in the dielectric. When this excess moves at a velocity greater than the velocity of light in the medium, Cerenkov radiation will be emitted and will be coherent for wavelengths longer than the transverse dimension of the shower, corresponding to  $\sim 1$  GHz. The electric field strength will be proportional to the shower energy. The radio attenuation length has been measured to be about 1600 m at 300 MHz, making the ice suitable for widely spaced detectors. This technique has been applied using either detectors observing the ice from balloons and satellites or with detectors placed right on the ground. The first technique allows the observation of large volumes of ice but will have higher detection thresholds. The second, due to the proximity of the detectors to the ice will have lower detection thresholds but will be limited to smaller detection volumes. The ANITA [63], Antarctic Impulsive Transient Antenna experiment, is a good example of the first technique. It used a balloon flying under the NASA Long Duration Balloon program at an altitude of 37 km which allows the observation of the whole antarctic ice sheet ( $1.5 \times 10^6 \text{ km}^2$ ). It flew 3 times for 35, 31 and 22 days respectively and used horizontal and vertical polarization antennas with a band width of 200–1200 MHz. The data was read with 2 GSamples/s digitization resulting in a 100 ns waveform per channel and per trigger. ANITA has been able to set the best limit on neutrinos for energies greater than  $10^{19.5} \text{ eV}$  as well as finding no neutrino coincident within 10 min of any of 12 Gamma Ray Bursts (GRBs). A possible extension of this technique would be EVA [64], the Exa Volt Antenna project, which would lower the energy threshold by a factor of 10 using the inner surface of a super-pressure balloon as a toroidal reflector antenna 115 m in diameter.

ARA [65], the Askaryan Radio Array, is a ground based radio array using several stations of 16 antennas each embeded 200 m deep into the South Pole ice. Three stations are operational with two more being installed. Each station consists of four strings separated by 10 m and each consisting of a mixture of horizontal and vertical polarization antennas. The trigger requires 3 out of 16 signal to exceed a power threshold within 110 ns. So far, the deployed stations have found no neutrino candidate including a search centred on 57 GRBs. Ground arrays have also been proposed on the Ross Ice Shelf, (ARIANNA [66]) and in Greenland (GNO [67]).

This technique has also been extended to neutrino interactions in underground Rock Salt which allows for better shielding from cosmic rays and also to interactions in the loose layer of regolith sands on the moon surface. Both of these materials have attenuation lengths for radio waves of the order of 100 m. While the rock salt experiments use ground based detectors, the lunar ones, use various radio telescopes pointed at the lunar limb. The first lunar radio experiment was the Parkes Lunar Radio Cherenkov experiment [68] and was followed by LUNASKA [69], GLUE [70] which constrained the neutrino flux above  $10^{21}$  eV and NuMoon [71] which constrained it above  $10^{23}$  eV.


Following the good performance of ANTARES off Toulon in the Mediteranean, a northern hemisphere kilometer cube detector, KM3Net, is currently being implemented [72]. It will consist of two modules. ORCA, at the ANTARES site will consist of 115 closely packed strings in order to address neutrino oscillations and the mass hierarchy in the energy range 3–50 GeV. The site will have a diameter of 200 m and a height of 100 m. ARCA, off Capo Passero in Sicily, will consist of two blocks of 115 widely spaced strings each block having a diameter of 1 km and a height of 600 m. ARCA's physics objectives are neutrinos from extra terrestrial sources above 1 TeV and the origin of high energy cosmic rays. Each string of both modules will consist of 18 optical modules, each housing 31 7.5 cm diameter photomultipliers. These yield a photocathode area that exceeds by a factor of three that of a single 25 cm photomultiplier, provides some directional information and a good separation between one and two photoelectron signals.

Lastly a detector, GVD [73], the Gigaton Volume Detector is under construction at Lake Baikal to observe cosmological neutrinos. It will consist of eight 120 m diameter clusters of 8 strings each, each string carrying 36 optical modules housing a 10" Hamamatsu photomultiplier. The cluster separation is 300 m. It is planned to extend GVD to 18 clusters to make it a cubic kilometer detector.

A detector immersed in the sea in the gulf of Taranto had also been proposed [74] to study  $\nu_\mu \rightarrow \nu_e$  oscillations in the then CNGS beam [75], the axis of which was at a height of 40 km above the surface, thus placing the detector in an off-axis location. In this case the array of photomultipliers consisted of a vertical plane facing the beam to observe Cerenkov light from the electrons and muons produced in charged current interactions and identify them from their light pattern.

### 8.3.2.1 Liquid Argon

The Liquid Argon Time Projection Chamber (TPC) is a detector technique that provides accurate imaging of interactions while providing a moderately dense medium ( $1.4 \text{ g} \cdot \text{cm}^{-3}$ ) thus very suitable for neutrino physics. It consists, Fig. 8.11, of a volume of liquid argon sandwiched between a cathode and anode providing a drift field of the order of  $500 \text{ V/cm}$ . Charged particles traversing this volume ionize it and the resulting electrons drift to the anode. To improve the uniformity of the electric field, a field cage surrounds the volume between the cathode and anode and is constructed with hoop-shaped electrodes held at potentials that increase from the cathode to the anode. The anode consists of a succession of wire planes, usually two or three. The first planes encountered by the drift electrons are biased such as to prevent them from being captured. Signals in these planes are by induction only. The last plane actually collects the electrons. The wires of the different planes are at differing orientations to the vertical. Associating signals in the different wire planes according to their timing allows the reconstruction of two of the coordinates of the drift electrons and therefore of the track portion they originated from. The third coordinate, along the drift field, is obtained using the drift velocity ( $\sim 1 \text{ mm}/\mu\text{s}$ ), and the time difference between the electron signal at the wire and the time of the neutrino interaction. The latter is in turn obtained from the timing



**Fig. 8.11** The principle of signal recording in a Liquid Argon Time Projection Chamber as depicted in [78]

of the scintillation light emitted in the argon by the products of the interaction and recorded by photomultipliers. A track deposits energy along its trajectory and this is recorded as pulse heights in the wires. The pulse height distribution provides particle identification through the ionization pattern whereas the pulse height sum is a measure of the particle energy. The latter can also be obtained by range.

ICARUS [76], was the first to develop and use this technique. It was located at the Gran Sasso LNGS laboratory and was exposed to neutrinos produced by the CNGS beam 732 km away at CERN. It consists of two 300 ton modules each  $3.6 \times 3.6 \times 19.9 \text{ m}^3$ . Each module includes a central high voltage plane and, along each of its long sides, three sets of detection wire planes, with orientation at  $0^\circ$  and  $\pm 60^\circ$ . Electrons drift over a maximum distance of 1.5 m in the electric field perpendicular to the wire planes. This very complete detector relies on long drift distances and therefore requires high purity liquid argon. The purity achieved [77] during a technical run was such as to allow an electron drift lifetime of 1.8 ms equivalent to an electron mean free path of 280 cm. The electron drift velocity at 89°K increased from 0.5 mm/ $\mu\text{s}$  at an electric field of 0.1 kV/cm to 2 mm/ $\mu\text{s}$  at 1.0 kV/cm.


In the US, the first liquid argon TPC used in a physics experiment was ArgoNeuT [82], a 0.35 ton detector installed upstream of the MINOS near detector in the NUMI beam line at Fermilab. It produced significant low energy neutrino energy results as well as providing a test bed for subsequent larger liquid argon detectors.

The liquid argon technique has since been adopted for SBN [83], the Short Baseline Neutrino beam program at Fermilab, intended to investigate the possibility of additional, sterile, neutrinos. It uses the Booster Neutrino beam and consists of three liquid argon TPC detectors: SBND at 110 m from the neutrino source, MicroBooNE at 470 m and ICARUS at 600 m. The first to be installed was MicroBooNE [84], approved to observe electrons and photons and determine the origin of the low energy electromagnetic excess observed by MiniBooNE (Sect. 8.3.1). Its good spatial resolution would allow it to distinguish converting photon showers which are not associated to the primary vertex and are twice minimum ionizing at the conversion point from prompt electrons which are connected to the vertex and are singly ionizing. This should result in a good electron/photon discrimination. The TPC is inserted in a foam insulated cylindrical cryostat. It is 10.4 m long, 2.3 m high and 2.5 m wide. Electrons drift horizontally over a maximum of 2.5 m (corresponding to a maximum drift time of 2.25 ms) in a 0.273 kV/cm electric field and are recorded by two induction and one collection successive wire planes inclined respectively at  $\pm 60^\circ$  and  $0^\circ$  to the vertical. The experiment was the first liquid argon TPC experiment to fill its cryostat without prior evacuation. It has achieved [85] an electron drift-lifetime of 18 ms corresponding to an  $\text{O}_2$  equivalent contamination of 17 ppt and a loss of signal of 12% over the 2.5 m drift length. It also placed pre-amplifiers and shapers in the cold to reduce connection lengths and hence electronic noise. The amplified signals exit the cryostat and are digitized in warm ADCs before entering the DAQ electronics for Huffman compression and storage. This is done in two independent streams. The first stream records all the data occurring over



8 ms encompassing the trigger. This long recording time relative to the maximum drift time allows the complete reconstruction of cosmic rays traversing the TPC before or after the trigger but having part of their tracks reaching the wires during the event drift time. The second stream, intended for non-beam physics studies such as supernovae neutrinos, records all the data continuously but applies a zero suppression algorithm. Both the fast (6 ns) and slow ( $\sim 1\mu\text{s}$ ) components of the argon ultra violet scintillation light are recorded by 35 Hamamatsu 5912-02MOD photomultipliers installed behind the anode wires and coated with Tetraphenyl Butadiene (TPB) to wave length shift the light from the ultra violet to the visible. The fast component is used to provide a trigger in time with the  $1.6\mu\text{s}$  beam spill and to tag cosmic ray tracks entering the detector during the event drift time. These are also tagged by a cosmic ray detector surrounding the cryostat and assembled out of scintillation bars read by Kuraray WLS Y11 (200) S-type multiclad wave length shifting fibers and Hamamatsu S12825-050P multi-pixel silicon photomultipliers. A UV laser is used to map the TPC electric field, especially in the regions of non-uniformity caused by space charge effects. MicroBooNE has been collecting data since 2015. It has developed algorithms to distinguish between cosmic rays entering the detector at a rate of 4 kHz (because of its surface location) and neutrino interactions. It is also in the process of developing reconstruction algorithms for electromagnetic showers that, at these low energies, can include gaps due to the propagation of low energy photons. Nonetheless liquid argon provides a remarkable visualization of events as depicted in Fig. 8.12. MicroBooNE is employing a Deep Learning technique [86] called semantic segmentation for the identification of the various classes of interactions.

The second detector to be installed will be ICARUS refurbished under the WA104/NP01 programme [87] at CERN. The following improvements were made to the detector:



**Fig. 8.12** A MicroBooNE Neutrino interaction event, showing charged particle tracks originating at the vertex and two photons, probably from a pizero, converting away from the vertex but pointing back to it



- The cathode plane was rebuilt to correct for up to 5 mm non-planarity.
- The optical system was upgraded to 360 8'' Hamamatsu 5912-mod (10 stage) cryogenic photomultipliers with TPB coating their face and read out by the CAEN V1730B 500 MHz 14 bit ADC system. The speed of this readout should allow the correlation of beam events with the Booster RF substructure, namely 1.15 ns pulses separated by 19 ns. If achieved, this correlation will reduce further the contamination of cosmic rays.
- The TPC electronics was modified as follows. The analogue signal shaping time was reduced to 1.5  $\mu$ s to match the electron transit time between wire planes and reduce undershoot in induction. Serial ADCs as well as a serial bus architecture with optical links were adopted. The feedthrough flange was used as the electronics backplane.

The third detector, SBND, described and referred to as LAr1-ND in [83], is a detector intended to measure the intrinsic beam composition, in particular of  $\nu_e$ , before oscillations can occur. However its closeness to the neutrino target ensures a large number of neutrino interactions and hence a rich cross-section measurement programme. Its dimensions are 5 m along the beam, 4 m in height and 4 m laterally. The electrons drift along this latter dimension which consists of two 2 m drift spaces placed side by side. The Cathode Plane Assembly, CPA, is located in the middle and one Anode Plane Assembly, APA, is placed on either side and consists of the same number of wire planes and orientation as MicroBooNE. Each APA is made up of two 2.5 m wire frames along the beam but the U and V wires are connected to ensure continuous coverage. Unlike MicroBooNE, the ADCs will be in the cold together with the front end pre-amplifiers and shapers. The digitized signals will be multiplexed out of the cryostat via an FPGA. This will reduce the electronic noise and reduce the size of feed throughs. Upon exiting from the cryostat the signals will be converted to optical signals and sent, over optical fibres, to the warm DAQ electronics which will be identical to the one used by MicroBooNE. A 100 kV high voltage will provide a 500 V/cm drift field, the uniformity of which will be ensured by a field cage constructed with roll-formed tubes. A cosmic ray tagger of similar construction to the MicroBooNE one and a membrane cryostat will encase the detector. The light detection system will use the same pmt type and readout system as ICARUS. SBND will pioneer several detector concepts such as APAs and CPAs intended to be applied to the DUNE detector.

The liquid argon technique has been chosen for DUNE [88], the Deep Underground Neutrino Experiment  $\nu_\mu \rightarrow \nu_e$  oscillation search intended to determine whether CP is violated in the neutrino sector and to measure the mass hierarchy. It will also address non-neutrino beam physics such as potential supernovae, proton decay and  $n\bar{n}$  oscillations. The liquid argon technique was chosen instead of that of water Cerenkov for its good electron/photon discrimination resulting in a higher electron efficiency and therefore the possibility to use a smaller detector to achieve the same sensitivity. DUNE will be located 1475 m underground at SURF [89], the Sanford Underground Research Facility, in Lead, South Dakota and will be observing neutrinos produced at Fermilab 1300 km away. It will consist of four


TPC modules each containing 17,000 tons (10,000 tons fiducial volume) of liquid argon. The construction of the first module will follow the APA, CPA concept being tested in SBND, the so-called single-phase (liquid) approach. Its TPC dimensions are 12 m high, 14.5 m wide and 58 m along the beam. Three rows of APAs will be interleaved with 2 rows of CPAs, all oriented parallel to the beam. The APA-CPA horizontal separation, or drift length, will be 3.6 m, necessitating a 180 kV high voltage system for a 500 V/cm drift field. Each row of APAs consists of 25 vertically stacked pairs. Each row of pairs of facing APA-CPA is surrounded by a field cage. An APA consists of 4 wire planes separated by 4.76 mm with biases of  $-655$  V,  $-365$  V,  $0$  V and  $+860$  V and orientation of  $0^\circ$ ,  $+35.7^\circ$ ,  $-35.7^\circ$  and  $0^\circ$  respectively. The wire separation is 4.7 mm. The TPC data is continuously digitized at 2 MHz by cold ADCs, serialized and transferred out of the cryostat on 12,000 high speed links per 10 kton module. They are received by Reconfigurable Computing Elements (RCEs) that buffer the raw data, zero-suppress it and pass it on to the trigger. While the zero-suppressed data is kept for non-beam physics, a second pass collects the full data set in regions of interest selected by the trigger. The photon detector system consists of light guides (2.2 m long, 83 mm wide and 6 mm thick) coated with TPB. The UV scintillation light impacting on the surface is re-emitted inside the bar at 430 nm and internally reflected in the guide to reach 12 SensL Cseries 6 mm<sup>2</sup> SiPMs. Ten such devices are mounted on each APA.

The second module will introduce a novel concept, the dual-phase approach first studied in [79] and tested as described in [80], in which the drifting electrons exit the liquid and are amplified in gaseous argon above the liquid. The concept is illustrated in Fig. 8.13 depicting the design of the dual phase prototype, ProtoDUNE-DP [81], currently under test at CERN. The DUNE dual phase module will consist of a 12 m wide, 12 m high and 60 m long homogeneous volume TPC. The electrons drift vertically over a maximum of 12 m in the 500 V/cm field provided by a segmented cathode at the bottom, the anode readout at the top and 60 stacked horizontal rectangular field rings. The reduction of the number of drift electrons reaching the wires due to absorption over the long drift space is compensated by the amplification in the gas. A 2 kV/cm field between a grid located just below the surface of the liquid and the Large Electron Multipliers (LEMs) charge amplification devices, causes the electrons to be extracted. The LEMs consist of a 1 mm thick printed circuit board with a micro-pattern of holes through its thickness and with one electrode on the top and one on the bottom surfaces. A 3 kV potential difference between the two electrodes results in a high field in each hole and the amplification of electrons entering them by about an order of magnitude through an avalanche process. The charge is collected in a two-dimensional  $x$ ,  $y$  readout plane above the LEMs.

The technology of subsequent modules will depend on the performance of the single phase and double phase prototypes currently being built and tested at the CERN neutrino platform. DUNE also plans to use a near detector located close to the Fermilab neutrino source.

The addition of a magnetic field to a liquid argon detector would greatly enhance its capabilities. This has been tested [90] with an 11 L chamber placed in a 0.55 T magnetic field and the drifting properties were found to be preserved. However

**Fig. 8.13** The design of the ProtoDUNE dual phase prototype under test at CERN, showing the various components of the detector and, in particular, the amplification LEMs at the anode



measuring the charge and momentum of electrons would be challenging due to the short, 14 cm, radiation length of liquid argon resulting in only the first few cms being useful for magnetic measurements.

### 8.3.3 Calorimeters

These detectors are well suited to investigations requiring a measure of the total energy of an event rather than energy measurements of individual particles other than muons.

#### 8.3.3.1 Iron-Scintillator

The CDHS detector [91] was used in the CERN SPS neutrino beam to study neutrino interactions in the energy range 30–300 GeV. It consisted of magnetized iron modules built from alternating layers of iron and scintillator and separated by drift chambers for a total of 1250 tons. Each module was constructed of circular iron plates 3.75 m in diameter and with a 30 cm central hole for the coil insertion. The coil consisted of 30 turns and was powered at 1000 A resulting in magnetic field of 1.65 T on average. It was uniform to  $\pm 1.5\%$  azimuthally and dropped by about 20% with increasing radius. Two types of modules were used. Seven modules used fifteen 5 cm thick plates and twelve modules used five 15 cm thick plates. The iron plates alternated with planes of eight 45 cm wide NE110 scintillators except for

the last four modules which used a single scintillator plane for triggering. The drift chambers were 4 m wide hexagons and drifted vertically or at  $\pm 60^\circ$  to the vertical in order to resolve ambiguities. The average efficiency was typically 99.5% and the spatial resolution 0.7 mm, which was adequate given the contribution of multiple scattering in the iron.

NuTeV/CCFR [92] used at Fermilab for a similar range of neutrino energies, differed from CDHS in that the calorimeter was separate from the magnetic spectrometer used to measure muon momenta. The 690 ton calorimeter consisted of  $168 \times 3 \times 5.15$  cm steel plates instrumented with Bicron 517L scintillator oil counters placed every two plates and drift chambers every four plates. This was followed by the magnetized iron toroidal spectrometer with an inner diameter of 25 cm to accommodate the four coils and an outer diameter of 350 cm. It consisted of three sections each followed by a drift chamber and two additional drift chambers downstream of the last section for improved momentum resolution. An important feature of this experiment was that a calibration beam was available in situ to determine the response [92] of the detector to electrons, muons and hadrons. The hadronic resolution was  $\sigma_E/E = 0.86/\sqrt{E(\text{GeV})} \oplus 0.022$  with an absolute scale uncertainty of 0.43%. The muon scale uncertainty was 0.7%, dominated by the field map determination in the iron. NuTeV performed a precise measurement of  $\sin^2 \theta_W$  necessitating the measurement of both neutral and charged current events. They discriminated between the two on the basis of event length defined as the number of scintillator planes with non-zero pulse height in an event.


The 5.4 kiloton, 31 m long MINOS detector [93], similar in concept to CDHS, consists of 486 2.54 cm thick iron plates interleaved with planes of scintillator strips read by wave length shifting fibres. It was exposed to the Fermilab NuMI beam and housed in the Soudan mine 735 km away from Fermilab. This detector is studying  $\nu_\mu$  disappearance and therefore the shape and magnitude of the beam energy distribution must be very well understood. To minimize its dependence on Monte Carlo calculations the experiment is equipped with a near detector, located 1015 m from the target, which measures the beam spectrum and composition before any significant oscillations can occur. A transfer matrix is then used to predict the flux at Soudan. The transformation does not depend simply on the inverse of the square of the distance to the detector as the near detector, being close to the target, is not exposed to an exact point source because of the finite (725 m) length of the decay tunnel. The extruded polystyrene scintillator strips are 4.1 cm wide and 1 cm thick and are read by a 1.2 mm wave-length shifting fibre housed in a groove. The fibres are read by multi-anode photomultipliers, structured as 16 pixel in the far detector and 64 pixel in the near one. The data is multiplexed to reduce the number of readout channels. The coil provides a toroidal magnetic field in the iron allowing the measurement of the momentum and charge of secondary muons.

### 8.3.3.2 Fine-Grained

CHARM II [94] was a detector designed to measure  $\sin^2 \theta_W$  in  $\nu_\mu - e$  scattering. The scattered electron is produced very forward unlike background events arising from  $\nu$ -nucleon events. A cut of  $E\theta^2 \leq 1 \text{ MeV}$  and  $\theta < 10 \text{ mrad}$ , where  $E$  and  $\theta$  are the scattered electron energy and production angle to the beam direction was used to reject this background, necessitating a very good angular resolution. Hence, glass, a low  $Z$  material to minimize multiple scattering, was selected as target material. Each of the 420 48 mm thick glass plates was followed by a plane of 352 plastic streamer tubes with a 1 cm pitch. The wires were readout in digital mode and 18 mm wide cathode strips glued to the outside of the tubes in a direction orthogonal to the wires were readout in analog mode to provide a measure of the energy and centroid of the electron showers. Consecutive modules had their strip and wire orientations rotated by  $90^\circ$  and consecutive modules with the same orientation had their wire spacing shifted by half the wire pitch. A scintillator plane was inserted in the detector after every 5 glass plates. The total mass of the calorimeter was 692 tons covering a volume of  $3.7 \times 3.7 \times 15.4 \text{ m}^3$ . An electron angular resolution,  $\sigma_\theta/\theta$ , varying between  $15\text{--}20(\text{mrad})/\sqrt{E(\text{GeV})}$  over the  $2\text{--}24 \text{ GeV}$  energy range of the experiment was achieved[95] as well as a vertex resolution of about 22 mm. The ability to discriminate the electrons from  $\nu_\mu - e$  scattering from background is demonstrated in Fig. 8.14. A muon spectrometer consisting of six of the CDHS modules followed the glass target and provided a momentum resolution of 14% at  $20 \text{ GeV}/c$  and an angular resolution at the vertex of  $18 \text{ mrad}/E(\text{GeV})$ .

### 8.3.4 Emulsions


Detectors based on the photographic emulsion technique have a sub-micron spatial resolution and are therefore the detector of choice when searching for secondary vertices related to charmed particles or  $\tau$  leptons. Until recently, this technique was limited because of the difficulty in scanning the emulsion. However recent developments in fast microscopes have revived it. Although  $\nu_\mu$  disappearance in atmospheric neutrinos was widely believed to be due to  $\nu_\mu \rightarrow \nu_\tau$  interactions, it needed to be demonstrated through  $\nu_\tau$  appearance in a  $\nu_\mu$  beam. This was undertaken, using emulsions, by E531 [96] at Fermilab and by CHORUS [97] at CERN. At the neutrino energy of these experiments the  $\tau$  travels only about 1–2 mm. CHORUS, the more sensitive of the two experiments used a 770 kg emulsion target built out of plates consisting of a  $90 \mu\text{m}$  plastic base holding  $350 \mu\text{m}$  thick emulsion layers on either side. The target was divided into four stacks each one followed by three interface emulsion sheets and a scintillating fibre tracker. Other sheets of emulsions were interleaved between the stacks and were changed several times throughout the data-taking in order to be exposed to fewer tracks and therefore ease the track reconstruction in the bulk emulsion. The fibres (more than 1 million) were read out by 58 optoelectronic readout chains each consisting of four image



**Fig. 8.14** The number of events as a function of  $E\theta^2$ , with  $E$  and  $\theta$  respectively the scattered electron energy and direction as obtained during the antineutrino running of CHARM II, demonstrating the ability to identify electrons from  $\nu_\mu - e$  scattering as evidenced by the sharp peak at small  $E\theta^2$

intensifiers and a CCD camera. The target was followed by a hadron spectrometer including an air-core hexagonal magnet, an electromagnetic calorimeter and a muon spectrometer. The pulsed hexagonal magnet provided a momentum resolution varying between 20 and 50% in the momentum range 0–10 GeV/c. The scanning of the emulsions was made with fully automated Ultra Track Selector microscopes based on the track selector principle [98]. A series of tomographic images (Fig. 8.15) are taken in the emulsion at successive depths along the beam direction. Tracks then appear as aligned grains when the images are shifted according to track angle. Although they failed to find oscillations because of the kinematic region which they were sensitive to, E531 [100] and CHORUS [99] were successful in observing secondary vertices from a large sample of charm decays.


The search was then taken over by OPERA [102]. This experiment collected data at the LNGS laboratory using the CNGS beam. The long baseline of OPERA allowed the search for  $\nu_\tau$  appearance in the  $\Delta m^2$  region then favoured by  $\nu_\mu$  disappearance. It used the emulsion cloud chamber technology as it is well suited to search for detached vertices or kinks over distances of the order of a mm. The 1766 ton detector was made up two supermodules. Each supermodule included 31 walls of bricks each 8.3 kg brick consisting of 57 plates of emulsions alternating



**Fig. 8.15** The principle of tomographic scanning of emulsions

with 1 mm sheets of lead. An emulsion plate was made of two  $44\mu\text{m}$  emulsion sheets on either side of a  $200\mu\text{m}$  plastic layer. Each wall was followed by two planes of scintillator trackers one with vertical strips and one with horizontal strips. Each supermodule was completed with a magnetized iron muon spectrometer constructed with iron plates interleaved with resistive plate chambers and preceded and followed by 7 m long drift tubes. Bricks identified by the scintillator tracker as likely candidates for having been the site of a neutrino interaction were removed from the setup on a daily basis using a robot. They were briefly exposed to cosmic rays to provide sheet to sheet alignment and were then dismantled. The emulsions were developed and scanned. They observed [103] five  $\nu_\tau$  CC interactions through the detection of the resulting  $\tau$  decaying to a single hadron in three events, to 3 hadrons in one event and to a muon in the fifth event. Their secondary vertex and kink detecting capability was again demonstrated by their observation of neutrino generated charm events [104].

Emulsions were also used in DONUT [101], the first experiment to observe  $\nu_\tau$  interactions albeit from  $\nu_\tau$ 's intrinsically present in the beam rather than from oscillations. In order to reduce the number of  $\nu_\mu$  and  $\nu_e$  in the beam, the experiment used a neutrino beam produced by impinging the Fermilab proton beam in a 1 m long tungsten beam dump. Most pions and kaons that normally give rise to  $\nu_\mu$  and  $\nu_e$  interacted before decaying, whereas the decay  $D_s \rightarrow \tau \nu_\tau$  followed by  $\tau \rightarrow \nu_\tau + \dots$  is fast enough to produce  $\nu_\tau$ 's before the  $D_s$  interacted. This resulted in a neutrino beam with a 5%  $\nu_\tau$  content. Four emulsion targets were interleaved with scintillating fibre trackers. An electromagnetic calorimeter and a muon spectrometer completed the detector. Three types of emulsion targets were



**Fig. 8.16** The identification of secondary vertices using the emulsion cloud chamber technique described in the text

used. In order to increase the number of  $\nu_\tau$  interactions and reduce the amount of emulsion needed, in some of the targets DONUT used the emulsion cloud chamber technology, in which emulsion sheets are interleaved with lead or stainless steel plates as shown in Fig. 8.16. DONUT chose to use 1 mm thick stainless steel plates. In these targets, two types of emulsion plates were used: 100  $\mu\text{m}$  emulsion sheets on either side of 200  $\mu\text{m}$  or 800  $\mu\text{m}$  plastic base. The remainder of the targets used bulk emulsion: 350  $\mu\text{m}$  emulsion layers on either side of a 90  $\mu\text{m}$  base. In the emulsion cloud chamber detectors the  $\tau$  vertex is predominantly in the iron and therefore unobserved. But the precision with which the neutrino interaction vertex and the  $\tau$  decay products can be reconstructed allows the identification of secondary vertices as described in Fig. 8.16. A plate to plate alignment accuracy of 0.2  $\mu\text{m}$  over a  $2.6 \times 2.6 \text{ mm}^2$  area was achieved by matching high momentum tracks in successive layers using position and direction information. This allowed a measure of the momentum of a particle using its multiple scattering, itself estimated using repeated changes of direction of the particle as it traverses the emulsion sheets.

This experiment also used external trackers to predict the position of interesting interactions in the emulsion. To facilitate this match each emulsion stack was followed by a changeable sheet, changed often to reduce its track density and facilitate the tracker-emulsion match. However it also used fast enough microscopes




to allow the use of stand alone techniques in which the emulsions were scanned without external information.

The latest scanning microscopes developed in Japan and Italy can measure at an average speed of 50 and 20 cm<sup>2</sup>/h respectively. In addition to providing very precise spatial reconstruction emulsions, because of the large number of measurements along a track, can provide momentum measurements, as described above, and particle identification using ionization measurements, especially near the end point of stopping tracks. Please refer to Chap. 5 of this volume which addresses emulsion techniques in more details.

### 8.3.5 Hybrid Detectors

NOMAD [105] was a detector, Fig. 8.17, built to search for  $\nu_\mu \rightarrow \nu_\tau$  oscillations by observing  $\nu_\tau$  interactions in the same  $\nu_\mu$  beam as used by CHORUS. However, unlike CHORUS, it intended to identify  $\tau$ 's not through the reconstruction of its separate decay vertex, but through kinematic criteria such as the missing transverse momentum arising from the unobserved neutrinos produced in  $\tau$  decay. This demanded very good momentum resolution. The  $\tau$  decay modes used in the analysis were  $\tau^- \rightarrow \nu_\tau + \text{hadrons}$  and  $\tau^- \rightarrow \nu_\tau + \bar{\nu}_e + e^-$ . The latter mode was particularly useful as its main background is CC interactions of the intrinsic  $\nu_e$  in the beam but this background is greatly suppressed because of the small  $\nu_e$  contamination,  $\sim 1\%$ . However this mode did require very good electron identification. This latter




**Fig. 8.17** The NOMAD hybrid detector used at CERN in the search for  $\nu_\mu \rightarrow \nu_\tau$  oscillations

requirement as well as the requirement of good momentum resolution dictated the use of a light detector. The detector consisted of 49 drift chamber modules each one providing three coordinates with sense wires at  $0^\circ$ ,  $+5^\circ$  and  $-5^\circ$  degrees to the vertical. The chambers were built out of honeycomb panels made of aramid fibres sandwiched between two kevlar skins. These panels provided the target material, 2.7 tons, for neutrino interactions. The average density of the target was  $0.1 \text{ g} \cdot \text{cm}^{-3}$ , close to that of a hydrogen bubble chamber and the drift chambers provided measurements every 2% of a radiation length ( $X_0$ ). The spatial resolution was  $150 \mu\text{m}$  providing a momentum resolution of  $\sigma_p/p = \frac{0.05}{\sqrt{L}} \oplus \frac{0.008p}{\sqrt{L^5}}$ , with the momentum  $p$  and the track length  $L$  expressed in  $\text{GeV}/c$  and meters respectively. The chambers were complemented by 9 modules of transition radiation detectors consisting of polypropylene foils and straw tubes containing an 80% xenon–20% methane gas mixture. These modules together with a  $1.6 X_0$  lead and proportional tube preshower and a  $19 X_0$  lead glass array provided the necessary  $e/\pi$  separation. These detectors were housed in a  $7.5 \times 3.5 \times 3.5 \text{ m}^3$  dipole magnet providing a 0.4 T horizontal magnetic field. The lead glass array, being inside the magnetic field, was read by tetrodes with a gain of 40 and had an energy resolution  $\Delta E/E = (3.02/\sqrt{E(\text{GeV})} + 1.04)\%$ . An iron-scintillator hadron calorimeter was located just outside the magnet coil and was followed by two muon detection stations consisting of large area drift chambers located after 8 and 13 interaction lengths.

Silicon is another technique that provides very precise track localization and hence, secondary vertex identification as has been proved repeatedly in hadronic interactions. NOMAD-STAR [106] was a 45 kg prototype of a possible application of this technology to neutrino interactions. It consisted of 4 plates of boron carbide providing the interaction mass interleaved with five planes of silicon detectors. Each plane consisted of ten 72 cm long ladders of 12 silicon-strip detectors with a pitch of  $50 \mu\text{m}$ . It was exposed to a neutrino beam within the NOMAD detector and, in conjunction with the rest of the detector, it was able to reconstruct 45 charm decays. The hit to noise ratio was 17:1 and the hit finding efficiency 98%. The impact parameter resolution of the  $\mu^-$  produced in a  $\nu_\mu$  CC interaction relative to a hadronic jet consisting of at least three charged particles was  $33 \mu\text{m}$ .

The magnet used by NOMAD is now being used in the T2K experiment as part of the hybrid near detector [107] located 280 m from the target. The magnet houses:

- Scintillator planes interleaved with lead or brass optimized for photon detection and  $\pi^0$  reconstruction.
- Three time projection chambers (TPC) using Micromegas modules for the drift electrons amplification and readout.
- These TPC's are interleaved with fine-grained detectors consisting of strips of scintillator providing target mass.
- A scintillator and radiator electromagnetic calorimeter.
- Scintillator planes housed in the slots located in the return yoke providing a muon range detector



**Fig. 8.18** The triangular scintillator strips and wave length shifting fibres of MINERvA

All scintillators in this near detector use Multi-Pixel Photon Counters as photosensors, a total of 50,000 channels. These are well suited as they operate in a magnetic field and provide single photon detection capability.

MINERvA, Main INjector ExpeRiment for  $\nu$ -A, [108], an experiment to make precision measurements of neutrino cross sections using several nuclear targets (carbon, iron and lead) uses the NuMI beam at Fermilab and is located in front of the MINOS near detector. It consists of a fully active central detector surrounded and followed by electromagnetic and hadronic calorimeters. The central detector is built out of planes of 128 scintillator strips of triangular cross section, Fig. 8.18, and the electromagnetic and hadronic calorimeters use the lead-scintillator and steel-scintillator technology respectively. Wave-length shifting fibres are embedded in the scintillator strips and the light is channelled via clear fibres to multi-anode photomultipliers. Muons are identified and measured using the MINOS near detector. The overall cross section of the detector is hexagonal.

### 8.3.6 Radiochemical Detectors

Solar neutrino interactions are recorded by radiochemical experiments using the reaction:  $\nu_e + (A, Z) \rightarrow e^- + (A, Z + 1)$ . The atoms of  $(A, Z + 1)$  produced are chemically extracted every few weeks, so this is not a real time process. They were first observed by the Homestake experiment [109] using  $^{37}\text{Cl}$  producing  $^{37}\text{Ar}$ . It was followed by three others, Gallex [110], Sage [111] and GNO [112], all of which used  $^{71}\text{Ga}$ , changing to  $^{71}\text{Ge}$ . Their characteristics are listed in Table 8.2. These experiments were housed underground to reduce cosmic ray background. In spite of the large flux of solar neutrinos on earth only a few such reactions occur

**Table 8.2** Characteristics of the radiochemical solar neutrino experiments

	Homestake	Gallex	SAGE	GNO
Location	South Dakota	Gran Sasso	Baksan mine	Gran Sasso
Material	C <sub>2</sub> Cl <sub>4</sub>	Gallium (solution)	Gallium (metallic)	Gallium (solution)
Initial isotope	<sup>37</sup> Cl	<sup>71</sup> Ga	<sup>71</sup> Ga	<sup>71</sup> Ga
Detected isotope	<sup>37</sup> Ar	<sup>71</sup> Ge	<sup>71</sup> Ge	<sup>71</sup> Ge
Mass [tons]	615.0	30.3	57.0	30.3
Threshold	0.814 MeV	0.233 MeV	0.233 MeV	0.233 MeV
Extraction rate	3–4 months	3–4 weeks	3–4 weeks	3–4 weeks
Half-life of detection reaction	34 days	16.5 days	16.5 days	16.5 days

daily even for detectors weighing about a hundred tons due to the small neutrino interaction cross section at these low energies.

The Homestake experiment was located in the mine of the same name in South Dakota at a depth of 4200 m.w.e. The detector consisted of a cylindrical tank containing 615 tons of C<sub>2</sub>Cl<sub>4</sub> and with 5% of its volume filled with helium gas at a pressure of 1.5 atmospheres. The argon produced was removed from the tank by bubbling helium through the tank and then trapping the argon in a cryogenically cooled charcoal absorber. Following several stages of purification the argon was transferred to a proportional counter after adding 7% of methane. The extraction efficiency was measured by inserting and extracting known amounts of either <sup>36</sup>Ar or <sup>38</sup>Ar. Because of the very low event rate possible radioactive contaminants in the tube material had to be minimized. The counters consisted of a highly refined iron cylindrical cathode and a 12–25 μm tungsten wire anode. The <sup>37</sup>Ar decays occur dominantly through K orbital electron capture depositing 2.82 keV of energy in the counter. This deposition is highly localized (100 μm) thus allowing it to be distinguished by pulse shape and rise time discrimination from background which is less localized.

Gallex and GNO, located in the LNGS, used 30.3 tons of Gallium containing 12 tons of <sup>71</sup>Ga in an aqueous solution acidified by the addition of HCl. This ensures that the <sup>71</sup>Ge produced is in the form of the highly volatile GeCl<sub>4</sub> in contrast with the non-volatile GaCl<sub>3</sub>. The extraction procedure, as exemplified by that of GNO, is as follows [113]:

- The atoms of GeCl<sub>4</sub> (approximately 16 per 3–4 week run) are extracted into water by pumping nitrogen gas through the system.
- They are then converted into a gas, GeH<sub>4</sub> and mixed with Xenon.
- This mixture is introduced into proportional tubes 32 mm long and 6.4 mm in diameter made of ultrapure Suprasil quartz. The cathode consists of a single silicon crystal with impurities limited to ≤2ppt <sup>238</sup>U, ≤0.2ppt <sup>232</sup>Th and ≤0.1ppb <sup>40</sup>K. The anode is a 13 μm tungsten wire. The efficiency for transferring the

Germanium nuclei to the counters is measured to be 95–98% using non-radioactive Germanium carriers.

- X-rays occurring through the reactions  $e^- + (A, Z + 1) \rightarrow (A, Z) + \nu_e$  are detected over a period of about 6 months although the mean life of the reaction is only 16.5 days, allowing a good estimate of the background.
- The  $^{71}\text{Ge}$  decays produce pulses of 10.4 keV or 1.1 keV for K and L captures respectively. The localized nature of this ionization allows the reduction of background using amplitude and shape analysis of the recorded pulses to a level of less than 0.1 event/day.
- The counters are calibrated 5 times during a 6-month exposure using a Gd/Ce X-ray source.

Gallex measured [114] their extraction efficiency using a 60 PBq  $^{51}\text{Cr}$  source of 750 keV neutrinos (90%) and 430 keV (10%) neutrinos. They found a ratio of measured/expected signal of  $0.93 \pm 0.08$ . Their extraction efficiency was also confirmed [115] to be as expected to within 1% by introducing several thousand atoms of  $^{71}\text{As}$  that decay to  $^{71}\text{Ge}$ .

The SAGE detector was built using up to 60 tons of metallic Gallium. It was housed in the Baksan Neutrino Observatory in the Caucasus at a depth of 4700 m.w.e. While the liquid gallium was stirred at a rate of 80 rpm the Germanium was extracted from it by oxidizing it using a weakly acidic aqueous solution. The subsequent steps are similar to the procedure described above. Their extraction efficiency was also measured with a Chromium source.

The reaction threshold in chlorine only allows the observation of the beryllium and boron neutrinos whereas the threshold in gallium allows, in addition, the observation of some of the pp neutrinos.

### 8.3.7 Bubble Chambers

Bubble chambers were heavily used in earlier studies of neutrino interactions and were instrumental in making significant advances in the understanding of the properties of neutrinos [116]. Their filling varied from liquid hydrogen to heavy liquids, the latter used to increase the overall target mass, to contain the secondary hadrons produced and to convert photons. They were placed within a magnetic field in order to measure the momenta of the charged particles produced in the neutrino interactions. Their time resolution was poor as they were sensitive to all events occurring within a beam spill of typically millisecond duration. This could be improved by associating them to external electronic detectors.

Gargamelle was a cylindrical chamber 4.8 m long and 1.8 m in diameter. It was situated in a 2 T magnetic field produced by two coils. Neutral currents were first identified using this chamber [117] with a heavy freon ( $\text{CF}_3\text{Br}$ ) filling resulting in a density of  $1.5 \text{ g} \cdot \text{cm}^{-3}$  and an interaction length of 58 cm. The identification of

neutral currents required the rejection of events containing muons in the final state. Muons were defined as satisfying one of the following categories:


- a particle leaving the visible volume without undergoing a nuclear scatter
- a particle which stops in the chamber and decays to an electron
- a negative particle stopping in the chamber without producing visible products (44% of negative muons are absorbed in the nucleus).

The hypothesis that the observed neutral current candidates could have been due to neutral hadrons entering the chamber was rejected as this would have resulted in a decrease of their number as a function of depth within the chamber, an effect that was not observed.

A 12-foot bubble chamber [118] (a  $26\text{ m}^3$  cylinder) was used in the neutrino beam at the Argonne Zero Gradient Synchrotron (ZGS). It was the first chamber to use a superconducting magnet.

The superconducting technology was then used in subsequent bubble chambers. BEBC (the Big European Bubble Chamber) was a  $33.5\text{ m}^3$  bubble chamber, Fig. 8.19, operating in a 3.5 T magnetic field and was exposed to the CERN neutrino beams. It was equipped with an External Muon Identifier (EMI) [119] consisting of proportional wire chambers placed behind an iron absorber and covering an area 6 m high and 25 m wide. A muon candidate track observed within the chamber was confirmed as a muon if, when its trajectory was extrapolated to the EMI, a hit was found within a distance consisting with multiple scattering. In order to reduce the miss-classification of events due to the random association of a neutral current event with a background muon in the EMI, an Internal Picket Fence [120] of proportional tubes placed between the chamber body and the magnet cryostat provided timing information for all events occurring within BEBC with a resolution of 230 ns full width. The tubes operated within the BEBC magnetic field. The chamber was used with fillings of liquid hydrogen, liquid deuterium or a neon-hydrogen mixture. This allowed the study of both  $\nu p$  and  $\nu n$  interactions. It was also used with a  $3\text{ m}^3$  Track Sensitive Target (TST) which, when filled with hydrogen within a neon-hydrogen environment provided a sample of clean interactions within the hydrogen which could be compared to interactions in the neon for which nuclear effects had to be taken into account. In addition the heavier neon allowed a more efficient detection of secondaries.

The 15-foot bubble chamber at Fermilab ran in a magnetic field of 3.0 T. When filled with a neon-hydrogen mixture [121] (61.7% atomic neon and 38.3% atomic hydrogen) it provided a 23 ton target with a density of  $0.75\text{ g} \cdot \text{cm}^{-3}$ , an interaction length of 125 cm and a radiation length of 40 cm. It was also fitted with an external muon identifier.



**Fig. 8.19** The layout of the Big European Bubble Chamber (BEBC) including the External Muon Identifier and Internal Picket Fence

## 8.4 Ongoing Development Efforts on Neutrino Beams

Research based on neutrinos is currently proceeding along two paths. The first is focussed on completing our knowledge of the oscillation parameters by determining the mass hierarchy and searching for CP violation in the neutrino sector by comparing  $\nu$  to  $\bar{\nu}$  oscillations. The second avenue is determining whether sterile neutrinos exist. The first goal could be achieved with the presently planned detectors and beams described above. Nonetheless more accurate measurements would greatly benefit from higher intensity beams and much larger detectors. This has motivated several avenues of research pursued in the US and in Europe.

### 8.4.1 *Beta Beams*

Beta beams [122, 123] are beams of neutrinos based on the production, storage and  $\beta$ -decay of radioactive ions. A possible European solution was studied in the context of the Eurisol project [124].  ${}^6\text{He}$  ions which, decaying via  $\beta^-$ , produce  $\nu_e$  and  ${}^{18}\text{Ne}$  ions, which decaying via  $\beta^+$ , yield  $\bar{\nu}_e$  would be stored. These ions would be accelerated to the energy required to produce decay neutrinos of the required energy and then stored in a race-track shaped storage ring. The Lorentz boost produces a well focussed forward beam which would illuminate one or more detectors in line with the straight sections of the storage ring. The search for CP violation would proceed by observing  $\nu_e \rightarrow \nu_\mu$  and  $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$  oscillations. These oscillations would result in the observation of muons produced via the charged current interaction of  $\nu_\mu$ 's and  $\bar{\nu}_\mu$ 's. These beams are particularly advantageous for the study of these oscillations as they do not, unlike present accelerator neutrino beams, have an intrinsic oscillated flavour component. Thus an important background is eliminated. The detector envisaged for this project was MEMPHYS, a water Cerenkov counter described in Sect. 8.3.2, located 130 km away from a potential beta beam source at CERN. This distance would require neutrino energies of a few hundred MeV to be at oscillation maximum. These low energies and distances would preclude any resolution of the mass hierarchy. As usual, an additional detector near the storage ring would be needed to study the beam before oscillations can occur. It is estimated that  $2.9 \times 10^{18}$   ${}^6\text{He}$  ions and  $1.2 \times 10^{18}$   ${}^{18}\text{Ne}$  ions decaying per year in the straight sections would be needed to meet the physics requirements. Whereas this seems achievable for  ${}^6\text{He}$ , new production methods [125] would be needed for  ${}^{18}\text{Ne}$ .

### 8.4.2 *Neutrino Factory*

A neutrino factory [126, 127] uses the decay of muons to produce neutrinos. The first step is to produce pions using a very high intensity proton beam impinging on a target. The decay of these pions then produce muons. Before they are injected into a storage ring their momentum and angular spread must be reduced to maximize their capture efficiency. This is done by phase rotation and ionization cooling. Longitudinal momentum spread would be reduced by phase rotation using the Neuffer scheme. This entails capturing multi bunches of muons with a very high Radio Frequency (RF) and rotating their phase with decreasing RF along the cooling channel. Angular spread (transverse momentum) would be reduced through ionization and subsequent longitudinal acceleration using RF cavities. The storage ring includes straight sections pointing to one or more detectors [129]. In the scheme studied in the context of the International Scoping Study [128] muons of both signs of about 20 GeV/c can be captured and stored simultaneously. Storage ring geometries have been identified that can deliver both neutrinos and antineutrinos to one or more detectors. In a race track geometry neutrinos and antineutrinos would



be identified in the same detector by time of arrival, itself related to the timing separation of  $\mu^+$  and  $\mu^-$  bunches in the storage ring. In a triangular geometry two straight sections could point to two detectors. The physics envisaged with this project is the observation of  $\nu_e \rightarrow \nu_\mu$  oscillations using the  $\nu_e$ 's produced in  $\mu^+$  decay. The signal is the observation of a  $\mu^-$  as opposed to the copious  $\mu^+$ 's produced by the interaction of the  $\bar{\nu}_\mu$  also produced in  $\mu^+$  decay. The identification of the charge of these wrong sign muons necessitates the use of a magnetic detector. A 50kton magnetized iron detector [130] coupled with scintillator or RPC's lends itself to this. Less dense detectors such as liquid argon TPC's and emulsion detectors [131] using the OPERA technology are also being considered to observe respectively electrons and  $\tau$  leptons. These would allow the observation of additional oscillation channels which would be useful in removing ambiguities in the determination of oscillation parameters. With the higher energies and distances being considered the resolution of the mass hierarchy could be envisaged in addition to the search for CP violation.

The very high intensity proton beams needed to produce an adequate neutrino flux impose strong restrictions on the type of material used for the proton target. MERIT [132] is an R&D experiment at CERN intending to investigate the effectiveness of a mercury jet in a solenoidal field as a target. The constant flow of mercury would circumvent the problems related to stress and heating of a solid target. Muon cooling is being studied by MICE [133] at RAL with its strong synergy with MUCOOL [134] at Fermilab, with a setup including capture solenoids, liquid hydrogen absorbers and RF cavities. Incoming and outgoing spectrometers measure the effectiveness of the cooling.

### 8.4.3 High Current Cyclotrons

The present accelerator based long baseline experiments intend to compare oscillations of  $\nu_\mu$  to oscillations of  $\bar{\nu}_\mu$ . However the  $\bar{\nu}_\mu$  beam has much more  $\nu_\mu$  background than the  $\nu_\mu$  beam has  $\bar{\nu}_\mu$  background. This is because of the  $\pi^+$  to  $\pi^-$  ratio at the proton target being larger than unity and because of the  $\nu$  interaction cross section being larger than the  $\bar{\nu}$  cross section. DAE $\delta$ ALUS [135] is an experiment aiming to remedy this situation by using the decay at rest of pions to produce a very pure source of  $\bar{\nu}$  that would illuminate a detector also exposed to a long base line  $\nu_\mu$  beam. A beam of 800 MeV protons produced by a high current cyclotron impinges on a thick target producing pions which stop in the target, with the  $\pi^-$  being captured before decaying resulting in a beam dominated by the decay of  $\pi^+$ . As a consequence there will be essentially no  $\pi^- \rightarrow \mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e$  and hence any  $\bar{\nu}_e$  interaction observed must be from a  $\bar{\nu}_\mu$  to  $\bar{\nu}_e$  oscillation. The DAE $\delta$ ALUS project proposes installing three sources of pions at rest: one at 20 km from the detector which, for an average neutrino energy of 45 MeV, would be at the maximum oscillation probability and at the same L/E as the long baseline beam, one at 8 km to observe the rise in  $\bar{\nu}_e$  appearance and one at 1.5 km for flux normalization.

The  $\bar{\nu}_e$  would be observed through IBD for which a liquid argon detector as planned for DUNE would not be suitable. However, a water Cerenkov detector such as T2K or Hyper-K (especially if containing gadolinium to enhance the neutron capture rate as described in Sect. 8.3.2) or a large liquid scintillator detector would be ideal. As the flux out of the cyclotron would be continuous unlike the flux from the long baseline accelerator, the two sources of events would be distinguishable through absolute timing. The DAE $\delta$ ALUS collaboration is currently involved in increasing the current capability of cyclotrons to reach the 10 mA of protons necessary. The 800 MeV cyclotron would be superconducting, accelerate  $H_2^+$  ions and would use as an injector the 60 MeV cyclotron described earlier in the context of IsoDAR. The  $H_2^+$  ions would be stripped at extraction.

## 8.5 Conclusions

Neutrino detectors use the whole range of detector technologies available to high energy physicists. The smallness of neutrino cross sections necessitates the use of very large detectors that have ranged up to 50 kilotons when man-made and even 1000 megatons when using sea water or antarctic ice. The exception is the recent observation of coherent neutrino-nucleus scattering, a much larger cross section process, which allows the detection of neutrinos with smaller, albeit complex, detectors. Future generations of neutrino detectors to be used in conjunction with Very Long Base Line beams will address the outstanding questions in neutrino oscillation physics, namely the determination of the mass hierarchy and of CP violation in the neutrino sector as well as the determination of the possible existence of sterile neutrinos. In addition neutrinos are being used as probes. Ultra high energy (PeV) neutrinos originating in regions of space undergoing very violent processes are now beginning to be detected thus providing a new tool to study these processes. At the other end of the scale, neutrinos of a few MeV allow us to study the Earth and monitor reactors. These issues will require a whole range of detector sizes, up to the megatons, while at the same time requiring the precise measurements of the energies of electrons and photons and the identification of the secondary vertices of charmed particles and  $\tau$  leptons. These detailed studies dictate the use of varied and complex detectors, thus ensuring that neutrino experiments will continue to use the very latest developments in detector technology.

## References

1. [http://www.ethbib.ethz.ch/exhibit/pauli/neutrino\\_e.html](http://www.ethbib.ethz.ch/exhibit/pauli/neutrino_e.html)
2. Reines F *et al* 1953 Phys. Rev. **92** 830
3. Danby G *et al* 1962 Phys. Rev. Lett. **9** 36
4. Kodama K *et al* 2001 Phys. Lett. **B504** 218

5. Koshiba M 2008 Experimental results on neutrino masses and mixings, in Handbook of Particle Physics, Editor H. Schopper Landolt-Bornstein, Volume 1/12A
6. Camilleri L, Lisi E and Wilkerson J 2008 Neutrino Masses and Mixings: Status and Prospects Annu. Rev. Nucl. Part. Sci. **58** 343
7. Bahcall JN, Pinsonneault 2004 Phy. Rev. Lett. **92** 121301 and references therein.
8. Gaisser TK, Honda M 2002 Annu. Rev. Nucl. Part. Sci. **52** 153
9. Learned JG, Mannheim K 2000 Annu. Rev. Nucl. Part. Sci. **50** 679
10. Bemporad C *et al* 2002 Rev. Mod. Phys. **74** 297
11. Astier P *et al* 2003 Nuc. Instr. and Meth. **A515** 800
12. van der Meer S 1961 CERN Yellow Report **61-07**
13. Bernstein R *et al* 1994 Sign Selected Quadrupole Train FERMILAB-TM-1884
14. Yu J *et al* 1998 NuTeV SSQT Performance FERMILAB-TM-2040
15. Beavis D *et al* Long Baseline Neutrino Oscillation Experiment at the AGS (Proposal E889), Physics Design Report BNL 52459 (1995)
16. Dydak F 1980 Beam-Dump Experiments CERN-EP/80-204
17. Wachsmuth H 1979 Neutrino and Muon Fluxes in the CERN 400 GeV Proton Beam Dump Experiments CERN/EP 79-125
18. De Rujula A *et al* 1993 Nucl. Phys. **B405** 80
19. Apollonio M *et al* 2003 Eur. Phys. J. **C27** 331
20. Boehm F *et al* 2001 Phys. Rev. **D64** 112001
21. ExxonMobil <https://ilrc.ucf.edu/documents/ILRC%2000000080/MSDS%2000000080.pdf>
22. Ardellier E *et al* 2006 (Double Chooz Collaboration) hep-ex/0606025v4
23. Choi J H *et al* (RENO Collaboration) 2016 Phys. Rev. Lett. **116** 211801
24. An F P *et al* (Daya Bay Collaboration) 2016 Nucl. Instr. and Meth. **A811** 133
25. Choi J H 2016 Phys. Rev. Lett. **116** 211801
26. Abe S *et al* 2008 Phys. Rev. Lett. **100** 221803
27. Araki T *et al* 2005 Nature **436** 499
28. Alonso J R and Nakamura K (IsoDAR Collaboration) 2017 arXiv:1710.09325v1[physics.ins-det]
29. Joo K K (RENO and RENO50 Collaborations) 2017 J. Phys. Conf. Ser. **888** 012012
30. Ashenfelter J *et al* 2016 J. Phys.G:Nucl.Part.Phys. **43** 113001
31. Abreu Y *et al* 2018 arXiv:1802.02884v1[physics.ins-det]
32. Adam T *et al* (JUNO Collaboration) 2015 arXiv:1508.07166v2[physics.ins-det]
33. Learned J *et al* 2008 Hanohano: a deep ocean anti-neutrino detector for unique neutrino physics and geophysics studies. arXiv:0810.4975v1[hep-ex]
34. Alimonti G *et al* (Borexino Collaboration) 2009 Nucl. Instr. and Meth. **A600** 568
35. Derbin A and Muratova V *et al* (Borexino Collaboration) 2016 arXiv:1605.06795v1[hep-ex]
36. Bellini G *et al* 2013 arXiv:1304.7721v2[physics.ins-det]
37. Aguilar A A *et al* (MiniBooNE Collaboration) 2008 arXiv:0806.4201v1[hep-ex]
38. Aguilar A *et al* 2001 Phys. ReV. **D64** 112007
39. Ayres DS *et al* (NOvA Collaboration) 2005 arXiv:0503053[hep-ex]
40. Nitta K *et al* 2004 Nucl. Instr. and Meth. **A535** 147
41. Freedman D Z 1974 Phys. Rev. **D9** 1389
42. Akimov D *et al* 2017 arXiv:1708.01294v1[nucl-ex]
43. Bionta R *et al* 1983 Phys. Rev. Lett. **51** 27
44. Hirata KS *et al* 1988 Phys. Rev. **D38** 448
45. Fukuda S *et al* 2003 Nucl. Instr. and Meth. **A501** 418
46. Itow Y *et al* arXiv:0106019[hep-ex] and <http://www2.phys.canterbury.ac.nz/~jaa53/presentations/Kato.pdf>
47. de Bellephon A 2006 *et al* arXiv:0607026[hep-ex]
48. Gerigk F *et al* *Conceptual Design of the SPL II* CERN Yellow Report CERN 2006 -006
49. Abe K *et al* (Hyper-Kamiokande Working Group) 2011 arXiv:1109.3262v1[hep-ex]
50. The NSF multi-disciplinary initiative for a deep underground laboratory. <http://www.lbl.gov/nsd/homestake>

51. Boger J *et al* 2000 Nucl. Instr. and Meth. **A449** 172
52. Watanabe H *et al* 2008 arXiv:0811.0735v2[hep-ex]
53. Aynutdinov V *et al* 2008 Nucl. Instr. and Meth. **A588** 99
54. Aslanides E *et al* (ANTARES Collaboration) 1999 arXiv:9907432[astro-ph]
55. Carr J (ANTARES Collaboration) 2008 Nucl. Instr. and Meth. **A588** 80
56. Belias A (2007) in Proceedings of the First workshop on Exotic Physics with Neutrino Telescopes, EPNT06, page 97 arXiv:0701333[astro-ph]
57. Simeone F (On behalf of the NEMO Collaboration) 2008 Nucl. Instr. and Meth. **A588** 119
58. Ackermann M *et al* 2005 Astropart. Phys. **22** 339
59. Aartsen M G *et al* (IceCube Collaboration) 2017 JINST **12** P03012
60. Aartsen M G *et al* (IceCube-Gen2 Collaboration) 2014 arXiv:1412.5106v2[astro-ph.HE]
61. Aartsen M G *et al* (IceCube-Gen2 Collaboration) 2017 J.Phys. **G44** 054006
62. Askaryan G 1962 *Soviet Physics JETP-USSR* **14** (2) 441
63. Gorham P W *et al* 2009 Astropart. Phys. **32** 10
64. Gorham P W *et al* 2011 Astropart. Phys. **35** 242
65. Allison P *et al* 2016 Phys. Rev. **D93** 082003
66. Barwick S W *et al* 2014 arXiv:1410.7369[astro-ph]
67. Wissel S A *et al* 2016 Published in PoS ICRC2015 1150
68. Hankins T H 1996 MNRAS **283** 1027
69. Bray J D *et al* 2015 Astropart. Phys. **65** 22
70. Gorham P W *et al* 2004 Phys. Rev. Lett **93** 041101
71. Buitink S *et al* 2010 Astron.Astrophys. **521** A47
72. Adrian-Martinez S *et al* (KM3Net Collaboration) 2016 J. Phys. **G43** 084001 and arXiv:1601.07459v2[ astro-ph.IM]
73. Avrorin A D *et al* 2014 Nuc. Instrum. Meth. **A742** 82
74. Ball A E *et al* 2007 Eur. Phys. J. **C49** 1117
75. Acquistapace G *et al* 1998 CERN Yellow Report 98-02, INFN-AE-98-05
76. Amerio S *et al* 2004 Nucl. Instr. and Meth. **A527** 329
77. Amoroso S *et al* 2004 Nucl. Instr. and Meth. **A516** 68
78. Adams C *et al* 2018 arXiv:1802.08709v2[ physics.ins-det]
79. Rubbia A 2004 arXiv:0402110[hep-ph] and <http://neutrino.ethz.ch/GLACIER/>
80. Badertscher A 2012 arXiv:1204.3530v3[physics.ins-det]
81. Manenti L (ProtoDUNE Collaboration) 2017 arXiv:1705.05669v2[physics.ins-det]
82. <http://t692.fnal.gov/> ArgoNeut: Mini LAr TPC Exposure to Fermilab's NuMI Beam
83. Acciarri R *et al* 2015 arXiv:1503.01520v1[physics.ins-det]
84. A Proposal for a New Experiment Using the Booster and NuMI Beamlines: MicroBooNE 2007 Fermilab Proposal P974. and Acciarri R *et al* 2017 J.Inst **12** P02017
85. The MicroBooNE Collaboration 2017 MicroBooNE Public Note <http://microboone.fnal.gov/wp-content/uploads/MICROBOONE-NOTE-1026-PUB.pdf>
86. He K *et al* 2015 arXiv:1512.03385v1[cs.CV] and Ronneberger O *et al* 2015 arXiv:1505.04597v1[cs.CV]
87. Bonesini M (WA104 Collaboration) 2015 J. Phys.:Conf.Ser. **650** 012015
88. Acciarri R 2017 *et al* arXiv:1601.02984v1[ physics.ins-det]
89. Heise J 2017 arXiv:1710.11584v1[ physics.ins-det]
90. Badertscher A *et al* 2005 Nucl. Instr. and Meth. **A555** 294
91. Holder M *et al* 1978 Nucl. Instr. and Meth. **148** 235
92. Harris D A *et al* 2000 Nucl. Instr. and Meth. **A447** 377
93. Michael D G *et al* 2006 Phys. Rev. Lett. **97** 191801 and references therein
94. De Winter K *et al* 1989 Nucl. Instr. and Meth. **A278** 670
95. Geiregat D *et al* 1993 Nucl. Instr. and Meth. **A325** 92
96. Ushida N *et al* 1986 Phys. Rev. Lett. **57** 2897
97. Eskut E *et al* (1997) Nucl. Instr. and Meth. **A401** 7
98. Aoki S *et al* 1990 Nucl. Instr. and Meth. **B51** 466
99. Onengut G *et al* 2005 Phys. Lett. **B613** 105

100. Ushida N *et al* (E531 Collaboration) 1988 Phys. Lett. **B206** 375
101. Kodama K *et al* 2002 Nucl. Instr. and Meth. A **493** 45 and references therein
102. Guler M *et al* CERN SPSC 2000-028 SPSC/P318 LNGS P25/2000
103. Agafonova N *et al* (OPERA Collaboration) 2015 Phys. Rev. Lett. **115** 121802
104. Agafonova N *et al* (OPERA Collaboration) 2014 Eur. Phys. J. **C74** no.8 2986
105. Altegoer J *et al* 1998 Nucl. Instr. and Meth. A **404** 96
106. Ellis M, Soler FJP 2003 J. Phys. G: Nucl. Part. Phys. **29** 1975
107. Lindner T, 2008 Status of the T2K 280 m Near Detector. arXiv:0810.2220v1[hep-ex]
108. 2004 Proposal to Perform a High-Statistics Neutrino Scattering experiment Using a Fine-grained Detector in the NuMI Beam arXiv:0405002v1[hep-ex]
109. Cleveland B T *et al* 1998 Astrophys. J. **496** 505
110. Hampel W *et al* 1999 Phys. Lett. **B447** 127
111. Aburashitov J N *et al* 1999 Phys. Rev. **D60** 055801
112. Altmann M *et al* 2005 Phys. Lett. **B616** 174
113. Altmann M *et al* 2000 Phys. Lett. **B490** 16
114. Hampel W *et al* 1998 Phys. Lett. **B420** 114
115. Hampel W *et al* 1998 Phys. Lett. **B436** 158
116. Haidt D 1994 Nucl. Phys. (Proc. Suppl.) **B36** 387
117. Hasert FJ *et al* 1974 Nucl. Phys. **B73** 1
118. Barish SJ *et al* 1974 Phys.Rev.Lett. **33** 448
119. Brand C *et al* 1976 Nucl. Instrum.Meth. **136** 485
120. H. Foeth *et al* 1987 Nucl. Instrum. Meth. **253** 245
121. Baker NJ *et al* 1989 Phys. Rev. **D40** 2753
122. Zucchelli P 2002 Phys. Lett.B **532** 166
123. Autin B *et al* 2003 J. Phys. G: Nucl. Part. Phys. **29** 1785
124. *European Isotope Separation Online* <http://ganinfo.in2p3.fr/eurisol/>
125. Rubbia C, Ferrar A, Kadi Y and Vlachoudis V 2006 Nucl. Instrum. Meth. **568** 475
126. Geer S 1998 Phys. Rev. D **57** 6989
127. De Rujula A, Gavela M B and Hernandez P 1999 Nucl. Phys. B **547** 21
128. Zisman M S (For the ISS Accelerator Working Group) 2008 J. Phys. Conf. Ser. **110** 112006
129. Abe T *et al* 2009 J. Inst. **4** T05001
130. Cervera-Villanueva A, MIND performance and prototyping, in the proceedings of the 9th International Workshop on Neutrino Factories, Superbeams and Betabeams-NuFact 07, editors O. Yasuda, C. Ohmori and N. Mondal American Institute of Physics, p 178.
131. Autiero D *et al* arXiv:0305185[hep-ph]
132. Bennett J *et al* CERN-INTC 2004-16
133. Kaplan D (For the MAP and MICE Collaborations) 2013 arXiv:1307.3891v1[physics.acc-ph]
134. Kochemirovskiy A *et al* 2016 in CNUM616-08-21 NuFact16 in Quy Nhon, Vietnam <http://Vietnam.in2p3.fr/2016/nufact>
135. Aberle C *et al* (The DAE $\delta$ ALUS Collaboration) 2013 physics.acc-ph arXiv:1307.2949v1 and Alonso J R 2016 arXiv:1611.03548v1[physics.acc-ph]

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 9

## Nuclear Emulsions



Akitaka Ariga, Tomoko Ariga, Giovanni De Lellis, Antonio Ereditato,  
and Kimio Niwa

### 9.1 Introduction

Among all tracking devices used in particle physics, nuclear emulsion particle detectors feature the highest spatial resolution in measuring ionizing particle tracks. Emulsions have contributed to outstanding achievements and discoveries in particle physics. Although there was a period of decline of the emulsion technique, the interest in the technique has moved into the front line of physics research because of the advances in digital read-out by high-speed automated scanning and the continuous development of emulsion gel design. In particular, they are unsurpassed for the topological detection of short-lived particles and for specific applications in neutrino physics and other emerging fields. Indeed, a huge potential of emulsion detectors in applied research will be shown in this study. In this chapter, we will mainly focus on developments in experimental techniques for particle physics and briefly present a selection of the main experimental results.

---

A. Ariga · A. Ereditato

Laboratory for High Energy Physics - Albert Einstein Center for Fundamental Physics, University of Bern, Bern, Switzerland

e-mail: [akitaka.ariga@hep.unibe.ch](mailto:akitaka.ariga@hep.unibe.ch); [antonio.ereditato@cern.ch](mailto:antonio.ereditato@cern.ch)

T. Ariga

Faculty of Arts and Science, Kyushu University, Fukuoka, Japan

e-mail: [tomoko.ariga@cern.ch](mailto:tomoko.ariga@cern.ch)

G. De Lellis (✉)

Dipartimento di Fisica, Napoli, Italy

e-mail: [Giovanni.de.Lellis@cern.ch](mailto:Giovanni.de.Lellis@cern.ch)


K. Niwa

Graduate School of Science, Nagoya University, Nagoya, Aichi, Japan

e-mail: [niwa@flab.phys.nagoya-u.ac.jp](mailto:niwa@flab.phys.nagoya-u.ac.jp)

A nuclear emulsion comprises a large number of small silver halide crystals, uniformly dispersed in gelatine. Each crystal has a typical diameter of 200 nm and works as an independent detection channel, which results in a very high detection channel density of  $O(10^{14})$  channels/cm<sup>3</sup> in emulsion detectors. This makes emulsion detectors unique as particle detectors. The latest knowledge of the general photographic process is described in [1]. Herein, we discuss the detection principle of nuclear emulsions for ionizing particles.


The recent nuclear emulsion is made from silver bromide with a small fraction of iodide ( $AgBr_{1-x}I_x$ ,  $x$  being the fraction of iodide, about a few mol%). The crystal structure of AgBr used for nuclear emulsions is face-centred cubic, and its shape is octahedral, as shown in Fig. 9.1. An AgBr crystal has a band gap of 2.684 eV. When a charged particle passes through the crystal, electrons in the valence band are transferred to the conduction band. Owing to shallow electron traps of 21–25 meV, the electrons diffuse inside the crystal until they are trapped in one of the sensitisation centres located at the surface of the crystal (electronic process). The sensitisation centre is artificially created via chemical sensitisation (e.g. sulphur-and-gold sensitisation), which is positively charged at the initial stage and works as an electron trap. The sensitisation centre, which traps an electron, is negatively charged; therefore, it attracts interstitial silver ions, which are ions migrating in the crystal lattice. The silver ion reacts with the trapped electron and forms a single silver atom ( $Ag^+ + e^- \rightarrow Ag$ , ionic process). The sensitisation centre is again positively charged, being ready to trap an electron. These electronic and ionic processes are repeated several times to form an aggregate of silver atoms,  $Ag_{n-1} + e^- + Ag^+ \rightarrow Ag_n$ , deepening its energy level. The energy level of an aggregate equal to or larger than  $Ag_4$  is sufficiently deep to be “developable”, and the sensitisation centre at this stage is called the “latent image centre”. This signal is chemically amplified during the development procedure. The emulsion film is soaked in a developing solution, namely a reduction chemical. The above-mentioned



**Fig. 9.1** Left: silver bromide crystals (0.2 μm linear size), as seen with an electronic microscope. Right: the track left by a minimum ionizing particle (10 GeV  $\pi^-$ ) in nuclear emulsions; about 36 grains/100 μm are detected. Compton electrons of approximately 100 keV are also visible on right-bottom of the view



electronic and ionic processes are repeated by receiving electrons from the reducer through the latent image centre because it is a deep electron trap. This repetition lasts until all the crystals are reduced. The reaction is expressed as follows:




where Red and Ox are the developing agent and the oxidized developing agent, respectively;  $n$  is the number of ions and  $m$  is the number of protons produced. Thus, a metallic silver filament remains at the position of the crystal with a latent image centre, whereas crystals without latent image centres remain unchanged. The gain of this amplification is very high,  $O(10^8)$ . After washing out the remaining AgBr crystals via the fixing procedure, particle tracks are ready to be observed under the microscope, as shown in the right image of Fig. 9.1.

The detection efficiency of a single crystal for minimum ionizing particles (MIP) is about 0.17 [2]. The sensitivity of nuclear emulsions is translated into the number of grains per unit length. A typical emulsion has a sensitivity of 30–50 grains per 100  $\mu\text{m}$  along the particle trajectory for minimum ionizing particles. Apart from the crystal size and chemical sensitisation, the sensitivity scales with the volume occupancy of the AgBr crystals with respect to the total volume of the emulsion layer, which ranges from 30 to 55%. The number of grains is proportional to the ionization power of the particle, which allows the measurement of local energy deposition ( $dE/dx$ ) of each track. The random noise, so called “fog”, is due to several reasons, such as thermal noise, gelatine impurity and over-sensitisation. In general, a fog density of  $<5$  grains/10- $\mu\text{m}$ -cubic is considered acceptable. In the process of producing nuclear emulsion as detectors, emulsion layers with thicknesses of 10–300  $\mu\text{m}$  are formed on a glass or plastic base. To track high-energy particles ( $>100$  MeV), a double-side coated emulsion film with a 50- $\mu\text{m}$ -thick emulsion layer on either side of a 200- $\mu\text{m}$ -thick plastic base is often employed. To observe both emulsion layers across the plastic base with optical microscopes, the plastic base material should not have double refraction; e.g. triacetyl cellulose and polymethylmethacrylate are appropriate.

The RMS resolution of a one-dimensional detector with a segmentation pitch of  $D$  is  $D/\sqrt{12}$ . Assuming that the silver halide crystal shape is approximately spherical, the resolution with a crystal diameter  $D$  is  $\sqrt{\pi}D/8$  (RMS). For example, this gives 44 nm for an emulsion with 200-nm-diameter crystals. In reality, these values are slightly larger owing to the delta-ray component. A measured resolution of 50 nm (RMS) was reported for an emulsion film with a 200-nm crystal size by using high-energy particles [3], as shown in Fig. 9.2. The one-dimensional intrinsic angular resolution of a double-sided emulsion film with 200-nm-diameter crystals and a base thickness of 200  $\mu\text{m}$  is therefore 0.35 mrad. Owing to the excellent position and angular resolution, one can build a vertex detector, while using a sampling calorimeter to reconstruct electromagnetic showers and also measuring the momentum of particles by the multiple Coulomb scattering, which will be discussed in Sect. 9.2. Nuclear emulsion detectors may be coupled with electronic detectors to add timing information and/or muon identification. Since emulsion

**Fig. 9.2** Distribution of the distances between grains and straight-line fits to the tracks of minimum ionizing particles, showing the emulsion intrinsic spatial resolution [3]



detectors can be produced in many different sizes and shapes, there is a large variety of possibilities for a hybrid detector system, depending on physics goals, which we shall discuss in next sections. One double-sided film with a size of  $10\text{ cm} \times 10\text{ cm}$  has approximately 1-cc emulsion layer and comprises  $O(10^{14})$  detection channels, as mentioned above. After chemical treatment, such a huge number of channels has to be read-out for physics analyses. This is the task of automated scanning microscopes, whose implementation is of fundamental importance in modern experiments making use of nuclear emulsions. This will also be discussed in the following sections.

## 9.2 Early Times of the Technique and the Emulsion Cloud Chamber

Thorough reviews of the basic properties and early applications of nuclear emulsions can be found in [4, 5]. The first notable examples of the use of photographic emulsion (plates) are the discovery of radioactivity by Becquerel in 1896 [6] and the measurements by Kinoshita [7], who in 1910 found records of alpha-particle radiation detected as tracks by means of optical microscopes. The emulsion technique greatly improved during the 1930s and 1940s thanks to the group of Bristol University led by Powell. He developed electron-sensitive nuclear emulsions produced by ILFORD and KODAK [8]. Powell and his group had further developed and greatly extended the seminal work of Marietta Blau. She is also known for the development of thick emulsions by a two-bath method [9].

The thickness of emulsions increased from the original  $50\text{--}100\text{ }\mu\text{m}$  used in 1946 to  $600\text{--}1000\text{ }\mu\text{m}$ . Even with a  $500\text{ }\mu\text{m}$  thickness, a large part of the tracks of charged particles originated in the emulsion were not contained. Although an exceptional attempt to process a  $2000\text{ }\mu\text{m}$  thick emulsion was reported in 1950 [10], the


difficulties of processing increased rapidly with the thickness and new difficulties appeared in the visual inspection, due to the larger scattering of light in the emulsions and the loss of optical contrast.

Plates were arranged in pairs with emulsions face to face, thus doubling the effective thickness. In 1952 a new approach was established [11–13]. Once a batch of plates was produced, the emulsions were stripped from the glass and packed together to form an almost solid sensitive mass, named stack. After exposure, the emulsions were dipped in a solution of glycerine with gelatine and then made to adhere to specially prepared glass plates. The use of a penetrating X-ray beam defined a reference frame to connect consecutive emulsion layers. With such a procedure, tracks of single particles could be quickly followed through the successive emulsions of a stack. The use of stripped emulsions became popular and allowed to make important contributions to many experiments in particle physics, as we will see in the following.

Photographic plates with  $600\mu\text{m}$  thickness were manufactured by means of newly produced emulsion gel able to record and detect the passage of ionizing particles. In parallel, dedicated microscopes were developed to observe and measure the particle tracks. With these emulsion detectors exposed to cosmic-rays Powell solved in 1947 the mystery of the Yukawa meson, by detecting the pion through its decay into a muon [14–16]. A picture of this decay as seen in nuclear emulsions is shown in Fig. 9.3. Powell was awarded the Nobel Prize for physics in 1950 for his discovery made possible by using nuclear emulsions. In the presentation of the Nobel Committee, the simplicity of the apparatus used to make such a discovery was underlined.

Few years later, in 1955, exotic hyper-nuclei were also identified by nuclear emulsions [17]. In a large balloon experiment in 1960, a 70 l emulsion chamber called “Bloury Stack” was exposed to high-altitude cosmic-rays to study their nature and the features of the induced high-energy interaction phenomena [18]. However, a crucial limitation of the technique (for those days) was met: due to the lack of scanning power, the experiment could not achieve the expected results.

A major breakthrough in the emulsion technique was the introduction of the so-called Emulsion Cloud Chamber (ECC) detector [19]. With the ECC a drastic



**Fig. 9.3** Photomicrographs of one example of  $\pi \rightarrow \mu$  decay taken from [15]


change in the detector design philosophy occurred: emulsions became a high-resolution tracking detector with three-dimensional reconstruction capabilities, rather than a visual and volume detector. This is obtained by sandwiching emulsion films or plates with passive material layers, usually made of plastic or metal plates. Today, we would call such a detector a very finely subdivided sampling-calorimeter, by means of which all charged tracks originating from the shower are reconstructed in space with high resolution. In the ECC, emulsion films are placed perpendicular to the incoming particles, so acting as a tracking detector featuring high spatial resolution (up to 1  $\mu\text{m}$ ).

The first design of the ECC consisted of a sandwich of brass plates and thin emulsion films. This type of detector was first developed by Kaplon and used to study heavy primaries in cosmic-ray interactions [19]. ECC detectors were applied to the study of the cosmic-ray spectrum and to very-high energy interaction processes. Nishimura, in particular, proposed the cascade shower analysis method to measure the energy of interacting  $\gamma$ -rays and predicted the capability of this detector to regulate the development of electron showers by selecting passive material plates on purpose [20].

Niu developed double-sided emulsion plates in which the sensitive emulsion layer is deposited on either side of a plastic substrate (see e.g. [21]). For this purpose FUJI developed a special 800  $\mu\text{m}$  thick plastic base to allow gel pouring on both sides of the layer. The emulsion layers were 50  $\mu\text{m}$  thick. With this new film design, two problems had to be solved: the availability of a plastic base with optical properties compatible with that of nuclear emulsions, and of a high-power objective lens with a working distance longer than 1 mm. The first problem was overcome with meta-acrylic (lucite) plates, the second was solved thanks to the efforts of Tiyo-da Optical Co. The use of a plastic base between the two emulsion layers allows a precise measurement of the track angle by connecting those grains closest to the base. These points indeed are not affected by distortions. The long lever arm available with such a thick base improves the angular resolution up to 1 mrad.

The ECC opened the way to a series of important experiments of large size, thanks to the use of the dense metal plates allowing the realization of large-mass detectors with unprecedented space resolution. For the study of high-energy cosmic-rays (10 TeV) and the determination of their power law spectrum, we mention in particular the Chacaltaya experiment [22] that allowed the study of the central core of air showers, and the relatively large-size Mt. Fuji experiment [23]. For even higher cosmic-ray energies (1000 TeV and more), the RUNJOB [24] and JACEE [25] experiments studied the spectrum of primary heavy ions.

As said above, the analysis methods of ECC events are based on the reconstruction of all tracks produced following a primary interaction, likely occurring in the dense passive material. Space angles are measured for all track segments. Shower reconstruction and identification (electromagnetic or hadronic) can be performed on the basis of the topological features of the shower. In the same way, one can also reconstruct particle decays. In addition to the topological studies, powerful kinematical analyses can be conducted with ECC detectors by exploiting Multiple



**Fig. 9.4** Schematic drawing of the first evidence for the production and decay of short-lived “X particles” (charmed particles) in cosmic-ray interactions

Coulomb Scattering and emulsion ionization measurements, which can lead to surprisingly accurate measurements of particle momenta and particle identification.

A notable example is given by the Niu’s discovery of the so-called X-particles in 1971 [21, 26]. Figure 9.4 shows a sketch of the observed topology for one event where two charged particles produced in the cosmic-ray interaction show a kink decay-topology. Today we know that this event had to be attributed to charmed meson production and decay. This happened 3 years earlier than the discovery of the  $J/\Psi$  particle by the groups of Richter [27] and Ting [28]. During those 3 years, several papers referring to that event were published by Japanese theorists [29], while there was no comparable response from the western high energy physics community. The reason for that could probably be attributed to the lack of confidence in the emulsion technique felt at that time by western scientists and also to the fact that the community carrying out cosmic-ray studies was quite apart from that employing particle accelerators, at that time not active in Japan. It is worth noting that the main distinctive features of charmed mesons were indicated already in 1974 as a result of the kinematical analysis conducted by Niu and coworkers [30], who had also re-analysed ECC data from cosmic-ray exposures carried out many years before. In addition, these authors realized that the lifetimes of charged and neutral charmed hadrons differed by a factor ranging from 2 to 3 [31].

ECC detectors allowed to design hybrid experiments combining emulsions and electronic detectors, the latter mainly used for two purposes: (1) to provide time resolution to the emulsion stack (trigger signal), (2) to pre-select the region of interest for the event occurring in the ECC, indicating the place where to start the emulsion scanning. The first hybrid experiments employed semi-automated video-

camera systems to read out the emulsion tracks and reconstruct three-dimensional vectors by measuring  $X, Y, \theta_X, \theta_Y$ , with  $Z$  the emulsion depth. Computers were only used to assist an operator in performing the track measurements and to provide the micro-metric movement of the microscope stage. Relative alignment was performed by fiducial X-ray marks combined with the precision measurement of the film edge positions. Typical thickness of the double-sided emulsion plate was 1 mm or larger, so allowing to follow-up tracks with a given angle *w.r.t.* the emulsion plane by varying the focal plane of the objective lenses. The video-camera was used to grab the image from the objective lens with a rather time-consuming procedure. An operator had to manually adjust the video-image on the visually detected track, while dark spots could be automatically detected. The TV screen also allowed to run graphic tools for measuring track positions and angles.

The mechanical stability of the ECC sandwich was ensured by a vacuum packing paper known as “origami”, also required to isolate the emulsion films from the external light, humidity and polluting gases. Plate-to-plate alignment was performed by X-ray lines and/or X-ray spots typically from a  $^{55}\text{Fe}$  source. The association between the ECC and the electronic detectors was accomplished by joining particle tracks, better if of high momentum and hence less affected by Multiple Coulomb Scattering. In this respect, the idea of using interface emulsion plates in between the ECC module and the electronic detectors has proven to be very effective. These interface films were called Changeable Sheets (CS) because they were frequently replaced during the physics run in order to limit the integration of background tracks and to easily identify tracks found in the electronic detectors. This concept was first applied in the E531 experiment at Fermilab [32], as we will see in the following, and it is presently being used in several applications also for large scale ECCs.

### 9.3 Notable Experiments Employing Nuclear Emulsions

During the 1970s, emulsion detectors of increasing mass and complexity were developed for applications to particle physics experiments conducted at particle accelerators with experimental setups also including electronic detectors (hybrid experiments). Emulsions are often employed as active targets with high space-resolution, and the electronic detectors, namely trackers, calorimeters and spectrometers, are used to pre-select or trigger specific events in the emulsions and to complement the kinematical information of the events.

In early experiments with accelerators, nuclear emulsions were coupled to spark and bubble chambers in order to reduce the total scanning time. We recall here the observation of the decay of a charmed particle produced in a high-energy neutrino interaction in a Fermilab experiment [33]. The latter was performed in the wide-band neutrino beam produced by 400 GeV protons, by using a detector made of spark chambers placed downstream of nuclear emulsion stacks. Stacks containing altogether 16 l of ILFORD X5 emulsion made up of pellicles of  $20\text{ cm} \times 8\text{ cm} \times 0.6\text{ mm}$  dimensions were placed in association with a double wide-gap spark

chamber followed by a detector of electromagnetic showers and a muon identifier. A veto counter upstream discriminated against interactions in the emulsion produced by charged particles. About 250 neutrino interactions were predicted by the spark chamber. Given its vertex position resolution, a volume of about  $0.7\text{ cm}^3$  was visually scanned around the prediction for about one third of the events; 16 of them were located and fully reconstructed in the emulsions and one of them was found with a topology consistent with that of charm.

A search for charmed particles in neutrino interactions was carried out at CERN in 1977 with stacks of nuclear emulsions placed in front of the entrance window of the Big European Bubble Chamber (BEBC) [34], filled with liquid hydrogen and placed in a magnetic field of 3.5 T. A veto-coincidence counter system was added in front of BEBC for this purpose. The emulsion stacks were made of 3150 pellicles of ILFORD emulsion, each  $600\text{ }\mu\text{m}$  thick. The quality of the emulsion as well as the high level of muon track background precluded any systematic scanning along the track. A “surface” scan was therefore carried out for the bulk of the events with  $200\times$  and  $300\times$  objective lenses, over an emulsion volume centred on the predicted vertex position of  $5\times 31\text{ mm}^2$  for 7 plates. A total of 206,000 BEBC pictures were analysed, leading to 935 neutrino interaction vertices inside the emulsion, 523 of which identified as charged current events. After kinematical and topological cuts, 169 charged current interactions were selected, 8 of them being identified as neutrino-induced charmed particles. The experiment reported the first direct observation of a charmed baryon decay [35] and of a neutral charmed particle [36].

The E531 experiment [32] was proposed in 1978 at Fermilab to study the properties of charmed particles and their production mechanism in neutrino interactions [37]. The neutrino beam was produced by 350 GeV protons for a first exposure ( $7.2 \times 10^{18}$  protons on target) and by 400 GeV protons for the second one ( $6.8 \times 10^{18}$  protons on target). The overall beam composition was 92.3%  $\nu_\mu$ , 7.0%  $\bar{\nu}_\mu$ , 0.5%  $\nu_e$  and 0.2%  $\bar{\nu}_e$ . The active neutrino target was made of nuclear emulsions where short-lived particles were detected with micrometer accuracy. The decay products were then measured by means of an electronic spectrometer, thus making E531 the first hybrid particle physics experiment.

The emulsion target consisted of 22.6 l in the first run and of 30 l in the second one; it was made of modules composed of plates with  $300\text{ }\mu\text{m}$  emulsion layers coated on both sides of  $70\text{ }\mu\text{m}$  thick polystyrene foils. Downstream of the emulsion modules, two large lucite plates  $800\text{ }\mu\text{m}$  thick, coated on both sides with  $75\text{ }\mu\text{m}$  emulsion layers, acted as interface emulsion films, so establishing the new detector concept of the Changeable Sheets (CS). Tracks reconstructed by electronic detectors were first searched for in these interface films and then followed back in the bulk target up to the neutrino interaction vertex. The CS were replaced every 2 or 3 days of data taking in order to limit the number of accumulated background tracks that would have affected the efficiency of finding the interaction vertex in the target.

Downstream of the target, a magnet equipped with high-resolution drift chambers provided the track prediction in the CS with an accuracy of about  $150\text{ }\mu\text{m}$  and 1 mrad. A time-of-flight detector made of two scintillator planes located 2.7 m



apart yielded a time resolution better than 1 ns. The setup was complemented by a lead glass array and a hadron calorimeter followed by a muon spectrometer. Three thousand eight hundred eighty-six neutrino interactions were located in the fiducial volume of the target. One hundred and twenty-two events were tagged by the presence of a secondary vertex in the target, 119 induced by neutrinos and 3 by anti-neutrinos. Events with a candidate charmed hadron in the final state were studied in detail in order to detect the presence of heavily ionizing particles (baryons) and fully reconstruct the kinematics at the decay vertex. Among those events, 57 were classified as  $D^0$  candidates.


The analysis of the charmed hadrons is reported in [38]. Re-analyses of these results were conducted later and removed some biases present in the original studies ([39, 40]). The result on the cross-section measurements are given in [41]. In this paper, the observation of one event with the  $D^0 - \bar{D}^0$  topology was reported, interpreted as associated charm production in neutral current interactions. The lifetime of charmed particles was extensively studied by E531 [42]. Limits were also set on  $\nu_\mu \leftrightarrow \nu_\tau$  oscillations [43].

After the discovery of the  $b$  quark in 1977 [44], experiments with nuclear emulsions aimed at the direct observation of the production and decay of  $B$  flavored hadrons. A successful search was first performed by the WA75 experiment at CERN by using a  $\pi^-$  beam of 350 GeV [45]. Eight hundred and one of nuclear emulsion, in the form of double-coated plates and stripped pellicles, was exposed in 1983 and 1984. The emulsion stacks were placed both parallel and perpendicular to the beam, so exploiting the advantages of both approaches. Emulsions held perpendicular to the beam in vertical position can in fact tolerate higher track densities, while those placed parallel are more sensitive to short particle lifetimes.

The emulsion was delivered in gel form by FUJI (75 l) and ILFORD (5 l) and the pouring was done in a facility set-up at CERN [46]. Each vertical stack was made of 25 double-coated plates (330  $\mu\text{m}$  thick emulsion, poured on both sides of a 70  $\mu\text{m}$  thick Lexan support), 25  $\times$  25  $\text{cm}^2$  wide and packed in vacuum. The horizontal stacks were made of 60 stripped emulsion pellicles, 11  $\text{cm} \times$  4  $\text{cm}$  (4  $\text{cm}$  along the beam) and 600  $\mu\text{m}$  thick, piled-up and clamped between two rigid Perspex plates. The processing of the films was carried out in Nagoya for double-coated plates, in Rome for pellicles, and at CERN for both. After processing, each double-coated plate was cut into 64 squares, 3  $\times$  3  $\text{cm}$  in size, so-called mini-modules. Twenty-five squares of a module were then stuck, in sequence, on a single Lucite foil. With such a technique, the corresponding areas of consecutive emulsion plates were adjacent, thus reducing the time needed to follow a track through the stack [47]. The size of the beam was so small that it was necessary to move the target during each beam spill in order to have a uniform irradiation, thus introducing the concept of target mover. The WA75 experiment observed one event [48], schematically depicted in Fig. 9.5 as recorded in the pellicles, where both  $B$  hadrons are observed to decay into a charmed particle. The experiment also made the first observation of the purely muonic  $D_s$  decay measuring the decay constant  $f_{D_s}$  [49].

The Fermilab E653 experiment [50] was designed to measure the lifetime of  $B$  hadrons. This detector was an extension of the hybrid emulsion technique developed





**Fig. 9.5** Schematic drawing of the first hadro-produced  $B\bar{B}$  pair event observed in nuclear emulsions by the WA75 experiment

for the E531 experiment and was optimized for a hadron beam. In fact, while in the E531 neutrino experiment charm was produced in one out of twenty charged current interactions, only one hadronic interaction in a thousand produces charm, and one in a million bottom. Thus, larger discrimination against non-heavy quark background was required to limit the emulsion scanning load. To achieve this, a high-resolution electronic spectrometer was placed downstream of the emulsions. Moreover, in order to cope with the large number of candidate events, the emulsion analysis required the development of computer-aided microscope techniques [51]. Events reconstructed in the spectrometer with a muon of high transverse momentum ( $p_{\perp} > 1.5 \text{ GeV}/c$ ) were selected for scanning in the emulsion. In the first run of 1985 a 800 GeV proton beam was used, mainly aiming at charm production. In a second run in 1987 a 600 GeV negative pion beam was exploited for the study of  $B$  mesons. Two types of target modules were employed; 55 were “vertical” and the rest “horizontal”. In the first run, vertical modules were exposed to  $1.5 \times 10^5$  protons/cm<sup>2</sup> and the horizontal ones to  $0.8 \times 10^5$  protons/cm<sup>2</sup>. The second-run exposures corresponded to  $3.0 \times 10^5$  pions/cm<sup>2</sup> and  $1.0 \times 10^5$  pions/cm<sup>2</sup>, respectively for the two orientations.

Forty nine and fifty six target modules were exposed, respectively in the first and second run, for a total of 71 l of FUJI nuclear emulsion. Each vertical module consisted of 20 thick emulsion plates (330  $\mu\text{m}$  emulsion layer on each side of a 25 cm  $\times$  25 cm  $\times$  70  $\mu\text{m}$  polystyrene plate) and a thin film (70  $\mu\text{m}$  emulsion layer on either side of a 25 cm  $\times$  25 cm  $\times$  500  $\mu\text{m}$  lucite plate). The thin film was separated from the main block of thick plates by a 10 mm thick honeycomb, the latter combination being considered as the analysing region, while thick plates made the target region.

The emulsion modules were mounted on a target mover and displaced through the beam during the slow spill, in order to have a uniform exposure. The movement of the target was digitally controlled and the positioning encoding system granted an accuracy of 10  $\mu\text{m}$  [52]. 18 silicon microstrip planes in the electronic vertex detector were located 5.7 cm downstream of the emulsion target. Secondary vertices were reconstructed by the silicon planes with typical resolutions of 6  $\mu\text{m}$  transverse

to and 350  $\mu\text{m}$  along the beam direction. The total fiducial decay region for bottom particles including emulsions and silicon planes was 12 cm long.

The emulsion analysis procedure first located the primary vertex. Six thousand five hundred forty-two events were selected within the fiducial volume of the emulsion and for all but 9 the primary vertex was found thanks to the excellent performance of the electronic detectors. The majority of the events were discarded by requiring a stringent angular agreement (2 mrad for tracks with a slope within 40 mrad) between the reconstructed spectrometer track and any track at the primary vertex. Three hundred and fifty-nine events in which the muon did not come from the primary vertex were retained for the secondary vertex search. Nine events met the selection criteria for bottom [53]. The  $b$  lifetime was also measured.

At the end of the 1980s, the production of charmed particles from quark-gluon plasma was expected to differ from that due to proton-nucleus interactions [54]. In particular, a large enhancement of the charmed quark pair creation was expected. From the experimental point of view, the major difficulty for charm detection in such nucleus-nucleus interactions came from the very short-path decay in a region close to the primary interaction where the particle density was extremely high. Two studies were carried out at CERN on this subject with emulsions, one within the NA34/2 emulsion-HELIOS programme [55] and the other one within the EMU09 Collaboration [56]. In NA34 [55], the production of charmed particles was detected in 200 GeV/nucleon  $^{16}\text{O}$ -emulsion interactions and its cross-section was measured. Stacks of FUJI gel were exposed vertically to the  $^{16}\text{O}$  beam. Each stack consisted of 8 double-coated plates with a surface of  $25 \times 15 \text{ cm}^2$  and a thickness of 700  $\mu\text{m}$  (70  $\mu\text{m}$  polystyrene base coated on both sides with a 315  $\mu\text{m}$  thick emulsion layer).

In order to study charmed particle production in central interactions of 200 GeV per nucleon  $^{32}\text{S}$  nuclei, the EMU09 Collaboration designed an emulsion-counter hybrid experiment at CERN [56]. The hybrid design was meant to reduce the background from secondary interactions in the emulsion, which would have spoiled the signal with heavier projectiles, differently from the case of  $^{16}\text{O}$ . A thin and pure target was made of 100  $\mu\text{m}$  thick Ag and Pb plates. Two emulsion plates, in the form of tapes, were placed downstream of the target and used as a tracking device, able to detect short-path decay vertices and producing very little secondary activity. The emulsion tape used in this experiment was derived from an Acetate base 200  $\mu\text{m}$  thick and FUJI gel poured on both sides of the base, to obtain 70  $\mu\text{m}$  thick layers. The emulsion analysis speed at that time did not allow to integrate a sufficiently large statistics for such rare events.

Nuclear emulsions have also played an important role in the study of multiquark systems and the quark confinement aspects of QCD. The hybrid emulsion experiment E176 [57] was carried out at KEK by using a 1.66 GeV/c  $K^-$  beam to study double-strangeness nuclei produced via  $\Xi^-$  hyperon capture at rest. Indeed, the  $K^- p \rightarrow K^+ \Xi^-$  interaction produces a  $\Xi^-$ , which at rest may be captured via the process  $\Xi^- p \rightarrow \Lambda\Lambda$ . In the hypothesis that the  $H$ -dibaryon (ssuudd) exist, the double  $\Lambda$  hypernucleus can decay by emitting a  $H$ -dibaryon, in turn decaying into  $\Sigma^- p$  within less than 1 mm from the  $\Xi^-$  stopping point. Unlike old-fashioned emulsion experiments where only emulsion stacks were exposed to  $K^-$  beams [58],


the hybrid design allowed the identification of the  $K^+$  meson and the accumulation of a large statistics. Emulsion stacks were exposed vertically, perpendicular to the beam. Emulsion plates were of two types: 550  $\mu\text{m}$  thick emulsion layers on both sides of a 70  $\mu\text{m}$  thick polystyrene base, and 70  $\mu\text{m}$  layers on both sides of a 500  $\mu\text{m}$  lucite base. Thinner films with thicker base were used to avoid the degradation of the angular resolution due to distortion effects. Three double- $\Lambda$  hypernuclei candidates were observed [57, 59]. However, no conclusive answer was provided on the  $\Lambda$ - $\Lambda$  interaction. With this aim, the E373 experiment at KEK [60, 61] searched for  $S = -2$  nuclei in nuclear emulsion with higher statistics. The apparatus was based on an emulsion-counter hybrid method, where a laser microscope performed the three-dimension graphic processing of the emulsion images, scintillating fiber blocks detected the decay products of strange particles, and a glass capillary tracker filled with liquid scintillator provided precise predictions of the  $\Xi^-$  emission angle and position. The experiment reported the observation of double hypernuclei and the  $\Lambda$ - $\Lambda$  interaction was finally measured [62, 63]. A follow-up experiment is planned for the new J-PARC hadron facility at Tokai, still employing the hybrid detector technique with an emulsion plate stack [64].

## 9.4 Nuclear Emulsion Detectors with Digital Technology

### 9.4.1 Automated Scanning Systems and Analysis Methods

A major breakthrough in the emulsion technique occurred in 1974 when the idea of a tomographic read out of the emulsion plates was introduced by the Nagoya group [65]. In the case of an emulsion layer about 20 times thicker than the focal plane depth, one can take multiple tomographic images by sampling the emulsion layer. Those images can then be superimposed according to a given value of the presumed track slope, looking for space coincidences of the grains. After applying a detection threshold needed to remove the accidental background, a track can be defined. A first implementation of this concept led to the development of a first generation system [47] where 16 tomographic images were superimposed and a TV tube used to grab the image. This concept was developed and successfully applied to the CS emulsion scanning of the E653 experiment at Fermilab [51].

This technique was further developed by the Nagoya group and led to the so-called Track Selector [66]. The TV video was replaced by a CCD camera, yielding to higher stability and better space resolution. An FPGA-based image processor handled the 16 tomographic images of each emulsion plate. The scanning speed was actually limited by the time required for the computer-controlled objective lenses to move to the 16 different focal positions, since for each step some time was needed to damp the stage vibrations. Another limiting factor was the size of the optics field of view. A tracking efficiency as large as 90% was reached, with the main source of noise given by short Compton electron tracks. A scanning system based on a



**Fig. 9.6** Different steps of the emulsion data processing in the net-scan method. On the left plot all base-tracks in 15 films of the volume under study are reconstructed; they participate in the alignment process from which tracks are reconstructed, as shown in the middle plot; on the right plot passing-through tracks are discarded and the interaction vertex is reconstructed

different approach was developed in Salerno [67]. This device exploited a multi-track approach without any angular preselection.

A further important step was the establishment of fully-automatic offline analysis methods, beyond the digitization of the individual tracks around a given angle performed with the Track Selector. This progress was mainly driven by the availability of faster electronics and CCDs and more performing stage mechanics. The so-called net-scan method (described in [68]) developed in Nagoya allowed the reconstruction of tracks by associating all detected track segments regardless of their angle. Obviously, the area over which net-scan could be realistically performed depended on the available scanning speed. The latter was about  $1 \text{ cm}^2$  per h with the UTS system [69] that exploited parallel data processing. The net-scan method allowed complete event reconstruction both at the interaction and decay vertices, precise measurements [70], search for downstream particle decays, momentum determination by Multiple Coulomb Scattering [71, 72], and electron identification by cascade shower analysis [73, 74]. Figure 9.6 shows the different steps of the emulsion data processing in the net-scan method.

### 9.4.2 Applications to Neutrino Experiments

In the early 1990s it became evident that the next generation of neutrino (oscillation) experiments would greatly profit from the use of the dense, high space-resolution emulsions to realize hybrid detectors well suited to the high sensitivity study of short decay topologies (charm,  $\tau$ ) with the possibility of a full reconstruction of the event kinematics, in turn required for background suppression. This approach was indeed motivated and justified by the advances in the emulsion technique in relation to the possibility of handling large quantities of emulsions, and also thanks to the above mentioned progress in the emulsion scanning and offline analysis, which could allow

reducing the analysis time of the emulsions by orders of magnitude as compared to the early times.

The CHORUS detector [75] is a good example of a large hybrid experimental setup combining a nuclear emulsion target with various electronic detectors. The detector was designed to search for  $\nu_\mu \leftrightarrow \nu_\tau$  oscillations in the CERN WANF neutrino beam with high sensitivity. At that time, a relatively massive  $\nu_\tau$  was a preferred candidate to explain the Dark Matter of the Universe. Since charmed particles and the  $\tau$  lepton have similar lifetimes, the detector was also well suited for the observation of the production and decay of charmed particles.

Also in CHORUS nuclear emulsions acted both as neutrino target and as a high space-resolution detector, allowing three-dimensional reconstruction of short-lived particles. The emulsion target had an unprecedented large mass of 770 kg and was segmented into four stacks, each consisting of eight modules, each in turn composed of 36 plates with a size of  $36 \times 72 \text{ cm}^2$ . Each plate had a  $90 \text{ }\mu\text{m}$  plastic support coated on both sides with a  $350 \text{ }\mu\text{m}$  emulsion layer [76]. Each stack was followed by a set of scintillating fibre tracker planes. Three Changeable Sheets with a  $90 \text{ }\mu\text{m}$  emulsion layer on both sides of a  $800 \text{ }\mu\text{m}$  thick plastic base were used as interface between the fibre trackers and the bulk emulsion. The accuracy of the fibre tracker prediction was about  $150 \text{ }\mu\text{m}$  in position and 2 mrad in the track angle. The electronic detectors downstream of the emulsion target and the associated trackers included a hadron spectrometer measuring the bending of charged particles in an air-core magnet, a calorimeter where the energy and direction of showers were measured and a muon spectrometer.


CHORUS represents a milestone in the history of nuclear emulsions for the size of the target and of the CS, which implied very labor-intensive procedures for emulsion gel production, pouring on the plastic bases, and development conducted in the CERN emulsion laboratory [46], as well as for the first massive use of automated scanning microscopes running in the Japanese and European laboratories of the Collaboration [75].

The operation of the experiment consisted of several steps. It is worth noting that the large-size emulsion target was replaced only once during the entire duration of the experiment, while the CSs were periodically exchanged with new detectors, therefore integrating tracks for a relatively short period. The best time resolution was obviously provided by the electronic detectors. With the CS scanning, the association between electronic detectors and emulsions took place, and tracks with position and angle compatible with that of the electronic trackers' predictions were searched for in the interface emulsions. If found, these tracks were further extrapolated into the bulk emulsion, with a much better resolution, up to the track stopping point, with a procedure called scan-back, consisting in connecting emulsion layers progressively more upstream. After that, a "volume scan" (net-scan) around the presumed vertex was accomplished and repeated for all stopping tracks until the neutrino interaction vertex was found.

In the search for charmed particle decays, a dedicated topological selection was applied to the collected net-scan data. The analysis procedure was complemented by the visual inspection of the selected event candidates, aimed at checking both

primary and secondary vertices making use of the “stack” configuration. Decay topologies could be well separated from ordinary nuclear interactions, since the latter usually exhibit fragments from nuclear break-up or so-called “blobs” from nuclear recoil.

More than 100,000 neutrino interactions were located in CHORUS. The search for oscillations was negative and an upper limit to the oscillation probability was eventually set [77]. CHORUS reported the first observation of the associated charm production in charged current interactions [78]. This first observed event is shown schematically in Fig. 9.7. It represents the production of two charmed particles in a charged current interaction induced by a muon neutrino. Apart from six tracks of high ionization coming from the nuclear break-up and not drawn in the sketch, at the primary neutrino interaction vertex there are two charged tracks: one is the negative muon and the other one, indicated as particle 1, is a charmed hadron. The charged charmed particle shows a 417 mrad kink angle after travelling 1010  $\mu\text{m}$ . The outgoing particle, indicated as particle 2, shows a flight length of 7560  $\mu\text{m}$  and a reinteraction with an outgoing particle (particle 3) of high ionization. In addition to the charged charmed hadron, the decay of a neutral charmed particle is visible 340  $\mu\text{m}$  downstream of the primary vertex. Two particles are generated from the neutral particle decay. The non-planarity of parent and daughter particles rules out the two-body decay and thus both the  $K_S^0$  and the  $\Lambda$  hypotheses for the neutral particle. A kinematical analysis confirmed the interpretation of the event as the associated charm production induced by a muon neutrino in a charged current interaction [78]. An unprecedented statistics of about 2000 fully reconstructed neutrino-induced charmed hadron event vertices was collected. With this statistics, CHORUS measured the  $\Lambda_c$  and  $D^0$  exclusive production cross-section [79] and the double-charm production cross-section in both neutral and charged current interactions [80]. The CHORUS emulsion data also provided an upper limit to the production of charmed pentaquark states [81].

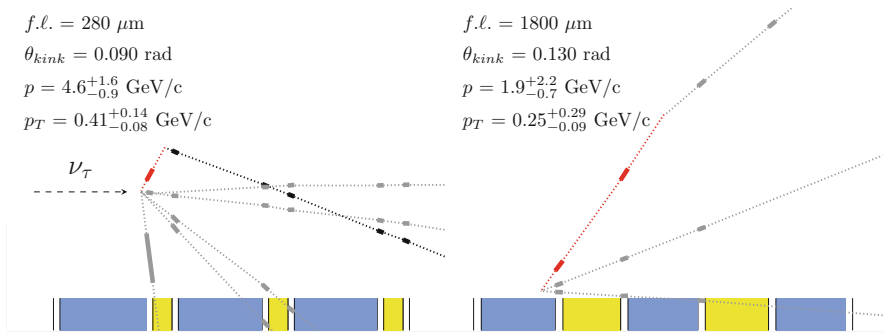


**Fig. 9.7** Schematic drawing of the first neutrino-induced associated charm production event observed in the emulsions of the CHORUS experiment (see the text for explanation)

Higher sensitivity follow-ups of the CHORUS experiment were proposed, with the purpose of increasing by more than one order of magnitude its sensitivity in the measurement of the oscillations (smaller mixing angle). We mention in particular the COSMOS proposal at Fermilab [82]. The use of emulsions as large-surface trackers for the high-resolution measurement of hadron and muon momenta was proposed in [83] and then applied for the proposal of the TOSCA experiment at CERN [84]. Eventually, all those experiments were not realized mainly due to the first strong indications for  $\nu_\mu \leftrightarrow \nu_\tau$  oscillations detected with atmospheric neutrinos, in disappearance mode, in a complementary region of the oscillation parameters.

The DONUT experiment at Fermilab aimed at the first direct detection of  $\nu_\tau$ s, in this case promptly produced in a 800 GeV proton beam dump and not coming from the possible oscillation mechanism as in CHORUS. The experimental apparatus and the detection techniques used in the experiment are described in [68, 85]. The DONUT Collaboration employed an iron/emulsion ECC target able to offer a sufficiently high mass to the interaction of the neutrinos and to provide the detection of the interaction vertex, as well as a clear observation of the short track of the  $\tau$  lepton (up to a few mm) produced in the  $\nu_\tau$  charged current interaction. The ECC was complemented by high-precision fiber trackers to drive the scan back in the emulsions.

The emulsion target eventually integrated a relatively high muon background. In a first analysis, 203 neutrino interactions were located in the ECC target, observing 4  $\nu_\tau$  candidate events with an estimated background of 0.34 events [86]. This represents the first direct detection of the  $\nu_\tau$ . Figure 9.8 shows a display of two candidate events. In the final analysis, 9  $\nu_\tau$  charged-current (CC) events were detected, with an estimated background of 1.5 events, from a total of 578 observed neutrino interactions and were used to estimate  $\nu_\tau$  CC cross section for the first time [87]. The main source of error in measuring the  $\nu_\tau$  cross section was due to the systematic uncertainties, whereas 33% of the relative uncertainty was due to the limited number of detected  $\nu_\tau$  events. Owing to the



**Fig. 9.8** Schematic drawing of two  $\nu_\tau$  induced events measured by the DONUT experiment. The kinks relative to the  $\tau$  decay are visible



lack of accurate measurements of the  $D_s$  differential production cross section, DONUT expressed its  $\nu_\tau$  cross-section measurement as a function of the parameter  $n$ , responsible for the differential production cross section of  $D_s$ , as  $\sigma_{\nu_\tau}^{\text{const}} = 2.51n^{1.52} \times 10^{-40} \text{ cm}^2 \text{ GeV}^{-1}$ . The cross section was estimated to be  $\sigma_{\nu_\tau}^{\text{const}} = (0.39 \pm 0.13(\text{stat.}) \pm 0.13(\text{syst.})) \times 10^{-38} \text{ cm}^2 \text{ GeV}^{-1}$ , when assuming the value of the parameter  $n$  as derived from PYTHIA 6.1 simulations.

## 9.5 Present Emulsion Detectors

### 9.5.1 Fast Scanning Systems and Large-Scale Film Production

As stated above, the advances in the scanning systems aimed at higher efficiency and speed have led in recent times to the rebirth of the emulsion detectors. A further generation of the Track Selector, called S-UTS (Super-Ultra Track Selector), was developed in Nagoya [88]. It is based on highly customized components. The main feature of this approach is the removal of the stop-and-go process of the stage in the image data taking, which is the mechanical bottleneck of traditional systems. To avoid the stop, the objective lens moves at the same constant speed of the stage while moving also along the vertical axis and grabbing images with a very fast CCD camera running at 3000 Hz. The optical system is driven by a piezoelectric device. The camera has a sensor with  $512 \times 512$  pixels that imposes a smaller field of view ( $\sim 120 \times 120 \mu\text{m}^2$ ) to ensure a comparable position resolution (about  $0.3 \mu\text{m}/\text{pixel}$ ). The high-speed camera provides a data rate of 1.3 GB/s. This is handled by a front end image processor that makes the zero-suppression and the pixel packing, reducing the rate to 150–300 MB/s. A dedicated processing board performs track recognition, builds micro-tracks and stores them in a temporary device with a rate of 2–10 MB/s. A computer algorithm links the micro-tracks of different emulsion layers and writes the resulting tracks in a database that is used as input for physics analysis. The routine scanning speed is  $20 \text{ cm}^2/\text{h}/\text{layer}$  while one of the S-UTS systems has reached the speed of  $72 \text{ cm}^2/\text{h}/\text{layer}$  by using a larger field of view, without deteriorating the intrinsic micrometric accuracy of the emulsion films.

In the framework of the OPERA experiment (see next section), a joint effort of several European laboratories allowed the development of an automated scanning system (ESS) that employs commercial subsystems in a software-based framework. The ESS, derived from a system developed in Salerno [67], is extensively described elsewhere [89–91]. The microscope is a Cartesian robot, holding the emulsion film on a horizontal stage movable in  $X - Y$  coordinates, with a CMOS camera mounted on the optical axis ( $Z$ ), along which it can be moved to change the focal plane with a step roughly equal to the focal depth of about  $3 \mu\text{m}$ . The control workstation hosts a motion control unit that directs the stage to span the area to be scanned and drives the camera along the  $Z$  axis to produce optical tomographic image sequences (with the  $X - Y$  stage holding steady). Areas larger than a single field of view ( $\sim 300$

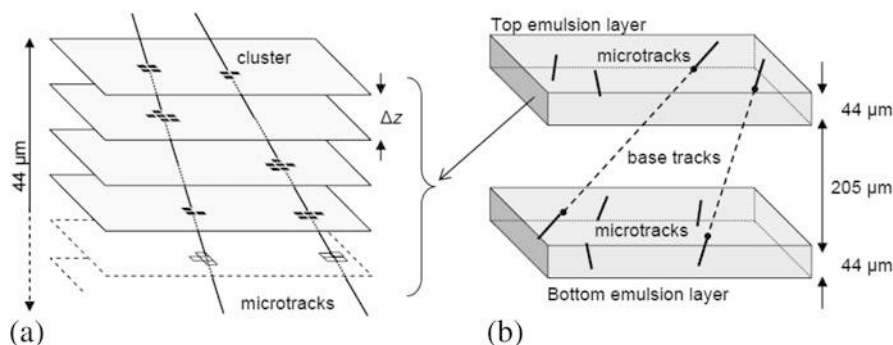


$\times 400 \mu\text{m}^2$ ) are scanned by repeating the data acquisition sequence on a grid of adjacent fields of view. The stage is moved to the desired position and the images are grabbed after it stops, with a stop-and-go algorithm. The images are grabbed by a Mpixel camera at the speed of 376 frames per second while the camera is moving in the Z direction. The whole system can work at a sustained speed of  $20 \text{ cm}^2/\text{h}/\text{layer}$ , 24 h/day, with an average data rate as large as 4 GB/day/microscope still preserving the intrinsic emulsion accuracy. A different setup of this system makes no use of immersion oil as interface between the objective lens and the film being scanned [92].

The track building method applied in both systems is schematically drawn in Fig. 9.9. The whole emulsion thickness is spanned by adjusting the focal plane of the objective lens and a sequence of 16 tomographic images is taken for each field of view at equally spaced depth levels, matching the focal depth of the objective. Emulsion images are then digitized, converted into a grey scale of 256 levels, sent to a vision processor board and analyzed to recognize sequences of aligned grains, i.e. clusters of dark pixels of given shape and size. Some of these spots are track grains; others, in fact the majority, are fog grains not associated to particle tracks. The three-dimensional structure of a track in an emulsion layer (microtrack) is reconstructed by combining clusters belonging to images at different levels and searching for geometrical alignments (Fig. 9.9a). Each microtrack pair is finally connected across the plastic base to form the so-called base track (Fig. 9.9b).

Figure 9.10 shows a S-UTS system and the scanning station in Bern employing the ESS system with dry objectives and with an automated emulsion film changer. The latter device allows fully unattended operation [93].

A second feature that significantly contributed to the rebirth of the emulsion detectors in recent times has been the realization of industrial emulsion films, optimized for micro-tracking applications. This is in particular the case of the FUJI R&D work conducted in collaboration with the Nagoya University [2] for the OPERA experiment that will be described later. Uniform automated machine



**Fig. 9.9** (a) Microtrack reconstruction in one emulsion layer by combining clusters belonging to images at different levels; (b) microtrack connections across the plastic base to form base tracks



**Fig. 9.10** Left: photograph of one of the Nagoya S-UTS scanning systems; right: the Bern scanning station equipped with five ESS microscopes with the associated automated film changers

coating of  $44\text{ }\mu\text{m}$  emulsion layers on either side of a plastic base was achieved for the unprecedented large-scale application of the OPERA ECC modules. The quality and the uniformity of the films is remarkable.

For machine coating in an industrial plant, dilution of the gel is required in order to reduce the viscosity. This implies a reduction of the grain density. In order to recover for this degradation, improvements in the gel sensitivity were applied, such as a controlled double jet method for the production of mono-dispersion of AgBr micro crystals. The crystal size is well controlled by this method. The number of crystals along a particle trajectory is increased, while the volume occupancy of AgBr and the average diameter of the crystals is kept constant.

In order to match the experimental requirements of a relatively thick layer with the limitations coming from the industrial process, a multi-coating method was adopted by FUJI. After the first layer ( $20\text{ }\mu\text{m}$  thick) is coated on both sides of a rolled plastic base, a second layer is coated over the first one. A thin  $2\text{ }\mu\text{m}$  gelatine spacer protects the emulsion layers. The resulting thickness is  $44\text{ }\mu\text{m}$ , well sufficient for automated track recognition. A glycerine bath is used to restore the thickness of the emulsion layers to its original value, thus recovering for the shrinkage induced by the development process.

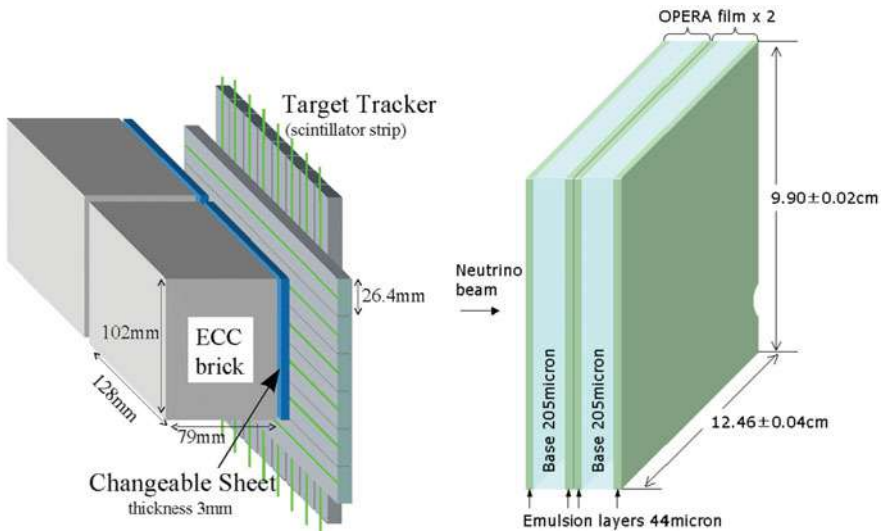
Another notable development related to the OPERA experiment has been the realization of the so-called emulsion refreshment. High temperature and high relative humidity enhance the latent image fading. This possibility is in particular useful when the exposure occurs much later than the film production and a low background is required, as in the case of OPERA. A good tuning of the fading features was achieved by introducing 5-methylbenzotriazole into the emulsion gel [2]. Absorption of this chemical by the silver specks induced by radiation lowers the oxidation reduction potential and makes the specks easy to oxidize. On the other hand, the sensitized centers (sulfur and gold) remain stable against the oxidation. Therefore, the recorded tracks are erased while the sensitivity remains sufficiently high. For example, keeping the films for 3 days at 98% relative humidity and  $27\text{ }^{\circ}\text{C}$ ,

the grain density of tracks accumulated before the refreshing goes from 30 to less than 10 grains/100  $\mu\text{m}$ , thus erasing about 96% of the stored tracks, including those from Compton electrons and cosmic-rays. The industrially produced films also feature a rather low track distortion induced by the development, as well as a limited level of fog density, with an initial value of 2.9 fog grains/1000  $\mu\text{m}^3$ ).

### 9.5.2 The OPERA Experiment

The OPERA experiment was designed to unambiguously prove  $\nu_\mu \rightarrow \nu_\tau$  oscillations in appearance mode. Indeed, studies of atmospheric neutrinos had shown the disappearance of muon neutrinos [94], later confirmed by accelerator experiments [95] and interpreted in terms of  $\nu_\mu \rightarrow \nu_\tau$  oscillations. Therefore, the appearance of tau neutrinos in a pure muon neutrino beam was the missing tile in the coherent scenario of neutrino mixing.

The conceptual design of the experiment was originally proposed in [96–98] and the detector is extensively described in [99, 100]. The distinctive feature of  $\nu_\tau$  charged-current interactions is the production of a short-lived  $\tau$  lepton ( $c\tau = 87 \mu\text{m}$ ). Thus, one has to accomplish the very difficult task of detecting sub-millimeter  $\tau$  decay topologies out of a huge background of  $\nu_\mu$  reactions in a target of more than a *kiloton*, as required to have a sufficient interaction rate. This is achieved in OPERA by employing a modern version of the ECC technology.



**Fig. 9.11** Schematic view of the ECC unit (brick) used in the OPERA experiment. A detail of the Changeable Sheet doublet is also shown

The OPERA experiment has been running from 2008 to 2012 at the underground LNGS laboratory in Italy, 730 km away from CERN where the CNGS neutrino beam was produced. OPERA is the first very large scale emulsion experiment, profiting from all the technological advances in the emulsion technology and in the scanning systems described in the previous section. To give a figure, the ECC target is made of films with a total surface of  $110,000\text{ m}^2$  and  $105,000\text{ m}^2$  lead plates. The industrially produced, machine-coated emulsion films by FUJI provided very uniform layer thickness and the possibility of erasing unwanted background tracks by the refreshing technique. The scanning of the events was performed with about 40 fully automated microscopes, each of them faster by about two orders of magnitude than those used in the CHORUS experiment [75].

The ECC target consisted of multi-layer arrays of target walls interleaved with pairs of planes of plastic scintillator strips. A target wall (with about  $10 \times 10\text{ m}^2$  cross-section) was an assembly of horizontal trays each loaded with ECC target units called bricks. A brick consisted of 57 emulsion films interleaved with 56 lead plates, 1 mm thick, light-tight packed. Brick dimensions were  $128 \times 102 \times 79\text{ mm}^3$  for a weight of 8.3 kg (Fig. 9.11). Interface Changeable Sheets (CS) were attached to the downstream face of each brick. The choice of the CS geometry was such to assemble two adjacent emulsion films as a doublet, coupled as an independent, detachable package to the downstream face of the brick (Fig. 9.11). The use of doublets provided the cancellation of random coincidences of tracks accumulated during the storage and transportation and unerased by the refreshing procedure.

There were 150,000 bricks in total for a target mass of 1.25 kton. This represents the largest ever ECC detector assembly and posed an unprecedented challenge for the production of emulsion films and bricks, as well as for the emulsion handling, development and analysis, i.e. scanning power. Just to give some numbers, more than nine million emulsion films were produced and the corresponding 150,000 bricks were built by a fully robotised chain assembling films and lead plates in an underground dark-room at LNGS. Large infrastructures were also realized at LNGS for brick manipulation (automatic extraction from the target matrix), X-ray marking, cosmic-ray exposure and emulsion development [100].

The principle of the experiment can be summarized as follows. At the occurrence of a neutrino interaction, the resulting charged particle tracks are detected by the scintillator counter planes placed behind each brick target wall, similarly to what happens in a sampling calorimeter. The reconstruction of the “shower axis” or the identification of a penetrating track (e.g. a muon) allows identifying the brick where the neutrino likely interacted. At this point, the brick is extracted from the wall, the attached CS doublet is removed and developed, while the brick, still packed, is placed in an underground storage area waiting for the response of the CS scanning.

It is important to stress the key roles accomplished by the CS in OPERA [101]: the first step is to confirm that the ECC brick contains the neutrino interaction; the second step is to provide event-related tracks to be used for the ECC scan-back analysis. By using Compton electrons from environmental radioactivity, the systematic uncertainties in the relative alignment between the two CS doublet films are reduced, thus bringing the position accuracy to the level of  $1\text{ }\mu\text{m}$  [102].


Such an accuracy allows using CS tracks made of only 3 out of the possible 4 track segments, thus increasing the track finding efficiency. Thanks to the CS, the bricks wrongly identified by the scintillator trackers are not disassembled but put back in the target with a fresh CS attached to them. This avoids useless film handling, processing and scanning of the misidentified bricks, and minimizes the corresponding waste of target mass. Moreover, whenever the electronic detector reconstruction is compatible with two or more “candidate” bricks, these are ordered by probability and their CS are scrutinized accordingly. This significantly increases the event finding efficiency.

If one or more “event related” tracks are found, the selected brick is exposed to cosmic-rays for about 12h, thus providing a set of tracks to be used for precise correction of local deformations as required for precision topological and kinematical measurements. The brick is then disassembled and its films are developed. The tracks measured in the CS analysis provide predictions for the so-called scan-back procedure. The latter consists of following a predicted track upstream in the ECC brick until it “disappear”. This procedure is initiated in the most downstream film of the brick.

The disappearance of a scan-back track indicates a possible neutrino interaction vertex. A wide area scan is performed over a volume of about  $1\text{ cm}^3$  around the track stopping point, looking for partner tracks and/or secondary decays with a dedicated decay search procedure [103]. This procedure, developed for the tau neutrino search, was successfully applied to the search for charmed hadron production induced by neutrinos. The latter process was indeed used as a control sample to check the efficiency for the detection of the tau lepton, given the similar lifetime of charmed hadrons (about  $10^{-12}\text{ s}$ ). The application of this procedure to muon neutrino interactions led to the observation of 50 decay candidates [103], in good agreement with the expected charmed hadron yield ( $54 \pm 4$ ), derived from the value measured by the CHORUS experiment [104]. Good agreement was found also in the shape of the relevant kinematical and topological variables, like the angle in the transverse plane between the charmed hadron and the muon and the impact parameter of the decay daughter particles with respect to the primary neutrino interaction vertex [103].

Unlike the experiments using “bulk” emulsions like CHORUS where the visual inspection of the primary and decay vertices allows rejecting most of the residual background, the ECC structure prevents the direct check of the vertices for the majority of the events. However, one can still exploit precise kinematical measurements for background suppression. For interesting event topologies, in fact, a detailed kinematical analysis is performed in OPERA by means of the electromagnetic shower energy measurement in the downstream part of the brick, the determination of the momentum by Multiple Coulomb Scattering measurement in the lead/emulsion structure [105], and the connection of tracks in consecutive target walls.

During the five CNGS production runs from 2008 to 2012, OPERA collected about  $1.8 \times 10^{20}$  protons on target and more than 19,000 neutrino interactions. The first tau neutrino candidate was reported in 2010 [106] and the display of its event




**Fig. 9.12** Display of the first  $\nu_\tau$  candidate. Top left: view transverse to the neutrino direction. Top right: same view zoomed on the primary and secondary vertices. Bottom: longitudinal view. Track 4 exhibits a kink topology with an angle of  $(41 \pm 2)$  mrad after a path length of  $(1335 \pm 35)$   $\mu\text{m}$  and produces track 8 and the two  $\gamma$ 's. Track 2 is identified as a proton and the other charged particles are all consistent with being hadrons [106]

reconstruction is shown in Fig. 9.12. The primary neutrino vertex consists of 7 tracks of which one shows a kink decay topology. None of the primary particles is consistent with neither a muon nor an electron. Two electromagnetic showers induced by  $\gamma$  conversions are visible in Fig. 9.12, indicated as  $\gamma_1$  and  $\gamma_2$ . These two  $\gamma$ s originate from the secondary vertex and their invariant mass is consistent with that of a  $\pi^0$ . From the kinematical analysis performed, the observed decay is consistent with the  $\tau \rightarrow \rho \nu_\tau$  channel (B.R.  $\simeq 25\%$ ), followed by  $\rho \rightarrow \pi^0 \pi$ .

The second [107] and third [108] tau neutrino candidates were reported in 2013, respectively in the  $\tau \rightarrow \pi \pi \pi \nu_\tau$  and  $\tau \rightarrow \mu \bar{\nu}_\mu \nu_\tau$  decay channels. The forth

**Table 9.1** Overall number of located neutrino interactions with the decay search procedure applied

	2008	2009	2010	2011	2012	Total
p.o.t. ( $10^{19}$ )	1.7	3.5	4.1	4.8	3.9	18.0
$0\mu$ events	150	255	278	291	223	1197
$1\mu$ events ( $p_\mu < 15$ GeV/c)	543	1024	1001	1031	807	4406
Total events	693	1279	1279	1322	1030	5603



**Fig. 9.13** Visible energy distribution of the 10 tau neutrino candidates found in the final sample [114]

candidate was reported in 2014 [109] while the discovery of  $\nu_\tau$  appearance was achieved in 2015 with the observation of a fifth tau neutrino candidate over an expected background of 0.25 events [110]. The OPERA discovery of tau neutrino appearance was explicitly mentioned in the Scientific Background of the 2015 Nobel Prize in Physics.

The emulsion handling was completed in 2015 while the emulsion film scanning was completed in 2016 when the detector was decommissioned. The final number of events passing all the analysis chain up to the decay search are shown in Table 9.1. Events are divided in two categories according to the presence ( $1\mu$ ) or absence ( $0\mu$ ) of a muon in the final state and undergo different selections: a momentum cut of 15 GeV/c is applied to the muons in order to reduce the background.

Given the data-driven validation of the simulation in all corners of the parameter space [103, 111, 112], the OPERA Collaboration decided to release the cuts and exploit the kinematical features of the events with a likelihood approach: this approach enlarges the selected sample, thus reducing the statistical uncertainty for the estimate of the oscillation parameters. Ten tau neutrino candidates were found with the new analysis strategy in the final sample. The distribution of the visible energy for the 10 candidates is shown in Fig. 9.13 together with the expected spectrum.

The number of expected tau neutrino events with looser cuts applied is reported in Table 9.2, together with the number of observed  $\nu_\tau$  candidates in each tau



**Table 9.2** Expected signal and background events for the analysed data sample

Channel	Expected background				$\nu_\tau$ expected	Observed
	Charm	Had. re-interaction	Large $\mu$ -scat.	Total		
$\tau \rightarrow 1h$	$0.15 \pm 0.03$	$1.28 \pm 0.38$	—	$1.43 \pm 0.39$	$2.96 \pm 0.59$	6
$\tau \rightarrow 3h$	$0.44 \pm 0.09$	$0.09 \pm 0.03$	—	$0.52 \pm 0.09$	$1.83 \pm 0.37$	3
$\tau \rightarrow \mu$	$0.008 \pm 0.002$	—	$0.016 \pm 0.008$	$0.024 \pm 0.008$	$1.15 \pm 0.23$	1
$\tau \rightarrow e$	$0.035 \pm 0.007$	—	—	$0.035 \pm 0.007$	$0.84 \pm 0.17$	0
Total	$0.63 \pm 0.10$	$1.37 \pm 0.38$	$0.016 \pm 0.008$	$2.0 \pm 0.4$	$6.8 \pm 1.4$	10

decay channel. The reported values assume  $\Delta m_{23}^2 = 2.50 \times 10^{-3} \text{ eV}^2$  [113] and  $\sin^2 2\theta_{23} = 1$ . The discovery of tau neutrino appearance is confirmed with a significance of  $6.1\sigma$ , evaluated by accounting for the features of the events with a likelihood analysis. The increased statistical sample was used to provide the first measurement of  $\Delta m_{23}^2$  in appearance mode with an improved accuracy, giving  $\Delta m_{23}^2 = (2.7_{-0.6}^{+0.7}) \times 10^{-3} \text{ eV}^2$  [114].

OPERA has demonstrated the capability of identifying all three neutrino flavours. Emulsion cloud chambers can clearly distinguish between electrons and  $\gamma$ s, given their micrometric accuracy emphasizing the displacement between the  $\gamma$  production and conversion vertices. Unlike other detectors, this feature makes the  $e/\pi^0$  separation particularly efficient and their selection pure: this translates into a very good separation between  $\nu_e$  charged-current interactions and  $\nu_\mu$  neutral-current ones with a  $\pi^0$  in the final state. OPERA has searched for the sub-dominant  $\nu_\mu \rightarrow \nu_e$  oscillations also to constraint the existence of sterile neutrinos. In the analysis of the 2008 and 2009 run data, 19 electron neutrino candidates were found and the results are summarised in [115]. The analysis of the final sample has collected 35  $\nu_e$  candidates and the constraints to sterile neutrinos are reported in [116]. Constraints to sterile neutrinos were set also with the analysis of  $\nu_\mu \rightarrow \nu_\tau$  oscillations [117].

## 9.6 Future Experiments and Applications

After more than 100 years since its first use, nuclear emulsions are still attractive in a wide range of scientific fields and applications. As it was the case for past developments, the future of nuclear emulsions will again rely on the parallel progress of high-performance readout systems as well as of innovative detector design. We discuss here the cutting edge technology and also review ongoing and emerging applications.



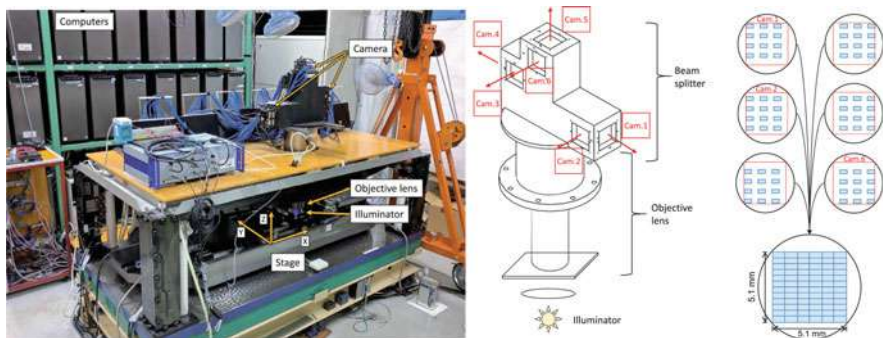
## 9.6.1 The State-of-the-Art Emulsion Technology

### 9.6.1.1 High-Performance Scanning Systems

Improvements of scanning systems in speed and quality are continuously progressing. One of the recent breakthrough was the appearance of GPGPU (General Purpose Graphic Processing Unit, or simply GPU). Up to the systems for OPERA, either FPGAs or CPUs were employed for image processing and track reconstruction. The FPGA has a big computing power, but also difficulties in implementing sophisticated algorithms and in flexibility. The CPU can process complicated algorithms but is limited in computing capability. On the other hand, the GPU provides both computing power and flexibility.

The effort to implement GPUs for scanning systems started soon after the release of CUDA [118], and it has quickly become the “standard” in the scanning system development nowadays. The early works were aiming at improving the angular acceptance in track reconstruction which was limited by the lack of computing power for online processing. The previously mentioned S-UTS, the scanning system for OPERA, could recognize tracks with their angle within  $30^\circ$  with respect to the normal of the film surface. This angular acceptance is equivalent to 14% of the entire solid angle. An extension of the S-UTS algorithm was translated into the GPU code, which reconstructed tracks up to  $72^\circ$  (68%) with a reasonable processing time [119]. In parallel, new algorithms suitable for parallel processing were developed to extend the track reconstruction to almost the entire  $4\pi$  solid angle [120, 121], which finally allowed to fully exploit the 3D tracking capability of nuclear emulsion. Examples of applications of such systems will be discussed further below.

In the data acquisition, there are two complementary ways for the fast readout of emulsion data: maximize the field of view or minimize the dead time due to microscope stage movement. An extreme case of the first approach was implemented in the HTS system (Hyper Track Selector, [122]) as shown in Fig. 9.14, which is the




**Fig. 9.14** Left: the fast emulsion readout system, Hyper Track Selector (HTS) [122]. Right: the optics and camera system for HTS. The optical path is divided into six mosaic camera modules. Each camera module consists of 12 2.2-Mpixel image sensors. In total 72 image sensors work in parallel to realize a large FOV of 5.1 mm  $\times$  5.1 mm with sub-micrometric resolution

fastest readout system at present. Conventional systems were using a field of view (FOV) of  $0.12 \text{ mm} \times 0.12 \text{ mm}$  (S-UTS) or  $0.3 \text{ mm} \times 0.4 \text{ mm}$  (ESS). HTS makes use of the custom made objective lens with a large FOV of  $5.1 \text{ mm} \times 5.1 \text{ mm}$  and a magnification of 12.1. The optical path is divided into six, correspondingly the image is projected on six “mosaic camera modules” as also schematically drawn in Fig. 9.14. Each mosaic camera module consists of 12 2.2-Mpixel image sensors. In total, 72 image sensors work in parallel to build the large FOV. The raw image data throughput from 72 image sensors amounts to 48 GBytes/s, which is then processed in real-time by 36 tracking computers with two GPUs each. The scanning speed has reached  $4700 \text{ cm}^2/\text{h}$ , which is clearly a big leap from the previous generations as shown in Fig. 9.15.


Another approach is to remove the dead time due to the microscope stage movement. In conventional systems, the data taking sequence is the so-called “stop-and-go” where the need to dump stage vibrations limits the repetition cycle up to 6 Hz. In order to minimize this effect, it was proposed to use tomographic image data taking without stopping the stage. In fact, S-UTS was the first system to implement the continuous motion as above mentioned. However, the camera resolution was relatively small ( $512 \times 512$  pixels) when compared to the market standard of today. The New Generation Scanning System (NGSS) was developed with a larger camera resolution ( $2336 \times 1728$  pixels) and with a different style of continuous motion that allowed running on normal motion hardware of ESS. The schematic of image taking sequence is shown in Fig. 9.16. By realizing a 12 Hz data taking, the scanning speed reached  $190 \text{ cm}^2/\text{h}$  [123].

The advances in scanning speed allows physicists to design experiments with a detector areas of  $1000 \text{ m}^2$  to be analysed in a year, to be compared to the total scanned area of OPERA of about  $500 \text{ m}^2$  in 5 years. The environment of emulsion readout is continuously changing as long as technologies grow. A new design of scanning system, so called HTS2, is going to combine the large field of view of HTS and the continuous motion [122], which might reach a scanning speed of  $5 \text{ m}^2/\text{h}$  in early 2020s. At this stage, the scanning speed would be no longer a bottleneck of

**Fig. 9.15** Time evolution of the scanning speed of the Track Selector system. The scanning speed progress in log scale



**Fig. 9.16** Schematic drawing of the Stop and Go (SG) motion and Continuous Motion (CM) of NGSS [123]




any experiment and new challenging experiments might be proposed, based on such a high-speed readout framework.

### 9.6.1.2 Fine-Grained Emulsion Production


Owing to its unbeatable position and angular resolution, the emulsion technique is being adopted in different applications in the fields of fundamental physics and applied science. The OPERA film [2], which was mass produced in industries, has been used for some applications, although the properties of the detector are tuned for the OPERA experiment. Following the increased interest in using emulsion detectors in a broad range of applications, the R&D of emulsion gel has become essential for optimising the detector for each application. However, conducting R&D for each small-scale experiment is difficult in industrial companies. This motivated the Nagoya University group to set up their own emulsion gel production facility in 2010. With the help of experts from FUJI Film Co. Japan, custom-made emulsion gels were successfully produced with an improved sensitivity to minimum ionizing particles with respect to OPERA films [3]. Moreover, some R&D programs were conducted to control silver halide crystal size, which defines spatial resolution and sensitivity. Fine-grained emulsions were produced with a crystal size of a few tens of nanometres, which is approximately one order of magnitude smaller than the conventional one (Fig. 9.17). They are called Nano Imaging Trackers (NIT) [124, 125]. The average size of NIT crystals was measured to be  $44.2 \pm 0.2$  nm, with a standard deviation of 6.8 nm. NITs are not sensitive to the minimum ionizing particles but are sensitive enough to low-velocity heavy ions. They are considered a possible detector for detecting the recoiled nuclei induced by dark matter.

### 9.6.1.3 Large Grain Emulsion Production

For certain applications such as muon radiography, large-scale detectors are required. An improvement in the readout speed is therefore crucial to make future large-scale applications possible, and the availability of a new type of emulsion featuring crystals of larger sizes is one way to pursue this goal. This would allow a lower magnification for the microscopes and, consequently, a larger field of view resulting in a faster data analysis. The size of the crystals used for the neutrino



**Fig. 9.17** Left: Silver halide crystals in the fine-grained emulsion [124, 125], as seen with a transmission electron microscope. Photolytic silver grains are also visible on the surfaces of silver halide crystals. Right: Tracks of Kr ions in such an emulsion, as seen with a scanning electron microscope



**Fig. 9.18** Electron microscope pictures of silver halide crystals (left) and electron tracks (right) in a conventional film and in the newly developed samples [126]

oscillation experiments mentioned above was 200 nm and has never been larger than 300 nm in previous experiments. The production of new types of emulsions with crystal sizes of 600–1000 nm, 3–5 times larger than those of standard films, has been studied and realised using the gel production machine at Nagoya University. The first results characterising newly produced emulsions have been reported [126], showing a sufficient sensitivity and a good signal to noise ratio (Fig. 9.18). This development will allow a 25 times faster readout speed by using lower magnification objective lenses. These new detectors will pave the way to future large-scale applications of the technology, e.g. 3D imaging using muon radiography or future neutrino experiments.

In close connection with the production of large crystals, there has also been a study to produce crystals slightly larger (350–400 nm) than 200 nm and to check

the dependence of crystal sensitivity on its size. This study was motivated by the interest to understand the phenomenology of the latent image formation, which predicts that the quantum sensitivity can be better at such a crystal size. Further studies are in progress with the aim of developing emulsions with higher sensitivity. The conditions of chemical sensitisation and development were optimised for each crystal size in the range of 200–800 nm. The increase in the crystal sensitivity depending on the crystal size was confirmed for crystals of 350–800 nm [127]. These R&D activities form the base for a broad range of future applications.


## 9.6.2 Projects in Fundamental Physics

### 9.6.2.1 Balloon Experiments

Balloon experiments employing emulsion detectors were reported in [128]. The use of emulsion technique for cosmic-ray physics experiments has recently attracted research interest after the significant technological advances in the last decades. In 2004, a balloon experiment using emulsions was performed to observe primary cosmic-ray electrons [129]. Various innovations such as the industrial emulsion films, the refreshing technique, the automated emulsion read-out system and the off-line analysis methods were introduced. In addition, a dedicated device was developed to distinguish between particles passing through a chamber at the balloon level altitude and those recorded during other periods. The mechanism of this device is such that it causes intentional shift of the upper block of the chamber with respect to the lower block, when the balloon reaches float altitude and again when the flight at float altitude is terminated. The working principle of the technique was successfully demonstrated.

Based on these techniques, a balloon-borne emulsion  $\gamma$ -ray telescope was proposed [130] and the Gamma-ray Astro-Imager with Nuclear Emulsion (GRAINE) project was developed for the observation of  $\gamma$ -rays in the energy range of 10 MeV–100 GeV. A precise, polarisation-sensitive, large-aperture-area emulsion telescope with repetitive long-duration balloon flights was employed. The electron and positron angles at the pair creation point can be measured in emulsions and the angular resolution for  $\gamma$ -rays (10 MeV–10 GeV) is about one order of magnitude higher than that of the Fermi Large Area Telescope (Fermi-LAT) (Fig. 9.19). The polarisation sensitivity of an emulsion-based telescope was demonstrated using a polarised  $\gamma$ -ray beam at SPring-8/LEPS [131].

An emulsion multi-stage shifter was used to develop an innovative solution capable of providing the event time-stamp and hence the  $\gamma$ -ray absolute direction [135]. The relative alignment between the automatically sliding emulsion films, each moving with a known different speed, provides the required time association of the event. This technique allows  $\gamma$ -ray detection with low energy threshold, minimising the electric power and also limiting the overall detector mass.



**Fig. 9.19** Angular resolution of the emulsion  $\gamma$ -ray telescope (lines show simulation results and dots with error bars show experimental results) [132]. The measurements were performed with  $\gamma$ -ray beams (LEP/SPring-8, UVSOR, and New SUBARU) and using a flight data [133]. Dotted lines show angular resolution with Fermi-LAT for the front section with thin radiation foils and the back section with thick foils [134]

In 2011, the first balloon-borne experiment was performed with a  $12.5 \times 10 \text{ cm}^2$  aperture area and 4.6-h flight duration for a feasibility test [133]. The chamber comprised three sections. The top section was made of an ECC with emulsion films interleaved with copper foils ( $50 \mu\text{m}$ ), meant to measure the  $\gamma$ -ray angle around the conversion point. The middle section included an emulsion multi-stage shifter, providing the time-stamp of the events. The bottom part contained a calorimeter comprising a lead/emulsion ECC for the  $\gamma$  energy measurement. With this flight data, systematic detection, energy reconstruction, and timestamping of  $\gamma$ -ray events were performed [133] and subsecond time resolution of the emulsion  $\gamma$ -ray telescope was demonstrated [136]. The second balloon-borne experiment was performed at the Alice Springs balloon-launching station in 2015 [137]. The telescope had a  $3780 \text{ cm}^2$  aperture and was taking data for a total of 14.4 h. The experiment aimed at demonstrating the overall performance of the emulsion  $\gamma$ -ray telescope. The improvements in the emulsion characteristics and handling applied to this experiment are summarised in [138]. The project plans a third balloon-borne experiment in 2018 for the celestial source detection and envisions scientific observations from 2021.

### 9.6.2.2 The NEWSdm Experiment

The nature of Dark Matter is one of the fundamental questions to be answered. Direct Dark Matter searches are focussed on the development, construction, and operation of detectors looking for the scattering of Weakly Interactive Massive

Particles (WIMPs) with target nuclei. The measurement of the direction of WIMP-induced nuclear recoils is a challenging strategy to extend the sensitivity of dark matter searches beyond the neutrino-induced background event rate and provide an unambiguous signature of the detection of Galactic dark matter [139]. Current directional experiments are based on the use of gas TPC whose sensitivity is strongly limited by the small achievable detector mass. Nuclear Emulsions for WIMP Search with directional measurement, NEWSdm, is an innovative directional experiment proposal based on the use of a solid target made by newly developed nuclear emulsion films and read-out systems capable to detect nanometric trajectories.

The approach proposed by the NEWSdm Collaboration [140] consists of using a nuclear emulsion-based detector acting both as target and as nanometric tracking device. The NEWSdm project foresees the employment of NIT. The detector is conceived as a bulk of NIT surrounded by a shield to reduce the external background. The detector is then placed on an equatorial telescope in order to absorb the Earth rotation, thus keeping fixed the detector orientation with respect to the incoming apparent WIMP flux. The angular distribution of the WIMP-scattered nuclei is therefore expected to be strongly anisotropic with a peak centred in the forward direction.

NIT have a linear density of about 11 crystals/ $\mu\text{m}$  [124], thus making the reconstruction of trajectories with path lengths as short as 100 nm possible, if analysed by means of microscopes with enough resolution. The presence in the emulsion gel of lighter nuclei such as carbon, oxygen and nitrogen, in addition to the heavier nuclei of silver and bromine, is a key feature of the NEWSdm project, resulting in a good sensitivity to WIMPs in the mass range between 10 to 100 GeV/ $c^2$ .

In the NEWSdm experiment a WIMP signal consists of short-path, anisotropically distributed, nuclear recoils over an isotropically distributed background. The search for signal candidates requires the scanning of the whole emulsion volume. The read-out system has therefore to fulfil two main requirements: a fast, completely automated, scanning system is needed to analyse the target volume over a time scale comparable with the exposure; the spatial resolution achievement has to go well beyond the diffraction limit, in such a way to ensure high efficiency and purity in the selection of signal candidates. The analysis of NIT emulsions is performed with a two-step approach: a fast scanning with a state-of-the-art resolution for the signal pre-selection followed by a pin-point check of preselected candidates with unprecedented nanometric resolution to further enhance the signal to noise ratio.

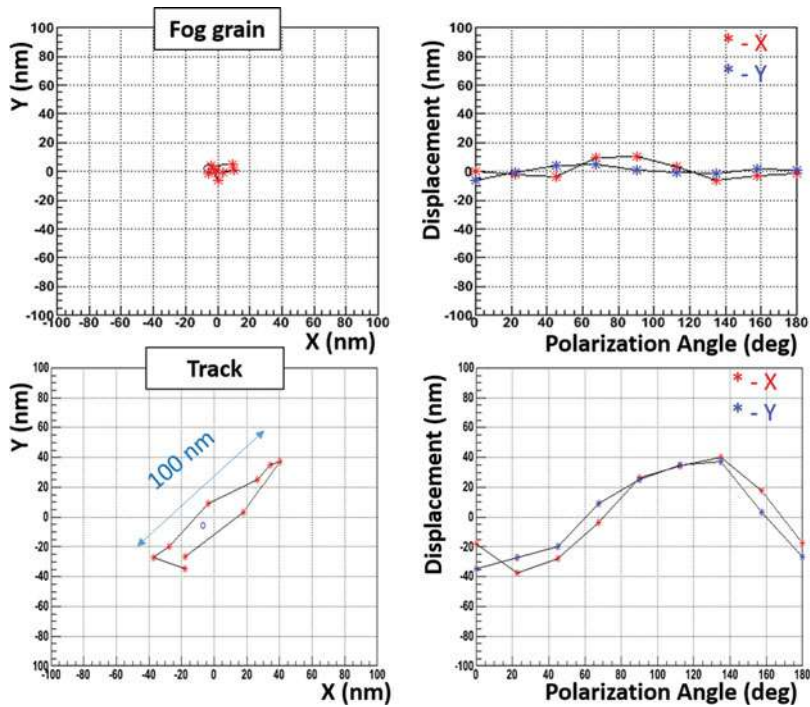
In the first analysis phase, a fast scanning is performed by means of an improved version of the optical microscope used for the scanning of the OPERA films [141]. An R&D program has achieved a speed of about 200 cm<sup>2</sup>/h [123, 142].

The starting point of the emulsion scanning is the image analysis to collect clusters making up silver grains. Given the intrinsic resolution of the optical microscope ( $\sim 200$  nm), the sequence of several grains making a track of a few hundred nanometers may appear as a single cluster. Nevertheless, a cluster made of several grains tends to have an elliptical shape with the major axis along the direction of the trajectory, while a cluster produced by a single grain tends to have a



spherical shape. The shape analysis with an elliptical fit is indeed the first approach to select signal. In order to simulate the effect of a WIMP-induced nuclear recoil and to measure the efficiency and the resolution of the new optical prototype, a test beam with low velocity ions was performed. Kr ion beams with energies of 200 and 400 keV [143] and C ion beams with energies of 60, 80 and 100 keV were used. Silver grains belonging to the tracks appear as a single cluster. An elliptical fit of the cluster shape allows a clear separation between fog grains and signal tracks [144].


The second analysis step at the microscope makes use of the plasmon resonance effect occurring when nanometric silver grains are dispersed in a dielectric medium [145]. The polarization dependence of the resonance frequencies strongly reflects the shape anisotropy and can be used to infer the presence of non-spherical nanometric silver grains within a cluster made of several grains. NEWSdm is using this technology to retrieve track information beyond the diffraction limit. Images of the same cluster taken with different polarization angles show a displacement of the position of its barycentre. The analysis of this displacement allows to distinguish clusters made of a single grain from those made of two or more grains building up a track, as shown in Fig. 9.20: unlike the single grain reported in the top



**Fig. 9.20** Displacement of the barycentre as a function of the light polarization angle. The response to a single grain (top) and to a C ion track (bottom) are compared. The ion track shows a clear displacement




**Fig. 9.21** Barycentre displacement of clusters made by single grains




plots, the Carbon ion track in the bottom plot shows a barycentre displacement of 100 nm length while changing the polarization angle. An unprecedented nanometric accuracy has been achieved in both coordinates with this method: Fig. 9.21 reports the displacement of the barycentre of clusters made of single grains, showing an RMS smaller than 10 nm. Such an achievement allows to detect path lengths where the barycentre displacement induced by the polarization change is only a few tens of nanometres. The actual threshold achievable on path lengths depends on the crystal size and can in principle be reduced to a few tens of nanometres as well.

The wavelength of the scattered light depends on the size of the grains where light is scattered off. In order to exploit this effect, the latest version of the optical microscope makes use of a colour camera, thus providing sensitivity to the sense of the track, since grains are expected to be larger at the end of the track range and therefore the scattered light shifts to the red colour. The prototype of this new system in operation in Naples is shown in the left plot of Fig. 9.22 while the image of an  $\alpha$  track is reported on the right: different grains show different colours due to their different size, that in turn can provide sensitivity to the particle sense.

The NEWSdm collaboration has installed at the Gran Sasso underground laboratory a facility for the emulsion handling and film production. Moreover, a dedicated structure was constructed in the Hall B of the underground Gran Sasso Laboratory early in 2017 to shield a detector of 10 g mass against the environmental background sources over an exposure time of about 1 month. The experimental setup consists of a shield from environmental backgrounds, made of a few tons of polyethylene and lead, and a cooling system to ensure the required temperature level to the NIT emulsion detector. The aim is to measure the detectable background from environmental and intrinsic sources and to validate estimates from



**Fig. 9.22** Left: new optical microscope equipped with colour camera in Naples. Right: the last few microns of an  $\alpha$  track path showing grains of different colours



**Fig. 9.23** Cross-sectional view of the detectors [149]. Detectors with  $\text{LiNO}_3$  doping (left) and the  $^{10}\text{B}_4\text{C}$  thin layer (right)

simulations [146]. The confirmation of a negligible background would pave the way for the construction of a pilot experiment with an exposure of about 10 kg year.

### 9.6.2.3 Development of Cold-Neutron Detector

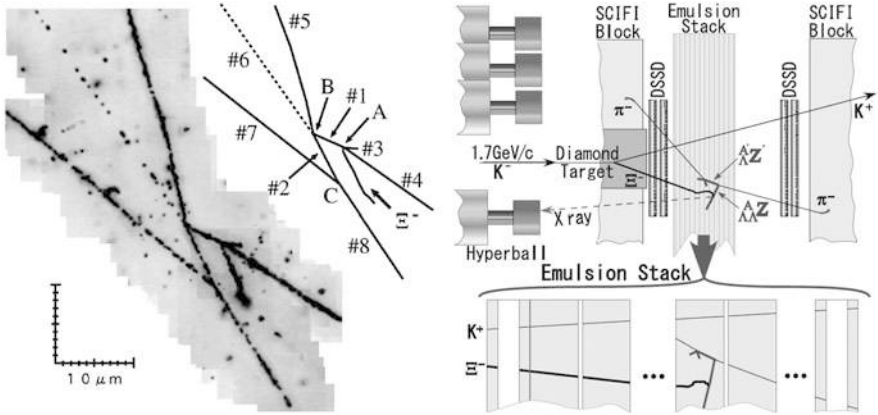
A new detector for detecting cold and ultra-cold neutrons has been recently developed. It employs fine-grained emulsion detectors with 35-nm-diameter crystals and nuclides with large neutron absorption cross sections, such as  $^6\text{Li}$  and  $^{10}\text{B}$ . One detector type is realised by doping  $\text{LiNO}_3$  into fine-grained emulsion detectors [147]. The cross-sectional view of the detector is shown in Fig. 9.23 (left). An  $\alpha$  particle and a tritium are emitted during the reaction:  $^6\text{Li} + n \rightarrow \alpha + t + 4.78 \text{ MeV}$ . Events of neutron absorption by  $^6\text{Li}$  were successfully observed by exposing the

detector to thermal neutrons at the Kyoto University Research Reactor Institute (KURRI). The spatial resolution achieved in the measurement of the absorption point was estimated to be  $0.34\text{ }\mu\text{m}$  from the average grain density of the track far from the end of its range. The detection efficiency was measured by exposing the detector to a cold neutron beam at BL05 port in the Materials and Life Science Experimental Facility (MLF) at J-PARC [148]. The measured efficiency of  $(3.3 \pm 0.6) \times 10^{-4}$  was consistent with the expectation.

The other detector type consists of a 50-nm-thick converter layer made of  $^{10}\text{B}_4\text{C}$  formed on a 0.4-mm-thick silicon substrate and coated by 10- $\mu\text{m}$ -thick, fine-grained emulsion [149]. The converter layer was covered by C (50 nm) and NiC (60 nm) layers. An  $\alpha$  particle or a  $^7\text{Li}$  nucleus will be detected in the emulsion, as shown in Fig. 9.23 (right). They are produced via the reactions:  $^{10}\text{B} + n \rightarrow \alpha + ^7\text{Li} + 2.79\text{ MeV}$  (6%) or  $^{10}\text{B} + n \rightarrow \alpha + ^7\text{Li} + 2.31\text{ MeV}$  (94%). The detector was exposed to cold and ultra-cold neutrons at J-PARC, and the events of neutron absorption by  $^{10}\text{B}$  were clearly observed. The position resolution of the absorption point in the  $^{10}\text{B}_4\text{C}$  layer depends on the track angle. By limiting the track angle, the expected position resolution is  $\sim 100\text{ nm}$  [150], which is 1–2 orders of magnitude higher than that of the conventional detectors used for detecting cold and ultra-cold neutrons. Further optimisation of the thickness of the converter layer and development of automatic track reconstruction are explored. The development of these detectors paves the way to future applications such as the precise measurement of the position distribution of quantised states of ultra-cold neutrons or neutron imaging with future neutron sources.

#### 9.6.2.4 Study of Double-Hypernuclei

The knowledge of  $\Lambda - \Lambda$  interaction is limited as only one out of the nine double hypernuclei events detected by E373 is fully analysable to extract information of the interaction (Fig. 9.24). In order to answer questions such as the nuclear mass dependence of  $\Lambda - \Lambda$  interaction, the E07 experiment at J-PARC is being carried out, which is aiming at studying  $\Lambda - \Lambda$  interactions with 100 double hypernuclei events, one order of magnitude larger statistics with respect to E373. As schematically shown in the right side of the Fig. 9.24, E07 uses a  $1.7\text{ GeV}/c$   $K^-$  beam hitting a diamond target to produce  $\Xi^-$  hyperons ( $dss$ ), subsequently stopped and captured by one of emulsion detector nuclei to produce double hypernuclei. The detector has a hybrid structure with a silicon strip detector and a spectrometer system for the  $K^+$  identification, needed to tag  $\Xi^-$  hyperons. The emulsion detector will be analysed by an automated scanning system dedicated to this experiment. E07 conducted the physics runs in 2016 and 2017. The emulsion readout and analysis are in progress.




**Fig. 9.24** Left: The so-called “Nagara” double hypernucleus event found in the E373 experiment at KEK.  $\Xi^-$  was captured at rest by a carbon nucleus in the emulsion detector and produced a double hypernucleus ( $^6_{\Lambda\Lambda}\text{He}$ ), which decayed in series, leaving a peculiar event topology in emulsion [60]. Right: A schematic of the experimental setup of the E07 experiment at J-PARC [151]

### 9.6.2.5 Measurements of Antimatter

Emulsion detectors have been recently considered as high-accuracy position sensitive detectors for low-energy antimatter studies. These studies include the AEgIS experiment at CERN [152, 153], with the goal of measuring the Earth’s gravitational acceleration on antihydrogen atoms to the ultimate precision of 1%. The vertical deflection of the  $\bar{H}$  atoms due to gravity will be detected by a setup comprising material gratings coupled with a position-sensitive detector. The position detector requires the best possible position resolution, which currently is provided by emulsion-based detectors.

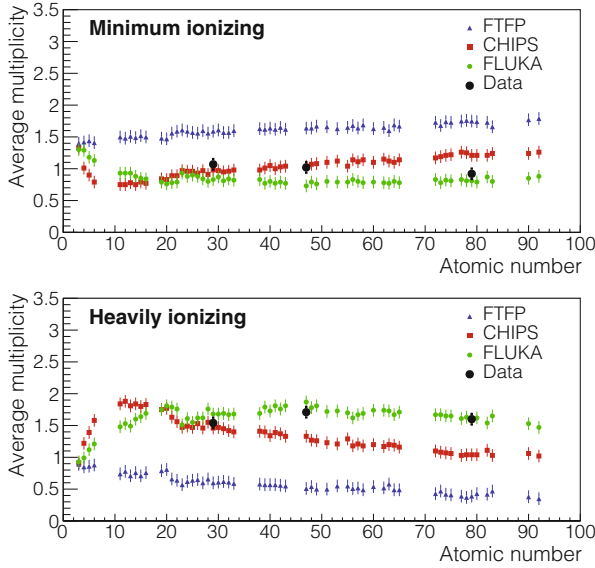
There were technical challenges for emulsion detectors to be operated in vacuum and at cryogenic temperatures. In vacuum, water loss leads to cracks in the emulsion layer and an increase in random noise due to mechanical stress caused by the drying process. Two ways of solving these problems were then established. One is to mix glycerin with emulsion gel to replace water with glycerin, and the other one is to put gas barrier films on emulsions to keep water in the films. Both approaches have proven to work in ordinary vacuum ( $10^{-7}$ – $10^{-5}$  mbar). At 77 K, the performance of emulsion detectors was not well known. The sensitivity of the emulsion at 77 K was studied and observed to be 43% of the value at 300 K. By optimising the track reconstruction, detecting minimum ionizing particles with such a sensitivity will be feasible since the tracking efficiency for tracks with more than 10 grains in an 50–100  $\mu\text{m}$  thick emulsion layer could be close to 100%. In 2012, an antiproton exposure to the emulsion detectors was performed for the feasibility study at the Antiproton Decelerator (AD [154]) at CERN. Fig. 9.25 shows an annihilation vertex on the bare emulsion surface. Annihilation vertices in the metal target were

**Fig. 9.25** An antiproton annihilation vertex in an emulsion layer [155]. The view is perpendicular to the antiproton beam direction



also reconstructed, demonstrating a resolution of  $1\text{ }\mu\text{m}$  on the vertical position. In addition, a proof-of-principle experiment with mini-moiré deflectometer was performed [156]. The periodic patterns were observed as expected in the emulsions, and the measured shift between antiprotons and light was consistent with the force from the magnetic field at the given position. The results are a crucial step toward the direct detection of the gravitational acceleration of antihydrogen. In 2014, measurements of the multiplicities of charged annihilation products on different target materials, namely copper, silver, and gold, were performed [157] at the CERN AD. Apart from the obvious applications in nuclear physics, this measurement can provide a useful check of the ability of standard Monte Carlo packages to reproduce fragment multiplicities and energy distributions. The measured fragment multiplicities were not well reproduced by the different models used in Monte Carlo simulation with the exception of FLUKA [158, 159], which is in good agreement with the particle multiplicities for both minimum and heavily ionizing particles (Fig. 9.26).

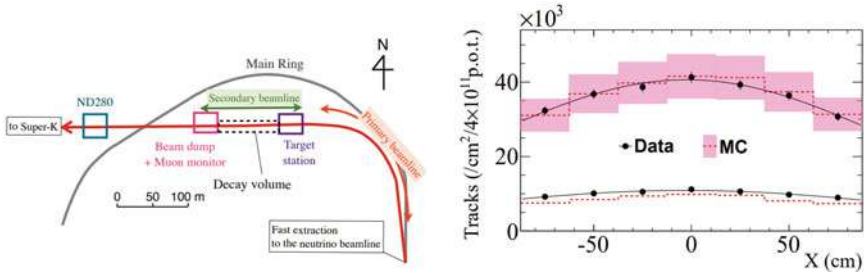
There is another proposal by the QUantum interferometry and gravity with Positrons and LASers (QUPLAS) project to use emulsions for their studies on positrons [160]. The sensitivity of the emulsion detectors was studied using a mono-energetic positron beam at energies as low as 9–18 keV. The obtained results prove that the emulsions are highly efficient at detecting positrons at these energies. This achievement paves the way to perform matter-wave interferometry with positrons using this technology.



**Fig. 9.26** Particle multiplicity from antiproton annihilations as a function of atomic number for minimum and heavily ionizing particles [157]

#### 9.6.2.6 Accelerator Beam Characterization: Muon Measurements at the T2K $\nu$ Beamline

The high spacial resolution of emulsion detectors turns out to have an advantage in the characterization of high-intensity accelerator beams, in particular in fast extraction mode where billions of particles arrive within a nanosecond: in these conditions electronic detector cannot identify particles on an event by event base given their limited occupancy. A notable example is the muon measurement at the T2K neutrino beam from J-PARC in Japan [161]. As neutrinos and muons are both produced by meson decays ( $\pi, K \rightarrow \mu \nu_\mu$ ), the understanding of the muons provides valuable information about neutrinos, such as the parent hadron production and momentum distribution. Nevertheless, low energy electromagnetic components highly contaminate the muon flux at the muon pit downstream of the decay volume ( $\mu^\pm = 53\%$ ,  $e^\pm = 7\%$ ,  $\gamma = 40\%$ , estimated by MC). This makes it difficult to extract meaningful information from the muon beam. The muons are regularly measured by the silicon photodiodes and ionization chambers at the muon pit to monitor the beam direction in each spill. These are charge-integration detectors not optimised to measure muon tracks. Therefore, a measurement of muons by means of the emulsion detectors was performed. The emulsion detector module was composed of 8 OPERA-type films with a cut-off momentum of 30 MeV/c for electrons given by the dedicated track recognition procedure, efficiently achieving a 99% purity of muons after reconstruction. The



**Fig. 9.27** Left: A schematic of the T2K neutrino beamline. The muon measurement was performed at the muon monitor pit behind the decay volume. Right: Comparison of the muon flux with the prediction at the horn current of 250 kA (top) and off (bottom). Figures from [161]

measurement was done at an intensity of the order of  $10^{11}$  protons on target, which yielded  $O(10^4)$  muons/cm<sup>2</sup> in the emulsion detectors. The measured profile is shown in Fig. 9.27 with the expected profile predicted by the FLUKA simulation with dedicated hadron production tuning. The absolute muon flux was measured for the first time at neutrino beamlines, thus characterising the beamline. In addition to the flux measurement, the momentum distribution of muons has been measured with the OPERA-like ECC with 25 films sandwiched with 24 1-mm-thick lead plates, which is to be published.

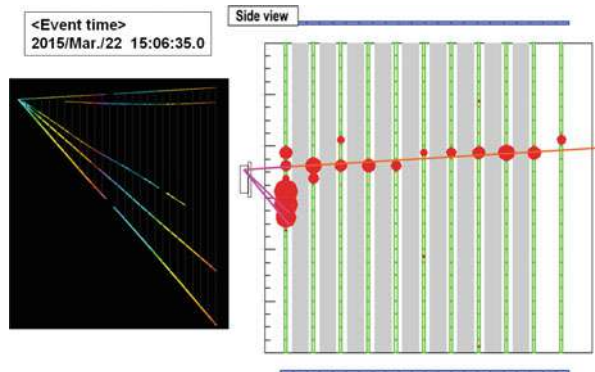
Such muon measurements can be performed at future neutrino beamlines e.g. the J-PARC neutrino beamline for the Hyper-K experiment [162] and the LBNF (Long-Baseline Neutrino Facility) for the DUNE experiment [163]. In general, nuclear emulsions provide unique capability to study high-intensity accelerator beamline operated in fast extraction mode. Thanks to the automated scanning system, this field is expected to grow in the future.

#### 9.6.2.7 The NINJA Project

The Neutrino Interaction research with Nuclear emulsion and J-PARC Accelerator (NINJA) project was initiated for the precise measurement of neutrino–nucleus interactions. The study of neutrino–nucleus interactions in the sub-multi-GeV region is important to reduce systematic uncertainties in present and future neutrino oscillation experiments. The emulsion detector can measure particles with a low energy threshold for various targets such as iron, carbon and water. It also exhibits good electron/gamma separation capability, allowing for precise measurements of electron–neutrino interactions. Given these capabilities, the future program includes searches for sterile neutrinos with a detector made of three components. The upstream part is made of an ECC with emulsion films interleaved by the target material, which is used for detecting neutrino interactions. The middle part includes an emulsion multi-stage shifter device [164], providing the timing information of



**Fig. 9.28** Hybrid analysis of ECC and INGRID (side view of an event) [166]



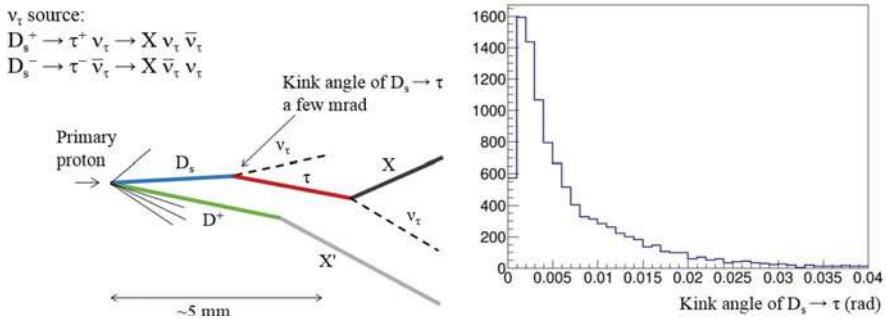
the events. The downstream part is the Interactive Neutrino Grid (INGRID), which is one of the near detectors for the T2K experiment [165] used to identify muons.

A test experiment (J-PARC T60) was implemented as a first step in the project to check the performance of the detector. Its neutrino event analysis is based on scanning the full area of the emulsion detectors by the HTS system. A more detailed analysis of the detected events could be performed by a dedicated scanning procedure with an extended angular acceptance [119]. The feasibility of the project was studied in the first exposure to a 2-kg iron target in 2015. The full area of the emulsion films ( $\sim 1.2 \text{ m}^2$ ) was scanned and a systematic analysis was performed to locate neutrino interactions. The neutrino candidate events located in the emulsions were matched to events observed by INGRID by employing the timing information from the multi-stage shifter. The hybrid analysis of ECC and INGRID was demonstrated [166]. An analysed event is shown in Fig. 9.28. Further, some other exposures of the anti-neutrino beam to the detectors were conducted, testing the first proto-type of water-target ECC or checking the detector performance with higher statistics [167]. The plan is to scale up the detector step-wise. Physics runs will be planned based on the results of these test experiments.

### 9.6.2.8 Tau-Neutrino Production Studies

At the CERN Super Proton Synchrotron (SPS), a new project called DsTau has been proposed to study tau-neutrino production [168] aiming at providing important information for future  $\nu_\tau$  measurements where high  $\nu_\tau$  statistics is expected. The results of DsTau are a prerequisite for measuring the  $\nu_\tau$  charged-current cross section, which has never been adequately measured (only the DONUT measurement was reported so far [87]). Precise measurement of the cross section would enable a search for new physics effects in  $\nu_\tau$ -nucleon CC interactions. It also has practical implications for neutrino oscillation experiments such as Super-K, Hyper-K [162] and DUNE [163], which suffer from a  $\nu_\tau$  background to their  $\nu_e$  measurements. As for the DONUT experiment, the dominant source of  $\nu_\tau$  is the sequential decay of  $D_s$





**Fig. 9.29** Topology of  $D_s \rightarrow \tau \rightarrow X$  events (left) and simulated kink angle distribution of  $D_s \rightarrow \tau$  (right) [168]

mesons,  $D_s^+ \rightarrow \tau^+ \nu_\tau \rightarrow X \nu_\tau \bar{\nu}_\tau$  and  $D_s^- \rightarrow \tau^- \bar{\nu}_\tau \rightarrow X \bar{\nu}_\tau \nu_\tau$  produced in high-energy proton interactions. The topology of such an event is shown in Fig. 9.29. Directly measuring  $D_s \rightarrow \tau$  decays will provide an inclusive measurement of the  $D_s$  production rate and the decay branching ratio to  $\tau$ . The  $D_s$  momentum will be reconstructed by combining the topological variables measured in the emulsion detector.

The project aims at detecting  $10^3$   $D_s \rightarrow \tau$  decays to study the differential production cross section of  $D_s$  mesons. For this purpose, emulsion detectors with a nanometer precision readout will be used. An emulsion detector with a crystal size of 200 nm has a position resolution of 50 nm [3], as shown in Fig. 9.2, allowing for kink detection with a threshold of 2 mrad at the  $4\sigma$  confidence level. The global analysis will be based on fast scanning of the full area by the HTS system [122]. After the  $\tau$  decay trigger, the events will be analysed by dedicated high-precision systems [120] using a piezo-based high-precision z-axis, allowing the emulsion hits to be measured with a nanometric resolution. Each detector unit consists of a 500  $\mu\text{m}$ -thick tungsten target, followed by 10 emulsion films interleaved with 200  $\mu\text{m}$ -thick plastic sheets acting as decay volumes for short-lived particles as well as high-precision particle trackers. Ten such units are used to construct a module, which is followed by an ECC to measure the momenta of the daughter particles. With this module,  $4.6 \times 10^9$  protons on target are needed to accumulate  $2.3 \times 10^8$  proton interactions in the tungsten plates. The data generated by this project will enable the  $\nu_\tau$  cross section measured by DONUT to be re-evaluated, which should significantly reduce the total systematic uncertainty. Once  $\nu_\tau$  production is established, the next stage will be to increase the number of  $\nu_\tau$  detected events. This could be achieved within the framework of the SHiP project [171] at CERN because its beamline (beam-dump type) is well suited for this task. The DsTau project aims to look for new physics effects in  $\nu_\tau$ -nucleon CC interactions with a total uncertainty of 10%. In addition to the main aim of measuring  $D_s$ , analysing  $2.3 \times 10^8$  proton interactions, combined with the high yield of  $10^5$  charmed decays produced as by-products, will enable the extraction of additional physical quantities. Based on the


results of beam tests undertaken in 2016 and 2017 for the feasibility study, a pilot run is scheduled for 2018 and physics runs are planned from 2021 after the upcoming long shutdown of the accelerator complex at CERN.

Another proposal, called the SHiP-charm project [169], aims at measuring the associated charm production by employing the SPS 400 GeV/c proton beam. Charmed hadrons are produced either directly from interactions of the primary protons or from subsequent interactions of the particles produced in the hadronic cascade showers. Recent detailed simulation studies of proton interactions in heavy and thick targets show a sizeable contribution from the cascade production to the charmed hadron yield [170]. This proposal includes a study of the cascade effect to be carried out using ECC techniques, i.e. slabs consisting of a replica of the SHiP experiment target [171] interleaved with emulsion films. The detector is hybrid, combining the emulsion technique with electronic detectors to provide the charge and momentum measurement of charmed hadron decay daughters and the muon identification. This allows a full kinematical reconstruction required by the double-differential cross-section measurement. According to the simulation performed, the delivery of  $2 \times 10^7$  protons on target would allow the detection of about 1000 fully reconstructed charmed hadron pairs. An optimisation run is scheduled for 2018 and the full measurement is planned after the long shutdown LS2 of the CERN accelerator complex, with  $5 \times 10^7$  protons on target and a charm yield of about 2500 fully reconstructed interactions.

These two approaches, DsTau and SHiP-charm, are complementary since DsTau will detect  $10^5$  charmed hadron pairs with good  $D_s$  selection capability, while SHiP-charm will study about 2500 fully reconstructed charmed hadron pairs including the hadronic cascade effect. The results of these approaches will provide essential input for future  $\nu_\tau$  measurements.

### 9.6.2.9 The SHiP Experiment

The discovery of the Higgs boson in 2012 has fully confirmed the Standard Model of particles and fields. Nevertheless, there are still fundamental phenomena, like the existence of dark matter, the baryon asymmetry of the Universe and the origin of neutrino masses, that could be explained by the discovery of new particles. Searches for new physics with accelerators are performed at the LHC, looking for very massive particles coupled to matter with ordinary strength. A new experiment, Search for Hidden Particles (SHiP), has been proposed [171], designed to operate at a beam dump facility to be built at CERN and to search for weakly coupled particles in the few GeV mass range. A beam dump facility using high intensity 400 GeV/c protons would be a copious source of such unknown particles in the GeV mass range. Since a high-intensity tau neutrino flux is produced by such a facility from  $D_s$  decays, the experimental apparatus foresees a neutrino detector to study the tau neutrino cross-section and discover the tau anti-neutrino. This detector is also suited to detect dark matter or any weakly interacting particle through its scattering off the



**Fig. 9.30** Layout of the SHiP project

atoms of the apparatus target. The physics case for such an experiment is widely discussed in [172].

Figure 9.30 shows the SHiP facility to be placed in the North Area. In 5 years, the facility will integrate  $2 \times 10^{20}$  400 GeV/c protons, produced by the SPS accelerator complex, impinging on a 12 interaction length ( $\lambda_{int}$ ) target made of Molybdenum and Tungsten, followed by a  $30 \lambda_{int}$  iron hadron absorber. Downstream of the target, the hadron absorber filters out all hadrons, therefore only muons and neutrinos are left. An active muon shield [173] is designed with two sections with opposite polarities to maximize the muon flux reduction: it reduces the muon flux from  $\sim 10^{10}$  down to  $\sim 10^5$  muons per spill. Approximately  $4 \times 10^{13}$  protons are extracted in each spill, designed to be 1 s long to reduce the detector occupancy. The tau neutrino detector is located downstream of the muon shield, followed by the decay vessel and the detector for hidden particles.

The neutrino detector is made of a magnetised target region, followed by a muon spectrometer. The neutrino target is based on the emulsion cloud chamber technology employed by the OPERA experiment, with a compact emulsion spectrometer, made of a sequence of very low density material and emulsion films to measure the charge and momentum of hadrons in magnetic field. Indeed, this feature would allow to discriminate between tau neutrinos and anti-neutrinos also in the hadronic decay channels of the tau lepton. The emulsion target is complemented by high resolution tracking chambers to provide the time stamp to the event and connect muon tracks from the target to the muon spectrometer. The muon spectrometer is based on the concept developed for the OPERA apparatus: a dipolar iron magnet where high precision tracking chambers provide the momentum and coarse resolution chambers provide the tracking within the iron slabs. About 10,000 tau neutrino interactions are expected to be observed in SHiP.

The emulsion target also acts as the target of very weakly interacting particles, like the dark matter, produced at the accelerator, if their mass is in the GeV range. Unlike the non-relativistic galactic dark matter producing nuclear recoils of the keV energy range, dark matter produced at the accelerator is ultra-relativistic and it could be observed through its scattering off the electrons of the emulsion target of the neutrino detector. The elastic interaction of dark matter particles with electrons produces one electron in the final state, thus mimicking elastic interaction of neutrinos that constitute the main background for this search. In [171] the sensitivity

to light dark matter shows to be very competitive with all the planned experiments in the next decade.

The SHiP Collaboration is preparing a Comprehensive Design Report to be submitted within 2018, in the framework of the Physics Beyond Colliders working group, that will be evaluated within 2020. The construction and installation is expected to start in 2021 with data taking to start in 2026.


### 9.6.3 *Projects in Applied Science*

#### 9.6.3.1 **Muon Radiography**

Muon radiography measures the absorption of cosmic-ray muons in matter, analogously to the conventional radiography that makes use of X-rays. The interaction of primary cosmic-rays with the atmosphere provides an abundant source of muons that can be used for various applications of muon radiography. Muon radiography was first proposed to determine the thickness of snow layers on a mountain [174]. The first application was realised in 1971 with the seminal work of Alvarez and collaborators searching for unknown burial cavities in Chephren's pyramid [175]. The pioneering work done in Japan for the radiography of the edifice of volcanoes by using quasi-horizontal cosmic-ray muons [176, 177] has opened new possibilities for the study of their internal structure.

Nuclear emulsions were used for the first time in 2006 for the muon radiography of the Asama volcano in Japan [178]. The main advantages of the emulsion technique are the simplicity and portability of the detector setup, and the absence of power supplies and electronic data acquisition systems, usually difficult to transport and operate on the summit of a volcano.

In 2012 an emulsion detector was installed on the Stromboli volcano to image its crater region. Despite of the strong influence that the crater area and the Sciara del Fuoco slope have on the volcanic dynamics of the Stromboli island, their internal structure is not well known because of the limited resolution of conventional geophysical methods. An emulsion detector of  $0.73 \text{ m}^2$  surface was exposed there and took data for about 5 months in 2012. Emulsion films were exposed in the form of two doublets separated by 5 mm iron slabs intended to reject the background induced by the soft component of cosmic-rays. Figure 9.31 shows an excess in the rate of muons in the crater region that is interpreted in terms of a lower density region [179]. This excess lies in the region B of Fig. 9.31, the one where the detector is sensitive to density variations, while A denotes the free sky region. In the region C, instead, the average thickness is larger than 800 m, such that the rate is too low to appreciate density changes. The data analysis provided an image of the crater area of Stromboli with a resolution of about 10 m in the center of the target area. The observed muon excess larger than 30% indicates an average density decrease along the muon path down to  $1.7 \text{ g/cm}^3$  with respect to the standard rock density




**Fig. 9.31** Excess of the muon rate seen with an emulsion detector in the crater region of the Stromboli volcano [179]. The colour scale indicates the number of muons over an angular range of  $10 \times 10 \text{ mrad}^2$  and a surface of  $0.73 \text{ m}^2$

of  $2.65 \text{ g/cm}^3$ . Further measurements campaigns are foreseen with larger detector surfaces at Stromboli as well as on other volcanoes.

An interdisciplinary project between the fields of geosciences and particle physics was also initiated. This project aims to image the bases of Alpine glaciers in three dimensions via cosmic-ray muon radiography using emulsion particle detectors. The results will be used to test the models for erosional processes and provide clues revealing how the Alpine glaciers have been shaped. The results also have an impact on society since they can be used to check the possibility of disasters caused by glacier retreats. However, studying the morphology of active Alpine glaciers has been a difficult task due to the lack of technology. Muon radiography is considered as a powerful tool to address this issue. The technique has been applied to map the bases of the Eiger Glacier and Aletsch Glaciers in Central Swiss Alps, where the Jungfrau railway tunnel provides a situation suitable for placing the detectors.

Recently, a measurement at the upper part of the Aletsch Glacier has been performed [180]. Muon detectors made of emulsion films were installed at three sites along the wall of the Jungfrau railway tunnel running through the bedrock underneath the Aletsch Glacier. The detectors had a total effective detection area of  $250 \text{ cm}^2$  for each site, and the data were collected for 47 days. The shape of the boundary between glacial ice and the bedrock at the upper part of the Aletsch Glacier was measured as shown in Fig. 9.32. This is the first successful application of this technology to a glaciated environment, which demonstrates that muon radiography can be a complementary method for determining the bedrock topography in such an environment when suitable detector sites are available. To image the bedrock topography underlying the Eiger glacier, another measurement is underway.


**Fig. 9.32** The three-dimensionally reconstructed ice–rock interface (blue surface) determined via muon radiography analysis [180]



This technique has also been applied to other fields such as investigations of archaeological sites (pyramids and tumuli). One recent success was the discovery of a large void in Khufu's Pyramid [181]. This void was first observed with emulsion detectors installed in the Queen's chamber and later confirmed with scintillator hodoscopes placed in the same chamber and with gas detectors outside the pyramid. Figure 9.33 shows that large known structures were observed as expected. In addition, an unexpected muon excess was observed, indicating that there is an additional void. This discovery demonstrated that this technique is useful for such investigations. The muon radiography technique with emulsions has been established and is further broadening our knowledge in several new fields, such as safety inspections, by looking for underground cavities or diagnosing furnace problems.

### 9.6.3.2 Medical Applications

Medical applications of the emulsion technique have also been attempted in the last decade. In the treatment of cancer by hadron-therapy, beams of carbon nuclei present therapeutic advantages over proton beams. The knowledge of the fragmentation of carbon nuclei when they interact with human tissues is important to evaluate the spatial profile of the energy deposition in the human body, thus maximizing the effectiveness in hitting the cancer with minimal damage to the




**Fig. 9.33** Two-dimensional histogram of the detected muon flux at a position (left) and the result of a simulation with the known inner structures (right) [181]. The large known structures (A: the King's chamber and B: the Grand Gallery) and a new void were observed. The colour scale indicates muons per square centimetre per day per steradian

neighbouring tissues. For this purpose, ECC detectors simulating human tissues have been realized and exposed to ion beams. The ECC technique, in fact, allows to integrate target and tracking devices in a very compact structure. This fact, together with the development of techniques of controlled fading of particle tracks in nuclear emulsions [2], has opened the way to measurements of the specific ionization over a very broad dynamic range. The application of several refreshing treatments to the emulsion films makes them sensitive to different ionization values. The combined analysis of several films allows to overcome saturation effects, so that films normally sensitive to minimum ionizing particles can be used to measure the charge of carbon ions and of their induced fragments. Details of the technique are reported in [182, 183].


Figure 9.34 shows the identification of fragments produced by the interaction of carbon ions with Lexan plates, which simulate the human body tissues given the similar electron density [183]. Such a charge identification capability allowed the measurement of the charge-changing cross-section of carbon ions with water by placing the target ECC inside a water tank [184] as well as the charge-changing cross-section of carbon ions with lexan [185].

In the framework of the FIRST (Fragmentation of Ions Relevant for Space and Therapy) experiment [186], two Emulsion Cloud Chambers were exposed to the fragments produced by a  $^{12}\text{C}$  beam (400 MeV/n) impinging on a composite target. The detectors were located in such a way to collect  $^{12}\text{C}$  fragments emitted at large angles with respect to the beam axis, as shown in the left plot of Fig. 9.35. Indeed, the characterization of secondary fragments produced by a  $^{12}\text{C}$  beam incident on a target is crucial to monitor the dose deposition inside the patient and to estimate the overall biological effectiveness due to the fragmentation of the incident beam. Data available in literature are rather scarce in this respect. Films used in this experiment were belonging to the same batch produced for the OPERA experiment.





**Fig. 9.34** Measurement of the electric charge of nuclear fragments produced by carbon ion interactions in an ECC detector [183]. Left: separation between hydrogen and helium ions. Right: separation of heavier fragments



**Fig. 9.35** Left: ECC used in the FIRST experiment. Right: angular distribution of protons [188]

The ECC structure was made of two sections [187]: the first section, consisting of six nuclear emulsion films, was meant to trigger all the incoming fragments entering the detector; the second section, consisting of 55 nuclear emulsion films interleaved with 1 mm thick lead plates, was optimised for the momentum measurements of fragments using the particle range. Given the peculiar geometry, tracks impinge on the emulsion films with rather large incident angles. Recent developments in the scanning technology [123, 141, 142] were essential to analyse these films. Almost 37,000 proton tracks were fully reconstructed and their angular and momentum spectra were measured in a wide angular range extending for the first time to more



than  $80^\circ$ , as shown in the right plot of Fig. 9.35. The momentum was also measured through the particle range as reported in [188].

The FOOT (FragmentatiOn Of Target) experiment [189] is designed to study target and projectile fragmentation processes. Target nuclei ( $^{16}\text{O}$ ,  $^{12}\text{C}$ ) fragmentation induced by 150–250 MeV proton beams will be studied via the inverse kinematic approach. The detector includes a magnetic spectrometer based on silicon pixel detectors and drift chambers, a scintillating crystal calorimeter with TOF capabilities, thick enough to stop the heavier fragments produced, and a  $\Delta E$  detector based on scintillating bars to achieve the needed energy resolution and particle identification. An alternative setup of the experiment will exploit the emulsion chamber capabilities. Dedicated emulsion chambers will be coupled with the interaction region to measure the interaction vertices within the target, tag the produced light charged fragments such as protons, deuterium, tritium and Helium nuclei, and measure their angular and momentum spectra. Given the very good identification capability of the emulsion technology for low  $Z$  fragments, the results from the emulsion chamber detectors are expected to be of particular interest also for the radio-protection in the space where Helium is relevant. The FOOT data taking is foreseen at the CNAO centre in Pavia starting from 2018 with the emulsion setup, while the electronic detectors will start data taking in 2019. Data taking in the major European laboratories such as HIT at Heidelberg and GSI at Darmstadt is also foreseen.

## References

1. Tadaaki Tani, Photographic Science, Advances in Nanoparticles, J-Aggregates, Dye Sensitization, and Organic Devices, Oxford University Press (2011), ISBN: 9780199572953.
2. T. Nakamura et al., Nucl. Instrum. Meth. **A556** (2006) 80.
3. C. Amsler et al., JINST **8** (2013) P02015.
4. P.H. Fowler, D.H. Perkins and C.F. Powell, The study of elementary particles by the photographic method, Pergamon Press (1959).
5. W.H. Barkas, Nuclear research emulsion, Academic Press, New York, 1973.
6. H. Becquerel, C.R. Academy of Science, **122** (1896) 501 and **122** (1896) 1086.
7. S. Kinoshita, Proc. Royal Society (London) **83A** (1910) 432.
8. C.F. Powell, G.P.S. Occhialini, D.L. Livesey and L.V. Chilton, Journal of Sci. Instrum. **23** (1946) 102.
9. M. Blau and J. A. De Felice, Phys. Rev. **74** (1948) 1198.
10. H. Yagoda, Phys. Rev. **79** (1950) 207.
11. P. Demers, Sci. and Indust. Phot. **23** (1952) 1.
12. B. Stiller, M. M. Shapiro and P.W. O'Dell, Rev. Sci. Instr. **25** (1954) 340.
13. C.F. Powell, Phil. Mag. **44** (1953) 219.
14. C.M.G. Lattes, H. Muirhead, G.P.S. Occhialini and C.F. Powell, Nature **159** (1947) 694.
15. C.M.G. Lattes, G.P.S. Occhialini and C.F. Powell, Nature **160** (1947) 453.
16. C.M.G. Lattes, G.P.S. Occhialini and C.F. Powell, Nature **160** (1947) 486.
17. W.F. Fry, J. Schneps, and M.S. Swami, Phys. Rev. **103** (1956) 1904–1905;  
S. Lokanathan, D.K. Robinson and S.J. St Lorant, Proc. of the Royal Society of London, Series A, Mathematical and Physical Sciences, **254**, No. 1279 (1960) 470.

18. N.A. Dobrotin, *Sov. Phys. Usp.* **2** (1960) 974;  
M.F. Kaplon and D. M. Ritson, *Phys. Rev.* **88** (1952) 386.
19. M.F. Kaplon, B. Peters, H.L. Reynolds and D.M. Ritson, *Phys. Rev.* **88** (1952) 295.
20. J. Nishimura, *Soryushiron Kenkyu* **12** (1956) 24.
21. K. Niu, *Proc. Jpn. Acad., Ser. B* **84** (2008) 1;  
K. Niu, *proc. of the I International Workshop on Nuclear Emulsion Techniques*, Nagoya, June 1998.
22. C.M.G. Lattes et al., *Prog. Theor. Phys. Suppl.* **47** (1971) 1.
23. I. Ohta, K. Kasahara, I. Mito, A. Ohsawa, T. Taira and S. Torii, *Proc. of the Int. Cosmic Ray Conf., Denver*, **3** (1973) 2250;  
M. Akashi et al., *Proc. of the Int. Cosmic Ray Conf., Munich*, **7** (1975) 2549.
24. A.V. Apanasenko et al., *Astropart. Phys.* **16** (2001) 13.
25. T.H. Burnett et al., *Nucl. Instrum. Meth.* **A251** (1986) 583.
26. K. Niu, E. Mikumo, Y. Maeda, *Prog. Theor. Phys.* **46** (1971) 1644.
27. J.E. Augustin et al., *Phys. Rev. Lett.* **33** (1974) 1406.
28. J.J. Aubert et al., *Phys. Rev. Lett.* **33** (1974) 1404.
29. S. Ogawa et.al, *Soryushiron Kenkyu* **43** (1971) 801;  
Z. Maki and T. Maskawa, *Prog. Theor. Phys.* **46** (1971) 1647;  
T. Hayashi et al., *Prog. Theor. Phys.* **47** (1972) 280.
30. K. Hoshino et al., *Proc. Int. Cosmic Ray Symp. on High Energy Phenomena* (1974) 161;  
H. Sugimoto et al., *Prog. Theor. Phys.* **53** (1975) 1541.
31. K. Hoshino et al., *Proc. of the Int. Cosmic Ray Conf., Munich*, **7** (1975) 2442.
32. N. Ushida et al., *Nucl. Instrum. Meth.* **224** (1984) 50.
33. E.H.S. Burhop et al., *Phys. Lett.* **B65** (1976) 299.
34. D. Allasia et al., *Nucl. Phys.* **B176** (1980) 13.
35. C. Angelini et al., *Phys. Lett.* **B84** (1979) 150.
36. D. Allasia et al., *Phys. Lett.* **B87** (1979) 287.
37. K. Niwa, *Physics and Astrophysics of Neutrinos*, Springer, Berlin, M. Fukugita and A. Suzuki (Eds.), (1994) 520.
38. N. Ushida et al., *Phys. Lett.* **B206** (1988) 375.
39. T. Bolton, *hep-ex/9708014*.
40. G. De Lellis, P. Miglioni and P. Santorelli, *Physics Reports* **399** (2004) 227.
41. N. Ushida et al., *Phys. Lett.* **B206** (1988) 380.
42. N. Ushida et al., *Phys. Rev. Lett.* **45** (1980) 1049;  
N. Ushida et al., *Phys. Rev. Lett.* **45** (1980) 1053;  
N. Ushida et al., *Phys. Rev. Lett.* **48** (1982) 844;  
N. Ushida et al., *Phys. Rev. Lett.* **56** (1986) 1767.
43. N. Ushida et al., *Phys. Rev. Lett.* **47** (1981) 1694;  
N. Ushida et al., *Phys. Rev. Lett.* **57** (1986) 2897.
44. S.W. Herb et al., *Phys. Rev. Lett.* **39** (1977) 252.
45. S. Aoki et al., *Nucl. Instrum. Meth.* **A274** (1989) 64.
46. K. Hoshino and G. Rosa, *Nucl. Tracks Radiat. Meas.* **12** (1986) 477.
47. S. Aoki et al., *Nucl. Tracks and Rad. Meas.* **12** (1986) 249.
48. J.P. Albanese et al., *Phys. Lett.* **B158** 186.
49. S. Aoki et al., *Prog. Theor. Phys.*, **89** (1993) 131.
50. K. Kodama et al., *Nucl. Instrum. Meth.* **A289** (1990) 146.
51. S. Aoki et al., *Nucl. Instrum. Meth. B* **51** (1990) 466.
52. K. Kodama et al., *Nucl. Instrum. Meth.* **B93** (1994) 340.
53. K. Kodama et al., *Prog. Theor. Phys.* **89** (1993) 679.
54. A. Shor, *Phys. Lett.* **B215** (1988) 375.
55. S. Aoki et al., *Phys. Lett.* **B224** (1989) 441.
56. N. Armenise et al., *Nucl. Instrum. Meth.* **A361** (1995) 497.
57. R. Tanaka and N. Ushida, *Nucl. Phys.* **A585** (1995) 323;  
S. Aoki et al., *Phys. Rev. Lett.* **65** (1990) 1729.

58. M. Danyysz et al., Nucl. Phys. **49** (1963) 121;  
D. Prowse, Phys. Rev. Lett. **17** (1966) 782.
59. S.Aoki et al., Prog. Theor. Phys. **85** (1991) 951;  
S.Aoki et al., Prog. Theor. Phys. **85** (1991) 1287.
60. H. Takahashi et al., Phys. Rev. Lett. **87** (2001) 212502.
61. K. Nakazawa, Nucl. Phys. **A585** (1995) 75.
62. H. Takahashi et al., Nucl. Phys. **A721** (2003) 951.
63. A. Ichikawa et al., Phys. Lett. **B500** (2001) 37.
64. T. Nakazawa, J. Soc. Photogr. Sci. Technol. Japan **71** No. 4 (2008) 245;  
E. Hayata et al., P07 Proposal to J-PARC (2006).
65. K. Niwa, K. Hoshino and K. Niu (1974) Proc. Int. Cos. Ray Simp. on High Energy Phenomena.  
(Cos. Ray Lab., Univ. of Tokyo) (1974) 149.
66. T. Nakano, PhD Thesis, University of Nagoya (1997).
67. G. Rosa et al., Nucl. Instrum. Meth. A **394** (1997) 357.
68. K. Kodama et al., Nucl. Instrum. Meth. **A493** (2002) 45.
69. T. Nakano, Butsuri **56** (2001) 411.
70. M. De Serio et al., Nucl. Instrum. Meth. **A554** (2005) 247.
71. M. De Serio et al., Nucl. Instrum. Meth. **A512** (2003) 539.
72. K. Kodama et al., Nucl. Instrum. Meth. **A574** (2007) 192.
73. K. Kodama et al., Rev. Sci. Instrum. **74** (2003) 53.
74. L. Arrabito et al., JINST **2** (2007) P02001.
75. E. Eskut et al., Nucl. Instrum. Meth. A **401** (1997) 7.
76. S. Aoki et al., Nucl. Instrum. Meth. A **447** (2000) 361.
77. E. Eskut et al., Nucl. Phys. **B793** (2008) 326.
78. A. Kayis-Topaksu et al., Phys. Lett. **B539** (2002) 188.
79. G. Onengut et al, Phys. Lett. **B613** (2004) 105;  
A. Kayis-Topaksu et al., Phys. Lett. **B555** (2003) 156;  
A. Kayis-Topaksu et al., Phys. Lett. **B575** (2003) 198.
80. A. Kayis-Topaksu et al., Eur. Phys. J. **C52** (2007) 543.
81. G. De Lellis et al., Nucl. Phys. **B763** (2007) 268.
82. K. Kodama, et al., FERMILAB-PROPOSAL-0803 (Oct 1993).
83. A. Ereditato, G. Romano and P. Strolin, Nucl. Phys. **54B** (1997) 139.
84. A.S.Ayan et al., CERN-SPSC/97-5, SPSC/I213 (1997).
85. K. Kodama et al., Nucl. Instrum. Meth. **A516** (2004) 21.
86. K. Kodama et al., Phys. Lett. **B504** (2001) 218.
87. K. Kodama et al., Phys. Rev. **D78** (2008) 052002.
88. T. Nakano, J. Soc. Photogr. Sci. Technol. Japan **71** No. 4 (2008) 229.
89. N. Armenise et al., Nucl. Instrum. Meth. **A551** (2005) 261.
90. L. Arrabito et al., Nucl. Instrum. Meth. **A568** (2006) 578.
91. L. Arrabito et al., JINST **2** (2007) P05004.
92. I. Kreslo et al., JINST **3** (2008) P04006.
93. K. Borer et al., Nucl. Instrum. Meth. **A566** (2006) 327.
94. Y. Fukuda et al., Phys. Rev. Lett. **81**, 1562 (1998);  
W.W.M. Allison et al., Phys. Lett. **B449** (1999) 137;  
M. Ambrosio et al., Phys. Lett. **B517** (2001) 59.
95. M.H. Ahn et al., Phys. Rev. **D74**, 072003 (2006);  
D.G. Michael et al., Phys. Rev. Lett. **97** (2006) 191801.
96. K. Niwa, Contrib. to the “Snowmass ’94” conference on particle and nuclear astrophysics and cosmology in the next millennium, Snowmass, 1994.
97. A. Ereditato, K. Niwa and P. Strolin, INFN/AE-97/06, Nagoya DPNU-97-07, Jan 27th 1997.
98. A. Ereditato, K. Niwa and P. Strolin, Nucl. Phys. Proc. Suppl. **66** (1998) 423.
99. M. Guler et al., CERN-SPSC-2000-028.
100. R. Acquafredda et al., JINST **4** (2009) P04018.
101. A. Anokhina et al., JINST **3** (2008) P07005.

102. S. Miyamoto et al., Nucl. Instrum. Meth. **A575** (2007) 466.
103. N. Agafonova et al. [OPERA Collaboration], Eur. Phys. J. **C74** (2014) no.8, 2986.
104. A. Kayis-Topaksu et al., New J. Phys. **13** (2011) 093002.
105. N. Agafonova et al. [OPERA Collaboration], New J. Phys. **14** (2012) 013026.
106. N. Agafonova et al. [OPERA Collaboration], Phys. Lett. B691 (2010) 138.
107. N. Agafonova et al. [OPERA Collaboration], JHEP 1311 (2013) 036.
108. N. Agafonova et al. [OPERA Collaboration], Phys. Rev. **D89** (2014) no.5, 051102.
109. N. Agafonova et al. [OPERA Collaboration], PTEP 2014 (2014) no.10, 101C01.
110. N. Agafonova et al. [OPERA Collaboration], Phys. Rev. Lett. 115 (2015) no.12, 121802.
111. H. Ishida et al., PTEP 2014 (2014) N. 9, 093C01.
112. A. Longhin et al., IEEE Trans. Nucl. Sci. **62** (2015) 2216–2225.
113. C. Patrignani et al., Review of Particle Physics, Chin. Phys. **C40** (2016) N. 10, 100001.
114. N. Agafonova et al., [OPERA Collaboration], Physical Review Letters 120 (2018) 211801.
115. N. Agafonova et al. [OPERA Collaboration], JHEP **07** (2013) 004.
116. N. Agafonova et al., [OPERA Collaboration], arXiv:1803.11400
117. N. Agafonova et al. [OPERA Collaboration], JHEP 1506 (2015) 069.
118. NVIDIA CUDA web page: <http://www.nvidia.com/cuda>.
119. T. Fukuda et al., JINST 9 (2014) P12017.
120. A. Ariga and T. Ariga, JINST **9** (2014) P04002.
121. A. Alexandrov, et al., JINST **10** (2015) P11006
122. M. Yoshimoto, T. Nakano, R. Komatani and H. Kawahara, PTEP **10** (2017) 103.
123. A. Alexandrov et al., Nature Scientific Reports **7** (2017) 7310.
124. N. Natsume al., Nucl. Instr. Meth. A **575** (2007) 439.
125. T. Naka et al., Nucl. Instrum. Meth. A **718** (2013) 519.
126. T. Ariga et al., JINST **11** (2016) P03003.
127. A. Nishio et al., PoS KMI2017 (2017) 057.
128. J. Nishimura, Adv. Space Res. **30** No. 5 (2002) 1071..
129. K. Kodama et al., Adv. Space Res. **37** (2006) 2120.
130. S. Aoki et al., J. Soc. Photogr. Sci. Technol. Japan **71** No. 4 (2008) 256.
131. K. Ozaki et al., Nucl. Instrum. Meth. **A833** (2016) 165.
132. S. Takahashi, S. Aoki for GRAINE collaboration, Adv. Space Res. (in press), <https://doi.org/10.1016/j.asr.2017.08.029>.
133. S. Takahashi et al., PTEP 2015 **4** (2015) 043H01.
134. The P7SOURCE\_V6 Instrument Response Functions, [https://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat\\_Performance.htm](https://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm).
135. S. Takahashi et al., Nucl. Instrum. Meth. **A620** (2010) 192.
136. H. Rokujo et al., Nucl. Instrum. Meth. **A701** (2013) 127.
137. S. Takahashi et al., PTEP 2016 **7** (2016) 073F01.
138. K. Ozaki et al., JINST **10** (2015), P12018.
139. J. B. R. Battat et al., Physics Reports **662** (2016) 1.
140. A. Alexandrov et al., LNGS-LOI 48/15, arXiv:1604.04199.
141. A. Alexandrov et al., JINST **10** (2015) P11006.
142. A. Alexandrov et al., JINST **11** (2016) P06002.
143. T. Naka et al., EAS Publ. Ser. **53** (2012) 51.
144. M. Kimura and T. Naka, Nucl. Instrum. Meth. A **680** (2012) 12.
145. H. Tamaru et al., Applied Phys. Lett. **80** (2002) 1826.
146. A. Alexandrov et al., Astroparticle Physics **80** (2016) 16.
147. N. Naganawa et al., Phys. Procedia **88** (2017) 224.
148. K. Mishima et al., Nucl. Instrum. Meth. **A600** (2009) 342.
149. N. Naganawa et al., PoS KMI2017 (2017) 077.
150. N. Naganawa, T. Ariga, S. Awano, M. Hino, K. Hirota, H. Kawahara, M. Kitaguchi, K. Mishima, H. M. Shimizu, S. Tada, S. Tasaki, A. Umemoto, arXiv:1803.00452.
151. T. Nakazawa, Few-Body Systems, 2013, Volume 54, Issue 7-10, pp1279–1282.
152. G. Drobychev et al., AEgIS Proposal, CERN-SPSC-P-334 (2007).

153. S. Aghion et al. [AEgIS Collaboration], JINST 8 (2013) P08013.
154. The Antiproton Decelerator webpage, <https://home.cern/about/accelerators/antiproton-decelerator>.
155. T. Ariga et al. [AEgIS Collaboration], Int. J. Mod. Phys. Conf. Ser. 30 (2014) 1460268.
156. S. Aghion et al. [AEgIS Collaboration], Nature Communications 5 (2014) 4538.
157. S. Aghion et al. [AEgIS Collaboration], JINST 12 (2017) P04021.
158. A. Ferrari, P.R. Sala, A. Fasso and J. Ranft, CERN-2005-010, CERN, Geneva Switzerland, (2005).
159. T.T. Bohlen et al., Nucl. Data Sheets 120 (2014) 211.
160. S. Aghion et al., JINST 11 (2016) P06017.
161. K. Suzuki et al., Prog. Theor. Exp. Phys. (2015) 053C01.
162. K. Abe et al. [Hyper-Kamiokande Proto-Collaboration], PTEP 2015 (2015) 053C02.
163. R. Acciarri et al. [DUNE Collaboration], FERMILAB-DESIGN-2016-01 (2016).
164. K. Yamada et al., PTEP 2017 6 (2017) 063H02.
165. K. Abe et al. [T2K Collaboration], Nucl. Instrum. Meth. A **659** (2011) 106.
166. T. Fukuda et al., PTEP 2017 6 (2017) 063C02.
167. T. Fukuda on behalf of the NINJA Collaboration, PoS KMI2017 (2017) 012.
168. S. Aoki et al. [DsTau Collaboration], CERN-SPSC-2017-029, SPSC-P-354 (2017).
169. A. Akmete et al. [SHiP Collaboration], CERN-SPSC-2017-033, SPSC-EOI-017 (2017).
170. H. Dijkstra and T. Ruf, CERN-SHiP-NOTE-2015-009. <http://cds.cern.ch/record/2115534/files/SHiP-NOTE-2015-009.pdf>
171. M. Anelli et al., CERN-SPSC-2015-016, SPSC-P-350, arXiv:1504.04956.
172. S. Alekhin et al., Rept. Prog. Phys. **79** (2016) no.12, 124201.
173. A. Akmete et al., JINST **12** (2017) no.05, P05011.
174. E.P. Georg, Commonw. Eng., July 1955.
175. L.W. Alvarez et al., Science **167** (1970) 832.
176. K. Nagamine et al., Nucl. Instrum. Meth. A**356** (1995) 585.
177. H.K.M. Tanaka et al., Nucl. Instrum. Meth. A**507** (2003) 657;  
H.K.M. Tanaka et al., Nucl. Instrum. Meth. A**555** (2005) 164.
178. H.K.M. Tanaka et al., Earth and Planetary Science Letters **263** (2007) 104;  
H.K.M. Tanaka et al., Nucl. Instrum. Meth. A**575** (2007) 489;  
H.K.M. Tanaka et al., Geophysical Research Letters **34** (2007) 389.
179. V. Tioukov et al., Annals of Geophysics 60 (2017) N. 1 S0111.
180. R. Nishiyama, A. Ariga, T. Ariga, S. Kaser, A. Lechmann, D. Mair, P. Scampoli, M. Vladymyrov, A. Ereditato, F. Schlunegger, Geophysical Research Letters 55933 (2017).
181. K. Morishima et al., Nature **552** (2017) 386–390.
182. T. Toshito et al., Nucl. Instrum. Meth. A**556** (2006) 482.
183. G. De Lellis et al., JINST **2** (2007) P06004.
184. T. Toshito et al., Phys. Rev. C**75** (2007) 054606.
185. G. De Lellis et al., Nuclear Physics A **853** (2011) 124–134.
186. R. Pleskac et al., Nucl. Instrum. Meth. A**678** (2012) 130–138.
187. A. Alexandrov et al., Meas. Sci. Technol. **26** (2015) 094001.
188. A. Alexandrov et al., JINST **12** (2017) P08013.
189. S. Argiro et al., The 26th International Nuclear Physics Conference, 11–16 September, 2016, Australia, Proceedings of Science, INPC2016, 128.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 10

## Signal Processing for Particle Detectors



V. Radeka

### 10.1 Introduction

This chapter covers the principles and basic limits of signal processing for detectors based on measurements of charge induced on predominantly capacitive electrodes. While this presents a very limited scope, it already includes a broad range of different detector technologies employed in experiments in several areas of science and in various imaging devices. Detector technologies of interest involve semiconductors, gas and liquid ionization media as well as photo detectors converting (scintillation light) into photo-electrons or ionization. One class of detectors not considered here are bolometric detectors (see Sect. 10.4).

The literature cited in this chapter is twofold: the textbooks, tutorials and review articles which may serve to provide a systematic introduction to the reader, and references to journal articles describing specific applications and technological solutions. The former, while very good and useful, are unfortunately few, the latter are only a small selection from the vast body of journal articles and conference records. The former are listed first [1–10], and the latter are cited along with the material presented in this chapter.

For many detectors, particularly very large scale detectors used at colliding beam machines in particle and nuclear physics, systems aspects require most of the attention in the design. In high precision measurements of energy, time arrival or position of the incident particle or photon the noise introduced in the measurement is of primary interest. Each area of science may impose greatly different requirements on various performance parameters of the detector and signal processing. A silicon pixel detector for particle tracking at a high luminosity collider requires very short

---

V. Radeka (✉)  
Brookhaven National Laboratory, Upton, NY, USA  
e-mail: [radeka@bnl.gov](mailto:radeka@bnl.gov)

pulse shaping (a few tens of nanoseconds) and it can tolerate a noise level of several hundred electrons rms. Consequently, a leakage (dark) current contributing shot noise may be  $1 \text{ nA/cm}^2$  or more. A silicon detector for x-ray spectroscopy in photon science must be read out with a total noise of less than ten electrons rms with a shaping time of the order of  $1 \text{ }\mu\text{s}$ . This allows a leakage current of only  $\sim 10 \text{ pA/cm}^2$  or less. This chapter should enable the reader to evaluate the relations among such detector and readout parameters.

Signal processing for particle detectors rests on understanding of signal formation and of the sources of noise and their effects on measurement accuracy. Signal formation in detectors is based on electrostatics and it is calculated relatively easily starting from the Shockley-Ramo theorem [11, 12]. It gets more involved in multi-electrode detectors [13] and in crosstalk analysis. Signal processing with time-invariant systems has been extensively covered in the literature and is well understood.

Most innovations in signal processing in recent years have been in circuit implementations using monolithic CMOS technology. This technology has brought about a significant shift in the circuit concepts to time variant circuits due to the use of switched capacitance circuits for which CMOS transistors are well suited.

The noise analysis of time variant circuits brings up the question of whether to perform the analysis in frequency domain or in time domain. Analysis in frequency domain provides sufficient insight into time invariant circuits and it has been used in most of the literature. Frequency domain is less well suited for time variant circuit analysis as it does not provide much insight into the system transfer function. This is where the concept of the weighting function and Campbell's theorem [14] provide the tools which are simpler to use and provide more insight. While both analytical methods provide the same results in noise calculations, we note that particle detector signals are best described and are observed in the time domain, and so is the system response and the weighting function. In contrast, the noise analysis of narrow band circuits is best done in the frequency domain. The time domain analysis is based on the representation of noise as a random sequence of elementary impulses. In spite of our thinking and observing in the time domain, the device and circuit noise sources are customarily characterized in the frequency domain, e.g., we talk about the "white noise", " $1/f$  noise", etc. Thus we switch our thinking between the two (Fourier transform related) domains depending on which one provides better insight and is easier to analyze in a particular case.

This chapter is intended to provide some insight into detector signal processing. Detailed circuit design, particularly of the monolithic circuits, has been rapidly developing and there has been a proliferation of publications. Advanced simulation tools are being used for noise analysis. Sometimes such an analysis provides numerical results without providing much insight into the role of various noise sources and into the overall weighting function of the signal processing chain. We concentrate here on the interpretation of the weighting function and on those aspects of signal processing and noise that have been less covered in the literature such as the induced signals in multi-electrode (strip and pixel) detectors, the " $kTC$ " noise, correlated sampling and basic properties of low noise charge amplifiers. Induced




signals are determined using the “weighting field” concept, and noise analysis is based on the “weighting function” concept. The former is based on electrostatics and the latter on superposition of noise impulse contributions to the variable (current, charge) measured by the readout system.

## 10.2 Charge Collection and Signal Formation in Detectors

### 10.2.1 Current Induced by the Moving Charge and the Weighting Field Concept

Figure 10.1 illustrates the Shockley-Ramo theorem for induced signals, current and charge.  $E_w$  is the weighting field in units of 1/cm, and it is a measure of electrostatic coupling between the moving charge and the sensing electrode. The procedure to calculate the induced current as a function of time is as follows. First, the weighting field is determined by solving Poisson’s equation analytically or numerically assuming unity potential on the sensing electrode of interest and zero potential on all other electrodes. Next, the velocity of the moving charge,  $v = dx/dt$ ,



**Fig. 10.1** Weighting potential (blue lines) and weighting field lines for planar strip electrode readout

as a function of position is determined from the operating (applied) field on the detector. This gives the induced current as a function of the position of the moving charge,

$$i = -q \vec{E} \cdot \vec{v}. \quad (10.1)$$

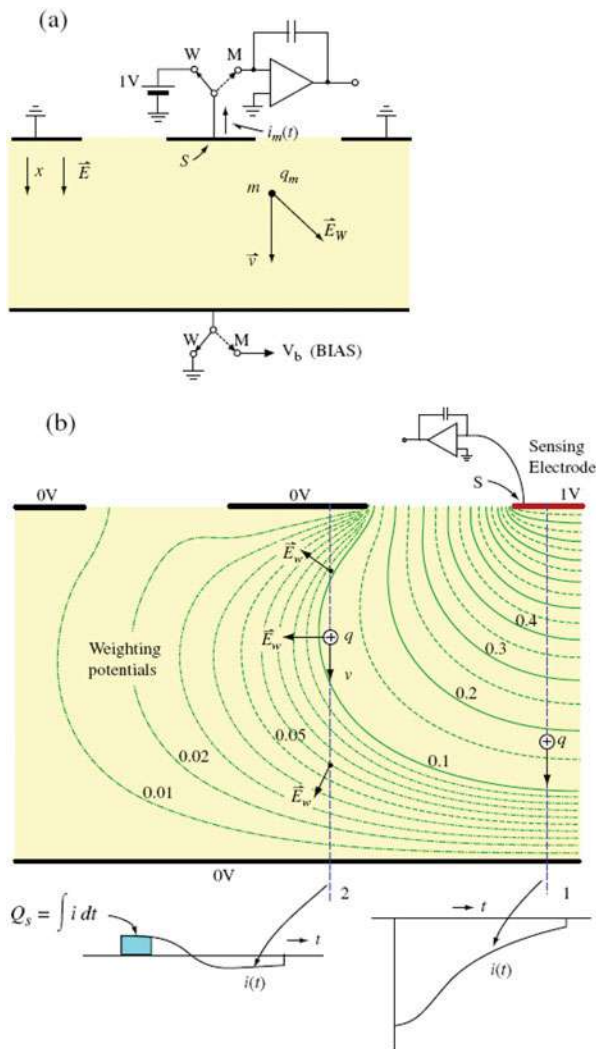
Third, the position of the moving *charge* as a function of time is determined by solving the equations of motion. This is necessary in the case of ballistic motion of charge, but it is simple in the case of transport by drift as the charge carriers follow the applied electric field. If we are interested only in the total induced charge and not in the waveforms, the induced charge is simply given by the difference in the weighting potentials between any two positions of the moving charge,

$$\begin{aligned} Q_s &= \int i dt = + \int \vec{E}_w d\vec{x}, \\ Q_s(1, 2) &= q (V_{w2} - V_{w1}). \end{aligned} \quad (10.2)$$

An example of the weighting-field (potential) profiles is illustrated by the plot of equipotential lines for planar geometry with a strip sensing electrode. The operating (applied field) in this case is uniform and perpendicular to the electrodes. The weighting field map is in general quite different from that of the operating field; the two field maps are identical only in some special cases. The minus sign in Ramo's equation (Eq. 10.1) for the induced current results from the arbitrary assumption of induced current *into* the electrode being positive.

The sketch in Fig. 10.2a shows conceptually how the weighting field (potential) is defined: the sensing electrode is connected to unity potential, and all other electrodes to zero potential. The equipotential lines in Fig. 10.2b illustrate the solution for this case, showing two strips next to the sensing electrode. A great variety of results for the induced current and charge may arise in an electrode structure, such as this, depending on the particle type detected (distribution of ionization) and on the ratio of the charge observation measurement (or integration) time and the charge carrier transit time. The current waveforms shown are drawn qualitatively for a simple example. The operating field is assumed uniform and perpendicular to the electrode planes. Charge  $q_m$  traversing the full distance between the electrodes along line 1 is observed as  $Q_1 = -q_m$ , while the current decreases with distance from the sensing electrode 1, as the electrostatic coupling decreases. For a charge moving along line 2, the induced charge (i.e., the difference between the weighting potentials) is zero, if the measurement time is longer than the transit time. For a short measurement time a net induced charge is observed. The induced current waveform (the "crosstalk signal") is bipolar, since the weighting field direction changes along the path.


Figure 10.3 illustrates a simple case where the real (operating) electric field and the weighting field have the same form  $= 1/d$ . The induced current waveforms shown are for a semiconductor detector with different electron and hole mobilities. For extended ionization the waveforms result from superposition of the waveforms for localized ionization, and the currents decrease as the carriers arrive at the electrodes from different initial positions within the bulk of the detector.



**Fig. 10.2** Definition of the weighting potential: Solution of the Laplace equation for unity potential at the sensing electrode and zero potential at all other electrodes. From Radeka [9] Annual Reviews, [www.annualreviews.org](http://www.annualreviews.org), by permission


### 10.2.2 Induced Current and Charge in Strip and Pixel Electrodes: Shielding Effect

The shielding effect is proportional to the ratio of the distance between the planar electrodes and the strip width (i.e., pixel radius). The shielding effect is more pronounced for pixels than for strips. The result of these configurations is that the



**Fig. 10.3** Induced currents in infinite planar electrodes for localized and extended ionization

**2D case (strip)**



**Fig. 10.4** Weighting field (potential) for strip electrode configuration

signal charge (integral of the induced current) is independent of the position of the origin of ionization for most of the volume of the detector except near the readout electrodes. This effect is used in detectors where only electrons are collected during the integration time, such as Cadmium Zinc Telluride (CZT), and some gas and noble liquid detectors. To illustrate this, histograms are shown in Figs. 10.4 and 10.5 for a strip and pixel illuminated by a beam of penetrating x-rays absorbed uniformly through the detector.




Fig. 10.5 Weighting field (potential) for pixel electrode configuration

### 10.2.3 Weighting Potential and Induced Charge in Co-Planar Electrodes


Coplanar grid readout was introduced for unipolar charge sensing by Luke [15] and it is commonly used with Cadmium Zinc Telluride (CZT) detectors. In such materials only electrons are collected (from the ionization produced by gamma rays or x-rays), the holes suffering from very low mobility and trapping. With parallel plane electrodes the induced charge for single carrier collection is dependent on the position (depth) where the ionization took place. In the coplanar grid concept one set of alternate strips is biased slightly more positively with respect to the other set of strips. This results in a drift field such that the signal electrons are collected on one set of strips only, Fig. 10.6. The weighting potential (field) for both sets of strips is identical, Fig. 10.7. The induced charges and currents are quite different, Figs. 10.8 and 10.9. Their respective differences are independent of the position, as can be concluded by following the weighting potential plot from any point on the planar sloped part of the plot to unity weighting potential for the collecting electrode, and (across the saddle) to zero potential for the non-collecting electrode, Fig. 10.9.

## 10.3 Noise: Origin and Properties


### 10.3.1 Noise Process and Noise Variance

The basis of a noise process can be represented as a sequence of randomly generated elementary impulses that has a Poisson distribution in time and mean rate of occurrence  $\langle n \rangle$ . Upon acting on a physical system with impulse response much

**Fig. 10.6** Drift field for coplanar electrodes




**Fig. 10.7** Weighting potential for coplanar electrodes




longer than  $\langle n \rangle^{-1}$ , the characteristic noise waveforms (e.g., such as those we observe on an oscilloscope) are produced as a superposition of responses to individual impulses. The noise variance at the output of the physical system (a simple RC filter or a complete readout system) is calculated by using Campbell's theorem [14], which states that the sum of mean square contributions of all preceding impulses equals the variance. The expressions for the variance are given after subtracting the mean value. The variance is determined by the rate of impulses  $\langle n \rangle$ , their area  $q$

**Fig. 10.8** Induced charges in coplanar electrodes:  
 $Q_1 - Q_2 = \text{const}$   
 independently of  $x_0$



**Fig. 10.9** Induced currents in coplanar electrodes



(charge), and by the impulse response  $h(t)$ , i.e., the weighting function  $w(t)$  of the measurement system, the preamplifier and the subsequent readout chain,

$$\sigma^2 = \langle n \rangle q^2 \int_{-\infty}^{\infty} h^2(t) dt = \langle n \rangle q^2 \int_{-\infty}^{\infty} w^2(t) dt . \quad (10.3)$$

The noise variance is determined by the noise process, the rate of impulses  $\langle n \rangle$ , and their area  $q$  (charge), and by the impulse response  $h(t)$ . If we measure the variance and the  $h(t)$ , we can determine  $\langle n \rangle q^2$ , but we cannot determine  $\langle n \rangle$  and  $q$ . Only when randomly generated carriers move in one direction, which results in a mean current  $I_0 = \langle n \rangle q$ , can the rate and the charge of impulses be determined from  $\sigma^2$  and  $I_0$ . It is shown in Ref. [9] that  $\langle n \rangle q^2$  equals the mathematical (two-sided) noise current spectral density, whereas  $\overline{i_n^2} = 2 \langle n \rangle q^2$  equals the physical (single-sided) one, to be used in calculations of the equivalent noise charge (ENC) in Sect. 10.4.2.



10.3.2 A Model for Generation of Noise Spectra

Almost any noise spectrum can be generated from a random sequence of impulses (i.e., white noise with “infinite bandwidth”) by using an appropriate filter, as illustrated in Fig. 10.10. These impulses may be either of only one polarity or of both polarities (current thermally generated in a p-n junction under reverse bias in the former, and with zero bias in the latter case). The mean value depends on the impulse polarities, but the variance does not.

“Infinite bandwidth” implies a noise spectrum which is flat over the frequency range where our measurement system has a non-zero response. Simple integration of white noise results in “random walk” with  $1/f^2$  spectrum. An elementary impulse response for generation of this noise is the step function  $U(t)$ . Generation of  $1/|f|$  noise is somewhat more elaborate. It requires fractional integration of order one half. The impulse response of the transforming filter is  $U(t)t^{1/2}$ , as shown in the figure. The basic feature of any noise generating mechanism for low frequency divergent noises is an “infinitely long memory”, i.e., very long memory, for individual independent elementary perturbations.

For a discussion of the basics of power-law spectra and of fractional integration see Ref. [16].

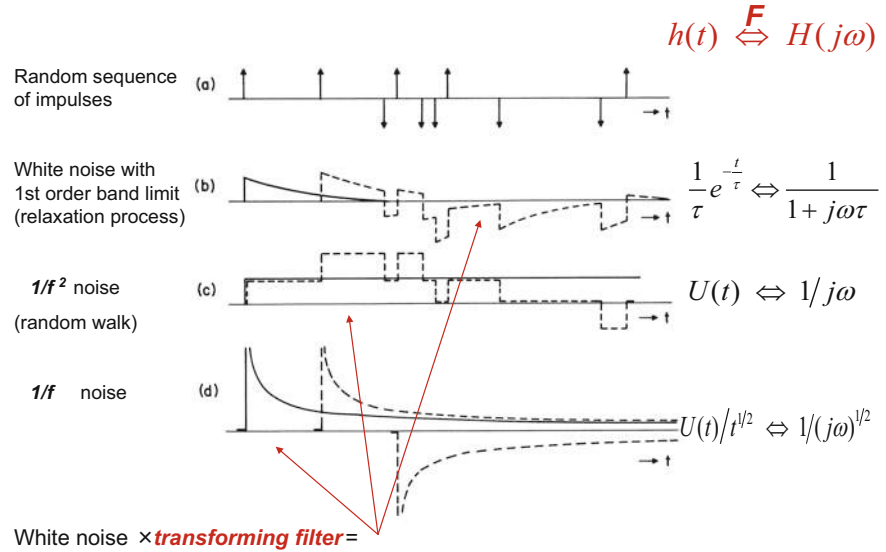
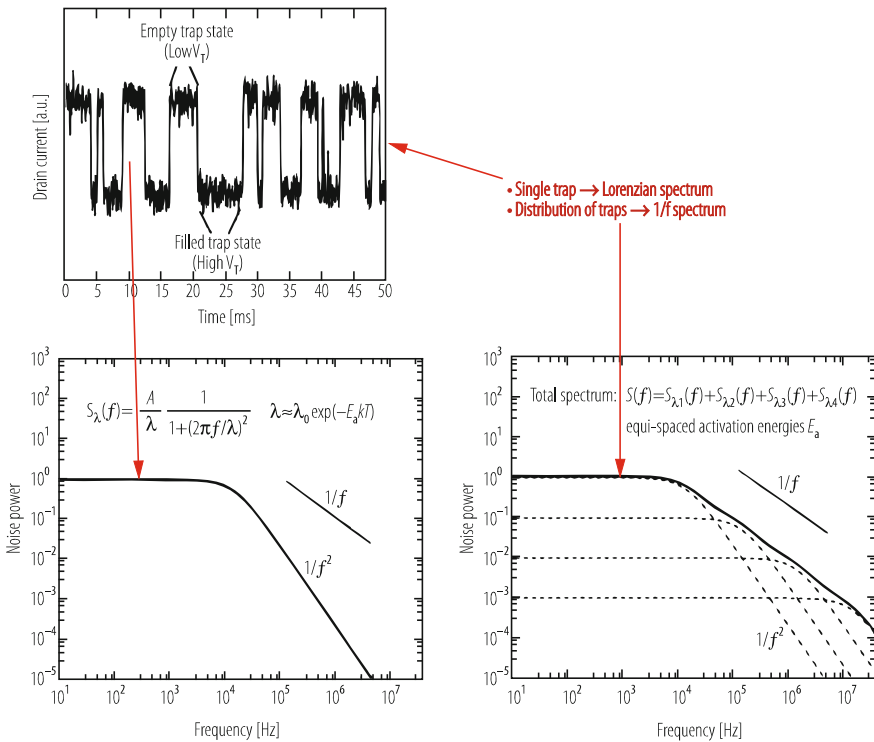


Fig. 10.10 Generation of some basic noise spectra from white noise by a transforming filter

### 10.3.3 Random Telegraph Noise and $1/f$ Noise

A noise spectrum very close to  $1/|f|$  can be generated by superposition of relaxation processes with uniform distribution of life times, as illustrated in Fig. 10.11. The relaxation process is described by  $U(t)\exp(-t/\tau)$ , which represents a step change with exponential decay. Trapping-detrapping in semiconductors is one such possible mechanism for generation of  $1/|f|$  noise. Since a simple RC integrator has the same response, a hardware filter which transforms white noise into  $1/|f|$  noise can be made requiring about one time constant (one RC circuit) per decade of frequency, as shown in Ref. [16].

A single trap in a very small (minimum size) MOS transistor results in a drain current modulation known as random telegraph noise (RTS). This noise presents a limit to sensitivity in imaging arrays with a pixel capacitance of a few femtofarads and other noise sources reduced to a few electrons rms. There is extensive literature on RTS, e.g., Refs. [17–20].



**Fig. 10.11** Generation of random telegraph noise (left) and of  $1/f$  noise (right) by trapping-detrapping. Adapted from Compagnoni et al. [17]

$1/f$  noise is one of the fractal processes, and its waveform preserves the same features independently of the time scale [16]. Another expression of this is independence of the measurement variance on the time scale of the measurement as long as the ratio of the high frequency and the low frequency cutoffs remains constant. As the bandpass moves along the frequency spectrum the spectral density integral (i.e., the measurement variance),

$$\sigma^2 \propto \int_{f_l}^{f_h} \frac{df}{f} = \ln\left(\frac{f_h}{f_l}\right) \quad (10.4)$$

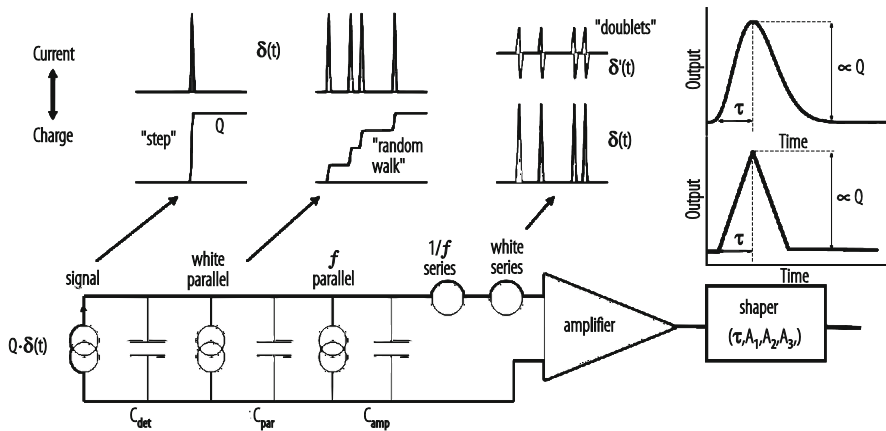
remains constant for  $f_h/f_l = \text{const.}$  In detector pulse processing it is well known that the contribution of  $1/f$  noise to the equivalent noise charge (ENC) remains independent of the shaping time, as will be shown in the following Sect. 10.4.2. Physical mechanisms of  $1/f$  noise are discussed in Ref. [21].

## 10.4 Noise in Charge Measurements

### 10.4.1 Sources of Noise in Charge Amplifiers

Principal noise sources in charge amplifiers and an equivalent diagram for calculation of the equivalent noise charge (ENC) are shown in Fig. 10.12.

Two elementary noise generators are included in the equivalent circuit, a series noise voltage generator representing the noise in the amplifying device, and a parallel noise current generator representing various noise sources not inherent to



**Fig. 10.12** An illustration of noise sources in charge amplifiers, as a Poisson sequence of elementary *current* pulses into the input capacitance, or *voltage* pulses on the capacitance

amplification (detector leakage current noise, parallel resistor noise, etc.). Both types of noise are assumed to have a white spectrum. Two forms of presentation in terms of a sequence of random pulses are shown, as charge (or voltage) at the input of the amplifier, and as a current injected into the input capacitance (comprised of the detector + amplifier parasitic capacitances). The presentation of the series noise in terms of a current into the detector input is the derivative of the charge (voltage) representation. The sequence of voltage impulses representing the amplifier series noise thus corresponds to an equivalent sequence of current doublets (derivatives of delta function) injected at the detector. The parallel noise is by its origin a current source in parallel with the detector, and it is presented by a sequence of impulses (delta functions). *It is this difference in the location of the two white noise sources with respect to the detector capacitance that makes their apparent noise spectra and their effect on the measurement quite different.* ENC due to the former is inversely proportional to the square root of the peaking time, and proportional to it due to the latter. The series  $1/|f|$  noise contribution to ENC is independent of the peaking time, as indicated in Fig. 10.14. The  $1/|f|$  noise due to a dissipative dielectric depends on the dielectric loss factor  $\tan(\delta)$ , as will be discussed in Sect. 10.4.7. It can be significant with detector-amplifier connections on glass fibre circuit boards.

### 10.4.2 Equivalent Noise Charge (ENC) Calculations

Calculation of equivalent noise charge (ENC) for a signal processing chain described by a weighting function  $w(t)$  is summarized in the following.

The noise calculation is performed in the time domain by using Campbell's theorem Eq. (10.3), that is by superposition of effects of all random *current* impulses illustrated in Fig. 10.12. The weighting function is normalized to unity so that the definition of *ENC* is *the noise charge which produces an output of the same magnitude as an impulse signal of equal charge*. For calculation of ENC due to the series *voltage* noise, we will use the representation in terms of an equivalent current generator connected in parallel with the detector. This requires differentiation of the sequence of voltage impulses. Each resulting doublet  $C_{in}\delta'(t)$  acting upon the weighting function  $w(t)$  produces by convolution  $C_{in}w'(t)$ . The equivalent noise charge (ENC) is then given by,

$$ENC^2 = \frac{1}{2}e_n^2 C_{in}^2 I_1 + \pi C_{in}^2 A_f I_2 + \frac{1}{2}i_n^2 I_3, \quad (10.5)$$

where  $\overline{e_n^2} = 4kTR_s$  is the physical (single-sided) noise *voltage* spectral density for series noise in  $V^2/\text{Hz}$  expressed in terms of an equivalent series noise resistance. The second term is due to the series  $1/f$  noise. The  $1/f$  noise physical spectral density is defined as  $A_f/f$  in  $[V^2/\text{Hz}]$ . The third term is due to the parallel noise, where  $i_n^2 = 2qI_0 = 4kT/R_p$  is the physical noise *current* spectral density due to either a current or a resistance in parallel with the detector.

$I_1$ ,  $I_2$ ,  $I_3$  are the noise integrals for the *series (voltage)* white noise and the *1/f noise*, and for the *parallel (current)* noise, respectively. The integrals are derived in the time domain from Campbell's theorem Eq. (10.3) and expressed in the frequency domain using Parseval's theorem [9],

$$I_1 = \int_{-\infty}^{\infty} [w'(t)]^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H(j\omega)|^2 \omega^2 d\omega = \frac{A_1}{\tau}, \quad (10.6)$$

$$I_2 = \int_{-\infty}^{\infty} [w^{(1/2)}(t)]^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H(j\omega)|^2 \omega d\omega = A_2, \quad (10.7)$$

$$I_3 = \int_{-\infty}^{\infty} [w(t)]^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H(j\omega)|^2 d\omega = A_3 \tau, \quad (10.8)$$


where  $\tau$  is the time width parameter of the weighting function, either the peaking time of the function, or some characteristic time constant of the filter implemented in hardware (or software in case of digital filtering).

Noise contributions for both types of white noise due to various segments (piece-wise linear approximation) of the weighting function are shown in Fig. 10.13 (expressions for integrals  $I_1$  and  $I_3$ ). In these calculations, either the *impulse response*  $h(t)$  of the system or the *weighting function*  $w(t)$  (the mirror image of the impulse response) can be used for time-invariant systems. For time-variant (gated or switched) systems, only a weighting function describes the performance correctly, while an apparent impulse response (waveform at the output) is not correct and can be misleading. Steepest parts of the weighting function contribute most to  $ENC_s$ , as they correspond to larger bandwidth. Flat parts do not contribute anything. In contrast,  $ENC_p$  is largely proportional to the width of the weighting function where it has any significant value.


A bipolar weighting function, i.e., impulse response  $h(k)$ , as shown in Fig. 10.13, with equal lobes would result in square root of two higher ENC than for a unipolar function (single lobe). If the amplitude of the second lobe is less than one half, its rms noise contribution becomes small (<12%).

The half-order integral  $I_2$  for 1/f noise is not amenable to such a simple interpretation, and it will be discussed in Sect. 10.4.6.

Equations (10.6, 10.7, and 10.8) provide an insight into the general behaviour of signal processing systems with respect to noise. The ENC due to the series white amplifier noise is proportional to the slope ( $\sim 1/t$ ) of the weighting function and therefore proportional to the bandwidth of the system. The ENC due to parallel white noise is proportional to the width of the weighting function and therefore to the overall integration time. If the weighting function form remains constant the ENC due to 1/f noise is independent of the width of the weighting function, since the ratio of the high frequency cutoff and low frequency cutoff remains constant,



**Fig. 10.13** An illustration of noise contributions for both locations of white noise sources (series and parallel) due to various segments of the weighting function. Such a piece-wise linear approximation of the weighting function provides an estimate of the noise within a few percent of accurately computed integrals  $I_1$  and  $I_2$



**Fig. 10.14** An example of general behaviour of equivalent noise charge (ENC) as a function of the width parameter  $\tau$  of the weighting function.  $1/f$  noise raises the noise minimum but does not affect its position on the time scale

Eq. (10.4). This is illustrated in Fig. 10.14. From Eqs. (10.5, 10.6, and 10.8) the optimum width parameter of the weighting function is given by,

$$\tau_{opt} = (R_S R_P)^{1/2} C_{in} \left( A_1 / A_3 \right)^{1/2}, \quad (10.9)$$


and it is not affected by  $1/f$  noise.

### 10.4.3 Weighting Function

The concept of weighting function is very useful for time domain noise analysis of time variant, sampled and switched systems. The role of “pulse shaping,” “signal filtering,” or “signal processing” is to minimize the measurement error with respect to the noise, various baseline offsets and fluctuations, and at high counting rates to minimize the effects of pulse overlap or pileup. The term “pulse shaping” implies that the amplifier-filter system is time invariant. In such a system the system parameters do not vary during the measurement and a single measurement of amplitude or time is performed. Such a system is described completely by its impulse response.

In signal filtering, we also use time-variant methods, such as capacitor switching and correlated multiple sampling of the signal. The filtering properties of a time-variant system are described by its weighting function  $w(t)$ . *The weighting function describes the contribution that a noise impulse, occurring at time  $t_i$ , makes at the measurement  $t_m$* , as illustrated in Fig. 10.15. It is essentially a measure of the memory of noise impulses (or any other signals) occurring before the observation time  $t_m$ . As shown, *the weighting function for time-invariant systems is simply a mirror image in time of the impulse response, with its origin displaced to  $t_m$* . For a time-variant system, the impulse response (output waveform) is generally quite different from its weighting function. In some cases time-invariant and time-variant processing could be devised to produce the same result, i.e., both methods will be described by the same weighting function, while their implementation will be quite different. The noise-filtering properties of any weighting function for detector signal processing can be most easily determined by the time domain analysis technique shown in Fig. 10.13. The time domain analysis method based on Campbell’s theorem was first introduced by Wilson [22], and subsequently elaborated in Refs. [9, 23, 24].


A composite weighting function for multiple correlated sampling is obtained by superposition of weighting functions for individual samples. This is illustrated



**Fig. 10.15** An illustration of the weighting function  $w(t)$  corresponding to impulse response  $h(t)$ . A unit noise impulse at  $t_i$  contributes  $w(t_i)$  at the time  $t_m$  of the peak of response to the signal impulse at  $t_0$ . The weighting function in this case of a simple time-invariant filter is a mirror image of the impulse response delayed by the sample (i. e., measurement) time  $t_m$



**Fig. 10.16** Composite weighting function for *correlated double sampling* (CDS). Time-invariant filtering (“pulse shaping”) described by an impulse response  $h(t)$  is assumed prior to sampling



for *correlated double sampling* (CDS), a technique commonly used for readout of CCD's and large pixel arrays, Fig. 10.16. Single sample processing is described by a symmetrical triangular impulse response approximating single RC differentiation and one or two RC integrations. The single sample weighting function with respect to the sampling time at  $t_m$  is shown (dashed), and it is a mirror image of the impulse response. It is assumed that a (delta function) signal of interest will arrive at time  $t_{0+}$ , and produce a response described by the impulse response. In double correlated sampling another sample is taken at  $t_0$ , just before the arrival of the signal. This sample, sometimes called “baseline sample”, is subtracted from the “signal or measurement sample”. The weighting function for the baseline sample is shown inverted and earlier in time by  $t_m - t_0$ .

The composite weighting function (thicker solid line) is bipolar and it has area balance. This is another way of saying that CDS has zero dc response and that it attenuates (but does not eliminate) baseline fluctuations at low frequencies. The ENC can be easily calculated from such a composite weighting function using the technique for time domain noise analysis shown in Fig. 10.13. A noise analysis of such a case in the frequency domain and without the use of a composite weighting function is considerably more time consuming.

#### 10.4.4 Simple ENC Calculation for Series White Noise

Following on the discussion in Sect. 10.4.2 and referring to Fig. 10.17, a simple relation for the equivalent noise charge ( $ENC_s$ ) due to *series white noise* follows,


$$ENC_s^2 = (1/2) e_n^2 C_{in}^2 I_1, \quad (10.10)$$

where

$$I_1 = \int_{-\infty}^{\infty} [w'(t)]^2 dt = 2/t_m = A_1/t_m, \quad (10.11)$$

$$ENC_s = e_n C_{in} / t_m^{1/2}. \quad (10.12)$$

**Fig. 10.17** Simplified equivalent circuit for calculation of equivalent noise charge (ENC) due to amplifier series white noise



It requires the knowledge of three parameters: noise spectral density  $e_n$ , total input capacitance (detector + amplifier)  $C_{in}$ , and peaking time  $t_m$  of the triangle approximating the weighting function. Such an approximation is useful for noise estimation, since the series noise coefficient for a fifth order semi-Gaussian weighting function with equal peaking time,  $A_1 = 2.2$ , differs by only  $\sim 10\%$  from  $A_1 = 2$  for the triangular function. In a preamplifier design, the expected  $e_n$  can be determined from the operating conditions (current and transconductance) of the first transistor, or from a more complete equivalent circuit of the input transistor shown in Sect. 10.5.2. A primary objective of low noise amplifier design is to make the noise contributions of all other circuit components negligible compared to the input transistor. Eq. (10.12) describes simply also the noise charge slope with respect to detector capacitance (in pF),

$$\partial (ENC_s/q_e) / \partial C = e_n / t_m^{1/2} \quad [\text{rms electrons/pf}] \quad (10.13)$$


#### 10.4.5 Simple ENC Calculation for Parallel White Noise

From Sect. 10.4.2 and Figs. 10.13 and 10.18 simple relations follow for  $ENC_p$  due to parallel shot noise and resistor (thermal) noise. The gated integrator case, where the weighting function equals unity for the duration of the gate, illustrates that the  $ENC_p$  for shot noise is simply the square root of the variance of a Poisson sequence of impulses counted for a time  $t_G$ .

$$ENC_p = (q_e I_0 t_G)^{1/2} = (q_e^2 n t_G)^{1/2} = q_e (n t_G)^{1/2} = q_e n_G^{1/2}. \quad (10.14)$$

By Campbell's theorem the contribution of each impulse to the variance is determined by the weighting function, and for a given weighting function the parallel noise integral  $I_3$  has to be determined. For an approximation by a triangle with a peaking time  $t_m$ ,  $I_3 = (2/3)t_m$ . The parallel noise contribution for the triangular weighting function is the same as for gated integration one third as wide.

The contribution by the parallel resistor thermal noise can be compared simply to the shot noise by the "50 mV rule": a dc current  $I_0$  causing a voltage difference of



**Fig. 10.18** Simplified equivalent circuit for calculation of equivalent noise charge (ENC) due to detector and amplifier parallel white noise (bias or feedback resistance, detector leakage current, tunneling gate current in MOS, base current bipolar junction transistor)

$\sim 50$  mV on a resistor  $R_p$  contributes a shot noise equal to the thermal noise of that resistor at room temperature; from  $4kT/R_p = 2q_e I_0$ :  $R_p I_0 = 2kT/q_e$ .


#### 10.4.6 Calculation and Estimation of ENC for $1/f$ Noise

$1/f$  noise becomes a limiting factor in many physical measurements. We can imagine reducing the series white noise in charge measurements to a very low level by continuing to increase the measurement (integration) time  $\tau$ , provided the parallel (leakage or dark current) noise is very low. We would eventually reach the “noise floor” due to the  $1/f$  noise. Once the  $1/f$  noise spectral density is determined experimentally and defined by the parameter  $A_f$  in  $[V^2]$  as in Eq. (10.5), ENC can be calculated by the integral  $I_2$ , Eq. (10.7). In the time domain this is an integral of a fractional-order (half-order) derivative squared of the weighting function (a mathematical operation which cannot be called “trivial” before one learns how to do it, and it can be considered “tedious” at best). In the frequency domain the calculation is somewhat easier for time-invariant systems, but for time-variant systems defining the transfer function  $H(j\omega)$  is more difficult and less intuitive than determining the weighting function.

We illustrate this here on the example of a commonly used weighting function of trapezoidal form as shown in Fig. 10.19. There are many different hardware implementations of this function in different applications. Time-invariant versions have used delay line clipping and higher order RC prefilters. Gated integrator and higher order prefilters have been used in several applications, starting with germanium gamma-ray detectors [25]. This function is widely used with CCDs in astronomy, implemented by correlated double sampling and dual-ramp integration.

We define the trapezoidal weighting function by the width of the ramp  $\tau_p$  and the flat top as a fraction of the ramp,  $\Delta\tau_p$ . The equivalent noise charge for  $1/f$  noise is then,

$$ENC_f^2 = \pi C_{in}^2 A_f A_2, \quad (10.15)$$



**Fig. 10.19**  $1/f$  noise coefficient  $A_2$ , Eq. (10.7), for trapezoidal weighting function. *Solid line*: analytical solution, Eq. (10.18); *dashed*: approximate solution, Eq. (10.22)

where

$$A_2 = \int_{-\infty}^{\infty} \left[ w^{(1/2)}(t) \right]^2 dt. \quad (10.16)$$

The half order derivative of the weighting function,  $w^{(1/2)}(t)$ , is obtained by convolution of  $w'(t)$  with the elementary pulse  $U(t)/t^{1/2}$  for  $1/f$  noise shown in Fig. 10.10d,

$$w^{(1/2)}(t) = \left( 1/\sqrt{\pi t} \right) * w'(t), \text{ for } t \geq 0+ \quad (10.17)$$

Using *Mathematica*, and after some manipulation, the result for  $A_2$  is, [26],

$$A_2 = \frac{1}{\pi} \left[ \Delta^2 \ln \Delta + (2 + \Delta)^2 \ln (2 + \Delta) - 2(1 + \Delta)^2 \ln (1 + \Delta) \right]. \quad (10.18)$$

The coefficient  $A_2$  vs the flat top  $\Delta$  of the trapezoidal weighting function is plotted in Fig. 10.19.

The effect of the series  $1/f$  noise is lowest for a triangular weighting function,  $\Delta = 0$ , which is also the case for the series white noise, Eqs. (10.10, 10.11, and 10.12). As the flat top is made longer,  $A_2$  increases, since such a trapezoidal function has a higher ratio of its cutoff frequencies, which results in integrating a wider band of the  $1/f$  noise spectrum, Eq. (10.4).

An almost identical result for  $A_2$  has been obtained by a calculation in the frequency domain and described in a study of CCD noise performance [27].

Since such exact calculations of  $ENC_f$  for any weighting function can be time consuming, we emphasize here a simple estimation method, which provides results sufficiently close to the exact calculations for most purposes. It has been pointed out by Gatti et al. [28] that the three integrals in Eqs. (10.6, 10.7, and 10.8) have to satisfy the Cauchy-Schwartz inequality,

$$\int_{-\infty}^{\infty} [w^{(1/2)}(t)]^2 dt \leq \left\{ \int_{-\infty}^{\infty} [w'(t)]^2 dt \bullet \int_{-\infty}^{\infty} [w(t)]^2 dt \right\}, \quad (10.19)$$

that is,

$$A_2 \leq (A_1 A_3)^{1/2}. \quad (10.20)$$

Thus there is an upper limit to  $A_2$  in relation to  $A_1$  and  $A_3$  which are easily calculated from Fig. 10.13, or Eqs. (10.6, 10.7, and 10.8). A study of the most commonly used weighting functions, [28], reveals that  $A_2/(A_1 A_3)^{1/2}$  falls between 0.64 and 0.87, a spread of less than  $\pm 8\%$  in the calculation of rms noise, so that for the estimation of  $1/f$  noise the following approximate relation can be used,

$$A_2 \approx 0.75(A_1 A_3)^{1/2}. \quad (10.21)$$

For the trapezoidal weighting function in Fig. 10.19,  $A_1 = 2$  and  $A_3 = \Delta + 2/3$ , and the approximation for this case is,

$$A_2 \approx 0.75 \left[ 2 \left( \Delta + \frac{2}{3} \right) \right]^{1/2}. \quad (10.22)$$

Figure 10.19 shows that this approximation is within a few percent of the exact analytical solution, Eq. (10.18).

In any noise analysis of charge amplifiers one will have already calculated, or otherwise determined the values of  $A_1$  and  $A_3$ , so that the information about the filtering (pulse shaping) effect on the series and parallel white noise will readily also provide an estimate of the  $1/f$  noise,

$$ENC_f^2 = \pi C_{in}^2 A_f A_2 \approx \pi C_{in}^2 A_f \left( 0.75 \sqrt{A_1 A_3} \right). \quad (10.23)$$

It is interesting to note that for a Gaussian weighting function  $A_2 = 1.00$ , for a triangular weighting function 0.88, for a fourth order semi-Gaussian 1.02, for CR-RC 1.18.

$A_f$  is a parameter resulting from a measured spectral density and it does not contain any specific information about the properties of the amplifying device unless other parameters are known.

For input transistor optimization a parameter which is to the first order independent of the device dimensions is more useful [29],  $K_f = A_f C_{gs}$  [J]. This constant

ranges from  $10^{-27}$  J for junction field-effect transistors (JFETs) to  $\sim 10^{-25}$  J for p-channel and  $\sim 10^{-24}$  J for n-channel MOS transistors.

For an accurate calculation of the noise charge for noise spectra departing from the three-term power-law representation (“white series voltage noise”, “ $1/f$  series voltage noise”, and “white parallel current noise”), circuit simulation and numerical calculation are the tools of choice to obtain accurate results. The discussion here was intended to provide some insight: *ENC<sub>f</sub> depends only on the shape of the weighting function but not on the time scale.*

### 10.4.7 Noise in Dielectrics

Thermal fluctuations in dielectrics generate a noise which is quantitatively related to the parameters describing dielectric losses. This type of noise and its importance for detectors was first studied in [30] and then summarized in [31]. For a dielectric with low losses, the dissipation factor or the loss factor  $D$  (equal to the imaginary part  $\varepsilon''$  of the permittivity  $\varepsilon = \varepsilon' + j\varepsilon''$ ) is independent of frequency in the range of interest for particle and photon detectors ( $\sim 10^4$  to  $10^8$  Hz). It can be defined as  $D = G(\omega)/(\omega C_{diel})$ , where  $G(\omega)$  and  $C_{diel}$  are the loss conductance and the capacitance of the dielectric as measured on an impedance bridge at an angular frequency  $\omega$ . According to the fluctuation-dissipation theorem [32, 33], and using the Johnson-Nyquist formula for thermal noise, a dissipative dielectric generates a noise current with a spectral density,

$$i_n^2 = 4kTG(\omega) = 4kTD\omega C_{diel}. \quad (10.24)$$

The equivalent noise charge  $ENC_{diel}$  due to dielectric noise can be calculated using Eq. (10.7),

$$ENC_{diel}^2 = 2kTDC_{diel}A_2. \quad (10.25)$$

Following on the discussion in the previous section we assume here  $A_2 = 1.2$ ,

$$ENC_{diel}^2 = 2.4kTDC_{diel}. \quad (10.26)$$

We note that the spectral density (Eq. 10.24) upon integration on the input capacitance becomes  $1/f$ , and therefore the equivalent noise charge due to dielectric noise is independent of the width (the time scale) of the weighting function.

The noise from lossy dielectrics may pose in some detectors a lower limit to total noise. If a lossy dielectric contributes 1 pF, such as a glass fibre board with  $D \approx 2 \cdot 10^{-2}$ , this alone would present a lower limit of  $ENC_{diel} \approx 86$  rms electrons. Best dielectrics (e.g. Teflon, polystyrene, quartz) have  $D \approx 5 \cdot 10^{-5}$ , which results in  $\sim 5$  rms electrons. This noise contribution to the charge measurement can be reduced only by reducing  $C_{diel}$  (and/or the temperature).

## 10.5 Gain Mechanisms and Noise in Transistors

### 10.5.1 Gain Mechanism, Electron Transit Time, Unity-Gain Frequency

The charge control concept as the basis for the gain mechanism in all three-terminal amplifying devices (transistors) was discussed in Ref. [9]. The “control charge”  $Q_c$  is illustrated in Fig. 10.22.

The relation among the control capacitance  $C_{gs}$ , the transconductance  $g_m$ , the electron transit time  $\tau_e$  and the unity gain frequency  $f_T$  is summarized by,

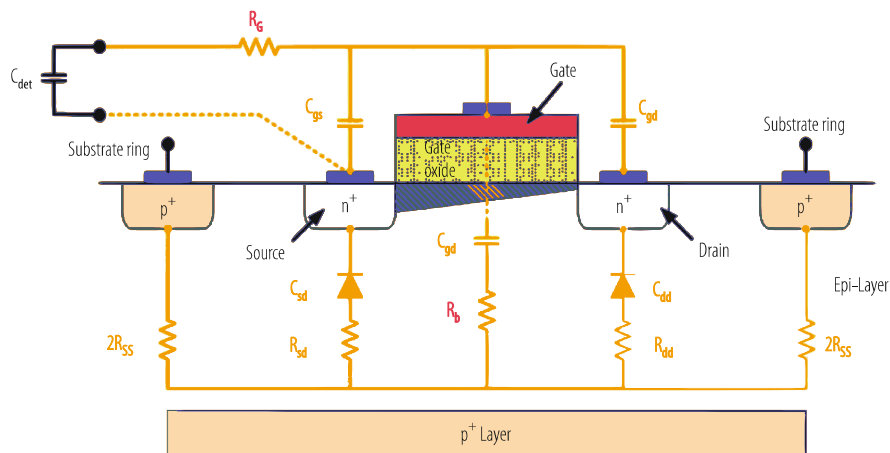
$$\begin{aligned} \Delta Q_c / \Delta I_d = \tau_e = C_{gs} \Delta V_{gs} / \Delta I_d = C_{gs} / g_m \rightarrow C_{gs} = g_m \tau_e, \\ f_T \approx 1 / (2\pi \tau_e) = \left( 1 / 2\pi \right) \left( g_m / C_{gs} \right). \end{aligned} \quad (10.27)$$

NMOS transistors in submicron range (channel length below ~0.25 microns) will have a unity gain frequency in the range 10 to 100 GHz when operated in strong inversion. These same devices will be operated in weak inversion to maximize the transconductance/current ratio and the power dissipation in detectors with large numbers of channels (pixels or strips). This means reduced transconductance with almost the same gate capacitance resulting in a unity gain frequency in the range of 1 GHz or less. This affects the speed of response and the stability considerations in the design of feedback amplifiers.

Equation (10.27) describe only a simplified basic relation among intrinsic device parameters. A very extensive treatment of charge control concepts for CMOS transistors including parasitic parameters and device operating conditions is given in Ref. [7].

### 10.5.2 Noise Sources in MOS Transistor

A brief overview of white noise sources in an NMOS transistor normalized to the intrinsic channel series noise resistance  $\gamma/g_{ms}$  is illustrated in Fig. 10.20 and Eq. (10.28) with  $\gamma$  typically in the range 0.5 to 1.0. The second term in Eq. (10.28) is the gate induced noise contribution [34]. The coefficient  $\delta/5\gamma$  depends on the bias conditions. For estimation purposes  $(\delta/5\gamma) \sim 1/3$ . With capacitive sources, such as most radiation detectors, this term is usually negligible. In particular, at operating conditions to minimize the power in the input transistor, the optimum ratio  $C_{gs}/C_{in}$  is small. The contributions by the gate resistance and substrate resistance can be made



**Fig. 10.20** An illustration of parasitic resistive noise sources in an NMOS transistor in addition to the channel noise  $\gamma/g_{ms}$

small by the device design. The equivalent series noise resistance of the NMOS transistor can be summarized referring to the notation in Fig. 10.20 as,


$$\frac{R_{eq}}{\gamma/g_{ms}} = 1 + \frac{\delta}{5\gamma} \left( \frac{C_{gs}}{C_{in}} \right)^2 + \frac{1}{\gamma} (R_g g_{ms}) + (R_b g_{mb}) \frac{g_{mb}}{g_{ms}} + \frac{1}{\gamma} (R_b g_{ms}) \left( \frac{C_{gb}}{C_{in}} \right)^2. \quad (10.28)$$

### 10.5.3 Charge Transfer from Detector to Transistor: Capacitance Matching

In all cases where the amplifier is connected directly to the detector via a resistive conductor the charge produced by ionization is distributed among the detector capacitance, amplifier capacitance and any stray capacitance according to the ratio of capacitances, Fig. 10.21. Due to this, only a fraction of the charge of interest (the signal) arrives where it matters—that is to the conduction channel of the input transistor where it controls the drain (collector) current, as illustrated in Fig. 10.22. (An exception to this is if the two capacitors are connected by an inductor in which case the charge is transferred periodically between the two capacitors.) In case of a CCD the ionization charge is moved peristaltically in a potential well formed and driven by appropriate clock voltages applied to the gate electrodes. The charge shifted a few hundred (or thousand) times arrives at the collection electrode (“floating diode”) which is connected to a source follower. In the CCD the charge arriving at the collection electrode is the original charge packet produced




**Fig. 10.21** Sharing of the induced (“collected”) charge between the detector and amplifier capacitance



**Fig. 10.22** An illustration of the control charge with respect to the control capacitance  $C_{gs}$ :

$$\Delta Q_c = Q_s / [1 + (C_{det} + C_{gsp} + C_{ds}) / C_{gs}]$$



by ionization except for a few electrons lost to trapping. The charge transport in a conductor is by a small displacement of a large number of free electrons. The CCD principle allows multiple measurements on the same charge packet as described in Ref. [35].

Optimization of the signal to noise ratio requires appropriate matching of the transistor active capacitance (which controls the current) to all other capacitances connected to the input—a subject addressed in some detail in Ref. [10].

The charge control concept expressed by Eq. (10.27) is at the basis of the gain mechanism in almost all electronic amplifying devices: it takes an increment of charge,  $\Delta Q_c$  in Fig. 10.22, to cause a steady state change of the current in the conducting channel. From Eqs. (10.12 and 10.27) we can determine the lower limit for the charge sensitivity due to the series noise in terms of the electron transit time in the conducting channel. Scaling the device width (with no power limitation) to achieve the best signal-to-noise ratio requires that the control electrode capacitance equal (match) the detector capacitance,  $C_{gs} = C_{det}$ . With this and Eq. (10.27), the lowest noise for an electronic amplifier that could be achieved under ideal circumstances is given by,

$$ENC_{s\ opt} \cong 2\sqrt{2}\sqrt{kTC_{gs}}\sqrt{\tau_e/t_m}, \quad (10.29)$$

where  $t_m$  is the weighting function zero-to-peak time, as in Fig. 10.13, also referred to as the “integration time” or the “measurement time”.

It is assumed here for simplicity that the equivalent series noise resistance is equal to the inverse of the device transconductance, i.e.,  $\gamma \sim 1$ . While Eq. (10.29) is useful for estimation purposes and for establishing a lower limit for the amplifier series noise, the electron transit time is rarely used directly for noise calculations.

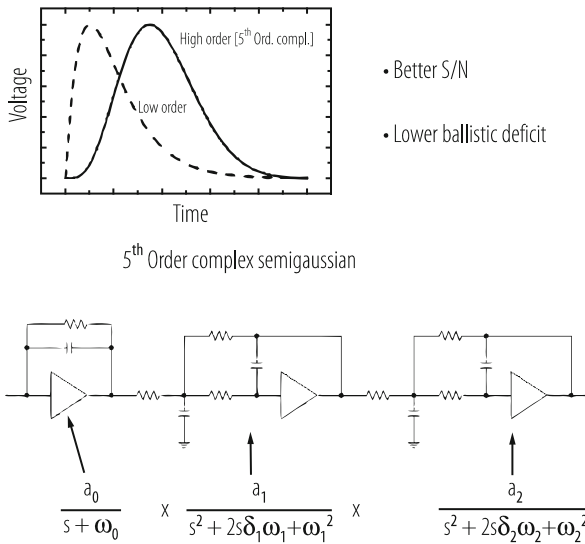
Accurate noise calculations usually rely upon the measured or known noise spectral density for a particular device and have to include parasitic capacitances and resistances, as indicated in Fig. 10.20.

## 10.6 Weighting Function Realizations: Time Invariant and Time Variant

### 10.6.1 Time-Invariant Signal Processing

A low order asymmetrical function (dashed), shown in Fig. 10.23, results in higher series noise (due to a steeper rise, see Fig. 10.13). A nearly symmetrical function requires a higher order signal processing chain, as shown.

A time-invariant circuit realization of a nearly symmetrical weighting function is described in Ref. [36].



**Fig. 10.23** A typical time-invariant signal processing chain resulting in a fifth order pseudo-Gaussian impulse response (weighting function), where  $s$  is the Laplace transform variable. This function results in lower noise coefficients  $A_1$ ,  $A_2$ ,  $A_3$ , Eqs. (10.6, 10.7, and 10.8), 2.2, 1.05, and 0.78, respectively, compared to a low-order CR-RC function of equal width at the base. See Ref. [10] for design considerations and other implementations

## 10.6.2 *Uncorrelated Sampling and Digital Filtering*

Sampled data digital signal processing has become prevalent in detector systems for gamma-ray and x-ray spectroscopy, for time projection chambers and in various forms in particle physics. One of the advantages is that it provides flexibility in the realization of mathematically optimal weighting functions. Optimal signal processing cannot be achieved without some analogue functions. An anti-aliasing filter is an essential part of the system. Its function is to limit the bandwidth prior to sampling so as to satisfy the Nyquist-Shannon sampling criterion: the bandwidth at the output of this filter must be no more than one half of the sampling frequency (the “Nyquist limit”). If this is not satisfied, the noise at frequencies higher than the Nyquist limit is shifted in frequency, i.e., (“aliased”) by undersampling, to the frequencies below this limit. *The resulting loss in S/N due to aliasing cannot be recovered by any subsequent processing.* The role of digital filtering is to create optimized weighting functions in spectroscopic systems, and to enable an optimal particle track measurement in tracking systems. In spite of the power of digital processing, it is most efficient to cancel any long tails in the detector-preamplifier response by analogue means. If the tail cancellation is performed digitally, much larger numbers of samples have to be processed (deconvolved) for each event. For asynchronous (uncorrelated) sampling in semiconductor detectors for gamma-ray and x-ray spectroscopy see Refs. [37, 38]. In such systems optimum weighting functions are of trapezoidal form as in Fig. 10.19, where the flat top allows uniform weighting for a variable charge collection time.

## 10.6.3 *Correlated Sampling*

*Correlated double sampling* (CDS) is being used with many detectors in various implementations and under different names. One of these is known as “baseline subtraction”—taking a sample prior to the arrival of the usually unipolar signal (a delayed signal or with the arrival time known). This case is illustrated in Fig. 10.16, and it results in a bipolar weighting function, which defines quantitatively the effect of CDS on the noise, as discussed in Sects. 10.4.2 and 10.4.3. Correlated double sampling is an essential part of CCD signal processing in astronomy.

Signal processing by *multiple correlated (synchronous) sampling* has been used for noble liquid calorimeters, such as the liquid argon electromagnetic calorimeter in the ATLAS experiment, Refs. [39, 40].


## 10.7 Equipartition and $kTC$ Noise

Integration of the power spectrum (spectral density) arising on a capacitance from the thermal noise current,  $i_n^2 = 4kT/R$ , of the resistor results in the *total fluctuation* of charge (and voltage), which is *independent of the value of the resistance  $R$* . The bandwidth (equivalent to an abrupt cutoff) of the RC circuit in Fig. 10.24 is  $1/(4RC)$ , and the total fluctuation in voltage and charge is

$$\begin{aligned}\sigma_v^2 &= kT/C, \\ \sigma_q^2 &= kTC.\end{aligned}\tag{10.30}$$

The resistance (with the capacitance  $C$ ) determines the bandwidth of the noise but not its magnitude. The  $kTC$  noise at 300 K is quite high even on small capacitances, as shown in Table 10.1. Most of this noise does not affect the measurement in systems where filtering and a very high parallel (feedback) resistance is used following a preamplifier, in other words where the noise corner time constant, Eq. (10.9), is much longer than the width of the filter weighting function. Such a system responds only to the portion of the spectrum where the spectral density is very low. *An example:* high resolution x-ray spectrometry with silicon detectors, where time-invariant or digital filtering is usually used. In contrast, when the measurement is performed by taking a sample directly on the detector capacitance and the filtering is not possible, the full  $kTC$  noise is included in the measurement, and it can be reduced only by correlated double sampling (CDS)—if applicable, as discussed in Fig. 10.27. CDS is just another way of excluding the noise at low frequencies from the measurement.

From the above discussion, which is based on circuit analysis, one is led to conclude that the  $kTC$  noise arises from the resistance, and yet its magnitude is independent of the value of the resistance. One may also be led to conclude that an “ideal” capacitance would have no noise. This is contradicted by the equipartition theorem which makes no direct assumption about the resistance. The equipartition



**Fig. 10.24** A simple circuit for calculation of  $kTC$  noise

**Table 10.1** Charge and voltage total fluctuation vs capacitance ( $T = 300$  K)

Capacitance	Charge fluctuation	Voltage fluctuation
$C$ [F]	$(kTC)^{1/2}/q_e$ [rms e]	$(kT/C)^{1/2}$ [ $\mu$ V]
1a	0.4	$6.4 \cdot 10^4$
10a	1.26	$2.0 \cdot 10^4$
100a	4	$6.4 \cdot 10^3$
1f	$1.26 \cdot 10$	$2.0 \cdot 10^3$
10f	$4.0 \cdot 10$	$6.4 \cdot 10^2$
100f	$1.26 \cdot 10^2$	$2.0 \cdot 10^2$
1p	$4.0 \cdot 10^2$	64
10p	$1.26 \cdot 10^3$	20
100p	$4 \cdot 10^3$	6.4

theorem states, that for a system in thermal equilibrium, the fluctuation energy per degree of freedom is  $kT/2$ . “*Per degree of freedom*” applies to any variable by which the energy of an energy storage object, or energy storage mode, can be defined. Thus for a capacitance,  $C\langle v^2 \rangle = kT/2$ , from which Eq. (10.30) follows. So the statistical mechanics gives the same result for the total fluctuation but without any details about the dissipative components and the noise spectrum. A practical consequence is that as a capacitor becomes closer to an ideal one, the noise spectrum shifts toward zero frequency, while the total fluctuation remains constant.

We add here parenthetically that in a resonant system (an inductance-capacitance circuit), where there are two degrees of freedom, the noise spectrum is concentrated around the resonant frequency, while the integral of the power spectral density (total charge fluctuation) on the capacitance equals  $kTC$ , and the total current fluctuation (variance) in the inductance equals  $kT/L$ .

The above considerations apply also to analogous mechanical systems.


*Transient behaviour* of  $kTC$  noise is of great interest for switched capacitance circuits and for pixel detectors, where pixels are read out directly without filtering, by being sampled in a matrix arrangement, either before or after simple amplifiers (source followers, or three transistor circuits as in, Figs. 10.32, 10.33, and 10.34).

Transient behaviour of noise on a capacitance after switching the resistance or capacitance can best be studied by applying Campbell’s theorem, as shown in Figs. 10.25 and 10.26. In this case the integration of the variance has to be carried out from the time when the switching takes place (zero) to the time  $t$  when the observation is made,


$$\sigma_v^2 = \left(\frac{1}{2}\right) 4kT \frac{1}{R} \int_0^t h^2(u) du = \frac{kT}{C} \left(1 - e^{-2t/\tau}\right). \quad (10.31)$$

The oscillogram shows build-up of noise after switching a white noise source onto an RC circuit,  $h(t) = (1/C)e^{-t/\tau}$ . The time constant  $\tau$  equals the product of the capacitance and the resistance after the switching. Such a build up occurs after

**Fig. 10.25** Model for calculation of the build-up of  $kTC$  noise on an RC circuit



**Fig. 10.26** Build-up of  $kTC$  noise on an RC circuit




a reset switch across a capacitor is opened and a much higher value of resistance appears in parallel with the capacitance. This is illustrated in Fig. 10.27.

Figure 10.27 illustrates what happens with  $kTC$  noise in active pixel sensors and CCDs. While the reset switch is closed, the  $kTC$  noise extends to very high frequencies corresponding to the very short time constant  $r_{ON}C$ . When the switch is “opened” the time constant increases by many orders of magnitude. A value of the “old”  $kTC$  noise is stored on the capacitance, and it decays very slowly with this very long time constant,  $CR_{OFF}$ . It is this sample of the “old” wide bandwidth noise that is often referred to as the “reset noise”, even though its origin is not in the reset action. During the same time after switching, the “new”  $kTC$  noise builds up also very slowly, but faster than the stored value decays, since the rms noise build-up proceeds as  $(1 - \exp[-2 t/CR_{OFF}])^{1/2}$ , Eq. (10.31) and Fig. 10.27. From this illustration one can see the conditions under which correlated double sampling may reduce significantly the  $kTC$  noise: *Sample 1 may be taken any time between opening of the reset switch and the arrival of the signal. Sample 2 may be taken any time after the arrival of the signal but before the “new”  $kTC$  noise has built up.*

Analysis of the effects of the  $kTC$  noise in some cases is not straightforward. Here are some general guidelines:

- *Fluctuation-dissipation theorem* with Johnson-Nyquist expression for thermal noise is essential for calculation of noise spectra and for detailed information on noise sources based on circuit analysis.



**Fig. 10.27** Transient behaviour of  $kTC$  noise caused by the reset action of the sense node in CCDs and in pixel detectors with matrix readout (e.g., hybrid CMOS detectors)

- *Equipartition theorem* provides no detailed information on the noise spectra, but provides a *check on the integrals* of noise spectra (the total fluctuation).
- *Transient behaviour* of noise in switched capacitor circuits and matrix readout pixel arrays is best understood by means of *Campbell's theorem*, which provides *noise variance vs time*, as shown in Figs. 10.25, 10.26, and 10.27 and by Eq. (10.31). The knowledge of the dissipative component (resistance) is necessary for the transient analysis.
- A charge reset and transfer by a switch result in  $kTC$  independently of the switch ON resistance. This noise can be subtracted only if the *first sample* in the CDS is taken *before the signal*.
- Transfer (i.e., direct transport) of charge without switching (as in a CCD) does not result in  $kTC$  noise. *Reset* of the sense node *does*.

A frequently asked question: Can the total charge fluctuation (variance) on a capacitance be reduced below  $kTC$ ? Yes, by “electronic cooling”, where the apparent noise temperature of the resistance in parallel with the capacitance is reduced by feedback, Sect. 10.8.1 and Ref. [45]. While this is useful in practice, a note should be made that a system with active elements (gain) cannot be considered as being in thermal equilibrium with the surrounding.

An important distinction between two classes of signal processing schemes should be emphasized: (1) when no filtering (band limiting) takes place before the measurement, the total  $kTC$  fluctuation will contribute fully to the charge


measurement, and usually plays a dominant role. In contrast, (2) with charge amplifiers followed by filtering, the contribution of this noise (then usually referred to as “parallel noise”) can be made negligible in most cases, by avoiding the low frequency part of the spectrum (by using a short peaking time, Fig. 10.14). Correlated double sampling (CDS) is one form of filtering.

## 10.8 Some Basic Signal Processing and Detector Readout Circuits

### 10.8.1 Charge Amplifier Configuration


In the most basic charge amplifier feedback configuration only two transistors are essential to realize a complementary cascode, as shown in Fig. 10.28. The current sources in positive and negative supplies can be realized by resistors or by low noise transistor current sources. There is only one significant pole ( $C_0 R_0$ ) in the first order solution for the response of this circuit. Higher order poles are given by the unity gain frequency of the transistors used. The cascode alone is an “operational transconductance amplifier” (very high output impedance). With the follower amplifier  $\times 1$  it becomes an operational amplifier.

Gain and input impedance relations for the feedback charge amplifier configuration are derived from Figs. 10.28 and 10.29. The frequency dependence of the open loop gain is inherent to a high gain single pole amplifier. It is described by two parameters, unity gain frequency  $\omega_h = g_m/C_0$ , and the gain “roll off” frequency (3 dB point)  $\omega_l = 1/(R_0 C_0)$ . The dc gain is then  $|G_0| = \omega_h/\omega_l = g_m/R_0$ , where  $g_m$  is the transconductance of the input transistor,  $C_0$  is the dominant pole capacitance and  $R_0$  is the dominant pole resistance. Input impedance with capacitive feedback has two terms, a resistance  $R_{in} = 1/\omega_h C_f$  in series with a capacitance  $C_f G_0$ . The resistance term  $R_{in}$ , in conjunction with the total input capacitance, determines the



**Fig. 10.28** Basic folded cascode charge amplifier feedback configuration





**Fig. 10.29** Input equivalent circuit of feedback charge amplifier

rise time of the detector-amplifier. The rise time constant of the output voltage (i.e., the transfer of charge from the detector capacitance to the feedback capacitance) is  $\tau_r = R_{in}C_{in} = (1/\omega_h)(C_{in}/C_f) = (C_0/g_m)(C_{in}/C_f)$ , where  $C_{in} = C_{det} + C_{ampl}$ .

The resistive input impedance has a noise corresponding to the amplifier series noise resistance  $R_{seq}$ , and it appears as a resistance with a noise temperature,  $T_{eff} = TR_{seq}/R_{in}$ . For values of  $R_{in}$  higher than  $R_{seq}$ , the amplifier can be used as a termination for delay lines with a noise lower than that of a termination with a physical resistor  $Z_0$  at temperature  $T$ .


The apparent noise temperature of the resistance  $Z_0$  realized by the capacitance in feedback is  $TR_{seq}/Z_0$ , and this is why it can be called “electronically cooled termination” or “electronically cooled damping”, Ref. [45]. The resistance in parallel with the feedback capacitance adds two more terms to the input impedance of the preamplifier: inductance  $R_f/\omega_h$  in series with a resistance  $R_f/G_0$ . It is important to note the condition to achieve an aperiodic (“damped”) response of the feedback amplifier.

The feedback configuration allows the ultimate in noise performance because the parallel noise sources can be made negligible by using a transistor with a very low gate leakage current and a very high feedback resistance (megaohms to gigaohms). The feedback resistor can be avoided altogether by the use of optoelectronic feedback or a transistor switch to maintain amplifier voltages in the operating range. Signal integration is performed on the feedback capacitance  $C_f$ . The long tail can be cancelled in subsequent pulse shaping by a simple pole-zero cancellation circuit (not shown in the figure). Pulse shaping at the preamplifier by reducing  $R_L$  or  $R_f$  would result in increased noise from the thermal noise of these resistors. The object of the design is to avoid dissipative components at the detector-amplifier input and thus to make  $R_f$  as large as possible.


### 10.8.2 Cascaded Charge Amplifier Chain with Pole-Zero Cancellation

While the basic charge amplifier concept of a simple cascode circuit with a single dominant pole and capacitive feedback has not changed since the days of vacuum tube technology, the charge restoration techniques to control the operating point of the input transistor have evolved, particularly with the advent and widespread application of CMOS monolithic circuit technology. The dc feedback in charge amplifiers via a resistor ( $R_f$  in Fig. 10.28) has always been a problem in applications striving to achieve the ultimate in noise performance. Very high values in the gigaohm range are required in x-ray and gamma-ray spectroscopy, as this resistor injects a noise current inversely proportional to this resistance directly into the input node (i.e., “parallel noise”, Sect. 10.4.5). In most applications the resistance values required are higher than the practical range of polysilicon resistors in CMOS technology. In some applications where the detector capacitance is very low it is best to avoid entirely any resistor and any continuous dc feedback. Various charge restoration techniques using switching or “reset” have been developed. An example is the CCD readout as shown in Fig. 10.35.


Continuous charge restoration is usually simpler to implement and better suited in many applications where its noise contribution is negligible, i.e., with higher detector capacitances and shorter weighting functions (peaking times). Most present gas, noble liquid and silicon particle detectors fall into that category. Additional considerations in the choice between switched and continuous charge restoration are the knowledge of the event arrival time and whether switching transients pose a problem. Continuous feedback in MOS technology is realized by a transistor with a long and narrow channel ( $L \gg W$ ). The resistance of such a device will depend on the signal amplitude and on the detector leakage resulting in a nonlinear response. An elegant solution for accurate nonlinearity compensation is shown in Fig. 10.30. The transistor-capacitor network between the two amplifiers is an exact replica of the feedback network but increased in width by a factor  $N$ . Both networks operate at equal voltages on the two transistors and this ensures the compensation of



**Fig. 10.30** Charge amplifier with continuous reset, pole-zero cancellation and compensation of non-linearity in the feedback transistor [10, 46]



**Fig. 10.31** An alternative configuration for a charge amplifier with pole-zero and transistor nonlinearity compensation [10]




**Fig. 10.32** Matrix readout of integrating pixel detectors. Transistors switches are integrated on the detector substrate

nonlinearity and pole-zero cancellation. The charge (current) gain in the first stage including the pole-zero network equals  $N$ .

An alternative configuration is shown in Fig. 10.31. An advantage of this configuration is that it separates the bias point of the transistors  $M_f$  and  $M_o$  from the virtual ground of the amplifiers. This results in a larger dynamic range. The configuration in Fig. 10.31 can be used as an input stage (charge amplifier) or as a second stage where several gain stages are needed. An analysis of both configurations is given in Ref. [10, 70, 71]. An overview of dc charge restoration circuits is given in Ref. [47].

### 10.8.3 Pixel Matrix Readout


Large pixel arrays can be conveniently read out by a matrix arrangement as illustrated in Fig. 10.32. The charge due to a photon or charged particle is stored on a pixel capacitance. The switches (one/pixel) connect a row of pixels to charge amplifiers located at the bottom of the columns. In this way, a multiplexing density is achieved between that for a separate readout for each pixel and the CCD with only



**Fig. 10.33** Basic circuit diagram of the matrix readout with a switch (and no amplification) in each pixel [41–43]

one readout for the entire array. Such an array can be used in single event counting mode at sufficiently low event rates, or in charge integrating mode at very high event rates. This type of readout allows the use of the same technology, or different technologies, for the detector and the switching transistors. For many applications this approach is the best compromise between interconnect complexity and the speed of readout. An example of such an imaging detector for x-ray radiography is described in Ref. [41], and a silicon detector with Junction FET switches in Ref. [42]. The equivalent circuit diagram of such a readout with correlated double sampling is given in Fig. 10.33. It is convenient to group the noise contributions in this case in two classes: those that can be reduced by correlated double sampling and those that cannot. Referring to the notation in Fig. 10.33, the former are the reset ( $kTC$ ) noise of capacitances  $C_f$  and  $C_s$ , and the latter are the reset noise of the charge collecting pixel  $C_d$ , shot noise from pixel dark current and amplifier series white noise. The reset ( $kTC$ ) noise on the pixel capacitance  $C_d$  cannot be reduced by double correlated sampling according to the discussion in Sect. 10.7 and as illustrated in Fig. 10.27, because *both* the signal and the  $kTC$  noise start building up after the charge transfer from the pixel to the charge amplifier. Thus for the simplest matrix readout the minimum noise is limited by the pixel capacitance, Table 10.1, e.g.,  $\sim 400$  electrons rms for  $C_d \sim 1$  pF at 300 K. There will also be a significant contribution by the amplifier noise if the connection capacitance along the column to the amplifier is large due to the size of the pixel array.

A much lower noise can be achieved in pixel arrays with “amplified pixels” as illustrated in Fig. 10.34. Each pixel has three or more transistors to perform the basic functions of reset, amplification and row selection. A CMOS matrix readout die can be bump-bonded (or in the future, directly bonded) to a pixel detector, or serve as a monolithic active pixel sensor (MAPS), Ref. [44]. Performing the reset in each pixel separately from the charge transfer (unlike in the simplest matrix readout without in-pixel amplification) allows almost complete cancellation of the pixel  $kTC$  noise by double correlated sampling, in applications where the time interval between the two samples is not too long (see Fig. 10.27). Noise levels in the range of 10–20 electrons rms have been achieved with small pixels (20–30 fF). The same level of noise has been achieved with larger pixels ( $\sim 1$  pF) and more conventional charge amplifiers and pulse shaping [48]. A lower noise with silicon pixel detectors has been achieved by integration of a field effect transistor on the high resistivity




**Fig. 10.34** Three transistor cell for readout of pixel detectors comprised of a source follower (SF), reset transistor switch and select transistor switch, illustrated for a bump-bonded hybrid detector. In monolithic active pixel CMOS detectors (MAPS), the readout cells are integrated with sensing diodes [44]. An overview is given in Ref. [43]

detector die (DEPFET, Ref. [49]). The lowest noise has been achieved with CCDs, Sect. 10.8.4. Each of these results was obtained after optimization for a particular application and each one involves a different set of parameters.

### 10.8.4 Charge Conserving (CCD) Readout

The charge coupled device (CCD), Fig. 10.35, is a device with the highest charge detection sensitivity among the photon and particle detectors. Low noise in the single electron rms range has been achieved [35, 69], thanks to a very low capacitance of the readout node and the integrated source follower ( $\sim 15\text{--}30$  fF), and the absence of any continuous conduction path to the sensing node. A switched reset is used after reading out the charge. Correlated double sampling is used to make the  $kTC$  noise negligible, as discussed in Sect. 10.7. Nearly complete  $kTC$  noise cancellation is achieved because the first sample is taken before the signal charge is transferred to the readout node (floating gate). Dual slope integration/trapezoidal weighting function is used for optimal filtering of the transistor (source follower) series noise, as discussed in Sect. 10.4.6.



**Fig. 10.35** An illustration of CCD read-out. The source follower (shown only schematically) is integrated on the CCD substrate as a buried channel enhancement mode MOSFET. The signal charge transferred from left to right is sensed at the floating electrode between the output gate and the reset gate.  $C_{sn}$  is the “node capacitance” = floating gate + gate to source capacitance of the transistor

## 10.9 Electronics Technology Outlook

### 10.9.1 Scaling of CMOS

The process of reduction of the size of CMOS transistors has continued for more than 40 years, and it is only very recently that it has started approaching fundamental physical limits. In this process the device is being scaled down in all three dimensions and in voltage, while the doping concentration is being increased. All dimensions are being reduced by a scaling factor  $\alpha$ . This reduction includes all geometrical parameters of the device such as the gate oxide thickness  $t_{ox}$ , channel length  $L$ , channel width  $W$ , and junction depth. The substrate doping concentration is increased by the same scaling factor. The voltages applied were also expected to be reduced by the same scaling factor. These scaling rules were very clearly described in some detail already in 1974 [50], and the resulting device properties as a function of  $\alpha$  are given in Table 10.2.

The scaling by a factor  $\alpha = \sqrt{2}$  every  $\sim 2$  years has followed the Moore’s law [50] remarkably well until recently, with one exception. In the last few steps below the channel length  $L \sim 0.13$  microns it is not possible to reduce the applied voltages by the same scaling factor. Downscaling has made the speed of the MOSFETs higher, the power dissipation per circuit lower, and it has enabled an ever-increasing level of integration (the number of transistors on a single chip). Aside from the enormous progress in all digital devices, it has made possible increasingly complex functions in the readout of detectors by integration of mixed signal (analogue and digital) circuits in Application Specific Integrated Circuits (ASICs).

**Table 10.2** Scaling rules: dependence of device properties on scale factor  $\alpha$

Device property	Scaling rule
Electric field $E$	const.
Conductance (transconductance) $= I/V$	const.
Current density	const.
Power density	const.
Capacitance	$1/\alpha$
Speed $= I/CV = g_m/C = f_T$	$\alpha$
Switching energy $= CV^2$	$1/\alpha^3$
Power/gate $= CV^2f$	$1/\alpha^2$
Circuits density (transistors/unit area)	$\alpha^2$

The principal impact on analogue ASICs has been:

- More, faster transistors
- Lower capacitance-lower noise readout of small detector pixels
- Better radiation resistance
- Prospects for vertical integration with high resistivity silicon detectors

However, the impact on analogue circuits has been very positive only up to a certain point of scaling. Scaling of CMOS into the deep submicron range (below 100 nm down to a few nm) has some undesirable consequences for low noise amplifier design:

- The low supply voltage “headroom” in scaled CMOS processes imposes limits on analogue circuit topologies. The increasing ratio of  $V_{TH}/V_{DD}$  rules out the use of many classical analogue design topologies.
- The *cascode connection*, useful in providing high gain loads and current sources, becomes difficult to realize once  $V_{DD}$  falls below  $\sim 1.2$  V.
- The electrostatic control of the channel by the gate is reduced, resulting in reduced ratio  $I_{don}/I_{doff}$ . In CMOS transmission gates, commonly found in sample/hold and switched capacitor circuits, self-discharge rates increase due to incomplete current cutoff.
- The reduced electrostatic control of the channel results also in a lower ratio of the drain conductance and the transconductance, i.e., the gain of the transistor. To achieve a certain gain more amplification stages may be needed.
- The dynamic range of capacitance based circuits, memories and sample & hold circuits, is reduced (the ratio of stored charge to the kTC noise), both due to the reduced  $V_{dd}$  and due to a lower storage capacitance.
- Gate tunneling current arising with thin oxides contributes shot noise.

While the scaling in the CMOS technology is driven by the entire information technology industry, some applications require higher operating voltages than dictated by the smallest feature sizes. This has been recognized, and the large semiconductor foundries offer options with a thicker gate oxide in their deep submicron platforms. For example, a nominal gate oxide thickness in the 65 nm platform (or “node”) is about  $65/50 \approx 1.2$  nm, but the process includes also on the

same wafer an oxide thickness of  $\sim 5$  nm, that requires a minimum channel length of  $\sim 250$  nm and allows a supply voltage of  $\sim 2.5$  V. This will make possible in a single ASIC both high density digital circuits with the minimum feature size, and higher voltage input/output circuits and precision analogue circuits with an effectively larger feature size.

Analog design below 100 nm becomes gradually more difficult. The complexity of ASIC design rules and the costs of the design tools increase steeply as the feature size is reduced. This is discussed in some detail in Ref. [51].

While the developments in electronics technology have already made possible very large and complex detectors, the power dissipation associated with increasingly complex signal and data processing has been a problem and will remain a principal limitation and challenge.

Along with the quest in the CMOS scaling continuing down to the nodes in the  $\sim 7$  to 28 nm range, an additional path to higher circuit densities and to higher speed of digital circuits (by shortening the interconnections) is three-dimensional integration of several thin layers of CMOS circuits [52, 53]. A significant breakthrough in particle and photon silicon detectors will be integration (by direct bonding) of a thick (50–500  $\mu\text{m}$ ) high resistivity p-i-n detector die to one or more thin ( $\sim 10$   $\mu\text{m}$ ) CMOS layers of readout electronics.

### ***10.9.2 Transistors at Low Temperatures: Applications in Future Detectors***

Most of the electronics for particle detectors has been based on silicon CMOS devices following the trends in electronics industry. Over the last two decades the technology of silicon–germanium heterojunction bipolar transistors (SiGe HBT) has been developed [54]. These devices have some key properties superior to the bipolar junction transistor (BJT), notably a much higher current gain and a much higher unity-gain frequency. The HBT, unlike the BJT can operate at liquid nitrogen temperature, and furthermore, with an increased current gain. Recently emerging cryogenic applications have generated renewed interest in low temperature properties of both CMOS and HBT technologies [54, 55]. As with HBTs, it has been found that various CMOS device properties improve at low temperatures. The mobility and the transconductance increase by a factor of two to three as the temperature decreases from 300 K to 43 K, and the (inverse) subthreshold slope decreases from  $\sim 90$  mV/decade to  $\sim 20$  mV/decade resulting in a higher ratio of the transconductance to the drain current. The transistor (series) thermal noise has been observed to decrease monotonically from 300 K to 40 K [56]. There is one *caveat*: CMOS transistors have to be operated (at any temperature) under conditions where hot electron generation, which is more pronounced at low temperatures, is minimized [57].



In particle physics, recently developed cold monolithic electronics has enabled scaling up of liquid argon time projection chambers to a very large size, required for studies of neutrino oscillations and nucleon decay. Such detectors, which originated with the ICARUS project [58], in the size range of up to few hundred tons, will have to be built in the range, unthinkable so far, of tens of kilotons. The number of sense electrodes (wires) and the readout amplifiers will be in the range of  $\sim 5 \cdot 10^5$  to  $3 \cdot 10^6$ , assuming a reasonable electron drift length in the time projection chamber (TPC) and the sense wire length. Readout by monolithic amplifiers (with multiplexing) placed at wire electrode frames, results in a significantly lower noise (due to a lower capacitance) than with long cables bringing the signals from each sense wire to the outside of the cryostat, and in a much lower number of cables and feedthroughs. This also allows the design of the cryostat to be freed from the signal cable constraints. The “cold electronics” in this case will have a beneficial impact on both the engineering and the physics, i.e., a higher signal-to-noise ratio allowing better background rejection and a higher precision in the track measurement and reconstruction increasing the sensitivity for detection of interesting phenomena [59, 60]. Such a detector, on a smaller scale ( $\sim 75$  tons of liquid argon with  $\sim 8300$  sense wires and electronic channels immersed in liquid argon), has been built and operated for two years. Uniformity and stability of the gain and noise has been demonstrated [61].

### 10.9.3 *Beyond the Moore’s Law*

As the MOSFET technology advances further into the nanoscale domain (gate widths down to 5–10 nm), Moore’s law, as defined for CMOS transistors, is running up against the physical, technical and economical limitations, and eventually against the granularity of matter (silicon lattice constant  $\sim 0.35$  nm). Physical phenomena associated with small dimensions, such as quantum mechanical effects and fluctuations in the decreasing number of dopants, take place. These effects cause leakage currents (incomplete drain current turn-off and gate tunnelling currents) and dispersion in device parameters (threshold voltage). Associated with the necessity for tighter control of all device and fabrication process parameters are increasing costs.

Recent research toward smaller, faster and lower power devices has been concentrated on “beyond CMOS” devices [62–64]. As with the CMOS, new developments are being driven by the needs for high density-high speed digital and computer circuits. Besides the physical variables considered as “computational variables” familiar in CMOS (current, voltage, charge), other variables are being considered (electric dipole, magnetic dipole, orbital state). Among possible device concepts being explored are:

- Tunneling FET
- Graphene nanoribbon FET

- Bilayer pseudo spin FET
- SpinFET
- Spin transfer torque/domain wall
- Spin torque oscillator logic
- All spin logic device
- Spin wave device
- Nanomagnet logic
- III–V tunnel FETs

Each of these devices may have some properties superior to CMOS, but none, so far, satisfies the set of simple criteria that CMOS does. For example, upon analysis, the spintronic devices have longer switching delays and higher switching energies, due to inherent time of magnetization propagation.

Any “Beyond CMOS” device should have many of the same characteristics as CMOS devices:

- Power gain  $>1$
- Ideal signal restoration and fan-out (output of one device can drive two or more devices)
- Feedback prevention (output does not affect input)
- High ON/OFF current ratio  $\sim 10^{5-7}$
- Low static power dissipation
- Compatibility with Si CMOS devices for mixed functions

The consensus about the future technology has been so far: No new device is yet on the horizon with a potential to completely replace CMOS. More likely, new devices may emerge by gradual evolution. New or special functions (e.g., memories, [34]) may become possible in the nanoscale devices by new physics and such devices may be merged into CMOS circuits to enhance overall performance. Impedance matching may be necessary from the quantum resistance values (kohms) down to the 50–100 ohm range. The overall logic operations and communications will still be based on CMOS. Future integrated circuits are likely to still contain a majority of CMOS devices with a few other beyond-CMOS devices performing various specialized functions. An in depth evaluation of the various device concepts under investigation is given in Ref. [65, 72].

Work on carbon nanotubes as active electronic devices and passive devices (interconnects) has been going on for the last three decades. Some interesting devices and phenomena have been described extensively in literature [64]. Carbon nanotube as a channel in an MOS transistor structure has higher mobility, reducing electron transit time. However, accurate placement of carbon tubes, and the need to form higher current channels by placing multiple carbon tubes in parallel, has been a challenge to fabricate uniform devices. Graphene based devices are analyzed in Ref [66].

In the quest for nanoscale devices and higher density of analogue and digital functions other limitations appear. Since all logic circuits (even spintronics circuits)

need electrical contacts at the terminals of gates and channels (source and drain) their size will be limited by the metallization dimensions.

As far as particle and photon detectors are concerned, further progress in digital circuit technology, based on nanoscale nodes, will result in increased functionality of the detector circuits, particularly those integrated with detectors, as long as they are economically available from multi-project foundry services or as commercial components. The development of multi-layer (3D) circuits for detectors has been challenged by limited access to this technology, given the relatively small quantities needed in physics experiments and high costs.

As far as the analogue front-end circuits for particle detectors are concerned, it appears that CMOS will be less useful below the  $\sim 65$  nm node (Sect. 10.9.1 and Ref. [51]). In higher precision circuits, such as analogue to digital converters and switched capacitor memories, the dynamic range is limited by the low power supply voltages (one volt or less) at the upper end and by the  $kTC$  noise at the lower end. Fortunately, the provision of thicker oxide devices on the same nanoscale platforms leaves the choice of the operating voltage, the gate length and width to the designer (as discussed in Sect. 10.9.1).

An interesting domain, outside of the CMOS mainstream, has been presented by *single-electron transistors* (SETs), [67, 68]. Single electron transistors are devices with a capacitance so small that a single electron can generate a measurable voltage (above the thermal voltage). To be observable, the mean energy of the electron on the capacitance must be several times larger than its thermal energy ( $kTC$  noise). This sets an upper limit on the capacitance of the control electrode (gate) of the transistor (and the sensor connected to the transistor) and on the operating temperature. From Sect. 10.7 and Table 10.1, an upper limit for the capacitance at room temperature would be 1 aF ( $10^{-18}$  F). Alternatively, the temperature would have to be 300 mK, to allow a capacitance in the range of  $\sim 1$  fF. SETs, although known for more than three decades, have not been considered suitable for integration due to large variability in their fabrication.

SETs might be suitable as very sensitive electrometers for *equally low capacitance* sensors. That sensitivity is for the most part due to their low capacitance and in part due to improved carrier transport properties. At 300 K they are not matched to even the smallest pixels of present particle and photon detectors, and will be useful only with an entirely new generation of very fine grained detectors operated at low temperatures.

**Acknowledgments** The author is indebted to his colleagues, Gianluigi De Geronimo, Paul O'Connor, Sergio Rescia, Bo Yu, and the late Pavel Rehak for many stimulating discussions and, over time, for contributions of material that has contributed greatly to this article. In particular, most of the induced signal simulations are due to Bo Yu, the circuits shown in Figs. 10.23, 10.30 and 10.31 originated in publications by Gianluigi De Geronimo and Paul O'Connor. Paul O'Connor provided much of the material on CMOS scaling, and Sergio Rescia worked with the author on  $1/f$  and  $kTC$  noise calculations. Special thanks are to Anand Kandasamy for help with editing of the manuscript and the figures.

## References

1. G. Knoll, *Radiation Detection and Measurement*, John Wiley & Sons (2000).
2. C. Grupen and B. Shwartz, *Particle Detectors*, 2nd edition, Cambridge University Press (2009).
3. G. Lutz, *Semiconductor Radiation Detectors*, Springer-Verlag (1999).
4. H. Spieler, *Semiconductor Detector Systems*, Oxford University Press (2005).
5. I. Iniewski (ed.), *Medical Imaging*, John Wiley & Sons (2009).
6. D.M. Binkley, *Tradeoffs and Optimization in Analog CMOS Design*, John Wiley & Sons (2008).
7. C.C. Enz, E.A. Vittoz, *Charge-based MOS Transistor Modeling*, John Wiley & Sons (2006).
8. E. Gatti, P.F. Manfredi, *Processing the Signals from Solid-State Detectors in Elementary Particle Physics*, *Rivista de Nuovo Cimento*, 9(1) (1986).
9. V. Radeka, *Low-Noise Techniques in Detectors*, *Annu. Rev. Nucl. Part. Sci.* 38 (1988) 217.
10. G. De Geronimo, *Low-Noise Electronics for Radiation Sensors*. In: *Medical Imaging*, K. Iniewski (ed.), John Wiley & Sons (2009) p. 127.
11. W. Shockley, *Currents to Conductors Induced by a Moving Charge*, *J. Appl. Phys.* 9 (Oct. 1938) 635.
12. S. Ramo, *Currents Induced by Electron Motion*, *Proc. IRE* 27(1939) 584-585.
13. E. Gatti, G. Padovini, V. Radeka, *Signal Evaluation in Multielectrode Detectors by Means of a Time Dependent Weighting Vector*, *Nucl. Instrum. Meth.* 193 (1982) 651.
14. N.R. Campbell, *Proc. Cambridge Philos. Soc.* 15 (1909) 117.
15. P.N. Luke, *Unipolar Charge Sensing with Coplanar Electrodes – Application to Semiconductor Detectors*, *IEEE Trans. Nucl. Sci.* 42 (1995) 207.
16. V. Radeka, *1/f Noise in Physical Measurements*, *IEEE Trans. Nucl. Sci.* NS-16 (1969) 17.
17. C.M. Compagnoni et al., *Statistical Model for Random Telegraph Noise in Flash Memories*, *IEEE Trans. Electron Devices* 55(1) (2008) 388.
18. K. Kandiah, M.O. Deighton, F.B. Whiting, *A Physical Model for Random Telegraph Signal Currents in Semiconductor Devices*, *J. Appl. Physics* 86(2) (1989) 937.
19. P. van der Wel et al., *Modeling Random Telegraph Noise Under Switched Bias Conditions Using Cylostationary RTS Noise*, *IEEE Trans. Electron Devices* 50(5) (2003) 1378.
20. A. Konczakowska, J. Cicchosz, A. Szewczyk, *A New Method for RTS Noise of Semiconductor Devices Identification*, *IEEE Trans. Instrum. Meas.* 57(6) (2008) 1199.
21. L.K.J. Vandamme, F.N. Hooge, *What Do We Certainly Know About 1/f Noise in MOST?*, *IEEE Trans. Electron Devices* 55(11) (2008) 3070.
22. R. Wilson, *Noise in Ionization Chamber Pulse Amplifiers*, *Philos. Mag., Ser. 7* Vol. xli, Jan. (1950) 66.
23. V. Radeka, *Optimum Signal Processing for Pulse Amplitude Spectrometry in the Presence of High-Rate Effects and Noise*, *IEEE Trans. Nucl. Sci.* NS-15 (1968) 455.
24. M. Konrad, *Detector Pulse Shaping for High Resolution Spectroscopy*, *IEEE Trans. Nucl. Sci.* NS-15 (1968) 268.
25. V. Radeka, *Trapezoidal Filtering of Signals from Large Germanium Detectors at High Rates*, *Nucl. Instrum. Meth.* 99 (1972) 535.
26. S. Rescia, V. Radeka, unpublished notes.
27. G.R. Hopkinson, D.H. Lumb, *Noise Reduction Techniques for CCD Image Sensors*, *J. Phys. E: Sci. Instrum.*, 15 (1982) 1214.
28. E. Gatti et al., *Suboptimal Filtering of 1/f Noise in Detector Charge Measurements*, *Nucl. Instrum. Meth. A* 297 (1990) 467.
29. V. Radeka, *Semiconductor Position Sensitive Detectors*, *Nucl. Instrum. Meth.*, 226 (1984) 209.
30. V. Radeka, *State of the Art of Low Noise Amplifiers for Semiconductor Radiation Detectors*, *Proc. Int'l. Symposium on Nuclear Electronics*, Versailles, 1968, Vol. 1 (1968) 46-1.
31. V. Radeka, *Field Effect Transistors for Charge Amplifiers*, *IEEE Trans. Nucl. Sci.* NS-20 (1973) 182; see also, V. Radeka, *The Field-Effect Transistor – Its Characteristics and Applications*, *IEEE Trans. Nucl. Sci.* NS-11 (1964) 358.

32. H.B. Callen, R.F. Greene, *On a Theorem of Irreversible Thermodynamics*, Phys. Rev. 86 (1952) 701.
33. H.B. Callen, T.A. Welton, *Irreversibility and Generalized Noise*, Phys. Rev. 83 (1951) 34.
34. D.B. Strukov, K.K. Likharev, *CMOL FPGA: A cell-Based, Reconfigurable Architecture for Hybrid Digital Circuits Using Two-Terminal Nanodevices*, Nanotechnology 16 (2005) 888.
35. R.P. Craft et al., *Soft X-ray Spectroscopy with Sub-Electron Readnoise Charge-Coupled Devices*, Nucl. Instrum. Meth., A 361 (1995) 372.
36. F.S. Goulding, D.A. Landis, N.W. Madden, *Design Philosophy for High-Resolution Rate and Throughput Spectroscopy Systems*, IEEE Trans. Nucl. Sci. NS-30 (1983) 301.
37. V.T. Jordanov, G.F. Knoll, *Digital Synthesis of Pulse Shapes in Real Time for High Resolution Radiation Spectroscopy*, Nucl. Instrum. Meth. A 345 (1994) 337. X-ray Instrumentation Associates, Appl. Note 970323-1.
38. O. Benary et al., *Liquid Ionization Calorimetry with Time-Sampled Signals*, Nucl. Instrum. Meth. A 349 (1994) 367.
39. B.T. Turko, R.C. Smith, Conf. Record of the 1991 IEEE Nucl. Sci. Symp., (1991) 711.
40. O. Benary et al., *Precision Timing with Liquid Ionization Calorimeters*, Nucl. Instrum. Meth. A 332 (1993) 78.
41. N. Matsuura et al., *Digital Radiology Using Active Matrix Readout: Amplified Pixel Detector Array for Fluoroscopy*, Med. Phys. 26(5) (1999) 672.
42. W. Chen et al., *Active Pixel Sensors on High Resistivity Silicon and Their Readout*, IEEE Trans. Nucl. Sci. 49(3) (2002) 1006.
43. V. Radeka, *CCD and PIN-CMOS Developments for Large Optical Telescopes*, Proc. SNIC Symposium, Stanford, Ca., April 3–6, 2006.
44. G. Deptuch et al., *Monolithic Active Pixel Sensors with In-pixel Double Sampling Operation and Column-level Discrimination*, IEEE Trans. Nucl. Sci. 51(5) (2004) 2313.
45. V. Radeka, *Signal, Noise, and Resolution in Position-Sensitive Detectors*, IEEE Trans. Nucl. Sci. NS-21(1) (1974) 51.
46. G. De Geronimo, P. O'Connor, *A CMOS Fully Compensated Continuous Reset System* IEEE Trans. Nucl. Sci. 47 (2000) 1458.
47. G. De Geronimo et al., *Front-End Electronics for Imaging Detectors*, Nucl. Instrum. Meth. A 471 (2001) 192.
48. P. O'Connor et al., *Ultra Low Noise CMOS Preamplifier-Shaper For X-ray Spectroscopy*, Nucl. Instrum. Meth. A 409 (1998) 315.
49. J. Kemmer, G. Lutz, *New Semiconductor Detector Concepts*, Nucl. Instrum. Meth. A 253 (1987) 365.
50. R.H. Denard et al., *Design of Ion Implanted MOSFETs with Very Small Physical Dimensions*, J. Solid-State Circuits SC-9 (1974) 256.
51. L.L. Lewyn et al., *Analogue Circuit Design in Nanoscale CMOS Technologies*, Proc. IEEE 97(10) (2009) 11687.
52. IBM J. Res. & Dev., Issue on *3D Chip technology*, 52(6) (2008).
53. Proc. IEEE, Issue on *3-D Integration*, 97(March) (2009).
54. J.D. Cressler, *On the Potential of SiGe HBTs for Extreme Environment Electronics*, Proc. IEEE 93(9) (2005) 1559.
55. T. Chen et al., *CMOS Reliability Issues for Emerging Cryogenic Lunar Electronics Applications*, Solid-State Electron. 50 (2006) 959.
56. G. De Geronimo et al., *Front-End ASIC for a Liquid Argon TPC*, IEEE Trans. Nucl. Sci., Vol. 58, No. 3, June (2011) 1376–1385.
57. S. Li et al., *LAr TPC Electronics CMOS Lifetime at 300 K and 77 K and Reliability Under Thermal Cycling*, IEEE Trans. Nucl. Sci. Vol. 60, No. 6, Dec. (2013) 4737–4743.
58. S. Amerio et al., *Design, Construction and Tests of the ICARUS T600 Detector*, Nucl. Instrum. Meth. A 527(2004) 329, and references therein.
59. V. Radeka et al., *Cold Electronics for “Giant” Liquid Argon Time Projection Chambers*, Journal of Physics: Conference Series, Vol. 308, No. 1, (2011) 012021.
60. H. Chen et al., *Cryogenic Readout Electronics R&D for MicroBooNE and Beyond*, Nucl. Instrum. & Meth. A 623 (2010) 391–393.

61. R. Acciarri et al., *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*, *JINST*, 12:P08003, 2017, 1705.07341.
62. T. Hiramoto et al., *Emerging Nanoscale Silicon devices Taking Advantage of Nanostructure Physics*, *IBM J. Res. & Dev.* 50(4/5) (2006) 411.
63. M. Haselman, S. Hauck, *The Future of Integrated circuits: A Survey of Nanoelectronics*, *Proc. IEEE* 98(1) (2010) 11.
64. Int. J. High Speed Electron. Syst., *Issue on Nanotubes and Nanowires*, 16(4) (2006).
65. D.E. Nikonov, *JSNM* 21, 497 (2008).
66. F. Schwierz, *Graphene Transistors: Status, Prospects, and Problems*, *Proc. IEEE*, 101(7) (2013) 1567.
67. D.D. Smith, *Single Electron Devices*, *Int. J. High Speed Electron. Syst.* 9(1) (1998) 165.
68. M.H. Devoret, R.J. Schoelkopf, *Amplifying Quantum Signals with the Single-Electron Transistor*, *Nature* 406(31 Aug) (2000) 1039.
69. R.P. Kraft et al., *Soft X-ray Spectroscopy with Sub-Electron Readnoise Charge-Coupled Devices*, *Nucl. Instrum. Meth. A* 361 (1995) 372.
70. G. De Geronimo, P. O'Connor, *A CMOS Detector Leakage Current Self-Adaptable Continuous Reset System: Theoretical Analysis*, *Nucl. Instrum. Meth. A* 421 (1999) 322.
71. P. O'Connor, G. De Geronimo, *Prospects for Charge Sensitive Amplifiers in Scaled CMOS*, *Nucl. Instrum. Meth. A* 480 (2002) 713.
72. D. E. Nikonov and I. A. Young, *Overview of Beyond-CMOS Devices and a Uniform Methodology for Their Benchmarking*, *Proc. IEEE*, 101(12) 2013.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 11

## Detector Simulation



**J. Apostolakis**

This chapter provides an overview of particle and radiation transport simulation, as it is used in the simulation of detectors in High Energy and Nuclear Physics (HENP) experiments and, briefly, in other application areas. The past decade has seen significant growth in the availability of large networked computing power and particle transport tools with increasing precision available to HENP experiments, and enabled the use of detailed simulation at an unprecedented scale.

After describing the uses of detector simulation and giving an overview of its components, we will examine selected cases and key uses of detector simulation in experiments and its impact.

### 11.1 Overview of Detector Simulation

#### *11.1.1 Uses of Detector Simulation*

Simulating the generation of particles in an initial collision, the interaction of these primaries and their daughter particles with the material of a detector and the response of the detector is a key element of recent experiments. Its importance has grown with each generation of experiments from LEP, B-factories, and the Tevatron, through to the current generation at the LHC, due to the increased precision requirements. It will be a significant element in planned experiments at super-B factories, the International Linear Collider (ILC), the Future Circular Collider (FCC), and also for numerous ongoing and smaller future experiments.

---

J. Apostolakis (✉)  
CERN, Geneva, Switzerland  
e-mail: [john.apostolakis@cern.ch](mailto:john.apostolakis@cern.ch)

Simulation serves many purposes at each point in the lifecycle of an experiment or a facility. Different types of simulation are used, typically with an increasing level of detail during a lifecycle. At first, fast simulation follows only the most energetic particles, typically the particles arising from the primary interaction, through simplified geometrical descriptions to obtain average values for energy deposition in the key volumes in which a signal is generated. Later, in more detailed simulation, the interactions create secondary particles, which in turn are tracked and create further descendant particles; interactions are treated with more detail and energy deposition includes fluctuations. This enables the estimation of many quantities including the measurement resolution of detectors, and correlations.

To prepare a proposal for an experiment, different versions of the setup are simulated. For each design a simplified version is constructed and simulated, typically using a tool such as SLIC [1] or DDG4 [2] which provide templates for detector components and hit generation. Putting this together with tailored digitization and reconstruction, many key characteristics of a design necessary for technical design report [3] can be evaluated.

The energy dispersion or resolution of calorimeters, the longitudinal and lateral leakage, corrections to momentum measurements, the backscattering of particles into trackers or other ‘upstream’ components can be estimated for different designs. Accurate simulation is a powerful quantitative tool for optimizing the performance, the size and cost of each sub-detector. In addition, its use extends to quantifying the tradeoffs between the performance of combinations of detectors, and finally for optimizing the global performance of a complex modern detector.

During the prototyping and calibration phase of a detector, simulation is utilized for test beam setups of single detectors or combinations to ensure that their performance is understood and can be accurately modeled. The accuracy expected in today’s high precision experiments requires agreement between simulation and test beam measurements at the level of 1% or better. This type of high-quality quantitative comparison is the basis for evaluating through simulation significant corrections to measurements in the experiment that cannot be obtained in a test beam or directly from in-situ data once the experiment starts operating.

The possibility for detailed modeling of the conversion of energy deposition into signal within a detector is a key strength of detector simulation. This requires detailed knowledge of many detector-specific effects. Simulation is an important tool also for the data analysis phase of an experiment. After its accuracy has been validated against single-particle benchmarks and test beams, an experiment’s simulation can be utilized to model tracks of different types, and even full signal and background events. Detailed quantitative aspects of simulated tracks are used in preparing methods for particle identification and measurement before the start of an experiment, and continue to be crucial in many methods that utilize data to calibrate complex quantities such as the jet energy scale. It can also be used in modeling the separation of signal from background contributions to measured data.

The impact of a performant detector simulation, whose predictions matched data, was demonstrated in the first years of the operation of the LHC experiments, during



which reconstruction and calibration in many channels were achieved within the first years of operation, a fraction of the time of previous hadron collider experiments [4].

A challenging aspect of the use of simulation in data analysis is the estimation of the systematic errors of the simulation. The accuracy of a simulation depends on the accuracy of the description of the detector's geometry and material composition, the knowledge of the beam or primary particles delivered and the intrinsic accuracy of the simulation tool(s) used. The detector simulation tools, in turn, are limited by several factors: the availability and known accuracy of measurements utilized to tune or validate the physics models, in particular of the cross-sections; the limitations of the physical models in reproducing the energy spectra and other properties of interactions; the approximations utilized to obtain adequate computational speed, to simulate the required number of events using the available computing resources.

There is an important tradeoff between the level of detail, both in the geometrical description of a setup and the choice of physics modeling options, and the computational cost of large-scale simulation. In the past 5 years the LHC experiments have been able to use detailed simulation to produce several billion events per year [5] providing unprecedented support of analysis in hadron collider physics. The increase in luminosity in the HL-LHC era will bring a need for much higher statistics of simulated events, whereas projected growth in computing power is forecast to be modest in comparison [6]. This is driving research to achieve substantial performance gains in full simulation in GeantV [7], and the expectation that faster approximate simulation will be relied upon once again for most analyses—leading to efforts to create new kinds of fast simulation which more accurately capture additional features of the full simulation, including fluctuations of key quantities.

## 11.2 Stages and Types of Simulation

- Event generators and detector simulation
- Scale from full detail to fast simulation
- Simulation of energy deposition or signal generation
- Assessment of radiation effects
- Key tools: Event generators, detector Monte Carlo, radiation transport
  - Detector Monte Carlo: GEANT, FLUKA, GEANT4
  - Radiation related MC: FLUKA, MARS, MCNP/MCNPX
  - Signal generation: Garfield

The simulation of the passage of particles through a detector and the response of the detector's sensitive elements typically proceeds in different stages. In the first stage an external event generator simulates the initial interaction and then decays short-lived particles; the results are the “primary” tracks. The second stage is detector simulation and involves the tracking of the primary particles through the structures of the detector, sampling interactions with their components, and creating

secondary particles. In the third stage the “hit” information is processed, to estimate the signals that result.

In detector simulation the secondary particles and their descendant particles are tracked in turn. Information about the passage of particles through the sensitive parts of the detector is recorded as ‘hit’ information. For tracking detectors usually the individual position, momentum and particle type or charge information of each track is recorded; in calorimeters, the energy deposition within a cell is kept as a sum over all tracks. A key characteristic of detector simulation is that tracks are treated independently. Each particle created is tracked in turn, until all have exited the setup, suffered a destructive interaction, or have been dropped as unimportant according to a user’s chosen criteria. The dropping can be triggered, for example, when the energy of a track falls below a threshold or arrives in a particular ‘unimportant’ region. Potential indirect interactions between particles are not treated as part of the detailed simulation. As such, the creation of space charge in a gaseous detector must be introduced in an experiment’s ‘user’ code or else treated separately.

### ***11.2.1 Tools for Event Generation and Detector Simulation***

The creation of the primary particles by the high-energy interaction is modeled using specialized event generators. The type of interaction, energy range and applicability of these generators differ significantly: whether they include hadron–nucleus and/or nucleus–nucleus interactions, or the type of physics beyond the standard model they provide. Typically event generators are independent programs: including the established PYTHIA [8] and FRITIOF [9], which use the Lund fragmentation model [10], and more recent ones such as HERWIG [11]/HERWIG++ [12]. Most provide users with tunable parameters and the ability to create sets of parameter values (‘tunes’) compatible with the most relevant reference data at the energies of interest.

Some Monte Carlo tools include high-energy event generators: e.g. DPMJET is available in FLUKA, and has been used to simulate ion–ion collisions at RHIC and the collisions of high-energy cosmic rays in the atmosphere.

Codes for the simulation of the detector must handle geometries of significant complexity and a large number of volumes and they must model the full set of hadronic, electromagnetic and weak interactions as accurately as required, potentially within constraints of CPU time. In the past 20 years different tools have been used for this purpose, including GEANT 3, FLUKA and GEANT4. Other multi-particle codes for particular applications including the MARS [13, 14] code, and the SHIELD code which focus on ion–ion interactions. Different codes share some physics models; for example PHITS and MARS share several models with MCNPX, an extension of the neutron-gamma gold-standard code MCNP.

GEANT version 3 [15] was utilized by LEP experiments (ALEPH, L3), the TeVatron experiments at Fermilab, numerous other experiments and also by the ALICE experiment at LHC as its main simulation engine. It includes detailed descriptions of electromagnetic interactions down to 10 keV. For hadronic physics it

relies on external packages: GHEISHA [16], GCALOR [17], which uses CALOR89 [18], and GFLUKA, which interfaces to the 1993 version of FLUKA [19].

FLUKA [20], after a major overhaul in the 1990s, offers microscopic models for 60 elementary particles, all types of ions at energies from 1 keV to 10,000 TeV/A. It has been used for detector simulation by the Opera and ALICE experiments, in radiation assessment, accelerator collimation and target tuning, and many applications beyond HEP. A key emphasis and strength have been its single, consistent, core hadronic model, PEANUT, with a Dual Parton Model (DPM) based high energy cascade above 5 GeV, a generalized intranuclear cascade and suite of models for the excited nucleus. For nucleus–nucleus interactions above 5 GeV, it utilizes interfaces to DPMJET-III [21] event generator for interactions. There is an option for the detailed treatment of neutrons down to thermal energies, which uses the multi-group approach involving energy bins, and weighted averages of cross sections and interaction production. Physics processes for electromagnetic interactions and lepto-nuclear interactions are included. FLUKA focuses on a single set of physics processes, which are curated and validated extensively by its authors. A small set of optional variations of physics processes are provided, e.g. for the simulation of low-energy neutrons. The majority of uses in HEP lie outside detector simulation. Examples include the estimation of radiation backgrounds in experimental areas, whether in accelerator facilities or underground halls, and the modeling of beam interactions with collimators in accelerators. Extensive studies of the LHC radiation environment have been carried out using it over the past decade, and a FLUKA model of the full LHC collider is the production simulation for radiation studies.

GEANT4 [13] is the basis of the simulations of BaBar, ATLAS, CMS, LHCb and a large number of smaller experiments. Its standard configuration provides electromagnetic interactions for charged particles and gammas down to 1 keV, hadronic processes for nucleons, mesons and ions, models of electro-nuclear, lepto-nuclear interactions, radioactive decay of nuclei and optical processes for photons at visible and near-visible energies. A variety of hadronic processes has been used to span all projectiles at energies up to 1 TeV, with recent extension to 100 TeV for Future Circular Collider (FCC) applications [22]. An option for neutron interactions from 20 MeV down to thermal energies is available using cross-sections for individual elements and isotopes (a technique called ‘point-wise’). GEANT4 takes a toolkit approach, enabling and requiring its users to choose the parts required for their application area, including the configuration of physics models. Recommendations of physics model configurations are provided for several established type of application and for a number of HEP and external application domains; validations for several HEP applications have been undertaken in collaboration with experiments. For other application domains users are invited to undertake the appropriate validation, potentially using their specific data, and interacting with GEANT4 experts.

A few other codes provide extensive modeling of multiple particle types, including the PHITS [23] code and MCNP family. The most recent MCNP version 6 [24, 25] was created from the merger of MCNP5, which focused more on traditional neutron-gamma applications including simulation of nuclear facilities

and reactors and the all-particle offshoot MCNPX. Its models will be contrasted with the capabilities of FLUKA and GEANT4, but it has not seen use in HEP detectors, due to lack of electron/gamma models above 1 GeV, restrictive licensing and export control.

In order to compare with measurements, the response of detectors to the energy deposited by an event's tracks must be estimated. One tool used for the detailed estimation of the energy deposition in a gas detector is Garfield [26]. It generates low-energy gammas and electrons, down to eV energy using HEED [27], uses the electric field calculated externally, and transports electrons and ions under the combined influence of its electromagnetic field and diffusion. The recent Garfield++ rewrite [28] and extension extended its capabilities, added refined electron transport and physics models for semiconductors, and enabled interfacing with GEANT4 [29]. Its computational cost is 2–3 orders of magnitude larger, so it is utilized sparingly: for studies and to generate an accurate parameterization of a detector's response [30] for use in large scale simulations.

### ***11.2.2 Level of Simulation and Computation Time***

The modeling of every physical interaction, from the initial particle's energy down to the interaction of eV scale photons and electrons—or even the interactions of neutrons down to thermal energies—is possible. The computing cost of such simulation is prohibitive for most practical applications, and simplifications are required. Yet in some cases it is necessary to simulate down to very low energies, for example in order to estimate the activation of materials by neutrons.

In complex detectors, such as in an LHC detector, the full simulation of each event takes between 0.1 and a few minutes on modern computers, depending on the type of event (minimum bias or t/t-bar) and the region simulated (rapidity coverage). This limits the number of events that can be simulated.

In some applications the simulation effort can be reduced for many events: by simulating first the particles that are involved in the trigger. Otherwise, one may seek to limit the number of secondary particles generated or the total track length simulated, or to simplify the treatment of the most frequent interactions.

Another alternative is fast simulation. This involves selecting only a fraction of tracks for simulation, and approximating the detector and key physics interactions in order to reduce the computation time per event by one, two or more orders of magnitude. Fast simulation is a powerful tool for modern experiments, as it allows speedy turnaround to address changing conditions or assumptions, and to explore different model parameters at an affordable computing cost. It can be calibrated using full simulation, data or both. However it is not capable to estimate resolutions and correlations, and it can be harder to obtain systematic errors.

Recently ATLAS has created a hybrid simulation mode by selecting for detailed simulation the conical regions of the detector around the most energetic primary particles, and using fast simulation models for the remainder [31].

### ***11.2.3 Radiation Effects and Background Studies***

The background in a modern detector can be due to many factors, including remnants of past events, accelerator generated backgrounds, and the backscattering of particles by the detector's surroundings. These can require simulation, in order to determine their level and characteristics. Also in many cases effects of the experiment on its surroundings or its constituents such as activation must be estimated. Simulation is an essential component.

Tools that are utilized for these tasks include FLUKA and the MARS code [32]. In addition to inclusive physics models, where the whole interaction is simulated, MARS contains exclusive models, where the leading particles and a sample of other secondaries is produced by an interaction.

Biasing is a technique in which some 'unlikely' trajectories are enhanced by a large numerical factor and assigned a weight inverse to this factor, in order to rapidly estimate their effect. It is an essential component of background applications. In many cases a result cannot be obtained without it; in others it improves greatly the accuracy of the result. Good statistical accuracy can be achieved within a fraction of the computation time required for an unbiased, so-called 'analogue', calculation for means and similar observables. Correlations, widths and other second order observables can be obtained only in some cases and by recording key additional information during simulation.

MARS has been used for accelerator and background calculations for many facilities [14] and experiments [33]. FLUKA also has seen wide application in this domain.

## **11.3 Components of Detector Simulation**

- Geometry description and navigation
- External fields
- Electromagnetic physics models
- Hadronic physics models
- Low-energy neutron interactions
- Accuracy of simulation
- Fast simulation
- Signal generation
- Biasing, production thresholds

A complete tool for simulating particle interactions and detector response must include the description of a detector's geometry and material, the input or selection of primary particles, the modeling of all relevant physical interactions and the extraction of information such as the energy deposition and particle passage (hits).

Most tools also account for the effects of external electromagnetic fields on charged particles, provide visualisation of the geometry and simulated events. They provide for tallies, output of key physical quantities calculated during the simulation, such as totals of energy deposition, dose, particle flux and fluence. They also provide the opportunity for the user's code to filter and record track quantities at each step.

The geometry module provides the ability to describe the material composition and the geometry of the setup in terms of volumes. The tool must be able to navigate inside this volume description, identify the volume in which a point is located and calculate the distance to the next boundary in a given direction. The capabilities of the geometry modeler determine the type of volumes and their relative placement: whether volumes are generated directly as finite shapes, or whether they are the result of the intersection of surfaces; whether all volumes must be placed within a single 'world' volume, or a hierarchy of volumes can be created. In order to simulate a large, complex detector with hundreds of thousands to millions of volumes, the geometry modeler needs to support hierarchical geometry definitions.

To ensure good performance the key geometry operations must be computationally efficient; in particular, the computation of the distance to intersecting a boundary is critical. Optimisation methods which rely on data precomputed at initialization inspired by ray-tracing are used to greatly reduce computation time wherever many candidate sub-volumes exist.

Some experiments have chosen to use a geometry modeler external to the simulation tool. They use the same geometry description and modeler inside a Virtual Monte Carlo framework [34]. This interfaces to different simulation tools for modeling interactions: GEANT, FLUKA and GEANT4; they are labeled 'physics engines' and can be selected at runtime.

### **11.3.1 External Fields**

The effects of external electromagnetic fields on the trajectory and energy of a charged particle track are modeled utilizing the Lorentz equation. The equations of motion for the position, the momentum and optionally the polarisation of the particle are integrated to obtain the position and state of a particle after a distance  $s$ . In special cases, such as a constant magnetic field, an analytical solution can be used. In the general case, numerical integration is used, typically with a Runge-Kutta method.

After integration the idealized curved path of a particle track in a magnetic field is propagated through the geometry of a detector. The curved trajectory is split into linear chord segments, which are used to navigate in the model geometry. The intersection of a chord segment is progressively refined to identify the location where the curved path crosses a geometry boundary.

### 11.3.2 Introduction to the Transport Monte Carlo Method

At each step of a simulation, the Monte Carlo method for particle transport needs:

- the cross sections in the current material of each physical interaction;
- an algorithm to select which interaction occurs next;
- a method to apply the effects of each interaction: to generate new particles, and change the state of the potential surviving projectile.

The Monte Carlo method [35], general techniques [36] and its application to particle transport for charged and neutral particle transport [37] are well described in the literature. We touch on a few of the essential features.

A key ingredient is a source of ‘pseudorandom’ numbers, distributed uniformly in an interval, usually [0,1). These are obtained from pseudorandom number generators [38] and must come with guarantees of non-correlation, such as those provided by the generators based on ergodic theory, MIXMAX [39, 40] and RANLUX [41, 42], or at least have survived a barrage of empirical tests [43] to suggest there are no correlations which affect the Monte Carlo estimates.

For a general particle the total interaction cross-section (summed over all interactions)

$$\sigma_{\text{total}} = \sum_i \sigma_i$$

is used to sample the step length  $s$ , using a random number  $r$  from the interval (0,1):

$$s = -(1/\mu) \ln r$$

where the absorption coefficient  $\mu$  is proportional to the cross-section  $\sigma_{\text{total}}$  and density  $\rho$ . Thereafter, the type of interaction that will occur at this step is chosen. The probability for one particular type of interaction to occur (in one step) is proportional to its cross section.

In the ideal case, all interactions would be sampled this way. However, in practice a different approach is needed, as the cross section diverges for the emission of soft photons and delta rays. A systematic treatment proposed by Berger [44], separates collisions that alter the state of the particle below a chosen threshold, typically for the momentum transfer. These are not sampled individually; only their collective effect is sampled. The collisions above the threshold are simulated individually.

In this approach, the part of the cross section corresponding to an interaction below this threshold is labeled the continuous part, and it does not contribute to limiting the step. Its effect is applied separately as an integral over the length of the step, to the state of a track.

The discrete part of an interaction contributes its cross-section to limiting the step

$$\sigma = \sigma_{\text{discrete}} + \sigma_{\text{continuous}}$$

Its cross section represents all interactions resulting in secondary particles with energy  $E$  above the threshold energy  $E_0$ :

$$\sigma_{discrete} = \int_{E_0}^{\infty} \frac{d\sigma}{dE} dE$$

This treatment is required for the Bremsstrahlung process and delta-ray production, due to the large number of secondaries produced with low energy.

### 11.3.3 *Electromagnetic Interactions and Their Modeling*

The modeling of physical processes can be separated into models of electromagnetic (EM) interactions, models of hadronic interactions (involving the strong nuclear force) and the decay of unstable particles mediated by the weak force. The Monte Carlo simulation of EM interactions of charged particles with atoms has been well established in HEP applications since the advent of EGS4 in the 1980s. EGS4 was able to produce and track photons, electrons and positrons down to 10 keV.

At typical HEP energies of 1–100 GeV, the number of particles of the electromagnetic shower is large. The full simulation of all resulting particles is costly in computational resources, and a selection of particles is undertaken to represent the shower. Typically, particles are tracked until they reach a certain energy threshold, the tracking cut, and discarded. In addition, secondary particles are emitted only if their energy is above a chosen energy, called the production threshold. For specific applications the high density of energy deposition near the endpoint of a track (Bragg peak) is relevant, and can be simulated.

Electromagnetic interactions of gamma-ray photons include Compton scattering, the photoelectric effect and, gamma conversion, the production of electron–positron pairs. Cross sections for each process are calculated directly from theoretical or empirical formulae, or parameterized. For example the Klein-Nishina formula is used for the cross section of Compton scattering. To improve execution speed the value of the cross sections are pre-calculated at several energies; the value at any other energy is obtained by interpolation.

The method used to model multiple scattering, in particular near boundaries, is a key feature of a simulation tool. Obtaining accurate results using less computing power, and obtaining results that are stable when varying parameters (such as the production threshold or tracking cut) are significant algorithmic challenges.

The EGS approach for the simulation of photons and electrons and its implementation were pioneered by Nagel. It was improved and shared within the HEP community as the EGS3 [45] and EGS4 [46] code systems. From these other HEP codes for EM interactions are descended, or inspired.

The underlying assumptions in Monte Carlo simulation of radiation transport are the same amongst these and modern codes: materials are assumed to be amorphous,



and beam particles do not interact. The methods for modelling transport of photons and electrons used in Monte Carlo codes are based on sampling of differential cross sections obtained from approximate theoretical calculations. A recent review offered a comprehensive description of the principles and approximations [47] for models of electron and photons up to 1 GeV, documenting widely used models and those of the precise modern electron–photon code Penelope [48].

### ***11.3.4 Interactions of Photons***

Photon interactions are ‘discrete’ interactions, that occur at a point and can be modeled this way. This makes them much simpler than modeling the interactions of charged particles. Interactions considered including photoelectric, Compton incoherent scattering, electron–positron pair production and potentially Rayleigh coherent scattering.

The cross section for each interaction is sampled from measured or theoretical distributions. In some cases a simplified form is used, to reduce the cost of computation with a simplified description of the energy and Z dependence. Else, the values for each material at particular values can be pre-computed and stored in tables for interpolation.

Once the type of interaction is chosen, its products are sampled from the appropriate distributions. Pseudorandom numbers are used to sample the energy, angles and momenta from the differential distributions [37]. The original particle’s state, if it survives, is altered to preserve energy and momentum.

The interaction of energetic photons with nuclei is discussed in the hadronic section below. Often specialized tools are used to simulate optical photons and their collection. It is possible, though, to generate optical photons and model reflection, refraction and absorption on different types of surfaces. GEANT4 is able to do this.

### ***11.3.5 Interactions of Charged Particles***

The simulation of the electromagnetic (EM) interactions of charged particles is complicated by the large cross section for elastic interactions and of ionization, which produces low-energy electrons (delta electrons).

In a few cases it is useful to simulate every single interaction of a charged particle in a medium, including its elastic collisions with nuclei, the ionization of atoms and creation of delta electrons, and the ‘hard’ interactions, which create photons or electron–positron pairs.

All production simulation tools estimate the cumulative effect of the elastic scattering off nuclei. It is modeled in several different ways. Many utilize the multiple-scattering approach pioneered by Goudsmit and Sanderson [49] as their basis. One simple way to sample angular deviation over a short step is Moliere’s

theory [50], which is used in GEANT 3, but is limited to small angles. The approach of GEANT4 borrows from Lewis' description [51].

Key effects of multiple scattering are angular deflections and straggling. The latter's most important effect is the shortening of the distance travelled in the direction of the initial momentum. This must be modeled in order to obtain correct energy loss for the passage through material, as is done in GEANT 3. The second effect of straggling, the displacement in the lateral directions is correlated with angular deflection. Similarly to EGS4, FLUKA and GEANT4 also sample this displacement, each using different algorithms. This enables longer steps while maintaining accuracy. The best algorithms allow longer steps, or more accurate modeling of the correlations between the affected changes in the state of the particle.

The algorithm for multiple scattering has a significant effect on the results obtained in many detectors and setups. Examples include the partition of energy in sampling calorimeters, the correlation between the deflection of muons and their positions after substantial material. In particular, many quantities are very sensitive to the details of its formulation and implementation. These include the fraction of low-energy electrons ( $T \ll 1$  MeV) scattering backwards at the interface between low and high  $Z$  materials and the correlation between the direction of a particle exiting a detector (e.g. muon) and its position.

In addition to the sampling of the final state, high accuracy for electron transport necessitates careful treatment of multiple scattering of low-energy electrons at boundary crossing [52]. New algorithms have been developed for exact electron transport without special treatment for boundary crossing [53, 54]. These algorithms have been implemented in electron-photon Monte Carlo codes: PENELOPE [55], EGSnrc [56] and EGS5 [57]. A comparison [58] in 2007 benchmarked the algorithms in several of these codes and in GEANT4, using data from custom setups with thin slabs. PENELOPE and EGSnrc demonstrated the best performance, while GEANT4 obtained good results only with specific settings.

The new GEANT4 GS model [59], available since GEANT4 release 10.3, implements Kawrakow's approach to provide angular deflections for any size step without free parameters and offers the option for accurate boundary crossing. It achieves the best agreement, amongst GEANT4 models, with a wide range of benchmark data including backscattering data.

Models for specialized processes, such as transition radiation, exist in some tools including GEANT4 [60].

### ***11.3.6 Hadronic Interactions and Their Modeling***

In contrast to the simulation of EM physics processes, the simulation of hadronic physics processes from first principles is not possible, except partially at the high energy limit. At all energies, the cross-sections and the models used are based, directly or indirectly, on measured data of hadron-nucleon and hadron-nucleus interactions, and on phenomenology.

The most common particles produced by hadronic interactions are nucleons, pions and kaons. The diversity of particles and interactions make modeling a great challenge. Specialised codes including HETC [61], GHEISHA, CALOR [18], FLUKA, and SHIELD [62] were developed for HEP and other application areas in the 1970s and 1980s. Few models and codes from other application domains have been available which cover a substantial part of the energy range (above the 1–2 GeV used in spallation applications) and the full set of particles needed for HEP. One exception is the MCNPX tool, an extension of the MCNP code for neutron/gamma radiation transport and reactor simulation. MCNPX hadronic models are shared with MARS.

Interactions are divided into elastic and inelastic, which produce new particles in the final state. The smaller cross-sections for inelastic interactions of hadrons with nucleons compared to EM interactions, and the growing multiplicity and variety of particles emitted in interactions above a few GeV, result in significantly different structure for hadronic showers.

The modeling approach depends on several factors: the availability of detailed experimental measurements; the complexity of final states of reactions for a particular combination of incident particle, energy and target; and the availability and suitability of theoretical or phenomenological descriptions. In many application domains there is a requirement for conservation of energy, momentum and quantum number in each interaction, and for the coincidence or correlation between the products of an interaction; in selected cases conservation of energy only as an average over different interactions may suffice. A small number of interaction models, including GHEISHA, and most low-energy neutron interactions sampling methods treat particle interactions only in the average, and do not conserve energy and momentum.

For many applications full energy conservation of individual interactions and the treatment of the correlations of particle tracks is required in order to obtain reliable results. For example, the estimation of the energy resolution of hadronic calorimeters is strongly affected by these factors.

In many cases a phenomenological model is supplemented by fits of model parameters with available data. In a few cases (evaluated) data libraries are used directly—typically for low-energy neutron transport. Another approach is to use parameterizations, either directly of data or indirectly for the parameters of simplified models, as in GHEISHA.

### ***11.3.7 Models of Interactions at Low Energies***

At the lowest energies, the largest hadronic cross section belongs to the elastic interaction, which is a coherent interaction of a hadron projectile with the full nucleus. Hadronic cross sections, including those for elastic scattering, are typically parameterized from data.

Inelastic interactions, which excite the nucleus, typically become relevant at energies of order MeV. They are modeled with a statistical approach for energies up to about 100 MeV. The original Weisskopf evaporation model [63] describes the emission of protons and neutrons from nuclei in thermodynamic equilibrium. It is supplemented by several additional de-excitation channels, which compete to occur. These include the multi-fragmentation model for highly excited nuclei [64], Fermi breakup of light nuclei, fission of heavy nuclei and photon evaporation.

In an alternative approach, following the Generalised Evaporation Model (GEM) of Furihata [65], nuclei with up to 28 nucleons are evaporated directly. This improves greatly the description of the emission of light and medium fragments, with an extra computational cost. GEM is an option in MCPNX, in GEANT4 and a similar approach is used in FLUKA. Fragments heavier than  $^4\text{He}$ , though emitted infrequently, are important for specific applications, such as the response of silicon devices and damage to them; using the GEM approach is recommended for these.

At energies up to about 300 MeV a simple algorithm can be used to count the number of excited nucleons and holes of missing nucleons in the Cascade Exciton Model (CEM) model [66]. Such models are called pre-compound (or pre-equilibrium) models. A pre-compound model of this type is implemented in GEANT4. And CEM.03 [67], which is included in MARS, MCNPX and MCNP 6, is an improved CEM descendant.

These models are also used to calculate the de-excitation of nuclei after interactions at the higher energies, important in many applications. For example, they determine the energy that is lost to nuclear breakup, and the partitioning of energy between the low-energy protons, neutrons and gammas that are produced. These processes produce the majority of neutrons, whatever the initial interaction, and as a result affect substantially the escaping energy, lateral shower profiles and compensation of calorimeters—amongst other observables.

### ***11.3.8 Cascade Models of Hadron–Nucleus Interactions at Intermediate Energy***

At energies above about 100 MeV an intranuclear cascade model is used for nucleon and pion projectiles. In a cascade the interaction is modeled as a succession of independent collisions of the projectile (and secondaries) with individual nucleons inside the target nucleus [68].

In a cascade, the nucleus is described in two ways. It can be an ensemble of nucleons positioned at random locations, sampled from a model of nuclear density—as used Quantum Molecular Dynamics (QMD) models such as UrQMD [69], in the GEANT4 Binary cascade model [70] and in the Liege cascade INCL. Else, the nucleus can be composed of a number of shells of constant density, as in the original cascade of Bertini, in FLUKA, and in INUCL and its descendant, the

GEANT4 Bertini-type cascade [71]. A correction factor is used for the depletion of nuclear shells by earlier interactions in both FLUKA and GEANT4 Bertini.

It uses ‘free-space’ cross-sections derived from hadron–proton measurements, or, in some cases cross-sections modified for the presence of the nuclear medium [72]. FLUKA accounts for nuclear medium effect on the  $\Delta$  resonance properties in the treatment of pion interactions [20].

Hadrons may move in curved trajectories according to a chosen nuclear potential, as in GEANT4 Binary and FLUKA, or in straight lines (GEANT4 Bertini). In both cases, the potential is used to update the momentum of all hadrons before interactions. Interaction products can use models or be sampled from data, must observe the Pauli exclusion principle, and are subject to a hard-core nucleon repulsion. Particles arriving at the nuclear boundary with enough energy are ejected, while others are reflected and continue to interact. The difference of the total energy of the remaining nucleons and the ground state energy that corresponds to them is the excitation energy.

After either a fixed time or once the excitation energy has dipped below a threshold, the remaining nucleons are handed to a pre-equilibrium or de-excitation module. A pre-equilibrium model, such as the Precompound model, can eject higher energy nucleons and is used with the GEANT4 Binary cascade. A similar model is used in FLUKA. The subsequent de-excitation module combines evaporation, Fermi-breakup for light nuclei, fission for heavy nuclei and other competitive channels.

A common de-excitation module is shared by all models in FLUKA. A custom simpler de-excitation module is used in GEANT4 Bertini; recent extensions enabled it to use the default “Preco” Pre Compound and de-excitation module, used after Binary cascade and the higher energy string models (QGS and FTF). In MCNPX, MARS and MCNP a common module is used with different parameters by the CEM cascade and the higher energy LAQGSM [67] model.

Most cascade models are expected to work up to 1.5–3 GeV, yet they can provide good results from 30 MeV up to 3–5 GeV. At higher energies their assumptions break down, because quark degrees of freedom become important. Including additional reactions with larger multiplicities of products, and effects such as formation time (a simplified treatment of quantum-mechanical effects as a time interval before secondary hadrons can interact) allows a cascade model to have a higher energy limit. This is the case for the GINC/PEANUT cascade in FLUKA and the GEANT4 Bertini cascade.

Early versions of GEANT4 until 9.6 also included the CHIPS model [73] applicable for intermediate energies. It described a nucleus in terms of nucleon clusters, and interactions as exchanges of quarks and was part of physics lists used in LHC Run 1.

The GEANT4 Bertini cascade underwent a substantial rewrite and upgrade [74]. As a result, all long-lived hadrons can be projectiles (adding K,  $\Lambda$ ,  $\Sigma$ ,  $\Xi$  and  $\Omega$ ). It also implements gamma- and lepto-nuclear reactions. Its energy range was extended up to 15–20 GeV, with the addition of final states with higher multiplicities, up to nine for proton–proton. Total and partial cross sections and final states were obtained from the CERN-HERA data compilations, and completed using symmetries and

general principles for unmeasured reactions and energies. The number of nuclear shells varies from one for the light, three for medium, and a maximum six for the heaviest elements.

As its main use is the simulation of LHC experiment detectors, a number of modeling and implementation choices were made to optimize computation speed. These include linear interpolation for the sampling of partial cross sections and large 10-degree bins in angular distributions—justified by the smoothing effect of additional interactions.

The newest cascade model in GEANT4, the INCL Liege cascade, is the one under the most active development. This event generator, under development since the 1990s [75], was developed to reproduce spallation data for reactions at 100 MeV–1.5 GeV using a parameter-free model. It interfaced with the ABLA code for deexcitation. Its original Fortran version up to version 4.6 [76] has interfaced to MCNPX and MCNP6. Recent development focused on the re-engineered INCL++ [77], which reproduced the performance of 4.6 and was extended to handle light ion projectiles (up to carbon16) and to higher energies, up to 15 GeV producing multiple pions.

Large suites of benchmark data are used to tune and verify the modeling of each cascade, covering neutron, proton and pion production at energies from 60 MeV to 3 GeV. Spallation data from inverse kinematic reactions on hydrogen targets at GSI with a range of projectiles from  $^{56}\text{Fe}$  [78] to  $^{238}\text{U}$  [79] at 1 GeV/nucleon provide different challenges for modeling and complement these data.

Comparisons of hadronic and in particular cascade models have been undertaken periodically under the auspices of the IAEA [80]. These benchmarks use a large set of thin target data to probe the accuracy and predictive capabilities of each model. INCL was found to be one of the competitive models. Similar test suites are utilized as a part of internal benchmarks and for tuning of models.

The details of the intranuclear cascade and the pre-equilibrium model determine all the emission of higher energy particles, but the details of the coupling and the quality of the de-excitation module are also critical to the performance of many applications from activation to calorimeter simulation. Good modelling of the resulting nuclei is required to ensure that the energy lost to nuclear breakup, a key component of non-compensation in hadronic calorimeters, is accurate. Recent results from RD52 [81] were interpreted as deficiencies in the modelling of nuclear breakup in GEANT4 version 9.4, showing that there is room for further improvement in this energy range.

### 11.3.9 High-Energy ‘String’ Models

Models for interactions at high energies (above 5–10 GeV) simulate quark-level interactions and rely on phenomenological descriptions of soft QCD interactions to generate low-energy hadrons from the remnants of the high-energy collisions. They are applicable to all hadron projectiles. Three variants are available: the Dual Parton

model [82] is implemented in DPMJET [21] and used in FLUKA. The Quark Gluon String (QGS) Model [83] in different variants is used in MARS/MCNPX, GEANT4 and in the QGSJET event generator [84]. A third model, the Fritiof model, is used and extended in the FTF model in GEANT4.

Hadrons are produced in the initial collisions and the decay of QCD strings, tubes composed of compressed gluonic fields [10] generated by the separation of colored quarks. Models implement a Lund-like string. The proto-hadrons generated by string decay, once past their formation time, together with outgoing nucleons, and the remaining cluster of nucleons are handed to a cascade for additional scattering in the nucleus (e.g. in FLUKA and in GEANT4 QGS\_BIC and FTF\_BIC physics lists using the Binary cascade), or passed directly to a pre-compound module (in the GEANT4 QGSP\_BERT and FTFP\_BERT physics list).

Developments over the last 10–15 years include a high-energy extension of the PEANUT cascade in FLUKA to undertake reactions handled previously by DPMJET; the connection of the Binary cascade in GEANT4 with the QGS model (as QGS\_BIC) to re-scatter the slow products of high-energy model. More recent developments include the extensive improvement of the Fritiof-based FTF model in GEANT4 to model light anti-nucleus interactions at low energies and at rest [85], and to include an internal Reggeon cascade, and changes in the production ratio of different types of diquarks [86].

### ***11.3.10 Treatment of Low-Energy Neutron Interactions***

Amongst particles created by hadronic interactions, neutrons survive a longer time and are among those which travel the furthest. Also, they are amongst the most numerous. This makes their treatment important for many applications and correspondingly expensive computationally. Most neutrons are emitted in the de-excitation phase of a reaction and have energies of order MeV.

Neutrons produced in high-energy interactions (above 20 MeV) also lose energy in elastic and inelastic interactions (which release protons, alpha particles or light ions—or create gamma rays) before being captured by nuclei. Only part of this process occurs on fast timescales and others are much slower ( $\mu\text{s}$ –ms). By tracking time, it is possible to emulate the time dependence of the signal. This also allows the simulation tool to use a time threshold, and abandon neutron tracks after this time, in order to save computation time.

Neutrons' contribution to the visible energy measured in a detector comes via the transfer of energy to charged particles and from capture and other reactions with nuclei that generate gammas. Elastic scattering is particularly important in organic scintillators, where interactions with hydrogen transfer significant parts of energy and momentum to the recoil proton.

In some cases, are treatment of neutron interactions at a greater level of detail is required, potentially down to thermal energies. In HEP it is needed in special cases, such as the study of activation of materials for radiation safety purposes.

The detailed treatment of energy deposition in scintillators may also require a more detailed treatment of neutrons than provided by the simpler interaction models.

Simulating neutrons below 20 MeV relies on measurements of cross-sections for key processes, which have been assembled into established data libraries. Libraries, such as JEFF [87], ENDF/B VII [88], ENDF/B VIII [89], JENDL [90] and CENDL [91], include evaluations of cross sections and distributions of secondaries for key reactions based on a combination of measurements and estimates from nuclear model codes. They cover all measured interactions, from inelastic and capture at low energies; through inelastic interactions resulting in the emission of one or two neutrons (plus gammas); to multi-neutron production at tens of MeV. For some elements, data is available for several individual isotopes, while for others only the values for the natural composition are measured. Cross sections have many resonances in the keV–MeV region, complicating precise treatment.

For reactions that result in more than two outgoing particles, in most cases only spectra are available; i.e. information on the correlations of products is not included. As a result, sampling secondaries for one interaction in a way that conserves energy and momentum requires complex algorithms and additional computation. This has recently been introduced in PHITS; other codes rely on uncorrelated sampling and conserve energy only on average.

This detailed treatment consumes significant CPU and memory, because the full set of cross sections for all isotopes of all elements is required. Variants of this approach are utilised in MCNP/MCNPX (the gold standard for neutron simulation) and in the GEANT4's NeutronHP package.

A simpler approach averages cross sections over chosen sets of nuclei and fixed bins of energy. This 'multi-group' approach provides savings in memory use compared with the detailed approach and is adopted in FLUKA, as the option for precise treatment of low-energy neutrons. Accuracy is determined by the number of sets of nuclei and of the intervals of energy—and also the choice of the grouping.

In addition to the purely electromagnetic interactions of charged particles, and the interactions of hadrons with matter, it is necessary also to simulate the interaction of electrons, positrons and gammas with nuclei which result in hadronic final states. These photo-nuclear and electro-nuclear interactions account for a small portion of the total cross-section of gammas or electrons, below one percent at its peak. Yet they are the only interactions that convert electromagnetic energy into hadronic final states in typical HEP experiments.

Models for photo-nuclear and electro-nuclear interactions are provided in all multi-particle codes discussed [92].

Given the diversity of hadronic interactions, there is a need to focus on essential aspects. What most influences the accuracy of the description of the energy response, energy resolution of calorimeters and missing energy for hadronic calorimeters?

For HEP applications (in particular calorimeters) some of the key features are:

- most energy is deposited by low-energy particles, and its spectrum is independent of the type of projectile particle [93];



- the production of  $\pi^0$  particles and the fluctuation of the energy fraction in this channel which leads to prompt EM energy deposition plays a determining role in the resolution of calorimeters [94, 95]; these  $\pi^0$ 's can be the result of charge exchange or other hadron–nucleon collisions, or formed by the soft fragmentation, e.g. in the QCD string view of high energy reactions;
- the simulation of neutron generation, transport and interaction, which contribute to prompt and delayed signal, activation, and escaping energy;
- a component of missing energy (that influences the resolution) is the energy lost to nuclear breakup. The accuracy of the modeling of all hadronic reactions, but in particular the de-excitation stage, determines the quality of its simulation;
- the simulation of leading particles which determines the shower profile—the profile of energy deposition—and in particular the sharing between longitudinal compartments and the amount exiting in the direction of the projectile.

The accuracy of the modeling of the time dependence of different interactions, both in a tool and in an experiment simulation, are also essential.

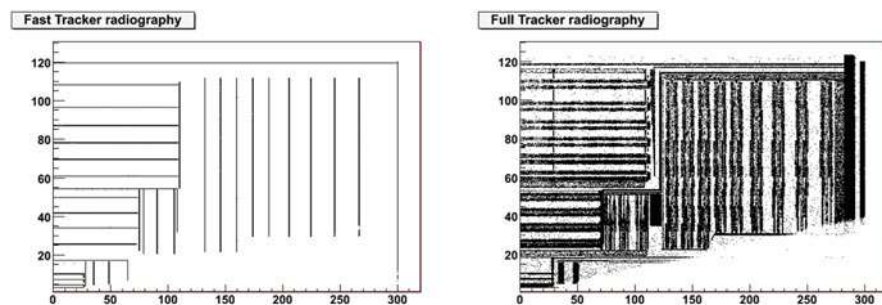
The fluctuation in the fraction of energy going into each type of secondary particle (gammas, charged hadrons, neutrons, neutrinos) in all reactions is an essential feature of a simulation tool. It is important for predicting and modeling the resolution and other aspects of detector performance.

### ***11.3.11 From Full to Parameterized ('Fast') Simulation***

In order to obtain high statistics, it is necessary for some applications to utilize simulation which is much faster (typically by two orders of magnitude) than the detailed full simulation. There is a spectrum of such simulations with different approximations and compromises.

The coarseness of simulation, from detailed to fast, is determined by a number of variables: whether the geometry is described approximately or in great detail; whether secondary particles are generated from the interaction of primary particles; the degree to which particles are eliminated during tracking—for example, the relation between the energy of primary particles and an energy threshold; and the type of physics models utilized. Choosing the least level of detail and following only the primary tracks distinguishes the fastest simulation. In other variants, some aspects are simulated in more detail in order to obtain more precision. Different ways can be used: simulating more particles, adding physics or geometry volumes.

In the first two LHC runs, it has been possible to produce billions of simulated events using detailed simulation. In some experiments the forward sub-detectors were simulated faster, e.g. use by ATLAS of frozen showers for forward calorimeters. Projections for the High Luminosity LHC era foresee an order of magnitude gap between the statistics possible using GEANT4-based detailed simulation on projected 2025 hardware and the needs of most analyses. This gap is driving the continued development of fast simulation methods, and the research into methods



**Fig. 11.1** Simulated radiography of a quarter of the CMS tracker geometry using the fast simulation FastSim (left) and full GEANT4-based simulation (right). This demonstrates the simplification of the geometry description used in fast tracker simulation, which projects the material onto smaller cylindrical shells, yet reproduces the hit structures and reconstructed tracks with the accuracy required in physics analyses. (Reproduced from [97])

for speeding up particle transport simulation. The GeantV [7] R&D aims to produce a prototype to demonstrate whether the core simulation work can be redesigned for more efficient use of current and forecast computer architectures with complex CPUs and deep memory architectures, with the goal of a speedup factor between 2 and 5.

The fastest type, so-called ‘parametric’ simulation involves the simplified propagation of tracks coupled with a reconstruction of fixed efficiency and idealized EM and hadronic calorimeters are given an input resolution. In one modern incarnation, Delphes [96], it is coupled with built-in reconstruction, and can be used to obtain first level estimates for some physics analyses from a simplified model of a detector. This type of simulation is used in the first feasibility stages of detector design, and to obtain a first understanding of physics analyses.

A more accurate type of fast simulation uses a simplified geometrical setup of a tracker device and/or sampling of showering using a parameterized distribution of a calorimeter to generate energy deposition hits and reconstruct events using about 100 times less computing resources than the full detailed simulation. An example of the simplified tracker geometry can be seen in Fig. 11.1, where the geometry of the CMS fast and full simulation of the tracker are compared.

The detailed geometry and physics of the full simulation, e.g. using GEANT4, is typically used as a yardstick for comparison of relevant physics quantities required by physics analyses, and sometimes to generate a library of pre-simulated showers at set energies for use in recreating realistic showers.

The LHC experiments have developed many different types of fast simulation, both specific to one part of a detector (tracker or calorimeter), and spanning the full detector.

One example is the mixing of detailed simulation for parts of the detector with simplified treatment within fixed regions or regions that depend on the particles inside an event or collision.

Recently ATLAS has produced an integrated simulation framework ISF [98], which includes detailed simulation and per-subdetector fast simulations. This enables the use of common modules for all elements of the simulation. A hybrid simulation mode is part of its design that selects conical regions around the most energetic primary particles for detailed simulation, and uses fast simulation models for the remaining parts of the detector. This capability has not yet been used in physics analyses.

## 11.4 Machine Learning for Fast Simulation

A completely new approach to parameterized simulation has emerged recently, exploring the potential of machine learning. One research avenue attempts to generate patterns of energy deposition that reproduced the distributions including fluctuations of key physical observables.

This has been the topic of interest and recent investigations using the generative-adversarial network (GAN) approach to the fast generation of patterns of hits. A first demonstration in a three layer LAr sampling calorimeter [99] was tested over an energy range from 1 to 100 GeV and a single particle direction. A potential speedup of  $O(100)$ – $O(10,000)$  was demonstrated. Most physical observables of interest for the classification and calibration of tracks were well reproduced, but a few showed clear differences. This promising avenue will clearly be an area of significant research in the next years.

### 11.4.1 Accuracy of Simulation

The accuracy of the simulation is determined not only by the artificial differences or defects introduced by such simplifications, but also by intrinsic factors. These factors include the accuracy of the cross-sections for particular interactions and the capability of the physical models. These can be explored by comparing with experimental data.

Criticisms of detector simulation focus on key limitations and question the predictive power of hadronic interaction modeling for use in designing and tuning hadronic calorimeters [100].

### 11.4.2 Signal Generation

In order to model the signal produced in a gaseous detector, all processes that contribute to the generation of charge and its collection in the cathode must be simulated. The detailed simulation of a small number of events modeled in full detail is used to understand the characteristics of a detector. For large-scale simulation, a

simple model or parameterization is produced for the signal generation given the energy deposition.

The simulation involves a level of detail beyond other Monte Carlo simulation for HEP detectors. Some of the important aspects include:

- modeling the generation of all secondaries, without an energy threshold, in every single inelastic collision in the gaseous volume;
- the effects of elastic collisions in the transport;
- for efficiency, pre-calculating the convolution of the effects of the resulting diffusion and the drift in the electromagnetic fields of the detector;
- the effect of potential build-up of charge on detector elements, e.g. the space charge in a gas.

Due to the need to simulate down to the eV scale, this simulation requires detailed knowledge of the excitations of the molecular constituents of the mixture. Specialised programs are necessary for this simulation. Garfield calculates the electric field in many regular cell geometries. Then it combines it with the energy deposition for each atomic shell via a specialized Photo Absorption Ionization model integrated from Heed [101] to generate all secondary gammas and electrons. Transport of the charged particles in the electromagnetic field is coupled with diffusion, using pre-calculated transport coefficients generated by the Magboltz code [102].

It is possible to simulate an avalanche near a sensor wire in order to accurately model the signal arriving at the detector's electronics. This is typically required only for the detailed understanding of the effects of the shape of the signal and the integration characteristics of the electronics, and requires three orders of magnitude more computation than the simulation of the energy deposition in the gas volume. Alternatively, a fraction of the track can be simulated, in order to determine key characteristics, such as the arrival time of the signal. This reduces the computing requirements by about an order of magnitude.

It is also possible to calculate the effect of charged particles on an integrated circuit element [103]. Using the energy deposition to create electron-hole pairs, an external Technical Computer Aided Design (TCAD) program simulates the detailed response of the circuit. Applications of this technique have focused on the simulation of single event upsets [104], in which a cosmic ray track results in the flipping of a bit in a silicon circuit. Typically, the circuit response involves proprietary TCAD programs.

### ***11.4.3 Production Thresholds and Other Biasing Techniques***

For many setups, computing resources for simulation in full detail are not available. For example, large-scale experiments can require millions of events in order to establish patterns related to rare processes, yet the full simulation of events is prohibitive. In this case the choice must be made how to discard particular tracks in

order to achieve the required computing time per event, while influencing the most important results as little as possible.

The simulation time for each particle type is proportional to the number of steps, and typically the energy of a particle (so long as it does not escape the setup). The large number of electron and gamma particles in an EM or hadronic shower necessitates that these tracks are key to reducing the computation time. For this, either the treatment of each track must be simplified, or the average track length must be greatly reduced, or the number of particles tracked must be reduced. The use of a tracking cut reduces the average track length, and the use of a production threshold reduces the number of tracks simulated.

A number of methods are widely practiced to reduce the computing time by simulating only the more important particles: to generate only particles whose energy is above a threshold energy (production threshold); to kill tracks once they fall below an energy threshold (tracking cut), or to treat neutrons via their average cross section (multi-group).

In other setups the interest is to estimate the fraction of particles passing through or around a shielding barrier, which e.g. could stop all but one particle in a million. In such cases, a method to speed up the simulation is needed. To estimate the flux of particles passing through such a barrier the transport mechanics must be changed. Most changes will favor paths which have already crossed part of the barrier or are likelier to cross—e.g. because they have higher energy [105]. There are many methods to achieve this, all of which assign a weight to a particle track. Most involve the creation of extra copies of tracks or the killing of tracks. Some of the most common are importance biasing; leading particle biasing and weight window.

Importance biasing involves separating the geometry into regions of high and low numerical importance. At the boundary between such regions particles that go from low importance to high importance are enhanced in number (splitting), and their weight is reduced in proportion. Particles moving in the opposite sense are reduced in number (Russian roulette) and each one's weight is increased.

Leading particle biasing involves sampling the results of an interaction, favoring particles that have the highest energy (and most chance to penetrate) while sampling other particles in a representative way. In all cases surviving particles from populations, which are suppressed, are given higher weight in proportion to the difficulty of survival. On the other hand, enhanced particles (where two particle tracks are created from a single one) are given reduced weight. A single event can be split into a large number of 'histories', trial tracks that carry a different weight. Physical observables must be estimated accounting for the weight of a track—which can be interpreted as a probability:

$$\langle O \rangle = \frac{1}{N} \sum w_i O_i$$

The mean value of an observable  $O$  when using event biasing is calculated using the weighted sum of the values for each particle track 'history'  $i$ , which contributes, and the total number  $N$  of events (or trial histories).

## 11.5 Case Studies

The discovery of the top quark [106, 107] involved detector simulation only in a minimal and indirect way. Detector simulation was undertaken by CDF and D0 in the optimization phase of the design of their detectors. In D0 it informed the design of the interface between the central and end calorimeters [108]. Subsequently, simplifications were made for the detector simulation used in production, to reduce the computing time per event. The individual calorimeter plates were replaced by a large block, which contained a mixture of the absorber and active material. Optionally the response of particles below 200 MeV was parameterized. Comparisons with test beam determined that the full detailed plate setup agreed well; the simplified mixture setup was found to agree less well with test beam measurements, but judged adequate for most purposes.

Events simulated using the D0 simulation program were used in the later observation of the production of single top quarks [109]. However, the detector simulation was found to have significant limitations; these appear to have stemmed both from the simplifications of the modeling of the detector and its response and from the intrinsic limitations of the simulation tool. Several corrections were required, including a factor for the efficiencies of the trigger reconstruction, and for the efficiency to identify and select particles and jets.

In the LHC experiments, simulation was first utilized to model the response of the calorimeters to muon, electron and hadron beams. One key application has been the detailed calibration of the energy response of the electromagnetic calorimeter, to obtain an estimate of the energy of the incoming particle as a function of the inferred visible energy measured as signal in the different calorimeter compartments.

The ATLAS calorimeter system is complex, utilizing different detector materials, geometrical structure and technologies for the different rapidity regions. The insensitive material between parts of a detector distorts the energy signal due to tracks crossing this region. Simulation is used to obtain correction factors for the energy of the incident particle or jet.

ATLAS has undertaken an extensive comparison of test beam results with Monte Carlo simulation. Simulation utilising GEANT 3.21 in the 1990s was used for detector design studies and the first test beam comparisons. Progressively from 2000 onwards, test beam results were compared mostly with GEANT4, and in 2004 a Geant4-based simulation was declared the official ATLAS ‘production’ detector simulation in 2004.

Comparisons were undertaken first with the test beam results for individual detectors. A typical comparison started with muons, as minimum ionizing particles, to make a first verification of the material description and geometry of active parts. Next, measurements of electrons were compared, verifying the detector description of passive parts. This determines the factor for conversion of the deposited energy into the signal measured by each detector, its electromagnetic scale, by comparison with a beam of a particular energy, typically 100 GeV. Key

observables compared include the linearity of the response, the energy resolution and the shower longitudinal and lateral shapes.

The comparisons identified areas where improvement was necessary: there were problems in the large dependence of the electron energy response on the production thresholds. These were also reported in studies related to the use of GEANT4 in medical applications [110] and were corrected in GEANT4 release 8.0.

### ***11.5.1 Calibration of EM Calorimeter Using Monte Carlo***

The calibration of the ATLAS liquid argon (LAr) Electromagnetic Calorimeter for electrons was developed utilizing a detailed Monte Carlo simulation of the detector and its test beam [111]. The simulation included a very detailed description of the complex geometry, collection of the energy deposition into hits and conversion to digitized signal.

The simulation was shown to describe well all the relevant measurements, including the mean reconstructed energy, the distributions of energy deposition in particular compartments and the energy profiles in the longitudinal compartments and lateral sections. For example, the mean reconstructed energy was described to within 2% in the Pre-Sampler and the first two (of three) compartments of the accordion. Also the distribution of the total reconstructed energy is described very well—within the uncertainty due to the upstream material ahead of the setup and in front of the PreSampler.

Based on this agreement, the simulation was used to correct for several effects, which could not be measured or could only be estimated indirectly. One involved the average energy deposited in the dead material between the pre-sampler and the first compartment; this provided an estimate of this energy deposit. Another effect was Bremsstrahlung in upstream material, due to which a fraction of events arrives at the detector with reduced energy. A quantitative description was made and compared with the measured total energy deposit.

Dedicated Monte Carlo simulations were undertaken to study the systematic uncertainties induced by each effect. As a result, the reconstructed energy response in the energy range from 15 to 180 GeV was found to be linear within 0.1% (an exception is at 10 GeV, where it was found to be 0.7% lower). The systematic uncertainties due to incomplete knowledge of the detector, the test beam, or the reconstruction were found to be about 0.1% at low energies and negligible at high energies. The effect of the non-linearity at about 40 GeV and above on the measurement of the  $W^+/W^-$  mass was found to match the aimed precision of 15 MeV—provided it can be extended from the section tested to the full calorimeter.

### ***11.5.2 Hadronic Calorimeters: Comparisons with LHC Test Beam Results***

For the hadronic calorimeters, after muons and electrons, the test beam results with pions were compared to simulation. This served to validate the simulation of hadronic interactions, since the conversion of the energy deposition into signal is common with the electron test beams. Additional observables compared included the longitudinal and lateral shower shape.

The ATLAS Hadronic Barrel Calorimeter's (TileCal) use of scintillator requires the accurate simulation of neutrons for those interactions that contribute significantly to energy deposition and which occur within its time window of 150 ns. This is treated in the simulation.

ATLAS has also undertaken extensive test beam measurements of the response of the TileCal to pions. The most recent comparisons in the Combined Test Beam setup (2004) involved pion energies up to 350 GeV. These tests [112] examined the energy response and resolution of the calorimeter, and compared them with the predictions of a simulation based on GEANT4 (version 9.1).


Events were selected based on several criteria, including the energy deposition in a cryostat scintillator (SC1) placed before the TileCal. This cut was made in order to enable comparisons with the previous test beam. Potential biases from this selection cut on the response and resolution were studied. An approximation of this cut was used in the simulation: the energy deposition in the surrounding dead material. The change in the energy response due to this cut ranged from  $-2.5\%$  to  $+0.5\%$  depending on the energy and eta value ( $0.25-0.65$ ) of the pion beam. This was reproduced within  $0.5\%$  at low energy (20 GeV) and within  $1.5\%$  at high energy (300 GeV).

The energy resolution was affected in a range from  $+10\%$  to  $-10\%$  between low and high energies respectively and reproduced within  $2\%$  for a large combination of angles and energies, except for a  $4\%$  deviation at one angle at 20 GeV. Comparing the final results for the energy response, the agreement obtained is within  $3\%$  for the full energy range studied (20–350 GeV). Typically, agreement at the  $1-2\%$  level is achieved for beam energies of 50–250 GeV; greater deviations are seen at 20 and 300 GeV (the latter, in particular, remains to be understood). For the energy resolution agreement at the  $10\%$  level is obtained.

The measurements of the reconstructed energy in the Atlas TileCal at energies from 20 GeV to 180 GeV have also been compared with simulation [113]. Figure 11.2 shows the ratio of simulation and data for the energy for the case of incident pions using different configurations of physics models in 2010 with GEANT4 version 9.2. Agreement for the mean energy between simulation and data ranges from  $-10\%$  (the legacy LHEP physics modeling) to  $+3\%$  (the production physics list QGSP\_BERT). The root mean square deviation (RMS) of the reconstructed energy agrees within  $-4\%$  (QGSP\_BERT) to about  $+15\%$  (LHEP).

Since a 1 MeV neutron could travel only 3 cm within 100 ns, even if it never interacted, the propagation of low energy neutrons—and thus their contribution to





**Fig. 11.2** Comparison of the reconstructed energy in the ATLAS Tile calorimeter (TileCal) with several different physics models. The ratio of the simulated (MC) response with the response reconstructed from test beam runs (Data) for pions of energies from 20 GeV to 180 GeV. GEANT4 9.2 was used, comparing the predictions of different physics lists. The normalization uncertainty is 1%. (Courtesy of the Atlas Collaboration. Reproduction with permission)


the signal—is limited to short distances near the point of their generation during the typical trigger window of an LHC detector.

Comparing with the same Tilecal test beam, with the beam at 90 degrees with different physics lists of release 10.1, the version used for Run 2 simulation by ATLAS, demonstrates the change in physics performance from the revision of physics modeling in GEANT4. The energy resolution (Fig. 11.3), longitudinal shape (Fig. 11.4) and lateral shapes (Fig. 11.5) are compared with four physics lists that combine the QGS or Fritiof FTF string models with the Bertini or Binary cascade.

### 11.5.3 Background Estimation for CMS

Simulations were used to assess the required shielding for the CMS detector, to reduce the background from the interaction region p-p collisions [33] and the accelerator tunnel [114]. These employed a combination of tools: the STRUCT code was used for simulation the accelerator lattice and the scoring of particles lost at collimators, MARS to generate the particles entering the experimental area and FLUKA to model their interactions and fluxes in the detector and surrounding area. The study confirmed the need for shielding from the accelerator background and evaluated the proposed solutions. Key aspects were the impact on muon physics, together with the flexibility of optionally tracking the products of muon-

**Fig. 11.3** Energy resolution of ATLAS Tilecal test beam compared to recent GEANT4 10.1 release, currently used in ATLAS simulation production. Lower panel is the ratio of simulated (MC) and data. Courtesy of the ATLAS collaboration (ATLAS public plot). The original data and comparisons of mean and RMS energy [113] were undertaken with GEANT4 version 9.2



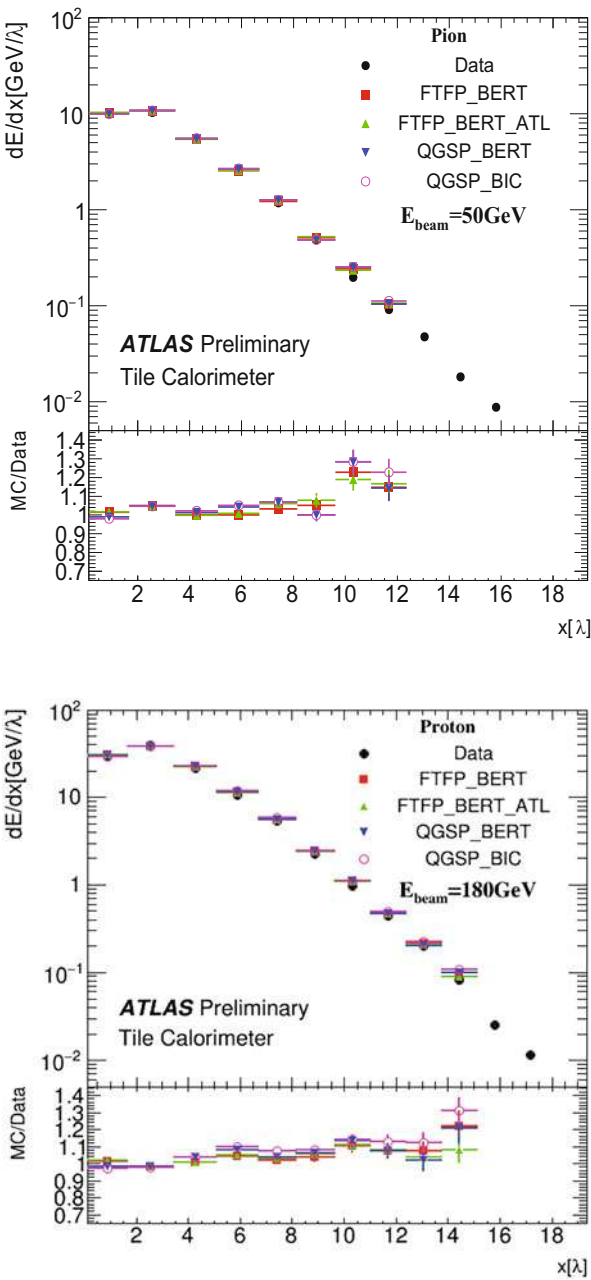
nucleus interactions. The study identified high-energy muons as the most important remaining background, affecting, in particular, the innermost barrel chambers, contributing to approximately 10% to the total trigger rate (which is very small).

#### 11.5.4 Validation from Comparisons with In-Situ Data


The accuracy of the simulation is determined using selected in-situ data from collisions, whenever possible. For example, ATLAS selected proton–proton collisions at 7 TeV (2010) and 8 TeV (2012) [115] to compare the energy deposition of charged hadrons with energies up to 30 GeV.

Decays of  $\Lambda$ , anti- $\Lambda$  and  $K_s^0$  were used to identify  $\pi^+$ ,  $\pi^-$ , protons and antiprotons and the ratio of their measured energy and momenta were compared with simulation using GEANT4 version 9.4. The ratio of  $\Lambda$  and anti- $\Lambda$  to  $K_s^0$  is 40% higher in data than in simulation, but the normalized distributions are well reproduced within statistical precision, as seen in Fig. 11.6. The tail beyond  $E/p > 1$  is due to the neutral background. The fraction of  $E/p \leq 0$  is due to interactions before the calorimeter; it is underestimated about 10% by the simulation across all particle species. The difference is taken between particle species of the mean values of  $E/p$  in order to reduce the effect of the neutral background. The difference between  $\pi^-$  and antiprotons, due to extra energy from the antiprotons annihilation, is described within uncertainties by the FTFP\_BERT physics list in GEANT4 version 9.4.

**Fig. 11.4** Longitudinal shower profile of 50 GeV pion and 180 GeV proton beams in ATLAS Tilecal, with the modules placed at 90 degrees in dedicated test beam setup, versus recent GEANT4 10.1. Lower panel is the ratio of simulated (MC) and data. Courtesy of the ATLAS collaboration (ATLAS public plot). The original data and earlier comparisons (vs. GEANT4 ver. 9.2) were in Ref. [113]



**Fig. 11.5** Lateral spread of 20, 50, 100 and 180 GeV pions incident on the ATLAS Tile Calorimeter at 90-degree angle. Black points represent data obtained in the period 2000–2003, and the colored points simulations using different physics lists of GEANT4 version 10.1. Lower panel is ratio of simulated (MC) and data. Courtesy of the ATLAS collaboration (ATLAS public plot). The original publication [113] compared data with the earlier version 9.2 of GEANT4



Inclusive spectra of isolated hadron tracks were used to compare  $E/p$  distributions with simulation. After the energy deposition of neutral particles is subtracted, a 5% discrepancy was found in the response to isolated charged hadrons between the modelling using two GEANT4 hadronic physics lists (FTFP\_BERT and QGSP\_BERT) in the central region of the calorimeter (Fig. 11.7).


When tracks that interact only in the hadron calorimeter were examined separately, the detector simulation was found to describe the response well. Tracks that interacted only in the EM calorimeter showed discrepancies 5–10% in  $E/p$ , which is consistent with being the origin of the difference of all tracks.

These comparisons are used in one of the methods of estimating the uncertainty of the jet energy scale. Compared with the most recent estimations of the jet energy scale, these estimates have larger uncertainties over most of the energy range, but confirm estimates from in situ beam data. However, they currently provide the only estimate for the largest momenta ( $p_T > 2$  TeV).

### 11.5.5 The Estimation of Jet Energy Scale in ATLAS and CMS

The earliest estimates of the jet energy scale, in the first years of the operation of ATLAS and CMS relied critically on detector simulation.

In CMS the detector simulation using GEANT4 was used in multiple stages of the initial calibration of the jet energy scale [116]. Initially, it was used to determine a base calibration factor  $C_{MC}(p_T^{\text{reco}})$  to account for the fraction of jet energy not

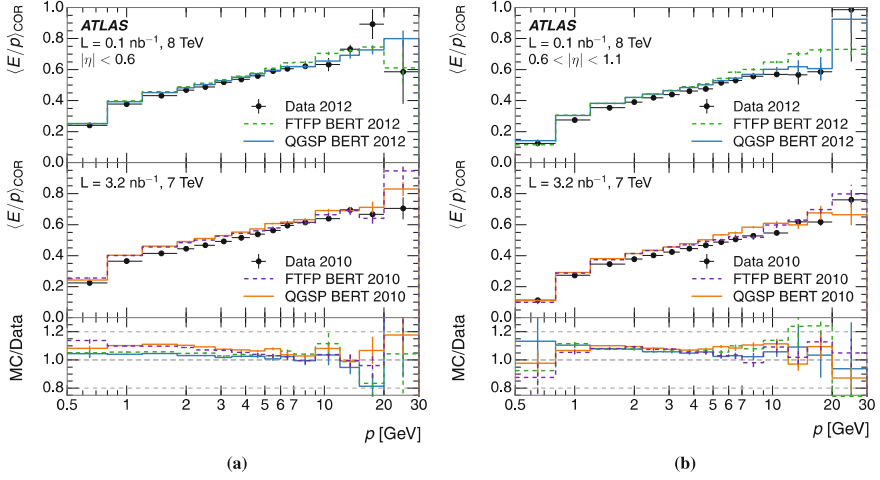


**Fig. 11.6** The  $E/p$  distribution for (a)  $\pi^+$ , (b)  $\pi^-$ , (c) protons and (d) antiprotons from selected ATLAS 8 TeV data of identified  $\Lambda$ , anti- $\Lambda$  and  $K_s$  decays [115] with  $|\eta| < 0.6$  and  $2.2 < p [\text{GeV}] < 2.8$ . The lower part of each panel shows the ratio of MC simulation (using GEANT4 ver. 9.4) to data. Reproduced from [115] under the Creative Commons License 4.0


observed due to the inactive parts, and to determine the variation of response for different types of particles.

Subsequently, the balance of the transverse momenta of  $\gamma$  + jets and  $Z$  + jets events was used to compare the well-measured electromagnetic response to the hadron/jet response predicted by simulation for different  $p_T$  and  $\eta$  values. CMS concluded that the Monte Carlo correction factor described the bulk of non-uniformity of  $C_{MC}$  in  $\eta$  and non-linearity in  $p_T$ . The estimates of the data/MC ratio of jet energy using the different samples (Fig. 11.8) were consistent, flat in  $p_T$  and its value was  $R_{\text{data/MC}} = 0.985 \pm 0.001$ .

The estimation of the jet energy scale with Run II data relies less on simulation. Comparisons versus simulation based on a newer GEANT4 version (10.2) found that a larger correction was required, partially ascribed to the migration to Fritiof-based hadronic models.



**Fig. 11.7** Comparison of the average ratio  $\langle E/p \rangle_{\text{COR}}$  of charged tracks of cluster energy in ATLAS calorimeters versus momentum  $p$  measured in the tracker. Comparisons within (a)  $|\eta| < 0.6$ , and (b)  $0.6 < |\eta| < 1.1$ , obtained after subtraction of corresponding estimates of neutral particle response, versus track momentum. Tracks with no matching energy cluster in the calorimeter are included. The bottom portion of the panels shows the ratio of simulation (using GEANT4 version 9.4 and two physics lists) to data. Error bars are statistical. Reproduced from [115] under the Creative Commons License 4.0



**Fig. 11.8** Correction factor from comparison of  $\gamma/W + \text{jet}$  events in CMS 2010 data at  $\sqrt{s} = 7 \text{ TeV}$ . From Ref. [116], reproduced under the Creative Commons license 3.0

### ***11.5.6 Fast Simulation in CMS During LHC Run I and II***

The fast simulations of the LHC experiments are sophisticated programs, combining the most important physics processes for electrons, gammas in the trackers and muons in the muon systems, with sampling from parameterized distributions for the showering of electrons, gammas and hadrons in the calorimeters. CMS's fast simulation FastSim [117] is a simplified geometrical description of the tracker, refined to obtain percent-level agreement for photon conversion, using around 30 thin nested cylinders.

The fast simulation in the tracker reconstructs each track using only its own generated hits, and cannot reproduce fake tracks that result from the incorrect association of hits. Though good agreement is seen in comparisons with Run-I data, the limitations of this approach are apparent in modeling the efficiency of track reconstruction and the fake rate [97].

Electron and hadron showers in calorimeters are turned into energy spot hits directly, distributed according to a  $\Gamma$ -function with parameters which fluctuate between showers using GLASH [118] or a similar approach.

Regular comparisons with the full detector simulation are used to monitor all the quantities used in physics analyses [119]. Agreement is observed at a level of 10%. Good agreement is particularly important for the missing transverse energy  $E_T$ .

During Run-I the CMS FastSim was used for the parameter scans for SuperSymmetry searches and samples of events used to evaluate systematic uncertainties [97], because the computing resources required for full simulation would have exceeded the available ones.

Once in-situ data is available, they are used as the final yardstick of the quality of both the fast and full simulation, and are used to address possible discrepancies between fast and detailed simulation.

### ***11.5.7 Future Detectors: Fine-Grained Calorimeters of CALICE***

Proposed detectors for next-generation collider experiments rely on particle flow reconstruction methods to obtain the required energy resolution for their physics programme. In order to obtain this performance, it is necessary to accurately model the energy deposition of charged hadrons, in order to subtract them from the observed signals.

Measurements with fine-grained calorimeters provide the most promising methods to validate the most important properties of detector simulation tools, and in particular their hadronic modeling. The CALICE experiment has undertaken test beam measurements of prototype calorimeters with many layers of scintillator tiles of fine granularity.

The prototype calorimeter with analog readout consisted of 38 layers, each containing a steel absorber plate and a scintillator layer. The  $30 \times 30 \text{ cm}^2$  core of the scintillator had granularity  $3 \times 3 \text{ cm}^2$ , and outer regions  $6 \times 6 \text{ cm}^2$  and  $12 \times 12 \text{ cm}^2$ . Data was collected with pions between 8 GeV and 100 GeV and compared with Monte Carlo simulation [120] using GEANT4 version 9.4.

The fine segmentation allows the estimation of the layer of the first hard interaction. This is used to obtain the shower profile relative to this starting layer. Averaging over showers starting in different layers reduces the effect of the variation of calibration, and is used to estimate uncertainties.

Comparisons of the mean longitudinal and lateral shower profiles for 8 GeV, 18 GeV and 80 GeV pions with GEANT4 physics lists including QGSP\_BERT and FTFP\_BERT were provided.

Figure 11.9 top row shows the longitudinal shower profile for pions of 8 GeV (left), 18 GeV (center) and 80 GeV (right) compared with GEANT4 physics lists FTFP\_BERT.

Normalization to unit total is used in each distribution. For FTFP\_BERT less energy is deposited in the early shower layers at all energies. At 80 GeV a difference is seen in the shower maximum of 10% (FTF versions) to 20% (QSGP versions), and the shower is more compact.

Similar comparisons for the radial shower profile show that all physics lists underestimate the radial extent of the showers and have a larger fraction of energy in the core. The effect is most pronounced at 80 GeV, see Fig. 11.10.

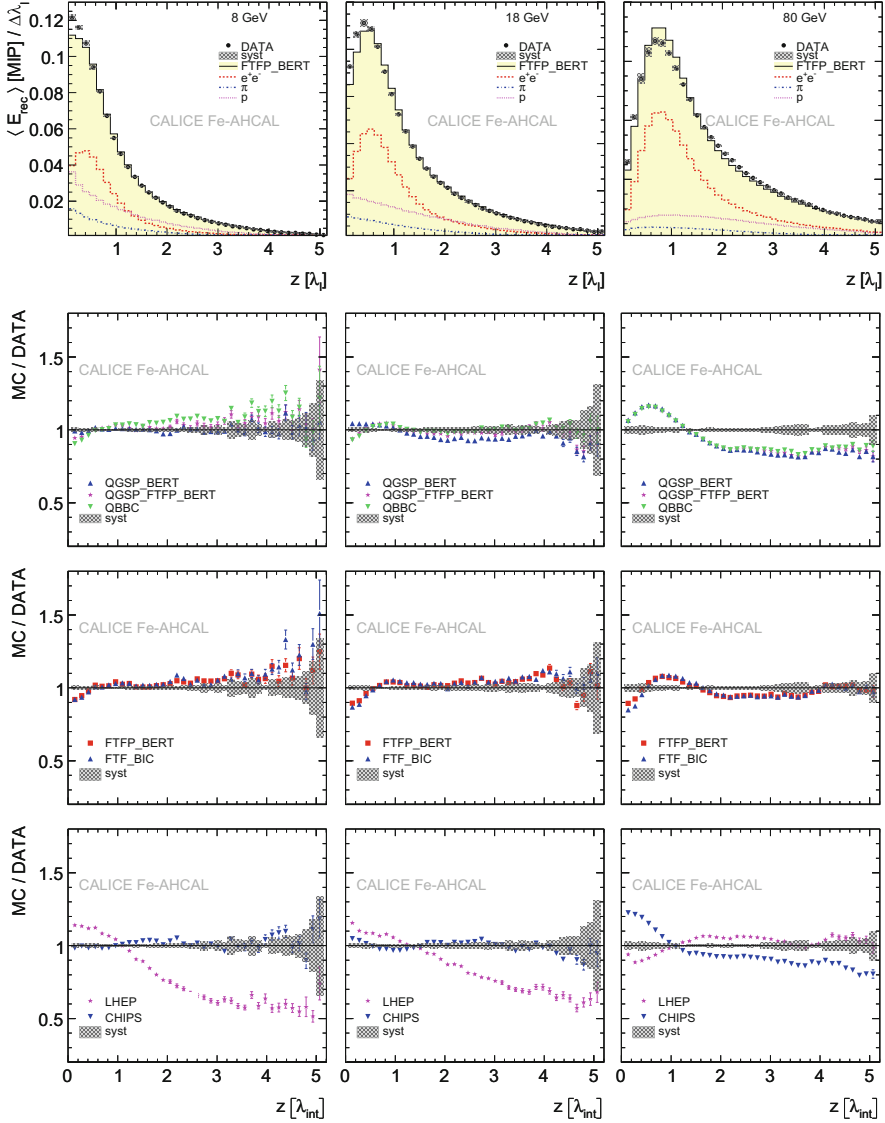
Either improvement of the relevant models or alternative physics models is needed to better describe these shower shapes, and provide the accuracy to use the full potential of simulation for future highly granular calorimeters.

## 11.6 Applications in Other Fields

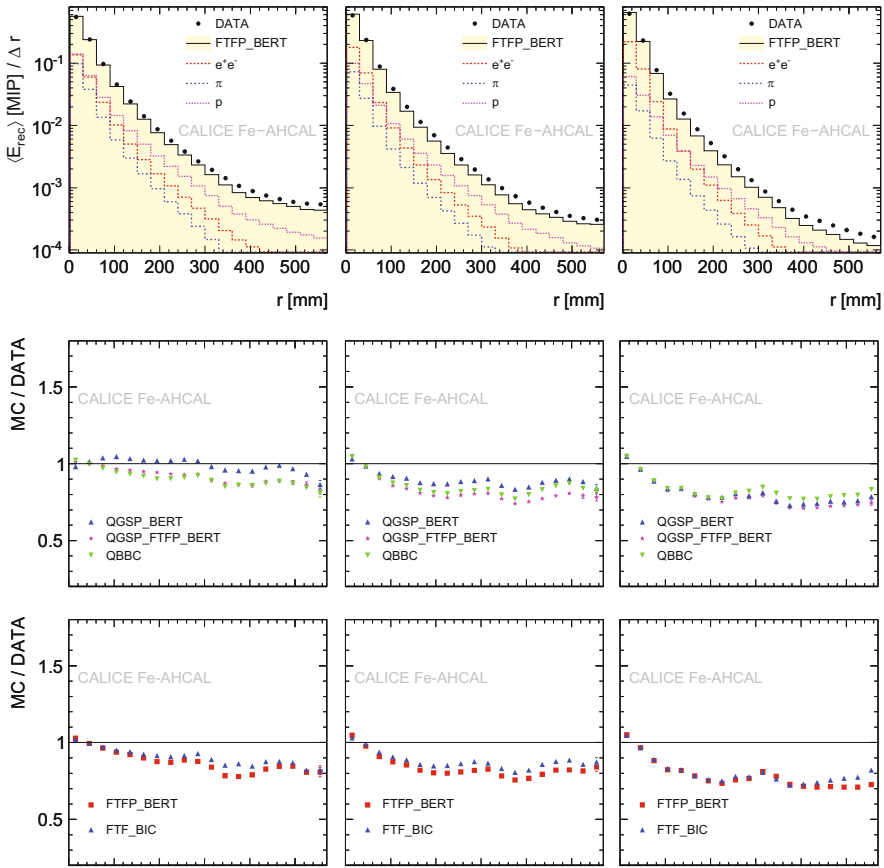
Particle transport simulation tools, including GEANT4 and FLUKA, have seen greatly increasing usage beyond High Energy and Nuclear Physics (HENP) experiments in the past decade.

In particular in medical physics, their application has seen spectacular growth, and has spanned several domains, especially the development and refinement of new methods and assessment of treatments in radiotherapy, and the simulation of medical imaging detectors.





**Fig. 11.9** Mean longitudinal shower profile, starting at the layer of the first interaction, for pion beams in CALICE iron-scintillator analog hadronic calorimeter. Pion energies are 8 GeV (left), 18 GeV (center) Top row: data (circles) compared with FTFP\_BERT physics list of GEANT4 version 9.4. The parts deposited by different particles (electrons/positrons, pions and protons) in the simulation are shown. Lower 3 rows: ratios between selected physics lists and data. (We note that the CHIPS and LHEP physics lists were withdrawn in Geant4 release 10.0). Reproduced from [120], under Creative Commons License 3.0



**Fig. 11.10** Mean radial shower profile for pion beams in CALICE iron-scintillator analog hadronic calorimeter. Pion energies are 8 GeV (left), 18 GeV (center) and 80 GeV (right). Top row: data (circles) compared with FTFP\_BERT physics list of GEANT4 version 9.4. Deposits by different particles (electrons/positrons, pions and protons) in the simulation are shown. Centre and lower row: ratios between selected physics lists and data, including QGSP\_BERT and FTFP\_BERT and variants. Reproduced from [120], under Creative Commons License 3.0. (Physics lists using models withdrawn in subsequent releases, CHIPS and LHEP, are omitted)

The simulation tools are used widely also in determining the effect of radiation on satellites and spacecraft from planetary radiation environments and the solar and galactic rays. Specialized tools have been developed for shielding studies [121] as have general purpose tools to evaluate effects of the space environments [122].

FLUKA is used for shielding and target design of accelerators, activation studies, and also for cosmic ray studies, due to its ability to simulate up to 20 PeV. A further application is the assessment of dose to aircrews flying in commercial aircraft.

### 11.6.1 *Medical Imaging*

A key application of Monte Carlo simulation in medical imaging is the development of novel instrumentation, e.g. progress in Position Emission Tomography (PET) and Single Photon Emission Computed Tomography (SPECT). Particle transport simulation enables an evaluation of new materials, geometries or system configurations in multiple versions without the expense of always creating a hardware prototype [123].

Dedicated particle transport tools have been developed specifically for the simulation of PET or SPECT devices. An early tool PET-EGS [124] used EGS, with GEANT4 used in the leading tool, GATE [125], and in GAMOS [126], and Penelope in PeneloPET [127].

GATE is one of the most commonly used dedicated simulation tools for PET [123]. This is due to its simplicity in generating setups and steering the simulation using text commands, and due to the benefits of the validation of the GEANT4 toolkit.

Particle transport simulation is a standard tool in many applications. However, one of its key drawbacks is the large computation times required. Methods have been therefore developed to mitigate this problem. As an example, networked computers are being employed to speed up the calculations. Alternatively, hybrid computational models are being employed, such as generating the initial photons using SIMSET [128] or using EGSnrc as its core simulation engine [129].

### 11.6.2 *Proton and Hadron Therapy*

As in photon radiotherapy, the recent advances of proton and ion beam therapy have heavily relied on radiation transport Monte Carlo tools [130]. Due to the need for short computation time, specialized Treatment Planning Systems (TPS) with analytical or simplified models for the fast estimation of dose delivery are the clinical standard.

GEANT4 validation for proton-therapy involves selection of the best performing physics models [131]. These tools are also essential in evaluating potential improvements in TPS methods [132].

Another critical aspect is the simulation of the effects of organ motion on dose delivery, e.g. with the GEANT4-based GATE and standalone applications [133].

Specialized applications were developed to use GEANT4 in particle therapy and provide tailored and validated physics configurations, interfaces to CT input, and tools including the reading and writing of snapshots of particles at specific interfaces as phase space files. In the past decade two dedicated applications, PTSim [134] and TOPAS [135] targeted easy use by clinical physicists in Japan and the US respectively. TOPAS emphasized reliable configuration, and modelling of the motion

of components. Both have seen increasing use for research and in selected clinical settings.

The therapy potential of light ion beams has been the topic of increasing research during the past decade. A key drawback is the energy deposition of nuclear fragments which extends beyond the Bragg peak. Early comparisons identified discrepancies of some tens of percent in non-differential quantities between data and MC [136].

In more recent studies better agreement in these tail dose depositions and dose profiles was obtained with GEANT4 and FLUKA, but differences in prompt gamma emission continue to be an issue [137]. The lateral beam widening is also well confirmed by FLUKA results [138].

A key need in ion therapy is the monitoring of the range of ions. The detection of positron emission in a PET detector has been in clinical use, and a newer method involves prompt gamma emission. Both techniques have been investigated using particle transport to quantify the location of emission and the spectra of clinically interesting gammas.

Modeling of ion–ion interactions at therapeutic energies is frequently at the edge of applicability for cascade models (below 150–200 MeV), and the resulting spectra are influenced by the details of many nuclear de-excitation processes. Discrepancies in secondary particle production in FLUKA were improved with the addition of a Boltzman Master Equation (BME) model and other modelling refinements [139]. New measurements have been made with ion projectiles to provide data for comparison and improvement of modelling. One set using lower energy (62 A MeV)  $^{12}\text{C}$  beam measured a large set of secondary spectra (p, d, t through to  $^{11}\text{B}$ ) [140]. Comparison with GEANT4 models revealed the need for improvement of the modeling used (binary cascade and QMD).

FLUKA's existing applications in particle therapy [141] include the production of data for Treatment Planning Systems (TPS), checking the plans created by TPS in selected cases for quality control and improvement of patient dose delivery, and in feasibility and sensitivity studies of prompt gammas for range and dose monitoring. Another use has been the monitoring of dose delivery in ion therapy through PET imaging of positron emitter production, undertaken either after treatment or through an integrated PET device during the patient treatment.

### ***11.6.3 Developments for Microdosimetry and Nanodosimetry***

New models and adaptation of physics models have been developed to extend their application to smaller energies. Dedicated Monte Carlo track structure codes have been used in the investigation of radiation effects at the micron scale and at scales appropriate for biological research [142] and modeling of radiotherapy outcomes [143]. GEANT4 has been extended to provide track structure modeling in liquid and gaseous water with the development of the GEANT4-DNA package [144]. This has enabled its use in many applications in these fields [145, 146].

The GEANT4-DNA package provides new physical models for the description of elastic and inelastic electromagnetic interactions of electrons and select ions (Li to O, plus Si and Fe) in liquid water, previously only available in dedicated ‘track structure’ codes such as PARTRAC [147]. In addition, GEANT4-DNA offers features to model the water radiolysis from ionizing radiation: ionized or excited water molecules and water anions are generated and tracked at a physicochemical stage up to a few picoseconds, and subsequently in a ‘chemical’ stage using models of generation, diffusion and reaction of new chemical species. This is part of an effort to model and understand the first stages of DNA damage.

## 11.7 Outlook

During the past decade the application of detector simulation tools has been significantly widened through the implementation of improved physics models. Code and models have become more accurate in describing benchmark data. The need for more accurate data for comparison and model improvement has been one motivation for some thin-target (HARP, MIPP) and thick target (CALICE) experiments.

The requirements that arise from the projected use of simulations as an integral part of the next generation detectors becomes ever stronger. Witness, e.g. the use of particle flow reconstruction [148] in proposed experiments at the energy frontier, including the Linear Collider and the Future Circular Collider, to address one of their major challenges.

The need for further development of physics models for high-energy hadron–nucleus interactions is evident. Several promising approaches are being pursued, including the extension and tuning of existing implementations of current models (Fritiof, Quark-Gluon String), the incorporation of alternative implementations of existing models, such as QGSJET, complementary modeling approaches (DPMJET) and the incorporation of new models (EPOS). The availability of high quality thin target experimental data at energies over a range of momenta from 20 to 158 GeV/c [149] is an important resource; lack of data for higher energies is a constraint.

In addition to physics improvements, the large increase in statistics of simulated events for the HL-LHC requires a large improvement in CPU performance, of approximately a factor of ten. Research in CPU-performance and emerging architectures in the GeantV R&D effort indicate a more realistic target of a factor of 2–4 in performance improvement may be within reach for detailed simulation.

These prospects strengthen the need for parameterized (fast) simulation methods which can reproduce the results of detailed simulation as accurately as possible, for use in a majority and potentially an ever-larger fraction of analyses. Hybrid methods combining parameterized and detailed simulation in innovative ways and the machine learning approach to parameterized/fast simulation appear amongst the options which will see significant development and research in the next years.

## Bibliography

1. Graf, N., McCormick, J.: Simulator For The Linear Collider (SLIC): A Tool For ILC Detector Simulations. In: AIP Conference Proceedings. pp. 503–512. AIP (2006)
2. Frank, M., Gaede, F., Nikiforou, N. et al.: DDG4 A Simulation Framework based on the DD4hep Detector Description Toolkit. J. Phys. Conf. Ser. 664, 072017 (2015). <https://doi.org/10.1088/1742-6596/664/7/072017>
3. Behnke, T., Brau, J.E., Burrows, P.N. et al.: The International Linear Collider Technical Design Report - Volume 4: Detectors. Batavia, IL (United States) (2013)
4. Daniel Elvira, V.: Impact of detector simulation in particle physics collider experiments. Phys. Rep. 695, 1–54 (2017). <https://doi.org/10.1016/J.PHYSREP.2017.06.002>
5. Albrecht, J., Alves, A.A., Amadio, G. et al.: A Roadmap for HEP Software and Computing R&D for the 2020s. (2017). <https://doi.org/10.1007/s41781-018-0018-8>
6. Alves, A.A., Amadio, G., Anh-Ky, N. et al.: A Roadmap for HEP Software and Computing R&D for the 2020s. (2017)
7. Amadio, G., Apostolakis, J., Bandieramonte, M. et al.: The GeantV project: Preparing the future of simulation. J. Phys. Conf. Ser. 664, (2015). <https://doi.org/10.1088/1742-6596/664/7/072006>
8. Sjöstrand, T., Ask, S., Christiansen, J.R. et al.: An introduction to PYTHIA 8.2. Comput. Phys. Commun. 191, 159–177 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024>
9. Andersson, B., Gustafson, G., Pi, H.: The FRITIOF model for very high energy hadronic collisions. Zeitschrift fr Phys. C Part. Fields. 57, 485–494 (1993). <https://doi.org/10.1007/BF01474343>
10. Sjöstrand, T., Bengtsson, M.: The Lund Monte Carlo for jet fragmentation and e+ e- physics - jetset version 6.3 - an update. Comput. Phys. Commun. 43, 367–379 (1987). [https://doi.org/10.1016/0010-4655\(87\)90054-3](https://doi.org/10.1016/0010-4655(87)90054-3)
11. Corcella, G., Knowles, I.G., Marchesini, G. et al.: HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). J. High Energy Phys. 2001, 10 (2001)
12. Bellm, J., Gieseke, S., Grellscheid, D. et al.: Herwig 7.0/Herwig++ 3.0 release note. Eur. Phys. J. C. 76, 196 (2016). <https://doi.org/10.1140/epjc/s10052-016-4018-8>
13. Agostinelli, S., Allison, J., Amako, K. et al.: Geant4—a simulation toolkit. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. 506, 250–303 (2003). [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)
14. Nakao, N., Mokhov, N.V.: MARS15 code in accelerator applications. <http://www-ap.fnl.gov/users/mokhov/papers/2007/Conf-07-416-APC.pdf> (2007)
15. Brun, R., Giani, S.: GEANT—Detector description and simulation tool. (1994)
16. Fesefeldt, H.S.: Simulation of hadronic showers, physics and applications. Physikalisches Institut, RWTH Aachen Physikzentrum, 5100 Aachen, Germany (1985)
17. Zeitnitz, C., Gabriel, T.A.: The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters. Nucl. Instr. Methods A. 349, 106–111 (1994). [https://doi.org/10.1016/0168-9002\(94\)90613-0](https://doi.org/10.1016/0168-9002(94)90613-0)
18. Gabriel, T.A., Bishop, B.L., Brau, J.E.: The physics of compensating calorimetry and the new calor89 code system. IEEE Trans. Nucl. Sci. 36, 14–22 (1989). <https://doi.org/10.1109/23.34394>
19. Fassò, A., Ferrari, A., Ranft, J., et al.: FLUKA: present status and future developments. In: Menzione, A. and Scribano, A.P.G. 493 (eds.) Proc. IV Int. Conf. on Calorimetry in High Energy Physics, La Biodola, Italy, 21–26 Sept. 1993. p. World Scientific. World Scientific (1993)
20. Battistoni, G., Boehlen, T., Cerutti, F. et al.: Overview of the FLUKA code. Ann. Nucl. Energy. 82, 10–18 (2015). <https://doi.org/10.1016/J.ANUCENE.2014.11.007>

21. Roesler, S., Engel, R., Ranft, J.: The Monte Carlo Event Generator DPMJET-III. In: A. Kling M. Nakagawa L. Távora & P. Vaz PG - 1033, F.B. (ed.) *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications*. p. 1038 (2001)
22. Allison, J., Amako, K., Apostolakis, J. et al.: Recent developments in Geant4. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 835, 186–225 (2016). <https://doi.org/10.1016/j.nima.2016.06.125>
23. Sihver, L., Sato, T., Gustafsson, K. et al.: Iwase, H., Niita, K., Nakashima, H., Sakamoto, Y., Iwamoto, Y., Matsuda, N.: An update about recent developments of the PHITS code. *Adv. Sp. Res.* 45, 892–899 (2010). <https://doi.org/10.1016/j.asr.2010.01.002>
24. Mashnik, S.G.: Validation and Verification of MCNP6 Against Intermediate and High-Energy Experimental Data and Results by Other Codes. *Eur. Phys. J. Plus.* (2011). <https://doi.org/10.1140/epjp/i2011-11049-1>
25. Goorley, T., James, M., Booth, T. et al.: Features of MCNP6. *Ann. Nucl. Energy.* 87, 772–783 (2016). <https://doi.org/10.1016/J.ANUCENE.2015.02.020>
26. Veenhof, R.: Garfield, recent developments. *Nucl. Instruments Methods Phys. Res. A.* 419, 726–730 (1998). [https://doi.org/10.1016/S0168-9002\(98\)00851-1](https://doi.org/10.1016/S0168-9002(98)00851-1)
27. Smirnov, I.B.: Modeling of ionization produced by fast charged particles in gases. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 554, 474–493 (2005). <https://doi.org/10.1016/j.nima.2005.08.064>
28. Veenhoff, R., Schindler, H.: Garfield++ – simulation of ionisation based tracking detectors, <http://garfieldpp.web.cern.ch/garfieldpp/>
29. Pfeiffer, D., De Keukeleere, L., Azevedo, C. et al.: A Geant4/Garfield++ and Geant4/Degrad Interface for the Simulation of Gaseous Detectors. (2018)
30. van der Ende, B.M., Rand, E.T., Erlandson, A. et al.: Use of SRIM and Garfield with Geant4 for the characterization of a hybrid 10B/3He neutron detector. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 894, 138–144 (2018). <https://doi.org/10.1016/J.NIMA.2018.03.056>
31. Lukas, W.: Fast Simulation for ATLAS: Atfast-II and ISF. *J. Phys. Conf. Ser.* 396, 022031 (2012). <https://doi.org/10.1088/1742-6596/396/2/022031>
32. Mokhov, N.V., Gudima, K.K., James, C.C. et al.: Recent enhancements to the MARS15 code. *Radiat Prot Dosim.* 116, 99–103 (2005)
33. Huhtinen, M., Aarnio, P.A.: Neutron and photon fluxes and shielding alternatives for the CMS detector at LHC. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 363, 545–556 (1995). [https://doi.org/10.1016/0168-9002\(95\)00444-0](https://doi.org/10.1016/0168-9002(95)00444-0)
34. Hrivnacova, I., Adamova, D., Berejnoi, V. et al.: The Virtual Monte Carlo. In: *Computing in High Energy and Nuclear Physics (CHEP03)*, La Jolla, CA. p. arXiv:cs/0306005 (2003)
35. Kalos, M.H., Whitlock, P.A.: *Monte Carlo methods*. Wiley - VCH (2008)
36. James, F.: Monte Carlo theory and practice. *Reports Prog. Phys.* 43, 1145–1189 (1980). <https://doi.org/10.1088/0034-4885/43/9/002>
37. Bielajew, A.F.: *Fundamentals of the Monte Carlo method for neutral and charged particle transport*. (2000)
38. Panneton, F., L'Ecuyer, P.: Resolution-stationary random number generators. *Math. Comput. Simul.* 80, 1096–1103 (2010). <https://doi.org/10.1016/j.matcom.2007.09.014>
39. Savvidy, G., Ter-Arutyunyan-Savvidy, N.: On the Monte Carlo simulation of physical systems. *J. Comput. Phys.* 97, 566–572 (1991). [https://doi.org/10.1016/0021-9991\(91\)90015-D](https://doi.org/10.1016/0021-9991(91)90015-D)
40. Savvidy, K., Savvidy, G.: Spectrum and entropy of C-systems MIXMAX random number generator. *Chaos, Solitons & Fractals.* 91, 33–38 (2016). <https://doi.org/10.1016/j.chaos.2016.05.003>
41. Lüscher, M.: A portable high-quality random number generator for lattice field theory simulations. *Comput. Phys. Commun.* 79, 100–110 (1994). [https://doi.org/10.1016/0010-4655\(94\)90232-1](https://doi.org/10.1016/0010-4655(94)90232-1)
42. James, F.: RANLUX: A Fortran implementation of the high-quality pseudorandom number generator of Lüscher. *Comput. Phys. Commun.* 79, 111–114 (1994). [https://doi.org/10.1016/0010-4655\(94\)90233-X](https://doi.org/10.1016/0010-4655(94)90233-X)

43. L'Ecuyer, P., Simard, R.: TestU01: A C library for empirical testing of random number generators. *ACM Trans. Math. Softw.* 33, 22–es (2007). <https://doi.org/10.1145/1268776.1268777>
44. Berger, M.J.: Monte Carlo calculation of the penetration and diffusion of fast charged particles. In: B. Alder S. Fernbach and Rotenberg, M. (eds.) *Methods in Computational Physics: Advances in Research and Applications*, Vol. 1. Statistical Physics. pp. 135–215. Academic, New York (1963)
45. Ford, R.L., Nelson, W.R.: *The EGS Code System: Computer Programs for the Monte Carlo Simulation of Electromagnetic Cascade Showers (Version 3)*. (1978)
46. Nelson, W.R., Hirayama, H., Rogers, D.W.O.: *The EGS4 code system*. (1985)
47. Salvat, F., Fernández-Varea, J.M.: Overview of physical interaction models for photon and electron transport used in Monte Carlo codes. *Metrologia*. 46, S112–S138 (2009). <https://doi.org/10.1088/0026-1394/46/2/S08>
48. Baró, J., Sempau, J., Salvat, F. et al.: PENELOPE: An algorithm for Monte Carlo simulation of the penetration and energy loss of electrons and positrons in matter. *Nucl. Instruments Methods Phys. Res. B* 100, 31–46 (1995). [https://doi.org/10.1016/0168-583X\(95\)00349-5](https://doi.org/10.1016/0168-583X(95)00349-5)
49. Goudsmit, S., Saunderson, J.L.: Multiple Scattering of Electrons. II. *Phys. Rev.* 58, 36–42 (1940). <https://doi.org/10.1103/PhysRev.58.36>
50. Molière, G.: Theorie der Streuung schneller geladener Teilchen I. Einzelstreuung am abgeschirmten Coulomb-Feld. *Zeitschrift Naturforsch. Tl. A*. 2, 133–+ (1947)
51. Lewis, H.W.: Multiple Scattering in an Infinite Medium. *Phys. Rev.* 78, 526–529 (1950). <https://doi.org/10.1103/PhysRev.78.526>
52. Bielajew, A.F., Rogers, D.W.O.: Presta: The parameter reduced electron-step transport algorithm for electron monte carlo transport. *Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms*. 18, 165–171, 174–181 (1986). [https://doi.org/10.1016/S0168-583X\(86\)80027-1](https://doi.org/10.1016/S0168-583X(86)80027-1)
53. Kawrakow, I., Bielajew, A.F.: On the condensed history technique for electron transport. *Nucl. Instruments Methods Phys. Res. B*. 142, 253–280 (1998). [https://doi.org/10.1016/S0168-583X\(98\)00274-2](https://doi.org/10.1016/S0168-583X(98)00274-2)
54. Bielajew, A.F., Salvat, F.: Improved electron transport mechanics in the PENELOPE Monte-Carlo model. *Nucl. Instruments Methods Phys. Res. B*. 173, 332–343 (2001). [https://doi.org/10.1016/S0168-583X\(00\)00363-3](https://doi.org/10.1016/S0168-583X(00)00363-3)
55. Salvat, F., Fernández-Varea, J.M., Sempau, J.: “PENELOPE, A Code System for Monte Carlo Simulation of Electron and Photon Transport.”, Barcelona (2009)
56. Kawrakow, I.: Accurate condensed history Monte Carlo simulation of electron transport. I. EGSnrc, the new EGS4 version. *Med. Phys.* 27, 485–498 (2000)
57. Hirayama, H., Namito, Y., Bielajew, A.F. et al.: *The EGS5 code system*. (2005)
58. Vilches, M., García-Pareja, S., Guerrero, R. et al.: Monte Carlo simulation of the electron transport through thin slabs: A comparative study of penelope, geant3, geant4, egsrc and mcnp. *Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms*. 254, 219–230 (2007). <https://doi.org/10.1016/j.nimb.2006.11.061>
59. Incerti, S., Ivanchenko, V., Novak, M.: Recent progress of Geant4 electromagnetic physics for calorimeter simulation. *J. Instrum.* 13, C02054–C02054 (2018). <https://doi.org/10.1088/1748-0221/13/02/C02054>
60. Grichine, V.M., Sadilov, S.S.: Geant4 models for X-ray transition radiation. *Nucl. Instruments Methods Phys. Res. A*. 522, 122–125 (2004). <https://doi.org/10.1016/j.nima.2004.01.031>
61. Armstrong, T.W., Chandler, K.G.: HETC - a high energy transport code. *Nucl. Sci. Eng.* 49, 110–111 (1972)
62. Dementyev, A. V., Sobolevsky, N. M.: SHIELD - Universal Monte Carlo Hadron Transport Code: Scope and Applications. *Radiat. Meas.* 50, 553–557 (1999). [https://doi.org/10.1016/S1350-4487\(99\)00231-0](https://doi.org/10.1016/S1350-4487(99)00231-0)
63. Weisskopf, V.F., Ewing, D.H.: On the Yield of Nuclear Reactions with Heavy Elements. *Phys. Rev.* 57, 472–485 (1940). <https://doi.org/10.1103/PhysRev.57.472>



64. Botvina, A.S., Iljinov, A.S., Mishustin, I.N. et al.: Statistical simulation of the break-up of highly excited nuclei. *Nucl. Phys. A.* 475, 663–686 (1987). [https://doi.org/10.1016/0375-9474\(87\)90232-6](https://doi.org/10.1016/0375-9474(87)90232-6)
65. Furihata, S.: The GEM Code - the Generalized Evaporation Model and the Fission Model. In: A. Kling, Barão, F., Nakagawa, M., Távora, L., and P. Vaz (eds.) *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications*. p. 1045–+ (2001)
66. Gudima, K.K., Mashnik, S.G., Toneev, V.D.: Cascade-exciton model of nuclear reactions. *Nucl. Phys. A.* 401, 329–361 (1983). [https://doi.org/10.1016/0375-9474\(83\)90532-8](https://doi.org/10.1016/0375-9474(83)90532-8)
67. Mashnik, S.G., Gudima, K.K., Prael, R.E. et al.: CEM03.03 and LAQGSM03.03 Event Generators for the MCNP6, MCNPX, and MARS15 Transport Codes. (2008). <https://doi.org/10.1016/j.nimb.2010.09.005>
68. Bertini, H.W.: Intranuclear-Cascade Calculation of the Secondary Nucleon Spectra from Nucleon-Nucleus Interactions in the Energy Range 340 to 2900 MeV and Comparisons with Experiment. *Phys. Rev.* 188, 1711–1730 (1969). <https://doi.org/10.1103/PhysRev.188.1711>
69. Bleicher, M., Zabrodin, E., Spieles, C. et al.: Relativistic hadron-hadron collisions in the ultra-relativistic quantum molecular dynamics model. *J. Phys. G Nucl. Part. Phys.* 25, 1859–1896 (1999)
70. Folger, G., Ivanchenko, V.N., Wellisch, J.P.: The Binary Cascade. *Eur. Phys. J. A.* 21, 407–417 (2004). <https://doi.org/10.1140/epja/i2003-10219-7>
71. Heikkinen, A., Stepanov, N., Wellisch, J.P.: Bertini intra-nuclear cascade implementation in Geant4. In: 13th Intern. Computing in High Energy and Nuclear Physics, (CHEP 2003): La Jolla, California, March 24–28, 2003. p. arXiv:nucl-th/0306008 (2003)
72. Duarte, H.: Particle production in nucleon induced reactions above 14 MeV with an intranuclear cascade model. *Phys. Rev. C.* 75, 24611 (2007). <https://doi.org/10.1103/PhysRevC.75.024611>
73. Degtyarenko, P.V., Kossov, M.V., Wellisch, H.-P.: Chiral invariant phase space event generator. *Eur. Phys. J. A.* 8, 217–222 (2000). <https://doi.org/10.1007/s100500070108>
74. Wright, D.H., Kelsey, M.H.: The Geant4 Bertini Cascade. *Nucl. Instrum. Methods A.* 804, 175–188 (2015). <https://doi.org/10.1016/j.nima.2015.09.058>
75. Boudard, A., Cugnon, J., Leray, S. et al.: Intranuclear cascade model for a comprehensive description of spallation reaction data. *Phys. Rev. C.* 66, 044615 (2002). <https://doi.org/10.1103/PhysRevC.66.044615>
76. Boudard, A., Cugnon, J., David, J.-C. et al.: New potentialities of the Liège intranuclear cascade model for reactions induced by nucleons and light charged particles. *Phys. Rev. C.* 87, 014606 (2013). <https://doi.org/10.1103/PhysRevC.87.014606>
77. Mancusi, D., Boudard, A., Cugnon, J. et al.: Extension of the Liège intranuclear-cascade model to reactions induced by light nuclei. *Phys. Rev. C.* 90, 054602 (2014). <https://doi.org/10.1103/PhysRevC.90.054602>
78. Napolitani, P., Schmidt, K.-H., Botvina, A.S. et al.: High-resolution velocity measurements on fully identified light nuclides produced in Fe 56 + hydrogen and Fe 56 + titanium systems. *Phys. Rev. C.* 70, 054607 (2004). <https://doi.org/10.1103/PhysRevC.70.054607>
79. Ricciardi, M. V., Armbruster, P., Benlliure, J. et al.: Light nuclides produced in the proton-induced spallation of U 238 at 1 GeV. *Phys. Rev. C.* 73, 014607 (2006). <https://doi.org/10.1103/PhysRevC.73.014607>
80. Leray, S., David, J.C., Khandaker, M. et al.: Results from the IAEA Benchmark of Spallation Models. *J. Korean Phys. Soc.* 59, 791 (2011). <https://doi.org/10.3938/jkps.59.791>
81. Akchurin, N., Bedeschi, F., Cardini, A. et al.: Lessons from Monte Carlo simulations of the performance of a dual-readout fiber calorimeter. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 762, 100–118 (2014). <https://doi.org/10.1016/J.NIMA.2014.05.121>
82. Capella, A., Sukhatme, U., Tan, C.-I. et al.: Dual parton model. *Phys. Rep.* 236, 225–329 (1994). [https://doi.org/10.1016/0370-1573\(94\)90064-7](https://doi.org/10.1016/0370-1573(94)90064-7)

83. Kaidalov, A.B.: Interactions of hadrons and nuclei at superhigh energies and small-x physics. Nucl. Phys. B Proc. Suppl. 75, 81–88 (1999). [https://doi.org/10.1016/S0920-5632\(99\)00218-2](https://doi.org/10.1016/S0920-5632(99)00218-2)
84. Ostapchenko, S.: QGSJET-II: towards reliable description of very high energy hadronic interactions. Nucl. Phys. B Proc. Suppl. 151, 143–146 (2006). <https://doi.org/10.1016/j.nuclphysbps.2005.07.026>
85. Uzhinsky, V., Apostolakis, J., Galoyan, A. et al.: Antinucleus and nucleus cross sections implemented in Geant4. Phys. Lett. B. 705, 235–239 (2011). <https://doi.org/10.1016/j.physletb.2011.10.010>
86. Uzhinsky, V., Galoyan, A.: Effect of  $u\bar{u}$  diquark suppression in proton splitting in Monte Carlo event generators. Phys. Rev. D. 91, 037501 (2015). <https://doi.org/10.1103/PhysRevD.91.037501>
87. Koning, A.J. et al.: The JEFF evaluated nuclear data project. In: O. Bersillon F. Gunsing, E.B.R.J. and S. Leray (eds.) International Conference on Nuclear Data for Science and Technology 2007 (April 22–27, 2007, Nice, France). pp. 194–199. EDP Sciences (2008)
88. Chadwick, M.B., Obložinský, P., Herman, M. et al.: ENDF/B-VII.0: Next Generation Evaluated Nuclear Data Library for Nuclear Science and Technology. Nucl. Data Sheets. 107, 2931–3060 (2006). <https://doi.org/10.1016/j.nds.2006.11.001>
89. Brown, D.A., Chadwick, M.B., Capote, R. et al.: ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data. Nucl. Data Sheets. 148, 1–142 (2018). <https://doi.org/10.1016/j.nds.2018.02.001>
90. Shibata, K., Iwamoto, O., Nakagawa, T. et al.: JENDL-4.0: A New Library for Nuclear Science and Engineering. J. Nucl. Sci. Technol. 48, 1–30 (2011). <https://doi.org/10.1080/18811248.2011.9711675>
91. Ge, Z.G., Zhao, Z.X., Xia, H.H. et al.: The Updated Version of Chinese Evaluated Nuclear Data Library (CENDL-3.1). J. Korean Phys. Soc. 59, 1052–1056 (2011). <https://doi.org/10.3938/jkps.59.1052>
92. Kossov, M.V.: Chiral-invariant phase space model. Eur. Phys. J. A - Hadron. Nucl. 14, 265–269 (2002). <https://doi.org/10.1140/epja/i2001-10211-3>
93. Gabriel, T.A., Groom, D.E., Job, P.K. et al.: Energy dependence of hadronic activity. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. 338, 336–347 (1994). [https://doi.org/10.1016/0168-9002\(94\)91317-X](https://doi.org/10.1016/0168-9002(94)91317-X)
94. Wigmans, R.: Calorimetry : Energy measurement in particle physics. Clarendon Press (2000)
95. Groom, D.E.: Energy flow in a hadronic cascade: Application to hadron calorimetry. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. 572, 633–653 (2007). <https://doi.org/10.1016/j.nima.2006.11.070>
96. de Favereau, J., Delaere, C., Demin, P. et al.: DELPHES 3: a modular framework for fast simulation of a generic collider experiment. J. High Energy Phys. 57 (2014). [https://doi.org/10.1007/JHEP02\(2014\)057](https://doi.org/10.1007/JHEP02(2014)057)
97. Giammanco, A.: The Fast Simulation of the CMS Experiment. J. Phys. Conf. Ser. 513, 022012 (2014). <https://doi.org/10.1088/1742-6596/53/2/022012>
98. Ritsch, E., Collaboration, the A.: Concepts and Plans towards fast large scale Monte Carlo production for the ATLAS Experiment. J. Phys. Conf. Ser. 523, 012035 (2014). <https://doi.org/10.1088/1742-6596/523/1/012035>
99. Paganini, M., de Oliveira, L., Nachman, B.: CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks. Phys. Rev. D. 97, 014021 (2018). <https://doi.org/10.1103/PhysRevD.97.014021>
100. Wigmans, R.: Toward Meaningful Simulations of Hadronic Showers. In: M. Albrow & R. Raja (ed.) Hadronic Shower Simulation Workshop. pp. 123–136 (2007)
101. Smirnov, I.B.: Modeling of ionization produced by fast charged particles in gases. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. 554, 474–493 (2005). <https://doi.org/10.1016/j.nima.2005.08.064>

102. Biagi, S.F.: Monte Carlo simulation of electron drift and diffusion in counting gases under the influence of electric and magnetic fields. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 421, 234–240 (1999). [https://doi.org/10.1016/S0168-9002\(98\)01233-9](https://doi.org/10.1016/S0168-9002(98)01233-9)
103. Fedoseyev, A.I., Turowski, M., Alles, M.L. et al.: Accurate numerical models for simulation of radiation events in nano-scale semiconductor devices. *Math. Comput. Simul.* 79, 1086–1095 (2008). <https://doi.org/10.1016/j.matcom.2007.09.013>
104. Schrimpf, R.D., Weller, R.A., Mendenhall, M.H. et al.: Physical mechanisms of single-event effects in advanced microelectronics. *Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms.* 261, 1133–1136 (2007). <https://doi.org/10.1016/j.nimb.2007.04.050>
105. Bielajew, A.F., Rogers, D.W.O.: Variance-reduction techniques. In: Jenkins, T.E., Nelson, W.R., Rindi, A., Nalum, A.E., and Rogers, D.W.O. (eds.) *Monte Carlo Transport of Electrons and Photons*. pp. 407–420. Plenum Press, New York (1990)
106. Abe, F., Akimoto, H., Akopian, A. et al.: Observation of Top Quark Production in  $p\bar{p}$  Collisions with the Collider Detector at Fermilab. *Phys. Rev. Lett.* 74, 2626–2631 (1995). <https://doi.org/10.1103/PhysRevLett.74.2626>
107. Abachi, S., Abbott, B., Abolins, M. et al.: Observation of the Top Quark. *Phys. Rev. Lett.* 74, 2632–2637 (1995). <https://doi.org/10.1103/PhysRevLett.74.2632>
108. Abachi, S., Abolins, M., Acharya, B.S. et al.: The DØ detector. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 338, 185–253 (1994). [https://doi.org/10.1016/0168-9002\(94\)91312-9](https://doi.org/10.1016/0168-9002(94)91312-9)
109. Collaboration, D., Abazov, V.M., Abbott, B. et al.: Evidence for production of single top quarks. *Phys. Rev. D.* 78, 12005 (2008). <https://doi.org/10.1103/PhysRevD.78.012005>
110. Poon, E., Verhaegen, F.: Accuracy of the photon and electron physics in GEANT4 for radiotherapy applications. *Med. Phys.* 32, 1696–1711 (2005). <https://doi.org/10.1118/1.1895796>
111. Aharrouché, M., Colas, J., Ciaccio, L. et al.: Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 568, 601–623 (2006). <https://doi.org/10.1016/j.nima.2006.07.053>
112. Khramov, E., Rusakovich, N., Carli, T. et al.: Study of the Response of the Hadronic Barrel Calorimeter in the ATLAS Combined Test-beam to Pions of Energies from 20 to 350 GeV for Beam Impact Points from 0.2 to 0.65. Geneva (2009)
113. Adragna, P., Alexa, C., Anderson, K. et al.: Measurement of pion and proton response and longitudinal shower profiles up to 20 nuclear interaction lengths with the ATLAS Tile calorimeter. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 615, 158–181 (2010). <https://doi.org/10.1016/j.nima.2010.01.037>
114. Drozhdin, A.I., Huhtinen, M., Mokhov, N.V.: Accelerator related background in the CMS detector at LHC. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 381, 531–544 (1996). [https://doi.org/10.1016/S0168-9002\(96\)00807-8](https://doi.org/10.1016/S0168-9002(96)00807-8)
115. Aaboud, M., Aad, G., Abbott, B. et al.: A measurement of the calorimeter response to single hadrons and determination of the jet energy scale uncertainty using LHC Run-1 pp-collision data with the ATLAS detector. *Eur. Phys. J. C.* 77, 26 (2017). <https://doi.org/10.1140/epjc/s10052-016-4580-0>
116. CMS: Determination of jet energy calibration and transverse momentum resolution in CMS. *J. Instrum.* 6, P11002–P11002 (2011). <https://doi.org/10.1088/1748-0221/6/11/P11002>
117. Rahmat, R., Kroeger, R., Giammanco, A.: The Fast Simulation of The CMS Experiment. *J. Phys. Conf. Ser.* 396, 062016 (2012). <https://doi.org/10.1088/1742-6596/396/6/062016>
118. Grindhammer, G., Rudowicz, M., Peters, S.: The fast simulation of electromagnetic and hadronic showers. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 290, 469–488 (1990). [https://doi.org/10.1016/0168-9002\(90\)90566-O](https://doi.org/10.1016/0168-9002(90)90566-O)
119. Sekmen, S., Collaboration, for the C.: Recent Developments in CMS Fast Simulation. (2017)

120. Adloff, C., Blaha, J., Blaising, J.-J. et al.: Validation of GEANT4 Monte Carlo models with a highly granular scintillator-steel hadron calorimeter. *J. Instrum.* 8, P07005 (2013). <https://doi.org/10.1088/1748-0221/8/07/P07005>
121. Santina, G., Nieminen, P., Evansa, H. et al.: New Geant4 based simulation tools for space radiation shielding and effects analysis. *Nucl. Phys. B - Proc. Suppl.* 125, 69–74 (2003). [https://doi.org/10.1016/S0920-5632\(03\)90968-6](https://doi.org/10.1016/S0920-5632(03)90968-6)
122. Santin, G., Ivanchenko, V., Evans, H. et al.: GRAS: a general-purpose 3-D Modular Simulation tool for space environment effects analysis. *IEEE Trans. Nucl. Sci.* 52, 2294–2299 (2005). <https://doi.org/10.1109/TNS.2005.860749>
123. Gillam, J.E., Rafecas, M.: Monte-Carlo simulations and image reconstruction for novel imaging scenarios in emission tomography. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 809, 76–88 (2016). <https://doi.org/10.1016/J.NIMA.2015.09.084>
124. Castiglioni, I., Cremonesi, O., Gilardi, M.C. et al.: Scatter correction techniques in 3D PET: a Monte Carlo evaluation. *IEEE Trans. Nucl. Sci.* 46, 2053–2058 (1999). <https://doi.org/10.1109/23.819282>
125. Strulab, D., Santin, G., Lazaro, D. et al.: GATE (geant4 application for tomographic emission): a PET/SPECT general-purpose simulation platform. *Nucl. Phys. B - Proc. Suppl.* 125, 75–79 (2003). [https://doi.org/10.1016/S0920-5632\(03\)90969-8](https://doi.org/10.1016/S0920-5632(03)90969-8)
126. Arce, P., Lagares, J.I., Harkness, L. et al.: GAMOS: An easy and flexible way to use GEANT4. In: 2011 IEEE Nuclear Science Symposium Conference Record. pp. 2230–2237. IEEE (2011)
127. España, S., Herraiz, J.L., Vicente, E. et al.: PeneloPET, a Monte Carlo PET simulation tool based on PENELOPE: features and validation. *Phys. Med. Biol.* 54, 1723–1742 (2009). <https://doi.org/10.1088/0031-9155/54/6/021>
128. Barret, O., Carpenter, T.A., Clark, J.C. et al.: Monte Carlo simulation and scatter correction of the GE Advance PET scanner with SimSET and Geant4. *Phys. Med. Biol.* 50, 4823–4840 (2005). <https://doi.org/10.1088/0031-9155/50/20/006>
129. Kawrakow, I., Mitev, K., Gerganov, G. et al.: SU-GG-I-109: Using EGSnrc Within GATE to Improve the Efficiency Of positron Emission Tomography Simulations. *Med. Phys.* 35, 2667–2667 (2008). <https://doi.org/10.1118/1.2961507>
130. Seco, J.: Monte carlo techniques in radiation therapy. CRC Press (2016)
131. Zacharatou Jarlskog, C., Paganetti, H.: Physics Settings for Using the Geant4 Toolkit in Proton Therapy. *IEEE Trans. Nucl. Sci.* 55, 1018–1025 (2008). <https://doi.org/10.1109/TNS.2008.922816>
132. Paganetti, H., Jiang, H., Parodi, K. et al.: Clinical implementation of full Monte Carlo dose calculation in proton beam therapy. *Phys. Med. Biol.* 53, 4825–4853 (2008). <https://doi.org/10.1088/0031-9155/53/17/023>
133. Paganetti, H., Jiang, H., Adams, J.A. et al.: Monte Carlo simulations with time-dependent geometries to investigate effects of organ motion with high temporal resolution. *Int. J. Radiat. Oncol. Biol. Phys.* 60, 942–50 (2004). <https://doi.org/10.1016/j.ijrobp.2004.06.024>
134. Aso, T., Yamashita, T., Akagi, T. et al.: Validation of PTSIM for clinical usage. In: IEEE Nuclear Science Symposium & Medical Imaging Conference. pp. 158–160. IEEE (2010)
135. Perl, J., Shin, J., Schümann, J. et al.: TOPAS: An innovative proton Monte Carlo platform for research and clinical applications. *Med. Phys.* 39, 6818–6837 (2012). <https://doi.org/10.1118/1.4758060>
136. Böhlen, T.T., Cerutti, F., Dosanjh, M. et al.: Benchmarking nuclear models of FLUKA and GEANT4 for carbon ion therapy. *Phys. Med. Biol.* 55, 5833–5847 (2010). <https://doi.org/10.1088/0031-9155/55/19/014>
137. Dedes, G., Pinto, M., Dauvergne, D. et al.: Assessment and improvements of Geant4 hadronic models in the context of prompt-gamma hadrontherapy monitoring. *Phys. Med. Biol.* 59, 1747–1772 (2014). <https://doi.org/10.1088/0031-9155/59/7/1747>

138. Mairani, A., Brons, S., Cerutti, F. et al.: The FLUKA Monte Carlo code coupled with the local effect model for biological calculations in carbon ion therapy. *Phys. Med. Biol.* 55, 4273–4289 (2010). <https://doi.org/10.1088/0031-9155/55/15/006>
139. Robert, C., Dedes, G., Battistoni, G. et al.: Distributions of secondary particles in proton and carbon-ion therapy: a comparison between GATE/Geant4 and FLUKA Monte Carlo codes. *Phys. Med. Biol.* 58, 2879–99 (2013). <https://doi.org/10.1088/0031-9155/58/9/2879>
140. De Napoli, M., Agodi, C., Battistoni, G. et al.: Carbon fragmentation measurements and validation of the Geant4 nuclear reaction models for hadrontherapy. *Phys. Med. Biol.* 57, 7651–7671 (2012). <https://doi.org/10.1088/0031-9155/57/22/7651>
141. Battistoni, G., Bauer, J., Boehlen, T.T. et al.: The FLUKA Code: An Accurate Simulation Tool for Particle Therapy. *Front. Oncol.* 6, 116 (2016). <https://doi.org/10.3389/fonc.2016.00116>
142. Nikjoo, H., Uehara, S., Emfietzoglou, D. et al.: Track-structure codes in radiation research. *Radiat. Meas.* 41, 1052–1074 (2006). <https://doi.org/10.1016/j.radmeas.2006.02.001>
143. El Naqa, I., Pater, P., Seuntjens, J.: Monte Carlo role in radiobiological modelling of radiotherapy outcomes, (2012)
144. Bernal, M.A., Bordage, M.C., Brown, J.M.C. et al.: Track structure modeling in liquid water: A review of the Geant4-DNA very low energy extension of the Geant4 Monte Carlo simulation toolkit. *Phys. Medica.* 31, 861–874 (2015). <https://doi.org/10.1016/J.EJMP.2015.10.087>
145. Incerti, S., Douglass, M., Penfold, S. et al.: Review of Geant4-DNA applications for micro and nanoscale simulations. *Phys. Medica.* 32, 1187–1200 (2016). <https://doi.org/10.1016/J.EJMP.2016.09.007>
146. Pedoux, S., Cugnon, J.: Extension of the Liège intranuclear cascade model at incident energies between 2 and 12 GeV. Aspects of pion production. *Nucl. Phys. A.* 866, 16–36 (2011)
147. Dingfelder, M., Ritchie, R.H., Turner, J.E. et al.: Comparisons of calculations with PARTRAC and NOREC: transport of electrons in liquid water. *Radiat. Res.* 169, 584–594 (2008). <https://doi.org/10.1667/RR1099.1>
148. Thomson, M.A.: Particle flow calorimetry and the PandoraPFA algorithm. *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 611, 25–40 (2009). <https://doi.org/10.1016/j.nima.2009.09.009>
149. Aduszkiewicz, A., Ali, Y., Andronov, E. et al.: Measurements of  $\pi^\pm$ ,  $K^\pm$ , p and p-bar spectra in proton-proton interactions at 20, 31, 40, 80 and 158 GeV/c with the NA61/SHINE spectrometer at the CERN SPS. *Eur. Phys. J. C.* 77, 671 (2017). <https://doi.org/10.1140/epjc/s10052-017-5260-4>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 12

## Triggering and High-Level Data Selection



W. H. Smith

### 12.1 Level-1 Trigger


#### 12.1.1 Introduction

The data taken by a particle physics collider detector consists of events, which are snapshots of the detector data at specific intervals in time. Usually these snapshots are taken at the frequency of the crossing of the colliding beams. For HERA this was 96 ns, for the Tevatron Run II this was 396 ns and for the LHC at design luminosity this is 25 ns. An individual bunch crossing may contain either no, one or many interactions between the particles in the colliding beams. The time during which beam collisions take place during a beam crossing is 1–2 ns. Even if there are multiple collisions in a single crossing the detector elements will make only one recording and the events will be superimposed. Therefore, each bunch crossing is individually evaluated. Not all of the detector data from an individual crossing is available immediately. Some may be stored as charge and need digitization. Other digital detector data may be inaccessible until further detector processing is complete.

The selection of bunch crossings is a highly complex function that involves a series of levels which take increasing amounts of time, process increasing amounts of data, use increasingly complex algorithms and make increasingly more precise determinations to reject increasing numbers of crossings. The first level(s) of the series usually involve(s) specific custom high-speed electronics. The subsequent level(s) involve more general CPU farms that run code similar to that found in the offline reconstruction. Due to this structure, the first level of trigger decision

---

W. H. Smith (✉)  
University of Wisconsin-Madison, Madison, WI, USA  
e-mail: [wsmith@hep.wisc.edu](mailto:wsmith@hep.wisc.edu)



**Fig. 12.1** Layout of the elements of the L1T

is based on particle identification (e.g. muon, electron, etc.) from local pattern recognition and energy evaluation. The higher trigger levels start by identifying the particle signature (e.g. Z, W, etc.), calculating kinematics for effective mass and event topology cuts and performing track reconstruction and detector matching (e.g. muon and tracking or calorimeter and tracking). The highest-level triggers perform identification of the physics process detected using event reconstruction and analysis. As shown schematically in Fig. 12.1 the Level-1 trigger<sup>1</sup> (L1T) inspects a subset of the detector information for each bunch crossing and provides the first in a series of decisions to either keep or discard it. The L1T system generally uses coarsely segmented data from calorimeter and muon detectors and in a few cases some rudimentary tracking detector information, while holding all the high-resolution data in pipeline memories in the front-end electronics. During the L1T decision time that is typically a few  $\mu\text{secs}$ , all of the data from all crossings are stored. Usually a good fraction of the L1T time is used in transmission of the L1T data from the detector front ends to a central location where trigger processing is performed and transmission of the L1T decision back to the front ends, leaving a fraction of the L1T decision time available for the trigger processing.

The need to process each new crossing of data requires that the L1T function in a pipelined mode, e.g. be composed of a series of steps each of which processes its input and produces its output result at the crossing frequency. As noted above this can range from 396 ns at the Tevatron to 96 ns at HERA to 25 ns at the LHC. In order to avoid dead time, the trigger electronics must itself be pipelined: every process in the trigger must be repeated at the beam-crossing rate. This has important consequences for the requirements on the structure of the trigger system. The fact that each piece of logic must accept new data at the beam-crossing rate means that no

<sup>1</sup>This is commonly called Level-1 but in such experiments as ALICE and LHCb this corresponds to the Level-0.

piece of individual data processing can take more than this time. This prohibits the use of iterative algorithms, such as jet finding based on finding a seed tower and then adding the surrounding towers to make a jet energy sum. This pipelined structure means that each step in the L1T logic must be completed within the time of the crossing frequency and the results output so that this step in the logic is available to process the data from the next crossing. The L1T logic therefore consists of a number of pipelined steps equal to the processing time multiplied by the crossing frequency.

The tight timing structure of the L1T presents a couple of challenges. Generally, the detector calorimeters have long pulse shapes that exceed the time between beam-crossings. This implies that particles produced in different bunch crossings can produce significant pulse-height in the bunch crossing of interest. Therefore, the detector systems that calculate the input information for the trigger need to correctly identify the energy associated with the correct bunch crossing, usually against a background of additional energy deposits from other bunch crossings. Typically, these systems use peak-finding algorithms and finite input response filters to perform this determination. The gaseous tracking detectors used in the muon systems can also have drift times or pulse widths exceeding the time between bunch crossings. These systems are also required to not only detect the passage of the charged track but also to identify the crossing that produced the track. Often this is resolved by combining and comparing the hits found in adjacent planes of chambers. Another challenge is that the physical extent of large HEP detectors produces times of flight to traverse them that exceed the time between bunch crossings. Therefore, at any particular point in time, the particles from interactions of more than one bunch crossing are present in the detector at different locations. This requires tight timing and synchronization of the detector trigger and readout systems.

The trigger is the start of the physics event selection process. A decision to retain an event for further consideration has to be made at the crossing frequency. This decision is based on the event's suitability for inclusion in one of the various data sets to be used for analysis. The data sets to be taken are determined by the experiment's physics priorities as a whole. Examples of data sets used in LHC experiments include di-lepton and multi-lepton data sets for top and Higgs studies, lepton plus jet data sets for top physics, and inclusive electron data sets for calorimeter calibrations. In addition, other samples are necessary for measuring efficiencies in event selection and studying backgrounds. The trigger has to select these samples in real time along with the main data samples.

The L1T is based on the identification of physics objects such as muons, electrons, photons, jets, taus and missing transverse energy. Each of these objects is typically tested against several  $p_T$  or  $E_T$  thresholds. The efficiency of a trigger is determined by dividing the number of events that pass the trigger by the number of actual events that would populate the final physics results plots if all of them passed the trigger. The trigger must have a sufficiently high and understood efficiency at a sufficiently low threshold to ensure a high yield of events in the final physics plots to provide enough statistics and a high enough efficiency for these events so that the correction for this efficiency does not add appreciably to the systematic error of the



measurement. The efficiency of the trigger is evaluated with respect to benchmark physics processes derived from the physics goals of the experiment. The criteria are a sharp turn-on curve of the efficiency at its threshold and an asymptote as close to 100% as possible. The L1T thresholds should be somewhat smaller than the offline physics analysis cuts. The reason for such a requirement is that the efficiency turn-on curves for the L1T will be somewhat softer than can be achieved with a full analysis including the best resolutions and calibration corrections.

Much of the logic in contemporary L1T systems is contained in custom Application Specific Integrated Circuits (ASICs), semi-custom or gate-array ASICs, Field Programmable Gate Arrays (FPGAs), Programmable Logic Devices (PLD), or discrete logic such as Random-Access Memories (RAM) that are used for memory Look-Up Tables (LUT). Given the remarkable progress in FPGA technology, both in speed and number of gates, the technology of many trigger systems has mostly moved towards full implementation in FPGAs.

The key to a good trigger system is flexibility. Not only should all thresholds be programmable, but also as mentioned above, algorithms are either implemented in FPGAs or LUTs. Reprogramming the FPGAs or downloading new LUT contents allows for revisions of the trigger algorithms. The only fairly fixed aspect of the trigger system is which data is brought to which point for processing. However, this is determined by the detector elements, size of showers and curvature of tracks, which are well known and basic features of the detectors and physics signals. There are new technologies being developed that are expected to provide flexibility in data routing, including backplanes and cards that use programmable cross-point switches.

The L1T system sustains a large dataflow. This is either carried on optical fibres, copper cables, or on backplanes within crates. At the LHC, the data carried by these means may be sent in parallel at either 40 MHz, or a higher multiple of this frequency, or converted from parallel to serial and transmitted at a higher rate on a single lines or pair of lines. Serial data transmission has the advantage of transmitting more data per cable wire or backplane pin but the disadvantage of extra latency for the parallel to serial and serial to parallel operations plus the risk of data errors involved with the encoding, high frequency transmission and link synchronization. In many cases this requires the overhead of monitoring and error detection bits. Copper cables in general avoid the necessity for optical drivers with their cost, size and power requirements, but have limited length capability, take up more volume and use more material.

### ***12.1.2 L1T Requirements***

The L1T has to be inclusive, local, measurably efficient, and fill the DAQ bandwidth with a high purity stream. The local philosophy of the trigger implies an initial trigger selection of electrons, photons, muons and jets that relies on local information tied directly to their distinctive signatures, rather than on global topologies. For

example, electron showers are small and extremely well defined in the transverse and longitudinal planes. Information from a few Electromagnetic and Hadronic calorimeter towers at the L1T, the corresponding elements of the preshower detector, and a small region of the tracking volume (at higher trigger levels) are sufficient for electron identification. The only global entities are neutrinos (from a global sum of missing  $E_T$ ).

For the trigger to be measurably efficient the tools to measure lepton and jet efficiencies must be built into the trigger architecture from the start. One such tool is overlapping programmable triggers so that multiple triggers with different thresholds and cuts that can run in parallel. A second tool is pre-scaled (e.g. random selection of a fraction) triggers of lower threshold or weaker criteria that run in parallel with the stricter triggers. A third tool is pre-scaling of a particular trigger with one of its cuts removed.

The requirement on the use of DAQ bandwidth implies two conditions. First, each level of the trigger attempts to identify leptons and jets as efficiently as possible, while keeping the output bandwidth within requirements. The selected event sample should include all events that would be found by the full offline reconstruction. Hence, the selection criteria in the trigger must be consistent with those of the offline. Second, since the bandwidth to permanent storage media is limited, events must be selected with care at the final trigger level.

The measurement of trigger efficiency requires the flexibility to have overlapping triggers so that efficiencies can be measured from the data. The overlaps include different thresholds, relaxed individual criteria, prescaled samples with one criterion missing, and overlapping physics signatures. For example, measurement of the inclusive jet spectrum uses several triggers of successively higher thresholds, with the lower thresholds prescaled by factors that allow a reasonable rate to storage. These triggers overlap in jet energy all the way down to minimum bias events so that the full spectrum can be reconstructed accurately. The efficiency and bias of each higher threshold can be measured from the data sets of lower threshold. A requirement for understanding the trigger efficiency is that the data used as input to the L1T system is also transmitted via the DAQ for storage along with the event readout data. In addition, all trigger objects found, whether they were responsible for the L1 trigger or not should also be sent.

The L1T accept rate is limited by the speed of the detector electronics readout and the rate at which the data can be harvested by the data acquisition system. Since it is pipelined and deadtimeless, the L1T renders a decision on every bunch crossing. The maximum L1T accept rate is set by the average time to read information for processing by the Higher Level Triggers (HLT) and the average time for completion of processing steps in the HLT logic.

The high operational speed and pipelined architecture also requires that specific data is brought to specific points in the trigger system for processing and that there cannot be fetching of data based on analysis of other data in an event. The data must flow synchronously across the trigger logic in a deterministic manner in the same way for each crossing. At any moment there are many crossings being processed in sequence in the various stages of the trigger logic. The consequence is that most


of the L1T operations are either simple arithmetic operations or functions using memory lookup tables where an address of data produces a result previously written into the memory.

The L1T requirements evolve with the experiment luminosity, energy and event pile-up (number of p–p collisions per beam crossing). For example for the LHC trigger systems [20], algorithms used by the ATLAS [13] and CMS [14] experiments at the LHC during the period before 2014 (Run-1) were optimized for 7–8 TeV center-of-mass energy, PU up to 40 due to the 20 MHz beam crossing frequency and luminosities up to  $7 \times 10^{-33} \text{ cm}^2 \text{ s}^{-1}$ , whereas afterwards (Run-2) these were optimized for a 13 TeV center-of-mass energy and PU above 50 due to the 40 MHz beam crossing frequency and luminosities exceeding  $15 \times 10^{-33} \text{ cm}^2 \text{ s}^{-1}$  [15, 16].

### 12.1.3 Muon Triggers

The design of L1T muon trigger logic depends on the detectors being used to generate the trigger information. These detectors include those with timing resolution and prompt signals that are generally less than the time between bunch crossings such as Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs). They also include special signal handling of detectors with individual signals and resolution greater than the bunch crossing time, such as Cathode Strip Chambers (CSCs) and Drift Tube Chambers (DTs). For these detectors, offset detector planes, front-end logic that processes over the drift time, and combinations of planes provide identification of the bunch crossing associated with the muon passage. Another important feature in muon trigger design is whether the muon chamber measuring stations are placed in a magnetic field in air or embedded in iron. In the former case, the muon momentum resolution is usually sufficient to provide an efficient threshold up to relatively high  $p_T$ . In the latter case, information from the tracking detectors is needed to provide a sufficiently sharp threshold.

L1T muon algorithms depend on comparison of tracks of hits with predefined geometrical patterns such as roads. For example, the ATLAS muon trigger employs RPCs and TGCs in an air-core magnetic field and the trigger algorithm uses Coincidence Windows that start with a hit in a central “pivot plane” and searches for time-correlated hits within an  $\eta$ – $\phi$  window in a “confirm plane” [1]. Different “confirm planes” are used for low and high  $p_T$  muons, as is shown in Fig. 12.2. The RPC barrel algorithm extrapolates hits in the middle RPC 2 station to a point and coincidence window in the innermost RPC 1 station along a straight line to the nominal interaction point. The size of this coincidence window depends on the muon’s bend in the magnetic field. A low- $p_T$  candidate is found if there is one hit in this window and hits in both views and planes of either RPC 1 or RPC 2. If there is also a hit in RPC 3, then a high- $p_T$  candidate has been found. For Run 2 ATLAS commissioned a fourth layer of barrel RPCs that improved the acceptance and added new trigger logic to the end-cap requiring additional coincidences with the TGC’s



**Fig. 12.2** ATLAS muon trigger algorithms

or the Tile hadronic calorimeter to reject particles not originating at the interaction point [15].

The CMS Detector uses Drift Tubes (DT), Cathode Strip Chambers (CSC) and overlapping Resistive Plate Chambers (RPC) for muon triggering in iron. The RPC readout strips are connected to pattern logic, which is projective in  $\eta$  and  $\phi$  and connected to segment processors that find the tracks and calculate the  $p_T$ . As shown in Fig. 12.3 the CSC logic forms Local Charged Tracks (LCT) from the charge distributions in the CSC planes, which are combined with the Anode wire information for bunch crossing identification and assignment of  $p_T$  and “quality”, which is an indicator of the number of planes hit. The CSC Track Finder combines the LCTs into full muon tracks and assigns  $p_T$  values to them. As is also shown in Fig. 12.3 the DTs are equipped with Bunch and Track Identifier (BTI) electronics that finds track segments from coincidences of aligned hits in four layers of one drift tube superlayer. The DT Track Finder combines the segments from different stations into full muon tracks and assigns  $p_T$  values to them. In Run 1, the Global Muon Trigger sorted and then correlated the RPC, DT and CSC muon tracks. In Run 2, the RPC, DT and CSC information were combined earlier, in the track-finding stage [12].

The LHCb Level-0 muon trigger searches for candidates in the quadrants of five stations of Multi-Wire Proportional Chambers separated by iron and sends the two highest  $p_T$  candidates from each quadrant to the Level-0 Decision Unit (L0DU) [2]. The ALICE dimuon trigger system is based on two stations of 18 RPCs each read out on both sides of the gas gap by  $X$ – $Y$  orthogonal strips with high resolution front-end electronics which feed local trigger electronics modules that find tracks in 3 out of the 4 detector planes in both  $X$  and  $Y$  [3]. The track is found and the magnetic deviation is calculated to enable a cut on a  $p_T$  threshold using memory Look-Up Tables (LUTs). Two unlike-sign muons are then required in the L1T.





Fig. 12.3 CMS muon chamber trigger algorithms

### 12.1.4 Calorimeter Electron and Photon Triggers

The calorimeter trigger begins with trigger tower energy sums formed by the detector electromagnetic calorimeter (ECAL) hadronic calorimeter (HCAL) and forward calorimeter. Experiments vary on whether these sums are performed by analog methods before digitization or by digital summation after an initial ADC.

For the ATLAS experiment, the calorimeter trigger begins with a Preprocessor (PPr) which sums analog pulses into  $0.1 \times 0.1$  ( $\eta \times \phi$ ) trigger towers, assigns their bunch crossing and adjusts for calibration. The Cluster Processor then identifies and counts electron/photon and tau candidates based on the energies and patterns of energy isolation found in overlapping windows of  $4 \times 4$  ECAL and HCAL trigger towers as shown in Fig. 12.4. For Run 2, the PPr was upgraded to provide improved Finite Input Response (FIR) filtering and dynamic bunch by bunch pedestal correction [15]. New cluster merging modules (CMX) were added that transmitted the location and energy of trigger objects, rather than the threshold multiplicities used in Run 1.

The CMS Calorimeter trigger algorithm for electron and photon candidates uses a  $3 \times 3$  trigger tower sliding window centered on all ECAL/HCAL trigger towers. A diagram of this electromagnetic algorithm is shown in Fig. 12.5. Two types of electromagnetic objects are defined. The non-isolated electron/photon identification is based on a large energy deposit in one or two adjacent ECAL 5-cell  $\phi$  strips in the trigger tower, the lateral shower profile in the central tower comparing maximum  $E_T$




**Fig. 12.4** ATLAS calorimeter electron/photon trigger algorithm

of each of four pairs of strips of 5 cells to the total tower level  $E_T$  of all 25 crystals (this “Fine Grain” veto uses a strip due to electron bending in the magnetic field), and the longitudinal shower profile defined by the ratio of  $E_T$  deposits in the HCAL and ECAL portions of the calorimeter (H/E veto). The isolated electron/photon has two additional requirements: the ECAL  $E_T$  deposited in one of the five trigger towers surrounding the central tower is below a programmable  $E_T$  threshold and the eight trigger towers surrounding the central tower in the  $3 \times 3$  region have passed the Fine Grain and H/E vetoes. For Run-2, the CMS Calorimeter Trigger hardware was upgraded so that more complex algorithms could be deployed [21]. The  $e/\gamma$  and  $\tau$  candidates started with a local maximum around which the trigger towers were dynamically clustered.

The LHCb Level-0 calorimeter trigger system combines the  $E_T$  measurement in clusters of  $2 \times 2$  cells in the electromagnetic (ECAL) and hadronic calorimeters (HCAL), as well as information from the Scintillator Pad Detector (SPD) and a Preshower (Prs) to indicate the charged and electromagnetic nature of the clusters. The calorimeter trigger system sends the highest  $E_T$  hadron, electron, photon and  $\pi^0$  candidates and the total HCAL  $E_T$  and SPD multiplicity to the Level 0 Decision Unit (LODU).





**Fig. 12.6** Organization of the ATLAS jet trigger system

refines the muon trigger with a sharper momentum threshold that is not affected by the backgrounds in the muon chambers. It also can be used to identify heavy flavour candidates. Both Tevatron experiments CDF and DØ employed level-1 tracking triggers. CDF used signals from the Central Outer Tracker (COT) open-cell drift chamber in the eXtremely Fast Tracker (XFT) to perform charged track reconstruction in the  $r$ - $\phi$  plane for the L1T [4]. Track segments were found by comparing hit patterns in a COT superlayer to a list of valid patterns or “masks”. These masks contained specific patterns of prompt and delayed hits on the 12 wire layers of an axial COT superlayer. Tracks were found by comparing track segment patterns in all four layers to a list of valid segment patterns or “roads”. The XFT had an efficiency  $>90\%$  for tracks with  $p_T > 1.5$  GeV/c, transverse momentum resolution of  $\delta p_T/p_T = 0.002 p_T$  and pointing resolution of  $\delta\phi = 0.002$  radians with respect to the beam line [5]. The XFT reported the highest  $p_T$  track in each of 288 azimuthal



segments ( $1.25^\circ$  each) to the XFT “Linker system” modules which cover  $15^\circ$  each and are matched to the segmentation of the trigger signals from the muon and calorimeter systems. The results from the linker system were passed to the Track Extrapolation System (XTRP), which sent one or more bits in  $2.5^\circ$  segmentation to the muon trigger systems set according to the calculated  $p_T$ ,  $\phi$  and multiple scattering. The XTRP also sent a set of 4 bits (for four momentum thresholds) for each  $15^\circ$  calorimeter wedge to the Level-1 calorimeter trigger. Finally the XTRP created a Level-2 tracking trigger based on the number of tracks and their  $p_T$  and  $\phi$  information.

The DØ experiment Central Tracking Trigger (CTT) used information from the Central Fiber Tracker (CFT) and the Central Preshower System (CPS). Hit information from each of the 80 axial sectors of the CFT/CPS detectors was fed through boards programmed with 16,000 Boolean equations that identified patterns of hits likely to be produced by a charged particle. A list of tracks in four momentum ranges between 1.5 and 10 GeV/c was then sent to the L1 muon trigger system [6]. The DØ L1 CTT also identified the number of tracks in each event for each of these four momentum ranges, whether a coincident CPS hit had been found, and whether the track was isolated. This information was also used in the DØ L1T decision. The DØ CTT had an efficiency of  $97.3 \pm 0.1\%$  for tracks with  $p_T > 10$  GeV/c [4].

Although both ATLAS and CMS are planning the use of Tracking information at Level-1 in their designs for the High Luminosity LHC (HL-LHC) project [22], this information was not included in Run-1 or Run-2.

### 12.1.7 Global Triggers

An experiment Global Trigger accepts muon, calorimeter and tracking (if available) trigger information, synchronizes matching sub-system data arriving at different times and communicates the Level-1 decision to the timing, trigger and control system for distribution to the sub-systems to initiate the readout. The global trigger decision is made using logical combinations of the input trigger data. Besides handling physics triggers, the Global Trigger provides for test and calibration runs, not necessarily in phase with the machine, and for prescaled triggers, as this is an essential requirement for checking trigger efficiencies and recording samples of large cross section data.

The ATLAS Level-1 Global trigger is called the Central Trigger Processor (CTP). It combines information on the multiplicities of calorimeter and muon trigger objects which have sufficiently high momentum. These are electrons/photons, taus, jets, and muons. These are also the “seeds” for the Level-2 trigger that are sent to the Region of Interest Builder (RoIB). In addition, threshold information on the global transverse energy and missing energy sums is also used in the Level-1 decision. In Run-1, the CTP discriminated the delivered multiplicities of the trigger objects against multiplicity conditions and then combined these conditions to form more complex triggers when multiple object triggers are needed. In Run-2, the ATLAS

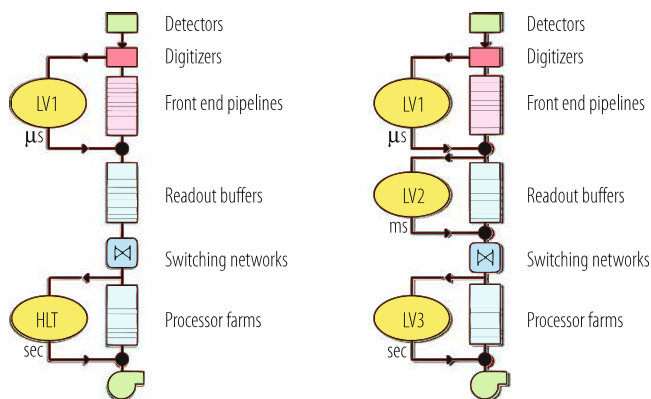
L1 Global trigger added a topological trigger (L1Topo) to allow geometrical or kinematic association between trigger objects received from the L1 Calorimeter or Muon Triggers [23].

The CMS L1 Global Trigger sorts ranked trigger objects, rather than histogramming objects over a fixed threshold. This allows all trigger criteria to be applied and varied at the Global Trigger level rather than earlier in the trigger processing. All trigger objects are accompanied by their coordinates in  $(\eta, \phi)$  space. This allows the Global Trigger to vary thresholds based on the location of the trigger objects. It also allows the Global Trigger to require trigger objects to be close or opposite from each other. In addition, the presence of the trigger object coordinate data in the trigger data, which is read out first by the DAQ after a Level-1 trigger accept (L1A), permits a quick determination of the regions of interest where the more detailed HLT analyses should focus. The Global L1 Trigger transmits a decision to either accept (L1A) or reject each bunch crossing. This decision is transmitted through the Trigger Throttle System (TTS) to the Timing Trigger and Control system (TTC). The TTS allows the reduction by prescaling or blocking L1A signals in case the detector readout or DAQ buffers are at risk of overflow. For Run-1, the Global L1 Trigger allowed up to 128 algorithms to contribute to the overall trigger decision. For Run-2, this was upgraded to a modular design capable of up to about 500 algorithms that was typically running about 300 [16].

## 12.2 Higher-Level Selection

### 12.2.1 Introduction

The design of the Higher-Level Selection of events after Level-1 takes place in a number of “trigger levels”. Generally, collider experiments use at least two additional trigger levels, referred to as the Level-2 and Level-3 trigger. Some experiments have a Level-4 trigger. The higher the number the more general purpose (or commercial) the implementation, with the Level-3 and Level-4 triggers being composed of farms of standard commodity computers. The physical implementation of the Level-2 trigger varies substantially between experiments from inclusion in the Level-3 farm of processors to an independent farm of processors to customized dedicated processing hardware. The Level-2 trigger has to operate at the output rate of the Level-1 trigger, generally with a subset of the higher resolution and full-granularity available to the full reconstruction code available at Level-3 and higher. Typically, the Level-1 output rate ranges between 1 and 100 kHz depending on the experiment. The Level-2 trigger is generally limited in execution time so that the full event data cannot be unpacked and processed. Instead, the higher resolution and full-granularity data is unpacked in “regions of interest” determined by the Level-1 trigger data.



**Fig. 12.7** Common architectures for collider detector trigger and data acquisition systems. Left: two physical levels. Right: three physical levels

The architectures of Level-2 trigger systems vary depending on the rejection factor required, the information provided as input, the interconnections with the front-end electronics, Level-1 and Level-2. Examples of two types of architecture presently employed by general-purpose collider detectors are shown in Fig. 12.7. Including Level-1, experiments such as H1, ZEUS, CDF, DØ and ATLAS have three physical levels of processing [18]. For Run-2, the ATLAS Higher Level Trigger Layers (HLT) were combined [15]. CMS has two layers of physical processing [19]. LHCb has three levels of processing, but the first level (Level-0) output trigger rate is 1.1 MHz, an order of magnitude higher than other collider experiments [7, 17]. The subsequent levels, HLT1 and HLT2, are software-based, running on the Event Filter Farm. In Run-2, HLT1 and HLT2 became two independent asynchronous processes on the same node and HLT2 was able to run a full reconstruction on real-time aligned and calibrated data [17].

There are more substantial differences in trigger architecture for experiments such as ALICE that is designed to study heavy ion collisions with a bunch spacing of 125 ns at a lower luminosity than the LHC experiments ATLAS and CMS. However, each Pb–Pb collision produces much higher multiplicities of secondary particles than a p–p collision, resulting in a much higher event size. Since the detectors in ALICE have different readout times, there are three parallel trigger systems, allowing readout from the faster detectors while slower detectors are occupied with reading out the data from earlier events [8]. The first decision is made 1.2  $\mu$ s after the event (Level-0), the Level-1 decision comes after 6.5  $\mu$ s, and the Level-2 trigger is issued after 88  $\mu$ s. The Level-1 and Level-2 decisions can veto trigger signals from Level-0. The ALICE Central Trigger Processor also checks the events for pile-up from events in a programmable time interval before and after the interaction at all three levels. For Run-2, an earlier L0 trigger decision time of 525 ns provides a pre-trigger for the TRD [24].

The algorithms deployed in the HLT are dynamic, reflecting continuing improvements in the offline reconstruction that represent the functions that the HLT is attempting to approach within the constraints of processing time. The descriptions below of algorithms in the LHC experiments represent a snapshot at the time of Run-1 processing. These considerably evolved during Run-1 into Run-2, although the general techniques shown continue to be applied.

### ***12.2.2 Tracking in Higher Level Triggers***

The principal new information in the higher level triggers is tracking information. Either it is introduced for the first time in the event selection process or it is greatly refined over rudimentary tracking used in the Level-1 trigger. There are two major sources of tracking information. A pixel detector provides the most inner tracking and some vertex information. Outside of the pixel detector, silicon strip and then in some cases drift chambers, fibers, or straw tube detectors provide additional information at larger radius. For example, ATLAS [9] uses space points found in the pixels and silicon central tracker (SCT) to find the  $z$ -vertex location, fit tracks into the Transition Radiation Tracker (TRT) and measure the  $\phi$  and  $p_T$  of the track above a  $p_T$  of 0.5 GeV/c. In the latter part of Run-2, ATLAS commissioned the Fast TracKer (FTK), a dedicated Associative-Memory hardware processor which delivers tracks with  $p_T > 1$  GeV/c for every L1A to the HLT within 100  $\mu$ s [25]. In CMS, two types of tracking are employed. Charged particle tracks are first quickly reconstructed using pixel hits and then more laboriously but more accurately reconstructed with additional hits from the silicon strip tracker. Generally, tracking is “seeded” by the confirmed Higher Level Trigger objects, which themselves are “seeded” by Level-1 trigger objects.

### ***12.2.3 Selection of Muons***

The first algorithms executed in Level-2 on Level-1 selected muons are refinements of the reconstruction of the tracks in the muon chambers. In the case of ATLAS, where only the RPC (Barrel) and TGT (Forward) chambers provide information for the L1T, the precision hit information from the Monitored Drift Tubes (MDTs) is added to the RPC and TGC determined candidates. This provides good track reconstruction in the muon spectrometer. Since the ATLAS Muon Chambers are mostly in air, there is little multiple coulomb scattering. The found tracks are extrapolated for combination with tracks found in the Inner Detector. Matching between muon tracks measured independently in the muon system and those in the Inner Detector selects prompt muons and rejects fake and secondary muons. The isolated muon triggers also use information from the calorimeter towers surrounding the found muon track.

In CMS, all of the muon chamber systems participate in the L1T. The L1T muon candidates are used to seed the reconstruction of tracks in the muon chambers in the Level-2 algorithm. First, an initial pattern recognition is performed on muon segments along the trajectory, then a second more precise fit using all hits on these segments is used to determine the muon parameters. Since the CMS chambers are surrounded by steel, the propagation of track parameters to adjacent muon stations must take into account material effects such as multiple Coulomb scattering, and energy losses due to ionization and bremsstrahlung in the muon chambers and the iron. To avoid excessive processing times, these are estimated from fast parameterizations. Muons passing this first reconstruction are then input to the Level-3 reconstruction that uses hits in the silicon tracker within a rectangular  $\eta \times \phi$  region. Pairs or triplets of hits in the innermost layers of the tracker form trajectory seeds that are required to be compatible with the  $\eta \times \phi$  region and the primary vertex constraints. These are then grown into tracks of about seven hits and optionally combined with the reconstructed hits from the Level-2 algorithm. In Level-2, the isolation variable is calculated from the weighted sums of energies deposited in the ECAL and in the HCAL in the region around the muon track. For the Level-3 isolation variable, only charged-particle tracks near the vertex of the candidate muon are selected for inclusion. This excludes tracks from pile-up of contributions from other pp collisions (which occur at another vertex location), making this isolation less sensitive to pile-up than calorimetric isolation.

### ***12.2.4 Selection of Electrons and Photons***

The first algorithms executed in Level-2 on Level-1 selected electrons and photons are refinements of the clustering algorithms. For example, in ATLAS [9], the energy deposited in windows of the electromagnetic LAr calorimeter cells and the energy-weighted position information, as well as the leakage energy into the hadronic calorimeter are calculated. CMS [10] also reconstructs energy in clusters of electromagnetic calorimeter cells corresponding to the Level-1 calorimeter triggers, adding a margin around the trigger region to ensure complete collection of energy. These clusters are then formed into “Super Clusters” which are groups of clusters along a road in the  $\phi$  direction, chosen due to bending in the magnetic field. These clusters are then required to be isolated in the electromagnetic calorimeter. The hadronic calorimeter energies are then reconstructed and the energies in the hadronic tower behind the cluster and the adjacent towers are required to be small with respect to the electromagnetic cluster energy.

The second tier of algorithms performed on electrons and photons confirmed by the first algorithms are tracking algorithms. The first or more local steps of these are generally called Level 2.5 algorithms. This involves establishing track isolation around the electromagnetic cluster and for electron triggers, associating the electromagnetic cluster with a track. For CMS electron triggers, the energy and position of the Super Cluster is used to search for hits in the pixel detector. These

hits are reconstructed and the track  $p_T$  is checked for consistency with the Super Cluster energy. For both electron and photon triggers, tracks are seeded from pairs of hits in the pixel layers in a rectangular  $\eta \times \phi$  region around the direction of the reconstructed electron or photon, where these seeds are required to be consistent with the nominal vertex spread (photons) or closest approach of the electron path to the beam line (electrons). Then for electrons a threshold is applied to the  $p_T$  sum of the tracks within a cone around the electron direction and on the number of tracks for the photon. In ATLAS [11] the electromagnetic cluster is identified as an electron by association with a track in the Inner Detector, which is found by independent searches in the SCT/Pixel and TRT detectors in the region identified by the L1T RoI. For electron candidates, matching in both position and momentum between the track and cluster is required.

### 12.2.5 Selection of Jets and Missing Energy

The primary processing of the jet candidates at Level-2 begins with the L1T jet candidates, which are used as seeds for the Level-2 jets. The first step is to recalculate the jet energy for these candidates using the full granularity and calorimeter energy resolution information, which is not available to the Level-1 jet energy calculation. In ATLAS, the Level-2 jet finding searches in the RoIs produced by the Level-1 calorimeter logic. In CMS, jets are reconstructed using an iterative cone algorithm with cone size  $R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.5$  that sums over all projected electromagnetic and hadronic calorimeter cells with energy greater than a threshold set above the level of noise (0.5 GeV). In addition, to be declared a jet, at least one seed tower must have  $E_T > 1$  GeV. After summation, the jet energy is adjusted by an  $\eta$ -dependent correction for the calorimeter response.

Missing energy is calculated by summing all towers with  $E_T$  above a noise threshold. For CMS, this threshold is 0.5 GeV. No energy corrections are applied to Missing  $E_T$ . Since Missing  $E_T$  is susceptible to noise because it is summing over many channels, an alternative is often considered. This is Missing  $H_T$ , which is Missing  $E_T$  calculated by summing over the jets in the event rather than the calorimeter cells. Since there are fewer cells involved in the computation of missing  $H_T$ , there is less noise included in this sum.

It is typical to ask for two or more jets in the HLT algorithms. It is also common to combine two or more jets with missing  $E_T$  or  $H_T$ . Also, topological constraints are often employed such as requiring forward jets or acoplanarity between multiple jets or jets and Missing  $E_T$ .

### 12.2.6 Selection of Hadronic Tau Decays

The Level-2 processing of tau jets relies only on calorimeter information. In ATLAS, the tau finding uses the same algorithms used for electron and photon candidates, but retuned for taus. The inputs are the Level-1 RoIs. A cluster summed over the full resolution data for the electromagnetic and hadronic cells is required to have  $E_T > 20$  GeV with at least 10 GeV required individually in the electromagnetic and hadronic cells. The position of the candidate cluster is required to be consistent with the Level-1 tau-jet candidate. Then shower shape variables are used to discriminate tau jets from regular jets. An example of one such variable is  $R_{37}$ , defined as the ratio of  $E_T$  contained in a  $3 \times 7$  cell cluster to the  $E_T$  contained in a  $7 \times 7$  cell cluster centred on the same seed cell calculated for the second electromagnetic layer of the LAr calorimeter. In CMS, the Level-1 tau jets are used as seeds for the Level-2 tau-jet reconstruction that employs an iterative cone algorithm with a radius of  $R = 0.5$ . Level-2 tau candidates are then these jets which have  $E_T > 15$  GeV and are tagged as isolated if the sum of electromagnetic calorimeter deposits in an annulus  $0.13 < R < 0.4$  around the jet direction,  $E_T < 5$  GeV.

The subsequent processing of tau candidates involves tracking. ATLAS requires a track formed from the pixel and SCT detector space points in the RoI to be within  $\Delta R < 0.3$  of the Level-2 tau candidate cluster direction. At Level-3 a requirement is made that the number of tracks within  $\Delta R < 0.3$  be either one or three. Additional detailed jet shape requirements also refine the identification. In CMS, at Level 2.5 (the higher level trigger processing following the initial Level-2 processing that uses calorimeter and muon information alone), tau selection is based on tracks with a  $p_T > 5$  GeV/c that are reconstructed from seeds from the pixel hits found in a small rectangle ( $\Delta\eta = \Delta\phi = 0.1$ ) around the tau-candidate direction. At Level 3 the rectangle is expanded to 0.5 and the  $p_T$  cut is reduced to 1 GeV/c. To save CPU time, these tracks are terminated when seven hits in the silicon strip tracker are acquired since the resolution with seven hits is close to final. Reconstructed tracks are associated with the tau-jet candidate if they are within a radius  $R < 0.5$  and originate from the primary vertex as determined by the pixel tracks. Tracks within a radius  $R < 0.1$  of the tau-jet candidate direction are classed as tau tracks. The leading tau track must have  $p_T > 3$  GeV and there must be no reconstructed tracks within an annulus  $0.07 < R < 0.3$  around this track.

### 12.2.7 Selection of $b$ -Jets

The  $b$ -jet selection is based on track reconstruction to tag displaced vertices associated with the jet. In ATLAS at Level-2,  $b$ -tagging uses reconstructed tracks from the silicon tracker within the Level-1 jet RoI. For each of these tracks the significance of the transverse impact parameter is computed and its error is

parameterized as a function of  $p_T$ . A b-jet discriminator is constructed using the likelihood ratio method to determine for each track in the jet the ratio of probability densities for the track to come from a b-jet or a u-jet. In CMS, Level-2 starts with events with 1, 2, 3 or 4 jets passing various thresholds or a high total  $E_T$  for the whole event. At Level 2.5 tracks are reconstructed using only pixel hits (at least three required), which are used to reconstruct the primary vertex. The b-tag algorithm runs on the four highest  $E_T$  jets with  $E_T > 35$  GeV and uses the pixel tracks and primary vertex to tag jets as b-jets if they have at least two tracks with a signed 3D impact parameter with large significance. Events pass Level 2.5 if they have at least one b-tagged jet. At Level 3, tracks of up to eight hits are reconstructed in a cone of size  $\Delta R = 0.25$  around the b-tagged jets. The level-3 filter selects events where there is at least one jet having at least two tracks with large impact parameter significance.

## 12.3 Outlook

Trigger and DAQ requirements will further evolve in the next decade with large increases in luminosity and the associated pile-up. ALICE will continuously read out the majority of its detectors with different latencies, busy times and technologies, differently optimized for pp, pA and AA running scenarios [26]. Triggered readout will be used by some detectors and for commissioning and some calibration runs. LHCb will run trigger-free at 30 MHz, reading every bunch crossing with inelastic collisions [27].

A major upgrade to the LHC, the HL-LHC [28], is planned to start in the middle of this decade and deliver a luminosity of  $5\text{--}7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$  at the LHC design centre of mass energy of 14 TeV, which corresponds to a pile-up of 140–200 at 25 ns bunch spacing. Present link technologies operable in the radiation and magnetic field environments of their inner detectors do not allow ATLAS and CMS to adopt a “triggerless” architecture with an acceptable detector power and material budget for their tracking detectors. Therefore, at the HL-LHC, both ATLAS and CMS will retain architectures with Level 1 triggers.

In order to maintain Run-2 physics sensitivity at the HL-LHC, ATLAS and CMS will add L1 tracking triggers for identification of tracks associated with calorimeter and muon trigger objects and will also feature a significant increase of L1 rate, L1 latency and HLT output rate. Additionally, ATLAS and CMS are also studying the use of fast timing information in the L1T. The ATLAS experiment will divide its L1T into two stages [29]. A L0 trigger with a rate of 1 MHz and latency of 6  $\mu\text{s}$  will use calorimeter and muon trigger information to produce seeds used with tracking and more fine-grained calorimeter and muon trigger information in the L1 trigger with an output rate of 400 kHz and latency of 30  $\mu\text{s}$ . This is processed by the HLT with an output storage rate of 5–10 kHz. The CMS L1T latency will increase to 12.5  $\mu\text{s}$  with an output range of 500–750 kHz for pileup ranging between 140 and 200 [30]. It will use an un-seeded L1 Track trigger along with finer granularity



calorimeter and muon triggers. The CMS HLT output rate to storage will range between 5 and 7.5 kHz for pileup ranging between 140 and 200.

The hardware implementations of the HL-LHC ATLAS and CMS L1T will use high-bandwidth serial I/O links for data communication and large, modern field-programmable gate arrays (FPGAs) for sophisticated and fast algorithms. The development and synthesis of FPGA firmware incorporating these algorithms is significantly enhanced in reliability, accessibility and performance with Higher Level Synthesis (HLS) tools [31]. The latest developments and expectations for future FPGAs not only include significant increases in the number of logic gates available and high-speed serial links, but also increases in the number of high-bandwidth serial links per device, more sophisticated and fast DSPs, embedded Linux, and integration with high speed networking. Fast Tracking Trigger devices such as the ATLAS FTK [25] use Associative Memories. The hardware framework will be designed following standards deployed in industry, such as the Advanced Telecommunications Architecture (ATCA) for backplanes, which offers substantial backplane bandwidth and flexibility and provides for users to extend the backplane connectivity using the spare I/O available on each card. Further interconnectivity technology developments such as optical backplanes and wireless data transmission may provide additional opportunities.

The increase in L1 output rate from 100 kHz to possibly as high as 1 MHz requires higher bandwidth into the DAQ system and more CPU power in the HLT. The addition of a tracking trigger and more sophisticated algorithms at L1 increases the purity of the sample of events passing the L1 trigger, but requires a higher sophistication and complexity of algorithms used at the HLT. This implies a greater CPU power than scaling with the L1 output rate but is somewhat mitigated by the availability of the L1 Tracking Trigger primitives in the data immediately accessible by the HLT. Without a L1 tracking trigger, the opportunity to access most of the tracker information at the first levels of the HLT is limited by the CPU time to unpack and reconstruct the tracking data. This is significantly improved in the ATLAS FTK that provides quick access to tracking information in the HLT. For the HL-LHC, the addition of the L1 tracking trigger means that the results from the L1T track reconstruction can be immediately used without the overhead of tracking data unpacking and reconstruction.

The evolution of the computing market towards different computing platforms and co-processors offers an opportunity to achieve substantial gains in HLT processing power at the price of adapting code to the new hardware. Examples include Graphical Processor Units (GPUs), such as the NVIDIA Tesla and GeForce (used by ALICE [32]), ARM processors, FPGAs (e.g. the Xeon/FPGA used by LHCb [33]) and the Intel Xeon Phi coprocessor. Additional HLT processing power may result from improved code such as machine learning algorithms for track reconstruction [34].

## References

1. M. J. Woudstra and the ATLAS Collaboration, Performance of the ATLAS muon trigger in pp collisions at  $\sqrt{s} = 8$  TeV, 2014 J. Phys.: Conf. Ser. 513 012040.
2. E. Aslanides et al., (LHCb Collaboration), The Level-0 muon trigger for the LHCb experiment, Nucl. Instr. and Meth. A579 (2007) 989–1004.
3. B. Forestier, Nucl. Instrum. Meth. A 533 (2004) 22–26.
4. R. Downing et al., Nucl. Instrum. Meth. A 570 (2007) 36–50.
5. E.J. Thompson et al., *Online Track Processor for the CDF Upgrade*. IEEE Trans. Nucl. Sci. 49(3) (2002).
6. Y. Maravin et al., *First Results from the Central Tracking Trigger of the DØ Experiment*. Proc. 2003 IEEE Nuclear Science Symposium, Portland, OR.
7. R. Aaij et al., (LHCb collaboration), *The LHCb Trigger and its Performance in 2011*, 2013 JINST 8 P04022.
8. M. Krivda et al. (ALICE Collaboration), *The ALICE trigger system performance for p-p and Pb-Pb collisions* 2012 JINST 7 C0105.
9. I. A. Christidi and the ATLAS Collaboration, *The tracking performance of the ATLAS High Level Trigger in pp collisions at the LHC*, 2011 J. Phys.: Conf. Ser. 331, 032006.
10. The CMS Collaboration, *The Trigger and Data Acquisition project*, Vol. II: *Data Acquisition and the Higher Level Trigger*. CERN/LHCC 2007-021.
11. Will Buttinger (ATLAS collaboration) *The ATLAS Level-1 Trigger System*, 2012 J. Phys.: Conf. Ser. 396 012010.
12. D. Acosta et al. (CMS Collaboration), *CMS Trigger Improvements Towards Run II*, Nuclear and Particle Physics Proceedings 273–275 (2016) 1008–1013.
13. The ATLAS Collaboration, *Performance of the ATLAS Trigger System in 2010*, Eur. Phys. J.C 72, 1849 (2012).
14. V. Khachatryan et al., (CMS Collaboration) *The CMS Trigger*, 2017 JINST 12 P01020.
15. The ATLAS Collaboration, *Performance of the ATLAS Trigger system in 2015*, Eur. Phys. J. C (2017) 77: 317.
16. L. Cadamuro et al. (CMS Collaboration), *The CMS Level-1 trigger system for LHC Run II*, 2017 JINST 12 C03021.
17. B. Sciascia (LHCb Collaboration), *LHCb Run 2 trigger performance*, PoS BEAUTY2016 (2016) 029.
18. The ATLAS TDAQ Collaboration, *The ATLAS Data Acquisition and High Level Trigger system* 2016 JINST 11 P06008.
19. The CMS Collaboration, *The CMS Trigger*, 2017 JINST 12 P01020.
20. W. H. Smith, *Triggering at the LHC*, Ann. Rev. Nucl. Part. Sci. 66, 123 (2016).
21. CMS collaboration, *Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II*, 2016 JINST 11 C02008.
22. ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*, CERN-LHCC-2015-020; CMS Collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, CERN-LHCC-2015-010.
23. M. zur Nedden (on behalf of the ATLAS Collaboration), *The LHC Run 2 ATLAS trigger system: design, performance and plans*, 2017 JINST 12 C03024.
24. M. Krivda et al. (The ALICE Collaboration), *The ALICE Central Trigger Processor (CTP) Upgrade* 2016 JINST 11 C03051.
25. Asbah, N. & ATLAS collaboration, *A hardware fast tracker for the ATLAS trigger*, Phys. Part. Nuclei Lett. (2016) 13: 527.
26. ALICE Collaboration, *Upgrade of the ALICE Readout & Trigger System*, CERN-LHCC-2013-019; M. Krivda and J. Pospisil for the ALICE Collaboration, *The ALICE Central Trigger Processor (CTP) Upgrade*, JINST 11, C03051 (2016); F Costa et al., *The detector read-out in ALICE during Run 3 and 4*, J. Phys.: Conf. Ser. 898 032011 (2017).

27. LHCb Collaboration, LHCb Trigger and Online Upgrade Technical Design Report, CERN-LHCC-2014-016; T. Szumlak, Real time analysis with the upgraded LHCb trigger in Run III, J. Phys.: Conf. Ser. 898 032051 (2017).
28. Apollinari, G. (ed.) et al., HL-LHC Preliminary Design Report, CERN-2015-005; P. Campana, M. Klute and P.S. Wells, *Physics Goals and Experimental Challenges of the Proton–Proton High-Luminosity Operation of the LHC*, Ann. Rev. Nucl. Part. Sci. 66, 273 (2016).
29. ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*, CERN-LHCC-2015-20 (2015).
30. CMS Collaboration, The Phase-2 Upgrade of the CMS Level-2 Trigger, CERN-LHCC-2017-013 (2017); CMS Phase-II Upgrade Scope Document, CERN-LHCC-2015-19 (2015).
31. N.P. Ghanathe, et al., *Software and firmware co-development using high-level synthesis*, JINST 12 C01083 (2017).
32. S. Gorbunov et al., *ALICE HLT High Speed Tracking on GPU*, in IEEE Transactions on Nuclear Science, 58, no. 4, 1845 (2011).
33. C. Färber, R. Schwemmer, J. Machen and N. Neufeld, *Particle identification on an FPGA accelerated compute platform for the LHCb upgrade*, 2016 IEEE-NPSS Real Time Conference (RT), Padua, 2016, pp. 1–2.
34. S. Farrell et al., *The HEP.TrkX Project: deep neural networks for HL-LHC online and offline tracking*, EPJ Web of Conferences 150, 00003 (2017).

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 13

## Pattern Recognition and Reconstruction



R. Frühwirth, E. Brondolin, and A. Strandlie

### 13.1 Track Reconstruction

#### 13.1.1 Introduction

Track reconstruction is the task of finding and estimating the trajectory of a charged particle, usually embedded in a static magnetic field to determine its momentum and charge. It involves pattern recognition algorithms and statistical estimation methods. Depending on the physics goals, not all charged tracks have to be reconstructed. For instance, in many cases there is a physically motivated lower limit on the momentum or transverse momentum of the particles to be found. Other examples are short-range secondary particles, such as  $\delta$ -electrons, that normally need not be reconstructed. It may also be useful to reconstruct electron-positron pairs from photon conversions in order to check the distribution of material in the detector. Track reconstruction frequently proceeds in several steps:

1. Pattern recognition or Track finding: Finds the detector signals (hits) that are generated by the same charged particle.
2. Track fitting: Estimates for each track candidate the track parameters and the associated covariance matrix.
3. Test of track hypothesis: Tests for each track candidate whether all hits do indeed belong to the track and identifies outliers.

---

R. Frühwirth (✉) · E. Brondolin

Institute of High Energy Physics of the Austrian Academy of Sciences, Vienna, Austria

e-mail: [Rudolf.Fruehwirth@oeaw.ac.at](mailto:Rudolf.Fruehwirth@oeaw.ac.at); [erica.brondolin@cern.ch](mailto:erica.brondolin@cern.ch)

A. Strandlie

NTNU—Norwegian University of Science and Technology, Gjøvik, Norway

e-mail: [are.strandlie@ntnu.no](mailto:are.strandlie@ntnu.no)

There are many different algorithms for track finding. A selection of them is described in Sects. 13.1.2.1 and 13.1.2.2. For an extended treatment of the subject, containing many examples, see the excellent exposition in [1]. The track fit takes a track candidate and estimates the track parameters (location, direction, momentum or curvature, see Sect. 13.1.3.2) from the detector hits (Sect. 13.1.3.4), taking into account the equation of motion (Sect. 13.1.3.2) in the magnetic field (Sect. 13.1.3.1) and the effects of the detector material on the trajectory (Sect. 13.1.3.3). In the test stage (Sect. 13.1.3.5) outliers are identified, i.e., hits which apparently do not belong to the track. If outliers are expected, the estimation procedure should be robust so that the estimated track is not significantly biased by the outliers. Some robust methods are discussed in Sect. 13.1.3.5. Section 13.1.4 treats track-based alignment, and Sect. 13.1.5 contains many useful formulas for determining the approximate momentum resolution of a tracking detector without extensive simulations.

## 13.1.2 *Pattern Recognition*

Pattern recognition or track finding methods can be divided into global and local methods. In a global method, all detector hits are treated on an equal footing, and all track candidates are found in parallel; in a local method, there is a privileged subset of hits which is used to find initial track candidates, which are then completed to full track candidates.

### 13.1.2.1 **Global Methods**

Typical global methods of track finding find the tracks in parallel, for instance by identifying peaks in a one- or two-dimensional histogram, or by observing the final state of a recurrent neural network.


#### Conformal Mapping

A popular method for finding circular particle tracks is the conformal mapping method [2]. It uses the fact that the mapping

$$u = \frac{x}{x^2 + y^2}, \quad v = \frac{y}{x^2 + y^2},$$

transforms circles going through the origin of an  $x$ - $y$  coordinate system into straight lines of the form

$$v = \frac{1}{2b} - u \frac{a}{b},$$



**Fig. 13.1** The original measurements (left) and the transformed measurements (right) of six circular tracks

where the parameters  $a$  and  $b$  are defined by the circle equation


$$(x - a)^2 + (y - b)^2 = R^2 = a^2 + b^2.$$

The distance of the line to the origin is equal to  $1/(2R)$ , so that for large radius  $R$  it passes very close to the origin. The lines can be found by a histogramming method. After transforming the measurements in the  $u-v$  plane to polar coordinates and collecting the polar angle  $\theta$  in a histogram, measurements belonging to the same particle will tend to create peaks in the histogram.

As an example, the measured points of six circular tracks are shown in the left hand panel of Fig. 13.1, while the transformed measurements are shown in the right hand panel of Fig. 13.1. The resulting histogram of the polar angle  $\theta$  is shown in Fig. 13.2.

### Hough Transform

In the general case of lines not passing close to the origin, a more general approach is needed in order to find the lines. A very popular method for this purpose is the Hough transform [3]. The principle of the Hough transform can be explained by noting that a straight line in an  $x-y$  coordinate system,  $y = cx + d$ , can also be regarded as a straight line in a  $c-d$  coordinate system by the transformation  $d = -xc + y$ . For a fixed point  $(x, y)$ , the line in  $c-d$  space (also denoted parameter space) corresponds to all possible lines going through this point in  $x-y$  space (also denoted image space). Measurements lying along a straight line in image space therefore transform into lines in parameter space which cross at the specific value of the parameters of the line under consideration in image space. In practice, parameter space is discretized, and each measurement  $(x, y)$  leads to



**Fig. 13.2** Histogram of  $\theta = \arctan(v/u)$

an increment of a set of histogram bins. Measurements lying along straight lines tend to create peaks in the histogram, and the lines can be found by searching for peaks in this histogram. The granularity of the discretization has to be optimized for each specific application, as it depends on the amount of noise present and the actual values of measurement uncertainties. A too fine-grained histogram can split or destroy peaks if the measurement uncertainties are non-negligible. On the other hand, a too course-grained histogram increases the sensitivity to noise, and nearby tracks may merge into a single peak.

The basic formulation of the Hough transform is an example of a divergent transform, i.e., one measurement in image space corresponds to a set of increments of histogram entries in parameter space. The Hough transform can also be made convergent by considering instead a pair of measurements in image space. A unique line passes through any such pair, and only one entry in the parameter space histogram needs to be incremented. A possible disadvantage of such an approach is that the number of pairs grows quadratically with the number of measurements in image space. In order to reduce computational complexity, one may consider only a randomly selected subset of all the pairs. This is the basic feature of probabilistic Hough transforms [4].

The Hough transform has turned out to be successful also for finding circles passing through the origin. With this constraint, two parameters are enough to uniquely describe the circle, and the task again amounts to finding peaks in a two-dimensional histogram. With three or more parameters, one has to search for clusters in multi-dimensional spaces, and in this case the Hough transform is in general less powerful than in the two-dimensional case.

For track finding in drift tubes, with their inherent left-right ambiguity, the drift circles can be transformed to sine curves in the  $(r, \theta)$  space by applying a Legendre transform [5]. The peaks at the intersections of several sine curves represent the common tangents to a set of several circles.

## Neural Networks

Recurrent neural networks of the Hopfield type [6] are used in finding solutions to certain kinds of combinatorial optimization problems, i.e., problems that can be formulated as finding the minimum of an energy function

$$E = -\frac{1}{2} \sum_i \sum_j T_{ij} S_i S_j$$

with respect to the configuration of  $n$  binary-valued neurons  $S_i, i = 1, \dots, n$  and fixed connection weights  $T_{ij}, i, j = 1, \dots, n$ . It was realized independently in [7] and [8] that the track finding problem can be formulated as a minimization problem of this kind. The neurons are links between measurements which potentially belong to the same track. The connection weights  $T_{ij}$  have a structure which favors links sharing a measurement and pointing in a similar direction. The standard network dynamics leads to a solution corresponding to a local minimum of the energy function. A better solution is to apply a mean-field annealing technique [9], which introduces a temperature parameter and thereby allows the neurons to take all values in the interval between the two original binary values. The network is initialized at a high temperature, the mean-field equations are iteratively solved as the network is cooled down, and the low-temperature limit is taken in the end. At a significantly lower computational effort, the approximate solutions found by the mean-field technique have been shown to be very close to the exact solutions [10]. For applications of the Hopfield network in experiments see e.g. [11–15].

The energy function of the Hopfield network can be generalized in order to take into account the track model (see Sect. 13.1.3.2), i.e., the known parametric form of the tracks. The resulting algorithm is called elastic tracking or elastic arms [16–19]. A related generalization is the elastic net, originally used to tackle the traveling salesman problem [20]. Applications to track finding are described in [21] and [22, 23].

### 13.1.2.2 Local Methods

A local track finding method finds the tracks sequentially, starting from an initial track segment or an initial collection of measured points.



## Track Road

The track road method starts out with a set of measurements that potentially belong to the same track, typically one close to the vertex area, one far out in the tracking detector, and one in the middle. The track model can then be used, either exactly or approximately, if speed is an important issue, to interpolate between the measurements and create a road around the hypothesized track. Measurements inside the road are then collected. The number of measurements and the quality of the subsequent track fit are used to determine whether the track candidate should be kept or discarded.


## Track Following

A track following procedure takes a track seed as a starting point. A seed is often a short track segment, potentially including a constraint of the position of the vertex region. Seeds can be generated at the inner part of the tracking detector, where the measurements frequently are of very high precision, or at the outer part, where the track density is lower. From the seed, the track is extrapolated to the next detector unit. As for the track roads method, this can be done either with the full track model or with an approximate, simplified model. The measurement closest to the predicted track is included in the track candidate, and the track is extrapolated again.

## Kalman Filter

The Kalman filter [24–26] can be regarded as a statistically optimal track following procedure. It works by alternating prediction and update steps. Starting from the seed, the track parameters and their covariance matrix are extrapolated to the next detector unit containing a measurement, using the full track model. If the measurement is compatible with the prediction, it is included in the track candidate, and the track parameters and their covariance matrix are updated with the information from the measurement. The procedure is repeated until too many detector units without compatible measurements are traversed or the end of the tracking detector is reached.

In the original formulation of the method, the measurement closest to the predicted track is included in the track candidate [27]. However, if the density of measurements is high, the closest measurement might originate from another particle or from noise in the detector electronics. Including the wrong measurement could therefore lead to a wrong subsequent prediction and ultimately to the loss of the track. The currently most popular approach, the combinatorial Kalman filter, avoids such losses by splitting the track candidate into several branches when several compatible measurements are found after the prediction [28]. In order to take into account detector inefficiencies an additional branch with a missing hit can be generated.



**Fig. 13.3** An example of the combinatorial Kalman filter (reprinted from R. Mankel [28], with permission from Elsevier)

All branches are extrapolated to the next detector layer containing compatible measurements. A branch is split again if several measurements are compatible with the branch prediction. Branches are removed if too many detector units without compatible measurements are traversed or if the quality of the track candidate, in terms of the value of a  $\chi^2$  statistic, is too low. If there are several surviving candidates after the end of the detector has been reached, the candidate with most measurements and the lowest value of the  $\chi^2$  statistic is kept and regarded as the final track candidate. An example is shown in Fig. 13.3.

A similar track finding method has been formulated in the language of cellular automata [23, 29]. The combinatorial problem can also be solved by using generalized, adaptive versions of the Kalman filter [30–32].

### 13.1.3 Estimation of Track Parameters

#### 13.1.3.1 Magnetic Field Representation

The presence of a magnetic field in a tracking detector causes a bending of the trajectory of a charged particle, and, hence, allows a measurement of the particle momentum. A precise knowledge of the magnetic field is therefore crucial for accurate estimates of the particle momenta.

The magnetic field can be calculated by solving Maxwell's equations, knowing the detailed configuration of the current sources and the magnetic materials in the detector volume. In the general case, a numerical solution of these equations in terms of a finite-element analysis is needed. In special cases, the field can be found by less general approaches. The simplest situation is a solenoidal magnet, providing a homogeneous field in a large volume. Also, it is known that the field inside a volume with no magnetic material can be determined by knowledge of the field on the volume boundary only [33]. Measurements of the field on the volume boundary allows an estimation of coefficients of polynomials obeying Maxwell's

equations. Field measurements inside the volume are used to evaluate the quality of the calculated field. If the measurements inside the volume are precise enough, they can be used to further refine the knowledge of the field by being included in the estimation procedure of the abovementioned coefficients [34].

In a track reconstruction application, fast access to the value of the magnetic field at any point inside the detector volume is crucial. For this purpose, a numerical representation of the field is needed. A frequently used approach is to create a table of the magnetic field values at a grid of points and to determine the field at points between the grid nodes by linear or quadratic interpolation. An alternative approach is to divide the detector volume into several sub-volumes and to fit the coefficients of low-order polynomials to the known field values inside each sub-volume [35, 36]. If the number of sub-volumes is large, potentially many coefficients have to be determined. On the other hand, once the coefficients are determined, the field access is very fast. Also, the derivatives of the field, which are needed by some track reconstruction algorithms, can be computed as fast as the field itself.

### 13.1.3.2 Track Models

Consider a charged particle with mass  $m$  and charge  $Q = qe$ ,  $e$  being the elementary charge. Its trajectory  $\mathbf{x}(t)$  in a magnetic field  $\mathbf{B}(\mathbf{x})$  is determined by the equations of motion given by the Lorentz force  $\mathbf{F} \propto q\mathbf{v} \times \mathbf{B}$ , where  $\mathbf{v} = d\mathbf{x}/dt$  is the velocity of the particle. In vacuum, Newton's second law reads [37]

$$\frac{d\mathbf{p}}{dt} = kq\mathbf{v}(t) \times \mathbf{B}(\mathbf{x}(t)), \quad (13.1)$$

where  $\mathbf{p} = \gamma m\mathbf{v}$  is the momentum of the particle,  $\gamma = (1 - \mathbf{v}^2/c^2)^{-1/2}$  is the Lorentz factor, and  $k$  is a unit-dependent proportionality factor. If  $\mathbf{p}$  is in GeV/c,  $\mathbf{x}$  is in meters, and  $\mathbf{B}$  is in Tesla,  $k = 0.29979 \text{ GeV/c T}^{-1} \text{ m}^{-1}$ . The trajectory is uniquely defined by the initial conditions, the six degrees of freedom specified for instance by the initial position and the initial velocity. If these are tied to a surface, five degrees of freedom are necessary and sufficient. Geometrical quantities other than position and velocity can also be used to specify the initial conditions. The collection  $\mathbf{q}$  of these quantities is called the initial track parameters or the initial state vector.

Equation (13.1) can be written in terms of the path length  $s(t)$  along the trajectory instead of  $t$ , giving [37]

$$\frac{d^2\mathbf{x}}{ds^2} = \frac{kq}{|\mathbf{p}|} \cdot \frac{d\mathbf{x}}{ds} \times \mathbf{B}(\mathbf{x}(s)) = F(s, \mathbf{x}(s), \dot{\mathbf{x}}(s)). \quad (13.2)$$

In simple situations this equation has analytical solutions. In a homogeneous magnetic field the trajectory is a helix; it reduces to a straight line in the limit of a vanishing field. In the general case of an inhomogeneous field, numerical

methods can be used, such as Runge–Kutta integration of the equations of motion or parametrization by polynomials or splines [37]. Among Runge–Kutta methods, the Runge-Kutta-Nyström algorithm is specially designed for second-order equations such as Eq. (13.2). In the fourth-order version a step of length  $h$ , starting at  $s = s_n$ , is computed by [37]

$$\mathbf{x}_{n+1} = \mathbf{x}_n + h\dot{\mathbf{x}}_n + h^2(k_1 + k_2 + k_3)/6, \quad \dot{\mathbf{x}}_{n+1} = \dot{\mathbf{x}}_n + h(k_1 + 2k_2 + 2k_3 + k_4)/6,$$

with

$$\begin{aligned} k_1 &= F(s_n, \mathbf{x}_n, \dot{\mathbf{x}}_n), \\ k_2 &= F(s_n + h/2, \mathbf{x}_n + h\dot{\mathbf{x}}_n/2 + h^2k_1/8, \dot{\mathbf{x}}_n + hk_1/2), \\ k_3 &= F(s_n + h/2, \mathbf{x}_n + h\dot{\mathbf{x}}_n/2 + h^2k_1/8, \dot{\mathbf{x}}_n + hk_2/2), \\ k_4 &= F(s_n + h, \mathbf{x}_n + h\dot{\mathbf{x}}_n + h^2k_3/2, \dot{\mathbf{x}}_n + hk_3), \end{aligned}$$

where  $\mathbf{x}_n$  is the position of the particle at  $s = s_n$  and  $\dot{\mathbf{x}}_n$  is the unit tangent vector. The magnetic field needs to be looked up for the calculation of  $k_2$ ,  $k_3$  and  $k_4$ , i.e., three times per step. If the field at the final position  $\mathbf{x}_{n+1}$ , which is the starting position of the next step, is approximated by the field used for  $k_4$ , only two lookups are required per step. If the field is (almost) homogeneous, as for example in a solenoid, the step size  $h$  can be chosen to be constant; otherwise a variable step size is more efficient. The step size can be optimized using an adaptive version of the Runge-Kutta-Nyström algorithm [38]. Note that the error of a step of length  $h$  may be larger than  $O(h^5)$  if the magnetic field does not have smooth derivatives, as is the case if it is computed by linear interpolation. If the field is represented by low-order polynomials in sub-volumes, Runge–Kutta steps should terminate at volume boundaries.

Different detector geometries often lead to different choices of the parametrization. However, the parametrization of the trajectory should comply to some basic requirements: the parameters should be continuous with respect to small changes of the trajectory; the choice of track parameters should facilitate the local expansion of the track model into a linear function; and the uncertainties of the estimated values of the parameters should follow a Gaussian distribution as closely as possible. For example, curvature should be used rather than radius of curvature, and inverse (transverse) momentum rather than (transverse) momentum.

The track model, given by the solution of the equations of motion, describes how the state vector  $\mathbf{q}_k$  at a given surface  $k$  depends on the state vector at a different surface  $i$ :

$$\mathbf{q}_k = \mathbf{f}_{k|i}(\mathbf{q}_i),$$

where  $\mathbf{f}_{k|i}$  is the track propagator from surface  $i$  to surface  $k$ . When analytical solutions of the equations of motion exist, the track propagator is also analytical. Even in

a homogeneous magnetic field, the path length can be determined analytically only for propagation to cylinders with symmetry axis parallel to the field direction or to planes orthogonal to the field direction. Otherwise, a Newton iteration or a parabolic approximation has to be used to find the path length.

For track reconstruction purposes, the covariance matrix of the estimated track parameters needs to be propagated along with the track parameters themselves. The track propagator is often a non-linear function of the track parameters at the initial surface, but the covariance matrix has to be transported under the assumption of a linear track model. This procedure, called linear error propagation, is based on a Taylor expansion of the track propagator, keeping only first-order terms. These first-order terms, defining the Jacobians of the track model, are given by

$$F_{k|i} = \left. \frac{\partial \mathbf{q}_k}{\partial \mathbf{q}_i} \right|_{\tilde{\mathbf{q}}_i},$$

where  $\tilde{\mathbf{q}}_i$  is the expansion point in surface  $i$ . For analytical track models, the Jacobian is also analytical. If, for example, the magnetic field is homogeneous, the general case of propagation to a plane of arbitrary spatial orientation uses a curvilinear coordinate frame moving along with the trajectory as a means of deriving the required Jacobians [39].

In the general case of a non-analytical track model, the Jacobians cannot be computed analytically either. The most straightforward approach is to calculate the relevant derivatives in a purely numerical way. The basis for these calculations is a reference trajectory corresponding to the expansion point. In addition, five other trajectories are created, corresponding to small variations in each of the track parameters. By propagating these five trajectories to the destination surface, numerical derivatives can be obtained. A potential disadvantage of such an approach is its computational complexity, as six trajectories have to be propagated instead of a single one. Much less computational load is introduced by transporting the Jacobian terms in parallel to the track parameters during the Runge–Kutta integration [40, 41], avoiding the need for propagating auxiliary trajectories.

The measurement model describes the functional dependence of the measured quantities on the state vector at a detector surface  $k$ :

$$\mathbf{m}_k = \mathbf{h}_k(\mathbf{q}_k).$$

The vector of measurements  $\mathbf{m}_k$  usually contains the measured coordinates, but may contain also other quantities, e.g. measurements of direction or even momentum. In a pixel detector or in a double-sided silicon strip detector,  $\mathbf{m}_k$  is two-dimensional; in a one-sided strip detector, it is one-dimensional. In a drift chamber or a multi-wire proportional chamber with several layers, the measurement may be a track segment resulting from an internal track reconstruction. In this case the vector  $\mathbf{m}_k$  may be four- or five-dimensional, depending on whether the curvature can be estimated or not.

In most cases the function  $\mathbf{h}_k(\mathbf{q}_k)$  includes a transformation of the state vector  $\mathbf{q}_k$  into the local coordinate system of the detector. For use in track reconstruction, the Jacobian of this transformation is needed:

$$\mathbf{H}_k = \left. \frac{\partial \mathbf{m}_k}{\partial \mathbf{q}_k} \right|_{\check{\mathbf{q}}_k}, \quad (13.3)$$

where  $\check{\mathbf{q}}_k$  is the expansion point in surface  $k$ . In many cases the Jacobian contains only rotations and projections, and thus can be computed analytically.

The measurement is always smeared by a measurement error:

$$\mathbf{m}_k = \mathbf{h}_k(\mathbf{q}_k) + \boldsymbol{\varepsilon}_k.$$

The mean value and the covariance matrix of  $\boldsymbol{\varepsilon}_k$  depend on the detector type and the detector geometry and have therefore in general to be calibrated for each detector unit independently. The measurement error is often assumed to follow a Gaussian distribution, but frequently exhibits tails which are incompatible with this assumption. In this case a Gaussian mixture is a more appropriate model.

### 13.1.3.3 Material Effects

A charged particle crossing a tracking detector interacts with the material of the detector. The most important types of interactions in track reconstruction are multiple Coulomb scattering, energy loss by ionization, and energy loss by bremsstrahlung. For an in-depth treatment of material effects see Chapter 2.

#### Multiple Coulomb Scattering

Elastic Coulomb scattering of particles heavier than the electron is dominated by the atomic nucleus. For small angles the differential cross-section is approximately equal to

$$\frac{d\sigma}{d\theta} = 2\pi \left( \frac{2Ze^2}{pv} \right)^2 \frac{1}{\theta^3},$$

where  $\theta$  is the polar angle of the scattering,  $Z$  is the charge of the nucleus in units of the elementary charge  $e$ ,  $v$  is the velocity of the scattered particle, and  $p$  is its momentum [42]. Because of screening effects and the finite size of the nucleus the differential cross-section is modified to [43]

$$\frac{d\sigma}{d\theta} = 2\pi \left( \frac{2Ze^2}{pv} \right)^2 \frac{\theta}{(\theta^2 + \theta_{\min}^2)^2}, \quad 0 \leq \theta \leq \theta_{\max}.$$

If the momentum  $p$  is given in GeV/ $c$ , the lower and upper limits are approximately equal to

$$\theta_{\min} \approx \frac{2.66 \cdot 10^{-6} Z^{1/3}}{p}, \quad \theta_{\max} \approx \frac{0.14}{A^{1/3} p}.$$

The average number of scattering processes in a layer of thickness  $d$  (in cm) is given by

$$N(d) = d\sigma \frac{N_A \rho}{A},$$

where  $\sigma$  is the integrated elastic cross section,  $N_A$  is the Avogadro constant,  $\rho$  is the density of the material (in g/cm<sup>3</sup>), and  $A$  is the atomic mass of the nucleus. In track reconstruction it is convenient to work with the projected scattering angles in two perpendicular planes. The projected multiple scattering angle  $\theta_P$  is equal to the sum of the projected single scattering angles, and its variance can be obtained by multiplying the variance of the projected single scattering angle by the average number of scatters, the projected single scattering angles being uncorrelated. With increasing thickness  $d$  the distribution of the projected scattering angle approaches a normal distribution, and the two projected angles become independent. For thin scatterers, however, the width of the Gaussian core is notably narrower than is indicated by the variance [42]. This is taken into account by Highland's formula for the standard deviation of the projected scattering angle [44]:

$$\sigma_P = E(\theta_P^2)^{1/2} = \frac{0.0136}{\beta p} \sqrt{d/X_0} [1 + 0.038 \ln(d/X_0)],$$

where  $X_0$  is the radiation length of the material in cm,  $\beta = v/c$  is the particle velocity in units of  $c$ , and  $p$  is the particle momentum in GeV/ $c$ . The logarithmic correction ceases to be applicable above  $d \approx X_0$ .

If a scatterer is sufficiently thin, the transverse offset of the track due to multiple scattering can be neglected. Only the track direction is affected in this case. If the direction is represented by the polar angle  $\theta$  and the azimuthal angle  $\varphi$ , their joint covariance matrix is given by

$$\text{var}(\Delta\theta) = \sigma_P^2, \quad \text{var}(\Delta\varphi) = \sigma_P^2 / \sin^2 \theta, \quad \text{cov}(\Delta\theta, \Delta\varphi) = 0.$$

If the direction is represented by the direction tangents  $t_x = dx/dz$  and  $t_y = dy/dz$ , the covariance matrix is [45]

$$\text{Var}[(\Delta t_x, \Delta t_y)^T] = \sigma_P^2 (1 + t_x^2 + t_y^2) \begin{pmatrix} 1 + t_x^2 & t_x t_y \\ t_x t_y & 1 + t_y^2 \end{pmatrix}.$$

If the direction is represented by the direction cosines  $c_x = dx/ds$  and  $c_y = dy/ds$ , the covariance matrix is [45]

$$\text{Var}[(\Delta c_x, \Delta c_y)^T] = \sigma_P^2 \begin{pmatrix} (1 - c_x)^2 & -c_x c_y \\ -c_x c_y & (1 - c_y)^2 \end{pmatrix}.$$

In all cases the projected variance  $\sigma_P^2$  takes into account the effective amount of material crossed by the track.

If the transverse offset cannot be neglected, its variance and its correlation with the angle have to be taken into account. Assume that the particle passes a scatterer of length  $d$ , traveling along the  $z$ -axis. Neglecting the curvature of the track, the joint covariance matrix of the offset  $\Delta x$  and the scattering angle  $\theta_x$  in the  $x$ - $z$  projection is

$$\text{var}(\Delta x) = \sigma_0^2 d^3/3, \quad \text{var}(\theta_x) = \sigma_0^2 d, \quad \text{cov}(\Delta x, \theta_x) = \sigma_0^2 d^2/2,$$

where  $\sigma_0^2$  is the variance of the projected scattering angle per unit length. If the particle enters the scatterer at  $z = 0$  with direction  $(t_x, t_y)$ , the joint covariance matrix of the offsets  $\Delta x, \Delta y$  and the angles  $\Delta t_x, \Delta t_y$  at  $z = d$  is

$$\text{Var} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta t_x \\ \Delta t_y \end{bmatrix} = \sigma_0^2 (1 + t_x^2 + t_y^2) \begin{pmatrix} (1 + t_x^2)D^3/3 & t_x t_y D^3/3 & (1 + t_x^2)D^2/2 & t_x t_y D^2/2 \\ t_x t_y D^3/3 & (1 + t_y^2)D^3/3 & t_x t_y D^2/2 & (1 + t_y^2)D^2/2 \\ (1 + t_x^2)D^2/2 & t_x t_y D^2/2 & (1 + t_x^2)D & t_x t_y D \\ t_x t_y D^2/2 & (1 + t_y^2)D^2/2 & t_x t_y D & (1 + t_y^2)D \end{pmatrix},$$

where  $D = (1 + t_x^2 + t_y^2)^{1/2} d$  is the effective thickness crossed. If the direction is represented by  $\theta$  and  $\varphi$ , the covariance matrix can be computed via the transformation

$$t_x = \tan(\theta) \cos(\varphi), \quad t_y = \tan(\theta) \sin(\varphi),$$

and linear error propagation with the Jacobian

$$\mathbf{T} = \frac{\partial(\theta, \varphi)}{\partial(t_x, t_y)} = \begin{pmatrix} \frac{t_x}{\sqrt{t_x^2 + t_y^2}(1 + t_x^2 + t_y^2)} & \frac{t_y}{\sqrt{t_x^2 + t_y^2}(1 + t_x^2 + t_y^2)} \\ -\frac{t_y}{t_x^2 + t_y^2} & \frac{t_x}{t_x^2 + t_y^2} \end{pmatrix} = \begin{pmatrix} \frac{\cos(\varphi)}{1 + \tan^2(\theta)} & \frac{\sin(\varphi)}{1 + \tan^2(\theta)} \\ -\frac{\sin(\varphi)}{\tan(\theta)} & \frac{\cos(\varphi)}{\tan(\theta)} \end{pmatrix}.$$

For analogous formulas in cylindrical coordinates, see [46].

If the curvature of the track cannot be neglected, the simplest approach is a stepwise integration of the equation of motion, assuming the validity of a helical track model within each step and considering each such step as a thin scatterer [39, 47].



## Energy Loss

For particles other than electrons the energy loss in material is almost exclusively due to scattering on electrons. The momentum correction  $\Delta p$  in a material layer of thickness  $d$  is calculated by integrating the Bethe-Bloch formula [37]:

$$\Delta p = \int_0^d \frac{dp}{dx} dx = \int_0^d \frac{1}{\beta} \frac{dE}{dx} dx = \int_0^d \frac{K}{\beta^3} \left[ \ln \frac{2m_e c^2 \beta^2 \gamma^2}{\langle I \rangle} - \beta^2 \right] dx, \quad (13.4)$$

where  $K$  is a constant depending on the material,  $m_e$  is the electron mass,  $\langle I \rangle$  is the average ionization potential of the material, and  $\beta = v/c$  and  $\gamma = E/mc^2$  are the usual kinematic parameters. The ratio  $\langle I \rangle/Z$  is about 20 eV for hydrogen and helium, between 12 and 16 eV for light nuclei, and around 10 eV for heavy nuclei [44]. For practical purposes, the differential energy loss  $dE/dx$  is a function only of  $\beta$ . For small  $\beta$ , it decreases like  $1/\beta^2$ . It has a minimum, the position of which drops with increasing  $Z$  from  $\beta\gamma \approx 3.5$  (carbon) to  $\beta\gamma \approx 3$  (lead). In terms of momentum, the minimum is at  $p = \beta\gamma mc$  and thus depends on the mass of the particle. This dependency is used for particle identification. The energy loss at the minimum can be parameterized for  $Z \geq 6$  by [44]:

$$(dE/dx)_{\min} = (2.35 - 0.64 \ln_{10} Z) \text{ MeV g}^{-1} \text{ cm}^2.$$

From this the constant  $K$  in Eq. (13.4) can be calculated. For large  $\beta\gamma$  the energy loss increases like  $\ln(\beta\gamma)$ ; this is called the relativistic rise. For momenta in the vicinity of the minimum  $dE/dx$  can be considered as constant, giving  $\Delta p \approx (dE/dx)_{\min} \cdot d\rho/\beta$ ,  $\rho$  being the density of the material.

## Bremsstrahlung

For an electron (or positron) passing through matter the most significant contribution to energy loss is bremsstrahlung, the emission of photons in the electric field of an atomic nucleus. In the Bethe-Heitler model [48] the relative energy loss is distributed independently of the energy. Let  $d$  be the path length in the material in units of radiation length, and  $z$  the fraction of energy remaining after the material is traversed. Then the distribution of  $z$  is given by the following probability density function:

$$f(z) = \frac{(-\ln z)^{c-1}}{\Gamma(c)}, \quad 0 \leq z \leq 1,$$

where  $\Gamma(x)$  is Euler's gamma function and  $c = d/\ln 2$ . For high energy electrons  $p \approx E$ , so the momentum correction is  $\Delta p \approx p(z - 1)$ . The first two moments of  $\Delta p$  are

$$E(\Delta p) = p(2^{-c} - 1), \quad \text{var}(\Delta p) = p^2(3^{-c} - 4^{-c}).$$

The moments can be used for a Gaussian representation of bremsstrahlung as an additional process noise in the Kalman filter (see Sect. 13.1.3.4). As this is a very crude approximation, more sophisticated methods have been developed that take into account the actual shape of the distribution. One of them is the Gaussian-sum filter [49, 50], see Sect. 13.1.3.5. A computationally less intensive approach is described in [51].

### 13.1.3.4 Estimation Methods

The main task of the track fit is to estimate the values of a set of parameters describing the state of a particle somewhere in the detector, often at a reference surface close to the interaction vertex. The information from the measurements created by the particle while traversing the tracking detector should be processed in an optimal manner. If the track model is truly linear, i.e., if the measurements are strictly linear functions of the track parameters, and all stochastic disturbances entering the estimation procedure are Gaussian, the linear least-squares method is the optimal one [37]. Since track parameter propagation in general is a nonlinear procedure, strict linearity holds very rarely in practice. The relation between the track parameter vector  $\mathbf{q}_0$  at a reference surface and the measurement vector  $\mathbf{m}_k$  at a detector layer  $k$  is a function  $\mathbf{d}_k$  given by

$$\mathbf{m}_k = \mathbf{d}_k(\mathbf{q}_0) + \boldsymbol{\gamma}_k,$$

where  $\boldsymbol{\gamma}_k$  is a noise term containing the measurement error of  $\mathbf{m}_k$  and all multiple scattering in front of  $\mathbf{m}_k$ . The function  $\mathbf{d}_k$  is a composition of the measurement model function  $\mathbf{h}_k$  and the track propagator functions  $\mathbf{f}_{i|i-1}$  (see Sect. 13.1.3.2):

$$\mathbf{d}_k = \mathbf{h}_k \circ \mathbf{f}_{k|k-1} \circ \cdots \circ \mathbf{f}_{2|1} \circ \mathbf{f}_{1|0}.$$

For the linear least-squares method  $\mathbf{d}_k$  has to be linearized around some expansion point, providing the Jacobian  $\mathbf{D}_k$  of each  $\mathbf{d}_k$ :

$$\mathbf{D}_k = \mathbf{H}_k \mathbf{F}_{k|k-1} \cdots \mathbf{F}_{2|1} \mathbf{F}_{1|0},$$

with  $\mathbf{H}_k$  from Eq. (13.3). The covariance matrix of  $\boldsymbol{\gamma}_k$  is obtained by linear error propagation:

$$\text{var}(\boldsymbol{\gamma}_k) = \mathbf{V}_k + \mathbf{H}_k (\mathbf{F}_{k|1} \mathbf{Q}_1 \mathbf{F}_{k|1}^T + \cdots + \mathbf{F}_{k|k-1} \mathbf{Q}_{k-1} \mathbf{F}_{k|k-1}^T + \mathbf{Q}_k) \mathbf{H}_k^T,$$

where  $\mathbf{V}_k$  is the covariance matrix of the measurement error  $\boldsymbol{\varepsilon}_k$  of  $\mathbf{m}_k$ , and  $\mathbf{Q}_j$  is the covariance matrix of multiple scattering after layer  $j-1$  up to and including layer  $j$ . The part of  $\mathbf{Q}_j$  originating from scattering between the layers has to be transported to layer  $j$  by the appropriate Jacobian. Because of the cumulative effect of multiple

scattering  $\boldsymbol{y}_i$  and  $\boldsymbol{y}_k$  are correlated. If  $i < k$ , the covariance is given by

$$\text{cov}(\boldsymbol{y}_i, \boldsymbol{y}_k) = \boldsymbol{H}_i (\boldsymbol{F}_{i|1} \boldsymbol{Q}_1 \boldsymbol{F}_{k|1}^T + \cdots + \boldsymbol{F}_{i|i-1} \boldsymbol{Q}_{i-1} \boldsymbol{F}_{k|i-1}^T + \boldsymbol{Q}_i \boldsymbol{F}_{k|i}^T) \boldsymbol{H}_k^T.$$

The observations  $\boldsymbol{m}_k$ , the functions  $\boldsymbol{d}_k$ , their Jacobians  $\boldsymbol{D}_k$ , and the noise  $\boldsymbol{y}_k$  are now collected in single vectors and a matrix:

$$\boldsymbol{m} = \begin{pmatrix} \boldsymbol{m}_1 \\ \vdots \\ \boldsymbol{m}_n \end{pmatrix}, \quad \boldsymbol{d} = \begin{pmatrix} \boldsymbol{d}_1 \\ \vdots \\ \boldsymbol{d}_n \end{pmatrix}, \quad \boldsymbol{D} = \begin{pmatrix} \boldsymbol{D}_1 \\ \vdots \\ \boldsymbol{D}_n \end{pmatrix}, \quad \boldsymbol{y} = \begin{pmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_n \end{pmatrix},$$

where  $n$  is the total number of measurement layers. This gives the following model:

$$\boldsymbol{m} = \boldsymbol{d}(\boldsymbol{q}_0) + \boldsymbol{y},$$

which now can be linearized into

$$\boldsymbol{m} = \boldsymbol{D}\boldsymbol{q}_0 + \boldsymbol{c} + \boldsymbol{y},$$

where  $\boldsymbol{c}$  is a constant vector. The global least-squares estimate of  $\boldsymbol{q}_0$  is given by

$$\tilde{\boldsymbol{q}}_0 = (\boldsymbol{D}^T \boldsymbol{G} \boldsymbol{D})^{-1} \boldsymbol{D}^T \boldsymbol{G} (\boldsymbol{m} - \boldsymbol{c}),$$

where  $\boldsymbol{V} = \boldsymbol{G}^{-1}$  is the non-diagonal covariance matrix of  $\boldsymbol{y}$ . The quality of the initial expansion point can be monitored by using the obtained estimate as a new expansion point, and the state vector estimate can hence be re-calculated. Such a procedure is repeated until convergence, defined by a suitable stopping criterion.

If the track model is a circle and multiple scattering and energy loss can be neglected, the estimation can be simplified substantially. Explicit estimators are given in [52] for the center and radius of the circle, and in [53] for the curvature, the direction and the distance from a fixed point. Other algorithms are based on conformal mapping in the plane [2] or on a mapping to the Riemann sphere [54–56].

If there is strong multiple scattering, the estimated track can be quite far away from the real track. In order to follow the actual track more closely, two projected scattering angles can be explicitly estimated at each detector layer or at a set of virtual breakpoints inside a continuous scatterer [45, 57]. The breakpoint method, also known as General Broken Lines [58], and the global least-squares method are equivalent, as far as the estimate of the state vector  $\boldsymbol{q}_0$  is concerned [59].

If the number of measurements or the number of breakpoints is substantial, the computational cost of these methods can be high due to the necessity of inverting large matrices during the estimation procedure. The Kalman filter, a recursive formulation of the least-squares method, requires the inversion of only

small matrices and exhibits the same attractive feature as the breakpoint method of following the actual track quite closely [26, 60].

As mentioned earlier, the Kalman filter proceeds by alternating prediction and update steps. The prediction step is the propagation of the track parameter vector from one detector layer containing a measurement to the next,

$$\mathbf{q}_{k|k-1} = \mathbf{f}_{k|k-1}(\mathbf{q}_{k-1|k-1}),$$

and the associated covariance matrix,

$$\mathbf{C}_{k|k-1} = \mathbf{F}_{k|k-1} \mathbf{C}_{k-1|k-1} \mathbf{F}_{k|k-1}^T + \mathbf{Q}_k.$$

The update step is the correction of the predicted state vector due to the information from the measurement in layer  $k$ :

$$\mathbf{q}_{k|k} = \mathbf{q}_{k|k-1} + \mathbf{K}_k [\mathbf{m}_k - \mathbf{h}_k(\mathbf{q}_{k|k-1})],$$

where the gain matrix  $\mathbf{K}_k$  is given by

$$\mathbf{K}_k = \mathbf{C}_{k|k-1} \mathbf{H}_k^T (\mathbf{V}_k + \mathbf{H}_k \mathbf{C}_{k|k-1} \mathbf{H}_k^T)^{-1}.$$

The update of the covariance matrix is given by

$$\mathbf{C}_{k|k} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{C}_{k|k-1}.$$

The information filter is a mathematically equivalent, but numerically more stable formulation of the Kalman filter. In the information filter, the update of the state vector reads

$$\mathbf{q}_{k|k} = \mathbf{C}_{k|k} \left[ (\mathbf{C}_{k|k-1})^{-1} \mathbf{q}_{k|k-1} + \mathbf{H}_k^T \mathbf{V}_k^{-1} \mathbf{m}_k \right],$$

whereas the update of the covariance matrix is given by

$$\mathbf{C}_{k|k} = \left[ (\mathbf{C}_{k|k-1})^{-1} + \mathbf{H}_k^T \mathbf{V}_k^{-1} \mathbf{H}_k \right]^{-1}.$$

The implementation of the Kalman filter requires the computation of the Jacobians  $\mathbf{F}_{k|k-1}$  and  $\mathbf{H}_k$ . A compilation of analytical formulas for two important cases (fixed-target configuration and solenoidal configuration) is given in [61].

Full information of the track parameters at the end of the track is obtained when all  $n$  measurements in the track candidate have been processed by the filter. The full information can be propagated back to all previous estimates by another iterative procedure, the Kalman smoother. A step of the smoother from layer  $k + 1$  to layer

$k$  is for the state vector

$$\mathbf{q}_{k|n} = \mathbf{q}_{k|k} + \mathbf{A}_k(\mathbf{q}_{k+1|n} - \mathbf{q}_{k+1|k}),$$

where the smoother gain matrix is given by

$$\mathbf{A}_k = \mathbf{C}_{k|k} \mathbf{F}_{k+1|k}^T (\mathbf{C}_{k+1|k})^{-1}.$$

The smoothed covariance matrix is

$$\mathbf{C}_{k|n} = \mathbf{C}_{k|k} - \mathbf{A}_k(\mathbf{C}_{k+1|k} - \mathbf{C}_{k+1|n})\mathbf{A}_k^T.$$

The smoother can also be realized by combining two filters running in opposite directions: a forward filter from  $\mathbf{m}_1$  to  $\mathbf{m}_n$  and a backward filter from  $\mathbf{m}_n$  to  $\mathbf{m}_1$ . The smoothed states are the weighted mean of the predicted states of one filter and the updated states of the other filter. This approach is numerically more stable than the gain matrix formulation of the smoother.

### 13.1.3.5 Track Quality and Robust Estimation

Robust estimators are insensitive to outliers, i.e., measurements that are biased or do not originate from the particle creating the majority of the hits in a track candidate. Some estimators are inherently robust by construction; other estimators can be made robust by finding and discarding outliers.

In the Kalman filter, the residual of the measurement in layer  $k$  with respect to the updated state vector is

$$\mathbf{r}_{k|k} = \mathbf{m}_k - \mathbf{h}_k(\mathbf{q}_{k|k}),$$

and the covariance matrix of this residual is

$$\mathbf{R}_{k|k} = \mathbf{V}_k - \mathbf{H}_k \mathbf{C}_{k|k} \mathbf{H}_k^T.$$

The chi-square increment in layer  $k$  is

$$\chi_{k,+}^2 = \mathbf{r}_{k|k}^T \mathbf{R}_{k|k}^{-1} \mathbf{r}_{k|k},$$

and the total chi-square of the track is found by summing up the chi-square increments for all measurements in the track candidate. The total chi-square is used to evaluate the quality of the track candidate. A too large value of this test statistic indicates that one or more of the measurements of the track candidate do not originate from the particle creating the majority of the measurements. Such measurements are called outliers.

An outlier rejection procedure can make use of the chi-squares of the measurements with respect to the smoothed predictions, i.e., a weighted mean of the predicted states of a forward and a backward Kalman filter. The measurement with the largest value of the chi-square is removed, and the total chi-square is again calculated. This procedure is repeated until the value of the total chi-square falls below a defined threshold.

In the presence of a potentially large fraction of outliers in a track candidate, the sequential outlier rejection procedure outlined above might become unstable, because the smoothed predictions may themselves be biased by outliers. An alternative approach is the Gaussian-sum filter [62]. This algorithm is based on the assumption that the probability distribution of the measurement error can be modeled as a two-component Gaussian mixture, where a narrow component represents the hypothesis that the measurement is real and a wider component represents the hypothesis that the measurement is an outlier. It takes the form of a set of Kalman filters running in parallel, each Kalman filter representing a specific hypothesis of a subset of the measurements that should be classified as outliers. A weight attached to each Kalman filter can be interpreted as the probability of correctness of the hypothesis. In the end, the Kalman filter with the largest weight or a weighted mean of the different filters can be taken as the final estimate.

The Gaussian-sum filter can also be used to deal with a mixture model of the process noise, i.e., the stochastic disturbance of the track because of interactions with the detector material [63]. In the case of bremsstrahlung, a successful application to the reconstruction of electrons is described in [50].

For the treatment of outliers, the Gaussian-sum filter has two disadvantages. First, it may create a large number of Kalman filters running in parallel because of poor knowledge of the track parameters in the early stages of the filter, making the approach expensive in terms of computing time. Second, an explicit outlier model is required. A faster and even more robust alternative is the Deterministic Annealing Filter [64]. This filter is an iterated Kalman filter with annealing, which assigns small weights to measurements far away from the track. A temperature parameter is introduced, facilitating convergence to the globally optimal solution. The iterations start at a high temperature, continue with a gradual lowering of the temperature and converge at the nominal value of the temperature. The procedure is easily generalized to the situation of several measurements being present in the same detector layer. In this case the measurements compete for inclusion in the track. As opposed to a standard outlier rejection approach, the assignment of measurements is soft. This means that several measurements in the same detector layer might contribute to the final estimate of the track parameters, each with a weight equal to the assignment probability. A further generalization is the multi-track filter, where several tracks are allowed to compete for compatible hits in all detector layers [65]. For an experimental application, see [66].

13.1.3.6 Jet Reconstruction

Jets are bundles of collimated hadrons, reflecting hard scattering processes at the parton level. In order to carry out detailed comparisons between parton-level predictions and hadron-level observations a well-defined “jet finder” is required. In the jet finding information from both the tracking devices and the calorimeters is used.

Jet finding can be understood as finding clusters in the set of reconstructed tracks, including neutral tracks. As in the case of vertex finding (see Sect. 13.2.2), various types of clustering methods have been proposed and investigated. The performance strongly depends on the underlying physics, and usually a jet finder is optimized for specific physics requirements. For instance, the widely used  $k_{\perp}$  clustering algorithm comes in several versions, for instance one for  $e^+e^-$  collisions [67], and one for hadron-hadron collisions [68].

Hierarchical cluster algorithms offer a large variety of jet finders, differing mainly by the definition of the measure of distance between objects (tracks and jets), but sometimes also by the order in which the objects are combined. Some examples of agglomerative clustering algorithms are described and studied in [69]. Table 13.1 gives a summary of the distance measures used. The names refer to the ones used in [69].  $E_i$  is the energy of cluster  $i$ ,  $p_i$  is its momentum,  $\theta_{ij}$  is the opening angle between the momentum vectors of the two clusters, and  $E_{\text{vis}}$  is the visible energy.

A divisive hierarchical clustering algorithm is described in [76]. It is based on the following measure of distance between two tracks:

$$d_{ij} = \frac{\theta_{ij}^2}{p_i p_j},$$

**Table 13.1** Some distance measures used for agglomerative jet finding with respective references

Name	Distance $d_{ij}$	References
Jade	$\frac{2 E_i E_j (1 - \cos \theta_{ij})}{E_{\text{vis}}^2}$	[70]
Durham, $k_{\perp}$	$\frac{2 \min(E_i^2, E_j^2) (1 - \cos \theta_{ij})}{E_{\text{vis}}^2}$	[67, 68, 71–73]
Luclus	$\frac{2 p_i^2 p_j^2 (1 - \cos \theta_{ij})}{(p_i + p_j)^2 E_{\text{vis}}^2}$	[74]
Geneva	$\frac{8 E_i E_j (1 - \cos \theta_{ij})}{9 (E_i + E_j)^2}$	[75]
Cambridge	$\frac{2 \min(E_i^2, E_j^2) (1 - \cos \theta_{ij})}{E_{\text{vis}}^2}$	[71]

but can be generalized to any other measure of distance. The method first constructs a minimum spanning tree [77] in the edge-weighted graph connecting all particles with each other and then proceeds to cut the tree along its longest edges. The procedure stops when the longest remaining edge is shorter than a fixed multiple of the median of all edge lengths.

Several non-hierarchical cluster algorithms have been proposed as well. Some of them employ general unsupervised learning methods, such as deterministic annealing [78] or  $k$ -means [79]. Others are specially designed for jet finding, for instance the cone algorithm described in [72]. It is an iterating procedure which constructs jets out of seeds. In contrast to the hierarchical clustering method the jets may overlap and a unique assignment has to be forced at the end. A modified cone algorithm suitable for the much larger multiplicity of heavy-ion collisions is proposed in [80]. A specialized jet finder for the reconstruction of hadronic  $\tau$ -decays is described in [81].

### 13.1.4 Detector Alignment<sup>1</sup>

Alignment is the general term used in experimental high energy physics to refer to the process of obtaining and applying corrections to the nominal setup of a given experiment. These corrections are typically related to geometrical displacements of devices with a spatial resolution, in contrast to calibrations, where the corrections are usually extracted from pedestal or reference measurements to compensate for offsets in scalar measurements. Misalignment compromises tracking and vertex finding [82] and thus directly affects physics measurements such as momentum and invariant mass resolutions, or the efficiency of  $b$ -tagging algorithms. There are various possibilities for the treatment of alignment corrections, ranging from simple translations and rotations, equivalent to those of a rigid body, to more complex deformations, like sags or twists.

To this end experiments typically use several independent strategies [83]. For testing the long-term stability or the alignment of sub-detectors with respect to each other, very often so-called hardware alignment is utilized, where special reference markers are measured directly e.g. via optical systems or photogrammetry. However, these techniques reach only a limited precision in the range of several tens to hundreds of microns. If the intrinsic resolution of a tracking device is smaller, an improved resolution can only be obtained with track-based alignment, where the information from recorded particle tracks is used to obtain the alignment parameters [83, 84]. For various examples of the track-based alignment methods used in experiments since the LEP era, see [85–100].

---

<sup>1</sup>The section on detector alignment was contributed by E. Widl (Institute of High Energy Physics, Vienna; now at Austrian Institute of Technology).



### 13.1.4.1 General Overview

The basis of all track-based alignment algorithms is an extended track model  $\mathbf{d}$ , where the measurements  $\mathbf{m}$  depend not only on the true track-parameters  $\mathbf{q}_0$ , but also on a set of alignment parameters  $\mathbf{p}_0$  that describe the effects of sufficiently small deviations from the ideal geometry:

$$\mathbf{m} = \mathbf{d}(\mathbf{q}_0, \mathbf{p}_0) + \boldsymbol{\gamma}, \quad \text{cov}(\boldsymbol{\gamma}) = \mathbf{V}.$$

The stochastic term  $\boldsymbol{\gamma}$ , which describes the intrinsic resolution of the tracking devices and the effects of multiple scattering, is dealt with via its covariance matrix  $\mathbf{V}$ . Since typically high momentum particles are used, energy-loss effects can be assumed to be deterministic and hence directly taken care of in the track model  $\mathbf{d}$  itself.

With an initial guess  $\check{\mathbf{q}}$  for the track parameters and  $\check{\mathbf{p}}$  for the alignment parameters, this model allows to define residuals that are functions of the unknowns  $\mathbf{q}$  and  $\mathbf{p}$ :

$$\mathbf{r}(\mathbf{q}, \mathbf{p}) = \mathbf{m} - \mathbf{d}(\mathbf{q}, \mathbf{p}) \approx \mathbf{m} - \check{\mathbf{d}} - \mathbf{D}_q \Delta \mathbf{q} - \mathbf{D}_p \Delta \mathbf{p} \quad (13.5)$$

with

$$\check{\mathbf{d}} = \mathbf{d}(\check{\mathbf{q}}, \check{\mathbf{p}}), \quad \Delta \mathbf{q} = \mathbf{q} - \check{\mathbf{q}}, \quad \Delta \mathbf{p} = \mathbf{p} - \check{\mathbf{p}},$$

$$\mathbf{D}_q = \left. \frac{\partial \mathbf{d}}{\partial \mathbf{q}} \right|_{\check{\mathbf{q}}, \check{\mathbf{p}}}, \quad \mathbf{D}_p = \left. \frac{\partial \mathbf{d}}{\partial \mathbf{p}} \right|_{\check{\mathbf{q}}, \check{\mathbf{p}}}.$$

The goal of a track-based alignment algorithm is to determine  $\mathbf{p}$  from the residuals  $\mathbf{r}$ , by minimizing the quadratic form  $\chi^2 = \mathbf{r}^T \mathbf{V}^{-1} \mathbf{r}$ , using a sufficiently large set of recorded tracks. The methods used are quite diverse, but can be grouped into two categories: biased and unbiased algorithms.

Biased algorithms initially ignore the fact that the initial guess of the track parameters  $\mathbf{q}_0$  is in general biased by the factual misalignment. In other words, by setting  $\mathbf{q} = \check{\mathbf{q}}$  for every track, the residuals become a function of  $\mathbf{p}$  alone, i.e.,  $\mathbf{r}(\mathbf{q}, \mathbf{p}) \rightarrow \mathbf{r}(\mathbf{p})$ . In general, the influence of the biased track information has to be compensated by iterating several times over the track sample, where at each iteration step the previously determined parameters are applied to the track reconstruction.

Unbiased algorithms on the other hand, minimize the residuals or the normalized residuals, respectively, estimating at the same time the track parameters. The problem with such an approach is the resulting huge number of parameters. In the presence of  $N$  alignment parameters and a sample of  $M$  tracks with  $m$  track parameters each, a total of  $N + m \cdot M$  parameters have to be dealt with. While the value of  $N$  depends on the experimental setup, and  $m$  usually equals 5, the number of tracks  $M$  has always to be of considerable size to acquire reasonable statistics.

On the other hand, unbiased algorithms usually do not require iterations, with the possible exception of problems like non-linearities or rejection of outliers.

Besides the differences between various algorithms it should be noted that the final result of any track-based alignment is always limited by the tracks used. Basic quality cuts, like the selection of high momentum tracks to minimize the influence of multiple scattering or cuts on the minimum number of hits, have a strong influence on the convergence. More subtle is the effect of an unbalanced mixture of tracks or the complete absence of some types of tracks, such as tracks from collisions and cosmic events or tracks taken with and without a magnetic field. This is due to the fact that any kind of tracks has several unconstrained degrees of freedom, usually referred to as weak modes, weakly defined modes or  $\chi^2$ -invariant modes. As an example, typical weak modes for straight tracks are shears but not bends, and vice versa for curved tracks. Combining the information of both kinds of tracks is therefore a reasonable strategy to avoid these deformations in the final result. The most obvious weak mode is a translation or rotation of the entire tracking device, which can be only fixed with some kind of reference frame, be it an external system or by definition. This, however, is less severe and sometimes even not considered at all, as it does not affect the internal alignment of the tracking device.

Once a set of alignment parameters is calculated, it should always be validated [83, chapter 11]. Apart from checking the improvement of the residuals, several physics measurements can be utilized, especially to probe for remaining weak modes. Known charge, forward-backward or  $\varphi$ -symmetries of distinct physics processes can be used. Distributions of the signed curvature or the signed transverse impact parameter are also sensitive observables.

### 13.1.4.2 Examples of Alignment Algorithms

Some modern experiments deploy large tracking devices that require a large number of alignment parameters, of the order of  $10^5$ . In such a case the computation of parameters by using straightforward recipes might become unreasonably slow or cause numerical problems. The two algorithms presented in this section are examples of how to cope with such challenging circumstances.

#### The HIP Algorithm

The HIP algorithm [101] is a straightforward and easy-to-implement biased alignment algorithm. It computes the alignment parameters for each alignable object separately. Only when iterating on the track sample a certain kind of indirect feedback between the alignable objects is established due to the track refit.

Since only individual alignable objects are regarded, Eq. (13.5) can be partitioned. This is simply done by evaluating the corresponding expressions for each alignable object  $i$  together with its associated parameters  $p_i$ :

$$r_i(p_i) = m_i - d_i(\check{q}, p_i) \approx m_i - \check{d}_i - D_{pi} \Delta p_i$$

with

$$\check{d}_i = d_i(\check{q}, \check{p}_i), \quad \Delta p_i = p_i - \check{p}_i, \quad D_{pi} = \left. \frac{\partial d_i}{\partial p_i} \right|_{\check{q}, \check{p}_i}.$$

The result is determined by minimizing the normalized squared residuals from a given set of tracks, again for each alignable object separately. The formal solution is given by

$$\Delta p_i = \left( \sum_{\text{tracks}} D_{pi}^T V_i^{-1} D_{pi} \right)^{-1} \left( \sum_{\text{tracks}} D_{pi}^T V_i^{-1} r_i(\check{p}_i) \right)$$

### The Millepede Algorithm

The Millepede algorithm [102] is an unbiased algorithm that minimizes the sum of the squared residuals of all tracks at once. To this end a system of linear equations, equivalent to the formal solution of an ordinary  $\chi^2$ -fit, is solved. However, to achieve this within a reasonable amount of time, only the solution for the alignment parameters is computed, while the computation of the improved track parameters is skipped. This is possible because of the special structure of the system: Firstly, the coefficient matrix is symmetric and, mostly due to the independence of the individual tracks, relatively sparse. Secondly, only the alignment parameters are common parameters for all track measurements, while the specific track parameters are only relevant for each corresponding track. Due to the latter, the solutions for the alignment and track parameters are only coupled via coefficient matrices of the form

$$G = D_p^T V^{-1} D_q.$$

To set up the reduced system of equations, for each track the following information has to be extracted:

$$\Gamma = D_q^T V^{-1} D_q, \quad \beta = D_q^T V^{-1} (m - \check{d} - D_p \Delta p').$$

Here  $\Delta p' = p' - \check{p}$  may already include an estimate  $p'$  on the actual alignment. Then compute

$$\Delta C = D_p^T V^{-1} D_p - G \Gamma^{-1} G^T, \quad \Delta g = D_p^T V^{-1} (m - \check{d} - D_p \Delta p' + D_q \Gamma^{-1} \beta).$$

Note the expression  $-\Gamma^{-1}\beta$  instead of  $\Delta q$ . These are all necessary terms, including implicitly the full information from all track parameters. The complete system of equations to determine the alignment parameters then reads

$$C \Delta p = -g,$$

with

$$C = \sum_{\text{tracks}} \Delta C, \quad g = \sum_{\text{tracks}} \Delta g.$$

The solution by matrix inversion is only feasible if the number of parameters is fairly small ( $N \leq 10^3$ ). The matrix  $C$  is usually relatively sparse, so that less time-consuming and more reliable methods can be used, such as the GMRES algorithm [103].

It is also possible to introduce constraints into the solution, which allows to align on various hierarchical levels at once. When aligning for instance on module- and layer-level at the same time, these constraints can remove redundant degrees of freedom by forcing the average movement of all modules within one layer to zero.

Millepede is a well-tested algorithm. To use it efficiently, some knowledge of its inner workings is of advantage. The HIP algorithm is simpler to implement, but less suitable for very large setups than Millepede. Another unbiased algorithm is the Kalman Alignment Algorithm [104]. It is a sequential method, derived from the Kalman filter (see also [105]).

### 13.1.5 Momentum Resolution

The momentum resolution that can be achieved by a tracking detector is determined by the magnetic field, the arrangement and precision of the tracking detectors, and the amount of material crossed by the particle. Simple approximate formulas can be obtained for two cases:

- (a) A spectrometer consisting of a central bending magnet and two arms of tracking detectors in front of and behind the magnet. This is a typical arrangement for a fixed-target experiment with small track multiplicities.
- (b) A set of cylindrical tracking detectors immersed in a homogeneous magnetic field. This is a typical arrangement for the barrel part of a collider experiment, for instance layers of silicon or a TPC.

The units are the same as in Sect. 13.1.3.2: momentum in GeV/c, length in meters, and magnetic field in Tesla.

### 13.1.5.1 Two-arm Spectrometer

We assume that the trajectory of the particle is parallel to the  $z$  axis and that  $B_y$  is the only significant component of the magnetic field. The angle of deflection is then given by [37]

$$\alpha \approx -\frac{kq}{p} \int_L B_y dz = -\frac{kq}{p} \bar{B}_y L,$$

where  $L$  is the length of the magnet,  $\bar{B}_y$  is the average value of the field along the trajectory,  $p$  is the momentum, and  $q, k$  are as in Eq. (13.2). Assuming that  $|q| = 1$ , linear error propagation gives

$$\frac{\sigma(p)}{p} = \frac{p\sigma(\alpha)}{k|\bar{B}_y|L}.$$

Assume that each arm consists of  $m$  identical position detectors spread over a length  $l$ , and that the standard deviation of the measurement error of  $x$  is equal to  $\delta$ . The best angular resolution is obtained if in each arm half of the detectors is placed at each end of the arm. Neglecting all multiple scattering, it is equal to

$$\sigma(\alpha) = \frac{2\delta}{l\sqrt{m/2}}.$$

The relative momentum resolution due to measurement errors is therefore

$$\frac{\sigma_{\text{me}}(p)}{p} = \frac{2p\delta}{l\sqrt{m/2}k|\bar{B}_y|L}.$$

Although this arrangement optimizes the precision in terms of geometry, it offers little redundancy for track finding and should be used only in setups with trivial pattern recognition requirements, for instance in the forward direction of fixed target experiments.

At low energies, multiple scattering can no longer be neglected. Whereas  $\sigma(p)/p$  arising from position measurement errors only is proportional to  $p$ , the term  $\sigma_{\text{ms}}(p)/p$  arising from multiple scattering is proportional to  $1/(\beta|\bar{B}_y|L)$ , which is large for small  $\beta$  and constant for high momenta ( $\beta \approx 1$ ). Under the same assumptions about the detector positions as above, the following formula is obtained:

$$\frac{\sigma_{\text{ms}}(p)}{p} = \frac{0.0136}{\beta k|\bar{B}_y|L} \left( \frac{md}{X_0} \right)^{1/2},$$

where  $d/X_0$  is the thickness of the detectors in units of radiation length. The total resolution is obtained by adding the corresponding variances and taking the square root,

$$\frac{\sigma(p)}{p} = \frac{\sigma_{\text{me}}(p)}{p} \oplus \frac{\sigma_{\text{ms}}(p)}{p} = ap \oplus b,$$

with  $a$  and  $b$  depending on the detector and the magnetic field.

### 13.1.5.2 Cylindrical Spectrometer

Assume that there are  $m$  cylindrical detectors immersed in a homogeneous magnetic field  $B_z$  parallel to the  $z$  axis. The projection of the track on the  $x$ - $y$  plane is a circle with curvature  $\kappa$ . For high momentum the circle can be approximated by a parabola, the detector cylinders can be approximated by planes, and multiple scattering can be neglected. For this case closed formulas for the joint covariance matrix of  $\kappa$  and the tangent  $t_\varphi = \tan \varphi$  of the initial track direction  $\varphi$  can be given [106, 107]. For equidistant detectors, uniform resolution  $\delta$  and  $t_\varphi = 0$  it is given by

$$\text{Cov} \begin{pmatrix} t_\varphi \\ \kappa \end{pmatrix} = \frac{\delta^2}{m(m+1)(m+2)} \begin{pmatrix} \frac{12(m-1)(2m-1)(8m-11)}{L^2(m-2)} & -\frac{360(m-1)^3}{L^3(m-2)} \\ -\frac{360(m-1)^3}{L^3(m-2)} & \frac{720(m-1)^3}{L^4(m-2)} \end{pmatrix},$$

where  $L$  is now the track length in the  $x$ - $y$  projection.  $L$  is approximately equal to the radial distance between the innermost and the outermost detector. As  $\kappa = k B_z / p_T$ ,

$$\frac{\sigma_{\text{me}}(p_T)}{p_T} = \frac{p_T}{k|B_z|L} \frac{\delta}{L} \left[ \frac{720(m-1)^3}{(m-2)m(m+1)(m+2)} \right]^{1/2}.$$

There is a high negative correlation between  $1/p_T$  and the direction tangent  $t_\varphi$ . For large  $m$ , the asymptotic values are

$$\begin{aligned} \frac{\sigma_{\text{me}}(p_T)}{p_T} &= \frac{p_T}{k|B_z|L} \frac{\delta}{L} \left[ \frac{720}{m+4} \right]^{1/2}, \quad \sigma(t_\varphi) = \frac{\delta}{L} \left[ \frac{192}{m+3.875} \right]^{1/2}, \\ \text{cov}(t_\varphi, 1/p_T) &= -\frac{\sqrt{15}}{4} = -0.968. \end{aligned}$$

More general closed formulas for  $t_\varphi \neq 0$  are given in [107].

**Table 13.2** Values of  $C_m$  in Eq. (13.6)

$m$	3	4	5	6	7	8	9	10	> 10 m
$C_m$	1.16	1.06	1.04	1.03	1.02	1.02	1.01	1.01	$\approx 1$ m

If one half of the detectors is placed at the center of the track and one quarter at either end, the variance of the curvature is minimal, and the covariance matrix reads

$$\text{Cov} \begin{pmatrix} t_\varphi \\ \kappa \end{pmatrix} = \delta^2 \begin{pmatrix} \frac{72}{L^2 m} & -\frac{128}{L^3 m} \\ -\frac{128}{L^3 m} & \frac{256}{L^4 m} \end{pmatrix},$$

which is considerably smaller than in the equidistant case. This arrangement, however, is not particularly well suited for track finding and moreover difficult to realize.

The contribution of multiple scattering to the transverse momentum resolution can be approximated by

$$\frac{\sigma_{\text{ms}}(p_T)}{p_T} = C_m \cdot \frac{s}{\beta k |B_z| L} \left( \frac{md}{X_0 \cos \lambda} \right)^{1/2}, \quad (13.6)$$

where  $d/X_0$  is the thickness of the detectors in units of radiation length,  $\lambda = \pi/2 - \theta$  is the dip angle of the track,  $s = 0.0136(1 + 0.038 \ln(d/X_0))$ ,  $k$  is as in Eq. (13.2), and  $C_m$  is a factor depending on  $m$ . Values of  $C_m$  for small  $m$ , obtained by the program described in [108], are given in Table 13.2. Note that the values are different from the ones given in [106]. In a time projection chamber  $md/X_0$  has to be replaced by  $L/X_0$ , where  $X_0$  is the radiation length of the gas. The factor  $\cos \lambda$  in the denominator accounts for the actual amount of matter traversed by a track with dip angle  $\lambda$ . Approximate formulas for the best possible resolution including multiple scattering can be found in [109].

The total transverse momentum resolution is calculated by quadratic addition,

$$\frac{\sigma(p_T)}{p_T} = \frac{\sigma_{\text{me}}(p_T)}{p_T} \oplus \frac{\sigma_{\text{ms}}(p_T)}{p_T},$$

which can be written in the form

$$\frac{\sigma(p_T)}{p_T} = \frac{a p_T}{\sqrt{m+4}} \oplus \frac{b \sqrt{m}}{\sqrt{\cos \lambda}}.$$

This shows that an optimal  $m$  exists for every  $p_T$  and  $\lambda$  if the projected track length  $L$  is kept fixed. Overinstrumentation will deteriorate the resolution for low momenta unless additional measurements can be included without increasing the amount of matter to be traversed.

In order to calculate the error of the momentum  $p = p_T / \cos \lambda$  the error in  $\lambda$  must be taken into account:

$$\sigma^2(p) = \sigma^2(p_T) / \cos^2 \lambda + \sigma^2(\lambda) p_T^2 \sin^2 \lambda / \cos^4 \lambda,$$

the correlation between  $p_T$  and  $\lambda$  being negligible in practice. Because of  $\sigma(p)/p = p\sigma(1/p)$  it follows that:

$$\frac{\sigma(p)}{p} = \frac{\sigma(p_T)}{p_T} \oplus \sigma(\lambda) \tan \lambda.$$

With the exception of very low momenta the track can be approximated by a straight line in the  $r$ - $z$  projection, where  $r = (x^2 + y^2)^{1/2}$ . For  $m$  equidistant detectors and uniform resolution  $\delta$ , the variance of the direction tangent  $t_\lambda = \tan \lambda$  due to the measurement errors is given by [106]:

$$\sigma_{\text{me}}^2(t_\lambda) = \frac{\delta^2}{L^2} \frac{12(m-1)}{m(m+1)} \frac{1}{\cos^4 \lambda}.$$

If the measurement error in  $z$  is very small, the variance of  $t_\lambda$  is dominated by multiple scattering. For equidistant layers of uniform thickness  $d$  an approximate formula can be given. Remarkably, it does not depend on the number of layers:

$$\sigma_{\text{ms}}^2(t_\lambda) \approx \frac{s^2}{p^2} \frac{d}{X_0 \cos \lambda} \frac{1}{\cos^4 \lambda} = \frac{s^2}{p_T^2} \frac{d}{X_0 \cos \lambda} \frac{1}{\cos^2 \lambda},$$

with  $s = 0.0136(1 + 0.038 \ln(d/X_0))$ .

In the design and optimization phase of the detector a precise evaluation of the resolution of all track parameters is mandatory. There are several software packages that allow a fast track simulation plus reconstruction in a general detector setup, for instance [110] (in FORTRAN), [111] (in Matlab/Octave), or [108] (in Java).

## 13.2 Vertex Reconstruction

### 13.2.1 Introduction

Vertex reconstruction is the task of finding and estimating the production point of a set of particles. The pattern recognition algorithms and statistical estimation methods involved are in many respects similar to the ones used in track reconstruction. For an overview of vertex reconstruction algorithms used in past or active experiments see for instance [112–116].



In practice it is useful to distinguish between several types of vertices:

1. The primary vertex is the point of collision of two beam particles (in a collider experiment) or of a beam particle and a target particle (in a fixed-target experiment).
2. A secondary decay vertex is the point where an unstable particle decays in the detector volume or in the beam pipe. An example is the decay  $K_S^0 \rightarrow \pi^+ \pi^-$ .
3. A secondary interaction vertex is the point where a particle interacts with the material of the detector. Examples are bremsstrahlung, pair production, and inelastic hadronic interactions.

Vertex reconstruction frequently proceeds in several steps:

1. Vertex finding: Finds the tracks that belong to a common primary or secondary vertex.
2. Vertex fitting: Estimates for each vertex candidate the location of the common vertex and computes the associated covariance matrix.
3. Test of vertex hypothesis: Tests for each vertex candidate whether all tracks do indeed belong to the vertex and identifies outliers.
4. Update: Uses the vertex constraint to improve the location and momentum estimate of the tracks belonging to the vertex.
5. Kinematic fit: Kinematic constraints such as momentum and energy conservation are imposed on the mother and daughter particles of a vertex, and mass hypotheses are tested. Kinematic fits are most frequently applied to secondary decay vertices.

Vertex finding can be accomplished in many different ways. A few of them will be described in Sect. 13.2.2. The vertex fit takes a vertex candidate and estimates the vertex location from the estimated track parameters of the outgoing particles (Sect. 13.2.3). As a rule, only charged particles are used, but sometimes also neutral particles contribute to the vertex fit. In the test stage (Sect. 13.2.3.2) outliers are identified, i.e., particles that apparently do not belong to the estimated vertex. As this can lead to a different assignment of particles to vertices, it can be considered as a method of vertex finding. If outliers are expected, the estimation procedure should be robust so that the estimated vertex is not significantly biased by the outliers (Sect. 13.2.3.3). Kinematic constraints (Sect. 13.2.4) are usually imposed via Lagrange multipliers. By repeating the kinematic fit under various mass hypotheses of the mother and/or daughter particles the most likely mass assignment can be found out.

### 13.2.2 Vertex Finding

Vertex finding is the process of dividing the reconstructed tracks in an event into classes such that presumably all tracks in a class are produced at the same vertex. The primary vertex in an event is usually easy to find, especially if prior information

about its location is available (beam profile, target position). On the other hand, secondary decay vertices of short-lived decays are hard to find, as some of the decay products may also be compatible with the primary vertex. Vertex finding methods can be roughly divided in three main types: generic clustering algorithms, topological methods, and iterated estimators. The latter can be considered as a special divisive clustering method.

### 13.2.2.1 Clustering Methods

As mentioned above in the context of jet finding (see Sect. 13.1.3.6), clustering methods are based on a distance matrix or a similarity matrix of the objects to be classified. A cluster is then a group with small distances (large similarities) inside the group and large distances (small similarities) to objects outside the group. The distance measure reflects only the geometry of the tracks.

Various clustering methods have been evaluated in the context of vertex finding, of both the hierarchical and the non-hierarchical type [117]. Hierarchical clustering can be agglomerative or divisive. In agglomerative clustering each track starts out as a single cluster. Clusters are merged iteratively on the basis of a distance measure. The shortest distance in space between two tracks is peculiar insofar as it does not satisfy the triangle inequality: if tracks  $a$  and  $b$  are close, and tracks  $b$  and  $c$  are close, it does not follow that tracks  $a$  and  $c$  are close as well. The distance between two clusters of tracks should therefore be defined as the maximum of the individual pairwise distances, known as complete linkage in the clustering literature. Alternatively, the distance between two clusters can be the distance between the two vertices fitted from the clusters. Divisive clustering starts out with a single cluster containing all tracks. Further division of this cluster can be based on repeated vertex estimation with outlier identification (see Sect. 13.2.2.3). Examples of non-hierarchical clustering methods used in vertex finding are vector quantization, the  $k$ -means algorithm and deterministic annealing [113].

### 13.2.2.2 Topological Methods

A very general topological vertex finder was proposed in [118]. It is related to the Radon transform, which is a continuous version of the Hough transform used for track finding (Sect. 13.1.2.1). The search for vertices is based on a function  $V(\mathbf{v})$  which quantifies the probability of a vertex at location  $\mathbf{v}$ . For each track a Gaussian probability tube  $f_i(\mathbf{v})$  is constructed. The function  $V(\mathbf{v})$  is defined taking into account that the value of  $f_i(\mathbf{v})$  must be significant for at least two tracks:

$$V(\mathbf{v}) = \sum_{i=0}^n f_i(\mathbf{v}) - \frac{\sum_{i=0}^n f_i^2(\mathbf{v})}{\sum_{i=0}^n f_i(\mathbf{v})}$$

Due to the second term on the right-hand side,  $V(\mathbf{v}) \approx 0$  in regions where  $f_i(\mathbf{v})$  is significant for only one track. The form of  $V(\mathbf{v})$  can be modified to fold in known physics information about probable vertex locations. For instance,  $V(\mathbf{v})$  can be augmented by a further function  $f_0(\mathbf{v})$  describing the location and spread of the interaction point. In addition,  $V(\mathbf{v})$  may be modified by a factor dependent on the angular location of the point  $\mathbf{v}$ .

Vertex finding amounts to finding local maxima of the function  $V(\mathbf{v})$ . The search starts at the calculated maxima of the products  $f_i(\mathbf{v})f_j(\mathbf{v})$  for all track pairs. For each of these points the nearest maximum of  $V(\mathbf{v})$  is found. These maxima are clustered together to form candidate vertex regions. The final association of the tracks to the vertex candidates can be done on the basis of the respective  $\chi^2$  contributions or by an adaptive fit (see Sect. 13.2.3.3). In [119] the topological vertex finder was augmented by a procedure based on the concept of the minimum spanning tree of a graph.

### 13.2.2.3 Iterated Estimators

Vertex finding can also be accomplished by iterated vertex fits (see Sect. 13.2.3). The procedure can be summarized in the following way:

1. Fit one vertex with all tracks
2. Discard all incompatible tracks
3. Repeat step 1 with all discarded tracks

The iteration stops when no vertex with at least two tracks can be successfully fitted. Step 2 might itself be iterative, especially if the vertex fit is not robust, so that the incompatible tracks have to be removed sequentially. Iterative vertex finders based on a least-squares fit (Sect. 13.2.3.1) and an adaptive fit (Sect. 13.2.3.3) are implemented in the RAVE toolbox [120, 121].

### 13.2.3 Vertex Fitting

The input to the vertex fit is a vertex candidate, i.e., a set of estimated track parameters  $\{\tilde{\mathbf{q}}_1, \dots, \tilde{\mathbf{q}}_n\}$  located at one or more reference surfaces, along with their covariance matrices  $\{\mathbf{C}_1, \dots, \mathbf{C}_n\}$ . For instance, in the primary vertex fit in a collider experiment the reference surface may be the beam tube. If possible, the reference surface(s) should be chosen such that multiple scattering between the vertex and the location of the track parameters is negligible.

The parameters to be fitted are the vertex position  $\mathbf{v}$  and the track momenta  $\mathbf{p}_i$  at the vertex. The functional dependence of the track parameters on the vertex parameters requires a track model, which depends on the shape of the magnetic field in the vicinity of the vertex. If the field is homogeneous, the track model is a helix;

if the field is zero, the track model is a straight line. In other cases the track model may have to be computed numerically (see Sect. 13.1.3.2).

### 13.2.3.1 Least-Squares Methods

The conventional approach to estimating the vertex position is the minimization of some quadratic objective function, yielding a least-squares estimate. There are two main flavors of least-squares estimation in vertex fitting, constrained and unconstrained minimization. In the first case the vertex constraint is introduced into the objective function via a Lagrange multiplier, in the second case the constraint is implicit in the track model.

As an example, consider a vertex fit with  $n$  straight tracks. The  $n$  straight tracks originating from the common vertex  $\mathbf{v} = (x_v, y_v, z_v)^T$  can be represented by  $n$  straight lines with parameters  $\lambda_i$ :

$$x = x_v + \lambda_i a_i, \quad y = y_v + \lambda_i b_i, \quad z = z_v + \lambda_i, \quad i = 1, \dots, n,$$

where  $a_i$  and  $b_i$  are the direction tangents at the vertex. At the reference surface  $z = z_{\text{ref}}$  track  $i$  is specified by its parameter vector  $\mathbf{q}_i = (x_i, y_i, a_i, b_i)^T$ , consisting of the intersection point  $(x_i, y_i)$  and the two direction tangents  $(a_i, b_i)$ . The track fit delivers estimates  $\tilde{\mathbf{q}}_i$  and information matrices  $\mathbf{G}_i$  for  $i = 1, \dots, n$ . In the constrained problem the sum of the squared residuals

$$M(\mathbf{q}_1, \dots, \mathbf{q}_n) = \sum_{i=1}^n \mathbf{e}_i^T \mathbf{G}_i \mathbf{e}_i, \quad \mathbf{e}_i = \tilde{\mathbf{q}}_i - \mathbf{q}_i, \quad (13.7)$$

must be minimized under the  $2n$  nonlinear constraints

$$x_v = x_i + a_i(z - z_{\text{ref}}), \quad y_v = y_i + b_i(z - z_{\text{ref}}), \quad i = 1, \dots, n$$

There are  $4n + 3$  unknowns,  $4n$  observations and  $2n$  constraints, giving  $4n + 2n - (4n + 3) = 2n - 3$  degrees of freedom. The resulting track parameters  $\tilde{\mathbf{q}}_i$  fit best, in the least-squares sense, to the track fit estimates  $\tilde{\mathbf{q}}_i$  and at the same time have a common vertex. For the solution of the constrained vertex fit see Sect. 13.2.4.2.

In the example, the constraints can be rewritten as

$$x_i = x_v + (z_{\text{ref}} - z)a_i, \quad y_i = y_v + (z_{\text{ref}} - z)b_i, \quad i = 1, \dots, n. \quad (13.8)$$

Insertion of Eq. (13.8) into Eq. (13.7) gives the objective function of the unconstrained nonlinear least-squares problem:

$$M(\mathbf{v}, a_1, b_1, \dots, a_n, b_n) = \sum_{i=1}^n \mathbf{e}_i^T \mathbf{G}_i \mathbf{e}_i, \quad \mathbf{e}_i = \tilde{\mathbf{q}}_i - \mathbf{q}_i.$$

There are now  $4n$  observations and  $2n + 3$  unknown parameters, namely the vertex position and the track directions at the vertex, giving again  $4n - (2n + 3) = 2n - 3$  degrees of freedom.

A generalization of this simple case to helix tracks can be found in [122–124]. In the general case, the unconstrained problem can be formulated in terms of the unknown vertex position  $\mathbf{v}$  and the unknown track momentum vectors  $\mathbf{p}_i$  at the vertex [26, 59]. The measurement equation reads

$$\mathbf{q}_i = \mathbf{h}_i(\mathbf{v}, \mathbf{p}_i), i = 1, \dots, n, \quad (13.9)$$

where the function  $\mathbf{h}_i$  incorporates the track model in the magnetic field. The objective function is equal to

$$M(\mathbf{v}, \mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{i=1}^n \mathbf{e}_i^T \mathbf{G}_i \mathbf{e}_i, \quad \mathbf{e}_i = \tilde{\mathbf{q}}_i - \mathbf{q}_i.$$

Minimization of the objective function can proceed in several ways. For a detailed exposition of non-linear least-squares estimation see e.g. [125].

#### Gauss-Newton Method

Assume that there are approximate values  $\check{\mathbf{v}}$  and  $\check{\mathbf{p}}_i$  for all  $i$ . Then Eq. (13.9) can be approximated by an affine function:

$$\mathbf{q}_i \approx \mathbf{h}_i(\check{\mathbf{v}}, \check{\mathbf{p}}_i) + \mathbf{A}_i(\mathbf{v} - \check{\mathbf{v}}) + \mathbf{B}_i(\mathbf{p}_i - \check{\mathbf{p}}_i) = \mathbf{c}_i + \mathbf{A}_i \mathbf{v} + \mathbf{B}_i \mathbf{p}_i,$$

with

$$\mathbf{A}_i = \left. \frac{\partial \mathbf{h}_i(\mathbf{v}, \mathbf{p}_i)}{\partial \mathbf{v}} \right|_{\check{\mathbf{v}}, \check{\mathbf{p}}_i}, \quad \mathbf{B}_i = \left. \frac{\partial \mathbf{h}_i(\mathbf{v}, \mathbf{p}_i)}{\partial \mathbf{p}_i} \right|_{\check{\mathbf{v}}, \check{\mathbf{p}}_i}, \quad \mathbf{c}_i = \mathbf{h}_i(\check{\mathbf{v}}, \check{\mathbf{p}}_i) - \mathbf{A}_i \check{\mathbf{v}} - \mathbf{B}_i \check{\mathbf{p}}_i.$$

The objective function then reads

$$M(\mathbf{v}, \mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{i=1}^n (\tilde{\mathbf{q}}_i - \mathbf{c}_i - \mathbf{A}_i \mathbf{v} - \mathbf{B}_i \mathbf{p}_i)^T \mathbf{G}_i (\tilde{\mathbf{q}}_i - \mathbf{c}_i - \mathbf{A}_i \mathbf{v} - \mathbf{B}_i \mathbf{p}_i).$$

As  $M$  is now quadratic in the unknown parameters, the minimum can be computed explicitly. The estimated vertex position and its covariance matrix are given by

$$\tilde{\mathbf{v}}_n = \mathbf{C}_n \sum_{i=1}^n \mathbf{A}_i^T \mathbf{G}_i^B (\tilde{\mathbf{q}}_i - \mathbf{c}_i), \quad \text{Var}(\tilde{\mathbf{v}}_n) = \mathbf{C}_n = \left( \sum_{i=1}^n \mathbf{A}_i^T \mathbf{G}_i^B \mathbf{A}_i \right)^{-1}, \quad (13.10)$$

with

$$\mathbf{G}_i^B = \mathbf{G}_i - \mathbf{G}_i \mathbf{B}_i \mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i, \quad \mathbf{W}_i = (\mathbf{B}_i^T \mathbf{G}_i \mathbf{B}_i)^{-1}.$$

In general, the procedure has to be iterated. The measurement equation is expanded at the new estimate, and the estimate is recomputed until convergence is obtained. The formulas required for the implementation of two important cases, fixed-target configuration and solenoidal configuration, are given in [61].

Once  $\tilde{\mathbf{v}}_n$  is known, the track momenta and the full covariance matrix can be computed:

$$\begin{aligned} \tilde{\mathbf{p}}_i^n &= \mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i (\tilde{\mathbf{q}}_i - \mathbf{c}_i - \mathbf{A}_i \tilde{\mathbf{v}}_n), \\ \text{Var}(\tilde{\mathbf{p}}_i^n) &= \mathbf{D}_i^n = \mathbf{W}_i + \mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i \mathbf{A}_i \mathbf{C}_n \mathbf{A}_i^T \mathbf{G}_i \mathbf{B}_i \mathbf{W}_i, \\ \text{Cov}(\tilde{\mathbf{p}}_i^n, \tilde{\mathbf{v}}_n) &= \mathbf{E}_i^n = -\mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i \mathbf{A}_i \mathbf{C}_n. \end{aligned} \quad (13.11)$$

The estimates can also be computed recursively, resulting in an extended Kalman filter [25, 26, 59]:

$$\begin{aligned} \tilde{\mathbf{v}}_i &= \mathbf{C}_i [\mathbf{C}_{i-1}^{-1} \tilde{\mathbf{v}}_{i-1} + \mathbf{A}_i^T \mathbf{G}_i^B (\tilde{\mathbf{q}}_i - \mathbf{c}_i)], \quad \mathbf{C}_i = (\mathbf{C}_{i-1}^{-1} + \mathbf{A}_i^T \mathbf{G}_i \mathbf{A}_i)^{-1} \\ \tilde{\mathbf{p}}_i &= \mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i (\tilde{\mathbf{q}}_i - \mathbf{c}_i - \mathbf{A}_i \tilde{\mathbf{v}}_i), \quad \mathbf{D}_i = \mathbf{W}_i + \mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i \mathbf{A}_i \mathbf{C}_i \mathbf{A}_i^T \mathbf{G}_i \mathbf{B}_i \mathbf{W}_i, \\ \mathbf{E}_i &= -\mathbf{W}_i \mathbf{B}_i^T \mathbf{G}_i \mathbf{A}_i \mathbf{C}_i. \end{aligned}$$

The associated smoother is tantamount to recomputing the track momenta using the last vertex estimate  $\tilde{\mathbf{v}}_n$ , i.e., Eq. (13.11).

### Newton–Raphson Method

This method uses a local quadratic approximation to the objective function. In order to simplify the notation we introduce  $\boldsymbol{\alpha} = (\mathbf{v}, \mathbf{p}_1, \dots, \mathbf{p}_n)^T$ ,  $\tilde{\mathbf{q}} = (\tilde{\mathbf{q}}_1, \dots, \tilde{\mathbf{q}}_n)^T$  and  $\mathbf{h} = (\mathbf{h}_1, \dots, \mathbf{h}_n)^T$ . Then the objective function can be written as

$$M(\boldsymbol{\alpha}) = [\tilde{\mathbf{q}} - \mathbf{h}(\boldsymbol{\alpha})]^T \mathbf{G} [\tilde{\mathbf{q}} - \mathbf{h}(\boldsymbol{\alpha})], \quad \mathbf{G} = \text{diag}(\mathbf{G}_1, \dots, \mathbf{G}_n).$$

If  $\check{\boldsymbol{\alpha}}$  is an appropriate expansion point,  $M(\boldsymbol{\alpha})$  is approximated by

$$M(\boldsymbol{\alpha}) \approx M(\check{\boldsymbol{\alpha}}) + \mathbf{g}^T (\boldsymbol{\alpha} - \check{\boldsymbol{\alpha}}) + \frac{1}{2} (\boldsymbol{\alpha} - \check{\boldsymbol{\alpha}})^T \boldsymbol{\Omega} (\boldsymbol{\alpha} - \check{\boldsymbol{\alpha}}),$$

where

$$\mathbf{g} = \frac{\partial M}{\partial \boldsymbol{\alpha}} = -2 \mathbf{H}^T \mathbf{G} [\tilde{\mathbf{q}} - \mathbf{h}(\check{\boldsymbol{\alpha}})], \quad \boldsymbol{\Omega} = \frac{\partial^2 M}{\partial \boldsymbol{\alpha} \partial \boldsymbol{\alpha}^T} = 2 \mathbf{H}^T \mathbf{G} \mathbf{H} - 2 \frac{\partial \mathbf{H}^T}{\partial \boldsymbol{\alpha}^T} \mathbf{G} [\tilde{\mathbf{q}} - \mathbf{h}(\check{\boldsymbol{\alpha}})]$$

are the gradient and the Hessian of  $M$ , respectively, evaluated at  $\tilde{\alpha}$ , and  $\mathbf{H}$  is the Jacobian of the track model  $\mathbf{h}(\alpha)$ . If  $\Omega$  is positive definite,  $M$  has a minimum when its gradient is zero, leading to

$$\tilde{\alpha} = \check{\alpha} - \Omega^{-1} \mathbf{g}.$$

If the second term of the Hessian is set to zero, the Gauss–Newton method is recovered. Clearly, the Newton–Raphson method is more complex, but it gives some additional information about the problem. In particular, a Hessian that is not positive definite indicates that the expansion point is too far from the true global minimum.

### Levenberg–Marquardt Method

In this method the matrix  $\mathbf{H}^T \mathbf{G} \mathbf{H}$  is inflated by a diagonal matrix  $k\mathbf{I}$ . As a consequence, the direction of the parameter update is intermediate between the direction of the Gauss–Newton step ( $k = 0$ ) and the direction of steepest descent ( $k \rightarrow \infty$ ). An example of a vertex fit with the Levenberg–Marquardt method is given in [122].

### Fast Vertex Fits

The estimated track parameters  $\tilde{\mathbf{q}}_i$  are frequently given at the innermost detector surface or at the beam tube. If the  $\tilde{\mathbf{q}}_i$  are propagated to the vicinity of the presumed vertex, the vertex estimation can be speeded up by applying some approximations.

The “perigee” parametrization for helical tracks was introduced in [123], with a correction in [124]. The track is parameterized around the point of closest approach (the perigee point  $\mathbf{v}^P$ ) of the helix to the  $z$ -axis. The variation of transverse errors along the track is neglected in the vicinity of the perigee, and the track direction and curvature at the vertex is considered to be constant. The approximate objective function of the vertex fit can then be written entirely in terms of the perigee points:

$$M(\mathbf{v}) = \sum_{i=1}^n (\mathbf{v}_i^P - \mathbf{v})^T \mathbf{T}_i (\mathbf{v}_i^P - \mathbf{v}), \quad (13.12)$$

where  $\mathbf{T}_i$  is a weight matrix of rank 2. The vertex estimate is then

$$\tilde{\mathbf{v}} = \left( \sum_{i=1}^n \mathbf{T}_i \right)^{-1} \left( \sum_{i=1}^n \mathbf{T}_i \mathbf{v}_i^P \right).$$

The Jacobians required to compute the  $\mathbf{T}_i$  are spelled out in [123, 124].

A further simplification was proposed in [126]. In the vicinity of the vertex the track is approximated by a straight line. The estimated track parameters are transformed to a coordinate system the  $x$ -axis of which is parallel to the track. The vertex is then estimated by minimizing the sum of the weighted transverse distances of the tracks to the vertex. The resulting objective function has the same form as in Eq. (13.12), again with weight matrices of rank 2. The estimate is exact for straight tracks.

A different type of a fast vertex fitting algorithm is described in [127]. It is based on approximating the tracks by straight lines both in the  $x$ - $y$  plane and in the  $x$ - $z$  plane. In either projection, the lines representing the tracks are Hough-transformed to points in the dual plane of line parameters. The vertex coordinates are then obtained by a weighted linear least-squares fit in the dual plane.

### Adding Prior Information

If the vertex to be fitted is the primary vertex, there may be prior information about the vertex position from the beam profile in a collider experiment or the target location in a fixed target experiment. The prior information usually comes in the form of a position  $\mathbf{v}_0$  plus a covariance matrix  $\mathbf{C}_0$ . The objective function is then augmented by an additional term

$$(\mathbf{v}_0 - \mathbf{v})^T \mathbf{C}_0^{-1} (\mathbf{v}_0 - \mathbf{v}).$$

For instance, the Gauss–Newton estimate Eq. (13.10) is modified in the following way:

$$\tilde{\mathbf{v}}_n = \mathbf{C}_n \left[ \mathbf{C}_0^{-1} \mathbf{v}_0 + \sum_{i=1}^n \mathbf{A}_i^T \mathbf{G}_i^B (\tilde{\mathbf{q}}_i - \mathbf{c}_i) \right], \quad \text{Var}(\tilde{\mathbf{v}}_n) = \mathbf{C}_n = \left( \mathbf{C}_0^{-1} + \sum_{i=1}^n \mathbf{A}_i^T \mathbf{G}_i^B \mathbf{A}_i \right)^{-1}.$$

Similar modifications apply to the Newton–Raphson estimate and the fast vertex fits.

#### 13.2.3.2 Vertex Quality and Outlier Removal

Some tracks used in the vertex fit may be outliers in the sense that they do not actually belong to the vertex. Also, the estimated track parameters may be distorted by outliers or distorted hits in the track fit. Both types of outliers distort the vertex estimate and need to be identified.

In the case of Gaussian errors and a linear model the contribution of each track to the minimum value of the objective function is distributed according to a  $\chi^2$ -distribution with two degrees of freedom. The contribution  $\chi_i^2$  of track  $i$  has to be computed relative to the vertex estimated without track  $i$ . For instance, in the



Gauss–Newton algorithm:

$$\chi_i^2 = \mathbf{r}_i^{nT} \mathbf{G}_i \mathbf{r}_i^n + (\tilde{\mathbf{v}}_n - \tilde{\mathbf{v}}_n^{-i})^T (\mathbf{C}_n^{-i})^{-1} (\tilde{\mathbf{v}}_n - \tilde{\mathbf{v}}_n^{-i}),$$

where  $\mathbf{r}_i^n = \tilde{\mathbf{q}}_i - \mathbf{c}_i - \mathbf{A}_i \tilde{\mathbf{v}}_n - \mathbf{B}_i \tilde{\mathbf{p}}_i^n$  is the residual of track  $i$  and  $\tilde{\mathbf{v}}_n^{-i}$  is the vertex estimate with track  $i$  removed:

$$\tilde{\mathbf{v}}_n^{-i} = \mathbf{C}_n^{-i} \left[ \mathbf{C}_n^{-1} \tilde{\mathbf{v}}_n - \mathbf{A}_i^T \mathbf{G}_i^B (\tilde{\mathbf{q}}_i - \mathbf{c}_i) \right], \quad \mathbf{C}_n^{-i} = \left( \mathbf{C}_n^{-1} - \mathbf{A}_i^T \mathbf{G}_i^B \mathbf{A}_i \right)^{-1}.$$

Analogous but somewhat simpler formulas hold for the fast vertex fits.

The test statistic  $\chi_i^2$  can be computed for all  $i$ , and the track with the largest  $\chi_i^2$  is a candidate for removal. This procedure can be repeated until all  $\chi_i^2$  are below the cut. Even if there is only a single outlier, all  $\chi_i^2$  are no longer  $\chi^2$ -distributed and the power of the test is impaired. This loss of power can be compensated by robust estimation of the vertex.

### 13.2.3.3 Robust and Adaptive Estimators

Robust estimators are less influenced or not influenced at all by outlying observations. This can be achieved by downweighting outliers or by excluding them from the estimate. For example, in the case of a one-dimensional location estimate, the M-estimator [128] downweights outliers, whereas the LMS (least median of squares) estimator [129] uses only one half of the sample (the one spanning the shortest interval) and ignores the other one.

Robust estimators tend to be statistically less efficient and computationally more expensive than least-squares estimators. On the other hand, estimation and outlier detection are performed in parallel, whereas a least-squares estimator has to be recomputed after an outlier has been identified and removed.

One of the earliest proposals for a robust vertex fit is in [130]. The method is an M-estimator with Huber's  $\psi$ -function [131]. It is implemented as a re-weighted least-squares estimator. The initial vertex estimate is a plain least-squares estimate. Then, for each track, the residuals are rotated to the eigensystem of the covariance matrix of the track, and weight factors are computed according to

$$w_i = \frac{\psi(r_i/\sigma_i)}{r_i/\sigma_i} = \begin{cases} 1, & |r_i| \leq c\sigma_i, \\ c\sigma_i/|r_i|, & |r_i| > c\sigma_i, \end{cases}$$

where  $r_i$  is one of the residuals in the rotated frame,  $\sigma_i$  is the standard deviation in the rotated frame, and  $c$  is the robustness constant, usually chosen between 1 and 3. The weight factors are applied and the estimate is recomputed. The entire procedure is iterated until convergence.

A different kind of re-weighted least-squares estimator is proposed in [132]. The weights are computed according to Tukey's bi-square function [128]:

$$w_i = \begin{cases} \left(1 - \frac{r_i^2/\sigma_i^2}{c^2}\right)^2, & |r_i| \leq c\sigma_i, \\ 0, & \text{otherwise,} \end{cases}$$

where  $r_i^2$  is the squared residual of track  $i$  with respect to the vertex,  $\sigma_i^2$  is its variance, and  $c$  is again the robustness constant. The estimator is now equivalent to a redescending M-estimator, and consequently less sensitive to outliers than Huber's M-estimator.

The combination of a redescending M-estimator with the concept of deterministic annealing [133] leads to the adaptive method of vertex fitting [113, 117, 134, 135]. The concept of the adaptive vertex fit is derived from the Deterministic Annealing Filter [64] (see Sect. 13.1.3.5). The weights are computed according to

$$w_i = \frac{\exp(-\chi_i^2/2T)}{\exp(-\chi_i^2/2T) + \exp(-\chi_{\text{cut}}^2/2T)},$$

where  $\chi_i^2$  is the  $\chi^2$ -contribution of track  $i$ ,  $\chi_{\text{cut}}^2$  is a cutoff value, and  $T$  is a temperature parameter. The computation of the redescending M-estimator can be interpreted as an EM (expectation-maximization) algorithm [136, 137]. Alternatively it can be viewed as the minimization of the energy function of an elastic arm algorithm [18, 19]. If annealing is employed, the iteration starts at high  $T$ . The temperature is then gradually decreased. At low  $T$  the weights approach either zero or one. The final weights can be used for classification of the tracks as inliers or outliers. A comparison of the adaptive method with other robust estimators can be found in [138]. The adaptive estimator has been extended to a multi-vertex estimator fitting several vertices simultaneously, including competition of all vertices for all tracks [139].

Iterated re-weighted least-squares estimators require a good starting point, in order to ensure convergence to the global minimum and to minimize the number of iterations required. In many cases a standard least-squares estimate is sufficient. In the presence of a large number of outliers also the starting point should be estimated robustly, preferably by an estimator with a high breakdown point [129]. Several such initial estimators have been proposed and studied in [117].

The M-estimators and the adaptive estimator presented above do not presuppose an explicit outlier model. If it is possible to describe the outliers by a Gaussian mixture model, estimation of the vertex can be carried out by the Gaussian-sum filter [140].

Several of the estimators described here are implemented in RAVE, a detector-independent toolkit for reconstruction of interaction vertices [120, 121].

### 13.2.4 Kinematic Fitting

Kinematic fitting imposes physical constraints on the particles participating in an interaction and thereby improves the measured track momenta and positions. At the same time hypotheses about the interaction and the participating particles can be tested.

#### 13.2.4.1 Lagrange Multiplier Method

The most commonly used method of imposing constraints on the measured tracks is by way of Lagrange multipliers [141]. Let  $\tilde{\mathbf{q}} = (\tilde{\mathbf{q}}_1, \dots, \tilde{\mathbf{q}}_n)^T$  be the unconstrained estimated parameters of a set of  $n$  tracks, along with their joint information matrix  $\mathbf{G} = \text{diag}(\mathbf{G}_1, \dots, \mathbf{G}_n) = \mathbf{V}^{-1}$ . The  $r$  functions describing the constraints can be written as  $\mathbf{g}(\mathbf{q}) = \mathbf{0}$ . Taylor expansion around a suitable point  $\check{\mathbf{q}}$  yields the linearized equation

$$\check{\mathbf{g}} + \mathbf{D}(\mathbf{q} - \check{\mathbf{q}}) = \mathbf{0},$$

where  $\mathbf{D}$  is the Jacobian of  $\mathbf{g}$  with respect to  $\mathbf{q}$ , evaluated at  $\check{\mathbf{q}}$ , and  $\check{\mathbf{g}} = \mathbf{g}(\check{\mathbf{q}})$ . The obvious expansion point is  $\check{\mathbf{q}} = \tilde{\mathbf{q}}$ . The constrained track parameters  $\bar{\mathbf{q}}_i$  are obtained by minimizing the objective function

$$M(\mathbf{q}, \boldsymbol{\lambda}) = (\mathbf{q} - \tilde{\mathbf{q}})^T \mathbf{G}(\mathbf{q} - \tilde{\mathbf{q}}) + 2\boldsymbol{\lambda}^T [\check{\mathbf{g}} + \mathbf{D}(\mathbf{q} - \check{\mathbf{q}})]$$

with respect to  $\mathbf{q}$  and  $\boldsymbol{\lambda}$ .  $\boldsymbol{\lambda}$  is a vector of  $r$  unknowns, the Lagrange multipliers. The solution is

$$\bar{\mathbf{q}} = \tilde{\mathbf{q}} - \mathbf{V} \mathbf{D}^T \bar{\boldsymbol{\lambda}}, \quad \text{with} \quad \bar{\boldsymbol{\lambda}} = \mathbf{G}_D [\check{\mathbf{g}} + \mathbf{D}(\tilde{\mathbf{q}} - \check{\mathbf{q}})] \quad \text{and} \quad \mathbf{G}_D = (\mathbf{D} \mathbf{V} \mathbf{D}^T)^{-1}.$$

The covariance matrix  $\bar{\mathbf{V}}$  and the  $\chi^2$  statistic are given by

$$\bar{\mathbf{V}} = \mathbf{V} - \mathbf{V} \mathbf{D}^T \mathbf{G}_D \mathbf{D} \mathbf{V}, \quad \chi^2 = \bar{\boldsymbol{\lambda}}^T \mathbf{G}_D^{-1} \bar{\boldsymbol{\lambda}} = \bar{\boldsymbol{\lambda}}^T [\check{\mathbf{g}} + \mathbf{D}(\tilde{\mathbf{q}} - \check{\mathbf{q}})].$$

If required, the constraint function  $\mathbf{g}$  can be re-expanded at the new point  $\check{\mathbf{q}} = \bar{\mathbf{q}}$ , and the constrained track parameters can be recomputed.

The Jacobian  $\mathbf{D}$  depends both on the parametrization of the tracks and on the type of constraint to be imposed. For kinematic constraints it is often convenient to choose a parametrization that uses physically meaningful quantities. In [142] it is proposed to use the four-momentum and a point in space, i.e.,  $\mathbf{q} = (p_x, p_y, p_z, E, x, y, z)$ . With this parametrization the following constraints can be formulated in a straightforward manner (for further examples see [142]).

1. Invariant mass constraint. The equation that constrains a track to have an invariant mass  $m_c$  is

$$E^2 - p_x^2 - p_y^2 - p_z^2 - m_c^2 = 0.$$

Expanding at  $\check{\mathbf{q}} = (\check{p}_x, \check{p}_y, \check{p}_z, \check{E}, \check{x}, \check{y}, \check{z})$  yields

$$\mathbf{D} = (-2\check{p}_x \ -2\check{p}_y \ -2\check{p}_z \ 2\check{E} \ 0 \ 0 \ 0), \quad \check{\mathbf{g}} = \check{E}^2 - \check{p}_x^2 - \check{p}_y^2 - \check{p}_z^2 - m_c^2.$$

2. Total energy constraint. The equation that constrains a track to have a total energy  $E_c$  is

$$E - E_c = 0.$$

It follows that

$$\mathbf{D} = (0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0), \quad \check{\mathbf{g}} = \check{E} - E_c.$$

3. Total momentum constraint. The equation that constrains a track to have a total momentum  $p_c$  is

$$\sqrt{p_x^2 + p_y^2 + p_z^2} - p_c = 0.$$

Expanding at  $\check{\mathbf{q}} = (\check{p}_x, \check{p}_y, \check{p}_z, \check{E}, \check{x}, \check{y}, \check{z})$  yields

$$\mathbf{D} = \left( \frac{\check{p}_x}{\check{p}} \ \frac{\check{p}_y}{\check{p}} \ \frac{\check{p}_z}{\check{p}} \ 0 \ 0 \ 0 \ 0 \right), \quad \check{\mathbf{g}} = \sqrt{\check{p}_x^2 + \check{p}_y^2 + \check{p}_z^2} - p_c.$$

#### 13.2.4.2 Vertex Constraint

If a vertex constraint is added to the kinematic constraints, the constraint functions depend on the unknown vertex position  $\mathbf{v}$  and are extended to  $\mathbf{g}(\mathbf{q}, \mathbf{v}) = \mathbf{0}$ . Taylor expansion around a suitable point  $(\check{\mathbf{q}}, \check{\mathbf{v}})$  yields the linearized equation

$$\check{\mathbf{g}} + \mathbf{D}(\mathbf{q} - \check{\mathbf{q}}) + \mathbf{E}(\mathbf{v} - \check{\mathbf{v}}) = \mathbf{0},$$

where  $\mathbf{E}$  is the Jacobian of  $\mathbf{g}$  with respect to  $\mathbf{v}$ , evaluated at  $\check{\mathbf{v}}$ , and  $\check{\mathbf{g}} = \mathbf{g}(\check{\mathbf{q}}, \check{\mathbf{v}})$ . It is assumed that there is prior information about the vertex position, represented by the position  $\check{\mathbf{v}}$  and the covariance matrix  $\mathbf{C}$ . The position  $\check{\mathbf{v}}$  can be used as the expansion point  $\check{\mathbf{v}}$ .

The constrained track parameters  $\tilde{q}_i$  and the estimated vertex position  $\tilde{v}$  are obtained by minimizing the objective function

$$M(\mathbf{q}, \mathbf{v}, \boldsymbol{\lambda}) = (\mathbf{q} - \tilde{\mathbf{q}})^T \mathbf{G} (\mathbf{q} - \tilde{\mathbf{q}}) + (\mathbf{v} - \tilde{\mathbf{v}})^T \mathbf{C}^{-1} (\mathbf{v} - \tilde{\mathbf{v}}) + 2\boldsymbol{\lambda}^T [\check{\mathbf{g}} + \mathbf{D}(\mathbf{q} - \check{\mathbf{q}}) + \mathbf{E}(\mathbf{v} - \check{\mathbf{v}})]$$

with respect to  $\mathbf{q}$ ,  $\mathbf{v}$ , and  $\boldsymbol{\lambda}$ . The solution is

$$\bar{\boldsymbol{\lambda}} = \mathbf{W} [\check{\mathbf{g}} + \mathbf{D}(\tilde{\mathbf{q}} - \check{\mathbf{q}}) + \mathbf{E}(\tilde{\mathbf{v}} - \check{\mathbf{v}})], \quad \bar{\mathbf{v}} = \tilde{\mathbf{v}} - \mathbf{C}\mathbf{E}^T \bar{\boldsymbol{\lambda}}, \quad \bar{\mathbf{q}} = \tilde{\mathbf{q}} - \mathbf{V}\mathbf{D}^T \bar{\boldsymbol{\lambda}},$$

with  $\mathbf{W} = (\mathbf{D}\mathbf{V}\mathbf{D}^T + \mathbf{E}\mathbf{C}\mathbf{E}^T)^{-1}$ . The covariance matrices are

$$\text{Var}(\bar{\mathbf{v}}) = \mathbf{C} - \mathbf{C}\mathbf{E}^T \mathbf{W} \mathbf{E} \mathbf{C}, \quad \text{Var}(\bar{\mathbf{q}}) = \mathbf{V} - \mathbf{V}\mathbf{D}^T \mathbf{W} \mathbf{D} \mathbf{V}, \quad \text{Cov}(\bar{\mathbf{q}}, \bar{\mathbf{v}}) = -\mathbf{V}\mathbf{D}^T \mathbf{W} \mathbf{E} \mathbf{C}.$$

The  $\chi^2$  statistic is

$$\chi^2 = \bar{\boldsymbol{\lambda}}^T \mathbf{W}^{-1} \bar{\boldsymbol{\lambda}} = \bar{\boldsymbol{\lambda}}^T [\check{\mathbf{g}} + \mathbf{D}(\tilde{\mathbf{q}} - \check{\mathbf{q}}) + \mathbf{E}(\tilde{\mathbf{v}} - \check{\mathbf{v}})],$$

with  $r$  degrees of freedom, where  $r$  is the number of constraint functions. If the vertex constraint is the only constraint imposed on the tracks, the  $\chi^2$  has  $2n$  degrees of freedom. If there is no prior information about the vertex, the prior vertex position is assigned an infinitely large covariance matrix, and  $\mathbf{W}$  is replaced by its limiting value:

$$\mathbf{W} = \lim_{\mathbf{C} \rightarrow \infty} (\mathbf{D}\mathbf{V}\mathbf{D}^T + \mathbf{E}\mathbf{C}\mathbf{E}^T)^{-1} = \mathbf{G}_D - \mathbf{G}_D \mathbf{E} (\mathbf{E}^T \mathbf{G}_D \mathbf{E})^{-1} \mathbf{E}^T \mathbf{G}_D.$$

The number of degrees of freedom is reduced to  $2n - 3$ .

## 13.3 Track Reconstruction in the LHC Experiments

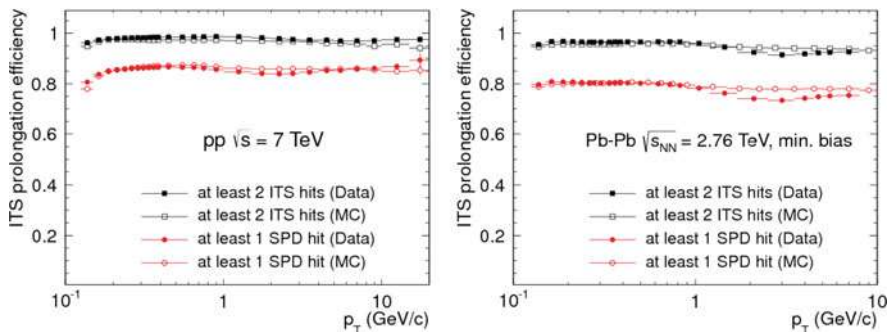
### 13.3.1 ALICE

ALICE [143] is the experiment at the LHC that is devoted to the physics of high energy ion collisions. Its main goal is to investigate the physics of strongly interacting matter and the quark-gluon plasma at extreme values of energy density and temperature in nucleus-nucleus collisions. Among the four experiments at the LHC, ALICE is equipped with the largest number of subdetectors in order to face the reconstruction complexity of ion physics events. In particular, three subdetectors focus on measuring the passage of charged particles using the bending power of the magnetic field. They are assembled in a cylindrical fashion: the Inner Tracking System (ITS) with six planes of high-resolution silicon pixel, drift, and strip detectors, the cylindrical Time-Projection Chamber (TPC) and the Transition Radiation Detector (TRD). The principal functions of the ITS are the

identification and reconstruction of secondary vertices, the track reconstruction of low- $p_T$  particles and the improvement of the impact parameter and momentum resolution. The TPC is the most important tracking sub-detector. Thanks to its time information, it can provide an efficient and robust tracking also in a very high multiplicity environments (in the order of 10,000 charged particles). Finally, the TRD is also used for tracking in the central region and for improving the  $p_T$  resolution at high momentum.

The first step in the track reconstruction in ALICE is the clusterization, which is performed separately for each of the three subdetectors [144]. Tracking then proceeds by determining the preliminary interaction vertex using tracklets defined as lines built with pairs of clusters in the first two layers of the ITS. The preliminary interaction vertex is thus found as a space point to which a maximum number of tracklets converge. In the next step, track finding and fitting is performed in three stages using an inward-outward-inward strategy:

- Initially, tracks in the TPC are searched for using the Kalman filter technique and the outermost layers of the TPC for the seed. A preliminary particle identification is also possible at this stage based on the specific energy loss in the TPC gas. Then, the reconstructed TPC tracks are propagated to the outermost ITS layer and become the seeds for finding tracks in the ITS. In Fig. 13.4 the ITS–TPC matching efficiency as a function of the transverse momentum for 2010–2013 data and Monte Carlo for pp and heavy ion collisions is shown. Finally, the last step is performed in order to recover tracks of particle with  $p_T$  down to 80 MeV. It performs a standalone ITS reconstruction with those clusters that were not used in the ITS–TPC tracks.
- All reconstructed tracks are then extrapolated to their point of closest approach to the preliminary interaction vertex, and are extrapolated from the innermost layer to the outermost one. Tracks are refitted by the Kalman filter using the clusters found at the previous stage. After the reconstruction in the TRD subdetectors, the track is matched with a possible TRD tracklet in each of the six TRD layers. In a



**Fig. 13.4** ITS–TPC matching efficiency vs.  $p_T$  for data and Monte Carlo for pp (left) for Pb-Pb (right) collisions in the ALICE experiment [144]

similar way, the tracks reaching the time-of-flight (TOF) detector are matched to TOF clusters.

- At the final stage of the track reconstruction, all tracks in both ITS and TPC subdetectors are propagated inwards and refitted one last time to determine the final estimate of the track position, direction, inverse curvature, and its associated covariance matrix.

The final interaction vertex is then re-determined using the all tracks reconstructed in TPC and ITS. The precise vertex fit is performed using track weighting to suppress the contribution of any remaining outliers. For data-taking conditions where a high pileup rate is expected, a more robust version of vertex finding inspired by the algorithm described in [132] is used. It is based on iterative vertex finding and fitting using Tukey bisquare weights to suppress outliers. The algorithm stops when no more vertices are identified in the scan along the beam direction. Once the tracks and the interaction vertex have been found, a search for photon conversions and decays of strange hadrons such as  $K_S^0$  and  $\Lambda_0$  concludes the central-barrel tracking procedure.

### 13.3.2 *ATLAS*

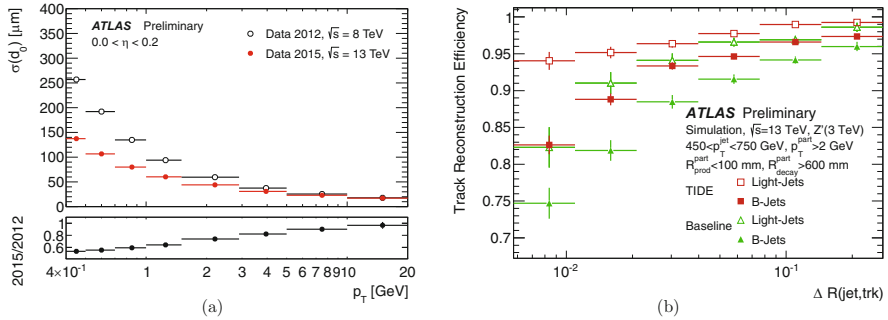
ATLAS [145] is the largest of the four LHC experiments, measuring 25 m in diameter and 44 m in length. Its magnet system is composed of a Central Solenoid Magnet with a 2 T field, a Barrel Toroid and an Endcap Toroids with 4 T each. The Inner Detector (ID) is very compact and highly sensitive in order to measure accurately the decay products of each collision. It consists of three different systems of sensors immersed in the solenoid magnetic field: the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT). The Pixel Detector is situated closest to the interaction point and has the highest granularity with about 80 million readout channels. The intrinsic spatial resolution of the Pixel Detector sensors is  $10\ \mu\text{m}$  in  $r-\phi$  and  $115\ \mu\text{m}$  in  $z$ . The SCT is a silicon microstrip detector surrounding the Pixel Detector. It provides eight measurements per track with an overall resolution of  $16\ \mu\text{m}$  in  $r-\phi$  and  $580\ \mu\text{m}$  in  $z$ . In the outermost region, the TRT is placed. It is a light-weight detector composed of proportional gas counters (70% Xe, 27%  $\text{CO}_2$  and 3%  $\text{O}_2$  straws) embedded in a radiator material and its operational drift radius accuracy is about  $130\ \mu\text{m}$ . The TRT contributes both to the track pattern recognition stage, featuring typically around 30 hits per track, and to particle identification.

The basic concepts of the ATLAS track reconstruction are described in [146, 147]. The tracking in the ID consists of two principal sequences: an initial inside-out tracking, and a subsequent outside-in tracking. Inside-out tracking starts with space point formation in the silicon part of the ID. Using the space points, track seeds are generated with or without a constraint on the longitudinal vertex position. The seeds are then followed through the SCT by a combinatorial Kalman filter/smoother.

After ambiguity solving, the remaining track candidates are extended into the TRT. Outside-in tracking first finds track segments in the TRT, using a Hough transform of the straw centers. A Kalman filter/smoothen using also the drift times builds the final track segments. These track segments are then extrapolated back into the SCT and the Pixel Detector. Muons are reconstructed in the ID like any other charged particles; for the standalone reconstruction of muons in the muon system and the combined reconstruction, see [148].


Based on the experience gained in Run 1, several improvements to track reconstruction were made for Run 2 [149]. For example, the tracking was adapted to the new insertable B-layer (IBL) [150], and track reconstruction in dense environments (TIDE) was optimized [151]. This included an artificial neural network based approach to identify pixel clusters created by multiple charged particles. The effect of these two developments is shown in Fig. 13.5. In Fig. 13.5a the transverse impact parameter as a function of track momentum resolution is shown for data taken in 2015 at 13 TeV with the inclusion of the IBL information and for data in 2012 at 8 TeV without the IBL. The data in 2015 was collected with a minimum bias trigger. The data in 2012 is derived from a mixture of jet, tau and missing  $E_T$  triggers [150]. Figure 13.5b shows the improvement of the track reconstruction efficiency in the jet core due to the TIDE optimization [151].

ATLAS track reconstruction efficiency as a function of pseudorapidity and transverse momentum with simulated data at a center-of-mass energy of 13 TeV is shown in Fig. 13.6 [152].



**Fig. 13.5** (a) Upper panel: unfolded transverse impact parameter resolution measured from data in 2015 at 13 TeV with the Inner Detector including the IBL, as a function of track  $p_T$  for values of  $0.0 < \eta < 0.2$ , compared to that measured from data in 2012 at 8 TeV [150]; lower panel: ratio of the resolution in 2015 over the resolution in 2012. (b) Improvement of the track reconstruction efficiency due to the TIDE optimization, as a function of the angular distance of the particle from the jet axis. The track selection is explained in [151]





**Fig. 13.6** The track reconstruction efficiency (a) as a function of pseudorapidity and (b) as a function of transverse momentum, as predicted by Pythia 8 A2 simulation. The statistical uncertainties are shown as black lines, the total uncertainties as green shaded areas [152]

### 13.3.3 CMS

CMS [153], together with ATLAS, is one of the two general-purpose experiments at the LHC. Its main distinguishing feature is a 3.8 T superconducting solenoid. With a length of 13 m and a diameter of 6 m, it provides a high bending power to precisely measure the momentum of charged particles. The solenoid magnetic field lines run parallel to the beam direction in the central region, where the tracking system is placed. The tracking system is designed to provide a precise and efficient measurement of particle trajectories using position-sensitive detectors. The CMS tracker is a silicon-based system [154]. It splits into two parts, the Pixel Tracker and the Strip Tracker and covers a pseudorapidity range up to  $|\eta| = 2.5$ . The Pixel Tracker is the innermost CMS detector sub-system and is composed of 66 million silicon pixels with dimensions  $100 \times 250 \times 250 \mu\text{m}$ , covering a total area of about  $1 \text{ m}^2$ . In the barrel layers the magnetic field induces a Lorentz angle which increases charge sharing between neighbouring pixels. Charge sharing in conjunction with analog readout allows to achieve  $10 \mu\text{m}$  position resolution for the  $(r, \phi)$  coordinate and  $15 \mu\text{m}$  in the  $z$  direction. The pixel detectors in the forward direction are tilted at an angle of  $20^\circ$  to induce charge sharing which allows to achieve  $15 \mu\text{m}$  and  $20 \mu\text{m}$  resolution respectively. This resolution is not only necessary for a precise track reconstruction, but also for the determination of both the vertices produced in the primary interaction and the decay vertices of short-lived particles.

The Strip Tracker constitutes the outer part of the tracking system. Its basic building blocks are silicon strip modules. Each module is equipped with one or two silicon sensors and a so-called Front-End hybrid containing readout electronics. In total, the CMS silicon strip tracker has 9.3 million strips and covers  $198 \text{ m}^2$  of active silicon area. The resolution in  $(r, \phi)$  is  $\simeq 30 \mu\text{m}$  in all layers. The inner layers of the strip tracker are equipped with double-sided sensors, one side of which is rotated by a stereo angle of  $100 \text{ mrad}$ , achieving a resolution along the  $z$  coordinate

of about  $230\text{ }\mu\text{m}$  and allowing the reconstruction of the hit position in 3-D. In the outer layers the sensors are single-sided, and the  $z$  resolution can be approximated by the strip length over  $\sqrt{12}$ , or about  $15\text{ mm}$ . In order to maintain excellent tracking performance until the Long Shutdown 3 of LHC, the Pixel Tracker was replaced in the year-end technical stop of 2016/2017 with a new Pixel Tracker composed of four barrel layers and six forward disks providing four-hit pixel coverage up to  $|\eta| = 2.5$ . After the Long Shutdown 3, the High Luminosity phase of the LHC (HL-LHC) is scheduled where the accelerator will provide an unprecedented instantaneous luminosity of  $5 - 7.5 \times 10^{34}\text{ cm}^{-2}\text{ s}^{-1}$ . In [155] the new CMS silicon tracker and its tracking and vertexing performance for different event types, pileup scenarios and detector geometries are presented.

The CMS track reconstruction algorithm is based on an iterative approach [156]. The main idea is to search for easier-to-find tracks first, to mask the hits associated to the found tracks, and to proceed to the next iteration. In this way the combinatorial problem is reduced, and the search for more difficult classes of tracks is simplified. Moreover, this approach introduces the possibility of developing special iterations that can improve track reconstruction in high-density environments such as jets, or to use the information from other subsystems such as muon chambers and calorimeters. In each iteration, the Combinatorial Track Finder is run. It can be divided into four different steps:

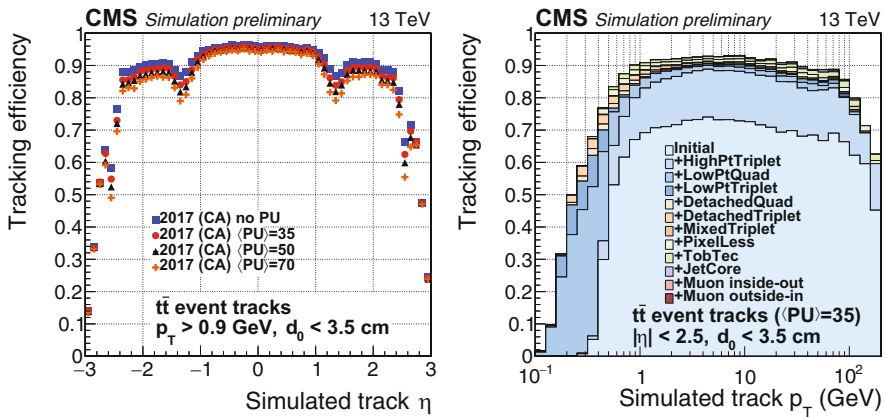
1. **Seed generation:** Using the information of three or four hits, the trajectory parameters and the corresponding uncertainties of the initial track candidates are computed.
2. **Track finding:** Starting from the seed, the current trajectory parameters and their uncertainties are extrapolated to the next layer and compatible hits are found. Each of them is added to a clone of the track candidate. Each of these candidates is again extrapolated to the next layer and compatible hits are found. This procedure is repeated for each candidate until there is more than one missing hit or the extrapolation does not find another tracker layer.
3. **Track fitting:** A Kalman filter or a Gaussian-sum filter/smoothener is performed to obtain the final estimate of the track parameters at the interaction point exploiting the full trajectory information.
4. **Track selection:** Tracks are grouped in classes according to different track quality criteria.

As an example, the twelve tracking iterations foreseen for 2017 data taking is listed in Table 13.3 [157]. The main difference between iterations is the configuration of the seed generation and the target tracks.

Figure 13.7 shows the tracking efficiency, using a standard sample of  $t\bar{t}$  events simulated with  $\sqrt{s} = 13\text{ TeV}$  with different superimposed pileup conditions. The contribution of different iterations for 2017 track reconstruction is also shown as a function of the  $p_T$  of the simulated particle. It can be seen how iterations targeting low- $p_T$  tracks are more efficient in the region between 100 and 500 MeV.

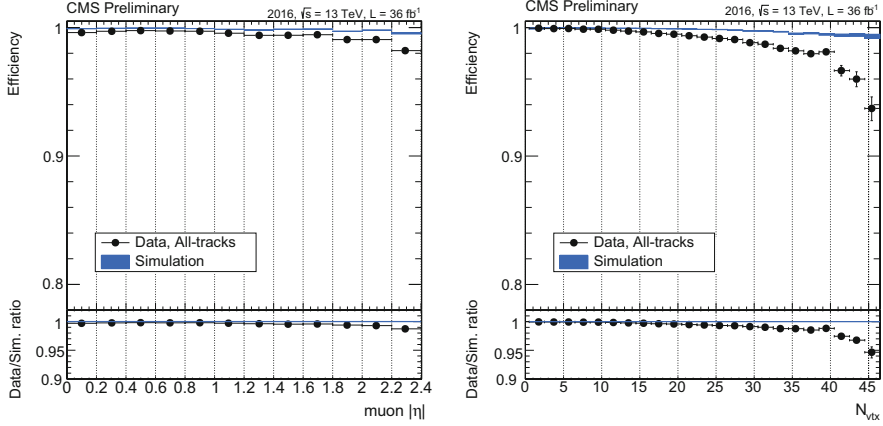
**Table 13.3** List of different tracking iterations used after the Pixel Tracker upgrade with the corresponding seeding configuration used and target tracks [157]

Iteration	Step name	Seeding	Target tracks
0	Initial	Pixel quadruplets	Prompt, high $p_T$
1	LowPtQuad	Pixel quadruplets	Prompt, low $p_T$
2	HighPtTriplet	Pixel triplets	Prompt, high $p_T$ recovery
3	LowPtTriplet	Pixel triplets	Prompt, low $p_T$
4	DetachedQuad	Pixel quadruplets	From b hadron decay, $r \leq 5$ cm
5	DetachedTriplet	Pixel triplets	From b hadron decay, $r \leq 10$ cm
6	MixedTriplet	Pixel+strip triplets	Displaced, $r \leq 7$ cm
7	PixelLess	Inner strip pairs	Displaced, $r \leq 25$ cm
8	TobTec	Outer strip pairs	Displaced, $r \leq 60$ cm
9	JetCore	Pixel pairs in jets	High- $p_T$ jets
10	Muon inside-out	Muon-tagged tracks	Muons
11	Muon outside-in	Standalone muons	Muons



**Fig. 13.7** Track reconstruction efficiency as a function of simulated track pseudorapidity for 2017 tracker at different pileup conditions (left) and cumulative contributions to the overall tracking performance from the twelve iterations in 2017 track reconstruction shown as a function of the simulated track  $p_T$  (right) [157]. The 2017 tracking reconstruction includes the Cellular Automaton-based Hit Chain-Maker (CA) seeding [158]

Figure 13.8 shows the muon tracking efficiency and the corresponding ratios between real and simulated data for 2016 collisions data coming from the Z resonance using the tag and probe method. The measured track efficiency as a function of  $|\eta|$  is found to be between 99.5% and 100% for the collection including all tracks. It degrades, however, with increasing number of primary vertices.



**Fig. 13.18** Data (black dots) and simulation (rectangles) tracking efficiency and respective ratio for muons coming from the Z decay as a function of the absolute pseudorapidity of the probe muon (left) and the number of primary vertices (right). The data are based on an integrated luminosity of  $36 \text{ fb}^{-1}$  [155]

### 13.3.4 LHCb

As its name indicates, LHCb [159] focuses on physics involving bottom quarks and investigates CP violation phenomena. These studies require the measurement of the rare decays of  $B_d$ ,  $B_s$ , and D mesons which are produced with a large cross-section at the LHC. Given the fact that b hadrons are predominantly produced in the forward or backward cone, the LHCb experiment is a single-arm spectrometer in contrast to the other three experiments. In order to exploit this large number of b hadrons, it requires a robust and flexible trigger and a data acquisition that allows high bandwidth data taking and provides powerful online data processing. Furthermore, superior vertex and momentum resolution are crucial to study the rapidly oscillating  $B_s - \bar{B}_s$  meson system. LHCb is thus equipped with the highly sophisticated silicon microstrip detector close to the interaction point, the Vertex Locator (VELO). It can be moved to a distance of only 7 mm from the proton beams and measures the position of the primary vertices and the impact parameters of the track with extremely high precision. A further silicon microstrip detector, the Tracker Turicensis (TT) is placed before the dipole magnet. Its task is to improve the momentum resolution of reconstructed tracks and reject pairs of tracks that in reality belong to the same particle. The magnet is placed behind the TT. It bends the flight path of the particles in the  $x - z$  plane and therefore allows the determination of their momenta. The tracking system is completed by the T stations (T1-T2-T3), which, together with the information from the VELO, determine the momentum and flight direction of the particles. The T stations are composed of silicon microstrip sensors close to the beam pipe and by straw tubes in the outer regions.

Track reconstruction uses hits in the VELO, TT and T stations. Depending on which detectors are crossed, different track types are defined [160, 161]:

- **Long tracks** traverse the full tracking system. They have hits in both the VELO and the T stations, and optionally in TT. They are the most important set of tracks for physics analyses.
- **Upstream tracks** pass only through the VELO and TT stations. In general their momentum is too low to traverse the magnet and reach the T stations.
- **Downstream tracks** pass only through the TT and T stations. They are important for the reconstruction of long-lived particles that decay outside the VELO acceptance.
- **VELO tracks** pass only through the VELO. These tracks are particularly important in the primary vertex reconstruction.
- **T tracks** pass only through the T stations. Like the downstream tracks, they are useful for particle identification in the Ring Imaging Cherenkov detectors.


Reconstruction of long tracks starts in the VELO. There are two complementary algorithms to add information from the downstream tracking stations to these VELO tracks. The first one combines the VELO tracks with information from the T stations. The second one combines the VELO tracks with track segments found after the magnet in the T stations, using a standalone track finding algorithm. The candidate tracks found by each algorithm are then combined, removing duplicates, to form the final set of long tracks used for analysis. Finally, hits in the TT consistent with the extrapolated trajectories of each track are added to improve their momentum determination.

Downstream tracks are found starting with T tracks, extrapolating them through the magnetic field and searching for corresponding hits in the TT. Upstream tracks are found by extrapolating VELO tracks to the TT where matching hits are then added in a procedure similar to that used by the downstream tracking. At least three TT hits are required to be present by these algorithms.

The found tracks are fitted using a Kalman filter, taking into account multiple scattering and energy loss due to ionisation. The  $\chi^2$ -statistic of the fit is used to determine the quality of the reconstructed track. If two or more tracks have many hits in common, only the one with most hits is kept.

The track reconstruction efficiency for the 2012 and the 2015 data as a function of the momentum can be seen in Fig. 13.9 [162]. The results of the two periods are compatible.

**Fig. 13.9** LHCb track reconstruction efficiency for the 2012 and the 2015 data as a function of the momentum. The efficiency is computed using the “Long method”, described in [161]



## 13.4 Conclusion

An overview of current methods in track and vertex reconstruction and alignment has been presented. Many of them have been developed in response to the requirements of the current experimental program at the Large Hadron Collider. The most difficult challenges are:

- Reliable reconstruction of signal events over a large background of non-signal events, pileup events, and low-momentum tracks;
- Reliable reconstruction of secondary vertices with very short distances from the primary vertex;
- Precise alignment of a large number of sensors.

Every experiment has to meet these challenges on its own terms. The outlines of the solutions found by the four major LHC experiments are described in Sect. 13.3 and, in more detail, in the references given there. In addition, the repertory of the methods discussed in this contribution can certainly not lay claim to completeness. We have tried to select widely applicable methods, thereby neglecting by necessity many experiment specific adaptations, improvements and innovations, for which we again refer to the references.

## References

1. R. Mankel, “Pattern recognition and event reconstruction in particle physics experiments,” *Rept. Prog. Phys.*, vol. 67, p. 553, 2004.
2. M. Hansroul, H. Jeremie, and D. Savard, “Fast circle fit with the conformal mapping method,” *Nucl. Instrum. Meth.*, vol. A270, no. 2, p. 498, 1988.
3. P. Hough, “Machine analysis of bubble chamber pictures,” in *Proceedings of the International Conference on High Energy Accelerators and Instrumentation*, (CERN, Geneva), p. 554, 1959.
4. H. Kälviäinen, P. Hirvonen, L. Xu, and E. Oja, “Probabilistic and non-probabilistic Hough transforms: overview and comparisons,” *Image and Vision Computing*, vol. 13, no. 4, p. 239, 1995.

5. T. Alexopoulos, M. Bachtis, E. Gazis, and G. Tsipolitis, "Implementation of the Legendre Transform for track segment reconstruction in drift tube chambers," *Nucl. Instrum. Meth.*, vol. A592, pp. 456–462, 2008.
6. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Nat. Acad. Sci.*, vol. 79, p. 2554, 1982.
7. C. Peterson, "Track Finding With Neural Networks," *Nucl. Instrum. Meth.*, vol. A279, p. 537, 1989.
8. B. Denby, "Neural Networks and Cellular Automata in Experimental High-energy Physics," *Comput. Phys. Commun.*, vol. 49, p. 429, 1988.
9. C. Peterson and J. Anderson, "A Mean Field Theory Learning Algorithm for Neural Networks," *Complex Systems*, vol. 2, p. 995, 1987.
10. M. Diehl, M. Junger, R. Frühwirth, and J. Scherzer, "Global optimization for track finding," *Nucl. Instrum. Meth.*, vol. A389, p. 180, 1997.
11. G. Stimpfl-Abele and L. Garrido, "Fast track finding with neural nets," *Comput. Phys. Commun.*, vol. 64, p. 46, 1991.
12. S. Baginian, A. Glazov, I. Kisel, E. Konotopskaya, V. Neskoromnyi, and G. Ososkov, "Tracking by a modified rotor model of neural network," *Comput. Phys. Commun.*, vol. 79, p. 165, 1994.
13. A. Badalà, R. Barbera, G. Lo Re, A. Palmeri, G. S. Pappalardo, A. Pulvirenti, and F. Riggi, "Neural tracking in ALICE," *Nucl. Instrum. Meth.*, vol. A502, pp. 503–506, 2003.
14. A. Pulvirenti, A. Badalà, R. Barbera, G. Lo Re, A. Palmeri, G. S. Pappalardo, and F. Riggi, "Neural tracking in the ALICE Inner Tracking System," *Nucl. Instrum. Meth.*, vol. A533, pp. 543–559, 2004.
15. A. Badalà, R. Barbera, G. Lo Re, A. Palmeri, G. S. Pappalardo, A. Pulvirenti, and F. Riggi, "Combined tracking in the ALICE detector," *Nucl. Instrum. Meth.*, vol. A534, p. 211, 2004.
16. M. Gyulassy and M. Harlander, "Elastic tracking and neural network algorithms for complex pattern recognition," *Comput. Phys. Commun.*, vol. 66, p. 31, 1991.
17. M. Gyulassy and M. Harlander, "High resolution multiparticle tracking without preprocessing via elastic tracking," *Nucl. Instrum. Meth.*, vol. A316, p. 238, 1992.
18. M. Ohlsson, C. Peterson, and A. Yuille, "Track finding with deformable templates: The Elastic arms approach," *Comput. Phys. Commun.*, vol. 71, p. 77, 1992.
19. M. Ohlsson, "Extensions and explorations of the elastic arms algorithm," *Comput. Phys. Commun.*, vol. 77, p. 19, 1993.
20. R. Durbin and D. Willshaw, "An analogue approach to the travelling salesman," *Nature*, vol. 326, no. 16, p. 689, 1987.
21. D. Bui, T. Greenshaw, and G. Schmidt, "A combination of an elastic net and a Hopfield net to solve the segment linking problem in the forward tracker of the H1 detector at HERA," *Nucl. Instrum. Meth.*, vol. A389, p. 184, 1997.
22. I. Kisel and V. Kovalenko, "Elastic net for broken multiple scattered tracks," *Comput. Phys. Commun.*, vol. 98, p. 45, 1996.
23. I. Kisel *et al.*, "Cellular automaton and elastic net for event reconstruction in the NEMO-2 experiment," *Nucl. Instrum. Meth.*, vol. A387, p. 433, 1997.
24. R. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, p. 35, 1960.
25. D. Catlin, *Estimation, Control, and the Discrete Kalman Filter*. New York: Springer, 1989.
26. R. Frühwirth, "Application of Kalman filtering to track and vertex fitting," *Nucl. Instrum. Meth.*, vol. A262, p. 444, 1987.
27. P. Billoir, "Progressive track recognition with a Kalman like fitting procedure," *Comput. Phys. Commun.*, vol. 57, p. 390, 1989.
28. R. Mankel, "A Concurrent track evolution algorithm for pattern recognition in the HERA-B main tracking system," *Nucl. Instrum. Meth.*, vol. A395, p. 169, 1997.
29. A. Glazov, I. Kisel, E. Konotopskaya, and G. Ososkov, "Filtering tracks in discrete detectors using a cellular automaton," *Nucl. Instrum. Meth.*, vol. A329, p. 262, 1993.

30. R. Frühwirth and A. Strandlie, "Application of adaptive filters to track finding," *Nucl. Instrum. Meth.*, vol. A559, p. 162, 2006.
31. A. Strandlie and R. Frühwirth, "Reconstruction of charged tracks in the presence of large amounts of background and noise," *Nucl. Instrum. Meth.*, vol. A566, p. 157, 2006.
32. A. Strandlie and R. Frühwirth, "Track and vertex reconstruction: From classical to adaptive methods," *Rev. Mod. Phys.*, vol. 82, pp. 1419–1458, May 2010.
33. H. Wind, "Evaluating a magnetic field component from boundary observations only," *Nucl. Instrum. Meth.*, vol. 84, p. 117, 1970.
34. M. Aleksa *et al.*, "Measurement of the ATLAS solenoid magnetic field," *JINST*, vol. 3, p. P04003, 2008.
35. R. Frühwirth, "Preparing Magnetic Field Measurements," *Comput. Phys. Commun.*, vol. 22, p. 223, 1981.
36. N. Amapane, V. Andreev, V. Drollinger, V. Karimäki, V. Klyukhin, and T. Todorov, "Volume-based representation of the magnetic field," in *Computing in high energy physics and nuclear physics. Proceedings, Conference, CHEP'04, Interlaken, Switzerland, September 27-October 1, 2004*, p. 310, 2004.
37. R. Frühwirth *et al.*, *Data Analysis Techniques for High-Energy Physics*. Cambridge: Cambridge University Press, 2 ed., 2000.
38. E. Lund, L. Bugge, I. Gavrilenko, and A. Strandlie, "Track parameter propagation through the application of a new adaptive Runge-Kutta-Nystroem method in the ATLAS experiment," *JINST*, vol. 4, p. P04001, 2009.
39. A. Strandlie and W. Wittek, "Derivation of Jacobians for the propagation of covariance matrices of track parameters in homogeneous magnetic fields," *Nucl. Instrum. Meth.*, vol. A566, no. 2, p. 687, 2006.
40. J. Myrheim and L. Bugge, "A Fast Runge-Kutta Method for Fitting Tracks in a Magnetic Field," *Nucl. Instrum. Meth.*, vol. 160, p. 43, 1979.
41. L. Bugge and J. Myrheim, "Tracking and Track fitting," *Nucl. Instrum. Meth.*, vol. 179, p. 365, 1981.
42. R. Frühwirth and M. Regler, "On the quantitative modelling of core and tails of multiple scattering by Gaussian mixtures," *Nucl. Instrum. Meth.*, vol. A456, no. 3, p. 369, 2001.
43. J. Jackson, *Classical Electrodynamics*. John Wiley & Sons, 2007.
44. C. Patrignani *et al.*, "Review of Particle Physics," *Chin. Phys.*, vol. C40, no. 10, p. 100001, 2016.
45. H. Eichinger and M. Regler, "Review of track-fitting methods in counter experiments," Tech. Rep. CERN-81-06, CERN, 1981.
46. P. Avery, "Applied fitting theory V: track fitting using the Kalman filter," Tech. Rep. CLEO Note CBX92-39, Cornell University, 1992.
47. V. Innocente and E. Nagy, "Trajectory fit in presence of dense materials," *Nucl. Instrum. Meth.*, vol. A324, p. 297, 1993.
48. H. Bethe and W. Heitler, "On the Stopping of fast particles and on the creation of positive electrons," *Proc. Roy. Soc. Lond.*, vol. A146, p. 83, 1934.
49. R. Frühwirth, "A Gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung," *Comput. Phys. Commun.*, vol. 154, no. 2, p. 131, 2003.
50. W. Adam, R. Frühwirth, A. Strandlie, and T. Todorov, "Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at the LHC," *J. Phys. G: Nuclear and Particle Physics*, vol. 31, no. 9, p. N9, 2005.
51. V. Kartvelishvili, "Electron bremsstrahlung recovery in ATLAS," *Nucl. Phys. Proc. Suppl.*, vol. 172, p. 208, 2007.
52. N. Chernov and G. Ososkov, "Effective Algorithms of Circle Fitting," *Comput. Phys. Commun.*, vol. 33, p. 329, 1984.
53. V. Karimäki, "Effective circle fitting for particle trajectories," *Nucl. Instrum. Meth.*, vol. A305, p. 187, 1991.
54. A. Strandlie, J. Wroldsen, R. Frühwirth, and B. Lillekjendlie, "Particle tracks fitted on the Riemann sphere," *Comput. Phys. Commun.*, vol. 131, p. 95, 2000.



55. A. Strandlie and R. Frühwirth, "Error analysis of the track fit on the Riemann sphere," *Nucl. Instrum. Meth.*, vol. A480, p. 734, 2002.
56. A. Strandlie, J. Wroldsen, and R. Frühwirth, "Treatment of multiple scattering with the generalized Riemann sphere track fit," *Nucl. Instrum. Meth.*, vol. A488, p. 332, 2002.
57. P. Laurikainen, W. Moorhead, and W. Matt, "Least squares fit of bubble chamber tracks taking into account multiple scattering," *Nucl. Instrum. Meth.*, vol. 98, p. 349, 1972.
58. C. Kleinwort, "General broken lines as advanced track fitting method," *Nucl. Instrum. Meth.*, vol. A673, p. 107, 2012.
59. P. Billoir, R. Frühwirth, and M. Regler, "Track Element Merging Strategy and Vertex Fitting in Complex Modular Detectors," *Nucl. Instrum. Meth.*, vol. A241, p. 115, 1985.
60. P. Billoir, "Track Fitting With Multiple Scattering: A New Method," *Nucl. Instrum. Meth.*, vol. A225, pp. 352–366, 1984.
61. R. Harr, "Calculation of track and vertex errors for detector design studies," *IEEE Trans. Nucl. Sci.*, vol. 42, p. 134, 1995.
62. R. Frühwirth, "Track fitting with non-Gaussian noise," *Comput. Phys. Commun.*, vol. 100, p. 1, 1997.
63. A. Strandlie and R. Frühwirth, "Discrimination between different types of material in track reconstruction with a Gaussian-sum filter," *IEEE Trans. Nucl. Sci.*, vol. 53, p. 3842, 2006.
64. R. Frühwirth and A. Strandlie, "Track fitting with ambiguities and noise: A study of elastic tracking and nonlinear filters," *Comput. Phys. Commun.*, vol. 120, p. 197, 1999.
65. A. Strandlie and R. Frühwirth, "Adaptive multitrack fitting," *Comput. Phys. Commun.*, vol. 133, p. 34, 2000.
66. R. Frühwirth, A. Strandlie, M. Winkler, and T. Todorov, "Recent results on adaptive track and multitrack fitting," *Nucl. Instrum. Meth.*, vol. A502, p. 702, 2003.
67. S. Catani, Y. Dokshitzer, M. Olsson, G. Turnock, and B. Webber, "New clustering algorithm for multi-jet cross-sections in  $e^+e^-$  annihilation," *Phys. Lett.*, vol. B269, p. 432, 1991.
68. S. Catani, Y. Dokshitzer, M. Seymour, and B. Webber, "Longitudinally invariant  $k_{\perp}$  clustering algorithms for hadron-hadron collisions," *Nucl. Phys.*, vol. B406, p. 187, 1993.
69. S. Moretti, L. Lönnblad, and T. Sjöstrand, "New and old jet clustering algorithms for electron-positron events," *J. High Energy Phys.*, vol. 08, p. 001, 1998.
70. W. Bartel *et al.*, "Experimental Studies on Multi-Jet Production in  $e^+e^-$  Annihilation at PETRA Energies," *Z. Phys.*, vol. C33, p. 23, 1986. [53(1986)].
71. Y. Dokshitzer, G. Leder, S. Moretti, and B. Webber, "Better jet clustering algorithms," *J. High Energy Phys.*, vol. 08, p. 001, 1997.
72. M. Seymour and C. Tevlin, "A Comparison of two different jet algorithms for the top mass reconstruction at the LHC," *J. High Energy Phys.*, vol. 11, p. 052, 2006.
73. M. Cacciari and G. Salam, "Dispelling the  $N^3$  myth for the  $k_t$  jet-finder," *Phys. Lett.*, vol. B641, p. 57, 2006.
74. T. Sjöstrand, "The Lund Monte Carlo for  $e^+e^-$  Jet Physics," *Comput. Phys. Commun.*, vol. 28, p. 229, 1983.
75. S. Bethke, Z. Kunszt, D. Soper, and W. Stirling, "New jet cluster algorithms: Next-to-leading order QCD and hadronization corrections," *Nucl. Phys.*, vol. B370, p. 310, 1992. [Erratum: *Nucl. Phys.* B523,681(1998)].
76. J. Dorfan, "A Cluster Algorithm for the Study of Jets in High-Energy Physics," *Z. Phys.*, vol. C7, p. 349, 1981.
77. J. Kruskal, "On the shortest spanning subtree of a graph and the traveling salesman problem," *Proc. Amer. Math. Soc.*, vol. 7, no. 1, p. 48, 1956.
78. L. Angelini, G. Nardulli, L. Nitti, M. Pellicoro, D. Perrino, and S. Stramaglia, "Deterministic annealing as a jet clustering algorithm in hadronic collisions," *Phys. Lett.*, vol. B601, p. 56, 2004.
79. S. Chekanov, "A New jet algorithm based on the k-means clustering for the reconstruction of heavy states from jets," *Eur. Phys. J.*, vol. C47, p. 611, 2006.
80. S.-L. Blyth *et al.*, "A Cone jet-finding algorithm for heavy-ion collisions at LHC energies," *J. Phys.*, vol. G34, p. 271, 2007.

81. F. Tarrade, "Reconstruction and identification of hadronic tau decays in ATLAS," *Nucl. Phys. Proc. Suppl.*, vol. 169, p. 357, 2007.
82. L. Barbone, N. De Filippis, O. L. Buchmüller, F. P. Schilling, T. Speer, and P. Vanlaer, "Impact of CMS silicon tracker misalignment on track and vertex reconstruction," *Nucl. Instrum. Meth.*, vol. A566, p. 45, 2006.
83. S. Blusk, O. Buchmüller, A. Jacholkowski, T. Ruf, J. Schieck, and S. Viret, eds., *Proceedings of the first LHC Detector Alignment Workshop, CERN, Geneva, Switzerland, 4–6 September 2006*, 2007.
84. V. Blobel, "Software alignment for tracking detectors," *Nucl. Instrum. Meth.*, vol. A566, p. 5, 2006.
85. W. Wiedenmann, "Alignment of the ALEPH tracking devices," *Nucl. Instrum. Meth.*, vol. A323, p. 213, 1992.
86. A. Andreazza and E. Piotto, "The Alignment of the DELPHI Tracking Detectors," Tech. Rep. DELPHI 99–153 TRACK 94, CERN, 1999.
87. A. Sopczak, "Alignment of the central D0 Detector," *Nucl. Instrum. Meth.*, vol. A566, p. 142, 2006.
88. Y. Fisyak *et al.*, "Overview of the inner silicon detector alignment procedure and techniques in the RHIC/STAR experiment," *J. Phys. Conf. Ser.*, vol. 119, p. 032017, 2008.
89. D. Brown, A. Gritsan, Z. Guo, and D. Roberts, "Local Alignment of the BABAR Silicon Vertex Tracking Detector," *Nucl. Instrum. Meth.*, vol. A603, p. 467, 2009.
90. P. Schleper, G. Steinbrück, and M. Stoye, "Alignment of the CMS silicon tracker using Millepede II," *J. Phys. Conf. Ser.*, vol. 119, p. 032040, 2008.
91. M. Gersabeck, "Alignment of the LHCb Vertex Locator," *Nucl. Instrum. Meth.*, vol. A598, p. 71, 2009.
92. S. Gonzalez-Sevilla, "Track-based alignment of the ATLAS inner detector," *J. Phys. Conf. Ser.*, vol. 119, p. 032019, 2008.
93. E. Widl and R. Frühwirth, "A large-scale application of the Kalman alignment algorithm to the CMS tracker," *J. Phys. Conf. Ser.*, vol. 119, p. 032038, 2008.
94. S. Chatrchyan *et al.*, "Alignment of the CMS tracker with LHC and cosmic ray data," *JINST*, vol. 9, p. P06009, 2014.
95. G. Mittag, "Alignment of the CMS Tracker: Latest results from LHC Run-II," *J. Phys. Conf. Ser.*, vol. 898, no. 4, p. 042014, 2017.
96. J. Schieck, "Track-based alignment for the ATLAS Inner Detector Tracking System," *JINST*, vol. 7, p. C01012, 2012.
97. G. Ripellino, "The alignment of the ATLAS Inner Detector in Run-2," *PoS*, vol. LHCP2016, p. 196, 2016.
98. M. Martinelli, "Novel real-time alignment and calibration of the LHCb detector in Run2," *J. Phys. Conf. Ser.*, vol. 898, no. 3, p. 032039, 2017.
99. J. Amoraal, "Alignment of the LHCb detector with Kalman filter fitted tracks," *J. Phys. Conf. Ser.*, vol. 219, p. 032028, 2010.
100. A. Dainese, "Alignment of the ALICE tracking detectors," *PoS*, vol. VERTEX2009, p. 021, 2009.
101. V. Karimäki, T. Lampen, and F. Schilling, "The HIP algorithm for track based alignment and its application to the CMS pixel detector," Tech. Rep. CMS-NOTE-2006-018, CERN, 2006.
102. V. Blobel and C. Kleinwort, "A New method for the high precision alignment of track detectors," in *Advanced Statistical Techniques in Particle Physics. Proceedings, Conference, Durham, UK, March 18–22, 2002*, 2002.
103. Y. Saad and M. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on scientific and statistical computing*, vol. 7, no. 3, p. 856, 1986.
104. E. Widl, R. Frühwirth, and W. Adam, "A Kalman filter for track-based alignment," tech. rep., CERN, 2006.
105. R. Frühwirth, T. Todorov, and M. Winkler, "Estimation of detector alignment parameters using the Kalman filter with annealing," *J. Phys.*, vol. G29, p. 561, 2003.

106. R. Gluckstern, "Uncertainties in track momentum and direction, due to multiple scattering and measurement errors," *Nucl. Instrum. Meth.*, vol. 24, p. 381, 1963.
107. M. Regler and R. Frühwirth, "Generalization of the Gluckstern formulas. I: Higher orders, alternatives and exact results," *Nucl. Instrum. Meth.*, vol. A589, p. 109, 2008.
108. R. Frühwirth and A. Beringer, "JDOT — a Java detector optimization tool," in *Nuclear Science Symposium Conference Record, 2008. NSS'08. IEEE*, pp. 3483–3487, IEEE, 2008.
109. W. Innes, "Some formulas for estimating tracking errors," *Nucl. Instrum. Meth.*, vol. A329, p. 238, 1992.
110. W. Innes, "TRACKERR: A Program for calculating tracking errors," Tech. Rep. SLAC-BABAR-NOTE-121, SLAC, 1993.
111. M. Regler, M. Valentan, and R. Frühwirth, "The LiC detector toy program," *Nucl. Instrum. Meth.*, vol. A581, p. 553, 2007.
112. S. Masciocchi, "Experience with HERA-B vertexing," *Nucl. Instrum. Meth.*, vol. A462, p. 220, 2001.
113. E. Chabanat, J. D'Hondt, N. Estre, R. Frühwirth, K. Prokofiev, T. Speer, P. Vanlaer, and W. Waltenberger, "Vertex reconstruction in CMS," *Nucl. Instrum. Meth.*, vol. A549, p. 188, 2005.
114. W. Erdmann, "Vertexing in the H1 experiment," *Nucl. Instrum. Meth.*, vol. A560, p. 89, 2006.
115. M. Costa, "Vertex and track reconstruction in ATLAS," *Nucl. Instrum. Meth.*, vol. A582, p. 785, 2007.
116. W. Adam, "Track and vertex reconstruction in CMS," *Nucl. Instrum. Meth.*, vol. A582, p. 781, 2007.
117. W. Waltenberger, *Development of vertex finding and vertex fitting algorithms for CMS*. PhD thesis, TU Wien, 2004.
118. D. Jackson, "A Topological vertex reconstruction algorithm for hadronic jets," *Nucl. Instrum. Meth.*, vol. A388, p. 247, 1997.
119. S. Hillert, "ZVMST: A Minimum spanning tree-based vertex finder," tech. rep., 2008.
120. W. Waltenberger and F. Moser, "RAVE — an Open, Extensible, Detector-Independent Toolkit for Reconstruction of Interaction Vertices," in *IEEE Nuclear Science Symposium Conference Record 2006*, vol. 1, p. 104, IEEE, 2006.
121. W. Waltenberger, W. Mitaroff, and F. Moser, "RAVE — a Detector-independent vertex reconstruction toolkit," *Nucl. Instrum. Meth.*, vol. A581, p. 549, 2007.
122. G. Patrick and B. Schorr, "Vertex Fitting of Several Helices in Space," *Nucl. Instrum. Meth.*, vol. A241, p. 132, 1985.
123. P. Billoir and S. Qian, "Fast vertex fitting with a local parametrization of tracks," *Nucl. Instrum. Meth.*, vol. A311, p. 139, 1992.
124. P. Billoir and S. Qian, "Erratum to Fast vertex fitting with a local parametrization of tracks," *Nucl. Instrum. Meth.*, vol. A350, p. 624, 1994.
125. D. Bates and D. Watts, *Nonlinear regression analysis and its applications*. Wiley & Sons, 1988.
126. V. Karimäki, "Effective Vertex Fitting," Tech. Rep. CMS-NOTE-1997-051, CERN, 1997.
127. E. Calligarich, R. Dolfini, M. Genoni, and A. Rotondi, "A Fast algorithm for vertex estimation," *Nucl. Instrum. Meth.*, vol. A311, p. 151, 1992.
128. F. Hampel, E. Ronchetti, P. Rousseeuw, and W. Stahel, *Robust statistics: the approach based on influence functions*. John Wiley & Sons, 2011.
129. P. Rousseeuw and A. Leroy, *Robust regression and outlier detection*. John Wiley & sons, 2005.
130. R. Frühwirth, P. Kubinec, W. Mitaroff, and M. Regler, "Vertex reconstruction and track bundling at the LEP collider using robust algorithms," *Comput. Phys. Commun.*, vol. 96, p. 189, 1996.
131. P. Huber and E. Ronchetti, *Robust statistics*. Wiley & Sons, 2 ed., 2009.
132. G. Agakichiev et al., "A new robust fitting algorithm for vertex reconstruction in the CERES experiment," *Nucl. Instrum. Meth.*, vol. A394, p. 225, 1997.

133. K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems," *Proceedings of the IEEE*, vol. 86, no. 11, p. 2210, 1998.
134. T. Speer, R. Frühwirth, P. Vanlaer, and W. Waltenberger, "Robust vertex fitters," *Nucl. Instrum. Meth.*, vol. A566, p. 149, 2006.
135. W. Waltenberger, R. Frühwirth, and P. Vanlaer, "Adaptive vertex fitting," *J. Phys.*, vol. G34, p. N343, 2007.
136. A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society. Series B (methodological)*, vol. 39, no. 1, p. 1, 1977.
137. R. Frühwirth and W. Waltenberger, "Redescending M-estimators and deterministic annealing," *Austrian J. Statistics*, vol. 37, no. 3&4, p. 301, 2008.
138. J. D'Hondt, R. Frühwirth, P. Vanlaer, and W. Waltenberger, "Sensitivity of robust vertex fitting algorithms," *IEEE Transactions on Nuclear Science*, vol. 51, no. 5, p. 2037, 2004.
139. R. Frühwirth and W. Waltenberger, "Adaptive Multi-vertex fitting," Tech. Rep. CMS CR 2004/062, CERN, 2004.
140. T. Speer and R. Frühwirth, "A Gaussian-sum filter for vertex reconstruction," *Comput. Phys. Commun.*, vol. 174, p. 935, 2006.
141. S. Lang, *Calculus of several variables*. Springer Science & Business Media, 2012.
142. P. Avery, "Applied fitting theory VI: Formulas for kinematic fitting," Tech. Rep. CLEO Note CBX98-37, Cornell University, 1999.
143. ALICE Collaboration, "The ALICE experiment at the CERN LHC," *JINST*, vol. 3, no. 08, p. S08002, 2008.
144. B. Abelev *et al.*, "Performance of the ALICE Experiment at the CERN LHC," *Int. J. Mod. Phys.*, vol. A29, p. 1430044, 2014.
145. G. Aad *et al.*, "The ATLAS Experiment at the CERN Large Hadron Collider," *JINST*, vol. 3, p. S08003, 2008.
146. T. Cornelissen *et al.*, "Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)," Tech. Rep. ATL-SOFT-PUB-2007-007, CERN, Geneva, 2007.
147. ATLAS Collaboration, "Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2," *Eur. Phys. J.*, vol. C77, no. 10, p. 673, 2017.
148. G. Aad *et al.*, "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at  $\sqrt{s} = 13$  TeV," *Eur. Phys. J.*, vol. C76, no. 5, p. 292, 2016.
149. H. Oide, "Improvements to ATLAS track reconstruction for Run-2," *PoS*, vol. EPS-HEP2015, p. 287, 2015.
150. K. Potamianos, "The upgraded Pixel detector and the commissioning of the Inner Detector tracking of the ATLAS experiment for Run-2 at the Large Hadron Collider," *PoS*, vol. EPS-HEP2015, p. 261, 2015.
151. ATLAS Collaboration, "The Optimization of ATLAS Track Reconstruction in Dense Environments," Tech. Rep. ATL-PHYS-PUB-2015-006, CERN, Geneva, Mar 2015.
152. G. Aad *et al.*, "Charged-particle distributions in  $\sqrt{s} = 13$  TeV pp interactions measured with the ATLAS detector at the LHC," *Phys. Lett.*, vol. B758, pp. 67–88, 2016.
153. S. Chatrchyan *et al.*, "The CMS Experiment at the CERN LHC," *JINST*, vol. 3, p. S08004, 2008.
154. CMS Collaboration, "The CMS tracker system project: Technical Design Report," Tech. Rep. CMS-TDR-005, CERN, 1997.
155. E. Brondolin, *Track reconstruction in the CMS experiment for the High Luminosity LHC*. PhD thesis, Technische Universität Wien, 2018.
156. S. Chatrchyan *et al.*, "Description and performance of track and primary-vertex reconstruction with the CMS tracker," *JINST*, vol. 9, no. 10, p. P10009, 2014.
157. CMS Collaboration, "2017 tracking performance plots," Tech. Rep. CMS-DP-2017-015, CERN, 2017.
158. F. Pantaleo, *New Track Seeding Techniques for the CMS Experiment*. PhD thesis, Universität Hamburg, 2017.
159. A. Augusto Alves *et al.*, "The LHCb Detector at the LHC," *JINST*, vol. 3, p. S08005, 2008.

160. R. Aaij *et al.*, “LHCb Detector Performance,” *Int. J. Mod. Phys.*, vol. A30, no. 07, p. 1530022, 2015.
161. LHCb Collaboration, “Measurement of the track reconstruction efficiency at LHCb,” *JINST*, vol. 10, no. 02, p. P02007, 2015.
162. LHCb Collaboration. Published online at [https://twiki.cern.ch/twiki/pub/LHCb/ConferencePlots/TrackEffPLong2015\\_2012.pdf](https://twiki.cern.ch/twiki/pub/LHCb/ConferencePlots/TrackEffPLong2015_2012.pdf), 2017.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 14

## Distributed Computing



**Manuel Delfino**

Distributed computing is an established discipline in computer science and engineering. It has evolved over the past 40 years to be one of the most important methodologies for implementing the data processing services needed by almost every activity in society. Distributed computing is still evolving rather rapidly, with major innovations introduced every few years. The aim of this chapter is to introduce the reader to the basic concepts of distributed computing in the context of particle physics, allowing a better understanding of what is behind the scenes when using distributed systems, and providing starting points in order to seek further information on the subject.

### 14.1 Usage by Particle Physics

Distributed computing, i.e. the coherent use of many computers to accomplish a given task, is extensively used in particle physics. It is particularly suited to simulating and analyzing data from particle collisions, where processing the data of one collision is largely independent from processing data of all other collisions and hence very little communication between the different computers is necessary. This type of distributed computing is called “High Throughput Computing” (HTC),<sup>1</sup> because there is a very large number of quasi-independent tasks to accomplish,

---

<sup>1</sup>Livny, Miron, et al. “Mechanisms for high throughput computing.” *SPEEDUP journal* 11.1 (1997): 36–40.

---

M. Delfino (✉)

Departament de Física, Facultat de Ciències, Universitat Autònoma de Barcelona, Barcelona, Spain

e-mail: [manuel.delfino@uab.cat](mailto:manuel.delfino@uab.cat)

and the performance perceived by the user is the rate of completion of tasks, or “throughput”. This type of computing is complimentary to “High Performance Computing” (HPC), i.e. the execution of a single task at maximum speed on a classical supercomputer. The particular case of event or collision processing requiring little inter-processor communication is an example of “loosely coupled parallel computing” or even “embarrassingly parallel computing”.<sup>2,3</sup> Certain kinds of accelerator simulations, particularly those that use ray-tracing techniques, can also be executed as distributed computing tasks. Multiple computer systems can also be interconnected in “tightly-bound” configurations via low-latency networks,<sup>4</sup> essentially yielding a supercomputer. Tightly-bound systems are used in particle physics, for example, for solving equations from theories using numerical methods (e.g. lattice-gauge theories); they are mentioned here for completeness and will not be covered further.

From the decade of the 2010s, the boundary between HTC and HPC has started to blur, for a number of disparate reasons:

*Multi-core Applications* Detectors have become more complex and have more channels. This impacts the memory footprint of reconstruction, simulation and analysis programs, which has grown substantially. In parallel, processors have become multi-core<sup>5</sup> with shared RAM.<sup>6</sup> The growth in number of cores per processor has been faster than the drop in price of RAM chips, resulting in an effective memory shortage. This means that running an independent copy of the operating system and the application in each core is not economical. The way out is to implement the safe execution of parallel threads<sup>7</sup> of a single copy of the program on each multi-core processor. This introduces a number of dependencies, for example competition for RAM and for input–output services, that break the loose coupling and make the applications behave more like HPC programs.

---

<sup>2</sup>Wilkinson, Barry, and Michael Allen. Parallel programming. Pearson India, 2004.

<sup>3</sup>Birritella, Mark S., et al. “Intel® Omni-path architecture: Enabling scalable, high performance fabrics.” High-Performance Interconnects (HOTI), 2015 IEEE 23rd Annual Symposium on. IEEE, 2015.

<sup>4</sup>Pfister, Gregory F. “An introduction to the infiniband architecture.” High Performance Mass Storage and Parallel I/O 42 (2001): 617–632.

<sup>5</sup>A core corresponds to the Central Processing Unit (CPU) of a classical computer, which is able to execute a single stream of instructions. Technology allows a growing number of ever smaller transistors to be placed on a single chip. However, profiting from these large number of transistors in a single core design would require prohibitively complex processor designs and impossible to achieve clock speeds. The alternative is to populate the chips with many copies of the same processor core, which work independently except for sharing external connections, leading to multi-core processors. Gepner, Pawel, and Michal Filip Kowalik. “Multi-core processors: New way to achieve high system performance.” Parallel Computing in Electrical Engineering, 2006. PAR ELEC 2006. International Symposium on. IEEE, 2006.

<sup>6</sup>Random Access Memory, the external solid state memory used by a processor.

<sup>7</sup>A thread is a stream of processing instructions coming from a shared program image which has its own private data instances in processor hardware registers, instruction and data stack and caches, and RAM.

*Supercomputers as High Core-Count Clusters* The original design of supercomputers was based on a rather small number of the fastest single core processors available at the time, interconnected with custom very low latency links. In addition, the design usually included a much larger amount of RAM than in standard computers. The design of supercomputers has changed completely over the last 20 years.<sup>8</sup> The main driver of this change is the difficulty in building processors with ever shorter clock cycles. The so-called “clock speed wall” has been hit, and practically all processors operate within a narrow range around a few GHz. Hence, the only way to make a supercomputer faster is through parallelism. Modern supercomputers are in fact huge clusters of relatively standard multi-core processors. Low latency links continue to be used and their hardware costs have become much lower. In addition, “latency hiding” techniques sometimes allow the use of standard networks, such as Ethernet. Finally, the amount of RAM per core does not differ much from standard computers. Hence, executing an HTC workload on a supercomputer targeted for HPC is no longer wasteful. At worst, only the low latency interconnect will be underutilized.

*Increase of Workloads with Low Input–Output* Particle physics has traditionally been a heavy user of HTC systems because they were the least expensive architecture for executing detector track reconstruction and analysis, applications with a relatively low CPU to input–output ratio. The needs for CPU for reconstruction and analysis have grown, however, as detectors have become much more complex. In parallel, the precision needed in the simulations has vastly increased. These high CPU, low input–output applications currently represent the largest computing demand of a modern particle physics detector. Hence, the global workload profile has moved closer to HPC in the last decades.

## 14.2 Functional Decomposition of a Distributed Computing Environment

It is useful to introduce a functional decomposition, or reference framework, to discuss distributed computing systems. The decomposition that has dominated particle physics data processing since the 1990s, is the “SHIFT Model”, introduced in the early 1990s by Robertson and collaborators.<sup>9</sup> It will be described here with some updates of terminology. The basic assumption in the SHIFT framework is that a Local Area Network (LAN) can be built with enough capacity and flexibility so that the rest of the elements in the distributed environment can

---

<sup>8</sup>Xie, Xianghui, et al. “Evolution of supercomputers.” *Frontiers of Computer Science in China* 4.4 (2010): 428–436.

<sup>9</sup>J.P. Baud, et al. “SHIFT, the Scalable Heterogeneous Integrated Facility for HEP Computing”, *Proc. Conference on Computing in the High Energy Physics, CHEP91, Tsukuba, Japan, Universal Academic Press.*



communicate freely amongst themselves. Furthermore, the environment is loosely coupled and is therefore not sensitive to relatively long round-trip time for network messaging (large network latency). Connected to the network we have the following elements:

- *CPU servers*: Elements that receive input data, perform calculations and produce output data. They do not implement any permanent storage, although they often provide volatile disk storage used to temporarily store data while executing a task.
- *Disk servers*: Elements that store data in a stable and reliable manner with an access latency (defined as time to open a file and receive the first byte of data) which is relatively low. They provide inputs to and receive outputs from the CPU servers. They can also send and receive data to Tape server elements.
- *Tape servers*: Elements that store data in a stable and reliable manner with an access latency which is relatively high. They send and receive data to Disk server elements.
- *Information servers*: Elements that maintain data in a stable and reliable manner about the status of the various elements of the distributed computing environment.
- *Control servers*: Elements that issue commands to trigger operations in other types of servers and coherently update the relevant Information servers.
- *Remote Data servers*: Elements that send and receive data from other security domains, often to Disk server elements via Wide-Area Network (WAN) connections.


The elements of the functional decomposition presented, illustrated in Fig. 14.1, together with the aforementioned powerful network and the maintenance of coherent security<sup>10</sup> domains, constitute a framework which is sufficient to analyze the most commonly used distributed computing environments.

The computing industry places a lot of emphasis on how storage is attached to the computers: Direct-Attached Storage (DAS) does not use a network, attaching storage hardware via an internal communication bus of a computer; Storage Area Network (SAN) uses a short-distance network (originally Fiberchannel, with an increasing tendency towards Ethernet) to couple storage hardware to computers; and Network Attached Storage (NAS) which uses computers hosting storage devices and connected to a LAN (or WAN configured to behave like a LAN) to present storage services to the other computers on the network. These distinctions are important in environments where few computers are used. In the case of particle physics, DAS and SAN are used to construct disk servers, which in turn are exposed almost exclusively in NAS mode.

---

<sup>10</sup>A more precise term is “Authentication and Authorization Infrastructure”, abbreviated as AAI or AA.

**Fig. 14.1** Functional decomposition of a distributed computing environment



### 14.3 Data Processing Clusters

The most common way to deploy a distributed computing environment is a “cluster”.<sup>11</sup> A cluster consists of a number of CPU, Disk and Control servers connected on a LAN<sup>12</sup> and sharing status information through local Information servers while maintaining a *single, coherent, security domain*. A cluster may also implement Tape servers and Remote Data servers.

Historically, the first complete and reliable commercial implementation of a cluster was the VAXcluster<sup>13</sup> (or VMScluster), implemented by Digital Equipment Corp. starting in 1983 on VAX computers and workstations using the VMS operating system. Today, essentially all clusters used by particle physics are implemented using the Linux operating system<sup>14</sup> as a basis, and adding a number of additional packages for handling authentication and storage.

<sup>11</sup>Unfortunately, after many decades of using the term “cluster” as in the present text, some commercial firms have taken to using the terms “Grid” or “Cloud” to mean “cluster”. This is considered by the author as incorrect and confusing.

<sup>12</sup>There are also implementations using several LANs transparently interconnected via WANs, but maintaining a single security domain. This configuration is often used in business environments to gain fault-tolerance and disaster recovery. Problems related to performance limitations of the WAN, as well as inter-institutional security management complications, limit the use of this configuration in particle physics data processing. Interconnection of clusters via a Grid or a federated Cloud is used instead.

<sup>13</sup>Kronenberg, Nancy P., Henry M. Levy, and William D. Strecker. “VAXcluster: a closely-coupled distributed system.” *ACM Transactions on Computer Systems (TOCS)* 4.2 (1986): 130–146.

<sup>14</sup>Linux (or more precisely GNU/Linux) is mentioned specifically as it has become the dominant Unix-like operating system, and has become better known than the original Unix operating system. Most of the remarks in this chapter, however, are equally applicable to other Unix-like systems, notably macOS and FreeBSD.

### 14.3.1 Authentication and User Identification

The machinery behind authentication and user identification is often poorly understood by users, which leads to usage and security problems. It is important to gain an understanding of the basic concepts, especially as more complex distributed environments, such as Grids and Clouds (see sections below), have come into widespread use.

A number of Information servers are deployed to integrate individual computers to form a cluster.<sup>15</sup> The most crucial ones are the ones that generate a single, coherent, security domain.<sup>16</sup> Security under Linux and most other operating systems is based on the concept of an “account” which, ideally, is used exclusively by a single trusted individual. An account is represented externally by an alphanumeric string (the “username”) and internally by a numeric code (the *userid* or *uid*). Although modern operating systems allow very long usernames, local restrictions may apply in order to maintain compatibility with older system software and utilities (the most common restriction is a maximum length of eight characters). Each account is associated to specific directives (called “rights” or “privileges”) to allow or disallow access to operating system services, and can also be used to control access to files and other resources directly or via Access Control Lists (“ACL”). Since people often work in teams, it is useful to associate an account with a “group” named by an alphanumeric string and represented internally by a numeric code (the *groupid* or *gid*). Certain rights and, very importantly, file access control can thus be quickly managed for all accounts belonging to a group.

The process of authentication involves an exchange of credentials which establishes the identity of the user wanting to access a system which results in the creation of a process running under the corresponding *uid*. Authentication has been traditionally accomplished by a “password”, an alphanumeric string which, ideally, is known only to the individual owner of an account. Password related issues are major contributors to computer security problems. Practically all computers are nowadays connected directly or indirectly to the Internet. An inherent weakness of a cluster (which is more than compensated by the gained functionality) is that having a single security domain means that gaining access to any component of the cluster grants access to all of it. The components of clusters directly connected to the Internet are under constant attacks, the most common being automated attempts to guess passwords (usernames are easily obtained from public information sources, such as Web pages listing email addresses). Attacks from within the cluster must also be considered, especially in clusters with many accounts. A

---

<sup>15</sup>The term Information server is used in a broad manner. For example, Domain Name Service (DNS) servers which translate alphabetic Internet addresses to numerical, and Network Time servers which ensure all cluster elements have their clocks synchronized, are considered Information servers.

<sup>16</sup>The actual implementation is a highly technical matter, using NIS, ldap, Microsoft Active Directory or other secure database sharing schemes.

password which is easily guessed is denominated “weak”. Simple measures exist to avoid weak passwords.<sup>17</sup> The password should be as long as possible and contain a mix of numbers and upper and lower case letters. Many installations require users to change their passwords periodically in order to improve computer security. A very dangerous practice which should be completely avoided, but which is unfortunately common practice in the particle physics community, is the “sharing” of personal accounts (and hence, their passwords) or the use of “service” accounts with passwords known by dozens of persons, written in documents and blackboards or even posted on Web sites. Modern operating system features make these practices completely unnecessary, as they can be configured so that each user can first authenticate with their personal account and then gain access to a common environment to perform the tasks required.

Portable devices (“laptops” and “smartphones”) can be deployed so that they are elements of a cluster and therefore part of the security domain. Since the portable computer must remain useable when temporarily disconnected from the network, the operating system will keep a local copy (or “cache”<sup>18</sup>) of the security data, including the passwords. This means that a stolen portable computer is a computer security threat. Care should be taken in configuring the “suspended” or “hibernated” modes in portable computers to ask for a password when they are turned back on.

Another major issue with authentication based on usernames and passwords is that each cluster (and each online service such as electronic mail, social networks, and document and file repositories), being a separate security environment, requires a separate username and password and has its own policies for requiring password changes. Users are becoming overwhelmed in keeping track of their usernames and passwords and their reaction is often to use weak passwords or keep a list of passwords in a file which can be stolen. In order to reduce username/password proliferation, many organizations are linking their clusters and online services to shared authentication servers, thus providing a “single-sign-on” environment.

In the 2000s, organizations, most notably the CERN Large Hadron Collider experiments, started to use an alternative authentication and authorization frame-

---

<sup>17</sup>Users should consult their local security rules for specific recommendations on password choices.

<sup>18</sup>The term “cache” is used in computer architecture to describe a local copy of a limited amount of data which is normally stored elsewhere. The introduction of this local copy serves to increase the performance and efficiency of use of the element connected to it, by smoothing out peaks and valleys in storing and retrieving information, or avoiding retrieves altogether if the same information is needed repetitively. The cache is transparent to the element accessing the information; if the data needed is in cache, it will be delivered from there (a cache “hit”) whereas if the data is not in cache, the normal access path to the normal storage place will be used (a cache “miss”). The key to successful implementation of a cache is the right choice of algorithm to choose which data to copy and keep in the cache. Examples are high-speed memory caches in CPU chips which hold local copies of instructions and data, and memory caches in disk controllers which hold local copies of disk blocks.

work based on digital certificates.<sup>19</sup> This was one of the pillars for constructing the Worldwide LHC Computing Grid (WLCG), later generalized to the European Grid Infrastructure (EGI), the U.S. Open Science Grid (OSG) and others. The certificates used are based on the X.509 standard, supported by all major Web servers and browsers. These schemes were very successful and are still being used, allowing tens of thousands of computers to provide coherent services to thousands of users. They are, however, extremely difficult to manage and expensive to operate. Hence, the tendency is to abandon these schemes in user-facing services, keeping them mostly for machine-to-machine services.

Industry and academia are putting a lot of effort on the emerging area of “distributed” authentication and authorization (AA) schemes.<sup>20</sup> The ultimate aim is to simplify the AA process and make it easier to use, enabling “single-sign-on” to a wide variety of resources, a functionality that users have come to expect from their experience with social networks. Unfortunately, progress has been rather slow and solutions proposed by industry and academia are often incompatible. Part of the problem is that there are two competing standards: SAML<sup>21</sup> and OpenID.<sup>22</sup>

The main problem, however, is that the usage by large collaborations, such as those in particle physics, ideally requires the simultaneous use of many different AA sources (something known as Identity Federation). The largest Identity Federation currently deployed is the *eduGAIN* federation,<sup>23</sup> where the AA sources come from the universities and research institutes who employ the users. *eduGAIN* is modeled after the successful *eduroam* system<sup>24</sup> used to grant worldwide access to academic WiFi networks worldwide. Solving the general access problem is much more difficult, however, as it requires the maintenance of many more attributes for each user. For example, each university and research institute would have to include and maintain in their databases which experiment or project each employee is participating in, something which is not practical. In order to solve this issue, hybrid schemes are being worked on, where the authentication would come from *eduGAIN* but the authorization information would come from an attribute server managed and operated by a specific experiment or project. This still leaves the problem of reliably operating a service build from thousands of independently managed AA servers with varying degrees of service quality.

---

<sup>19</sup>Thompson, Mary R., Abdelilah Essiari, and Srilekha Mudumbai. “Certificate-based authorization policy in a PKI environment.” *ACM Transactions on Information and System Security (TISSEC)* 6.4 (2003): 566–588.

<sup>20</sup>These schemes centralize AA information and then make it available to a distributed set of heterogeneous resources, hence the term “distributed”.

<sup>21</sup>Rosenberg, Jonathan B., and David L. Remy. *Securing web services with WS-security: Demystifying WS-security, WS-policy, SAML, XML signature, and XML encryption*. Sams, 2004.

<sup>22</sup><https://openid.net/developers/specs/>

<sup>23</sup>López, D. “eduGAIN: Federation interoperation by design.” *TERENA Networking Conference*. 2006.

<sup>24</sup>López, Gabriel, et al. “A proposal for extending the eduroam infrastructure with authorization mechanisms.” *Computer Standards & Interfaces* 30.6 (2008): 418–423.

Some interesting alternatives are starting to emerge, which avoid Identity Federation. One approach is to identify a single reliable provider for AA, for example the OpenID service from ORCID<sup>25</sup> combined with an appropriate attribute server. Another approach is for a project or experiment to deploy a single AA scheme available to services *for that project* around the world. This can be accomplished in a simple and economical manner by re-using the project's personnel database and exposing it through the Internet in a secure manner using an *ldap* server.<sup>26</sup> These type of implementations of single-sign-on are ideal for service providers that serve a single project. On the other hand, they do shift the diversity problem to the service providers, which must locally adjust crucial attributes such as *uid* and *gid* in order to avoid duplications in multi-project or multi-experiment environments.

### 14.3.2 Processing and Storage

Two configurations of CPU servers are usually deployed:

- “Batch workers”, which execute tasks (jobs) that don't require user interaction and are scheduled using a “Batch” system<sup>27</sup> or a more sophisticated resource harvester, such as HTCondor.<sup>28</sup>
- “Interactive nodes”, where users can connect (or “log-in”) and perform work that requires interaction via a screen, keyboard and mouse.

Several configurations of Disk servers are usually deployed:

- “Network File servers”, based on industry-standard or widely used protocols, such as NFS,<sup>29</sup> smb2<sup>30</sup> or http with WebDAV,<sup>31</sup> are used to provide “home” and “project” directories holding text and binary files of relatively small size. These directories appear in the CPU servers as “mount points” or “network shares” that largely behave as a local, conventional disk resource. The POSIX input–output

<sup>25</sup>Haak, Laurel L., et al. “ORCID: a system to uniquely identify researchers.” *Learned Publishing* 25.4 (2012): 259–264.

<sup>26</sup>Johner, Heinz, et al. *Understanding LDAP*. Vol. 6. IBM, 1998.

<sup>27</sup>See for example Yoo, Andy B., Morris A. Jette, and Mark Grondona. “Slurm: Simple linux utility for resource management.” *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer, Berlin, Heidelberg, 2003.

<sup>28</sup>Fajardo, E. M., et al. “How much higher can HTCondor fly?” *Journal of Physics: Conference Series*. Vol. 664. No. 6. IOP Publishing, 2015.

<sup>29</sup>Shepler, Spencer, et al. *Network file system (NFS) version 4 protocol*. No. RFC 3530. 2003.

<sup>30</sup>French, Steven M., and Samba Team. “A New Network File System is Born: Comparison of SMB2, CIFS and NFS.” *Linux Symposium*. sn, 2007.

<sup>31</sup>Goland, Yaron, et al. *HTTP Extensions for Distributed Authoring – WEBDAV*. No. RFC 2518. 1999.

standard<sup>32</sup> is the most often used definition of this behavior. The implementations are far from trivial, but fortunately they are widely used and hence available as packages fully supported by operating systems, or even as part of the operating system. These protocols support operations such as simultaneous access by multiple users in read/write/append mode as well as record-level locking. This forces the deployment in CPU servers of some Control and Information server functionality, with information being constantly exchanged between all nodes accessing a mount point in order to keep file system coherency and distribute lock information. In order to fulfill performance expectations, these Control and Information functionalities must often be implemented as drivers, often executing inside the operating system kernel. The price of this is the risk of all CPU servers becoming inoperative (or “hung”) if a network file server has a fault or the status information becomes corrupted. Some Network File servers, such as the CernVM-FS<sup>33</sup> system used for software installation, serve very specialized purposes. They implement only part of the functionalities, for example serving files only in read-only mode, with the benefit of having a smaller impact on the operating system environment.

- “Database servers”, are disk servers<sup>34</sup> which organize their data using database products such as Oracle, MySQL or PostgreSQL. They are mostly used to store and manage detector configuration and calibration data, though on occasion they can store high-level analysis data such as the “Event Tags” used by some experiments.<sup>35</sup> CPU servers read these data through the network by issuing database queries and retrieving the results. Database servers must be configured in accordance to the type of most-often used queries (a process called “tuning”) in order to achieve the required performance. Because the overall cost of purchasing and maintaining a Database server is high, situations often arise where the deployed capacity is insufficient to serve all the CPU servers unless specific caches or buffers<sup>36</sup> are deployed. An emerging alternative type of database

---

<sup>32</sup>Gallmeister, Bill. POSIX. 4 Programmers Guide: Programming for the real world. “O’Reilly Media, Inc.”, 1995.

<sup>33</sup>Blomer, Jakob, et al. “Status and future perspectives of CernVM-FS.” *Journal of Physics: Conference Series*. Vol. 396. No. 5. IOP Publishing, 2012.

<sup>34</sup>It may seem odd to classify a database server as a Disk server. It is considered correct in this context, as their usage pattern is unusual in that the record update frequency is very low compared to the record reading frequency. Note that the technical database servers used to implement Data and Tape server features are considered Information servers.

<sup>35</sup>See for example Cranshaw, Jack, et al. “Event selection services in ATLAS.” *Journal of Physics: Conference Series*. Vol. 219. No. 4. IOP Publishing, 2010.

<sup>36</sup>The term “buffer” is used in computer architecture to mean “temporary storage location”. Buffers are used to accumulate information and then transmit it as a block, or to receive information as a block and then distribute it. The difference between a buffer and a cache is that in the case of buffers all the information is temporarily stored, whereas a cache implements an algorithm to store only the most relevant information. In addition, whereas a cache is always transparent, a buffer may be transparent or it may allow explicit buffer manipulation by the application.



servers are those based on “Big Data” tools, such as NoSQL databases<sup>37</sup> with underlying *Hadoop* storage management systems.<sup>38</sup>

- “Data servers”, based on products such as dCache,<sup>39</sup> GPFS,<sup>40</sup> Lustre,<sup>41</sup> EOS<sup>42</sup> and DPM,<sup>43</sup> are used to hold the large-size binary files which contain physics data. They use a set of specific Control and Information Servers to present to the other elements of the cluster, in particular to the CPU servers, an interface approximating a conventional (POSIX-compliant) disk resource, but with some important differences in order to achieve the large capacities and high performance needed for particle physics applications (and data-intensive applications in other fields). Much of the optimization is possible because of the particular usage pattern of particle physics data, which is essentially written once and then read many times.<sup>44</sup> The key point is to separate<sup>45</sup> the “namespace” information (directory tree or file folder structures, file attributes such as name, creation date, etc.), the “file storage” information (physical location of the blocks that comprise a file, whether it is open for read, write or append, etc.) and the “file data” information (the actual data blocks that comprise the file). This separation makes it possible to serve the data blocks from a large number of Data servers with a combined capacity far exceeding what can be provided by a conventional file system. In addition, the data blocks of a given file may be stored in multiple servers in order to enhance read performance or provide high-availability (for reading).<sup>46</sup> The cost of this separation (apart from the complexity of the multiple

---

<sup>37</sup>Han, Jing, et al. “Survey on NoSQL database.” Pervasive computing and applications (ICPCA), 2011 6th international conference on. IEEE, 2011.

<sup>38</sup>Shvachko, Konstantin, et al. “The hadoop distributed file system.” Mass storage systems and technologies (MSST), 2010 IEEE 26th symposium on. Ieee, 2010.

<sup>39</sup>Millar, Paul, et al. “Storage for advanced scientific use-cases and beyond.” Parallel, Distributed and Network-based Processing (PDP), 2018 26th Euromicro International Conference on. IEEE, 2018.

<sup>40</sup>Schmuck, Frank B., and Roger L. Haskin. “GPFS: A Shared-Disk File System for Large Computing Clusters.” FAST. Vol. 2. No. 19. 2002.

<sup>41</sup>Schwan, Philip. “Lustre: Building a file system for 1000-node clusters.” Proceedings of the 2003 Linux symposium. Vol. 2003. 2003.

<sup>42</sup>Peters, A. J., E. A. Sindrilariu, and G. Adde. “EOS as the present and future solution for data storage at CERN.” Journal of Physics: Conference Series. Vol. 664. No. 4. IOP Publishing, 2015.

<sup>43</sup>Alvarez, Alejandro, et al. “DPM: future proof storage.” Journal of Physics: Conference Series. Vol. 396. No. 3. IOP Publishing, 2012.

<sup>44</sup>Aside from the write-once/read-many access pattern, particle physics datasets almost never have data appended to them nor have their already existing records updated. For an extreme contrast, consider the access pattern of bank account datasets, which constantly have data appended and updated.

<sup>45</sup>Conventional file systems, whether local or networked, tightly bind the management of the three types of information described in order to optimize file open/close operations and implement sharing of files in read/write/append mode.

<sup>46</sup>This is conceptually similar to the “striping” of files across many disks performed by RAID controllers on DAS devices.



Information Servers) is a relatively high overhead and latency for file open/close operations and the loss of full POSIX compliance. In some cases, applications have to use input–output methods specific to the data server product, such as *Xrootd*,<sup>47</sup> whereas in other cases standard protocols, such as NFS or http or subsets of POSIX I/O can be used. Historically, namespaces have been confined to single clusters. There are, however, some interesting implementations of global namespaces, using for example a hierarchy of distributed name servers coupled through redirection<sup>48</sup> techniques offered by various protocols such as *Xrootd*, WebDAV or http.<sup>49</sup> Many commercial systems use similar implementations, but with different design criteria, for example the requirement of serving hundreds of thousands of users at relatively low performance in a social network.

Tape servers are hidden from the user,<sup>50</sup> except in very special cases such as data recording of experimental data. Nevertheless, users and system managers should be vigilant, as the user may generate access patterns which make extremely inefficient use of the underlying tape cartridge system. Tape servers and their related tape cartridge systems are quite complex technologically. Exploring this is beyond the scope of this chapter, but a few remarks are appropriate. The main motivation for introducing Tape servers is that they provide storage which is less expensive than Disk servers, with additional features such as guaranteed read-only access, protection against erroneous deletion by users, lower probability of data loss due to hardware failure and lower power usage. As their name implies, tape servers have been historically implemented using storage on magnetic tape cartridges. Implementations for particle physics are of the “Active Tape” nature, using specific software to manage the tape in a much more agile manner than traditional tape backup systems, as well as providing file by file access. Different computer centers use different software packages for Active Tape management. Examples are HPSS<sup>51</sup> and Tivoli Storage Manager,<sup>52</sup> both developed by IBM, Enstore<sup>53</sup> developed by Fermilab and Castor2<sup>54</sup> developed by CERN. Decreasing prices for commodity disks, however, are driving an emerging area where Tape

---

<sup>47</sup>Dorigo, Alvise, et al. “XROOTD-A Highly scalable architecture for data access.” WSEAS Transactions on Computers 1.4.3 (2005).

<sup>48</sup>For example, redirection is a standard feature in the http protocol used for the World Wide Web.

<sup>49</sup>See for example Bloom, Kenneth, and Cms Collaboration. “Cms use of a data federation.” Journal of Physics: Conference Series. Vol. 513. No. 4. IOP Publishing, 2014.

<sup>50</sup>Conventional backup to tape of files in home and project directories is usually performed as part of the general operation of clusters using standard commercial or open-source packages. This is an important function but, being completely conventional, is not considered further in this context.

<sup>51</sup>Teaff, Danny, Dick Watson, and Bob Coyne. “The architecture of the high performance storage system (hpss).” (1998).

<sup>52</sup>IBM Corporation, Tivoli Storage Manager; see <http://www.ibm.com/software/tivoli/products/storage-mgr/>

<sup>53</sup>Bakken, Jon, et al. “Enstore Technical Design Document.” Fermilab-JP0026 (1999).

<sup>54</sup>Presti, Giuseppe Lo, et al. “CASTOR: A Distributed Storage Resource Facility for High Performance Data Processing at CERN.” MSST. Vol. 7. 2007.

servers are implemented on inexpensive commodity computers with large disks. This is in turn stimulating multi-layer implementations, for example the D2D2T (disk-to-disk-to-tape) scheme, where Disk servers store and retrieve data on disk-based pseudo-tape servers, which in turn store and retrieve data on “real” magnetic cartridge Tape servers. Both disk-based tape servers and D2D2T schemes are important because the large growth in data volumes in particle physics is the result of a very modest growth in file size multiplied by a large growth in the number of files, resulting in a situation which can be very inefficient for traditional tape systems (the so-called small file problem in tape cartridge storage).

A “Hierarchical Storage Manager” (HSM) is implemented between the Data servers and Tape servers using specific Control and Information servers. HSM systems are not a common feature of conventional computing environments. HSM systems that fulfill the needs of particle physics data processing are even less common. Furthermore, the interaction between Data servers and Tape servers can be rather complex and no clean, practical interfaces are available. Therefore, each data center implements its own customized solution using a combination of community-written and commercial software. Examples of Data server/HSM implementations are dCache interfaced to Tivoli or Enstore; GPFS interfaced to IBM HPSS; and CERN’s Castor-2 which has its own data server and HSM implementations. Depending on the solution, the HSM system may be more or less intertwined with the Data server software, although the tendency is to separate as much as possible the HSM and Data server implementations.

Experiment specific Information servers are deployed in order to manage and make accessible a view of all their files containing custom details about their characteristics, for example trigger conditions, detector configuration or beam type. This is often called a “File Catalog” and may in itself be a quite complex database.

## **14.4 The Main Workflows of Data Processing in an Experiment**

Most experiments implement two “organized” data processing workflows: one for data coming from the actual detector (awkwardly referred to as “real” data) and another for simulations resulting from the techniques described in the contribution “Detector Simulation” (awkwardly referred to as “Monte Carlo” data). An additional very important “organized” workflow derives calibration constants from detector data taken either in normal mode or in special conditions for calibration (cosmic rays, single beam, sending pulses to the front-end electronics, etc.). Altogether, these three workflows are referred to as “production” and constitute the process by which “raw” data is transformed into “analysis datasets”. Workflows related to the use of these analysis datasets by individual physicists are described in a later section.

The production workflows involve a number of tasks,<sup>55</sup> such as event reconstruction, simulated event production, data reduction, event filtering, and calibration data processing.

### 14.4.1 *Event Reconstruction*

Data arriving from an experiment's data acquisition system is called "raw" data. It is organized in "event records", each record corresponding to a particle interaction "event" which is relevant in the context of the experiment. An event record may or may not correspond to a single collision between particles, though particle physicists often refer to data arising from a single collision as "an event". In a collider, an event record may correspond to data collected from a single beam crossing, which depending on the beam type, luminosity and angular coverage may contain data arising from several particle collisions.<sup>56</sup> In a fixed target experiment an event record may correspond to data collected over a full beam spill. In an astroparticle experiment an event record may correspond to data collected starting from a trigger signal for a certain length of time. Furthermore, due to performance requirements of data acquisition systems, raw data is often compressed, encoded or packed into rather complex data structures, which are often not very convenient for direct use in reconstruction programs. Therefore, a raw event record is usually read from disk and immediately "unpacked" into simpler but larger data structures where the digitized signals are represented by standard integer or floating point numbers.

The output of the most complete simulations (often called "full" simulations, see below) can be stored in "packed raw" format identical to the detector data, or in "unpacked raw" format in which case the unpacking step is omitted in reading simulated event records.<sup>57</sup> Thereafter, simulated event records are treated identically to detector "real" data.

A number of event records are written by the data acquisition system to a raw data file, which is the main data item delivered from the "online" to the "offline" environment. The organization of event records into raw data files is very often unrelated to the experimental conditions and simply reflects technical constraints such as the size of disk buffers in the online and offline clusters.

---

<sup>55</sup>The names of these tasks are not standardized. Therefore readers will have to identify the corresponding task names for their respective project.

<sup>56</sup>In electron-positron collisions there is a single collision most of the time, but there may be occasionally an annihilation event together with a two-photon collision, for example. In proton-proton collisions there are on average 20 collisions per beam crossing at the nominal luminosity of the original LHC configuration.

<sup>57</sup>Even though it would be more economical to store simulated data in "packed raw" format, this is often not possible. For example, the simulations may be needed before the detector data acquisition system is fully designed and implemented, and therefore the exact format for packed raw data will not be known.

The output of the reconstruction arises from the application of the techniques described in the contribution “Pattern Recognition and Reconstruction” and consists of data structures containing tracks, clusters, etc. These data structures are written as event records, which may or may not have a one-to-one correspondence to the event records read on input. For example, the reconstruction program may disentangle data from different collisions contained in a single raw data event record and write the corresponding reconstruction output as separate event records for each collision.

The reconstruction program needs detector configuration and calibration data that often varies with time, though reasonably slowly compared to the event rate.<sup>58</sup> Therefore, the reconstruction program accesses a “conditions database” in order to fetch the configuration and calibration data matching the conditions under which a particular event was acquired. Alternatively, if stable and reliable calibration data are available at the time of acquisition, it may be stored within the raw data files, interspersed with the event records.

Event records are, in principle, all independent from one another and therefore suitable for loosely-coupled parallel processing. Early implementations<sup>59</sup> would process one raw data file at a time, in essence treating the raw data file as a buffer and implementing an “event server” to distribute individual events via the network to CPU servers. This methodology has been largely abandoned for event reconstruction, although it remains a good choice for parallel processing of analysis data sets. Modern event reconstruction clusters or “farms” are configured so that each CPU server executes jobs that process one or more complete data files at a time. The CPU servers read raw and calibration data and write reconstructed data in an autonomous fashion to the Disk and Database servers. Hundreds or even thousands of jobs can run simultaneously in a single cluster.

Care must be taken to handle aspects that may introduce coupling and break the highly parallel processing scheme. Consider, for example, one of the most common pitfalls that breaks parallelism: simultaneous access to the same file by many jobs. For example, the executable image of the reconstruction program is the same for all jobs and may be quite large. Therefore, each job that starts must read the image file from a Network File server. In an environment where thousands of jobs are executed, the number of jobs simultaneously reading the image file may be

---

<sup>58</sup>There are exceptions to this rule, especially for neutrino beam or non-accelerator events, where calibration data may in fact be updated faster than the event rate.

<sup>59</sup>Event processing in parallel dates back to the 1970s when physicists used their experience in building digital electronic boards for detectors to design inexpensive processor boards, called emulators, which executed the instruction set of the IBM System 370 mainframe computer and interfaced to its input-output channel by emulating a tape drive. Starting in the mid-1980s, parallel event processing was mostly done using specially configured clusters (called “farms”) of scientific workstations with Reduced Instruction Set CPUs (RISC) running the VMS operating system from Digital Equipment Corp. and a number of variants of Unix, interconnected by proprietary networks or by daisy-chained Ethernet networks. Starting in the mid-1990s to date, essentially all reconstruction is run on clusters of widely marketed, industry standard “commodity” computers using the x86 instruction set supported by Intel and Advance Micro Devices CPU chips and interconnected by switched Ethernet networks.

high enough to overload the Network File server, therefore creating a performance bottleneck.<sup>60</sup> Caching of the image file on a local disk of each CPU server can be used to restore parallelism more economically compared to deploying a more powerful Network File server. A similar situation arises from the access by all jobs to a unique conditions database. Parallelism is harder to restore in this case and requires tuning of the database server and the caching algorithm, or deployment of very powerful (and expensive) replicated database servers.

Raw data files must often be retrieved from tape, and parallelism may be broken from the fact that a single tape cartridge is used to store multiple raw data files. Many HSM systems are unable to handle multiple, asynchronous, closely spaced requests for files on the same tape. They queue the requests and mount and dismount the tape cartridge in order to read each of the files, a slow procedure as it involves mechanical actions. As individual raw file sizes have historically increased more slowly than total raw data volume, the number of files per tape cartridge has increased and parallelism breakage by tape access has become quite common. In order to circumvent these problems, production managers run special jobs to “pre-fetch” in an efficient way lists of specific datasets from tape to Disk servers. Efficient pre-fetching requires detailed knowledge of the HSM software and the contents of the tape cartridges in order to minimize the number of tape mounts.

Most modern experiments strive to run a production workflow for reconstructing the raw data that keeps up with the data acquisition, called “quasi-online” or “prompt” reconstruction. The quality of the calibration constants which are available for this workflow is often limited and therefore the reconstructed output may be of limited value for final physics results; the output is very valuable, however, to monitor detector performance and to derive more definitive calibration constants.

Frequently, reconstruction of a particular raw data file must be repeated a few times, in order to incorporate improvements or correct errors in calibration constants and reconstruction algorithms. This is called “reprocessing” and applies to both detector and simulated data.

### 14.4.2 *Event Simulation*

Production of simulated events usually starts with the generation of physics final-states using an “event generator”. The CPU time needed for this step is usually

---

<sup>60</sup>To understand how parallelism is broken, consider that the Network File server handles read requests in parallel up to a certain limit and will serialize into a queue all requests exceeding this limit, breaking parallelism. To understand how coupling is introduced, consider that the event records are coupled through the fact that they are all being reconstructed using the same reconstruction program image file. This can be exacerbated by additional coupling at the hardware level, for example when many cores compete for the same Ethernet network interface in modern multi-core computers.

small, as is the size of the output file (which is often called a “four-vector file”).<sup>61</sup> A single (non-parallel) job can usually produce enough four-vectors for a whole simulation campaign. The situation is completely different in the next step, where the four-vectors go through a full detector simulation which is very CPU time consuming. Many jobs must be run in parallel reading four-vectors and writing simulated event records<sup>62</sup> to Disk servers. The data rate per job is not very high, but the integrated output rate can be substantial since there are many CPU servers running simultaneously, an issue that should be taken into account. Simulation is often run on many separate clusters and the oversight of simulation productions can consume substantial amounts of time from a large number of physicists. Deployment of Grid interfaces on clusters (see below) has simplified these tasks somewhat. A recent tendency is to use “production management systems” based on databases to keep track of simulation jobs. The most tedious part, however, is the detection and correction of errors coming from the execution of the simulation program itself, which currently is not automated in part due to lack of a systematic approach to error signaling.

### ***14.4.3 Data Reduction and Event Filtering***

Appropriate data representations for physics analysis are quite different from the raw data format. They can also be quite different from the reconstruction output format for several reasons, some related to enabling limited reprocessing capabilities from the reconstruction output itself, while others are related to signal-to-noise issues. Reconstruction output may have a format optimized for limited reprocessing without access to the raw data, for example recalculating vertices from tracks or track and calorimeter cluster matches. This format may not be optimal for analysis. A simple example can illustrate this: a typical analysis condition is to access all tracks that have a given minimum number of points (or “hits”); hence, the optimum data structure is a list of tracks, each track having the property “number of points”. The reconstruction output, however, will most likely be stored as a list of points, each point having a property of belonging to a track, a data structure optimized for reprocessing the assignment of points to tracks, but completely “backwards” from the one needed for optimal analysis. Reconstruction output will often have additional information stored, related to enabling reprocessing, which makes it too bulky for practical use in analysis. It is often organized keeping a

---

<sup>61</sup> A single generation step is described for simplicity. Note that in some cases the output of an event generator may need to be post-processed by another physics simulation program, for example to turn quarks and gluons into jets.

<sup>62</sup> A single step is described for simplicity. Note that in some cases the output of the detector simulation may have to be post-processed to generate simulated event records, for example by adding simulated pile-up events for proton–proton collisions in the LHC.

one-to-one relationship to raw data files, which in turn may have little to do with physics analysis. Furthermore, it usually has output for all raw data records, which depending on detector triggering and physics signal-to-noise conditions may be very large compared to the optimal for physics analysis.

These considerations lead to the definition of “reduced” or “analysis” datasets along two complementary lines. Data reduction takes the reconstruction output and generates data structures optimized for analysis by removing information related to reprocessing and “turning around” data structures related to tracks, vertices, calorimeter clusters, etc. to be optimal for analysis. Particle identification by correlation between different detector elements is often done (or re-done) during data reduction and, when combined with Event Filtering, may be optimized for particular physics processes (heavy quark identification, rejection of electrons, etc.). The historical name for the output of data reduction is “Data Summary Tape” or “DST” and successively smaller outputs have been called “mini”, “micro” or even “nano” DSTs.<sup>63</sup> A nano-DST may simply consist of a list of “reconstructed particles”, each stored with a few properties such as particle ID probabilities and momentum components. A nano-DST is in some sense the experimental counterpart to the event generator output.

DST files often do not mirror the raw and reconstructed file organization scheme. At a minimum, DST files contain the DST records that correspond to many raw files, in order to ease the management of large numbers of datasets and avoid handling very small files. In accelerator experiments, DSTs may be organized by running conditions. In fixed-target experiments where on-beam and off-beam data are alternatively taken, separate DSTs for each may be produced. In astroparticle experiments, which often behave like telescopes, DSTs will often be organized by observed source, merging data from many observation days or nights. These reorganizations of data are special cases of event filtering.

Event filtering<sup>64</sup> may be performed before, after or during production of any DST-like output. The purpose is to group together in particular files the DST output most relevant for a particular analysis, thereby avoiding repetitive reading of the rest of the data by analysis programs. Event filtering can be done according to many criteria, but is most often performed as a pre-selection for particular physics analyses, especially in situations of unfavorable signal-to-noise ratio such as the LHC experiments. Output datasets are often referred to by the name of the physics process pre-selected and are written in an identical format as the original DST-like datasets.

---

<sup>63</sup>The names of Data Reduction output datasets are not standardized. Therefore readers will have to identify the corresponding dataset names for their respective project.

<sup>64</sup>Event Filtering in this context should not be confused with the actions taken by high-level trigger “Event Filters”.

#### 14.4.4 Processing of Calibration Data and Calculation of Calibration Constants

Reconstruction of data from detectors requires a large amount of calibration constants, many of them time-varying. Raw calibration data, such as temperatures, pressures and voltages are recorded periodically during data taking. Calibration data related to response and alignment of different detector elements can be obtained by generating artificial events with specialized calibration devices, for example pulsed lasers. Alternatively, they can be calculated from the raw physics data itself, using better understood reactions or additional events occurring in parallel to the physics, such as the passage of cosmic rays.

Data input to calibration procedures often comes from diverse sources, and handling it poses challenges that are often underestimated. For example, laser pulses for calibration may be recorded in parallel to the physics data and stored by the online system interspersed in the raw data files, thereby requiring the offline to access all raw data files in order to process the laser pulses, a rather tedious procedure. Conversely, sometimes calibration data are stored in a *potpourri* of small files and databases of diverse formats, introducing unnecessary complexity in the handling of these data.

There is no standardized framework to handle the output of all calibration procedures, although the recent trend is to centralize it into a single calibration (or “conditions”) database that is accessed through heavily cached methods by instances of the reconstruction, DST generation and analysis programs. Care must be taken to serve this database with adequate resources, and to ensure efficient and cached access to it in order not to break parallelism as described earlier.

The flow of raw and derived calibration data and the related data processing tasks should be studied with equal care as that of the physics data, and solutions should be implemented to ensure efficient and agile data access for calibration purposes.

### 14.5 Interactive Analysis Using Clusters

The last stages of physics analysis are performed by individual users using interactive tools such as ROOT,<sup>65</sup> Jupyter Notebooks<sup>66</sup> or SWAN.<sup>67</sup> They apply

---

<sup>65</sup>See <http://root.cern.ch>

<sup>66</sup>See <https://jupyter.org/about>. Although Jupyter Notebooks are language agnostic, they are closely associated with the Python language. They are increasingly popular in particle physics and many other scientific disciplines, given the rise of the use of Python by them. Jupyter Notebooks must be combined with suitable data-access modules. For example, the SWAN project at CERN combines Jupyter Notebook with ROOT input–output modules to create an interactive analysis platform.

<sup>67</sup>Danilo Piparo, Enric Tejedor, Pere Mato, Luca Mascetti, Jakub Moscicki, Massimo Lamanna, SWAN: A service for interactive analysis in the cloud, Future Generation Computer Systems, Volume 78, Part 3, 2018, Pages 1071–1078, ISSN 0167-739X, <https://doi.org/10.1016/j.future.2016.11.035>



selection criteria to events in DST-like files or in personal or group files derived from them and written directly in the format of the analysis tools. The basic record in these files is often a collection of ordered lists of properties, or “n-tuples”.

Distributed computing is used for preparing the input files for the interactive tools, using the same methodologies described earlier for running reconstruction and data reduction jobs. In addition, since the application of selection criteria and the generation of histograms from n-tuple files can often be mapped to highly parallel tasks (if each n-tuple record is independent of any other), efforts have been made to provide parallel “back-ends” to interactive analysis tools, for example the PROOF back-end to ROOT or the use of *Hadoop* in conjunction with Jupyter Notebooks. The idea is simple in principle: the user issues a command that specifies some actions to be performed on all records; the back-end automatically generates and executes in parallel a number of tasks reading the data and writing partial result files; when all parallel tasks have finished, the partial results are merged to a global result which is presented to the user via the interactive front-end. Speed is gained, in principle, by using many CPUs and, more importantly, many input data channels in parallel. In addition, since the interactive user pauses to examine the results, resources can be used more efficiently by serving many users from a single parallel back-end. Practical implementation with good performance, however, can be extremely complex, as they involve running a cluster at very large peak input–output rates and clever scheduling of bursts of parallel tasks.

## 14.6 Multiple Sites and Grid Computing

Grids were defined in the late 1990s as a new manner of integrating services from clusters across multiple sites.<sup>68</sup> Practical development and deployment of Grids was originally led by the particle physics community, especially the LHC collaborations. Grids are convenient for two main purposes: data distribution to many end-users who are geographically dispersed, and data processing load sharing across many computing centers (“resource” centers).

The main idea of a Grid is to abstract the most commonly used interfaces for authentication, access control, job submission and data access into a unique “meta-interface”, which is the only interface exposed to end-users. Therefore, multiple sites are integrated for the end-user under the illusion of a single “meta-cluster”. Specialized Grid Information and Control servers are deployed in order to translate from the abstract meta-interface into the specific interface used by each cluster. Grid interfaces and protocols have not been standardized. After an intense period of development during the 2000s, a handful of partially compatible Grid schemes are now deployed. Many claims of success in providing resources on the Grid

---

<sup>68</sup>“The grid: blueprint for a new computing infrastructure”, Ian Foster and Carl Kesselman, editors. San Francisco: Morgan Kaufmann Publishers, c1999. ISBN 1558604758.

come from providing simple CPU-intensive services with little input or output. Providing reliable data-centric services has proven to be much more difficult, both from the point of view of the operation of resource centers themselves and of the Information and Control servers which provide the Grid interface for each center. Nevertheless, the CERN Computer Center, the WLCG Tier-1 centers and a subset of the WLCG Tier-2 centers do provide the needed reliable data-centric services needed for nucleating the rest of resources (see below).

The European Grid Infrastructure (EGI) program,<sup>69</sup> financed in part by the EU, has deployed the largest production Grid Infrastructure to date in collaboration with the Worldwide LHC Computing Grid (see next section). Other Grid Infrastructures provide support for particle and astroparticle physics, for example the Open Science Grid (OSG) in the USA. The EGI Grid will be described as an example of a working Grid used for daily work by particle physicists. EGI defines a Grid Infrastructure as a set of services deployed over Internet that allows a large number of Resource Centers, each with a different security and management domain, to be used in a coherent fashion by a large number of users from different institutions grouped in virtual organizations. Usually, a virtual organization corresponds to a project.

The basic services offered by the EGI Grid Infrastructure are based on the UMD middleware distribution,<sup>70</sup> which is deployed in a managed way, including monitoring and fault detection, with the support of a distributed operations organization provided by resource centers, many of which are associated to LHC computing or national Research and Education networks. Resource centers deploy Remote Data Servers using a variety of Disk Servers (possibly with magnetic tape back ends) which expose a uniform interface to the Grid via UMD, whose instances are called “Storage Elements”. Authorization and access control is accomplished using the certificate method described earlier with tokens generated by the Virtual Organization Membership Service (VOMS) upon presentation of X.509 certificates issued by national research agencies and declaration of the virtual organization (e.g. project) under which the user wishes to be authorized. The actual data flow is predominantly through the *ftp* protocol, wrapped in tools such as *gridftp* (which incorporates certificate authentication and multiple parallel transmission streams).<sup>71</sup>

Each resource center implements a “Computing Element” (CE) service, a set of Information and Control Servers which map a specific batch system to an abstract

---

<sup>69</sup>See <http://www.egi.eu>. EGI operates the international coordination structure for Grid resource centers across Europe. EGI also coordinates the maintenance and security support for the Grid software originally developed in the 2000s by the EU Enabling Grids for E-science (EGEE) series of projects and the EU DataGrid project.

<sup>70</sup>The Unified Middleware Distribution (UMD) maintained by EGI is based on software developed in the EU European Middleware Infrastructure (EMI) project. The EMI software is the result of integration and compatibility work done on the gLite middleware, developed within the EGEE Project, the ARC middleware, developed by NorduGrid, and various other packages. See <https://wiki.egi.eu/wiki/Middleware>

<sup>71</sup>A more sophisticated Grid data access tool known as *srm* was used in the past but is largely being abandoned due to maintenance and management issues.

Grid batch system. Directives may be given to the CE specifying job requirements such as minimum amount of memory, length of execution, etc. A challenge in configuring very large Grids is to reach agreement on the normalizations used to measure cluster resources, for example the speed of processors. Users submit their jobs through an additional server, called a Resource Broker, which accesses information in tables maintained by all available CEs and submits the jobs as appropriate to balance the load amongst the multiple sites conforming the Grid. Large projects, such as the LHC experiments, have developed job management services which in practice replace the Resource Brokers, interacting directly with CEs.

The Grid data access services described above offer functionality similar to the low levels of operating systems, but with the great advantage of generating coherent behavior across many Resource Centers. In practice, these services are often too low-level for practical use by applications. Hence, additional software, services and APIs are provided by middleware distributions at a higher level. One example is the File Transfer Service (FTS),<sup>72</sup> designed to support applications which require automated management, monitoring and error correction of replication of large numbers of files and large volumes of data, such as the LHC experiments. It uses a database back end to store replication requests and instances of finite state machines to reliably accomplish the corresponding data replications, including error correction and reporting. In addition, FTS implements the concept of a channel (a virtual path between two Resource Centers) which may be managed by operators, adjusting parameters such as the number of simultaneous files being transferred, the maximum bandwidth used, etc.

Large scientific projects, such as particle physics experiments, can have specific needs that go beyond the scope or functionality of basic middleware (such as UMD) and basic Grid services (such as those offered by EGI or OSG). In these cases, the projects themselves provide their own enhancements.<sup>73</sup>

## 14.7 The Worldwide LHC Computing Grid (WLCG)

The largest data processing capacity deployed so far using clusters linked via Grid infrastructures is the Worldwide LHC Computing Grid (WLCG or LCG).<sup>74</sup> WLCG links some 170 clusters located in 42 different countries into a single

---

<sup>72</sup>Andrey Kiryanov, Alejandro Alvarez Ayllon and Oliver Keeble, “FTS3/WebFTS – A Powerful File Transfer Service for Scientific Communities”, *Procedia Computer Science*, Volume 66, 2015, Pages 670–678, ISSN 1877-0509, <https://doi.org/10.1016/j.procs.2015.11.076>

<sup>73</sup>Readers should inquire within their own projects about these enhancements. They usually fall into two categories: Job management enhancements, for example the DIRAC package used by the LHCb experiment, and File or Dataset management enhancements, such as the AAA package used in CMS and the Rucio package used by ATLAS.

<sup>74</sup>See <http://www.cern.ch/lcg>

**Table 14.1** Capacity  
deployed in the tiers of the  
LHC computing grid in 2018

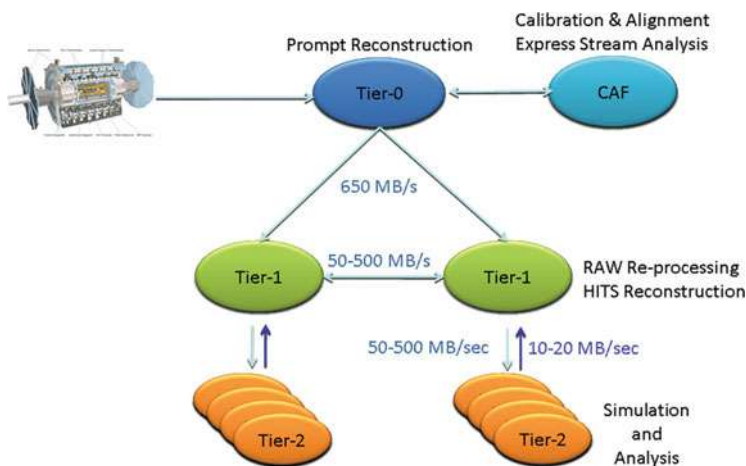
		2018
CPU (KHS06 <sup>a</sup> )	Tier-0 <sup>b</sup>	1272
	All Tier-1	2061
	All Tier-2	2562
Disk (TeraBytes)	Tier-0	90700
	All Tier-1	192214
	All Tier-2	185455
Tape (TeraBytes)	Tier-0	284700
	All Tier-1	461050

<sup>a</sup>Thousands of HEPSPEC06 units. HEPSPEC06 is a measure of CPU speed maintained by the particle physics community. See <http://w3.hepfix.org/benchmarking.html>. As a very rough guideline, a current 2 GHz CPU delivers about 15 HS06

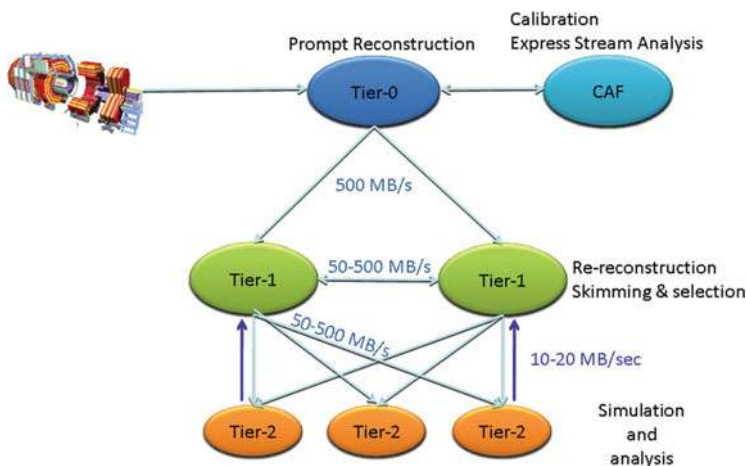
<sup>b</sup>In addition to the Tier-0, CERN also deploys an Analysis Facility (CAF) which is not included in this table

coherent distributed system which supports all data processing operations of the LHC experiments. WLCG defines a global service supported in part by Grid Infrastructure services in various countries and regions, for example EGI in Europe and OSG in the USA. The large amount of data handled by WLCG (over 100 Petabytes per year exchanged between sites) requires the sites to be organized in three “Tiers”. Centers belonging to a given Tier are dedicated to perform certain specific tasks within the computing model of a given LHC experiment. There is one Tier-0 center at CERN, which handles Data Recording, Prompt Reconstruction and Remote Data Distribution (mostly to the Tier-1 centers). There are 13 Tier-1 centers, which perform Data Recording (between them they hold a second copy of all raw data stored at the Tier-0), Reprocessing, Data Reduction and Event Filtering. Hundreds of Tier-2 centers receive from the Tier-1 centers DST-type files (reduced and filtered) which physicists then use for analysis. Tier-2 centers also run event simulation, transmitting the outputs to Tier-1 centers for data recording and reconstruction. FTS, coupled to some experiment specific tools, is used to automate all data transmission.<sup>75</sup> The capacities deployed are quite large: about 400 thousand computing cores, 250 Petabytes of disk and nearly an Exabyte of tape. Table 14.1 shows the capacities deployed in 2018. These capacities have grown at an impressive rate of about 37% year-to-year since 2008. An even larger growth is expected for Run-3 of the LHC starting in 2021.

<sup>75</sup>The Wide Area Network requirements between the Tier-0 center and the 13 Tier-1 centers are so demanding that a special network based on point-to-point 10 Gbps links has been deployed (the LHC Optical Private Network, or LHCOPN). As of this writing, the LHCOPN links are being upgraded to 100 Gbps.



**Fig. 14.2** Schematic view of the original ATLAS computing model



**Fig. 14.3** Schematic view of the original CMS computing model

As mentioned earlier, the detailed breakdown of tasks performed at each Tier depends on the computing model of each experiment. On the other hand, the Tier-0 and most Tier-1 sites serve multiple experiments, whereas most Tier-2 sites only serve a single experiment. This leads to quite complex optimization issues, especially in regards to data transmission, which depending on the experiment may involve few or many site pairings. For example, Figs. 14.2 and 14.3 show schematically the original computing models for the ATLAS and CMS experiments, and the different connectivity patterns between Tiers can clearly be seen. Originally, ATLAS closely linked a given Tier-1 to a group of Tier-2 sites, and hence the data flows in a more “hierarchical” fashion. The original CMS model had a focus on

coherent operation of all Tier-1 sites and allows data to flow between any pairing of Tier-1 and Tier-2 centers. The actual situation is even more complex, as there can be large size differences between Tier-1 or Tier-2 centers at different sites.

The computing models of the LHC experiments have undergone substantial evolution in the last 10 years, influenced by the operational experience that has been gained. The most important lesson is that reliably managing storage at sites is a difficult task, requiring specialized skills not always available at resource centers. In addition, the management of data and replica placement by the central production teams of each experiment requires substantial human resources, and the effort needed scales with the number of sites as opposed to the overall amount of storage. Hence, a model where a limited number of large sites provide highly reliable 10–100 Petabyte storage platforms, with the rest of the sites providing CPU power with input–output via Remote Data Servers, is globally more manageable and economical.

The original, static Tier hierarchy was motivated in part by the high costs and technical difficulties of deploying Gigabit per second (Gbps) WAN networking. However, networking in general and WAN in particular have had the fastest evolution of all computer industry technologies. In addition, Research and Education Networks provide a wider variety of connection options, including 10 and 100 Gbps links for specially demanding sites. This means that currently some Tier-2 sites have equal networking capabilities to a Tier-1 site, whereas they may or may not achieve the reliability levels of a Tier-1 site. It is convenient, then, to use these sites in a more flexible manner through a more dynamic hierarchical model.

Hence, the WLCG model has evolved towards differentiating two types of sites: Nucleation sites, which have high reliability and provide Petabyte-level storage, and satellite sites, which are mostly expected to deliver CPU capacity. WAN based input–output is accomplished through the custom *xrootd* or the standard *http* protocols, and the storage is remotely managed by the central experiment project teams via standard *WebDAV* protocols interacting with the Remote Data Servers at nucleation sites. Sites are continuously monitored and their status as nucleation points is defined dynamically according to their current reliability. In this model, most of the Tier-1 sites and larger Tier-2 sites act as nucleation sites. A site with reliability or availability problems will be temporarily removed as a nucleation site and restored when monitoring reveals the issues have been resolved. Tier-1 sites continue to have the differentiating feature of operating the robotic magnetic tape libraries used for bulk-archiving of the data.

## 14.8 Current State-of-the-Art Technologies

The computing industry continues to evolve at a rapid pace. The decade of the 2010s has seen the consolidation of a number of technologies and the emergence of an important change in business models for the provision of computing services. These are being integrated into the computing environments for particle physics.

### 14.8.1 *Multi-core CPUs and Virtualization*

Multi-core CPUs have been discussed in Sect. 14.1. Ideally, all applications would by now be multi-threaded and capable of efficiently utilizing the full capacity of the cores in a CPU. Unfortunately, thread-safe programming is a specialized skill and the vast majority of programs used in the world are not thread-safe. Hence, a different approach is needed to fully use multi-core CPUs.

In addition, applications running in distributed computing environments often have intricate dependencies on the underlying libraries and on the operating system. Continuously porting applications to newer versions of the underlying environment and verifying their proper operation requires much larger human resources than can be afforded by scientific communities (and even by many business communities). Hence, deployment of applications in these environments enters into conflict with system administration good practices, which require deploying new operating system versions soon after they become available.

These two factors have led to the re-emergence of two techniques from the past which enable operating system virtualization:<sup>76</sup> hypervisors and virtual machines (first deployed in the 1970s by IBM in its VM operating system<sup>77</sup>), and containers (also known as jails, partitions or zones, first deployed commercially in the early 2000s by Sun Microsystems<sup>78</sup>).

Virtual machines are implemented by deploying over the hardware a special type of operating system called a hypervisor. The hypervisor is designed to efficiently run multiple copies of a special type of application called a virtual image. In addition, newer versions of hardware implement special features to accelerate virtual image execution. The virtual image contains one or more applications together with the full operating system and library environment required by them. In this way, every time the hypervisor starts running an image it effectively bootstraps a new virtual computer, which is referred to as a virtual machine. Each image can have a different operating system and different libraries. If properly deployed, virtual machines allow to run older, less secure operating systems without compromising security, as the hypervisor can limit access to hardware and file system resources.

Virtual machines allow the implementation of many interesting features. One of the best known is live-migration,<sup>79</sup> which combines virtual machines with checkpoint/restart and high performance multi-connection file servers to be able

---

<sup>76</sup>For a short, rather technical overview, see “Virtualization and Containerization of Application Infrastructure: A Comparison”, Thijs Scheepers, 21st Twente Student Conference on IT June 23rd, 2014, Enschede, The Netherlands. Available from the author at <https://thijs.ai/papers/scheepers-virtualization-containerization.pdf>

<sup>77</sup>“The Origin of the VM/370 Time-Sharing System”, R. J. Creasy, IBM J. Res. Develop., Vol. 25, No. 5 (September 1981).

<sup>78</sup>Price, Daniel, and Andrew Tucker. “Solaris Zones: Operating System Support for Consolidating Commercial Workloads.” LISA. Vol. 4. 2004.

<sup>79</sup>Sapuntzakis, Constantine P., et al. “Optimizing the migration of virtual computers.” ACM SIGOPS Operating Systems Review 36.SI (2002): 377–390.

to move virtual machines between different hardware units. This can be useful for providing uninterruptable services or for saving energy in environments with large load variations.

Traditional offline batch computing in particle physics, however, predominantly uses virtualization for simply packaging together an application and its environment. This reveals a weakness of traditional virtualization, as the many copies of operating systems in the virtual machines require large amounts of memory (RAM) to execute efficiently, increasing the costs of the computing clusters. This has recently been circumvented by a new virtualization package, *Singularity*,<sup>80</sup> which has been rapidly adopted by the scientific community, and large portions of WLCG now support the execution of *Singularity* images. An additional advantage of *Singularity* is that it can be easily supported on supercomputers.

Containers, in their present form, represent a slightly different approach to virtualization. The current de-facto standard is *Docker*,<sup>81</sup> which is widely used in commercial applications. Containers are closely associated to Linux operating systems and to Cloud Computing environments (see below). Containers implement virtual machines which are less hermetic than traditional ones and which have certain limitations. On the other hand, containers can be rapidly created and destroyed, and packages exist, such as *Kubernetes*,<sup>82</sup> for orchestrating a set of containers which together can implement sophisticated applications using micro-service architectures.

In a somewhat unexpected way, *Singularity* and container environments have recently emerged as the leading approach for preserving data processing environments in order to ensure reproducibility of scientific results.

## 14.8.2 Cloud Computing and the Use of Commercial Data Processing Services

Cloud Computing<sup>83</sup> is the currently used term<sup>84</sup> to describe the deployment of computing resources and higher level computing services in a shareable, user configurable way. It is the new paradigm that replaced client-server computing in the

---

<sup>80</sup>Kurtzer GM, Sochat V, and Bauer MW (2017) Singularity: Scientific containers for mobility of compute. PLoS ONE 12(5): e0177459. <https://doi.org/10.1371/journal.pone.0177459>

<sup>81</sup>Boettiger, Carl. “An introduction to Docker for reproducible research.” ACM SIGOPS Operating Systems Review 49.1 (2015): 71–79.

<sup>82</sup>Bernstein, David. “Containers and cloud: From lxc to docker to kubernetes.” IEEE Cloud Computing 3 (2014): 81–84.

<sup>83</sup>Mell, Peter, Grance, Tim, “The NIST Definition of Cloud Computing”, Special Publication 800-145, National Institute of Standards and Technology, U.S. Department of Commerce (September 2011) <https://doi.org/10.6028/NIST.SP.800-145>

<sup>84</sup>Regalado, Antonio. “Who coined ‘cloud computing’.” Technology Review 31 (2011). <https://www.technologyreview.com/s/425970/who-coined-cloud-computing/> Retrieved 3 February 2019.



2000s, starting a new cycle in the development of outsourced computing services. The first large scale commercialization was done in 2006, when Amazon launched the Elastic Compute Cloud (EC2) service. The first large-scale production cloud services in the academic domain were NASA's OpenNebula<sup>85</sup> and the services deployed in the RESERVOIR EU-funded project.<sup>86</sup>

Cloud computing is still under heavy evolution, and a full description of the technology is beyond the scope of this work. It deploys interfaces and toolkits that allow users to configure virtual computers, storage servers, clusters and even networks. The ultimate goal is to use software to define the characteristics of a distributed computing service and then automatically map it to the needed hardware. This process is called provisioning.

Cloud computing can be used to implement private research data centers. CERN has migrated essentially all of its data center platforms to be managed by the *OpenStack*<sup>87</sup> cloud management system. However, the benefits of introducing cloud-style management in smaller data centers are not evident at present, as standard clusters are probably sufficient to fulfill the needs. In addition, there is at present a lack of personnel trained in cloud computing deployment and operation.

Cloud computing technologies started in the academic world, but are now driven by the cloud computing industry, which has grown in a decade to have sales of over 200 G\$/year. Cloud infrastructure companies deploy data centers which are much larger than academic computing centers, with the exception of a few academic supercomputing centers. Computer room floor areas above one hundred thousand square meters, power feeds above 50 MW<sup>88</sup> and capabilities for hosting tens of thousand of servers and close to a million cores are typical.

These huge deployments bring with them large economies of scale, making them attractive for scientific computing. In some cases, U.S. funding agencies have started to give grants for the purchase of commercial cloud computing services, replacing funding for purchasing research computing clusters. Gateways have been developed to include commercial clouds as part of WLCG, especially for CPU intensive tasks such as simulations. In the European Union, the HNSciCloud<sup>89</sup> pre-commercial procurement project in 2016–2018 was aimed to stimulate industry developments to enable hybrid clouds between research and commercial data centers capable of

---

<sup>85</sup>Nebula Cloud Computing Platform (20 November 2012) <https://www.nasa.gov/open/nebula.html>

<sup>86</sup>Rochwerger, Benny, et al. "The reservoir model and architecture for open federated cloud computing." IBM Journal of Research and Development 53.4 (2009): 4–1.

<sup>87</sup><https://www.openstack.org/software/>

<sup>88</sup>For comparison, CERN's Meyrin data center power was 3.5 MW in 2018. "Data Centre: Key Info & Numbers". [http://information-technology.web.cern.ch/sites/information-technology.web.cern.ch/files/CERNDDataCentre\\_KeyInformation\\_November2018V1.docx.pdf](http://information-technology.web.cern.ch/sites/information-technology.web.cern.ch/files/CERNDDataCentre_KeyInformation_November2018V1.docx.pdf). Retrieved 3 February 2019.

<sup>89</sup>Helix Nebula – The Science Cloud with Grant Agreement 687614 is a Pre-Commercial Procurement Action funded by H2020 Framework Programme. More information at <https://hnscicloud.eu>

executing data-intensive tasks. A similar project named ARCHIVER<sup>90</sup> will explore in 2019–2021 the commercial provision of cloud mass storage services for scientific data archiving and preservation.

Cloud computing deployment, especially on commercial services, is quite complex at the infrastructure and platform level. Fortunately, past investments in Grid computing can be re-used in order to hide these complexities from the vast majority of users. Job and data management packaged used for Grid computing have already been deployed on private and commercial clouds, as well as portals for analysis using higher level tools such as *Jupyter* notebooks and *Hadoop* data servers.

## 14.9 Future Challenges and Directions

Particle physics has historically been at the forefront of innovation in the adaptation of computing systems to fulfill its growing needs, as well as to accommodate tight budgets. It is rare, however, that the configurations developed in the particle physics context can directly be taken over by commercial systems. They should be considered as “precursors” or “early adoptions” of advances to come into general use. This situation has become a challenge with the growing time spans for experiments. If not properly managed, it may result in huge (often hidden) costs for maintenance and the lack of adoption of other innovations.

A first consideration for the future development is the number of scientific projects requiring wide-area distributed processing. In the short term, LHC will continue to produce unprecedented amounts of data and will be joined by a few experiments with similar needs, such as the Square Kilometer Array (SKA)<sup>91</sup> radiotelescope. However, many other experiments, such as neutrino detectors and astroparticle experiments, may produce orders of magnitude smaller data volumes, due to their own instrumental nature or to advances in data reduction within the data acquisition platforms. For example, even though 100 Gbps WAN will become common place by the 2020s, the baseline design of the Cherenkov Telescope Array (CTA)<sup>92</sup> requires only a 1 Gbps network interface for the sites (a capacity already available to home users in many countries). This is because CTA plans to use powerful clusters housed in compact containerized data centers placed at the sites to reduce the data at the instrument, thus avoiding the recording of large amounts of

---

<sup>90</sup>ARCHIVER – Archiving and Preservation for Research Environments project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 824516. More information at <https://archiver-project.eu>

<sup>91</sup>Dewdney, Peter E., et al. “The square kilometre array.” *Proceedings of the IEEE* 97.8 (2009): 1482–1496.

<sup>92</sup>Actis, M., et al. “Design concepts for the Cherenkov Telescope Array CTA: an advanced facility for ground-based high-energy gamma-ray astronomy.” *Experimental Astronomy* 32.3 (2011): 193–316.

raw data. Developments in the LHC domain, such as the LHCb turbo stream,<sup>93</sup> may also reduce the need for large-scale wide-area distributed processing. In parallel, as discussed below, data centers, particularly commercial ones, are growing in capacity. It is quite conceivable that the future needs of many experiments could be serviced by just a few data centers.

Another issue is the foreseen disappearance of magnetic tape technology. Tape has already disappeared from the commodity market, driven by ever lower disk prices and alternative backup media such as DVDs. If electrical energy considerations are excluded, the overall cost of high-end tape, associated robotic mechanisms and periodic migration to new media is becoming less attractive. Nevertheless, as of this writing, tape storage continues to be more economical than disk storage, especially if electrical costs are taken into account. The prospects are worrisome, however, as fewer companies continue to do R&D in tape technology. Hence, the particle physics community should be preparing for a hypothetical tape-less future by adapting its data management and cluster architectures as described above, allowing for multiple disk copies of the same dataset on a cluster or across a Grid or in several Clouds.

Grid and Cloud computing have brought along a resurgence of “timesharing”,<sup>94</sup> with research data centers deploying large clusters which give services to many projects of which only a minority are particle or astroparticle physics projects. The early adoption of Grids by particle physics, however, has resulted in development and deployment of a number of cluster and Grid tools which require extremely specific, non-commercially supported operating system or data server configurations which are of no interest to other projects or to the managers of general purpose clusters or Clouds. This does not mean that the developments led by particle physics are not worthwhile or of good quality; it simply means that further steps must be taken outside the particle physics scope to ensure a sustainable future and general applicability. Therefore, the particle physics community has to share knowledge about its developments, but also continuously evaluate alternative solutions. A case in point may be the use of non-standard data serving packages by LHC sites: it is perceived as the best solution (especially for LHC Tier-1 centers), but smaller particle or astroparticle projects with much lower data volumes often prefer what they perceive as a simpler, more standard solution such as NFS. Standard deployment of NFS on a cluster with thousands of CPU cores will however not be favored by system managers or commercial providers for a number of practical operational reasons, and they will propose newer technologies such as Ceph<sup>95</sup>

---

<sup>93</sup>Benson, Sean, et al. “The LHCb turbo stream.” *Journal of Physics: Conference Series*. Vol. 664. No. 8. IOP Publishing, 2015.

<sup>94</sup>The term “timesharing” was coined in the 1960s to describe the simultaneous use of a single “mainframe” computer for multiple batch or interactive tasks by multiplexing processes into the mainframe CPU.

<sup>95</sup>Weil, Sage A., et al. “Ceph: A scalable, high-performance distributed file system.” *Proceedings of the 7th symposium on Operating systems design and implementation*. USENIX Association, 2006.

or OpenStack swift<sup>96</sup> Object Storage. Therefore, a likely scenario is that current storage access techniques used by particle physics will become unsupported in the future. Of course, virtualization and Cloud techniques will make it possible for the particle physics community to deploy these legacy technologies themselves, with the corresponding expenses in personnel. A better way forward is to collaborate with industry, and between various Resource Centers, to test alternative tools and methods. Realistic evaluations are far from trivial, however, as they may require substantial investments by all parties, for example to provide UMD Grid or Cloud interfaces.

Another looming issue is the efficiency of use of the CPUs under data-intensive conditions. One basic assumption of the distributed architecture used by the particle physics community, as described above, is the existence of an infinitely powerful network connecting all elements, which in effect assumes infinitely powerful data servers as well as data ingestion by CPUs. Ever increasing data processing requirements and more and more powerful multi-core CPU nodes, however, are revealing data serving performance bottlenecks in clusters coming from the network and disk and CPU server limitations. Removing these bottlenecks may require more sophisticated architectures to be deployed.

A much more severe issue is related to the simplistic, often naïve, manner in which large particle physics projects pretend to use Grids and Clouds. They essentially desire them to behave as a very large cluster perfectly tuned to their needs, ignoring that imposing unreasonably high requirements in peak network and data serving rates within and between clusters can be very expensive. A Grid would better be viewed as a loosely connected federation of largely autonomous, self-sufficient clusters, with some of these clusters possibly hosted on commercial services.

These and other future challenges are being addressed in a coherent manner by the particle physics community through the HEP Software Foundation<sup>97,98</sup> and will certainly require a vigorous new cycle of research and development, in collaboration with computer scientists and engineers, the supercomputing and HPC communities and industry.

---

<sup>96</sup>Arnold, Joe. Openstack swift: Using, administering, and developing for swift object storage. "O'Reilly Media, Inc.", 2014.

<sup>97</sup>Alves Jr, Antonio Augusto. A Roadmap for HEP Software and Computing R&D for the 2020s. No. HSF-CWP-2017-001; HSF-CWP-2017-01; FERMILAB-PUB-17-607-CD; arXiv: 1712.06982. Fermi National Accelerator Lab. (FNAL), Batavia, IL (United States); Brookhaven National Laboratory (BNL), Upton, NY (United States); Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States); SLAC National Accelerator Lab., Menlo Park, CA (United States); Thomas Jefferson National Accelerator Facility (TJNAF), Newport News, VA (United States); Argonne National Lab. (ANL), Argonne, IL (United States), 2017.

<sup>98</sup>Slides from the Computing in High Energy Physics 2018 conference (proceedings to be published) [https://indico.cern.ch/event/587955/contributions/3012294/attachments/1681524/2708636/CHEP18\\_-\\_CWP\\_Lessons\\_and\\_Future\\_Work.pdf](https://indico.cern.ch/event/587955/contributions/3012294/attachments/1681524/2708636/CHEP18_-_CWP_Lessons_and_Future_Work.pdf)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 15

## Statistical Issues in Particle Physics



Louis Lyons

### 15.1 Introduction

In recent years there has been a growing awareness by particle physicists of the desirability of using good statistical practice. This is because the accelerator and detector facilities have become so complex and expensive, and involve so much physicist effort to build, test and run, that it is clearly important to treat the data with respect, and to extract the maximum information from them. The PHYSTAT series of Workshops and Conferences[1–12] has been devoted specifically to statistical issues in particle physics and neighbouring fields, and many interesting articles can be found in the relevant Proceedings. These meetings have benefited enormously from the involvement of professional statisticians, who have been able to provide specific advice as well as pointing us to some techniques which had not yet filtered down to Particle Physics analyses.

Analyses of experimental data in Particle Physics have, perhaps not surprisingly, tended to use statistical methods that have been described by other Particle Physicists. There are thus several books written on the subject by Particle Physicists[13]. The Review of Particle Physics properties[14] contains a condensed review of Statistics.

Another source of useful information is provided by the statistics committees set up by some of the large collaborations (see, for example, refs. [15–18]). Some conferences now include plenary talks specifically on relevant statistical issues (for example, Neutrino 2017[19], NuPhys17 and NuPhys18[20]), and the CERN

---

L. Lyons (✉)  
Physics, University of Oxford, Oxford, UK  
e-mail: [louis.lyons@physics.ox.ac.uk](mailto:louis.lyons@physics.ox.ac.uk)

Summer Schools for graduate students regularly have a series of lectures on statistics for Particle Physics[21].

This article is a slightly updated version of the one that appeared in ref. [22] in 2012.

### ***15.1.1 Types of Statistical Analysis***

There are several different types of statistical procedures employed by Particle Physicists:

- Separating signal from background: Almost every Particle Physics analysis uses some method to enhance the possible signal with respect to uninteresting background.
- Parameter determination: Many analyses make use of some theoretical or empirical model, and use the data to determine values of parameters, and their uncertainties and possible correlations.
- Goodness of fit: Here the data are compared with a particular hypothesis, often involving free parameters, to check their degree of consistency.
- Comparing hypotheses: The data are used to see which of two hypotheses is favoured. These could be the Standard Model (SM), and some specific version of new physics such as the existence of SuperSYmmetry (SUSY), or the discovery of the Higgs boson[23].
- Decision making: Based on one's belief about the current state of physics, the value of possible discoveries and estimates of the difficulty of future experiments, a decision is made on what should be thrust of future research. This subject is beyond the scope of this article.

### ***15.1.2 Statistical and Systematic Uncertainties***

In general any attempt to measure a physics parameter will be affected by statistical and by systematic uncertainties. The former are such that, if the experiment were to be repeated, random effects would result in a distribution of results being obtained. These can include effects due to the limited accuracy of the measurement devices and/or the experimentalist; and also from the inherent Poisson variability of observing a number of counts  $n$ . On the other hand, there can be effects that shift the measurements from their true values, and which need to be corrected for; uncertainties in these corrections contribute to the systematics. Another systematic effect could arise from uncertainties in theoretical models which are used to interpret the data. Scientists' systematics are often 'nuisance parameters' for statisticians.

Consider an experiment designed to measure the temperature at the centre of the sun by measuring the flux of solar neutrinos on earth. The main statistical

uncertainty might well be that due to the limited number of neutrino interactions observed in the detector. On the other hand, there are likely to be systematics from limited knowledge of neutrino cross-sections in the detector material, the energy calibration of the detector, neutrino oscillation parameters, models of energy convection in the sun, etc. If some calibration measurement or subsidiary experiment can be performed, this effectively converts a systematic uncertainty into a statistical one. Whether this source of uncertainty is quoted as statistical or systematic is not crucial; what is important is that possible sources of correlation between uncertainties here and in other measurements (in this or in other experiments) are well understood.

The magnitude of systematic effects in a parameter-determination situation can be assessed by fitting the data with different values of the nuisance parameter(s), and seeing how much the result changes<sup>1</sup> when the nuisance parameter value is varied by its uncertainty. Alternatively the nuisance parameter(s) for systematic effects can be incorporated into the likelihood or  $\chi^2$  for the fit; or a Bayesian method involving the prior probability distribution for the nuisance parameter can be used. (See Sects. 15.4.5 and 15.7.6 for ways of incorporating nuisance parameters in upper limit and in  $p$ -value calculations respectively).

How to assess systematics was much discussed at the first Banff meeting[6] and at PHYSTAT-LHC[24–26]. A special session of the recent PHYSTAT $\nu$  meeting at CERN[12] was devoted to systematics. Many reviews of this complex subject exist and can be traced back via ref. [27].

In general, much more effort is involved in estimating systematic uncertainties than for parameter determination and the corresponding statistical uncertainties; this is especially the case when the systematics dominate the statistical uncertainty.

Cowan[35] has considered the effect of having an uncertainty in magnitude of a systematic effect. As Cox has remarked[36], there is a difference in knowing that a correction has almost precisely a 20% uncertainty, or that it is somewhere between 0% and 40%.

### 15.1.3 *Bayes and Frequentism*

These are two fundamental approaches to making inferences about parameters or whether data support particular hypotheses. There are also other methods which do not correspond to either of these philosophies; the use of  $\chi^2$  or the likelihood are examples.

Particle physicists tend to favour a frequentist method. This is because in many cases we really believe that our data are representative as samples drawn according to the model we are using (decay time distributions often are exponential; the counts

---

<sup>1</sup>If the simulation yields a change in the result of  $a \pm b$ , there is much discussion about how the contribution to the systematic uncertainty should be assessed in terms of  $a$  and  $b$ —see ref. [27].



in repeated time intervals do follow a Poisson distribution; etc.), and hence we want to use a statistical approach that allows the data “to speak for themselves”, rather than our analysis being sensitive to our assumptions and beliefs, as embodied in the assumed Bayesian priors. Bayesians would counter this by remarking that frequentist inference can depend on the reference ensemble, the ordering rule, the stopping rule, etc.

With enough data, the results of Bayesian and frequentist approaches usually tend to agree. However, in smallish data samples numerical results from the two approaches can differ.

### 15.1.3.1 Probability

There are at least three different approaches to the question of what probability is. The first is the mathematical one, which is based on axioms e.g. it must lie in the range 0–1; the probabilities of an event occurring and of it not occurring add up to 1; etc. It does not give much feeling for what probability is, but it does provide the underpinning for the next two methods.

Frequentists, not surprisingly, define probability in terms of frequencies in a long series of essentially identical repetitions<sup>2</sup> of the relevant procedure. Thus the probability of the number 5 being uppermost in throws of a die is 1/6, because that is the fraction of times we expect (or approximately observe) it to happen. This implies that probability cannot be defined for a specific occurrence (Will the first astronaut who lands on Mars return to earth alive?) or for the value of a physical constant (Does Dark Matter contribute more than 25% of the critical density of the Universe?).

In contrast, Bayesians define probability in terms of degree of belief. Thus it can be used for unique events or for the values of physical constants. It can also vary from person to person, because my information may differ from yours. The numerical value of the probability to be assigned to a particular statement is determined by the concept of a ‘fair bet’; if I think the probability (or ‘Bayesian credibility’) of the statement being true is 20%, then I must offer odds of 4-to-1, and allow you to bet in either direction.

This difference in approach to probability affects the way Bayesians and frequentists deal with statistical procedures. This is illustrated below by considering parameter determination.

### 15.1.3.2 Bayesian Approach

The Bayesian approach makes use of Bayes’ Theorem:

$$p(A|B) = p(B|A) \times p(A)/p(B), \quad (15.1)$$

---

<sup>2</sup>Bayesians attack this concept of ‘essentially identical trials’, claiming that it is hard to define it without using the concept of probability, thus making the definition circular.

where  $p(A)$  is the probability or probability density of  $A$ , and  $p(A|B)$  is the conditional probability for  $A$ , given that  $B$  has happened. This formula is acceptable to frequentists, provided the probabilities are frequentist probabilities. However Bayesians use it with  $A = \text{parameter (or hypothesis)}$  and  $B = \text{data}$ . Then

$$p(\text{parameter}|\text{data}) \propto p(\text{data}|\text{parameter}) \times p(\text{parameter}), \quad (15.2)$$


where the three terms are respectively the Bayesian posterior, the likelihood function and the Bayesian prior. Thus Bayes' theorem enables us to use the data (as encapsulated in the likelihood) to update our prior knowledge ( $p(\text{parameter})$ ); the combined information is given by the posterior.

Frequentists object to the use of probability for physical parameters. Furthermore, even Bayesians agree that it is often hard to specify a sensible prior. For a parameter which has been well determined in the past, a prior might be a gamma function or log-normal or a (possibly truncated) Gaussian distribution of appropriate central value and width, but for the case where no useful information is available the choice is not so clear; it is easier to parametrise prior knowledge than to quantify prior ignorance. The 'obvious' choice of a uniform distribution has the problem of being not unique (Should our lack of knowledge concerning, for example, the mass of a neutrino  $m_\nu$  be parametrised by a uniform prior for  $m_\nu$  or for  $m_\nu^2$  or for  $\log m_\nu$ , etc?). Also a uniform prior over an infinite parameter range cannot be normalised. For situations involving several parameters, the choice of prior becomes even more problematic.

It is important to check that conclusions about possible parameter ranges are not dominated by the choice of prior. This can be achieved by changing to other 'reasonable' priors (sensitivity analysis); or by looking at the posterior when the data has been removed.

### 15.1.3.3 Frequentist Approach: Neyman Construction

The frequentist way of constructing intervals completely eliminates the need for a prior, and avoids considering probability distributions for parameters. Consider a measurement  $x$  which provides information concerning a parameter  $\mu$ . For example, we could use a month's data from a large solar neutrino detector ( $x$ ) to estimate the temperature at the centre of the sun ( $\mu$ ). It is assumed that enough is known about solar physics, fusion reactions, neutrino properties, the behaviour of the detector, etc. that, for any given value of  $\mu$ , the probability density for every  $x$  is calculable. Then for that  $\mu$ , we can select a region in  $x$  which contains, say, 90% of this probability. If we do this for every  $\mu$ , we obtain a 90% confidence band; it shows



**Fig. 15.1** The Neyman construction for setting a confidence range on a parameter  $\mu$ . At any value of  $\mu$ , it is assumed that we know the probability density for obtaining a measured value  $x$ . (For example,  $\mu$  could be the temperature of the fusion reactor at the centre of the Sun, while  $x$  is the solar neutrino flux, estimated by operating a large underground solar neutrino detector for 1 month.) We can then choose a region in  $x$  which contains, say, 90% of the probability; this is denoted by the solid part of the horizontal line. By repeating this procedure for all possible  $\mu$ , the band between the curved lines is constructed. This confidence band contains the likely values of  $x$  for any  $\mu$ . For a particular measured value  $x_2$ , the confidence interval from  $\mu_l$  to  $\mu_u$  gives the range of parameter values for which that measured value was likely. For  $x_2$ , this interval would be two-sided, while for a lower value  $x_1$ , an upper limit would be obtained. In contrast, there are no parameter values for which  $x_0$  is likely, and for that measured value the confidence interval would be empty

the values of  $x$  which are likely results<sup>3</sup> of the experiment for any  $\mu$ , assuming the theory is correct (see Fig. 15.1). Then if the actual experiment gives a measurement  $x_2$ , it is merely necessary to find the values of  $\mu$  for which  $x_2$  is in the confidence band. This is the Neyman construction.

Of course, the choice of a region in  $x$  to contain 90% of the probability is not unique. The one shown in Fig. 15.1 is a central one, with 5% of the probability on either side of the selected region. Another possibility would be to have a region with 10% of the probability to the left, and then the region in  $x$  extends up to infinity. This choice would be appropriate if we always wanted to quote upper limits on  $\mu$ . Other choices of ‘ordering rule’ are also possible (see, for example, Sect. 15.4.3).

The Neyman construction can be extended to more parameters and measurements, but in practice it is very hard to use it when more than two or three parameters are involved; software to perform a Neyman construction efficiently in several dimensions would be very welcome. The choice of ordering rule is also very important. Thus from a pragmatic point of view, even ardent frequentists

<sup>3</sup>The adjective ‘likely’ is appropriate for central intervals. For upper limits on  $\mu$ , however, the accepted values of  $x$  for a given  $\mu$  extend to infinity, and so ‘preferred results for the given ordering rule’ would be more appropriate.

are prepared to use Bayesian techniques for multidimensional problems (e.g. with systematics). They would, however, like to ensure that the technique they use provides parameter intervals with reasonable frequentist coverage.

#### 15.1.3.4 Coverage

One of the major advantages of the frequentist Neyman construction is that it guarantees coverage. This is a property of a statistical technique<sup>4</sup> for calculating intervals, and specifies how often the interval contains the true value  $\mu_t$  of the parameter. This can vary with  $\mu_t$ .

For example, for a Poisson counting experiment with parameter  $\mu$  and observed number  $n$ , a (not very good) method for providing an interval for  $\mu$  is  $n \pm \sqrt{n}$ . Thus an observed  $n = 2$  would give a range 0.59–3.41 for  $\mu$ . If  $\mu = 2.01$ , observed values  $n = 2, 3$  and 4 result in intervals that include  $\mu = 2.01$ , while other values of  $n$  do not. The coverage of this procedure for  $\mu = 2.01$  is thus the sum of the Poisson probabilities for having  $n = 2, 3$  or 4 for the given  $\mu$ .

For a discrete observable (e.g. the number of detected events in a search for Dark Matter), there are jumps in the coverage; in order to avoid under-coverage, there is necessarily some over-coverage. However, for a continuous observable (e.g. the estimated mass of the Higgs boson) the coverage can be exact.

Coverage is not guaranteed for methods that do not use the Neyman construction (see Sect. 15.2.1). Interesting plots of coverage as a function of the parameter value for the simple case of a Poisson counting experiment can be found in ref. [32].

#### 15.1.3.5 Likelihoods

The likelihood approach makes use of the probability density function (*pdf*) for observing the data, evaluated for the data actually observed.<sup>5</sup> It is a function of any parameters, although it does not behave like a probability density for them. It provides a method for determining values of parameters. These include point estimates for the ‘best’ values, and ranges (or contours in multi-parameter situations) to characterise the uncertainties. It usually has good properties asymptotically, but a major use is with sparse multi-dimensional data.

The likelihood method is neither frequentist nor Bayesian. It thus does not guarantee frequentist coverage or Bayesian credibility. It does, however, play a central role in the Bayesian approach, which obtains the posterior probability

---

<sup>4</sup>It is important to realise that coverage is a property of the **method**, and not of an **individual measurement**.

<sup>5</sup>The *pdf*  $f(x, \mu_0)$  gives the probability density for obtaining various data  $x$  when the parameter has some specified value  $\mu_0$ . The likelihood is the same function of two variables  $f(x_0, \mu)$ , but now with  $x_0$  fixed at the data actually obtained, and  $\mu$  regarded as the variable.

density by multiplying the likelihood by the prior. The Bayesian approach thus obeys the likelihood principle, which states that the only way the experimental data affects inference is via the likelihood function. In contrast, the Neyman construction requires not only the likelihood for the actual data, but also for all possible data that might have been observed.

Because the likelihood is not a probability density, it does not transform like one. Thus the value of the likelihood for a parameter  $\mu_0$  is identical to that for  $\lambda_0 = 1/\mu_0$ . This means that ratios of likelihoods (or differences in their logarithms) are useful to consider, but that the integration of tails of likelihoods is not a recognised statistical procedure.

A longer account of the Bayesian and frequentist approaches can be found in ref. [28]. Reference [29] provides a very readable account for a Poisson counting experiment.

## 15.2 Likelihood Issues

In this section, we discuss some potential misunderstandings of likelihoods.

### 15.2.1 $\Delta(\ln L) = 0.5$ Rule

In the maximum likelihood approach to parameter determination, the best value  $\lambda_0$  of a parameter is determined by finding where the likelihood maximises; and its uncertainty is estimated by finding how much the parameter must be changed<sup>6</sup> in order for the logarithm of the likelihood to decrease by 0.5 as compared with the maximum.<sup>7</sup> From a frequentist viewpoint, this should ideally result in the parameter range having 68% coverage. That is, in repeated use of this procedure to estimate the parameter, 68% of the intervals should contain the true value of the parameter, whatever its true value happens to be.

If the measurement is distributed about the true value as a Gaussian with constant width, the likelihood approach will yield exact coverage, but in general this is not so. For example, Garwood[31] and Heinrich[32] have investigated the properties of the likelihood approach (and other methods too) to estimate  $\mu$ , the mean of a Poisson, when  $n_{obs}$  events are observed. Because  $n_{obs}$  is a discrete variable, the coverage is

---

<sup>6</sup>If there are more than just one parameter, the likelihood must of course be remaximised with respect to all the other parameters when looking for the  $\Delta(\ln L) = 0.5$  points. Alternatively, a region in multi-parameter space can be selected by finding the contour at which  $\Delta(\ln L)$  decreases from its maximum by an amount which depends on the number of parameters.

<sup>7</sup>This (like several other methods) can give rise to asymmetric uncertainties. Techniques for dealing with this have been discussed by Barlow[30].

a discontinuous function of  $\mu$ , and varies from 100% at  $\mu = 0$  down to 30% at  $\mu \approx 0.5$ .<sup>8</sup>

### 15.2.2 Unbinned Maximum Likelihood and Goodness of Fit

With sparse data, the unbinned likelihood method is a good one for estimating parameters of a model. In order to understand whether these estimates of the parameters are meaningful, we need to know whether the model provides an adequate description of the data. Unfortunately, as emphasised by Heinrich[33], the magnitude of the unbinned maximum likelihood is often independent of whether or not the data agree with the model. He illustrates this by the example of the determination of the lifetime  $\tau$  of a particle whose decay distribution is  $(1/\tau) \exp(-t/\tau)$ . For a set of observed times  $t_i$ , the maximum likelihood  $L_{max}$  depends on the data  $t_i$  only through their average value  $\bar{t}$ . Thus any data distributions with the same  $\bar{t}$  would give identical  $L_{max}$ , which demonstrates that, at least in this case,  $L_{max}$  gives no discrimination about whether the data are consistent with the expected distribution.

Another example is fitting an expected distribution  $(1 + \alpha \cos^2 \theta)/(1 + \alpha/3)$  to data  $\theta_i$  on the decay angle of some particle, to determine  $\alpha$ . According to the expected functional form, the data should be symmetrically distributed about  $\cos \theta = 0$ . However, the likelihood depends only on the **square** of  $\cos \theta$ , and so would be insensitive to all the data having  $\cos \theta_i$  negative; this would be very inconsistent with the expected symmetric distribution.

In contrast Baker and Cousins[34] provide a likelihood method of measuring goodness of fit for a data **histogram** compared to a theory. The Poisson likelihood  $P_{Pois}(n|\mu)$  for each bin is compared with that for the best possible predicted value  $\mu_{best} = n$  for that bin. Thus the Baker-Cousins likelihood ratio

$$LR_{BC} = \Pi \frac{e^{-\mu_i} \mu_i^{n_i} / n_i!}{e^{-n_i} n_i^{n_i} / n_i!} = \Pi e^{(n_i - \mu_i)} (\mu_i / n_i)^{n_i} \quad (15.3)$$

is such that asymptotically  $-2 \ln LR_{BC}$  is distributed as  $\chi^2$ .<sup>9</sup> For small  $\mu$ , the Baker-Cousins likelihood ratio is better than a weighted sum of squares for assessing goodness of fit.

<sup>8</sup>It is of course not surprising that methods that are expected to have good asymptotic behaviour may not display optimal properties for  $\mu \approx 0$ .

<sup>9</sup>The binned Poisson likelihood is not a measure of fit. This is because, for example,  $\mu_i = n_i = 1$  and  $\mu_i = n_i = 100$  both correspond to perfect agreement between data and prediction, but  $P_{Pois}(1|1.0)$  is much larger than  $P_{Pois}(100|100.0)$ .

### 15.2.3 Profile Likelihood

In many situations the likelihood is a function not only of the parameter of interest  $\phi$  but also other parameters. These may be other physics parameters (for example, in neutrino oscillation experiments where the mixing angles and differences in mass-squared of the various neutrinos are relevant), but can also be nuisance parameters  $\nu$  associated with systematic effects (e.g. jet energy scales, particle identification efficiencies, etc.). To make statements about  $\phi$ , the likelihood  $L(\phi, \nu)$  is often ‘profiled’ over the nuisance parameters, i.e. at each value of  $\phi$ , the likelihood is remaximised with respect to  $\nu$ . Thus

$$L_{prof}(\phi) = L(\phi, \nu_{max}(\phi)) \quad (15.4)$$

Then  $L_{prof}(\phi)$  is used much as the ordinary likelihood when there are no nuisance parameters.

A profile likelihood is in general wider than the likelihood for a fixed value of the nuisance parameter  $\nu$ ; this results in the uncertainty in the parameter of interest  $\phi$  being larger when allowance is made for the systematic uncertainties.

In the standard profile likelihood,  $\nu$  is a continuous variable. An extension of this has been used by Dauncey et al. [38], to allow for uncertainties in the choice of functional form of the background parametrisation in searches for new particles as peaks above background in a mass spectrum. Here the systematic is discrete, rather than continuous.

An alternative way of eliminating nuisance parameters (known as marginalisation) is to use  $L(\phi, \nu)$  as part of a Bayesian procedure, and then to integrate the Bayesian posterior over  $\nu$ , i.e.

$$P_{marg}(\phi) = \int P_{post}(\phi, \nu) d\nu \quad (15.5)$$

Of course, both profiling and marginalisation result in the loss of information. Reference [37] provides a very trivial example of this for profile likelihoods.

### 15.2.4 Punzi Effect

Sometimes we have two or more nearby peaks, and we try to fit our data in order to determine the fractions of each peak. Punzi [39] has pointed out that it is very easy to write down a plausible but incorrect likelihood function that gives a biased result. This occurs in situations where the events have experimental resolutions  $\sigma$  in the observable  $x$  that vary event-by-event; and the distributions of  $\sigma$  are different for the two peaks.

For a set of observations  $x_i$ , it is tempting but wrong to write the unbinned likelihood as

$$L(f)_{wrong} = \Pi\{f * G(x_i, 0.0, \sigma_i) + (1 - f) * G(x_i, 1.0, \sigma_i)\} \quad (15.6)$$

where  $f$  is the fraction of the first peak (labelled  $A$  below) which is parametrised as  $G(x_i, 0.0, \sigma_i)$ , a Gaussian in  $x_i$ , centred on zero, and with width  $\sigma_i$ , and  $i$  is the label for the  $i$ th event; and similarly for the second peak (labelled  $B$ ), except that it is centred at unity.

Application of the rules of conditional probability shows that the correct likelihood is

$$L(f)_{right} = \Pi\{f * G(x_i, 0.0, \sigma_i) * p(\sigma_i|A) + (1 - f) * G(x_i, 1.0, \sigma_i) * p(\sigma_i|B)\} \quad (15.7)$$

where  $p(\sigma_i|A)$  and  $p(\sigma_i|B)$  are the probability densities for the resolution being  $\sigma_i$  for the  $A$  and  $B$  peaks respectively. We then see that  $L(f)_{wrong}$  and  $L(f)_{right}$  give identical values for  $f$ , provided that  $p(\sigma_i|A) = p(\sigma_i|B)$ . If however, the distributions of the resolution differ,  $L(f)_{wrong}$  will in general give a biased estimate.

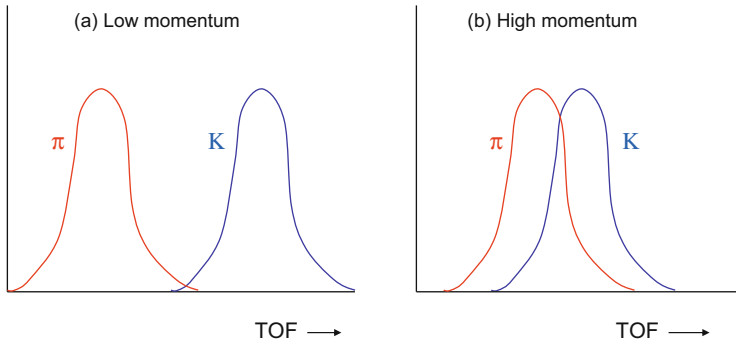
Punzi investigated the extent of this bias in a simple Monte Carlo simulation, and it turns out to be surprisingly large. For example, with  $f = 1/3$ , and  $p(\sigma_A)$  and  $p(\sigma_B)$  being  $\delta$ -functions at 1.0 and at 2.0 respectively (i.e.  $\sigma = 1$  for all  $A$  events, and  $\sigma = 2$  for all  $B$  events), the fitted value of  $f$  from  $L(f)_{wrong}$  turned out to be 0.65. Given that  $f$  is confined to the range from zero to unity, this is an enormous bias.

The way the bias arises can be understood as follows: The fraction  $f$  of the events that are really  $A$  have relatively good resolution, and so the fit to them alone would assign essentially all of them as belonging to  $A$  i.e. these events alone would give  $f \approx 1$  with a small uncertainty. In contrast the  $1 - f$  of the events that are  $B$  have poor resolution, so for them the fit does not mind too much what is the value of  $f$ . But the fit uses all the events together, and so assigns a single  $f$  to the complete sample; this will be a weighted average of the  $f$  values for the  $A$  and for the  $B$  events. Because the  $A$  events result in a more accurate determination of  $f$  than do the  $B$  events, the fitted  $f$  will be biased upwards (i.e. it will over-estimate the fraction of events corresponding to the peak with the better resolution).

The Punzi effect can also appear in other situations, such as particle identification. Different particle types (e.g. pions and kaons) would appear as different peaks in the relevant particle-identification variable e.g. time of flight, rate of energy loss  $dE/dx$ , angle of Cherenkov radiation, etc. The separation of these peaks for the different particle types depends on the momentum of the particles (see Fig. 15.2). The incorrect  $L$  is now

$$L_{wrong}(f_K) = \Pi\{(1 - f_K) * G(x_i, x_\pi(p_i), \sigma_i) + f_K * G(x_i, x_K(p_i), \sigma_i)\} \quad (15.8)$$





**Fig. 15.2** The Punzi effect in particle identification. The diagrams show the expected (normalised) distributions of the output signal from a particle identifier, for pions and for kaons **(a)** at low momentum where separation is easier, and **(b)** at high momentum where the distributions overlap. Because kaons are heavier than pions, they tend to have larger momenta. Because it is hard at high momentum to distinguish pions from kaons, the likelihood function is insensitive to whether these tracks are classified as pions or kaons, and hence the fraction of high momentum tracks classified as kaons will have a large uncertainty. In contrast, low momentum tracks will be correctly identified. Thus if the plausible but incorrect likelihood function that ignores the pion and kaon momentum distributions is used to determine the overall fraction of kaons, it will be biased downwards towards the fraction of low momentum particles that are kaons

where  $x_\pi(p_i)$  and  $x_K(p_i)$  are the expected positions of the particle identification information for a particle of momentum  $p_i$ , and  $x_i$  is the observed value for the  $i$ th event. So here the Punzi bias can arise even with constant resolution, because the momentum spectra of pions and kaons can be different. To avoid the bias, the likelihood needs to incorporate information on the different momentum distributions of pions and of kaons. If these momentum distributions are different enough from each other, it could be that the likelihood function bases its separation of the different particle types on the momenta of the particles rather than on the data from the detector's particle identifier. Catastini and Punzi[40] avoid this by using parametric forms for the momentum distributions of the particles, with the parameters being determined by the data being analysed.

The common feature potentially leading to bias in these two examples is that the ratio of peak separation to resolution is different for the two types of objects. For the first example of separating the two peaks, it was the denominators that were different, while in the particle identification problem it was the numerators.

The Punzi bias may thus occur in situations where the templates in a multi-component fit depend on additional observations whose distributions are not explicitly included in the likelihood.

## 15.3 Separating Signal from Background

Almost every Particle Physics analysis uses some technique for separating possible signal from background. First some simple ‘cuts’ are applied; these are generally loose selections on single variables, which are designed to remove a large fraction of the background while barely reducing the real or potential signal. Then to obtain a better separation of signal from background in the multi-dimensional space of the event observables, methods like Fisher discriminants, decision trees, artificial neural networks (including Bayesian nets and more recently deep neural nets), support vector machines, etc. are used[41, 42]. Extensions of these methods involve bagging, boosting and random forests, which have been used to achieve improved performance of the separation as seen on a plot of signal efficiency against background mis-acceptance rate. A description of the software available for implementing some of these techniques can be found in the talks by Narsky[43] and by Tegenfeldt[44] at the PHYSTAT-LHC Workshop.

More recently, deep learning techniques are rapidly becoming popular. In Particle Physics, they have been used for on-line triggering, tracking, fast simulation, object identification, image recognition, and event-by-event separation of signal from background. Reference [45] provides good introductions to the use of these methods for Particle Physics. There are now regular workshops and lectures on Machine Learning at CERN and at Fermilab (see refs. [46] and [47]), as well as at many universities.

The signal-to-background ratio before this multivariate stage can vary widely, as can the signal purity after it. If some large statistics study is being performed (e.g. to use a large sample of events to obtain an accurate measurement of the lifetime of some particle), then it is not a disaster if there is some level of background in the finally selected events, provided that it can be accurately assessed and allowed for in the subsequent analysis. At the other extreme, the separation technique may be used to see if there is any evidence for the existence of some hypothesised particle (the potential signal), in the presence of background from well-known sources. Then the actual data may in fact contain no observable signal.

These techniques are usually ‘taught’ to recognise signal and background by being given examples consisting of large numbers of events of each type. These may be produced by Monte Carlo simulation, but then there is a problem of trying to verify that the simulation is a sufficiently accurate representation of reality. It is better to use real data for this, but the difficulty then is to obtain sufficiently pure samples of background and signal. Indeed, for the search for a new particle, true data examples do not exist. However, it is the accurate representation of background that is likely to pose a more serious problem.

The way that, for example, neural networks are trained is to present the software with approximately equal numbers of signal and background events<sup>10</sup> and then

---

<sup>10</sup>For searches for rare processes, it is clearly inappropriate to use the actual fractions expected in the data to determine the ratio of signal to background Monte Carlo events to be used as the

to minimise a cost function  $C$  for the network. This is usually defined as  $C = \sum (z_i - t_i)^2$ , where  $z_i$  is the trained network's output for the  $i$ th event;  $t_i$  is the target output, usually chosen as 1 for signal and zero for background; and the summation is over all testing events presented to the network. The problem with this is that  $C$  is only loosely related to what we really want to optimise. For a search for a new particle this could be the sensitivity of the experimental upper limit in the absence of signal, while for a high statistics analysis measuring the properties (such as mass or lifetime) of some well-established particle, we would be interested in minimising the uncertainty (including systematic effects) on the result, without the training procedure biasing the measurement.

As with all event separation methods, it is essential to check the performance of a trained procedure by using a set of events that are independent of those used for training. This is to ensure that the network does not use specific but irrelevant features of the training events in its learning process, but can achieve good performance on unseen data.

Some open questions are:

- How can we check that our multi-dimensional training samples for signal and background are reliable descriptions of reality; and that they cover the region of multi-dimensional space populated by the data?
- How should the ratio of the numbers of signal and background training events be chosen, especially when there are several different sources of background?
- What is the best way of allowing for nuisance parameters in the models of the signal and/or background?[25, 48]
- Are there useful and easy ways of optimising on what is really of interest?[49]

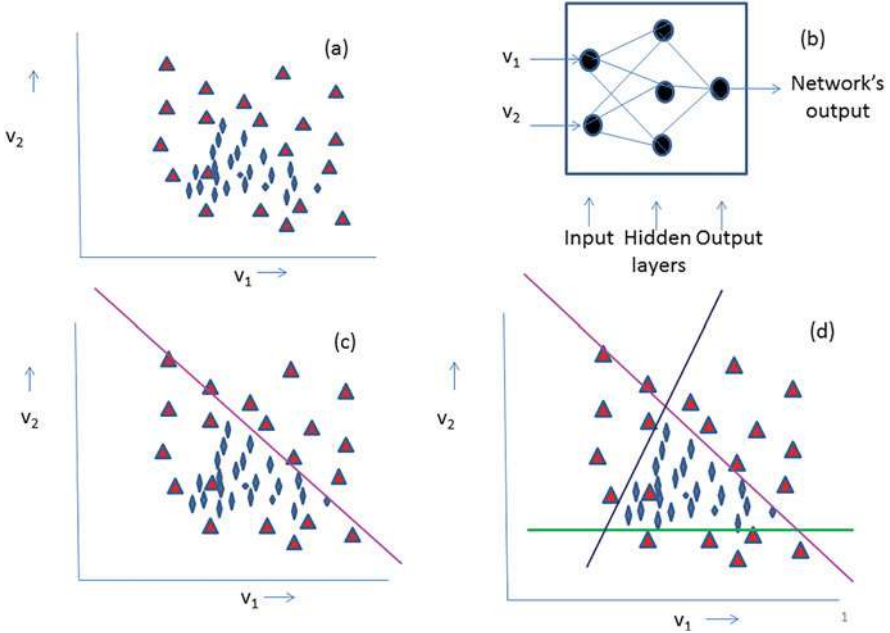
### 15.3.1 *Understanding How Neural Networks Operate*

It is useful to appreciate how neural networks operate in providing a good separation of signal and background, as this can help in choosing a suitable architecture for the network.

Figure 15.3a shows some hypothetical signal and background events in terms of two measured variables  $x$  and  $y$  for each event. A network with two inputs ( $x$  and  $y$ ), a single hidden layer with 3 nodes, and a single output is used; it aims to give 1 for signal and zero for background events (see Fig. 15.3b). This is achieved by training the network with  $(x, y)$  values for known examples of signal and background; and allowing the network to vary its internal parameters to minimise a suitably defined cost function e.g.  $\sum (z_e - t_e)^2$ , where the summation is over the training events, and  $z_e$  and  $t_e$  are the network's output and its target value (0 or 1) respectively.

---

training sample, because the network could then achieve a very small cost  $C$  simply by classifying everything as background.



**Fig. 15.3** (a) A 2 –  $D$  plot showing the regions of the variables  $v_1$  and  $v_2$  for the signal (dots) and background (triangles). (b) The neural network used for separating signal and background. (c) The top hidden node receives inputs from  $v_1$  and  $v_2$ . With suitable weights and threshold and a large value of  $\beta$ , the node's output will be on for  $(v_1, v_2)$  values below the diagonal line. (d) Similarly for the other two hidden nodes, their outputs can be on for  $(v_1, v_2)$  below the other diagonal line, and above the horizontal one, respectively. A further choice of weights to the output node and its threshold can ensure that the whole network's output will be on only if all three hidden nodes' outputs are on, i.e. if  $(v_1, v_2)$  values are within the triangle in (d)

The input  $q_i$  to a given hidden node  $i$  is a linear combination of the input variables  $x$  and  $y$

$$q_i = w_{xi}x + w_{yi}y + t_i \quad (15.9)$$

where the weights  $w$  and threshold  $t$  are varied during the fitting process. The output  $r$  from any hidden node is determined from its input  $q$  by something like a sigmoid function e.g.

$$r = 1/(1 + e^{-\beta q}) \quad (15.10)$$

This switches from zero for large negative  $q$  to unity for large positive  $q$ . The switch occurs around  $q = 0$ , and width of the region depends on the network parameter  $\beta$ . For very large  $\beta$ , there is a rapid switch from zero to unity. In terms of  $x$  and  $y$ , this

means that the hidden node  $i$  is ‘on’ (i.e.  $r_i = 1$ ) if

$$w_{xi}x + w_{yi}y + t_i > 0, \quad (15.11)$$

or ‘off’ otherwise. Thus the boundary between events having  $r_i$  on or off is a straight line in the  $(x, y)$  plane (see Fig. 15.3c). With suitable values for the weights and thresholds for the three hidden nodes, there will be three straight line boundaries in the  $(x, y)$  plane shown in Fig. 15.3d. Finally, to produce the “and” of these three conditions, the weights  $w_{jo}$  (from the hidden node  $j$  to the output node  $o$ ) and the output threshold  $t_o$  can be set as

$$w_{1o} = w_{2o} = w_{3o} = 0.4 \quad t_o = -1.0 \quad (15.12)$$

to ensure that the output will be “on” only if the three hidden layers are all “on”, i.e. that the selected input values are inside the triangular region in the  $(x, y)$  plane. With  $\beta$  set at a lower level, the contour for the selected region will be smoother with rounded corners, rather than being triangular.

It would be useful to have a similar understanding of how deep networks operate. Tishby[50] has provided some insight on what happens in the hidden layers of a deep neural network during the training procedure.

## 15.4 Parameter Determination

For a single parameter (e.g. the branching ratio for  $H \rightarrow \mu^+\mu^-$ ) the parameter range could be either a 2-sided interval or just an upper limit, at some confidence level (typically 68% for 2-sided intervals, but usually 90% and 95% for upper limits). For two parameters (e.g. mass and production rate for some new particle  $X$  that decays to a top pair), their acceptable values could be those inside some 2-dimensional confidence region. Alternatively an upper limit or 2-sided region for one parameter as a function of the other could be defined; these are known as a Raster Scan.

An upper limit on 2-variables is not a well-defined concept.

### 15.4.1 Upper Limits

Most recent searches for new phenomena have not found any evidence for exciting new physics. Examples from particle physics include searches for SUSY particles, dark matter, etc.; attempts to find substructure of quarks or leptons; looking for extra spatial dimensions; measuring the mass of the lightest neutrino; etc. Rather than just saying that nothing was found, it is more useful to quote an upper limit on the sought-for effect, as this could be useful in ruling out some theories. For example in

1887, Michelson and Morley[52] attempted to measure the speed of the Earth with respect to the aether. No effect was seen, but the experiment was sensitive enough to lead to the demise of the aether theory.

A simple scenario is a counting experiment where a background  $b$  is expected from conventional sources, together with the possibility of an interesting signal  $s$ . The number of counts  $n$  observed is expected to be Poisson distributed with a mean  $\mu = \epsilon s + b$ , where  $b$  is the expected number of events from background, and  $\epsilon$  is a factor for converting the basic physics parameter  $s$  into the number of signal events expected in our particular experiment; it thus allows for experimental inefficiency, the experiment's running time; etc. Then given a value of  $n$  which is comparable to the expected background, what can we say about  $s$ ? The true value of the parameter  $s$  is constrained to be non-negative. The problem is interesting enough if  $b$  and  $\epsilon$  are known exactly; it becomes more complicated when only estimates with uncertainties  $\sigma_b$  and  $\sigma_\epsilon$  are available.

An extension of the simple counting scenario is when a search for a new particle is carried out over a range of masses. This is usually dealt with by performing separate searches at a series of masses over a specified range. This 'Raster Scan' is in contrast with a method that regards the sought-for new particle's mass and its production rate as two parameters to be estimated simultaneously. The relative merits of these two approaches are described in ref. [51].

Even without the nuisance parameters, a variety of methods is available. These include likelihood,  $\chi^2$ , Bayesian with various priors for  $s$ , frequentist Neyman constructions with a variety of ordering rules for  $n$ , and various *ad hoc* approaches. The methods give different upper limits for the same data.<sup>11</sup> A comparison of several methods can be found in ref. [53]. The largest discrepancies arise when the observed  $n$  is less than the expected background  $b$ , presumably because of a downward statistical fluctuation. The following different behaviours of the limit (when  $n < b$ ) can be obtained:

- Frequentist methods can give **empty** intervals for  $s$  i.e. there are no values of  $s$  for which the data are likely. Particle physicists tend to be unhappy when their years of work result in an empty interval for the parameter of interest, and it is little consolation to hear that frequentist statisticians are satisfied with this feature, as it does not necessarily lead to undercoverage.

When  $n$  is not quite small enough to result in an empty interval, the upper limit might be **very small**.<sup>12</sup> This could confuse people into thinking that the experiment was much more sensitive than it really was.

- The Feldman-Cousins frequentist method[54] (see Sect. 15.4.3) that employs a likelihood-ratio ordering rule gives upper limits which **decrease** as  $n$  gets smaller

<sup>11</sup>By coincidence, the upper limits obtained by the Bayesian approach with an (improper) flat prior for  $s$  and by the appropriate Neyman construction agree when  $b = 0$ .

<sup>12</sup>Bayesian methods that use priors with part of the probability density being a  $\delta$ -function at  $s = 0$  can result in a posterior with an enhanced  $\delta$ -function at zero, such that the upper limit contains only the single point  $s = 0$ .

at constant  $b$ . A related effect is the growth of the limit as  $b$  decreases at constant  $n$ —this can also occur in other frequentist approaches. Thus if no events are observed ( $n = 0$ ), the upper limit of a 90% Feldman-Cousins interval is 1.08 for  $b = 3.0$ , but 2.44 for  $b = 0$ . This is sometimes presented as a paradox, in that if a bright graduate student worked hard and discovered how to eliminate the expected background without much reduction in signal efficiency, the ‘reward’ would be a weaker upper limit.<sup>13</sup> An answer is that although the actual limit had increased, the sensitivity of the experiment with the smaller background was better. There are other situations—for example, variants of the random choice of voltmeter (compare ref. [55])—where a measurement with better sensitivity can on occasion give a less precise result.

- In the Bayesian approach, the dependence of the limit on  $b$  is **weaker**. Indeed when  $n = 0$ , the limit does not depend on  $b$ .
- Sen et al. [56] consider a related problem, of a physical non-negative parameter  $\lambda$  producing a measurement  $x$ , which is distributed about  $\lambda$  as a Gaussian of variance  $\sigma^2$ . As the observable  $x$  becomes more and more negative, the upper limit on  $\lambda$  **increases**, because it is deduced that  $\sigma$  must in fact be larger than its quoted value.

In trying to assess which of the methods is best, one first needs a list of desirable properties. These include:

- Coverage: Even though coverage is a frequentist concept, most Bayesian particle physicists would like the coverage of their intervals to match their reported credibility, at least approximately. Because the data in counting experiments is discrete, it is impossible in any sensible way to achieve exact coverage for all  $\mu$  (see Sect. 15.1.3.4). However, it is not completely obvious that even Frequentists need coverage for every possible value of  $\mu$ , since different experiments will have different values of  $b$  and of  $\epsilon$ . Thus even for a constant value of the physical parameter  $s$ , different experiments will have different  $\mu = \epsilon * s + b$ . Thus it would appear that, if coverage in some average (over  $\mu$ ) sense were satisfactory, the frequentist requirement for intervals to contain the true value at the requisite rate would be maintained. This, however, is not the generally accepted view by particle physicists, who would like not to undercover for **any**  $\mu$ .
- Not too much overcoverage: Because coverage varies with  $\mu$ , for methods that aim not to undercover anywhere, some overcoverage is inevitable. This corresponds to having some upper limits which are high, and this leads to undesirable loss of power in rejecting alternative hypotheses about the parameter’s value.

---

<sup>13</sup>The  $n = 0$  situation is perhaps a special case, as the number of observed events cannot decrease as further selections are imposed to reduce the expected background. For non-zero observed events, if  $n$  decreases with the tighter cuts (as expected for reduced background), the upper limit is likely to go down, in agreement with intuition. But if  $n$  stays constant, that could be because the observed events contain signal, so it is perhaps not surprising that the upper limit increases.

- Short and empty intervals: These can be obtained for certain values of the observable, without resulting in undercoverage. They are generally regarded as undesirable for the reasons explained above.

It is not obvious how to incorporate the above desiderata on interval length into an algorithm that would be useful for choosing among different methods for setting limits. For different experiments studying the same phenomena (e.g. Dark Matter searches, neutrino oscillation experiments, etc.) it is worthwhile to use the same technique for calculating allowed parameter ranges.

### 15.4.2 Two-Sided Intervals

An alternative to giving upper limits is to quote two-sided intervals. For example, a 68% confidence interval for the mass of the top quark might be  $172.6\text{--}173.4\text{ GeV}/c^2$ , as opposed to its 95% upper limit being  $173.6\text{ GeV}/c^2$ . Most of the difficulties and ambiguities mentioned above apply in this case too, together with some extra possibilities. Thus, while it is clear which of two possible upper limits is tighter, this is not necessarily so for two-sided intervals, where which is shorter may be metric dependent; the first of two intervals for a particle's lifetime  $\tau$  may be shorter, but the second may be shorter when the ranges are quoted for its decay rate ( $= 1/\tau$ ). There is also scope for choice of ordering rule for the frequentist Neyman construction, or for choosing the interval from the Bayesian posterior probability density.<sup>14</sup>


### 15.4.3 Feldman-Cousins Approach

Feldman and Cousins' fully frequentist approach[54] exploits the freedom available in the Neyman construction of how to choose an interval in the data that contains a given fraction  $\alpha$  of the probability, by using their 'ordering rule'. This is based on the likelihood ratio  $L(x, \mu)/L(x, \mu_{best})$ , where  $\mu_{best}$  is the physically-allowed value of  $\mu$  which gives the largest value of  $L$  for that particular  $x$ . For values of  $\mu$  far from a physical boundary, this makes little difference from the standard central Neyman construction, but near a boundary the region is altered in such a way as to make it unlikely that there will be zero-length or empty intervals for the parameter  $\mu$ ; these can occur in the standard Neyman construction (see Fig. 15.4).

---

<sup>14</sup>A Bayesian statistician would be happy with the posterior as the final result. Particle physicists like to quote an interval as a convenient summary. For a parameter that cannot be negative and for which the exclusion of zero is interesting (e.g. testing whether the production rate of some hypothesised particle is non-zero), an upper limit would always include zero, a lower limit or a central interval would exclude it and a maximum probability density one would not be invariant with respect to changes in the functional form of the parameter.





**Fig. 15.4** The Feldman-Cousins 90% confidence band (solid curves) for the mean  $\mu$  of a Gaussian probability density function of unit variance for a measurement  $x$ . The straight dashed lines show the confidence band for the central Neyman construction. The Feldman-Cousins ordering rule pulls the interval to the left at small  $\mu$ , and hence, even for negative observed  $x$ , the  $\mu$  interval is not empty, as happens for central frequentist intervals when  $x$  is below  $-1.6$

The original Feldman-Cousins paper also considered how to extend their method when there is more than one parameter and one measurement. They describe an idealised neutrino oscillation experiment with the data being the energy spectrum of the interacting neutrinos, and the parameters are  $\sin^2(2\theta)$  and  $\Delta m^2$  (see Eq. 15.15). A practical problem of having many parameters is the CPU time required to compute the results.

Feldman and Cousins also point out that an apparently innocuous procedure for choosing what result to quote may lead to undercoverage. Many physicists would quote an upper limit on any possible signal if their observation was less than 3 standard deviations above the expected background, but a two-sided interval if their result was above this. With each type of interval constructed to give 90% coverage, there are some values of the parameter for which the coverage for this mixed procedure drops to 85%; Feldman and Cousins refer to this as ‘flip-flop’. Their ‘unified’ approach circumvents this problem, as it automatically yields upper limits for small values of the data, but two-sided intervals for larger measurements, while avoiding undercoverage for all possible true values of the signal.

### 15.4.4 Sensitivity

It is useful to quote the sensitivity of a procedure, as well as the actual upper limit as derived from the observed data.<sup>15</sup> For upper limits or for uncertainties on measurements, this can be defined as the median value that would be obtained if the procedure was repeated a large number of times.<sup>16</sup> Using the median is preferable to the mean because (a) it is metric independent (i.e. the median lifetime upper limit would be the reciprocal of the median decay rate lower limit); and (b) it is much less sensitive to a few anomalously large upper limits or uncertainty estimates.

It is common to present not only the median of the expected distribution, but also values corresponding to 16th and 84th percentiles (commonly referred to as  $\pm 1\sigma$ ) and also the 2.5% and 97.5% ones ( $\pm 2\sigma$ ). This enables a check to be made that the observed result is reasonable.

Punzi [57] has drawn attention to the fact that this choice of definition for sensitivity has some undesirable features. Thus designing an analysis procedure to minimise the median upper limit for a search in the absence of a signal provides a different optimisation from maximising the median number of standard deviations for the significance of a discovery when the signal is present. Also there is only a 50% chance of achieving the median result or better. Instead, for pre-defined levels  $\alpha$  and confidence level  $CL$ , Punzi determines at what signal strength there is a probability of at least  $CL$  for establishing a discovery at a significance level  $\alpha$ . This is what he quotes as the sensitivity, and is the signal strength at which we are sure to be able either to claim a discovery or to exclude its existence. Below this, the presence or otherwise of a signal makes too little difference, and we may remain uncertain (see Fig. 15.5).

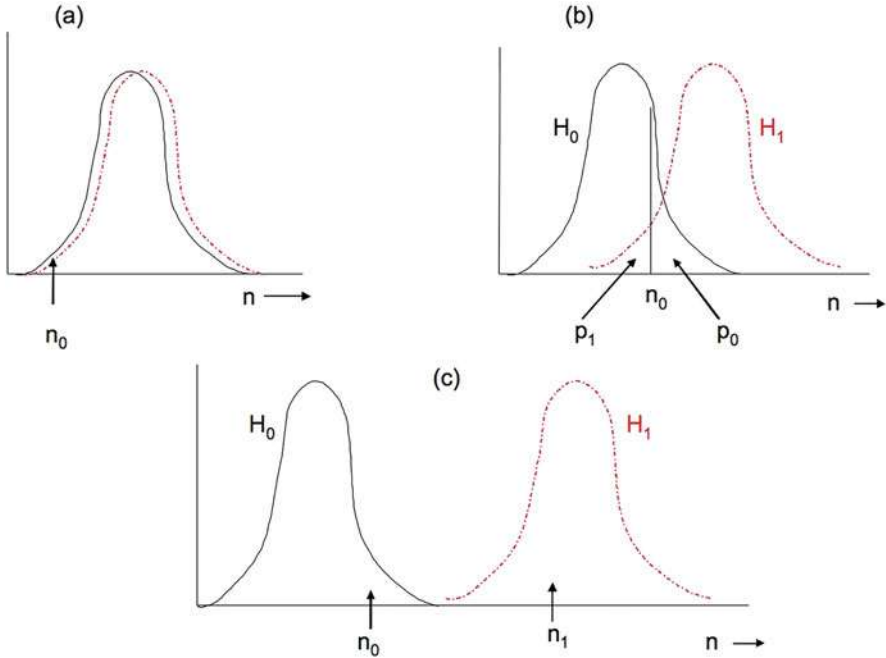
### 15.4.5 Nuisance Parameters

For calculating upper limits in the simple counting experiment described in Sect. 15.4.1, the nuisance parameters arise from the uncertainties in the background rate  $b$  and the acceptance  $\epsilon$ . These uncertainties are usually quoted as  $\sigma_b$  and  $\sigma_\epsilon$  (e.g.  $b = 3.1 \pm 0.5$ ), and the question arises of what these uncertainties mean. Sometimes they encapsulate the results of a subsidiary measurement, performed to estimate  $b$  or  $\epsilon$ , and then they would express the width of the Bayesian posterior or of the frequentist interval obtained for the nuisance parameters. However, in

---

<sup>15</sup>The sensitivity on its own will not do, because it is independent of the data.

<sup>16</sup>Instead of using a large number of simulations in order to extract the median, sometimes the ‘Asimov’ data set is used. This is the single data set that would be obtained if statistical fluctuations were suppressed. i.e. if a model predicted 11.3 events in a particular bin, the Asimov data set for that model would contain 11.3 events in that bin. The Asimov data set and the median of the toys usually but not always produce similar results.



**Fig. 15.5** Punzi definition of sensitivity. Expected distributions for a statistic  $t$  (which in simple cases could be simply the observed number of events  $n$ ), for  $H_0$  = background only (solid curves) and for  $H_1$  = background plus signal (dashed curves). In (a), the signal strength is very weak, and it is impossible to choose between  $H_0$  and  $H_1$ . As shown in (b), which is for moderate signal strength,  $p_0$  is the probability according to  $H_0$  of  $t$  being equal to or larger than the observed  $t_0$ . To claim a discovery,  $p_0$  should be smaller than some pre-set level  $\alpha$ , usually taken to correspond to  $5\sigma$ ;  $t_{crit}$  is the minimum value of  $t$  for this to be so. Similarly  $p_1$  is the probability according to  $H_1$  for  $t \leq t_0$ . The power function is the probability according to the alternative hypothesis that  $t$  will exceed  $t_{crit}$ . As the separation of the  $H_0$  and  $H_1$  pdfs increases, so does the power. According to Punzi, the sensitivity should be defined as the expected production strength of the signal such that the power exceeds another predefined CL, e.g. 95%. The exclusion region corresponds to  $t_0$  in the 5% lower tail of  $H_1$ , while the discovery region has  $t_0$  in the 5% upper tail of  $H_0$ ; in (b) there is a “No decision” region in between, as the signal strength is below the sensitivity value. The sensitivity is thus the signal strength above which there is a 95% chance of making a  $5\sigma$  discovery. i.e. The distributions for  $H_0$  and  $H_1$  are sufficiently separated that, apart possibly for the  $5\sigma$  upper tail of  $H_0$  and the 5% lower tail of  $H_1$ , they do not overlap. In (c) the signal strength is so large that there is no ambiguity in choosing between the hypotheses

many situations, the uncertainties may involve Monte Carlo simulations, which have systematic uncertainties (e.g. related to how well the simulation describes the real data) as well as statistical ones; or they may reflect uncertainties or ambiguities in theoretical calculations required to derive  $b$  and/or  $\epsilon$ . In the absence of further information the posterior is often assumed to be a Gaussian, usually truncated so as to exclude unphysical (e.g. negative) values. This may be at best only approximately

true, and deviations are likely to be most serious in the tails of the distribution. A log-normal or gamma function may be a better choice.

There are many methods for incorporating nuisance parameters in upper limit calculations. These include:

- Profile likelihood (see also Sect. 15.2.3)

The likelihood, based on the data from the main and from the subsidiary measurements, is a function of the parameter of interest  $s$  and of the nuisance parameters. The profile likelihood  $L_{prof}(s)$  is simply the full likelihood  $L(s, b_{best}(s), \epsilon_{best}(s))$ , evaluated at the values of the nuisance parameters that maximise the likelihood at each  $s$ . Then the profile likelihood is simply used to extract the limits on  $s$ , much as the ordinary likelihood could be used for the case when there are no nuisance parameters.

Rolke et al. [59] have studied the behaviour of the profile likelihood method for limits. Heinrich[32] had shown that the likelihood approach for estimating a Poisson parameter (in the absence of both background and of nuisance parameters) can have poor coverage at low values of the Poisson parameter. However, the profile likelihood seems to do better, probably because the nuisance parameters have the effect of smoothing away the fluctuating coverage observed by Heinrich.

- Fully Bayesian

When there is a subsidiary measurement for a nuisance parameter, a prior is chosen for  $b$  (or  $\epsilon$ ), the data are used to extract the likelihood, and then Bayes' Theorem is used to deduce the posterior for the nuisance parameter. This posterior from the subsidiary measurement is then used as the prior for the nuisance parameter in the main measurement (this prior could alternatively come from information other than a subsidiary measurement); with the prior for  $s$  and the likelihood for the main measurement, the overall joint posterior for  $s$  and the nuisance parameter(s) is derived.<sup>17</sup> This is then integrated over the nuisance parameter(s) to determine the posterior for  $s$ , from which an upper limit can be derived; this procedure is known as marginalisation.

Numerical examples of upper limits can be found in ref. [60], where a method is discussed in detail. Thus assuming (somewhat unrealistically) precisely determined backgrounds, the effect of a 10% uncertainty in  $\epsilon$  can be seen for various measured values of  $n$  in Table 15.1. A plot of the coverage when the uncertainty in  $\epsilon$  is 20% is reproduced in Fig. 15.6.

It is not universally appreciated that the choice for the main measurement of a truncated Gaussian prior for  $\epsilon$  and an (improper) constant prior for non-negative  $s$  results in a posterior for  $s$  which diverges[61]. Thus numerical estimates of the relevant integrals are meaningless. Another problem comes from the difficulty of choosing sensible multi-dimensional priors. Heinrich has pointed out the


---

<sup>17</sup>This is usually equivalent to starting with a prior for  $s$  and the nuisance parameters, and the likelihood for the data from the main and the subsidiary experiments together, to obtain the joint posterior.

**Table 15.1** Bayesian 90% confidence level upper limits for the production rate  $s$  as a function of  $n$ , the observed number of events

$n$	$b = 0.0$	$b = 3.0$
0	2.35 (2.30)	2.35 (2.30)
3	6.87 (6.68)	4.46 (4.36)
6	10.88 (10.53)	7.80 (7.60)
9	14.71 (14.21)	11.56 (11.21)
20	28.27 (27.05)	25.05 (24.05)

The Poisson parameter  $\mu = \epsilon * s + b$ , where the expected background  $b$  is either 0.0 or 3.0, and is precisely known; and  $\epsilon$ , whose true values is 1.0, is estimated in a subsidiary measurement with 10% accuracy. The numbers in brackets are the corresponding upper limits when  $\epsilon$  is known precisely. At large  $n$ , the limits for  $b = 3.0$  are 3 units lower than those for  $b = 0.0$ ; the latter are approximately  $n + 1.28\sqrt{n}$  at large  $n$ . The effect of the uncertainty in  $\epsilon$  is to increase the limits, and by a larger amount at large  $n$ . For  $n = 0$ , these Bayesian limits are independent of the expected background  $b$



**Fig. 15.6** The coverage  $C$  for the 90% confidence level upper limit as a function of the true parameter  $s_{true}$ , as obtained in a Bayesian approach. The background  $b = 3.0$  is assumed to be known exactly, while the subsidiary measurement for  $\epsilon$  gives a 20% accuracy. The discontinuities are a result of the discrete (integer) nature of the measurements. There is no undercoverage

- problems that can arise for the above Poisson counting experiment, when it is extended to deal with several data channels simultaneously[62].
- Fully frequentist
- In principle, the fully frequentist approach to setting limits when provided with data from the main and from subsidiary measurements is straightforward: the Neyman construction is performed in the multidimensional space where the parameters are  $s$  and the nuisance parameters, and the data are from all the relevant measurements. Then the region in parameter space for which the observed data was likely is projected onto the  $s$ -axis, to obtain the confidence region for  $s$ .

In practice there are formidable difficulties in writing a program to do this in a reasonable amount of time. Another problem is that, unless a clever ordering rule is used for producing the acceptance region in data space for fixed values of the parameters, the projection phase leads to overcoverage, which can become larger as the number of nuisance parameters increases. Good ordering rules have been found for a version of the Poisson counting experiment[63], and also for the ratio of Poisson means[64], where the confidence intervals are tighter than those obtained by conditioning on the sum of the numbers of counts in the two observations.

For the fully frequentist method, it is guaranteed that there will be no undercoverage for any combination of parameter true values. This is not so for any other method, and so most particle physicists would like assurance that the technique used does indeed provide reasonable coverage, at least for  $s$ . There is usually lively debate between frequentists and Bayesians as to whether coverage is desirable for all values of the nuisance parameter(s), or whether one should be happy with no or little undercoverage when experiments are averaged, for example, over the nuisance parameter true values.

- **Mixed**

Because of the difficulty of performing a fully frequentist analysis in all but the simplest problems, an alternative approach[65] is to use Bayesian averaging over the nuisance parameters, but then to employ a frequentist approach for  $s$ . The hope is that for most experiments setting upper limits, the statistical uncertainties on the low  $n$  data are relatively large and so, provided the uncertainties in the nuisance parameters are not too large, the effect of the systematics on the upper limits will not be too dramatic, and an approximate method of dealing with them may be reasonable.

Although such an approach cannot be justified from fundamentals, it provides a practical method whose properties can be checked, and is often satisfactory.

#### **15.4.6 Banff Challenges**

Given the large number of techniques available for extracting upper limits from data, especially in the presence of nuisance parameters, it was decided at the Banff meeting[6] that it would be useful to compare the properties of the different approaches under comparable conditions. This led to the setting up of the ‘Banff Challenge’, which consisted of providing common data sets for anyone to calculate their upper limits. This was organised by Joel Heinrich, who reported on the performance of the various methods at the PHYSTAT-LHC meeting[66].

At the second Banff meeting[8], the challenge was set by Tom Junk and consisted of participants trying to distinguish between histograms, some of which contained only background and others which contained a background and signal, which appeared as a peak (compare Sect. 15.7.5)

### 15.4.7 Recommendations

It would be incorrect to say that there is one method that must be used. Many Particle Physicists' ideal would be to use a frequentist approach if viable software were available for problems with several parameters and items of data. Otherwise they would be prepared to settle for a Bayesian approach, with studies of the sensitivity of the upper limit to the choice of priors, and of the coverage; or for a profile likelihood method, again with coverage studies. What is important is that the procedure should be fully defined before the data are analysed; and that when the experimental result and the sensitivity of the search are reported, the method used should be fully explained.

The CDF Statistics Committee [67] also suggests that it is useful to use a technique that has been employed by other experiments studying the same phenomenon; this makes for easier comparison. They tend to favour a Bayesian approach, chiefly because of the ease of incorporating nuisance parameters.

## 15.5 Combining Results

This section deals with the combination of the results from two or more measurements of a single (or several) parameters of interest. It is not possible to combine upper limits (UL). This is because an 84% UL of 1.5 could come from a measurement of  $1.4 \pm 0.1$ , or  $0.5 \pm 1.0$ ; these would give very different results when combined with some other measurement.

The combination of  $p$ -values is discussed in Sect. 15.7.9.

### 15.5.1 Single Parameter

An interesting question is whether it is possible to combine two measurements of a single quantity, each with uncertainty  $\pm 10$ , such that the uncertainty on the combined best estimate is  $\pm 1$ ? The answer can be deduced later.

To combine  $N$  different uncorrelated measurements  $a_i \pm \sigma_i$  of the same physical quantity  $a$ <sup>18</sup> when the measurements are believed to be Gaussian distributed about the true value  $a_{true}$ , the well-known result is that the best estimate  $a_{comb} \pm \sigma_{comb}$  is given by

$$a_{comb} = \Sigma(a_i * w_i) / \Sigma w_i, \quad \sigma_{comb} = 1 / \sqrt{\Sigma w_i}, \quad (15.13)$$

where the weights are defined as  $w_i = 1/\sigma_i^2$ . This is readily derived from minimising with respect to  $a$  a weighted sum of squared deviations

$$S(a) = \Sigma(a_i - a)^2 / \sigma_i^2 \quad (15.14)$$

The extension to the case where the individual measurements are correlated (as is often the case for analyses using different techniques on the same data) is straightforward:  $S(a)$  becomes  $\Sigma \Sigma (a_i - a) * H_{ij} * (a_j - a)$ , where  $H$  is the inverse covariance matrix for the  $a_i$ . It provides **Best Linear Unbiased Estimates** (BLUE)[70].

There are, however, practical details that complicate its application. For example, in the above formula, the  $\sigma_i$  are supposed to be the **true** accuracies of the measurements. Often, all that we have available are **estimates** of their values. Problems arise in situations where the uncertainty estimate depends on the measured value  $a_i$ . For example, in counting experiments with Poisson statistics, it is typical to set the uncertainty as the square root of the observed number. Then a downward fluctuation in the observation results in an overestimated weight, and  $a_{comb}$  is biased downwards. If instead the uncertainty is estimated as the square root of the expected number  $a$ , the combined result is biased upwards—the increased uncertainty reduces  $S$  at larger  $a$ . A way round this difficulty has been suggested by Lyons et al. [71]. Alternatively, for Poisson counting data a likelihood approach is preferable to a  $\chi^2$ -based method.

Another problem arises when the individual measurements are very correlated. When the correlation coefficient of two uncertainties is larger than  $\sigma_1/\sigma_2$  (where  $\sigma_1$  is the smaller uncertainty),  $a_{comb}$  lies outside the range of the two measurements. As the correlation coefficient tends to +1, the extrapolation becomes larger, and is sensitive to the exact values assumed for the elements of the covariance matrix. The situation is aggravated by the fact that  $\sigma_{comb}$  tends to zero. This is usually dealt with by selecting one of the two analyses, rather than trying to combine them. However, if the estimated uncertainty increases with the estimated value, choosing the result with the smaller **estimated** uncertainty can again produce a downward bias. On the other hand, using the smaller **expected** uncertainty can cause us to ignore an analysis which had a particularly favourable statistical fluctuation, which produced a result

---

<sup>18</sup>It is of course much better to use all the **data** in a combined analysis, rather than simply to combine the **results**.



that was genuinely more precise than expected<sup>19</sup> How to deal with this situation in general is an open question. It has features in common with the problem (inspired by ref. [55]) of measuring a voltage by choosing at random a voltmeter from a cupboard containing meters of different sensitivities.

Another example involves combining two measurements of a cross-section with small statistical uncertainties, but with large correlated uncertainties from the common luminosity. With this luminosity uncertainty included in the covariance matrix, BLUE can result in the combined value being outside the range of the individual measurements. For this situation, it is preferable to exclude the luminosity uncertainty from the covariance matrix, and to apply it to the combined result afterwards.

### 15.5.2 Two or More Parameters

An extension of this procedure is for combining  $N$  pairs of correlated measurements (e.g. the gradient and intercept of a straight line fit to several sets of data, where for simplicity it is assumed that any pair is independent of every other pair). For several pairs of values  $(a_i, b_i)$  with inverse covariance matrices  $\mathbf{M}_i$ , the best combined values  $(a_{comb}, b_{comb})$  have as their inverse covariance matrix  $\mathbf{M} = \Sigma \mathbf{M}_i$ . This means that, if the covariance matrix correlation coefficients  $\rho_i$  of the different measurements are very different from each other, the uncertainty on  $a_{comb}$  can be much smaller than that for any single measurement.

This situation applies for track fitting to hits in a series of groups of tracking chambers, where each set of close chambers provides a very poor determination of the track; but the combination involves widely spaced chambers and determines the track well. Using the profile likelihoods (e.g. for the intercept, profiled over the gradient) for combining different measurements loses the correlation information and can lead to a very poor combined estimate[37]. The alternative of ignoring the correlation information is also strongly discouraged.

The importance of retaining covariances is relevant for many combinations, e.g. for the determination of the amount of Dark Energy in the Universe from various cosmological data[73].

---

<sup>19</sup>For example, the ALEPH experiment at LEP produced a tighter-than-expected upper limit on the mass of  $\nu_\tau$  because they happened to observe  $\tau$  decay configurations which were particularly sensitive to the  $\nu_\tau$  mass.

### 15.5.3 Data Consistency

The standard procedure for combining data pays no attention to whether or not the data are consistent. If they are clearly inconsistent, then they should not all be combined. When they are somewhat inconsistent, the procedure adopted by the Particle Data Group[14] is to increase all the uncertainties by a common factor such that the overall  $\chi^2$  per degree of freedom equals unity.<sup>20</sup>

The Particle Data Group prescription for expanding uncertainties in the case of discrepant data sets has complications when each of the data sets consists of two or more parameters[72].

## 15.6 Goodness of Fit

### 15.6.1 Sparse Multi-Dimensional Data

The standard method loved by most scientists uses the weighted sum of squares, commonly called  $\chi^2$ . This, however, is only applicable to binned data (i.e. in a one or more dimensional histogram). Furthermore it loses its attractive feature that its distribution is model-independent when there is not enough data, which is likely to be so in the multi-dimensional case.

Although the maximum likelihood method is very useful for parameter determination with **unbinned data**, the value of  $L_{max}$  usually does not provide a measure of goodness of fit (see Sect. 15.2.2).

An alternative that is used for sparse one-dimensional data is the Kolmogorov-Smirnov (KS) approach[68], or one of its variants. However, in the presence of fitted parameters, simulation is again required to determine the expected distribution of the KS-distance. Also because of the problem of how to order the data, the way to use it in multi-dimensional situations is not unique.

The standard KS method uses the maximum deviation between two cumulative distributions; because of statistical fluctuations, this is likely to occur near the middle of the distributions. In cases where interesting New Physics is expected to occur at extreme values of some kinematic variable (e.g.  $p_T$ ), variants of KS such as Anderson-Darling[69] that give extra weight to the distributions' tails may be more useful.

---

<sup>20</sup>This is somewhat conservative, in that even if there are no problems, about half the data sets would be expected to have this larger than unity.

## 15.6.2 Number of Degrees of Freedom

If we construct the weighted sum of squares  $S$  between a predicted theoretical curve and some data in the form of a histogram, provided the Poisson distribution of the bin contents can be approximated by a Gaussian (and the theory is correct, the data are unbiased, the uncertainty estimates are correct, etc.), **asymptotically**<sup>21</sup>  $S$  will be distributed as  $\chi^2$  with the number of degrees of freedom  $\nu = n - f$ , where  $n$  is the number of data points and  $f$  is the number of free parameters whose values are determined by minimising  $S$ .

The relevance of the asymptotic requirement can be seen by imagining fitting a more or less flat distribution by the expression  $N(1 + 10^{-6} \cos(x - x_0))$ , where the free parameters are the normalisation  $N$  and the phase  $x_0$ . It is clear that, although  $x_0$  is left free in the fit, because of the  $10^{-6}$  factor, it will have a negligible effect on the fitted curve, and hence will not result in the typical reduction in  $S$  associated with having an extra free parameter. Of course, with an enormous amount of data, we would have sensitivity to  $x_0$ , and so asymptotically it does reduce  $\nu$  by one unit, but not for smaller amounts of data.

Another example involves neutrino oscillation experiments[54]. In a simplified two neutrino scenario, the neutrino energy spectrum is fitted by a survival probability  $P$  of the form

$$P = 1 - \sin^2 2\theta \sin^2(C * \Delta m^2), \quad (15.15)$$

where  $C$  is a known function of the neutrino energy and the length of its flight path,  $\Delta m^2$  is the difference in mass squared of the relevant neutrino species, and  $\theta$  is the neutrino mixing angle. For small values of  $C * \Delta m^2$ , this reduces to

$$P \approx 1 - \sin^2 2\theta (C * \Delta m^2)^2 \quad (15.16)$$

Thus the survival probability depends on the two parameters only via their product  $\sin 2\theta \Delta m^2$ . Because this combination is all that we can hope to determine, we effectively have only one free parameter rather than two. Of course, an enormous amount of data can manage to distinguish between  $\sin(C * \Delta m^2)$  and  $C * \Delta m^2$ , and so asymptotically we have two free parameters as expected.

## 15.7 Discovery Issues

Searches for new particles are an exciting endeavour, and continue to play a large role at the LHC at CERN, in neutrino experiments, in searches for dark matter, etc. The 2007 and 2011 PHYSTAT Workshops at CERN[7, 9] were devoted specifically

---

<sup>21</sup>The examples in this section go beyond the requirement that we need enough events for the Poisson distribution to be well approximated by a Gaussian.

to statistical issues that arise in discovery-orientated analyses at the LHC. Ref [74] deals with statistical issues that occur in Particle Physics searches for new phenomena; as an example, it includes the successful search for the Higgs boson at the LHC. A more detailed description of the plans for the Higgs search before its discovery is in ref. [75].

### 15.7.1 $H_0$ , or $H_0$ Versus $H_1$ ?

In looking for new physics, there are two distinct types of approach. We can compare our data just with the null hypothesis  $H_0$ , the SM of Particle Physics; alternatively we can see whether our data are more consistent with  $H_0$  or with an alternative hypothesis  $H_1$ , some specific manifestation of new physics, such as a particular form of quark and/or lepton substructure. The former is known as ‘goodness of fit’, while the term ‘hypothesis testing’ is often reserved for the latter.

Each of these approaches has its own advantage. By not specifying a specific alternative,<sup>22</sup> the goodness of fit test may be capable of detecting any form of deviation from the SM. On the other hand, if we are searching for some specific new effect, a comparison of  $H_0$  and  $H_1$  is likely to be a more sensitive way for that particular alternative. Also, the ‘hypothesis testing’ approach is less likely to give a false discovery claim if the assumed form of  $H_0$  has been slightly mis-modelled.

### 15.7.2 $p$ -Values

In order to quantify the chance of the observed effect being due to an uninteresting statistical fluctuation, some statistic is chosen for the data. The simplest case would be the observed number  $n_0$  of interesting events. Then the  $p$ -value is calculated, which is simply the probability that, given the expected background rate  $b$  from known sources, the observed value would fluctuate up to  $n_0$  or larger. In more complicated examples involving several relevant observables, the data statistic may be a likelihood ratio  $L_0/L_1$  for the likelihood of the null hypothesis  $H_0$  compared with that for a specific alternative  $H_1$ .

To compute the  $p$ -value of the observed or of possible data, the distribution  $f(t)$  of the data statistic  $t$  under the relevant hypothesis is required. In some cases this can be obtained analytically, but in more complicated situations,  $f(t)$  may require simulation. For  $t$  being  $-2 \ln L_0/L_{best}$ , Cowan et al have given useful asymptotic

---

<sup>22</sup>Even a test of the null hypothesis may not be completely independent of ideas about alternatives. Thus in an event counting experiment, new physics usually results in an **increase** in rate, unless we are looking for neutrino oscillations, in which case a **decrease** would be significant. Also, sometimes the statistic used for a goodness of fit test of  $H_0$  may be the likelihood ratio for  $H_0$  as compared with a specific alternative  $H_1$ .

formulae for  $f(t)$ [76]; here  $L_{best}$  is the value of the likelihood when the parameters in  $H_0$  are set at their best values.

A small value of  $p$  indicates that the data are not very compatible with the theory (which may be because the detector's response or the background is poorly modeled, rather than the theory being wrong).

Particle Physicists usually convert  $p$  into the number of standard deviations  $\sigma$  of a Gaussian distribution, beyond which the one-sided tail area corresponds to  $p$ ; statisticians refer to this as the  $z$ -score, but physicists call it significance. Thus  $5\sigma$  corresponds to a  $p$ -value of  $3 * 10^{-7}$ . This is done simply because it provides a number which is easier to remember, and not because Gaussians are relevant for every situation.

Unfortunately,  $p$ -values are often misinterpreted as the probability of the theory being true, given the data. It sometimes helps colleagues clarify the difference between  $p(A|B)$  and  $p(B|A)$  by reminding them that the probability of being pregnant, given the fact that you are female, is considerably smaller than the probability of being female, given the fact that you are pregnant. Reference [77] contains a series of articles by statisticians on the use (and misuse) of  $p$ -values.

Sometimes  $S/\sqrt{B}$  or  $S/\sqrt{S+B}$  or the like (where  $S$  is the number of observed events above the estimated background  $B$ ) is used as an approximate measure of significance. These approximations can be very poor, and their use is in general not recommended.<sup>23</sup>

### 15.7.3 $CL_s$

This is a technique[58] which is used for situations in which a discovery is not made, and instead various parameter values are excluded. For example the failure to observe SUSY particles can be converted into mass ranges which are excluded (at some confidence level).

Figure 15.5 (again) illustrates the expected distributions for some suitably chosen statistic  $t$  under two different hypotheses: the null  $H_0$  in which there is only standard known physics, and  $H_1$  which also includes some specific new particle, such as a SUSY neutralino. In Fig. 15.5c, the new particle is produced prolifically, and an experimental observation of  $t$  should fall in one peak or the other, and easily distinguishes between the two hypotheses. In contrast, Fig. 15.5a corresponds to very weak production of the new particle and it is almost impossible to know whether the new particle is being produced or not.

<sup>23</sup>For example, if selections to enhance signal with respect to background were optimised using  $S/\sqrt{B}$ , extremely hard cuts might be chosen, yielding expected numbers of events  $S = 0.1$  and  $B = 10^{-3}$ . This results in  $S/\sqrt{B} = 10$ , which sounds very good, but in fact this selection is disastrous.

The conventional method of claiming new particle production would be if the observed  $t$  fell well above the main peak of the  $H_0$  distribution; typically a  $p_0$  value corresponding to  $5\sigma$  would be required (see Sect. 15.7.7). In a similar way, new particle production would be excluded if  $t$  were below the main part of the  $H_1$  distribution. Typically a 95% exclusion region would be chosen (i.e.  $p_1 \leq 0.05$ ), where  $p_1$  is by convention the left-hand tail of the  $H_1$  distribution, as shown in Fig. 15.5b.

The  $CL_s$  method aims to provide protection against a downward fluctuation of  $t$  in Fig. 15.5a resulting in a claim of exclusion in a situation where the experiment has no sensitivity to the production of the new particle; this could happen in 5% of experiments. It achieves this by defining<sup>24</sup>

$$CL_s = p_1 / (1 - p_0), \quad (15.17)$$


and requiring  $CL_s$  to be below 0.05. From its definition, it is clear that  $CL_s$  cannot be smaller than  $p_1$ , and hence is a conservative version of the frequentist quantity  $p_1$ . It tends to  $p_1$  when  $t$  lies above the  $H_0$  distribution, and to unity when the  $H_0$  and  $H_1$  distributions are very similar. The reduced  $CL_s$  exclusion region is shown by the dotted diagonal line in Fig. 15.7; the price to pay for the protection provided by  $CL_s$  is that there is built-in conservatism when  $p_1$  is small but  $p_0$  has intermediate values i.e. there are more cases in which no decision is made. Most statisticians are appalled by the use of  $CL_s$ , because they consider that it is meaningless to take the ratio of two  $p$ -values.

It is deemed not to be necessary to protect against statistical fluctuations giving rise to discovery claims in situations with no sensitivity, because that should happen only at the  $3 \times 10^{-7}$  rate (the one-sided  $5\sigma$  Gaussian tail area).

Figure 15.7 is also useful for understanding the Punzi sensitivity definition (see Sect. 15.4.4). For any specified distributions of the statistic  $t$  for  $H_0$  and  $H_1$ , the possible  $(p_0, p_1)$  values lie on a curve or straight line which extends from (0,1) to (1,0). With more data, the  $t$  distributions separate, and the curve moves closer to the  $p_0$  and  $p_1$  axes. The amount of data required to satisfy the Punzi requirement of always claiming a discovery or an exclusion is when no part of the curve is in the “no decision” region of Fig. 15.7.

<sup>24</sup>Given the fact that  $CL_s$  is the ratio of two  $p$ -values, the choice of symbol  $CL_s$  (standing for ‘confidence level of signal’) is not optimal. Another source of confusion is that in definitions of  $CL_s$  the ways the  $p$ -values are defined vary, so the formulae can look different but the underlying concept is the same.

A subtlety with Eq. (15.17) is that  $p_0$  there is the probability of obtaining a measurement **greater than** the observed one, rather than the usual ‘greater than or equal to’. This is to make  $1 - p_0$  the probability of a value smaller than or equal to the observed one, in analogy with the definition of  $p_1$ . It makes a difference when the observation is a small discrete number.



**Fig. 15.7** Plot of  $p_0$  against  $p_1$  for comparing a data statistic  $t$  with two hypotheses  $H_0$  and  $H_1$ , whose expected  $pdf$ 's for  $t$  are given by two Gaussians of peak separation  $\Delta\mu$ , and of equal width  $\sigma$ . For a given pair of  $pdf$ 's for  $t$ , the allowed values of  $(p_0, p_1)$  lie on a curve or straight line (shown solid in the diagram). The expected density for the data along a curve is such that its projection along the  $p_0$ -axis (or  $p_1$ -axis) is expected to be uniform for the hypothesis  $H_0$  (or  $H_1$  respectively). As the separation increases, the curves approach the  $p_0$  and  $p_1$  axes. Rejection of  $H_0$  is for  $p_0$  less than, say,  $3 \times 10^{-7}$ ; here it is shown as 0.05 for ease of visualisation. Similarly exclusion of  $H_1$  is shown as  $p_1 < 0.1$ . Thus the  $(p_0, p_1)$  square is divided into four regions: the largest rectangle is when there is no decision, the long one above the  $p_0$ -axis is for exclusion of  $H_1$ , the high one beside the  $p_1$ -axis is for rejection of  $H_0$ , and the smallest rectangle is when the data lie between the two  $pdf$ 's. For  $\Delta\mu/\sigma = 3.33$ , there are no values of  $(p_0, p_1)$  in the "no decision" region. In the  $CL_s$  procedure, rejection of  $H_1$  is when the  $t$  statistic is such that  $(p_0, p_1)$  lies below the diagonal dotted straight line

### 15.7.4 Comparing Two Hypotheses Via $\chi^2$

Assume that there is a histogram with 100 bins, and that a  $\chi^2$  method is being used for fitting it with a function with one free parameter. The expected value of  $\chi^2$  is  $99 \pm 14$ . Thus if  $p_0$ , the best value of the parameter, yields a  $\chi^2$  of 85, this would be regarded as very satisfactory. However, a theoretical colleague has a model which predicts that the parameter should have a different value  $p_1$ , and wants to know what the data have to say about that. This is tested by calculating the  $\chi^2$  for that  $p_1$ , which yields a value of 110. There appear to be two contradictory conclusions:

- $p_1$  is satisfactory: This is based on the fact that the relevant  $\chi^2$  of 110 is well within the expected range of  $99 \pm 14$ .

- $p_1$  is ruled out: The uncertainty on  $p$  is estimated by seeing how much it must change from its optimum value in order to make  $\chi^2$  increase by 1. For this data,  $\chi^2(p_1)$  is 25 units larger than  $\chi^2(p_0)$ , and so, assuming that the behaviour of  $\chi^2$  in the neighbourhood of the minimum is parabolic,  $p_1$  is ruled out at the  $\sim 5$  standard deviation level.

Unfortunately, many physicists, over-impressed by the fact that  $\chi^2(p_1)$  appears to be satisfactory, are reluctant to accept that  $p_0$  is strongly favoured by the data.

A similar argument applies to comparing a given set of data with 2 separate hypotheses e.g. fitting a histogram with an exponential or a straight line. Again the **difference** between the  $\chi^2$  quantities provides better discrimination between the hypotheses than do the **individual**  $\chi^2$  values[78]. Another example of using the difference in  $\chi^2$ 's is given in the next section.

There are of course other ways available for comparing two hypotheses. e.g. likelihood ratio, Bayes factor, Bayesian information criterion, etc. For a fuller discussion, see ref. [79]. A description of their application in cosmology can be found in ref. [80]. Problems in choosing priors for the Bayes factor approach for selecting among hypotheses are discussed by Heinrich[81].

### 15.7.5 *Peak Above Smooth Background*

When comparing two hypotheses with our data, we can use the numerical values of the two  $\chi^2$  quantities with a view to making some decision about the hypotheses. For example, we may be fitting a smooth distribution by a power series, and wonder whether we need a quadratic term, or whether a linear expression would suffice. Alternatively we may want to assess whether a mass spectrum favours the existence of a peak on top of a smooth background, as compared with just the smooth background. Qualitatively, if the extra term(s) are unnecessary, they will result in a relatively small reduction in  $\chi^2$ , while if they really are required, the reduction could be larger.

It is sometimes possible to be quantitative about the expected reduction when the extra terms are not needed[82]. If we are in the asymptotic regime, and if the hypotheses are nested,<sup>25</sup> and if the extra parameters of the larger hypothesis are defined under the smaller one, and in that case do not lie on the boundary of their allowed region, then the difference in  $\chi^2$  should itself be distributed as a  $\chi^2$ , with the number of degrees of freedom equal to the number of extra parameters.

An example that satisfies this is provided by the different order polynomials. The hypotheses are nested, in that the linear situation is a special case of a quadratic, where the coefficient of the quadratic term is zero. Thus the extra parameter is defined and within the (infinite) allowed range. Then, provided we have a large

---

<sup>25</sup>This means that for suitable values of the parameters the larger hypothesis reduces to the smaller one.



amount of data, we expect the difference in  $\chi^2$  to have one degree of freedom, so a value larger than around 5 would be unlikely.

A contrast is provided by a smooth background  $C(x)$  compared with a background plus peak,  $C(x) + A \exp[-0.5 * (x - x_0)^2 / \sigma^2]$ . The extra parameters for the peak are its amplitude, position and width:  $A$ ,  $x_0$  and  $\sigma$  respectively. Again the hypotheses are nested, in that  $C(x)$  is just a special case of the peak plus background, with  $A = 0$ . However, although  $A$  is defined in the background only case,  $x_0$  and  $\sigma$  are not, as their values become completely irrelevant when  $A = 0$ . Furthermore, unless the peak plus background fit allows  $A$  to be negative, zero is on the boundary of its allowed region. We thus should not expect the difference of the  $\chi^2$  quantities itself to be distributed as a  $\chi^2$  [83–85]. To assess the significance of a particular  $\chi^2$  difference, this unfortunately means that we have to obtain its distribution ourselves, presumably by Monte Carlo. If we want to find out probabilities of statistical fluctuations at the  $10^{-6}$  level, this requires a lot of simulation, and probably needs us to use something better than brute force.

The problem of non-standard limiting distributions for  $\chi^2$  tests has a substantial statistical literature (see, for example, refs. [86] and [87].)

### 15.7.6 *Incorporating Nuisance Parameters*

The calculation of  $p$ -values is complicated in practice by the existence of nuisance parameters. (For the simple situation described in Sect. 15.7.2, there could be some uncertainty in the estimated background.) There are numerous ways of incorporating them. These include:

- **Conditioning:** For example, with a single nuisance parameter, it may be possible to condition on the sum of the number of counts in the main and the subsidiary experiments, and then to use the binomial distribution to obtain the  $p$ -value.
- **Plug-in  $p$ -value:** The best estimate of the nuisance parameter under the null hypothesis is used to calculate  $p$ .
- **Prior predictive  $p$ -value:** The  $p$ -values are averaged over the nuisance parameters, weighted by their prior distributions. This is in the spirit of the Cousins and Highland approach[65] for upper limits.
- **Posterior predictive  $p$ -value:** This time, the posterior distributions of the nuisance parameters are used for weighting.
- **Supremum  $p$ -value:** The largest  $p$ -value for any possible value of the nuisance parameter is used. This is likely to be useful only when the nuisance parameter is forced to be within some range; or when there is only a small number of possible alternative theoretical interpretations.
- **Confidence interval:** A region of frequentist confidence  $1 - \gamma$  is used for the nuisance parameter(s), and then the adjusted  $p$ -value is  $p_{\max} + \gamma$ , where  $p_{\max}$  is the largest  $p$ -value as the nuisance parameters are varied over their confidence

region. Clearly if it is desired to establish a discovery from  $p$ -values around  $10^{-7}$  or smaller, then  $\gamma$  should be chosen at least an order of magnitude below this.

The properties of these and other methods are compared by Demortier [84], while Cranmer [88] and Cousins et al. [89] have discussed some of them in the context of searches at the LHC.

The role of systematic effects is likely to be more serious here than for upper limits discussed in Sect. 15.4.5. This is because in upper limit situations the number of events is usually small, and so statistical uncertainties dominate. In contrast, discovery claims have  $p$ -values of  $3 \times 10^{-7}$  or smaller, and so tails of distributions are likely to be important.

### 15.7.7 Why $5\sigma$ ?

Unfortunately the usually accepted criterion for claiming a discovery in Particle Physics is that  $p$  should correspond to at least  $5\sigma$ . Statisticians almost invariably ask why such a stringent level is used. One answer is past experience: all too often interesting effects at the  $3\sigma$  or  $4\sigma$  level have gone away as more data are collected. Another is the multiple comparison problem, or “Look Elsewhere Effect” (LEE). While the chance of obtaining a  $5\sigma$  effect in one bin of a particular histogram (“local  $p$ -value”) is really small, it is to be remembered that histograms have many bins,<sup>26</sup> they could be plotted with different selection criteria and different binning,<sup>27</sup> and there are very many other histograms that were or could have been looked at in the course of the experiment.<sup>28</sup> Thus the chance of a  $5\sigma$  fluctuation occurring somewhere in the data (“global  $p$ -value”) is much larger than might at first appear. Calculating a global  $p$ -value may require an excessive amount of Monte-Carlo simulation. Reference [90] circumvents this for asymptotic situations by providing a formula for extrapolating the LEE correction factor from a lower significance level; this requires considerably less simulation.

Finally, physicists subconsciously incorporate Bayesian priors in assessing how likely they feel that they have discovered something new, and hence whether they

---

<sup>26</sup>In calculating a  $p$ -value in such a case, it is very desirable to take into account the number of chances for a statistical fluctuation to occur anywhere in the histogram (or anywhere in the search procedure, for more complicated analyses). At very least, it should be made clear what the basis of the calculated  $p$ -value is.

<sup>27</sup>If a blind analysis is performed, such decisions are made before looking at the data, and so this aspect of the “look elsewhere” effect is reduced.

<sup>28</sup>The extent to which other people’s searches should be included in an allowance for the “look elsewhere” effect depends on the implied question being addressed. Thus are we considering the chance of obtaining a statistical fluctuation in any of the analyses we have performed; or by anyone analysing data in our experiment; or by any Particle Physicist this year? Because of the ambiguity of which specific question is being addressed, which is often not explicitly mentioned, we recommend not including an extra “look elsewhere” factor for this.

should claim a discovery. Thus, in deciding between the possibilities of a new discovery or of an undetected systematic effect, our priors might favour the latter, and hence strong evidence for discovery is required from the data.<sup>29</sup>

However it is not necessarily equitable to use a uniform standard for large general-purpose experiments and for small ones with a specific aim; or for looking for a process which is expected (e.g.  $H^0 \rightarrow \mu^+ \mu^-$ ), as compared with a more speculative search, such as lepton substructure[91]. But physicists and especially journal editors seem to like a defined rule rather than a flexible criterion, so this bolsters the  $5\sigma$  standard. In any case, it is largely a semantic issue, in that physicists finding a  $4.5\sigma$  effect would clearly report it, using judiciously chosen wording to describe the interpretation of their observation.

Statisticians also ask whether models can really be trusted to describe the extreme tails of distributions. In general, this may be so—counting experiments are expected to follow Poisson distributions, with small corrections for possible long time-scale drifts in detector calibrations; and particle decays usually are described by exponential distributions in time. However, the situation is much less clear for nuisance parameters, where uncertainty estimates may be less rigorous, and their distribution is often assumed to be Gaussian (or truncated Gaussian) by default. The effect of these uncertainties on very small  $p$ -values needs to be investigated case-by-case.

It is important to remember that  $p$ -values merely test the null hypothesis. There are more sensitive ways of looking for new physics when a specific alternative is relevant. Thus a very small  $p$ -value on its own is usually not enough to make a convincing case for discovery.

### 15.7.8 *Repetitions in Time*

Often experiments accumulate data over several years. The same search for a new effect may typically be repeated once or twice each year as more data are collected. Does this constitute another factor of  $\sim 20$  in the number of opportunities for a statistical fluctuation to appear? Our reply is “No”. If there had been a  $6\sigma$  signal with the early data (which resulted in a claim for discovery), which had then become only  $3\sigma$  with more data, this would be grounds for downplaying the earlier discovery claim. Thus at any time, there is essentially only one set of data (everything) that is relevant.

For a  $p$ -value to be meaningful, it is important that the time at which the experiment stops collecting data is determined not by the significance of the observed signal but by external factors (e.g. accelerator being decommissioned, ending of funding, etc.). Indeed there is a theorem that states that, provided data is

---

<sup>29</sup>If I were performing an experiment to look for violations of energy conservation, I would require more than  $5\sigma$ , because my prior for energy being conserved is very large.

collected for long enough, it is possible to reach any arbitrary level of significance against a hypothesis that is in fact true.

### 15.7.9 Combining $p$ -Values

In looking for a given new effect, there may be several separate and uncorrelated analyses which are relevant. These could correspond to different decay modes for the new particle; or different experiments looking for the same signal. Thus, if the  $p$ -values for the null hypothesis (i.e. no new physics) for the separate analyses were  $10^{-6}$  and 0.1, what is the corresponding  $p$ -value for the pair of results?<sup>30</sup>

The unambiguous answer is that there is no unique recipe for combining them[92, 93]. There is no single way of taking a uniform distribution in two variables, and finding a transformation  $p_{comb}(p_1, p_2)$  that converts it into a uniform distribution of the single variable  $p_{comb}$ .

Two popular recipes involve asking what is the probability that the smaller  $p$ -value will be  $10^{-6}$  or smaller; or that the product is below  $p_1 * p_2 = 10^{-7}$ . (Note that these probabilities are **not**  $10^{-6}$  and  $10^{-7}$  respectively.) None of the possible methods has the property that in combining three  $p$ -values, the same answer is obtained if  $p_1$  is first combined with  $p_2$ , and then the result is combined with  $p_3$ ; or whether some different ordering is used.

Another problem is the lack of other information that might be relevant. For example, the  $p$ -values might arise from  $\chi^2$ 's with different numbers of degrees of freedom  $\nu$  e.g.  $\chi_1^2 = 90$  for 100 degrees of freedom, and  $\chi_2^2 = 20$  for  $\nu = 1$ . The second has a very small  $p$ -value, so many combination methods (including the two mentioned above) would conclude that overall the data do not look consistent with the null hypothesis. However, another plausible-sounding method is to add the separate  $\chi^2$  values and also the individual  $\nu$ ,<sup>31</sup> to obtain a total  $\chi^2 = 110$  for  $\nu = 101$ , which sounds perfectly satisfactory. The resolution of this discrepancy of interpretation depends on the nature of the two tests. If the second analysis with  $\chi^2 = 20$  corresponded to just one extra measurement like the previous 100, then it seems reasonable to combine the  $\chi^2$  values and the  $\nu$ , and to conclude that overall there is indeed nothing surprising. But on the other hand, if the second measurement was genuinely different, and an alternative way of looking for some discrepancy, then it may be more appropriate to combine the  $p$ -values by one of the earlier methods, which suggest that the overall consistency with theory is not good. It

<sup>30</sup>Rather than combining  $p$ -values, it is of course much better to use the complete sets of original data (if available) for obtaining the combined result.

<sup>31</sup>The method described earlier involving the product of the  $p$ -values is equivalent to converting each  $p$  to a  $\chi^2$ , assuming that  $\nu = 2$ , regardless of whether this was the actual number of degrees of freedom, and then adding the  $\chi^2$  and also the  $\nu$ .

is this extra information about the nature of the two tests that determines which combination method might be appropriate.

It is clearly important to decide in advance what combination method should be used, without reference to the specific data being analysed.

## 15.8 Blind Analyses

These are becoming increasingly popular as a means of avoiding personal bias affecting the result. They involve keeping part of the data unseen by the analysers, until the data selection procedure and the analysis method have been completely defined, all correction procedures specified, etc.

One of the early suggestions to use a blind analysis in a Particle Physics experiment was due to Luis Alvarez. An experiment at Stanford had looked for quarks, by measuring the residual charge on small spheres that were levitated in a superconducting magnet. If a single free quark were present in a sphere, the residual charge would be a third or two-thirds of the electron's charge. Several of the balls tested indeed yielded such values[94]. A potential problem was that large corrections had to be applied to the raw data in order to extract the final result for the charge. The suspicion was that maybe the experimenters were (subconsciously) applying corrections until the value turned out to be 'satisfactory'. The blind approach involved the computer adding a random number to the raw value of the charge, which would then be corrected until the experimentalists were satisfied, and only then would the computer subtract the random number to reveal the final answer for that sphere.<sup>32</sup>

There are various methods of performing blind analyses[95] most of which aim to allow the experimentalists to look at some of the real data, in order to perform checks that nothing is terribly wrong. Some of these are:

- The computer adds a random number to the data, which is only subtracted after all corrections are applied. This was the method suggested by Alvarez.
- Use only Monte Carlo to define the procedure. This completely avoids the danger of allowing the data to determine the procedure to be used, but suffers from the drawback that the data cannot be compared with the Monte Carlo, to check that the latter is reasonable.
- Use only a fraction of the data for defining the procedure, which then is held fixed for the remainder of the data. In principle, an optimisation can be employed to determine the fraction to be kept open, but in practice this is often decided by choosing a semi-arbitrary time after which the future data is kept blind.

---

<sup>32</sup>This suggestion was implemented, but in fact no subsequent results were published. The current consensus is that this 'discovery' of free quarks is probably spurious.

- The signal region is defined by a certain part of multi-dimensional space, and this is kept hidden, but all other regions, including those adjacent to the signal, are available for inspection.
- Keep the Monte Carlo parameters hidden. This is a technique suggested by the TWIST experiment in their high statistics precision determination of parameters associated with muon decay. The procedure involves comparing the data with various simulated sets, generated with a series of different parameter values. The data and the simulations are both visible, but the parameter values used to generate the simulations are kept hidden.
- Keep visible only a fraction of the contents of each bin of a histogram. This is used by the MINOS experiment searching for neutrino oscillations; these would affect the energy distribution of the observed events. By keeping visible different unknown fractions of the data in each bin, the energy spectral shape cannot be determined from the visible part of the data.

If several different groups within the same collaboration are performing similar analyses for extracting some specific parameter, then it is desirable to fix the procedure for selecting which result to present, or alternatively how to combine the separate results. This should be done before the results are seen, and is worth doing even if the individual analyses were not “blind”.

A question that arises with blind analyses is whether it should be permitted to modify the analysis after the data had been unblinded. It is generally agreed that this should not be done, unless everyone would regard it as ridiculous not to do so. For example, if a search for rare events yielded 10 candidates over the course of a year’s run, all of which occurred on Sunday mornings at precisely 1.17 a.m., it would be prudent to do some further investigation before publishing. If ‘post-unblinding’ modification of the procedure is performed, this should be made clear in any publication.

## 15.9 Topics that Deserve More Attention

### 15.9.1 *Statistical Software*

Particle physicists tend to write their own software for performing statistical computations. Although this has educational merits, it is inefficient use of one’s time. The data-manipulation system of programmes ROOT/RooFit/RooStats contains many useful statistical routines[96]. Tools also exist for implementing many methods for separating signal from background[43, 44].

A problem with these is that they are too easy to use. In the hands of a non-critical user, the required input data instructions may contain some error, with the consequence that they will produce the solution to a different procedure than the intended one. It is very important to check that the result obtained is not unreasonable.

### 15.9.2 *Deep Learning*

This involves the use of sophisticated techniques[45] for achieving nearly optimal extraction of information from data, but which are still relatively unfamiliar to many scientists. It is important to develop a set of protocols to ensure that they perform in a reliable manner, and are not introducing subtle biases of which users are unaware.

### 15.9.3 *Unfolding Data or Smearing Theory?*

Observed experimental distributions are almost always smeared versions of ‘the true distributions of Nature’. It is simpler to compare theory and data by smearing the theory, rather than trying to unfold the experimental effects from the data, as the latter is a less stable procedure and also introduces correlations among the bins of the unfolded distribution. Some fields tend to favour deconvolution; this is partly because it is rarer for them to have a dominant theoretical model with which the data is to be compared. Unfolding does have the advantage that it provides an estimate of the ‘true’ distribution, with which any future theory can be compared. Also it can be looked at by a physicist, but we are not accustomed to readily interpreting data where the contents of the histogram bins are highly correlated.

There are some situations where unfolding is desirable. For example, it allows the comparison of distributions from different experiments, with different resolutions. Another is using experimental data for tuning Monte Carlo generators; smearing the data at each step of the optimisation increases the computation time too much.

Even for checking in future whether new theories are compatible with data does not necessarily require unfolding. Provided that the smearing matrix of the detector is provided, the future data can be smeared, and then compared with the actual (not unfolded) data. However, including the effects of systematics can be a complication.

Sessions at the 2011 PHYSTAT workshop[9] and at CERN’s PHYSTAT $\nu$  meeting[12] were devoted to unfolding. Blobel[97] has reviewed the topic, while ref. [98] contains a statistician’s view of the statistical issues involved in unfolding.

### 15.9.4 *Visualisation*

The combination of the human eye and brain is very powerful at detecting patterns in data (even if sometimes they are not there!) This can be useful in deciding how to analyse the data; as a check on whether the result of an analysis is plausible; whether a machine learning method for separating signal from background is performing sensibly; etc. Such human inspection of data is feasible if there are only a small number (below 4) of relevant variables. Techniques for inspecting multi-dimensional data would be valuable.

### 15.9.5 *Non-parametric Methods*

These are so unknown to most Particle Physicists that they are usually unaware when they are using them. Simple examples include:

- A histogramme as an estimate of the density distribution of a variable of interest.
- Kernel density estimation.
- Kolmogorov-Smirnov or Anderson-Darling methods, to test whether distributions are consistent.
- Classification schemes based on  $k$  nearest neighbours.
- Neural networks

These all avoid the need to specify a particular parametric form, and hence the values of any parameters. In general such a method is less powerful than a parametric one, if the latter were available and relevant.

### 15.9.6 *Collaboration with Statisticians*

Other scientists seem to be better than particle physicists about involving statisticians in the analysis of their data. This is partly due to the fact that we like to try out statistical techniques ourselves; that we consider our data is too complicated for other people to deal with; and that we are somewhat over-protective of our data, and are reluctant to share it with others. None of this is particularly convincing, and it is clear that we would benefit from the involvement of professional statisticians. The advantages of having them participating in the recent PHYSTAT meetings have been obvious.

In the past, Particle Physicists have on occasion asked rather specific questions to Statisticians they happened to know. Statisticians prefer to be much more directly involved with the data itself. With analyses becoming more and more complex, it will be highly desirable for them to be affiliated with experimental groups.

## 15.10 Conclusion

Although the statistical aspects of many particle physics analyses are already at a sophisticated level, it is clear that there are many practical statistical issues to be resolved. With the increasing complexity of scientific investigations, more active collaboration with statisticians and machine learning experts will result in a better understanding of the relevant techniques and improved analyses in the future.

**Acknowledgements** I wish to acknowledge the patience and expertise of David Cox, Brad Efron, Jerry Friedman and David van Dyk, and also of other Statisticians too numerous to list, in explaining statistical issues to me; the ones who have contributed to the PHYSTAT meetings have



been particularly helpful. My understanding of the practical application of statistical techniques has improved considerably as a result of discussions with many experimental Particle Physics colleagues, and in particular with the members of the CDF and CMS Statistics Committees. I especially wish to thank Bob Cousins, Luc Demortier and Joel Heinrich for their careful reading and valuable comments on the original version of this article. To all of you, I am most grateful.

The Leverhulme Foundation kindly provided a grant which partially supported the original version this work.

## References

1. Workshop on Confidence Limits, CERN Yellow Report 2000-05.
2. FNAL Confidence Limits Workshop (2000), <http://conferences.fnal.gov/CLW/>.
3. Advanced Statistical Techniques in Particle Physics, Durham (2002) IPPP/02/39.
4. Proceedings of PHYSTAT2003, eConf C030908, SLAC-R-703.
5. "PHYSTAT05: Statistical Problems in Particle Physics, Astrophysics and Cosmology", Imperial College Press (2006), <http://www.physics.ox.ac.uk/phystat05/>.
6. BIRS Workshop on "Statistical inference Problems in High Energy Physics and Astronomy", Banff (2006), [http://www.birs.ca/birspages.php?task=displayevent&event\\_id=06w5054](http://www.birs.ca/birspages.php?task=displayevent&event_id=06w5054).
7. PHYSTAT-LHC Workshop on "Statistical Issues for LHC Physics" (2007), <http://phystat-lhc.web.cern.ch/phystat-lhc/2008-001.pdf>.
8. BIRS Workshop on "Statistical issues relevant to significance of discovery claims (10w5068)" (2010) <https://www.birs.ca/events/2010/5-day-workshops/10w5068>
9. PHYSTAT-LHC Workshop, "Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding", <https://cdsweb.cern.ch/record/1306523/files/CERN-2011-006.pdf>
10. PHYSTAT $\nu$  in Japan (2016), <https://indico.cern.ch/event/735431/>
11. PHYSTAT $\nu$  Workshop on Statistical issues in experimental neutrino physics, FNAL (2016), <https://indico.fnal.gov/event/11906/>
12. PHYSTAT $\nu$  in CERN (2019), <https://indico.cern.ch/event/735431/>
13. Roger Barlow, "Statistics: a Guide to the Use of Statistical Methods in the Physical Sciences", Wiley (1989).  
O. Behnke et al (eds), "Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods", Wiley (2013).  
Glen Cowan, "Statistical Data Analysis", Oxford University Press (1998).  
F. E. James, "Statistical Methods in Experimental Physics", World Scientific Publishing Co (2007).  
L. Lista, "Statistical Methods for Data Analysis in Particle Physics", Springer (2017).  
Louis Lyons, "Statistics for Nuclear and Particle Physics", Cambridge University Press (1986). See also <https://www-cdf.fnal.gov/physics/statistics/notes/Errata2.pdf> for an Update.  
Byron Roe, "Probability and Statistics in Experimental Physics", Springer Verlag (1991).
14. M. Tamabashi et al., "Review of Particle Physics", Phys. Rev. D98 030001 (2018).
15. BaBar Statistics Working Group, <http://www.slac.stanford.edu/BFROOT/www/Statistics/>.
16. CDF Statistics Committee, [http://www-cdf.fnal.gov/physics/statistics/statistics\\_home.html](http://www-cdf.fnal.gov/physics/statistics/statistics_home.html)
17. ATLAS Statistics Forum, [https://twiki.cern.ch/twiki/bin/view/Atlas/StatisticsTools#Statistics\\_Forum](https://twiki.cern.ch/twiki/bin/view/Atlas/StatisticsTools#Statistics_Forum).
18. CMS Statistics Committee, <https://twiki.cern.ch/twiki/bin/view/CMS/StatisticsCommittee>.
19. D. van Dyk, "Statistical quantification of discovery in neutrino physics", Neutrino2016 XXVII Int Conf on Neutrino Physics and Astrophysics, <http://neutrino2016.iopconfs.org/home>
20. L. Lyons, "Lessons learned from PhysStat-nu", NuPhys2016: Prospects in Neutrino Physics, <https://indico.ph.qmul.ac.uk/indico/conferenceDisplay.py?confId=170>; and

- “Statistical issues towards PHYSTAT $\nu$  2019”, NuPhys2018: Prospects in Neutrino Physics, <https://indico.ph.qmul.ac.uk/indico/conferenceDisplay.py?confId=289>
21. CERN European Schools, [https://physicschool.web.cern.ch/physicschool/ESHEP/previous\\_eshep.html](https://physicschool.web.cern.ch/physicschool/ESHEP/previous_eshep.html);  
CERN Latin-American Schools, [https://physicschool.web.cern.ch/physicschool/CLASHEP/previous\\_clashep.html](https://physicschool.web.cern.ch/physicschool/CLASHEP/previous_clashep.html); and  
CERN Asia-Pacific Schools, <http://aepshep.org/previous-schools.html>.
  22. L. Lyons, “Statistical Issues in Particle Physics” in “Elementary Particles: Detectors for Particles and Radiation, Part 1: Principles and Methods” Eds C. Fabjan and H. Schopper, (Landolt-Bornstein, **21B1** 2011).
  23. ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, Phys. Lett. **B716** (2012) 1;  
CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”, Phys. Lett. **B716** (2012) 30.
  24. N. Reid, “Some aspects of design of experiments”, ref. [7], p. 94.
  25. R. Neal, “Computing likelihood functions when distributions are defined by simulations with nuisance parameters”, ref. [7], p. 101; and in ref. [6].
  26. J. Linnemann, “A pitfall in estimating systematic errors”, ref. [7], p. 94.
  27. J. Heinrich and L. Lyons, Annual Reviews of Nuclear and Particle Science **57** (2007) 145.
  28. L. Lyons, “Bayes and Frequentism: a Particle Physicist’s perspective”, (2013) <https://arxiv.org/pdf/1301.1273.pdf>
  29. R. D. Cousins, Am. J. Phys. **63** (1995) 398.
  30. R. Barlow, “Asymmetric errors”, ref. [4], p. 250; and ref. [5], p. 56.
  31. F. Garwood, “Fiducial limits for Poisson the distribution”, Biometrika **28** (1936) 437.
  32. J. Heinrich, “Coverage of error bars for Poisson data” (2003),  
[http://www-cdf.fnal.gov/publications/cdf6438\\_coverage.pdf](http://www-cdf.fnal.gov/publications/cdf6438_coverage.pdf).
  33. J. Heinrich, “Pitfalls of Goodness-of-Fit from Likelihood”, ref. [4], p. 52.
  34. S. Baker and R. D. Cousins, “Clarification of the use of  $\chi^2$  and likelihood functions in fits to histograms”, NIM **221** issue 2 (1984) 437.
  35. G. Cowan, Eur Phys J C (2019) 79:133.
  36. D. Cox, private communication
  37. L. Lyons and E. Chapon, “Combining parameter values or  $p$ -values” (2017) <https://arxiv.org/pdf/1704.05540.pdf>
  38. P. Dauncey et al, “Handling uncertainties in background shapes: the discrete profiling method”, JINST **10** no.04 (2015) 04015
  39. G. Punzi, “Comments on likelihood fits with variable resolution”, ref. [4], p. 235.
  40. P. Catastini and G. Punzi, “Bias-free estimation of multicomponent maximum likelihood fits with component-dependent templates”, ref. [5], p. 60.
  41. H. B. Prosper, “Multivariate methods: a unified perspective”, ref. [3], p. 91.
  42. J. H. Friedman, “Recent advances in predictive (machine) learning”, ref. [4], p. 196; and  
“Separating signal from background using ensembles of rules”, ref. [5], p. 127.
  43. I. Narsky, “StatPatternRecognition in analysis of HEP and Astrophysics data”, ref. [7], p. 188.
  44. A. Hocker et al, “TMVA, Toolkit for Multi-Variate data Analysis with ROOT”, ref. [7], p. 184.
  45. D. Guest, K. Cranmer and D. Whiteson, “Deep Learning and Its Application to LHC Physics”, Annual Review of Nuclear and Particle Science **68** (2018) 161;  
A. Radovic et al, “Machine learning at the energy and intensity frontiers of particle physics”, Nature **560** (2018) 41;  
A. J. Larkoski, I. Moulton and B. Nachman “Deep Learning and Its Application to LHC Physics” (2017) <http://arxiv.org/abs/arXiv:1709.04464>;  
I. Goodfellow, Y. Bengio and A. Courville, “Deep Learning” (2016), MIT Press.
  46. CERN’s “Inter-Experimental LHC Machine Learning Working Group (IML)”, <https://iml.web.cern.ch/>
  47. Fermilab’s “ML at the Intensity and Cosmic Frontiers”, <https://machinelearning.fnal.gov/>
  48. L. Lyons, Nucl. Inst. Meth. **A324** (1993) 565.

49. S. Whiteson and D. Whiteson, “Stochastic Optimization for Collision Selection in High Energy Physics,” IAAI 2007: Proceedings of the Nineteenth Annual Innovative Applications of Artificial Intelligence Conference (July 2007) 1819.
50. N. Tishby and N. Zaslavsky, “Deep Learning and the Information Bottleneck Principle” (2015), <https://arxiv.org/abs/1503.02406>;  
R. Shwartz-Ziv and N. Tishby, “Opening the Black Box of Deep Neural Networks via Information” (2017), <https://arxiv.org/abs/1703.00810>
51. L. Lyons, “Raster scan or 2-D approach?”, <https://arxiv.org/pdf/1404.7395.pdf>
52. A. Michelson and E. Morley, “On the Relative Motion of the Earth and the Luminiferous Ether”, *American Journal of Science*, **34** (1887) 203: 333.
53. I. Narsky, “Comparison of upper limits”, in ref. [2].
54. G. J. Feldman and R. D. Cousins, *Phys. Rev.* **D57** (1998) 3873.
55. D. R. Cox, “Some problems connected with statistical inference”, *Annals of Mathematical Statistics* **29** (1958) 357.
56. B. Sen, M. Walker and M. Woodroffe, “On the Unified Method with Nuisance Parameters”, *Statistica Sinica* **19** (2009) 301
57. G. Punzi, “Sensitivity of searches for new signals and its optimisation”, ref. [4], p. 235.
58. A. L. Read, “Modified frequentist analysis of search results”, ref. [1], p. 81; “Presentation of search results—the  $CL_s$  method”, ref. [3], p. 11.  
T. Junk, “Confidence level computation for combining searches with small statistics”, *NIM A* **434** (1999) 435.
59. W. A. Rolke, A. M. Lopez and J. Conrad, *Nuclear Instruments and Methods* **A551** (2005) 493.
60. J. Heinrich et al. “Interval estimation in the presence of nuisance parameters. 1. Bayesian approach”, CDF note 7117 (2004), [https://www-cdf.fnal.gov/physics/statistics/notes/cdf7117\\_bayesianlimit.pdf](https://www-cdf.fnal.gov/physics/statistics/notes/cdf7117_bayesianlimit.pdf)
61. L. Demortier, “A fully Bayesian computation of upper limits for Poisson processes”, CDF note 5928 (2004).
62. J. Heinrich, “The Bayesian approach to setting limits: what to avoid”, ref. [5], p. 98.
63. G. Punzi, “Ordering algorithms and confidence intervals in the presence of nuisance parameters”, ref. [5], p. 88.
64. R. Cousins, *Nuclear Instruments and Methods* **A417** (1998) 391.
65. R. D. Cousins and V. L. Highland, *Nuclear Instruments and Methods* **A320** (1992) 331.
66. J. Heinrich, “Review of Banff challenge on upper limits”, ref. [7], p. 125.
67. CDF Statistics Committee, “Recommendations concerning limits” (2005), <http://www-cdf.fnal.gov/physics/statistics/recommendations/limits.txt>.
68. A. Kolmogorov, “Sulla determinazione empirica di una legge di distribuzione”, *G. Ist. Ital. Attuari*, **4** (1933) 83.  
N. Smirnov, “Table for estimating the goodness of fit of empirical distributions”, *Annals of Mathematical Statistics*. **19** (2) (1948) 279, [doi:10.1214/aoms/1177730256](https://doi.org/10.1214/aoms/1177730256)
69. T. W. Anderson and D. A. Darling, “Asymptotic theory of certain ‘goodness-of-fit’ criteria based on stochastic processes”, *Annals of Mathematical Statistics*. **23** (1952) 193, [doi:10.1214/aoms/1177729437](https://doi.org/10.1214/aoms/1177729437); and “A Test of Goodness-of-Fit”, *Journal of the American Statistical Association*, **49** (1954) 765, [doi:10.2307/2281537](https://doi.org/10.2307/2281537)
70. L. Lyons, D. Gibaut and P. Clifford, *Nuclear Instr. Meth.* **270** (1988) 210.
71. L. Lyons, A. Martin and D. Saxon, *Phys Rev* **D41** (1990) 982.
72. T. Trippe and Particle Data Group, private communication.
73. N. Suzuki et al, “Hubble Space Telescope cluster supernova study. V: Improving the Dark Energy constraints above  $z > 1$  and building an early-type-hosted supernova sample”, *Astrophys J* **746** (2012) 85, [arXiv:1105.3470\[astro-ph.CO\]](https://arxiv.org/abs/1105.3470)
74. L. Lyons and N. Wardle, “Statistical issues in searches for new phenomena in High Energy Physics”, *J Phys G Nucl Part Phys* **48** (2018) 033001

75. ATLAS Collaboration, CMS Collaboration and LHC Higgs Combination Group, “Procedure for the LHC Higgs boson search combination in Summer 2011”, [http://cds.cern.ch/record/1379837/files/NOTE2011\\_005.pdf](http://cds.cern.ch/record/1379837/files/NOTE2011_005.pdf)
76. G. Cowan, K. Cranmer, E. Gross and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics”, *Eur. Phys. J. C* **71** (2011) 1554.
77. “Statistical Inference in the 21st Century: A World Beyond  $p < 0.05$ ”, *American Statistician* **73** (2019)
78. L. Lyons, “Comparing two hypotheses” (1999), [http://www-cdf.fnal.gov/physics/statistics/statistics\\_recommendations.html](http://www-cdf.fnal.gov/physics/statistics/statistics_recommendations.html).
79. L. Lyons, “Methods for comparing two hypotheses”, [http://www.physics.ox.ac.uk/users/lyons/R\\_H\\_2009.pdf](http://www.physics.ox.ac.uk/users/lyons/R_H_2009.pdf).
80. R. Trotta, *Contemporary Physics* **49** (2008) 71.
81. J. Heinrich, “A Bayes factor example: Poisson discovery”, CDF note 9678 (2009), <http://newton.hep.upenn.edu/~heinrich/bfexample.pdf>.
82. S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses”, *Annals of Math. Stat.* **9** (1938) 60.
83. R. Protassov et al., “Statistics: Handle with care. Detecting multiple model components with the likelihood ratio test”, *Astrophysics Journal* **571** (2002) 545.
84. L. Demortier, “p-values and nuisance parameters”, ref [7], p. 23.
85. L. Demortier, “Setting the scene for p-values” (2006), [http://birs.pims.math.ca/~06w5054/Luc\\_Demortier.pdf](http://birs.pims.math.ca/~06w5054/Luc_Demortier.pdf).
86. S. G. Self and K. Y. Liang, *JASA* **82** (1987) 605.
87. M. Drton, “Likelihood ratio tests and singularities”, *Annals of Statistics* **37** No. 2 (2009) 979, <http://arxiv.org/abs/math/0703360>.
88. K. Cranmer, “Statistics for LHC: progress, challenges and future”, ref. [7], p. 47.
89. R. Cousins, J. Linnemann and J. Tucker, *Nuclear Instr. Meth. A* **595** (2008) 480.
90. E. Gross and O. Vitells, “Trial factors for the look elsewhere effect in high energy physics”, *E Phys J C* **70** (2010) 525.
91. L. Lyons, “Discovering the Significance of  $5\sigma$ ” (2013), <https://arxiv.org/abs/1310.1284>
92. CDF Statistics Committee, “Frequently asked questions”, [http://www-cdf.fnal.gov/physics/statistics/statistics\\_faq.html#iptn4](http://www-cdf.fnal.gov/physics/statistics/statistics_faq.html#iptn4).
93. R. Cousins, “Annotated bibliography on some papers on combining significances or p-values”, arXiv:0705.2209 (2007)
94. G. S. LaRue, J. D. Phillips and W. M. Fairbank, *Phys. Rev. Lett.* **46** (1981) 967.
95. J. R. Klein and A. Roodman, *Annual Review of Nuclear and Particle Physics* **55** (2005) 141.
96. I. Antcheva et al., “ROOT: A C++ framework for petabyte data storage, statistical analysis and visualization”, *Computer Physics Communications, Anniversary Issue*; **180** Issue 12 (2009) 2499;  
W. Verkerke and D. Kirkby, “The RooFit toolkit for data modeling” (2003), [arXiv:physics/0306116](https://arxiv.org/abs/physics/0306116);  
L. Moneta et al., “The RooStats Project”, 13<sup>th</sup> Int. Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT2010), [arXiv:1009.1003](https://arxiv.org/abs/1009.1003).PoSACAT:057
97. V. Blobel, “Unfolding” in ‘Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods’ page 187 in [13].
98. M. Kuusela, “Uncertainty quantification in unfolding elementary particle spectra at the Large Hadron Collider”, (2016) PhD thesis at EPFL Lausanne, [https://infoscience.epfl.ch/record/220015/files/EPFL\\_TH7118.pdf](https://infoscience.epfl.ch/record/220015/files/EPFL_TH7118.pdf)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



# Chapter 16

## Integration of Detectors into a Large Experiment: Examples from ATLAS and CMS



Daniel Froidevaux

### 16.1 Introduction

#### 16.1.1 *The Context*

The Large Hadron Collider (LHC) is the proton-proton accelerator which began operation in 2010 in the existing LEP tunnel at CERN in Geneva, Switzerland. It represents the next major step in the high-energy frontier beyond the Fermilab Tevatron (proton-antiproton collisions at a centre-of-mass energy of 2 TeV), with its design centre-of-mass energy of 14 TeV and luminosity of  $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ . The high design luminosity is required because of the small cross-sections expected for many of the benchmark processes (Higgs-boson production and decay, new physics scenarios such as supersymmetry, extra dimensions, etc.) used to optimise the design of the general-purpose detectors over a period of 15 years or so. To achieve this luminosity and minimise the impact of simultaneous inelastic collisions occurring at the same time in the detectors (a phenomenon usually called pileup), the LHC beam crossings are 25 ns apart in time, resulting in 23 inelastic interactions per crossing on average at design luminosity. Two general-purpose experiments, ATLAS and CMS, were proposed for operation at the LHC in 1994 [1], and approved for construction in 1995. The experimental challenges undertaken by these two projects of unprecedented size and complexity in the field of high-energy physics, the construction and integration achievements realised over the years 2000–2008, and the expected performance of the commissioned detectors are described in a variety of detailed documents, such as the detector papers [2, 3]. In this chapter, much of the description of the lessons learned based on this huge effort, and of

---

D. Froidevaux (✉)  
CERN, Geneva, Switzerland  
e-mail: [Daniel.Froidevaux@cern.ch](mailto:Daniel.Froidevaux@cern.ch)

the comparisons in terms of expected performance have been taken and somewhat updated from a recent review [4]. For completeness, it is important to mention also the two more specialised and smaller experiments, ALICE [5] and LHCb [6]. In 2019, at a moment when the accelerator and experiments have just completed very successfully the so-called run-2 with 4 years of operation at a centre-of-mass energy of 13 TeV, and after run-1 with operation at lower energies topped with the discovery of the Higgs boson, it is interesting to look back not only on the period of construction and integration with its great expectations, a period which is the main focus of this chapter, but also on almost 10 years of operation and data-taking with its own challenges and of course with the excitement stemming from the analysis of real data.

The prime motivation of the LHC is to elucidate the nature of electroweak symmetry breaking, for which the Higgs mechanism is presumed to be responsible. The experimental study of the Higgs mechanism can also shed light on the consistency of the Standard Model at energy scales above 1 TeV. The Higgs boson is generally expected to have a mass below about 200 GeV [7]. This expectation could be relaxed if there are problems in the interpretation of the precision electroweak data [8] or if there are additional contributions to the electroweak observables [9]. A variety of models without Higgs bosons have also been proposed more recently, together with mechanisms of partial unitarity restoration in longitudinal vector boson scattering at the TeV scale [10]. All these possibilities may appear to be remote, but they serve as a reminder that the existence of a light Higgs boson cannot be taken for granted.


Theories or models beyond the Standard Model invoke additional symmetries (supersymmetry) or new forces or constituents (strongly-broken electroweak symmetry, technicolour). It is generally hoped that discoveries at the LHC could provide insight into a unified theory of all fundamental interactions, for example in the form of supersymmetry or of extra dimensions, the latter requiring modification of gravity at the TeV scale. There are therefore several compelling reasons for exploring the TeV scale and the search for supersymmetry is perhaps the most attractive one, particularly since preserving the naturalness of the electroweak mass scale requires supersymmetric particles with masses below about 1 TeV.

### ***16.1.2 The Main Initial Physics Goals of ATLAS and CMS at the LHC***

There have been many studies of the LHC discovery potential as a function of the integrated luminosity and the ones released just before data-taking [11, 12] have focussed on the first few years, over which about  $10\text{ fb}^{-1}$  of integrated luminosity were expected to be accumulated by each experiment.

With some optimism that the performance of the ATLAS and CMS detectors would be understood rapidly and would be close to expectations, the expectations

**Fig. 16.1** Integrated luminosity required per experiment as a function of the mass of the Standard Model Higgs boson for a  $5\sigma$  discovery or an exclusion at the 95% confidence level, combining the capabilities of ATLAS and CMS




at the time were that a Standard Model Higgs boson could be discovered at the LHC with a significance above  $5\sigma$  over the full mass range of interest and for an integrated luminosity of only  $5\text{ fb}^{-1}$ , as shown in Fig. 16.1. This discovery potential should, however, be taken with a grain of salt, since the evidence for a light Higgs boson of mass in the 110–130 GeV range would not only have to be combined over both experiments but also over several channels with very different final states ( $H \rightarrow \gamma\gamma$  decays in association with various jet topologies,  $t\bar{t}H$  production with  $H \rightarrow b\bar{b}$  decay and  $q\bar{q}H$  production with  $H \rightarrow \tau\tau$  decay). Achieving the required sensitivity in each of these channels would require an excellent understanding of the detailed performance of most elements of these complex detectors and would therefore require sufficient experimental data and time.

The discovery potential for supersymmetry was expected to be very substantial in the very first months of data-taking, since only  $100\text{ pb}^{-1}$  of integrated-luminosity would be sufficient to discover squarks or gluinos with masses below about 1.3 TeV [1, 11, 13], a large increase in sensitivity with respect to that ultimately achieved at the Tevatron. This sensitivity would increase to 1.7 TeV for an integrated luminosity of  $1\text{ fb}^{-1}$  and to about 2.2 TeV for  $10\text{ fb}^{-1}$ , as shown in Fig. 16.2.

The few examples above illustrate the wide range of physics opened up by the seven-fold increase in energy from the Tevatron to the LHC. Needless to say, all Standard Model processes of interest, QCD jets, vector bosons and especially top quarks, would be produced in unprecedented abundance at the LHC, as illustrated in Table 16.1, and would therefore be studied with high precision by ATLAS and CMS.



**Fig. 16.2** Discovery potential for supersymmetry, expressed as lines corresponding to integrated luminosities ranging from 1 to 300 fb<sup>-1</sup> in the (*m*<sub>0</sub>, *m*<sub>1/2</sub>) parameter plane, shown as an example for the CMS experiment. Also shown are lines representing constant squark or gluino masses. The discovery potential depends only weakly on the values assumed for tanβ, *A*<sub>0</sub> and the sign of μ




**Table 16.1** For a variety of physics processes expected to be the most abundantly produced at the LHC, expected numbers of events recorded by ATLAS and CMS for an integrated luminosity of 1 fb<sup>-1</sup> per experiment

Physics process	Number of events per 1 fb <sup>-1</sup>
QCD jets with $E_T > 150$ GeV	10 <sup>6</sup> (for 10% of trigger bandwidth)
$W \rightarrow \mu\nu$	$7.0 \cdot 10^6$
$Z \rightarrow \mu\mu$	$1.1 \cdot 10^6$
$tt \rightarrow e/\mu + X$	$1.6 \cdot 10^5$
Gluino-gluino production (mass about 1 TeV)	$10^2 \dots 10^3$

### 16.1.3 A Snapshot of the Current Status of the ATLAS and CMS Experiments

From the year 2000 to end of 2009, the experiments have had to deal in parallel with a very complex set of tasks requiring a wide diversity of skills and personnel:

- the construction of the major components of the detectors was complete or nearing completion at the end of 2006, after a very long period of research



**Fig. 16.3** Left: picture of the ATLAS barrel toroid superconducting magnet with its eight coils of 25 m length and of the ATLAS barrel calorimeter with its liquid Argon electromagnetic calorimeter and its scintillating tile hadronic calorimeter, as installed in the experimental cavern. Right: picture of the first end-cap LAr cryostat, including the electromagnetic, hadronic and forward calorimeters, as it is lowered into its docking position on one side of the ATLAS pit



**Fig. 16.4** Left: picture of the CMS superconducting solenoid, as integrated with the barrel muon system (outside) and with the barrel hadron calorimetry (inside). Right: picture of the insertion of the CMS silicon-strip tracker into the barrel crystal calorimeter

and development, including validation in terms of survival to irradiation and preparation of industrial manufacturing;

- the integration and installation phase began approximately in 2003 and extended all the way to 2007 for the last major components. ATLAS was being installed and commissioned directly in its underground cavern (see Fig. 16.3). In contrast, CMS is modular enough that it could be assembled above ground (see Fig. 16.4).
- the commissioning of the experiments with cosmic rays began in 2006, with the biggest campaigns in 2008 and 2009. These have yielded a wealth of initial results on the performance of the detectors in situ, a very important asset to ensure a rapid commissioning of the detectors for physics with collisions;
- the next commissioning step was achieved in an atmosphere of great excitement with first collisions at the injection energy of 900 GeV of the LHC machine and with very low luminosities of the order of  $10^{26} - 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ . All detector

components were able to record significant samples of data, albeit at low energy and with insufficient statistics to fully commission the trigger and reconstruction algorithms dedicated to provide the signatures required for the initial Standard Model measurements and searches for new physics.

In parallel with the rapidly evolving integration, installation and commissioning effort at the experimental sites, the collaborations have also reorganised themselves to evolve as smoothly and efficiently as possible from a distributed construction project with a strong technical co-ordination team to a running experiment with the emphasis shifting to monitoring of the detector and trigger operation, understanding of the detector performance in the real LHC environment and producing the first physics results. A small but significant part of the human and financial resources are already focusing on the necessary upgrades to the experiments required by the LHC luminosity upgrade programme.

This chapter has been structured in the following way: Sect. 16.2 presents an overview of the ATLAS and CMS projects in terms of their main design characteristics, describes briefly the magnet systems, and summarises the main lessons learned from the 15-year long research and development and construction period. The next three sections, Sects. 16.3–16.5, describe in more detail the main features and challenges related respectively to the inner tracker, to the calorimetry and to the muon spectrometer, in the specific case of the ATLAS experiment. The subsequent two sections, Sects. 16.6 and 16.7, discuss in broad terms the various aspects of, respectively, the trigger and data acquisition system and the computing and software, again in the context of the ATLAS experiment. The next section, Sect. 16.8, summarises and compares briefly the expected performances at the time of beginning of data-taking of the main ATLAS and CMS systems. The last and final section, Sect. 16.9, gives a very brief overview of the performance and physics results achieved over the past 10 years.

## 16.2 Overall Detector Concept and Magnet Systems

This section presents an overview of the ATLAS and CMS detectors, based on the main physics arguments which guided the conceptual design, and describes the magnet systems, which have driven many of the detailed design aspects of the experiments.

### 16.2.1 Overall Detector Concept

Figures 16.5 and 16.6 show the overall layouts respectively of the ATLAS and CMS detectors and Table 16.2 lists the main parameters of each experiment. Both experiments are designed somewhat as cylindrical onions consisting of:




Fig. 16.5 Overall layout of the ATLAS detector




Fig. 16.6 Overall layout of the CMS detector

**Table 16.2** Main design parameters of the ATLAS and CMS detectors

Parameter	ATLAS	CMS
Total weight [tons]	7000	12,500
Overall diameter [m]	22	15
Overall length [m]	46	20
Magnetic field for tracking [T]	2	4
Solid angle for precision measurements ( $\Delta\phi \times \Delta\eta$ )	$2\pi \times 5.0$	$2\pi \times 5.0$
Solid angle for energy measurements ( $\Delta\phi \times \Delta\eta$ )	$2\pi \times 9.6$	$2\pi \times 9.6$
Total cost (MCHF)	550	550

- an innermost layer devoted to the inner trackers, bathed in a solenoidal magnetic field and measuring the directions and momenta of all possible charged particles emerging from the interaction vertex;
- an intermediate layer consisting of electromagnetic and hadronic calorimeters absorbing and measuring the energies of electrons, photons and hadrons;
- an outer layer dedicated to the measurement of the directions and momenta of high-energy muons escaping from the calorimeters.

To complete the coverage of the central part of the experiments (often called barrel), so-called end-cap detectors (calorimetry and muon spectrometers) are added on each side of the barrel cylinders.

The sizes of ATLAS and CMS are determined mainly by the fact that they are designed to identify most of the very energetic particles emerging from the proton-proton collisions and to measure as efficiently and precisely as feasible their trajectories and momenta. The interesting particles are produced over a very wide range of energies (from a few hundred MeV to a few TeV) and over the full solid angle. They need therefore to be detected down to very small polar angles ( $\theta$ ) with respect to the incoming beams (a fraction of a degree, corresponding to pseudorapidities  $\eta$  of up to 5, where  $\eta = -\log[\tan(\theta/2)]$ ; pseudorapidity is more commonly used at hadron colliders because the rates for most hard-scattering processes of interest are constant as a function of  $\eta$ ). Most of the energy of the colliding protons is however dissipated in shielding and collimators close to the focussing quadrupoles (on each side of the experimental caverns, which house the experiments). The overall radiation levels will therefore be very high: many components in the detectors will become activated and will require special handling during maintenance, particularly near the beams.

For all the above reasons, both experiments have been designed following similar guiding principles:

- No particle of interest should escape unseen (except neutrinos, which will therefore be identified because their presence will cause an imbalance in the energy-momentum conservation laws governing the interactions measured in the experiments). The consequences of this simple statement are profound and far-

reaching when one goes beyond simple sketches and simulations to the details of the real experiment:

- successful operation of detectors able to measure the energies of particles with polar angles as small as one degree with respect to the incoming beams has required quite some inventiveness in material technology and a lot of detailed validation work to qualify the so-called forward calorimeters in terms of the very large radiation doses and particle densities encountered so close to the beams. Similar issues have been addressed of course very early on for the trackers, the main concerns being damage to semi-conductors (sensors and integrated circuits) and ageing of gaseous detectors. Even the muon detectors, to the initial surprise of the community, were confronted with irradiation and high-occupancy issues from neutron-induced cavern backgrounds pervading the whole experimental area;
  - avoiding any cracks in the acceptance of the experiment (especially cracks pointing back to the interaction region) has been a challenge of its own in terms of minimising the thickness of the LAr cryostats in ATLAS and of properly routing the large number of cables required to operate the ATLAS and CMS inner trackers;
  - if no particle can escape from the large volumes occupied by the experiments, then it becomes very hard for human beings to enter for rapid maintenance and repair. The access and maintenance scenarios for both experiments are quite complex and any major operation will only be feasible during long shutdowns of the accelerators. The detector design criteria have therefore become close to those required for space applications in terms of robustness and reliability of all the components.
- The high particle fluxes and harsh radiation conditions prevailing in the experimental areas have forced the collaborations to foresee redundancy and robustness for the measurements considered to be most critical. A few of the most prominent examples are described below:
    - CMS has chosen the highest possible magnetic field (4 T) combined with an inner tracker consisting solely of Silicon pixel detectors (nearest to the interaction vertex) and of Silicon microstrip detectors providing very high granularity at all radii. The occupancy of these detectors is below 2–3% even at the LHC design luminosity and the impact of pile-up is therefore minimal;
    - ATLAS has invested a very large fraction of its resources into three superconducting toroid magnets and a set of very precise muon chambers, constantly monitored with optical alignment devices, to measure the muon momenta very accurately over the widest possible coverage ( $|\eta| < 2.7$ ) and momentum range (4 GeV to several TeV). This system provides a stand-alone muon momentum measurement of sufficient quality for all benchmark physics processes up to the highest luminosities envisaged for the LHC operation;



- Both experiments rely on a versatile and multi-level trigger system to make sure the events of interest can be selected in real time at the highest possible efficiency.
- Efficient identification with excellent purity of the fundamental objects arising from the hard-scattering processes of interest is as important as the accuracy with which their four-momenta can be determined. Electrons and muons (and to a lesser extent photons and  $\tau$ -leptons with their decay products) provide excellent tools to identify rare physics processes above the huge backgrounds from hadronic jets. The requirements at the LHC are far more difficult to meet than at the Fermilab Tevatron: for example, at a transverse momentum of 40 GeV, the electron to jet production ratio decreases from almost  $10^{-3}$  at the Tevatron to a few  $10^{-5}$  at the LHC, because of the much larger increase of the production cross section for QCD hadronic jets than for W and Z bosons.

For reasons of size, cost and radiation hardness, both experiments have limited the coverage of their lepton identification and measurements to the approximate pseudorapidity range  $|\eta| < 2.5$  (or a polar angle of  $9.4^\circ$  with respect to the beams). The implementation of these requirements has also had a very large impact on the design and technology choices of both experiments:

- the length of the ATLAS and CMS super-conducting solenoids has been largely driven by the choices made for the lepton coverage;
- ATLAS has chosen a variety of techniques to identify electrons, based first and foremost on the electromagnetic calorimeter with its fine segmentation along both the lateral and longitudinal directions of shower development, then on energy-momentum matching between the calorimeter energy measurement and the inner tracker momentum measurement, but enhanced significantly over most of the solid angle by the transition radiation tracker ability to separate electrons from charged pions. In contrast, CMS relies on the fine lateral granularity of its crystal calorimeter and on the energy-momentum matching with the inner tracker;
- CMS has privileged the accuracy of the electron energy measurement with respect to the identification power with their choice of crystal calorimetry. The intrinsic resolution of the CMS electromagnetic (EM) calorimeter is superb with a stochastic term of 3–5.5% (see Sect. 16.8.2.1 for quantitative plots illustrating the performance) and the electron identification capabilities are sufficient to extract the most difficult benchmark processes from the background even at the LHC design luminosity.
- The overall trigger system of the experiments must provide a total event reduction of about  $10^7$  at the LHC design luminosity, since the number of inelastic proton-proton collisions will occur at a rate of about  $10^9$  Hz, whereas the storage capabilities will correspond to approximately 100 Hz for an average event size of 1–2 MBytes. Even today's state-of-the art technology is however far from approaching the performance required for taking a trigger decision in the very small amount of time between successive bunch crossings (25 ns).

The first level of trigger (or L1 trigger) in the ATLAS and CMS experiments is based on custom-built hardware extracting as quickly as possible the necessary information from the calorimeters and muon spectrometer and provides a decision in 2.5 to 3  $\mu\text{s}$ , during which most of the time is spent in signal transmission from the detector (to make the trigger decision) and to the detector (to propagate this decision back to the front-end electronics). This reduces the event rate to about 100 kHz with a very high efficiency for most of the events of interest for physics analysis. During this very long (for relativistic particles) time, the hundreds of thousands of very sensitive and sophisticated radiation-hard electronics chips situated throughout the detectors have to store the successive waves of data produced every 25 ns in pipelines and keep track of the time stamps of all the data so that the correct information can be retrieved when the decision from the L1 trigger is received. The synchronisation of a vast number of front-end electronics channels over very large volumes has been a major challenge for the design of the overall trigger and timing control of the experiments.

## 16.2.2 Magnet Systems


The magnet systems of the ATLAS and CMS experiments [14] were at the heart of the conceptual design of the detector components and they have driven many of the fundamental geometrical parameters and of the broad technology choices for the components of the detectors. The large bending power required to measure muons of 1 TeV momentum with a precision of 10% has led both collaborations to choose superconducting technology for their magnets to limit the size of the experimental caverns and the overall costs. The choice of magnet system for CMS was based on the elegant idea of fulfilling at the same time with one magnet a high magnetic field in the tracker volume for all precision momentum measurements, including muons, and a high enough return flux in the iron outside the magnet to provide a muon trigger and a second muon momentum measurement for the experiment. This is achieved with a single solenoid of a large enough radius to contain most of the CMS calorimeter system. In contrast, the choice of magnet system for ATLAS was driven by the requirement to achieve a high-precision standalone momentum measurement of muons over as large an acceptance in momentum and  $\eta$ -coverage as possible. This is achieved using an arrangement of a small-radius thin-walled solenoid, integrated into the cryostat of the barrel electromagnetic calorimeter, surrounded by a system of three large air-core toroids, situated outside the ATLAS calorimeter systems and generating the magnetic field for the muon spectrometer. The main parameters of these magnet systems are listed in Table 16.3 and their stored energies are compared to those of previous large-scale magnets in high-energy physics experiments in Fig. 16.7.

In CMS, the length of the solenoid was driven by the need to achieve excellent momentum resolution over the required  $\eta$ -coverage and its diameter was chosen such that most of the calorimetry is contained inside the coil. In ATLAS, the



**Table 16.3** Main parameters of the CMS and ATLAS magnet systems

Parameter	CMS solenoid	ATLAS		
		Solenoid	Barrel toroid	End-cap toroids
Inner diameter	5.9 m	2.4 m	9.4 m	1.7 m
Outer diameter	6.5 m	2.6 m	20.1 m	10.7 m
Axial length	12.9 m	5.3 m	25.3 m	5.0 m
Number of coils	1	1	8	8
Number of turns per coil	2168	1173	120	116
Conductor size [mm <sup>2</sup> ]	64 × 22	30 × 4.25	57 × 12	41 × 12
Bending power	4 T · m	2 T · m	3 T · m	6 T · m
Current	19.5 kA	7.6 kA	20.5 kA	20.0 kA
Stored energy	2700 MJ	38 MJ	1080 MJ	206 MJ



**Fig. 16.7** Ratio of stored energy over mass,  $E/M$ , versus stored energy,  $E$ , for various magnets built for large high-energy physics experiments

position of the solenoid in front of the barrel electromagnetic calorimeter has demanded a careful optimisation of the material in order to minimise its impact on the calorimeter performance and its length has been defined by the design of the overall calorimeter and inner tracker systems, leading to significant non-uniformity of the field at the end of the tracker volume.

The main advantages and drawbacks of the chosen magnet systems can be summarised as follows, considering successively the inner tracker, calorimeter and muon system performances (see Sect. 16.8):

- the higher field strength and uniformity of the CMS solenoid provide better momentum resolution and better uniformity over the full  $\eta$ -coverage for the inner tracker;

- the position of the ATLAS solenoid just in front of the barrel electromagnetic calorimeter limits to some extent the energy resolution in the region  $1.2 < |\eta| < 1.5$ ;
- the position of the CMS solenoid outside the calorimeter limits the number of interaction lengths available to absorb hadronic showers in the region  $|\eta| < 1$ ;
- the muon spectrometer system in ATLAS provides an independent and high-accuracy measurement of muons over the full  $\eta$ -coverage required by the physics. This requires however an alignment system with specifications an order of magnitude more stringent (few tens of  $\mu\text{m}$ ) than those of the CMS muon spectrometer. In addition, the magnetic field in the ATLAS muon spectrometer must be known to an accuracy of a few tens of Gauss over a volume of close to  $20,000\text{ m}^3$ . The software implications of these requirements are non-trivial (size of map in memory, access time);
- the muon spectrometer system in CMS has limited stand-alone measurement capabilities and this affects the triggering capabilities for the luminosities envisaged for the LHC upgrade.

In terms of construction, the magnet systems have each turned out to be a major project in its own right with very direct and strong involvement from the Technical Coordination team [15] and from major national laboratories and funding agencies. A detailed account of the construction of these magnets is beyond the scope of this review and this section can be concluded by simply stating that during the course of the past few years, all these magnets have undergone very successfully extensive commissioning steps, sustained operation at full current, in particular for cosmic-ray data-taking in 2008/2009, and stable operation with beam in the LHC machine at the end of 2009.

### 16.2.2.1 Radiation Levels

At the LHC, the primary source of radiation at full luminosity comes from collisions at the interaction point. In the tracker, charged hadron secondaries from inelastic proton-proton interactions dominate the radiation backgrounds at small radii while further out other sources, such as neutrons, become more important. Table 16.4 shows projected radiation levels in key areas of the detector.

In ATLAS, most of the energy from primaries is dumped into two regions: the TAS (Target Absorber Secondaries) collimators protecting LHC quadrupoles and the forward calorimeters. The beam vacuum system spans the length of the detector and in the forward region is a major source of radiation backgrounds. Primary particles from the interaction point strike the beam-pipe at very shallow angles, such that the projected material depth is large. Studies have shown that the beam-line material contributes more than half of the radiation backgrounds in the muon system. The deleterious effects of background radiation fall into a number of general categories: increased background and occupancies, radiation damage and ageing of

**Table 16.4** The 1 MeV neutron equivalent fluence ( $F_{\text{neq}}$ ) and doses in key areas of the ATLAS detector after  $500\text{ fb}^{-1}$  of data (estimated to be approximately 7 years of operation at design luminosity)

Inner detector					
Location	$F_{\text{neq}}$		Dose		Charged-particle flux above 10 MeV [Hz/cm <sup>2</sup> ]
	[10 <sup>14</sup> cm <sup>−2</sup> ]		[kGy]		
Pixel layer 0	13.5		790		40 · 10 <sup>6</sup>
SCT layer 1	0.8		38		1.5 · 10 <sup>6</sup>
SCT disk 9	0.6		23		10 <sup>6</sup>
TRT outer radius	0.25		3.5		10 <sup>5</sup>
Calorimeters					
Location	η		Maximum dose [kGy]		
EM barrel	1.475		1.2		
EM end-cap	3.2		150		
Tile	1.2		0.15		
HEC	3.2		30		
FCal	4.9		1000		
Muon spectrometer					
Location	Flux				Single-plane rates [Hz/cm <sup>2</sup> ]
	[kHz/cm <sup>2</sup> ]		[Hz/cm <sup>2</sup> ]		
	<i>n</i>	<i>γ</i>	<i>μ</i>	<i>p</i>	
Barrel chambers	2.6–4.0	1.0–1.5	0.3–4.5	0.4–3.2	6.0–11.0
Inner edge of inner wheel	79	25	21	64	347
Inner edge of outer wheel	2.7	1.5	3	0.9	12

Also given are the charged-particle fluxes in the tracker and fluxes and single-plane rates in the muon spectrometer

detector components and electronics, single-event upsets and single-event damage, and creation of radionuclides which will impact access and maintenance scenarios.

### 16.2.3 Lessons Learned from the Construction Experience

It is fair to say that most of the physicists and engineers involved in the ATLAS and CMS construction were faced with a challenge of this scope and size for the first time. It seems therefore appropriate to put some emphasis in this article on the lessons learned from the construction of these detectors. This section describes the general lessons learned and the next sections will give more explicit examples in many cases when describing the experience from the construction of the detector components.

The lessons learned are of varying nature, many are organisational, many are technical and some are sociological. Some are specific to the LHC, some are specific to the way international high-energy physics collaborations work, and some are of a

general enough nature that they might well apply to any complex high-tech project of this size. It is therefore hard to classify them in a clear logical order, and this review has attempted to rank them from the general and common to the specific and unique to the LHC.

### 16.2.3.1 Time-Scales, Project Phases and Schedule Delays

If there has been one lesson learned from the days in the early 1990s when ATLAS and CMS came into being as detector concepts, it is certainly that the research and development phase of projects of this complexity are impossible to plan with real certainty about the time-scales involved. Modern tools for project management are of little help here because the vagaries of the initial phase do not generally obey the simple laws of project schedules and charts. These can be a posteriori explained of course:

- the research and development phase for new high-tech detector elements, such as radiation-hard silicon sensors and micro-electronics, crystals grown from a new material, large-scale electrodes for operation at high voltage in liquid Argon, etc., will always be a phase to which one has to allocate as much time as feasible within the overall project schedule constraints. The justification for this is basically that the potential rewards are enormous, as was exemplified by the late but striking success of the deep sub-micron micro-electronics chips pioneered by CMS and now used throughout all LHC experiments, and by the late but successful operation of CMS PbWO<sub>4</sub> crystals with their avalanche photodiode readout and associated electronics. Making the appropriate research and development choices at the right time will however always remain a challenge for any new project of this scope and complexity.
- less known to many colleagues in our community is the phase during which the components for producing complex detector modules are launched for manufacturing in industry. This phase can indeed be planned correctly if the required physicist/engineering experience is available, if the funding allows for multiple suppliers to mitigate potential risks, and if the physicists agree quickly to moderate their usually very demanding specifications to adapt them to the actual capabilities of industry.

Experience has shown however that success was far from guaranteed in this phase, with causes for delays or outright initial failures ranging from being forced to award contracts to the lowest bidder, to incomplete technical specifications, to handling and packaging issues during manufacturing, particularly for polyimide-based products, of which there are many thousands of m<sup>2</sup> in both experiments. This material shows up under various forms (especially in flexible printed circuit boards for various applications) and is a basic insulating material with excellent electrical and mechanical properties, with very high tolerance to radiation, but unfortunately also with a high propensity to absorb moisture and thereby lead to unexpected changes in even the course of a well-defined manufacturing process.

Serious technical problems in this area have affected the manufacturing schedule of major components of both experiments (hybrids for semi-conductor detectors, flexible parts of printed-circuit boards, large-size electrodes for electromagnetic calorimetry), but other issues such as welding, brazing and general integrity and leak-tightness of thin-walled cooling pipes have also been a concern for several of the components in each experiment.

In addition, several of the more significant contracts were seriously affected by changes in the industrial boundary conditions (insolvency, change of ownership). The recommended purchasing strategy of having multiple suppliers for large contracts, to minimise the consequences from a possible failure in the case of a single supplier, has not always been the optimal one (high-quality silicon sensors are perhaps the most prominent example).

The detailed construction planning can be consulted in the various Technical Design Reports (TDR), most of which were submitted from 1996 to 1998 to seek approval for construction of the major detector components. This called for completion of this construction phase by mid-2001 to mid-2003. At the time when a big schedule and financial crisis shook the LHC project in fall 2001 (see below), it was already clear that many detector components would not be on schedule by a significant margin.

The 2-year delay in the completion of the accelerator resulting from this crisis was also needed by the experiments, as can be seen from Table 16.5, which illustrates the major construction milestones originally planned at the time of the TDRs and actually achieved. When trying to assess the significance of the differences between the dates achieved for the delivery of major components of

**Table 16.5** Main construction milestones for the ATLAS and CMS detectors

Detector system	ATLAS		CMS	
	TDR	Actual	TDR	Actual
Pixels	06/03	03/07	03/05	12/07
Silicon micro-strips (barrel)	12/02	07/05	03/04	10/06
Silicon micro-strips (end-caps)	12/02	06/06	03/04	10/06
Transition radiation tracker	03/04	12/05		
Electromagnetic calorimeter (barrel)	06/03	07/04	12/03	03/07
Electromagnetic calorimeter (end-caps)	01/04	09/05	06/04	03/08
Hadronic calorimeter	12/02	02/04	12/03	12/04
Muon chambers	12/04	12/05	12/03	06/06
Solenoid magnet	01/02	09/01	03/03	12/05
Barrel toroid magnet	06/02	06/05		
End-cap toroid magnet	12/03	11/06		

Shown are the milestones for the delivery of major components to CERN, as planned at the time of the Technical Design Reports (TDR), and the actual delivery milestones achieved

the experiments and those planned 9 years ago, it is important to remember the prominent events, at CERN and within the collaborations, which happened during these years:

- at the time of the submission of the various TDRs for ATLAS and CMS, the construction and installation schedule was worked out top-down, based on a ready-for-operation date of summer 2005 for the LHC machine and the experiments;
- in 1999, the CMS collaboration decided to replace the micro-strip gas chamber baseline technology for the outer part of their Inner Detector by “low-cost” silicon micro-strip detectors. This is probably the most outstanding example of decisions, which the collaborations had to take after the TDRs were submitted and which have affected the construction schedule in a major way;
- in 2001, when the CERN laboratory management announced significant cost overruns, mostly in the machine, but also in the ATLAS and CMS experiments, it also announced a 2-year delay in the schedule for the machine, which obviously led to a readjustment of the construction and installation schedule of the experiments. By that time, both in ATLAS and CMS, the Technical Co-ordination teams had worked out a realistic installation schedule, which still needed to be fleshed out substantially in areas such as services installation, commissioning of ancillary equipment for operation of the huge devices to be operated underground, etc.;
- the ATLAS experimental cavern was delivered more or less on time in spring 2003, whereas the CMS experimental cavern suffered considerable delays and was delivered only towards the end of 2004.

### 16.2.3.2 Physicists and Engineers: How to Strike the Right Balance?

This is a very delicate issue because there exists no precise recipe to solve this problem. The ATLAS and CMS experiments were born from the dreams of physicists but are based today on the calculations and design efforts from some of the best teams of engineers and designers in the world. One should not forget that, originally (in 1987), even the physicists thought that only a muon spectrometer behind an iron dump was guaranteed to survive the irradiation and that most tracking technologies were doomed at the highest luminosities of the LHC [16].

Although a strong central and across-board (from mechanics to electronics, controls and computing) engineering effort would have been desirable from the very start (i.e. around 1993), a standard centralised and very systematic engineering approach alone, as is frequently used in large-scale astronomy projects, could not have been used for several reasons:

- the cost would have been prohibitive;
- only the physicists can actually make the sometimes difficult choices and decisions when faced with problems requiring certain heart-wrenching changes in the fundamental parameters of the experiment (number of layers in the

tracking detectors, number of cells in the electromagnetic calorimeter, overall strength and uniformity of the magnetic field, etc.). The number of coils to be constructed in the ATLAS superconducting toroid and the peak field of the CMS central solenoid are two examples of early and fundamental parameters of the experiments, which were studied for quite some time and had a significant bearing on the overall cost of the experiments;

- some of the usual benefits of such an approach, such as optimised production costs for repetitive manufacturing of the same product, are not there to be reaped when considering the experiments as a whole rather than looking at individual components, such as the micro-strip silicon modules, which number in many thousands and did indeed benefit in many aspects from a systematic engineering approach;
- the overall technological scope of these nascent experiments required creativity and novel approaches in areas as far apart as 3D-calculations of magnetic fields and forces over very large volumes containing sometimes unspecified amounts of magnetic materials and radiation-dose and neutron-fluence calculations of unprecedented complexity in our field to evaluate the survival of a variety of objects, from the basic materials themselves to complex micro-electronics circuits. Only a well-balanced mix of talented and dedicated designers, engineers and physicists could have tackled such issues with any chance of success;
- the decision-making processes in our community cannot be too abrupt. Consensus needs to be built, especially between physicists but also between engineers from sometimes widely different cultures and backgrounds.

In retrospect, however, there has emerged as a clear lesson, that the management of the experiments should have evolved at an earlier stage the decision-making process from a physicist-centric one at the beginning, when little was known about the detailed design of all the components, to a more engineer-centric one, as the details were fleshed out more and more. Establishing engineering envelopes and assembly drawings for the different systems, routing the very large and diverse amount of services needed to operate complex detectors distributed everywhere across the available space, and designing, validating and procuring common solutions for many of the electronics and controls components are examples, which clearly illustrate this need. The collaborations have indeed encountered difficulties to recognise such needs and to react to them at the appropriate moment in time.

### **16.2.3.3 International and Distributed: A Strength or a Weakness?**

ATLAS and CMS are truly international and distributed collaborations, even if the engineering and/or manufacturing of some of the major components of both experiments have been entrusted to large laboratories situated all across the world. Modern technology (web access to document servers, video-conferencing facilities, more uniform standards, such as the use of the metric system, for drawings, specifications and quality assurance methods, electronic reporting tools) has been instrumental in

improving the efficiency of the various strands of these collaborations, an admittedly weak point of such organisations. There are two major weaknesses intrinsic to collaborations structured as ATLAS and CMS with distributed funding resources:

- one is that it is not simple to converge on the minimum required number of technologies once the research and development phase is over. One example of perhaps unnecessary multiplication of technologies are the precision chambers in the ATLAS muon spectrometer, where the highest- $\eta$  part of the measurements are covered by cathode strip chambers rather than the monitored drift tube technology used everywhere else. A similar example can be found in the CMS muon spectrometer, which is also equipped with two different chamber technologies in the barrel and end-cap regions (see Sect. 16.5).
- the decision-making process is sometimes skewed by the difficulty of conveying a global vision of the best interests of the project, which should be weighed against the more localised and focussed interests of particular funding agencies, some of which operate within a rather inflexible legal framework.

The strengths of this international and distributed approach far outweigh however its deficiencies over a much more centralised one, such as that adopted for the Super-Conducting Super Collider with a centralised funding and management in Waxahachie (Texas) about 15 years ago:

- the flexibility achieved has often provided solutions to the inevitable problems, which have shown up during the design and construction phase. Whenever a link in the chain was shown to falter or even to be totally missing, the collaboration has often been able to find alternate solutions. If a large laboratory had difficulties in meeting a complex technological challenge alone because of limitations in funding and human resources, other laboratories with similar expertise could be sought out and integrated into the effort with minimal disruption. If the production line for certain detectors did not churn out the required number of modules per unit time because of yield issues or of an underestimate of the human resources required, other production lines, often on different continents with cheaper labour costs, were launched and operated successfully.
- many concrete examples have shown that motivation and dedication to the project go together with the corresponding responsibilities, both technical and managerial. It is worthwhile also to note here that it surely would have been beneficial for the overall LHC project if the management of the ATLAS and CMS experiments would have been integrated as a real partner into the CERN management structure at the highest level right from the beginning. Both experiments were severely handicapped by a cost ceiling without contingency defined top-down more than 10 years ago.

It is fair to say that, without the motivation and dedication of many of our colleagues all over the world, who fought and won their own battles at all required levels (technical, funding, human resources, organisational), and of their funding agencies, the construction of ATLAS and CMS would not have reached its astounding and successful completion with only small parts of each experiment



deferred. Dealing with significant deferrals has always been damaging to the atmosphere of large collaborations of this type and the fact that both experiments are now essentially complete should certainly be attributed to the credit of all their participants.

A particular mention should go here to our Russian colleagues, who have not only strongly contributed intellectually to the experiments, as all the others, from the very beginning, but who also staffed continuously, together with other Eastern European colleagues and also colleagues from Asia, a very large fraction of the teams needed to assemble, equip, test and commission the major detector components. This was quite striking during the installation period from just listening to the conversations occurring in the lifts bringing people and equipment up and down the experimental shafts.

- the concept of deliverables has also turned out to the advantage of the projects. Each set of institutes in each country have been asked to deliver a certain fraction of specific components of the detector systems, ranging from a modest (but critical!) scope, such as the fabrication of the C-fibre cylinders for the barrel semi-conductor tracker in ATLAS, to a very large (and very visible to the whole collaboration!) scope, such as the CMS crystal production in several commercial companies, or as the ATLAS super-conducting solenoid built in Japanese industry, in close collaboration with institutes from the same country, which are full-fledged members of the collaboration.

This concept has certainly maximised the overall funding received by ATLAS and CMS, because each funding agency has to a certain extent been asked and has agreed to take responsibility for the delivery of certain detector components without assigning to these a specific cost, since the real costs vary from country to country, and even the ratios of costs between different countries inevitably vary, because of the approximately uniform costs of raw materials as compared to the wildly differing costs of skilled and unskilled labour. Since the infrastructure of the experiments is a mixture of low and high technology components, most participating countries have in the end been able to contribute efficiently in kind to the common projects of interest to the whole collaboration.

- the scheme based on deliverables rather than raw funding could not have worked however without being completed by a sizable set of common projects, to which the funding agencies had to contribute, either through funds to be handled by the management of the experiments, either through in-kind contributions, the cost of which was determined in the context of the same scheme as for the deliverables. Examples of these common projects are the magnets of both experiments, the LAr cryostats and cryogenics of ATLAS, and much of the less high-tech infrastructure components of both experiments.
- finally, the computing operations of the experiments and the analysis of the data taken over the next 10 years do and will require a very distributed and international style of working also. This is not really new to our community, it is just of an unprecedented scale in size and duration. The collaborations are evolving now from an organisational model focussed initially on research and development and then on construction to a new model, which is focussed more on

detector operation, monitoring of the data quality and data preparation, leading to the analysis work required to understand precisely the behaviour of the detectors and extract as efficiently as possible the exciting physics ahead of us. The years spent together and the difficulties overcome over a 15-year long period of design and construction have certainly cemented the collaborations in a spirit of respect and mutual understanding of all their diverse components. This will surely turn out to be an excellent preparation for the forthcoming challenges when faced with real experimental data.

#### **16.2.3.4 A Well Integrated and Strong Technical Co-ordination Team**

It is clear that without such a team the experiments would most probably have faced insurmountable construction delays and integration problems. The Technical Co-ordination team must in a sense be perceived as the strong backbone of the experiment by all the physicists in the community. This was indeed the case in the installation phase of the experiments, at a time when it had to smoothly execute a complex suite of integration and installation operations for detector components arriving from all over the world. But this was less the case 10–15 years ago, at a time when the physicists and engineers in this team were sometimes perceived as a nuisance disrupting the delicate balance of the collaboration and were criticised in different ways:

- many physicists and engineers had great trouble when asked to specify all the details of cables, pipes and connectors, at a very early time (15 years ago) when they were desperately trying to move into mass production;
- strong resistance to reviews was encountered, based on partially correct, but also partially fallacious, arguments that all the expertise in a given area was already available in the project under review;
- the multiplicity of reviews also caused sometimes considerable friction and frustration, especially since an overall co-ordination between funding agency reviews and internal project reviews was almost impossible to put into place.

In retrospect, these reviews are indeed necessary, whether or not all of their recommendations and outcomes have turned out to be of a specific concrete usefulness, because they have usually forced the project teams to collect documentation, take stock, step back and think about issues sometimes obscured by the more immediate and pressing problems at hand.

Although the construction of the individual detector components can be argued to have been quite successful under the umbrella of deliverables and in the absence of a fully centralised management of the experiment resources, there are obviously a variety of tasks, which have to be solved by a strong centralised team of designers, engineers and physicists. As in any such process, this team is much better accepted if it is built up at least partially from people within the collaboration, who are already well integrated in and known to the collaboration. Despite all the grumbling and

moaning, the efforts of the Technical Co-ordination team have been crucial to the success of the ATLAS and CMS projects:

- finding common (often commercial) solutions does not come easily to large numbers of inventive and often opinionated physicists. Common solutions across the experiments are even harder to achieve, although they have turned out to be profitable to all parties in a number of areas. Clearly the strong research and development programme launched in 1989 by CERN for the development of the LHC detector technologies has been a key element in the definition of the various detector concepts (radiation-hard silicon detectors and electronics, electromagnetic and hadronic calorimetry, various tracking technologies, etc.).

In the areas where such common (often commercial) solutions have been adopted in many cases in the past, the successes of the research and development programme have been less spectacular (data transmission, specialised trigger processors, various offline software developments), most probably because the solutions emerging today were not easy to predict from the technology trends of 20 years ago, when the worldwide web, mobile phones, inexpensive desktop computing and high-speed networks did not exist.

The Technical Co-ordination team has certainly been very instrumental in encouraging the collaboration to adopt common technical solutions and has also delegated to the appropriate persons in the collaboration the mandate to negotiate and agree these common solutions across the experiments: the frame contracts with major micro-electronics suppliers, the gas systems, the power supplies, the electronics crates and racks and the slow controls infrastructure hardware and database software can be quoted as some of the more prominent examples.

- establishing a strong quality assurance and review process across the whole collaboration is a must at an early stage in such complex projects, where standard commercial products have often failed, sometimes for multiple reasons owing to the boundary conditions in the experimental caverns (radiation background and magnetic field).

As stated above, the review process (from conceptual engineering design reviews, to production readiness and production advancement reviews) can be very beneficial and even well accepted within the collaboration if it is kept lightweight and perceived as executed by people involved in the project as all the others rather than by an elite breed of top-level managers.

Most of the ATLAS and CMS Technical Design Reports quoted as references in this review address quality assurance with ambitions and specifications, which are fully justified on paper but much harder to implement in reality when facing time pressure and the inevitable lack of human resources to fulfill every aspect of the task. In relation to industry in particular, the effort required in monitoring production of delicate components had been totally underestimated or even ignored in the design phase. The reviews put in place by the Technical Co-ordination team have played an important role in keeping all aspects related to schedule, resources and quality assurance under control during the detector construction. They have also ensured that large groups with significant project

responsibilities were not allowed to operate for too long in a stand-alone mode without synchronising with and reporting back to Technical Co-ordination, the management of the experiments and the collaboration at large. The risks involved in letting things go astray too much are simply unacceptable for projects of this complexity and size.


- As stated above, one weakness perhaps of the multiple dimensions under which ATLAS and CMS are viewed is that the funding agencies have often conducted their own necessary review processes in a way largely decoupled from the review process operated by the management of the experiments. This weakness stems from the lack of central control of expenditures because of the distributed funding and spending responsibilities. This can obviously lead to inefficiencies in the actual execution of the project and, worse, sometimes to conflicting messages given to the institutes concerning priorities, since those of a given funding agency may not always coincide with those of the experiment. The common funds necessary to the construction of significant components of the experiments, such as magnets, infrastructure, shielding, cryostats, etc., are a prominent example which comes to mind, when assessing which of the components of the experiments had the most difficulty in dealing with the multi-threaded environment, in which the detector construction has been achieved.

Finally, it is in the very recent phase of assembly, installation and commissioning of the ATLAS and CMS detectors that the enormous efforts and contribution from the Technical Co-ordination teams have been most visible: they have had to organise the vast teams of sub-contractors and specialised personnel from the collaborating institutes and they have had to deal with the daily burden of making sure all the tasks were executed as smoothly as possible with safety as one of the paramount requirements.

## 16.3 Inner Tracking System

### 16.3.1 Introduction

The ATLAS tracker is designed to provide hermetic and robust pattern recognition, excellent momentum resolution and both primary and secondary vertex measurements [17] for charged tracks above a given  $p_T$  threshold (nominally 0.5 GeV, but as low as 0.1 GeV in some ongoing studies of initial measurements with minimum-bias events) and within the pseudorapidity range  $|\eta| < 2.5$ . It also provides electron identification over  $|\eta| < 2.0$  and a wide range of energies (between 0.5 and 150 GeV). It is contained within a cylindrical envelope of length  $\pm 3512$  mm and of radius 1150 mm, within the solenoidal magnetic field of 2 T. Figures 16.8 and 16.9 show the sensors and structural elements traversed by 10 GeV tracks in respectively the barrel and end-cap regions.




**Fig. 16.8** Drawing showing the sensors and structural elements traversed by a charged track of 10 GeV  $p_T$  in the ATLAS barrel inner detector ( $\eta = 0.3$ ). The track traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of  $50 \times 400 \mu\text{m}^2$ , the four cylindrical double layers (one axial and one with a stereo angle of 40 mrad) of barrel silicon-microstrip sensors (SCT) of pitch  $80 \mu\text{m}$ , and approximately 36 axial straws of 4 mm diameter contained in the barrel transition-radiation tracker modules within their support structure

The ATLAS tracker consists of three independent but complementary sub-detectors. At inner radii, high-resolution pattern recognition capabilities are available using discrete space-points from silicon pixel layers and stereo pairs of silicon micro-strip (SCT) layers. At larger radii, the transition radiation tracker (TRT) comprises many layers of gaseous straw tube elements interleaved with transition radiation material. With an average of 36 hits per track, it provides continuous tracking to enhance the pattern recognition and improve the momentum resolution over  $|\eta| < 2.0$  and electron identification complementary to that of the calorimeter over a wide range of energies.

Table 16.6 lists the main parameters of the ATLAS tracker:

- the radial position of the innermost measurement is essentially determined by the outer diameter of the beam pipe, which has been manufactured using expensive and delicate Beryllium material over an overall length of 7 m. The active part of the tracker has a half-length of 280 cm, slightly longer than that of its solenoid, resulting in significant field non-uniformities and momentum resolution degradation at each end.



**Fig. 16.9** Drawing showing the sensors and structural elements traversed by two charged tracks of 10 GeV  $p_T$  in the ATLAS end-cap inner detector ( $\eta = 1.4$  and  $2.2$ ). The end-cap track at  $\eta = 1.4$  traverses successively the beryllium beam-pipe, the three cylindrical silicon-pixel layers with individual sensor elements of  $50 \times 400 \mu\text{m}^2$ , four of the disks with double layers (one radial and one with a stereo angle of 40 mrad) of end-cap silicon-microstrip sensors (SCT) of pitch  $\sim 80 \mu\text{m}$ , and approximately 40 straws of 4 mm diameter contained in the end-cap transition radiation tracker wheels. In contrast, the end-cap track at  $\eta = 2.2$  traverses successively the beryllium beam-pipe, only the first of the cylindrical silicon-pixel layers, two end-cap pixel disks and the last four disks of the end-cap SCT. The coverage of the end-cap TRT does not extend beyond  $|\eta| = 2$

**Table 16.6** Main parameters of the ATLAS tracker system

Item		Radial extension [mm]	Length [mm]
Overall tracker envelope		$0 < R < 1150$	$0 <  z  < 3512$
Beam-pipe		$29 < R < 36$	
Pixel	Overall envelope	$45.5 < R < 242$	$0 <  z  < 3092$
3 cylindrical layers	Sensitive barrel	$50.5 < R < 122.5$	$0 <  z  < 400.5$
$2 \times 3$ disks	Sensitive end-cap	$88.8 < R < 149.6$	$495 <  z  < 650$
SCT	Overall envelope	$255 < R < 549$ (barrel)	$0 <  z  < 805$
		$251 < R < 610$ (end-cap)	$810 <  z  < 2797$
4 cylindrical layers	Sensitive barrel	$299 < R < 514$	$0 <  z  < 749$
$2 \times 9$ disks	Sensitive end-cap	$275 < R < 560$	$839 <  z  < 2735$
TRT	Overall envelope	$554 < R < 1082$ (barrel)	$0 <  z  < 780$
		$617 < R < 1106$ (end-cap )	$827 <  z  < 2744$
73 straw planes	Sensitive barrel	$563 < R < 1066$	$0 <  z  < 712$
160 straw planes	Sensitive end-cap	$644 < R < 1004$	$848 <  z  < 2710$

- the total power required for the tracker front-end electronics will increase from approximately 62 to 85 kW from initial operation to high-luminosity operation after irradiation. Bringing this amount of power to the detector requires large amounts of copper; the resulting heat load is very uniformly distributed across the entire active volume of the tracker and has to be removed using innovative techniques (fluor-inert liquids to mitigate the risks from possible leaks, thin-walled pipes made from light metals, evaporative techniques for optimal heat removal in the case of the silicon-strip and pixel detectors). There is also considerable heat created by the detectors themselves: the silicon-strip modules will dissipate about 1 W each from sensor leakage currents at the end of their lifetime, and the highest-occupancy TRT straws dissipate about 10 mW each at the LHC design luminosity.
- for all of the above reasons, it has been well known since the early 90's in the LHC community that the material budget of the tracker systems as built would pose serious problems in terms of their own performance (see Sect. 16.8.1) and even more so in terms of the intrinsic performance of the electromagnetic calorimeter and of the overall performance for electron/photon measurements (see Sect. 16.8.2). Despite the best efforts of the community, the material budget for the tracker has risen steadily over the years and reached values of two radiation lengths ( $X_0$ ) and close to 0.6 interaction lengths ( $\lambda$ ) in the worst regions (see Sect. 16.3.2.1 for more details and plots).

The high-radiation environment imposes stringent conditions on the inner-detector sensors, on-detector electronics, mechanical structure and services. Over the 10-year design lifetime of the experiment, the pixel inner vertexing layer must be replaced after approximately 3 years of operation at design luminosity. The other pixel layers and the pixel disks must withstand a 1 MeV neutron equivalent fluence  $F_{\text{neq}}$  [18] of up to  $\sim 8 \times 10^{14} \text{ cm}^{-2}$ . The innermost parts of the SCT must withstand  $F_{\text{neq}}$  of up to  $2 \times 10^{14} \text{ cm}^{-2}$ . To maintain an adequate noise performance after radiation damage, the silicon sensors must be kept at low temperature (approximately  $-5$  to  $-10^\circ\text{C}$ ) implying coolant temperatures of  $\sim -25^\circ\text{C}$ . In contrast, the TRT is designed to operate at room temperature.

The above operating specifications imply requirements on the alignment precision which are summarised in Table 16.7 and which serve as stringent upper limits on the silicon-module build precision, the TRT straw-tube position, and the measured module placement accuracy and stability.

This leads to:

- (a) a good construction accuracy with radiation-tolerant materials having adequate detector stability and well understood position reproducibility following repeated cycling between temperatures of  $-20$  and  $+20^\circ\text{C}$ , and a temperature uniformity on the structure and module mechanics which minimises thermal distortions;
- (b) an ability to monitor the position of the detector elements using charged tracks and, for the SCT, laser interferometric monitoring [19];

**Table 16.7** Intrinsic measurement accuracies and mechanical alignment tolerances for the tracker sub-systems, as defined by the performance requirements of the ATLAS experiment

Item	Intrinsic accuracy [ $\mu\text{m}$ ]	Alignment tolerances [ $\mu\text{m}$ ]		
		Radial ( $R$ )	Axial ( $z$ )	Azimuth ( $R - \phi$ )
Pixel				
Layer-0	10 ( $R$ - $\phi$ ) 115 ( $z$ )	10	20	7
Layer-1 and Layer-2	10 ( $R$ - $\phi$ ) 115 ( $z$ )	20	20	7
Disks	10 ( $R$ - $\phi$ ) 115 ( $R$ )	20	100	7
SCT				
Barrel	17 ( $R$ - $\phi$ ) 580 ( $z$ ) <sup>a</sup>	100	50	12
Disks	17 ( $R$ - $\phi$ ) 580 ( $R$ ) <sup>a</sup>	50	200	12
TRT	130			30 <sup>b</sup>

The numbers in the table correspond to the single-module accuracy for the pixels, to the effective single-module accuracy for the SCT and to the drift-time accuracy of a single straw for the TRT

<sup>a</sup>Arises from the 40 mrad stereo angle between back-to-back sensors on the SCT modules with axial (barrel) or radial (end-cap) alignment of one side of the structure. The result is pitch-dependent for end-cap SCT modules

<sup>b</sup>The quoted alignment accuracy is related to the TRT drift-time accuracy

- (c) a trade-off between the low material budget needed for optimal performance and the significant material budget resulting from a stable mechanical structure with the services of a highly granular detector.

The design and construction of systems, capable of meeting the physics requirements and of providing stable and robust operation over many years, has been perhaps the most formidable challenge faced by the experiment because of the very harsh radiation conditions to be faced near the interaction point and of the conflicting requirements in terms of material budget between the physics and the design constraints. The latter arise mostly from the on-detector high-speed front-end electronics, which require a lot of power to be fed into a limited volume and therefore a large amount of heat to be removed from a very distributed set of local heat sources across the whole tracker.

This section describes briefly the ATLAS tracker and its main properties and discusses a few salient aspects from the construction experience and from the measured performance in laboratory and test beam of production modules in the various technologies. A few examples of the overall performance expected in the actual configuration of the experiment are presented in Sect. 16.8.1, where it is also compared to the expected performance of the CMS tracker.



## **16.3.2 Construction Experience**

### **16.3.2.1 General Aspects**

The ATLAS tracker system has evolved considerably since the submission of the Technical Proposal in 1994 and even since the corresponding Technical Design Reports in 1997/1998. The evolution was dictated by many factors, some of which have already been alluded to in Sect. 16.2.3 and some of which are related to the specific design challenges posed:

- the rapid development of radiation-hard silicon sensors and of their front-end electronics led many physicists and engineers in the community to focus for a long time on the single module scale and, as a consequence, to perhaps address some of the systems issues, especially for the readout and cooling aspects, too late.
- the legitimate concerns throughout the collaborations about the material budget of the tracker systems resulted in huge pressures on the engineering design effort in terms of materials at a very early stage. This effort has been largely successful in terms of mechanics, as can be seen from the very light and state-of-the-art structures used to support and hold the detector components in the tracker system. The already considerable experience from the space industry across the world turned out to be invaluable, including in terms of thermal behaviour and of resistance to radiation and to moisture absorption.
- the tracker macro-assemblies, once completed as operational devices, are the sum of a large number of diverse and tiny components. Many of these components were not built into the design from the very beginning and only general assumptions based on past experience were made concerning their manufacture. Several of these assumptions turned out to be incorrect: for example, the use of silver in the electrical connections and cables has had to be minimised because of activation issues. The pressure on the material budget led to the choice of risky technical solutions for cooling and power, involving hard-to-validate thin-walled Aluminium, copper/Nickel or Titanium pipes and polyimide/Aluminium tapes rather than the less risky but heavier stainless steel pipes and polyimide/copper tapes.
- many of the systems aspects were discovered as the detailed design progressed, rather than foreseen early on, and this has led to difficult retrofitting exercises and sometimes to technical solutions more complex and risky than those which would be devised from a clean slate today. Some of the substrates for the electronics of the silicon modules barely existed in terms of conceptual design at a time when the front-end electronics chip was ready for production. This is one example of a specific and critical component, which was not always incorporated into the detailed design of the system from the very beginning.


Another more general example stems from the engineering choices made for the implementation of the on-detector and off-detector cooling systems: there are as many on-detector cooling schemes and pipe material choices as there are

- detector components. The cooling systems themselves are all operating under severe space limitations on-detector and at high pressure (from 3 to 6 bars). These systems range from room-temperature monophase  $C_6F_{14}$  for the TRT to cold evaporative  $C_3F_8$  for the SCT and pixels. Many problems have been encountered during the commissioning in situ and early operation of these systems, and it is fair to say a posteriori that this is one area where a stronger and more centralised engineering effort would have probably come up with a more uniform, more robust and redundant, and less risky implementation.
- Table 16.8 shows how optimistic the estimate of the material budget of the ATLAS tracker was at the time of the Technical Proposal in 1994 and how it has evolved since then to reach the values quoted in early 2008, after completion of the installation of all of its components. These values cannot be claimed to be final yet, although most of the remaining uncertainties are small and related to the exact routing details of the various services and of patch-panels for cable and pipe connections. These are situated within the tracker volume, but not always in the fiducial region where the detectors expect to perform precision tracking and electromagnetic calorimetry measurements (for example, the patch-panels for the pixel detector are outside this fiducial region). The material budget for the tracker has risen steadily over the years and the only significant decrease seen (from 1997 to now) is due to the rerouting of the pixel services from a large radius along the LAr barrel cryostat to a much smaller radius along the pixel

**Table 16.8** Evolution of the amount of material expected in the ATLAS tracker from 1994 to 2007

Date	ATLAS tracker material budget estimate [ $X/X_0$ ]	
	$ \eta  \approx 0$	$ \eta  \approx 1.7$
1994 (Technical Proposal)	0.20	0.70
1997 (Technical Design Report)	0.25	1.50
End 2005 (End of construction)	0.40	1.35
Summer 2007 (End of installation)	0.47	2.40

The numbers are given in fractions of radiation lengths ( $X/X_0$ ). Note that, for ATLAS, the reduction in material from 1997 to 2006 at  $\eta \approx 1.7$  is due to the rerouting of pixel services from an integrated barrel tracker layout with pixel services along the barrel LAr cryostat to an independent pixel layout with pixel services routed at much lower radius and entering a patch panel outside the acceptance of the tracker (this material appears now at  $\eta \approx 3$ ). Note also that the numbers do not represent all the material seen by particles before entering the active part of the electromagnetic calorimeter, since particles see in addition the barrel LAr cryostat and the solenoid coil (amounting to approximately  $2X_0$  at  $\eta = 0$ ) or the end-cap LAr cryostat at the larger rapidities



**Fig. 16.10** Material distribution ( $X_0, \lambda$ ) at the exit of the ATLAS tracker, including the services and thermal enclosures. The distribution is shown as a function of  $|\eta|$  and averaged over  $\phi$ . The breakdown indicates the contributions of external services and of individual sub-detectors, including services in their active volume. These plots do not include additional material just in front of the electromagnetic calorimeter, which is quite large in ATLAS (LAR cryostats and, for the barrel, solenoid coil)

support tube, a significant change in the ATLAS tracker design, which occurred in 1999.

Figure 16.10 shows how this material budget is distributed as a function of pseudorapidity. The material closest to the beam (pixel detectors) is clearly the one most critical for the performance of the tracker and of the electromagnetic calorimetry: this amounts to between 10 and 50%  $X/X_0$ . The material budget can also be broken down in terms of its functional components: a large contribution to the material budget arises from cooling and cables in areas where these services accumulate to be routed radially outwards, towards the cracks in the electromagnetic calorimetry foreseen for their passage. It is therefore not surprising that, until all the details of the granularity, technical components, routing, fixation schemes, etc., were known and incorporated into assembly drawings and detailed spreadsheets, the material budgets announced for this tracker of unprecedented scope and complexity were largely underestimated.

### 16.3.2.2 Silicon-Strip and Straw Tube Trackers

The ATLAS SCT contains a total of 4088 modules corresponding to 6.3 million channels, of which 99.7% have been measured to be fully operational in terms of electrical and thermal performance in situ. The ATLAS TRT comprises approximately 350,000 channels, of which about 98.5% fully meet the operational specifications in terms of noise counting rate and of basic efficiency and high-voltage behaviour.

The ATLAS tracker was installed in three successive stages, from summer 2006 (barrel SCT/TRT tracker), to end 2006 (end-cap SCT/TRT trackers), and to spring 2007 (pixels). It is impossible to properly give credit here to all the work performed

over the past 15 years to validate the design choices involving each and every one of the delicate components composing these tracking detectors. Only a few of the most prominent examples are quoted below:

- all the front-end electronic designs had to be submitted to stringent specifications in terms of survival to very high ionisation doses and neutron fluences and of robustness against single-event upsets. The performance of fully irradiated and operational modules equipped with the latest iteration in the design had to be repeatedly measured and characterised in laboratory tests and particle beams of various types and intensities [20].
- each component in contact with the active gas of the ATLAS TRT straws has had to be validated in a well-controlled set-up over many hundreds of hours of accelerated ageing tests using the gas mixture chosen for operation in the experiment. This was necessary because impurities of only a few parts per billion, picked up somewhere in the system, could be deposited on the wires and thereby destroy the gas gain in an irrecoverable way [21]. One critical component in the barrel TRT modules, a glass bead serving as wire joint to separate the two halves of each wire, actually failed the ageing tests with the originally chosen gas mixture ( $\text{Xe}-\text{CO}_2-\text{CF}_4$ ) and the collaboration had to eventually change the gas mixture to the current one ( $\text{Xe}-\text{CO}_2-\text{O}_2$ ), in which the fluorine component has been removed. This gas mixture reduces the direct risk to the wire joints, but is somewhat less stable operationally and does not have the same self-cleaning properties as the original one.

### 16.3.2.3 Pixel Detectors

The ATLAS pixel detector has been one of the last elements installed in the experiment, in great part for practical reasons, but also because this is the detector which has undergone the most difficult development path. It can perhaps be considered as the most striking example of the marvels achieved during the long and painstaking years of research and development: the pixel detector will survive over many years in the most hostile region of the experiment and deliver some of the most important data required to understand in detail what will be happening within a few tens of microns from the interaction point.

Fifteen years ago, at the time of the ATLAS Technical Proposal, very few physicists believed that these detectors could be built within the specifications required in terms of radiation hardness and of readout bandwidth and speed. Today, the data collected using cosmic rays (in 2008 and 2009) and early collisions (end of 2009) have demonstrated that the pixel detector works as expected. The future will tell how long the innermost layer will survive, but the collaboration is already proceeding towards a strategy of “replacement” of the innermost pixel layer on the timescale of 2015. This innermost layer is not expected to survive over the full timespan of the operation of the experiment, which should lead to integrated luminosities

**Table 16.9** Main parameters of the ATLAS pixel system


Number of hits per track	3
Total number of channels	$80 \cdot 10^6$
Pixel size ( $\mu\text{m}$ in $R\phi$ ) $\times$ ( $\mu\text{m}$ in $z/R$ )	$50 \times 400$
Lorentz angle [degrees], initial...end	$12 \dots 4$
Tilt in $R\phi$ [degrees]	20 (only barrel)
Total active area of silicon [ $\text{m}^2$ ]	$1.7 (n^+/n)$
Sensor thickness [ $\mu\text{m}$ ]	250
Total number of modules	1744 (288 in disks)
Barrel layer radii [mm]	50.5, 88.5, 122.5
Disk layer min...max. radii [mm]	88.8...149.6
Disk positions in $z$ [cm]	49.5, 58.0, 65.0
Signal-to-noise ratio for m.i.p. (initial)	120
Total fluence at $L = 10^{34}$ ( $n_{eq}/\text{cm}^2/\text{year}$ ) at radius of 5 cm (innermost layer)	$3 \cdot 10^{14}$
Signal-to-noise ratio (after $10^{15}$ $n_{eq}/\text{cm}^2$ )	80
Resolution in $R\phi$ ( $\mu\text{m}$ )	$\approx 10$
Resolution in $z/R$ ( $\mu\text{m}$ )	$\approx 100$

of close to  $300 \text{ fb}^{-1}$ . Table 16.9 shows the most relevant parameters concerning the ATLAS pixel system.


Finally, Fig. 16.11 shows the results of test-beam measurements of the  $R\phi$  accuracy of production modules of the ATLAS pixel detector before and after being irradiated with a total equivalent fluence corresponding to about  $10^{15}$  neutrons per  $\text{cm}^2$  [22]. These results are somewhat optimistic because they were obtained with analogue readout and at an ideal incidence angle, but they nevertheless demonstrate the extreme robustness of the pixel modules constructed for ATLAS. This is one striking example of the painstaking validation work done in the early phase of the construction years.

## 16.4 Calorimeter System

The design of the ATLAS calorimeter system is to a large extent the end product of about 25 years of development and experience gained over several generations of high-energy colliders and general-purpose experiments, all of which have brought major advances in the understanding of the field. These advances range from the concept of full coverage in total transverse energy at UA1, to that of precision hadron calorimetry at ZEUS, and to that of very high granularity of the electromagnetic calorimeters and the use of energy-flow techniques in the LEP detectors [23].



**Fig. 16.11** Residuals from  $R\phi$  measurements of production-grade ATLAS pixel module before irradiation (left) and after being irradiated with a total equivalent fluence corresponding to about  $10^{15}$  neutrons per  $\text{cm}^2$  (right), as obtained from test-beam data taken in 2004. The contribution of the track extrapolation to the width of the residuals is about  $5\text{ }\mu\text{m}$  (it should be subtracted in quadrature from the overall residual widths quoted in the figure to obtain the intrinsic resolution of the tested module)



**Fig. 16.12** Cut-away view of the ATLAS calorimeter system. The various calorimeter components are clearly visible, from the LAr barrel and end-cap electromagnetic calorimeters, to the scintillating tile barrel and extended barrel hadronic calorimeters, and to the LAr end-cap and forward hadronic calorimeters

The ATLAS calorimeter system, as depicted overall in Fig. 16.12, will play a crucial role at the LHC for two main reasons: first, its intrinsic resolution improves with energy, in contrast to magnetic spectrometers; second, it will provide the trigger primitives for all the high- $p_T$  objects of interest to the experiments except for the muons.

The integration of a hermetic and high-precision calorimeter system into the overall design of the ATLAS detector and its magnet systems has been a task of high complexity where compromises have had to be made, as will be shown in the first part of this section, which describes the basic requirements and features of the calorimeters. As illustrated in the second part, which highlights some aspects of the construction of the most critical element, namely the electromagnetic calorimeter, and of its measured performance in test beam, the impact of the main design choices and of the technology implementations on the performance has been very significant. A few examples of the overall performance expected in the actual configuration of the experiment are presented in Sect. 16.8.2, where it is also compared to the expected performance of the CMS calorimeter system.

## 16.4.1 General Considerations

### 16.4.1.1 Performance Requirements

The main performance requirements from the physics on the calorimeter system can be briefly summarised as follows:

- excellent energy and position resolution together with powerful particle identification for electrons and photons within the relevant geometrical acceptance (full azimuthal coverage over  $|\eta| < 2.5$ ) and over the relevant energy range (from a few GeV to several TeV). The electron and photon identification requirements are particularly demanding at the LHC, as already explained in Sect. 16.2.1. These considerations induce requirements of high granularity and low noise on the calorimeters. One has to add to this the operational requirements of speed of response and resistance to radiation (the electromagnetic calorimeters will have to withstand neutron fluences of up to  $10^{15}$  n/cm<sup>2</sup> and ionising radiation doses of up to 200 kGy over 10 years of LHC operation at design luminosity).
- excellent jet energy resolution within the relevant geometrical acceptance, which is similar to that foreseen for the electron and photon measurements (see above). The quality of the jet energy resolution would play an important role in the case of discovery of supersymmetric particles with cascade decays into many hadronic jets [24].
- good jet energy measurements over the coverage required to contain the full transverse energy produced in hard-scattering collisions at the LHC. A calorimetry coverage over  $|\eta| < 5$  is necessary to unambiguously ascribe the observation of significant missing transverse energy to non-interacting particles, such as neutrinos from W-boson decay or light neutralinos from supersymmetric particle cascade decays. With adequate calorimetry coverage providing precise measurements of the missing transverse energy, the experiments will be able to reconstruct invariant masses of pairs of hadronically decaying  $\tau$ -leptons produced for example in the decays of supersymmetric Higgs bosons. They