

Chapter 8

The Large Hadron Collider (LHC): The Energy Frontier

Giorgio Brianti and Peter Jenni

8.1 Introduction

The accelerator

The CERN Large Hadron Collider (LHC) complex, which includes the 27 km long superconducting collider itself and several huge particle detectors, is the largest and most sophisticated research instrument ever built. We describe here the main features of this marvel of technology [1] and the highlights of the developments which paved its way, from conception in the early 1980s to the first years of exploitation from 2009 to 2013. This period was marked by the discovery of the Higgs boson, announced in July 2012.

In the second half of the 1970s the European Committee for Future Accelerators (ECFA) chose the Large Electron Positron Collider (LEP) as CERN's flagship machine for a thorough experimental study of the Standard Model. The planned beam energy was about 100 GeV, the tunnel circumference about 27 km. Already at that time the sentiment prevailed that such a construction represented a major and perhaps unique undertaking, also in view of a possible subsequent use of the tunnel for a proton machine. While the construction of the LEP tunnel was underway, in the USA the design and construction of the Superconducting Super Collider (SSC) for proton-proton collisions with a beam energy of 20 TeV was proceeding at high speed. The question arose as to what CERN should do for a proton/hadron machine. CERN had not only a tunnel (in construction) but also an existing well-proven chain of injectors (linear accelerator, booster, PS and SPS), which, after some modifications, could be used for injecting protons and ions into a collider in the LEP tunnel [2]. In this context, we started the feasibility study of a proton-proton collider in the LEP tunnel at the end of 1982. Though modest, it prompted the interest of the user community and led to the Lausanne ECFA-CERN workshop in 1984 devoted to the Large Hadron

Collider in the LEP tunnel [3]. Three years later the Workshop on physics at future accelerators in La Thuile, Italy [4] and the report to the CERN Council of the long-range planning committee [5] chaired by Carlo Rubbia, 1984 Nobel Prize laureate and strong supporter of hadron colliders, put the LHC squarely on the map as the next CERN project after LEP.

The project received the first approval by the CERN Council in December 1994 and final approval for the present configuration in 1996 [6]. Lyndon Evans was appointed project leader and led the construction to completion. Commissioning started in 2008, but was soon interrupted by a major technical incident. After the repair, the first collisions were observed in November 2009, and a successful three-year run started in 2010. Globally, it has represented a considerable advance in the technology of large superconducting accelerators.

In a proton–proton collider the two counter-rotating beams travel in separate channels in the arcs around the ring and consist of many bunches of particles separated by a short distance; bunches are brought into collision at the centre of the experimental facilities with a small crossing angle. For the LHC the maximum number of bunches is 2808, the separation 7.5 m and the crossing angle $285 \mu\text{rad}$. The layout of the collider is shown in Fig. 8.1. There are eight arcs, where the beams are bent (the bending radius is 2804 m), and eight straight-sections used for utilities or collision insertions; four insertions are equipped with detectors.

The two most important parameters of a proton–proton collider are the energy of the collisions and the luminosity. The energy is related to the discovery potential of new particles of higher mass and the luminosity to the production of a relevant number of the desired events in a reasonable running time. The luminosity is a measure of the quality of the beams and is proportional to the rate of collisions [Box 6.1].

The same basic collider structure can be used as a proton–proton collider with centre-of-mass energy 14 TeV and luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ or as a lead ion–lead ion collider with a centre-of-mass energy of 1148 TeV (2.76 TeV per nucleon) and luminosity of $10^{27} \text{ cm}^{-2} \text{ s}^{-1}$. Protons can also be made to collide with Pb ions.

The *beam energy* is proportional to the product of B (the magnetic field of the main dipole magnets) and ρ (the bending radius of the arcs). As ρ was determined by the dimensions of the pre-existing LEP tunnel, it was mandatory to aim for the highest possible field: the magnets had to be superconducting [Boxes 4.3 and 8.1]. Prior to the LHC, three large accelerators were based on such magnets: the Tevatron (FNAL), HERA (DESY) and RHIC (BNL). All of these made use of well-proven Nb–Ti cable cooled with helium to a temperature slightly above 4.2 K (He I). In these machines the magnetic field was $\leq 5 \text{ T}$ [6]. The question was how to obtain a higher field with commercially available material and technology.

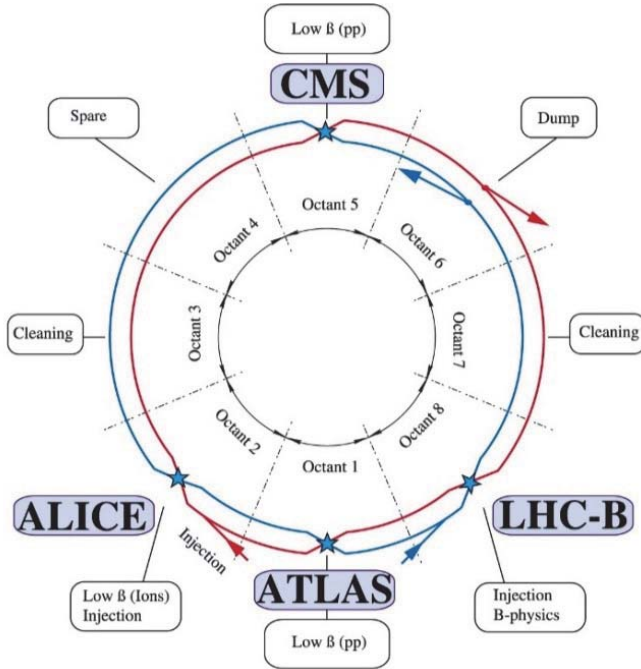


Fig. 8.1. Layout of LHC Collider.

As explained below, the choice was to use Nb-Ti superconductors cooled to a lower temperature, namely 1.9 K provided by superfluid helium (He II). At the required current density the gain due to the lower temperature is about 1.5 T, or an additional 1.3 TeV in energy — over and above the 6.8 T being proposed for the SSC at the time (itself based on intervening improvements of Nb-Ti superconducting material). Before the LHC, the technology of Nb-Ti cooled to 1.9 K by superfluid helium had already been used for a large magnet system (though of a different and more compact geometry), the Tokamak TORE SUPRA, at CEA-Cadarache (France). It should be stressed that the solution adopted for the LHC reconciled two requirements, the quest for the highest possible magnetic field and a design such that the magnets could be constructed in existing industry. With a dipole field of 8.3 T this design provides a nominal beam energy of 7 TeV. Of the almost ten thousand magnets required for the collider, the 1232 main dipoles (length 16.5 m, weight 34 t, max. field 8.3 T) represented a particularly challenging effort of design and construction [Highlight 8.2].

Operating Nb-Ti conductors at 1.9 K offers another important advantage at the expense of a more complex cryogenic system. Due to the particular properties of superfluid helium such as high heat capacity, low viscosity, and exceptionally good thermal conductivity, the long string of magnets can be cooled via a static bath of helium by heat exchange with a very small flow of liquid. This avoids depending on pumps — components of limited reliability and efficiency — for producing a large continuous flow of cold helium. The total cold mass operating at 1.9 K is 37500 t. This temperature is lower than the background temperature of the Universe [7] [Highlight 8.3].

The considerable kinetic energy of each beam (362 MJ) and the large electromagnetic energy stored in the magnets (11 GJ for the complete system) require a sophisticated protection system to prevent damage in case of beam guiding problems or resistive transition (magnet “quench”). Each of the two beams has sufficient energy to melt 500 kg of copper, and the electromagnetic energy of the magnets, if not properly discharged, can provoke substantial damage as in the case of the 2008 incident.

The nominal excitation current of the strings of dipoles and quadrupoles in the arcs is 12 kA. To optimize system cost the individually powered quadrupoles of the matching sections were designed to operate at a more modest 6 kA. Nevertheless, the total current to be transferred from room temperature down to the helium environment of the superconducting magnets was still about 3 MA. To reduce cryogenic load it was therefore decided to use High Temperature Superconducting (HTS) current leads. Since over a thousand such leads of three current ratings were required, a design suitable for mass production was developed at CERN. It has been the first large scale application of HTS [Highlight 8.4].

Another important aspect is the beam vacuum, for which requirements are imposed by beam lifetime and background in the experiments. The interactions between the protons and rest gas are driven by two processes, i.e. single proton-nuclear collisions in which a proton is lost, and multiple small-angle Coulomb scattering, which provokes an increase of the size of the beam and a corresponding decrease of the luminosity. To ensure a beam lifetime of about a day, the residual pressure in the vacuum chamber should not exceed 10^{-8} – 10^{-9} Pa. The vacuum chamber in the LHC consists basically of two types: (i) chambers made from stainless steel at 1.9 K in the arcs; (ii) chambers made from copper operating at ambient temperature in the total of 6 km of straight sections.

The vacuum system has to address a dynamic effect created by electrons, the electron cloud phenomenon. These electrons, produced by synchrotron radiation or by the ionization of the rest gas, can be accelerated to the walls of the vacuum chamber, producing secondary electrons, which in turn can be accelerated by the

subsequent bunches, impinging on the chamber walls and releasing further electrons with the risk of a run-away situation. This can give rise to beam blow-up and instability, and deposit heat in the cryogenic system, so it has to be controlled. This is done by lowering the secondary emission yield of the chamber walls. Different strategies are applied in the cold and warm sections.

In the cold sections the vacuum chamber is at the same temperature as the superconducting dipoles enclosing the chamber (“cold bore design”). For the pumping, one takes advantage of the fact that the vapour pressures of all residual gases are lower than 10^{-10} Pa at 1.9 K, a feature already explored in the ISR. Such pressures are sufficiently low for machine operation. The electron cloud is fought by so-called beam scrubbing, i.e. using gradually increasing beam currents until the walls have reached the desired degree of cleanliness. However, the circulating beam would deposit up to 1.7 W/m on the vacuum chamber (due to synchrotron radiation, image current, residual electron cloud effects, etc.) which should not be absorbed at 1.9 K. Therefore, the magnet cold bore is shielded with an inserted beam screen kept at 5 to 20 K by separate circulation of helium. The beam screen has holes (about 4% of the screen surface) to allow the gas molecules to reach the surrounding cold bore.

For the warm sectors, the successful performance of Non Evaporable Getter (NEG) pumping for LEP was a strong motivation to envisage NEG pumping also for these LHC sectors. As the LHC chambers are too small to host the LEP getter pump, a way was found to coat the chamber walls with a film of NEG to provide pumping and control of the electron cloud effect [Highlight 8.5].

A view of a section of the installed LHC machine is shown in Fig. 8.2.

As stated previously, besides the energy of the circulating particles, the luminosity L is an important parameter of the LHC, but it is the luminosity integrated over a given period, say one year, that is crucial for data-taking [Box 6.1]. The luminosity varies with time, since the beams degrade during the collision run. The integrated luminosity is simply the integral $\int L(t)dt$. These parameters are the key figures of merit of a collider.

The LHC was commissioned in 2010 with 3.5 TeV protons; operating at 4 TeV it has reached a luminosity of $0.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and provided up to 23 fb^{-1} per experiment in 2012. Since mid-2015 the LHC is operating with a beam energy of 6.5 TeV and has exceeded the nominal peak luminosity of $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The luminosity will be gradually doubled in the coming years promising an integrated luminosity of 300 fb^{-1} before the accelerator is shut-down for a substantial upgrade turning it into the High Luminosity LHC (HL-LHC) by around 2025 (Chapter 12). The expected peak luminosities at 7 TeV with the nominal 25 ns between bunches




Fig. 8.2. Partial view of the LHC, showing a string of dipoles and an interconnection.

leads to 30 to 60 inelastic interactions per bunch crossing, which determines approximately the number of events the detector has to be able to resolve. For very sensitive experiments this number is decreased by reducing the luminosity in their interaction region. A small vertical separation is introduced which changes with time to “level” the luminosity, hence optimizing the integrated luminosity.

The number of circulating proton bunches is very large (2808) with a bunch spacing of 7.5 m, corresponding to 25 ns. This means that the beams have to be brought into collision at a crossing angle of 285 μrad . This angle, while introducing a moderate reduction of the luminosity and a negligible loss in centre-of-mass energy, eliminates the unwanted bunch collisions near the collision point.

Lead-lead or lead-proton collisions have also been part of the programme since 2010. This mode of operation is regularly scheduled, though only for a fraction of the time foreseen for proton-proton operation. In 2015, the design luminosity of $1 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ and an integrated luminosity of 30 pb^{-1} was reached with 5.0 TeV centre-of-mass energy; these numbers refer to lead nucleon-pair collisions.

The joy of witnessing the first proton collisions following the repair of the accelerator, and the anticipation after so many years of work on both the collider and detectors, is recorded on the faces of scientists in the CERN Control Centre (CCC) in March 2010 [Fig. 8.3].



Fig. 8.3. The CERN Control Centre (CCC) witnesses the first 7 TeV collisions in March 2010.

The physics programme and the detectors

Enthusiasm for the extraordinary physics potential of a LHC arose in the community at the Lausanne ECFA-CERN workshop in 1984 [3]. Studying the Brout–Englert–Higgs (BEH) mechanism [Box 8.2] and searching for the associated so-called “Higgs boson” was clearly a priority for the future of particle physics, but this was expected to be very challenging. The boson mass is not predicted by the Standard Model and could have been as high as 1 TeV. This required exploring a broad range of mass — well-suited for a high energy p–p collider where the energy spectrum of the constituents of the protons (quarks and gluons) allow all possible Higgs masses to be searched for simultaneously. As detectable cross-sections were predicted to be very low, a very high luminosity collider was mandatory.

The Higgs search was by far not the only reason for the great interest generated by the LHC. Already in the 1980s there was clear motivation from speculative theories predicting physics Beyond the Standard Model (BSM), the most popular being Supersymmetry (SUSY) with its characteristic missing transverse energy signatures (E_T^{miss}) due to the escaping lightest neutral SUSY particle (LSP) [Box 7.2]. Other hypothetical models predicted new heavy resonances, or leptoquarks (particles with both lepton and baryon numbers), or even substructure to quarks,

A theory of superconductivity

Box 8.1

In 1957 John Bardeen, Leon Cooper and Robert Schrieffer (Nobel prize 1972) published the first microscopic theory of superconductivity (BCS Theory [1]) based on quantum mechanics treatment of the behaviour of conduction electrons in a metal in relation with lattice vibrations (phonons). Originally, superconductivity (SC) was explained by macroscopic descriptions such as those of London [2] and Ginzburg and Landau [3]. In the BCS theory, below a critical temperature T_c pairs of electrons (of opposite spin) form so-called “Cooper pairs”. Pairing is mediated by phonons in the ion lattice, to assume an energy state which is lower than that of the sum of the single electrons, and therefore favoured. As the pairs have zero net spin, other such pairs, unlike single electrons, can accumulate in the same state to form a condensate. Below T_c the energy of vibrations of the metal lattice is insufficient to change the state of the condensate which can then carry the current freely through lattice without expending energy — in contrast to normal unpaired electrons that experience resistance via encounters with the phonons in the lattice. The process of formation and dissolution of Cooper pairs is continuous, the net result being resistance-free transport of current (DC) through the metal (up to the critical magnetic field $\mu_0 H_c$). For RF SC see Box 7.4.

Superconductors come in two types, *Type I* (most pure metals), which are perfectly diamagnetic, and exclude magnetic field up to the critical field H_c , and *Type II* (alloys, compounds, ceramics and some metals), which feature two critical fields H_{c1} and H_{c2} ; up to H_{c1} they behave like Type I materials, but between H_{c1} and H_{c2} magnetic flux can penetrate locally in the form of *vortices*. While $\mu_0 H_{c1}$ is typically a fraction of a tesla, $\mu_0 H_{c2}$ ranges from about 15 T for Nb-Ti to about 24 T for Nb₃Sn, and up to 100 T for some ceramic materials. Only Type II material is useful for large-scale applications.

Whereas for Type I superconductivity the material should be pure, Type II materials must feature discontinuities (e.g. grain boundaries, dislocations and point defects) to encourage the *pinning* of magnetic vortices. In order to increase the critical current I_c , the manufacturing process therefore includes steps to decrease grain size (e.g. by cold work). Such pinning centres may also be introduced by doping with additives.

The BCS theory was first developed to explain Type I superconductivity. With a few tweaks the model could be applied to the Type II materials of the time, up to $T_c \approx 24$ K, beyond which temperature superconductivity was reckoned to be impossible. So there was great surprise when high temperature superconductors were discovered in 1986. BCS theory is both elegant and attractive, but despite many efforts it has not yet been possible to adapt it convincingly to explain the behaviour of these new materials, and there may be a place for a different theory that covers the phenomenon in a more satisfactory way [4]. However, the BCS theory of superconductivity, derived from the reasoning of solid state physics, also inspired particle physics theorists to postulate the BEH field and its associated boson [Box 8.2], and therein also resides major worth.

[1] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Phys. Rev.* **108**, 1175-1204 (1957).

[2] F. London and H. London, *Proc. Royal Soc. A: Math., Phys. & Eng. Sc.* **149**, 856 (1935).

[3] V. L. Ginzburg and L. D. Landau, *Zh. Eksperim. I Teor.* **20**, 1064 (1950).

[4] J. E. Hirsch, *Physica Scripta* **80** 3, 1-11 (2009).

The Brout–Englert–Higgs (BEH) mechanism**Box 8.2**

At the heart of the SM [Box 6.4] is the freedom to perform at any point of space-time an arbitrary rotation in internal spaces, called a “gauge” transformation (GT), leaving the physics invariant. Such an operation implies replacing one particle by another, e.g. e^- by its neutrino, or a strange quark by a charm quark. Intuitively, the invariance can be expected to work better if one swaps similar particles, e.g. of zero mass. Indeed, the first version of the SM worked only for massless particles, not describing our world: this simple form of Gauge Invariance (GI) does not provide a good description.

Can one reconcile the concept of GI with a quantum field theory describing Nature?

QED with its massless photon as force carrier shows the property of GI and provided a very helpful hint as did the theory of superconductivity, which is based on electro-magnetism. Inside a superconductor (SC) there exists a field due to the Cooper pairs [Box 8.1], made of two electrons interacting through the lattice. This field locks the phase of an electron over a macroscopic distance and thus effectively prevents the possibility to perform an arbitrary local GT. The BEH mechanism is similar, positing the existence of a new field permeating all of our universe and acting in the weak isospin internal space to impose there an orientation, preventing arbitrary rotations in that space. The analogy between the BEH mechanism and SC can be pushed quite far [1]. In the BCS theory of superconductivity the photon becomes massive by interacting with the field of Cooper pairs: this is the Meissner effect, i.e. the inability of a magnetic field to penetrate the SC beyond the penetration depth d , as if its quantum, the photon, had become massive in proportion to $1/d$. Similarly, the W and Z bosons acquire mass from the BEH field, whose non-zero value (“vacuum expectation value” VEV) plays the role of the Cooper pair field, described in a similar mathematical framework. The BEH field gives also the mass to the fermions, which depends on the strength of their interaction with the BEH field. However, we lack an understanding of these couplings, and the fundamental origin of the mass values remains a mystery.

It is now believed that some 10^{-11} seconds after the Big Bang the BEH field underwent a phase transition (Fig. 1.4), which changed its VEV from zero to a finite value, responsible for giving mass to all elementary particles from this moment onwards. Implementing the BEH mechanism into the early SM version was the decisive step in turning the SM into a description of our world.

The Higgs boson is the particle manifestation of an excitation of this BEH field. Its mass is linked to its property of self-interaction (SI). Not knowing the strength of SI, the SM only provides general guidelines, but cannot predict its mass (see however Box 7.2). It took the energy and luminosity of the LHC collisions to produce and observe this boson with a mass of ~ 130 protons at CERN in 2012 (see 8.1). Following this discovery F. Englert and P. Higgs were awarded the physics Nobel Prize in 2013 [2].

The importance of the discovery of the Higgs field goes beyond the SM. Unlike the gravitational or e-m field, it is a novel “scalar” field, not characterized by a direction in our space. Its detailed study may contribute to our understanding beyond the SM.

[1] <http://www.slac.stanford.edu/pubs/beamline/26/1/26-1-dixon.pdf> (1996).

[2] http://www.nobelprize.org/nobel_prizes/physics/laureates/2013/.

not to mention many other exotic ideas. It was also realized early on that the LHC would produce huge numbers of heavy flavour particles, opening a new frontier in precision flavour physics. Unprecedented exploratory steps could be made by studying the quark-gluon plasma at high density and temperatures when colliding heavy ions, such as fully ionized lead nuclei. Motivated by these physics prospects, CERN opted ultimately for an experimental programme with two very large general-purpose detectors (ATLAS and CMS), two large detectors optimized respectively for flavour physics (LHCb) and for heavy ion collisions (ALICE), complemented later by three much smaller specialized experiments (TOTEM, total cross-section and forward physics; LHCf, measuring forward neutral particle production; MoEDAL, monopole search).

The detection of the Higgs boson played a particularly important role in the design of the *general-purpose experiments* ATLAS and CMS. In the region of low mass ($114 < m_H < 150$ GeV), the two channels considered most suited for unambiguous discovery were the decay to two photons and the decay to two Z bosons, each decaying in turn into e^+e^- or $\mu^+\mu^-$, where one or both of the Z bosons could be virtual. As the natural width of the low-mass Higgs boson is less than 10 MeV, the width of any observed peak would be entirely dominated by the instrumental mass resolution. This meant that in designing the general-purpose detectors, considerable importance was placed on the magnetic field strength, on the precision tracking systems and on high-resolution electromagnetic (e-m) calorimeters. The high-mass region, as well as the signatures from supersymmetry, required excellent resolution for jets and missing transverse energy (E_T^{miss}), implying also almost complete 4π hadronic calorimetry coverage.

The choice of the magnetic field configuration determined the overall design for these experiments. It was also well understood that optimizing the chance of making discoveries at the new “magic” energy scale of the LHC - and in the harsh conditions generated by about a billion protons interactions every second, and several tens of events every bunch crossing — would require the invention of new technologies while at the same time pushing existing ones to their limits [8]. In fact, a prevalent saying was “We think we know how to build a high energy, high luminosity hadron collider — but we don’t have the technology to build a detector for it”. In reality of course both turned out to be difficult and demanded technological breakthroughs. That the general-purpose experiments have worked so well since the start-up of the LHC is a testimony to the enthusiasm and competence of many thousands of engineers and physicists, and to the difficult yet correct technological choices and decisions made during the construction of these experiments. Indeed, the very same elements mentioned above were crucial in the recent discovery of a Higgs boson.

Experimentation at the LHC**Box 8.3**

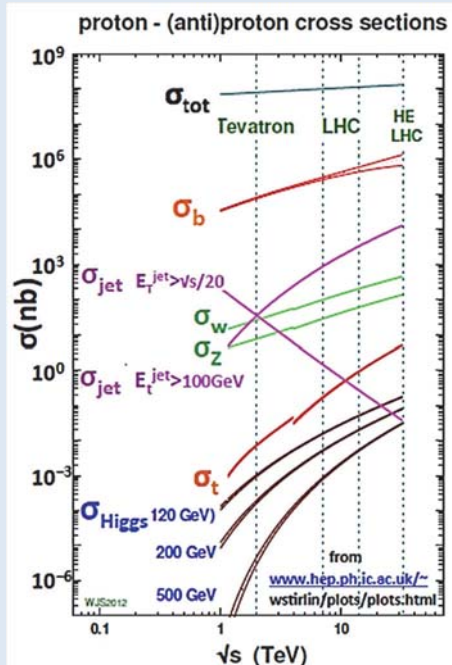
Protons have radius $R = 0.88$ fermi ($1 \text{ fermi} = 10^{-13} \text{ cm}$) and interact strongly with each other when $\lesssim 2R$ apart, corresponding to an area of $100 \times 10^{-27} \text{ cm}^2$, or a cross-section of 100 millibarn ($1 \text{ barn (b)} = 10^{-24} \text{ cm}^2$). Total pp cross section is $\sim 70 \text{ mb}$. This fact, and the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, impose severe constraints on the detectors:

- Every 25 ns bunches of protons collide, producing about 1000 particles. These must be measured individually, so the detectors must be finely subdivided;
- 31×10^6 bunch crossings/s (2808 bunches \times 11.2 kHz) must be examined. Within a few μs , the electronics select about 10^5 events based on their potential interest for physics (first level trigger). In a second computer-based selection about 200 events are kept. With $\sim 1 \text{ MB/event}$, a volume of 100 GB/s is analysed and 200 MB/s archived;
- Per year of data taking, typically 10^7 s , 2×10^9 events, corresponding to 2 petabytes (PB), are recorded and analysed. Such analysis takes about 100 s per event for a PC, implying 20000 equivalent PC units operating continuously. For the 4 experiments, data (re)processing and Monte Carlo data generates 30 PB/year;
- High radiation levels are a strong concern: up to 2×10^7 tracks per cm^2/s traverse the detectors. Each deposits $\sim 2 \text{ MeV/g cm}^2$ [Box 5.2]. This leads to an irradiation of about 100 kGy/year. Detectors and electronics are designed to work for ≥ 10 years.

Trackers must have low mass to avoid perturbing particles — another major constraint.

At the LHC, the focus is on energetic collisions of the point-like (size $< 10^{-4}$ fermi) proton constituents, the partons, gluons and quarks, for which the associated wavelength is about 10^{-3} fermi at the 100 GeV energy scale. So cross sections 10^{-6} times smaller than the total cross section must be considered (see figure, to be compared with Fig. 12.12 for an e^+e^- collider). The process rate also depends on the properties of the partons within the proton and how they interact (strongly or weakly). To produce heavy or energetic particles, the colliding partons must carry a large fraction x of the proton momentum — less probable to occur for larger x . Producing the Higgs boson [Box 8.2] proceeds through a loop of heavy virtual particles [Box 5.1] and is very infrequent. To beat background, clean rare decay modes are used. To observe one Higgs event in the 4-lepton channel means registering $\sim 10^{13}$ pp interactions.

To address such challenges, in 1987 CERN launched a world-wide R&D programme and set up an external, independent LHC Committee with the mandate to scrutinize, monitor and evaluate progress (Chapters 8 and 11).



Very different challenges were faced by the two *special-purpose experiments* LHCb and ALICE, which are reflected in their totally different designs. They are operated at lower luminosity, avoiding multiple collisions per bunch crossing, and employing a range of detectors optimized for particle identification.

These four sophisticated instruments were developed in a worldwide coordinated R&D effort, culminating in large scale prototype measurement campaigns in particle beams at CERN and other accelerator laboratories, over the 1990s and early 2000s. Construction of the various components took place over about 10 years, starting in the second half of the 1990s, in universities, national laboratories, and industry. Typically, after initial local testing, components were delivered to CERN; following reception tests they would be assembled and installed in the experimental caverns where they would be finally commissioned. Figure 8.4 shows the four detectors during their final installation phase.

Below the four detectors are briefly introduced, followed by selected technological highlights. These may address an aspect of a specific experiment or one of concern to all of them, such as electronics and signal processing. The selection of highlights is not exhaustive: they are meant to give a flavour of the many challenges that the experiments had to tackle in realizing the LHC research agenda. The four collaborations have published detailed and comprehensive technical descriptions of their detectors as finally built and operated during the first years of LHC [9–12].

ATLAS [9]

The design of the ATLAS detector is based on a large superconducting air-cored toroidal magnet system for the muon spectrometer complemented with a smaller superconducting solenoid around the inner tracking cylinder. The momentum of the muons are precisely measured as they travel unperturbed by material for more than 5 m in the toroid. About 1200 large muon drift tube chambers of various shapes, with a total area of 5000 m², measure the impact positions with an accuracy of better than 0.1 mm. Another set of about 4200 timing chambers are used to provide a fast muon trigger (resistive plate chambers in the barrel, thin gap chambers in the end-caps).

In the field-free region between solenoid and toroids there is a highly granular e-m calorimeter using a novel lead–liquid argon (LAr) sampling technique [Highlight 8.8], complemented by full-coverage hadronic sampling calorimeters. For the latter a plastic scintillator–iron sampling technique, also of novel geometry, is used in the barrel part of the experiment, and LAr calorimeters cover the end-cap region near the beam axis where particle flux and exposure to radiation is highest. The calorimeters have 200 000 e-m cells and 10 000 hadronic cells.

Reconstruction of all charged particles, including those of displaced secondary vertices, is done in the inner tracking detector, which combines highly granular pixel ($50\ \mu\text{m} \times 400\ \mu\text{m}$, 80 million channels) and microstrip ($13\ \text{cm} \times 80\ \mu\text{m}$, 6 million channels) silicon semiconductor sensors close to the beam axis [Highlights 8.6 and 8.7], and a ‘straw tube’ gaseous detector (4 mm diameter, 350 000 channels) that provides about 35 signal hits per track. The latter also uses information from transition radiation for the identification of electrons.

Figure 8.4(a) shows one end of the cylindrical barrel detector after about 4 years of *in situ* installation work in the underground cavern. The ends of four of the barrel toroid coils are visible, illustrating the eightfold symmetry of the structure. The relatively lightweight overall structure of the detector adds up to 7000 t, covering a volume of 22 m in diameter and 46 m in length.

CMS [10]

The CMS detector features prominently a state-of-the-art superconducting 4 tesla solenoid. This single magnet [Highlight 8.12] serves both the high precision inner tracking chambers as well as the external muon detector, which is instrumented with large gaseous drift chambers in the barrel and cathode strip chambers in the end-caps complemented by resistive plate trigger chambers, embedded in the return yoke. With this configuration a compact overall design could be achieved, with the diameter limited to 15 m. The magnet yoke makes up most of the total detector weight of 12 500 t.

Muon detection, and their most accurate measurement, was a priority criterion for the CMS design, in addition to precision measurements for photons and electrons, achieved with a new type of radiation-hard e-m crystal calorimeter, the largest ever built. The challenging, but very successful development and manufacture of the 75 848 lead tungstate scintillating crystals in the final set-up took more than a decade [Highlight 8.8].

For charged particle tracking the collaboration opted for precise position measurements for each trajectory (13 layers with a position resolution of $\sim 15\ \mu\text{m}$ per measurement) leading to a large number of cells distributed inside a cylindrical volume 5.8 m long and 2.5 m in diameter: 66 million $100 \times 150\ \mu\text{m}^2$ silicon pixels and 9.3 million silicon microstrips from about $10\ \text{cm} \times 80\ \mu\text{m}$ to $20\ \text{cm} \times 180\ \mu\text{m}$. With a silicon detector area of $198\ \text{m}^2$ the CMS tracker is the largest ever built [Highlights 8.6 and 8.7].

The hadron calorimeter, comprising ~ 3000 small solid angle projective towers covering almost the full solid angle, is built from alternate plates of $\sim 5\ \text{cm}$ brass absorber and $\sim 4\ \text{mm}$ thick scintillator plates that sample the shower energy. The

scintillation light is measured with photo-detectors (hybrid photodiodes) that can operate in the strong magnetic field, as the calorimeters are inside the solenoid.

The iron yoke of the CMS detector is sectioned into five barrel wheels and three endcap disks at each end. This enabled the detector to be assembled and tested in a large surface hall while the underground cavern was being prepared. The sections, weighing between 350 t and 2000 t were then lowered into the cavern (Fig. 8.4(b)) between October 2006 and January 2008, using a gantry crane system equipped with strand jacks: a pioneering use of this technology to simplify the underground assembly of large experiments.

LHCb [11]

The LHCb detector concept takes advantage of the dominant rate of beauty production close to the beam directions, for practical reasons on just one side of the interaction region. The layout of LHCb therefore resembles a fixed target spectrometer, however with very unique features.

A silicon strip vertex locator (VELO) detector can be positioned during stable beams very close to the interaction region and beam line in order to measure accurately primary and secondary vertices, selecting events with b-quarks. Particle identification to identify unambiguously the various final states is achieved by two ring imaging Cherenkov detectors [Highlight 8.10], whereas the momentum measurements are based on a large-aperture warm dipole magnet generating an integrated field of 4 T m for trajectories going through all tracking stations (silicon strips and straw tube drift chambers) of the spectrometer. Calorimetry uses sampling scintillator lead (e-m) and coarser scintillator iron technology. The muon detector behind the calorimeters, an absorber with iron plates providing a total of 20 interaction lengths, sampled by four chamber layers, completes the LHCb spectrometer. Figure 8.4(c) shows a picture of the fully installed LHCb detector in LHC Point 8.

ALICE [12]

The ALICE detector has to cope with events having extremely high multiplicity, characteristic of heavy ion collisions, with charged particle measurements down to the lowest technically possible momenta. Novel particle identification has been developed for many of the specific heavy ion physics signatures.

The ALICE Collaboration has reused the huge former L3 warm solenoid magnet providing a field of 0.5 tesla over a large central volume [Highlight 7.10]. Inside the magnet with its 10 000 ton iron yoke is located, as central tracking detector, the world's largest Time Projection Chamber (TPC) of 5.6 m in diameter and 5.4 m in length, which provides precision tracking as well as particle

identification by dE/dx ionization. The innermost region, inside the TPC, facing the most extreme particle density region around the collision point, is instrumented with an inner tracking system consisting of silicon pixel, silicon drift and double-sided silicon strip detectors. Several detector systems dedicated to particle identification cover the volume outside of the TPC: transition radiation and Cherenkov radiation detectors, and a state of the art time-of-flight system [Highlight 8.9]. High resolution e-m calorimetry for photon measurements is implemented with lead tungstate scintillating crystals, similar to those of CMS.

A muon spectrometer starting with a massive 4 m long sophisticated hadron absorber cone, and featuring a large classical dipole magnet, covers on one side the solid angle from 2 to 9 degrees with respect to the beam direction. A front view of the ALICE detector nearing installation completion is shown in Fig. 8.4(d).




Fig. 8.4. The four LHC detectors during the installation phase. Upper row left (a) ATLAS, right (b) CMS, lower row left (c) LHCb, and right (d) ALICE.

Table 8.1. The expected and observed quality of measurements in ATLAS and CMS, expressed in standard deviations of the excess (signal) from the background-only hypothesis, for $m_H = 125$ GeV, for various decay modes.

Experiment	ATLAS		CMS	
Decay mode	Expected (σ)	Observed (σ)	Expected (σ)	Observed (σ)
$\gamma\gamma$	4.6	5.2	5.3	5.6
ZZ	6.2	8.1	6.3	6.5
WW	5.8	6.1	5.4	4.7
bb	2.6	1.4	2.6	2.0
$\tau\tau$	3.4	4.5	3.9	3.8

Complementarity of technologies

As pointed out above, rather different technological choices have been made by the LHC experiments, and yet all performed at or even beyond specifications. Particularly striking are the different technologies implemented on major detector systems by ATLAS and CMS, thereby offering significant complementarity for the instrumental approaches to the searches for the Higgs boson and for BSM physics. This aspect was an important consideration in CERN's approval procedure for the LHC experiments, with the aim of addressing the difficult LHC research program with different and largely independent experimental strategies.

Table 8.1 illustrates that indeed the two experiments have achieved very similar results for the Higgs boson discovery and observation in different final states, and this despite the large conceptual differences in the detector designs. Shown in the table [13] are the observed and expected significances of the signals expressed in the number of standard deviations (σ).

8.2 Superconducting Magnets: Powerful, Precise, Plentiful

Romeo Perin

The LHC dipole magnet (Fig. 8.5) is one of the most complex objects designed and fabricated on a large scale (1232 installed) in the history of particle accelerators. Major challenges for the designers were the choice of materials, assembled at room temperature, shrinking by 2 to 5 mm/m during cool-down to the working temperature of 1.9 K (-271°C), and the precision required for the magnetic field, B , in the presence of large electromagnetic forces (proportional to B^2) tending to open and flatten the coils. Sophisticated methods and equipment also had to be developed for handling the large magnetic stored energy in case of a transition from the superconducting to the normal conducting state (*quench*).




Fig. 8.5. Cut-away view of the 15 m long superconducting LHC dipole magnet.

The dipole must provide an accurate field from 0.5 to 8.3 tesla (T). For comparison the magnetic field in the core of a transformer is about 1 T. Previous large-scale production of superconducting magnets for accelerators was for devices operating at up to about 5 T, so R&D was clearly necessary. Initially, with laboratory resources focused on the construction of LEP, this was done with limited means and contributions from institutes and industry, but could continue more vigorously following project approval in 1994 [14]. The presence at CERN of experienced staff in superconducting magnets and cryogenics was vital for developing the technology. It was necessary to extend the state-of-the-art for practically every feature of the magnet, requiring exhaustive calculations and an extensive mock-up and model programme to arrive at cost-effective solutions. Importantly, thanks to good cooperation with scientists in the USA, CERN could benefit from the R&D on the conductors and 6 T magnets that had been done in the framework of the Superconducting Super Collider (SSC) project.

Already early in the programme a twin-aperture magnet (Fig. 8.6) was favoured, due both to the limited transverse space in the LEP tunnel and to the efficiency of the design [15]. In this design concept, first proposed by John Blewett of BNL in 1971, the magnetic flux in one aperture is partly returned via the other. For the LHC magnet this resulted *inter alia* in a cost reduction of about 25%.

After a series of tests, the important decision on the superconductor was taken. Well-proven niobium-titanium alloy (Nb-Ti) was chosen rather than niobium-tin (Nb₃Sn). Due to the brittleness of Nb₃Sn, coils have to be wound before the heat treatment required for forming the superconductor — a risky procedure, and costly for mass production. So despite being less performant at high fields, Nb-Ti, was




Fig. 8.6. Field and forces in the twin aperture magnet. The direction of the field is opposite in the two apertures — required for bending the trajectories of counter-rotating beams of particles of like sign. The horizontal components of the forces on the adjacent coils of the two apertures cancel; at the design field of 9 T, $2 F_h = 400$ t/m and $F_v = 100$ t/m. Pre-stress on the coils and structure stiffness must ensure minimising movements due to these forces.

adopted — cooled to 1.9 K in superfluid helium (He II) to improve its current-carrying capacity. He II is also a very good conductor of heat (~ 1000 times better than copper), enabling very efficient cooling of the coils.

The stability of the beam in the collider depends crucially on the field quality of the dipole and quadrupole magnets, determined by the position of the conductors around the aperture rather than the accuracy of the surface of the poles as for normal iron-dominated magnets. The position of the conductor must be computed, the magnets manufactured to these specifications and maintained during operation in the presence of the strong electromagnetic forces [16].

The cross-section of the dipole magnet [17] is shown in Fig. 8.7. The coils are wound around a bore of diameter 56 mm and the force and flux containment structure consist of laminated stainless steel collars, a laminated iron yoke split vertically and a stainless steel shrinking cylinder. The latter also functions as the helium vessel, and is slightly curved in the horizontal plane to align the magnet apertures with the beam trajectory. The iron yoke serves to stiffen the structure and to shield the stray field; it also enhances the useful field by about 10 %. Static superfluid helium at atmospheric pressure permeates the magnet, cooled by vaporization of superfluid helium in the heat exchanger tube at the top of the iron yoke [Highlight 8.3]. Heat in-leak is reduced to a minimum by enclosing the cold mass in thermal shields covered with superinsulation under vacuum. The cold mass diameter is 570 mm and the complete magnet in its vacuum vessel has diameter 914 mm, and length 16.5 m. It weighs 34 t.

The 40 turns of cable are grouped in blocks separated by copper wedges to approximate a pure dipole field [6, 16]. For the design field of 9 T the current




Fig. 8.7. Model showing part of the dipole magnet cross-section.

required is 12840 A. This highlights the value of superconductivity: if the cable were to be normal conducting copper it would require a cross-section a thousand times larger than the 25 mm^2 of the superconducting cable and power consumption would be at least 50 times greater (with superconducting magnet systems it is the refrigeration plants that consume most of the power — independently of field level).

If a quench occurs the current circulating in the stabilizing copper of the cables generates a turn-to-turn voltage of 20 V and voltage to ground of up to 2 kV. Following extensive studies and tests, polyimide tape was chosen for the insulation. Each coil layer was cured in a mould under pressure at 190°C . Tolerances on conductor position and coil dimensions were 0.05 mm.

The integrated bending strength of each of the 1232 magnets had to be equal to within one part in 1000. This called for careful positioning of components, with checks at every stage. After collaring, field quality was verified by (room-temperature) magnetic measurement, and if necessary corrected by shimming.

When powered, the coils must not experience tensile stress and must therefore be pre-stressed to resist the electromagnetic forces [16]. Pre-stress is provided by compressing the coils with accurately stamped collars. The collars are locked by inserting calibrated rods (Fig. 8.7). To avoid over-stressing the insulation, compression at the top of the coil is chosen to be zero at design field; a maximum pressure of 90 MPa (about 9 kg/mm^2) occurs on the median plane.

At 9 T the twin-aperture dipole stores 8×10^6 joules (8 MJ) of magnetic energy. The energy stored in a string of 154 magnets is 1200 MJ — enough to lift the Eiffel Tower by about 15 m! Quenches happen, and could destroy a magnet and endanger the whole string if not controlled. Quenches are detected by monitoring the appearance of resistive voltages; on detection the quench is spread in the coil by

firing heaters that drive it above the critical temperature. Each magnet is equipped with a large semiconductor diode which conducts at the voltage induced by the resistance of the quenching coil; the current in the string by-passes the quenching magnet. Concurrently, the LHC beams are dumped, the power converter is turned off and a resistor is switched into the circuit to extract the energy of the magnet string which is completely discharged in about 100 seconds [18].

The magnets were produced in industry to build-to-print specifications and with sustained follow-up by CERN. To ensure uniformity of the magnets assembled by three manufacturers and to benefit from economy of scale, CERN procured directly most of the critical components (cables, insulation, spacers, steel laminations, collars, cold bore tubes, shrinking cylinder halves, diodes, etc.) and supplied them to the three companies [19].

Before installation in the LHC tunnel, the magnets were cooled down to 1.9 K and tested [20]. Field quality was measured by means of a specially developed equipment [21] consisting of 16 m long sets of rotating ceramic shafts carrying the measuring coils. Two of these shafts are inserted in room temperature enclosures installed in the magnet apertures and driven by a twin rotating unit.

Superconducting wires and cables

The conductors for the inner and the outer coil layers are flat cables made of composite strands consisting of thousands of Nb-Ti filaments embedded in pure copper (Fig. 8.8). For example, the strand for the cable of the inner layer of the coil has a diameter of 1.065 mm and contains 8900 Nb-Ti filaments. The copper serves to stabilize the superconductor, separate the filaments and to conduct the current briefly in the event of a quench. Fine filaments are important to limit the effect of currents that can be induced in the superconductor and will therefore persist. Such currents are proportional to the diameter of the filaments and are a source of field errors. Filament sizes of 6 μm to 7 μm were chosen for the strands as being the smallest that could be guaranteed without compromising the yield of the manufacturing process, and hence the cost.

Wire is made by stacking Nb-Ti bars and copper rods in a copper can, co-extrusion and drawing. The industrial process starts with 900 mm long, 200 mm diameter bars of Nb-Ti alloy wrapped by a thin sheet of pure niobium; each bar is turned into 30 km long filaments in a sequence of wire drawing and heat treatments. To limit eddy currents in the cable, yet provide sufficient contact for current sharing, strands are thinly coated with an alloy having lower conductivity than copper. Cabling had to be strictly controlled, to guarantee dimensional tolerances and smooth edges (to avoid damaging the insulation). A machine was developed for on-line optical control.




Fig. 8.8. A flat LHC cable, a strand and groups of Nb-Ti filaments in a strand [1].

The critical current of the 28-strand inner cable at 10 T is ≥ 13750 A.

The total quantity of cable procured amounted to about 7000 km (1200 t). The length of strand produced was about 270 000 km. To maintain the precision in the transverse dimensions of the strands and cables to some μm for such a production was an extraordinary industrial achievement [22].

In order to detect deviations from the specified process that could affect the strand and cable quality, strict Statistical Process Control (SPC) was applied at all steps of the production. To ensure uniformity, CERN handled the procurement of the 470 t of Nb-Ti alloy and 26 t of Nb sheet for the six suppliers of strand.

8.3 LHC Cryogenics: Quantum Fluids at Work

Philippe Lebrun

The LHC cryogenic system [23] had to face several unprecedented challenges, both qualitative and quantitative. Here we present the technological solutions developed to meet these challenges reliably, economically and efficiently.

The first challenge is the low temperature of operation of the magnets — 1.9 K — driven by the need to improve the current-carrying capacity of the Nb-Ti superconductor at high field. As this temperature is below that of the “lambda transition” between the two liquid phases of helium (Fig. 8.9), this implies operation in helium II, which exhibits particular transport properties associated with its superfluid state: very low viscosity, high heat capacity and excellent heat conductivity. These properties have been studied in detail since the discovery of superfluid helium by Piotr Kapitsa, John Allen and Don Misener in 1937, and can be profitably applied to the cooling of superconducting magnets [24]. Moreover, provided the superfluid helium can permeate the magnet coils via percolation channels designed into the electrical insulation, the stability of the superconductor




Fig. 8.9. Phase diagram of helium: the magnets operate in baths of pressurized He II, cooled by saturated He II flowing inside the heat exchanger tube.

against thermal disturbances is greatly enhanced, an essential benefit as its specific heat is an order of magnitude smaller at 1.9 K than in normal helium at 4.5 K.

It can be seen in the phase diagram (Fig. 8.9) that the saturation pressure of superfluid helium is below 5 kPa (50 mbar), which raises two concerns: (i) cryostats containing superfluid helium at saturation operate far below atmospheric pressure, with risk of air in-leaks, contamination of the helium and clogging by solid air; (ii) the dielectric strength of gaseous helium is greatly reduced at such pressure, increasing the risk of electrical breakdown. To circumvent these shortcomings, the LHC superconducting magnets operate in baths of “pressurized” He II, i.e. above saturation pressure, in fact in sub-cooled liquid at about atmospheric pressure [25]. The very large (but finite) heat conductivity of helium II ensures that each bath is quasi-isothermal; it cannot however be used to transport heat over long distances in the accelerator tunnel. The LHC magnet cooling scheme is thus composed of a heat exchanger tube [26] threading its way along the magnet chain, in which a small two-phase flow of saturated helium II gradually vaporizes as it absorbs heat from the pressurized helium II baths (Fig. 8.10).




Fig. 8.10. Principle of the LHC superfluid helium cooling scheme [1].

This scheme combines the benefit of magnets operating in static, isothermal superfluid helium close to atmospheric pressure, with that of fixed temperature imposed by the saturated helium II inside the tube. Moreover, the small, stratified two-phase flow inside the tube can be driven solely by gravity (in the sloping LHC tunnel) and/or by flowing vapour, thus avoiding the use of circulation pumps. Finally, in case of resistive transition of a magnet, the rise in temperature of the helium bath leads to dry-out inside the heat exchanger tube, thus limiting thermal propagation to neighbouring magnets.

The second challenge of the LHC cryogenic system, a novelty for hadron accelerators, stems from the dynamic heat loads induced by the circulation of the particle beams, due to synchrotron radiation, dissipation of image currents in the resistive wall of the beam pipes, loss of particles scattered by the residual gas as well as by the collisions in the interaction points (only affecting the sections of the accelerator close to the experiments), and deposition of energy by electrons resonantly accelerated in the electric field of the circulating bunches, the so-called “electron cloud”. In nominal operation, the beam-induced heat loads can be up to 1.7 W/m, i.e. an order of magnitude higher than the allowed heat in-leak at 1.9 K. To reduce their thermodynamic impact, it is therefore necessary to absorb most of these losses on beam screens [27] fitted in the magnet apertures and cooled at higher temperature, from 5 K to 20 K by forced circulation of supercritical helium. The beam screen tubes, made of non-magnetic austenitic stainless steel to resist eddy-current forces at resistive transitions of the magnets, are internally coated with 75 μm of high-conductivity copper to reduce the wall impedance seen by the beam and slotted to act as cryopump baffles, shielding the 1.9 K pumping surface of the cold bore from radiation and particles lost from the beams (Fig. 8.11).




Fig. 8.11. The LHC beam screen.

The large size of the collider, impacting heat loads and refrigeration duties, and its underground implantation, limiting access to eight points around its perimeter, with sector lengths of 3.3 km in-between access points and depths exceeding 100 m below ground, constitute a third challenge. The LHC ring is cooled by eight large-capacity cryogenic plants, located at five points around the perimeter, each normally cooling a 3.3 km sector, but able to serve two adjacent sectors to provide some redundancy at partial load. Each plant (Fig. 8.12) provides a mix of liquefaction and refrigeration duties at 50–75 K and 4.5–20 K, amounting to an equivalent entropic capacity of 18 kW at 4.5 K with excellent efficiency reaching 28% of the Carnot cycle [28]. The cryogenic plants are complemented at their cold end by eight 1.8 K stages producing each 2.4 kW of refrigeration power by expansion to the 1.6 kPa (16 mbar) saturation pressure of helium. Low-pressure helium vapour is compressed by a train of cold hydrodynamic compressors up to a fraction of atmospheric pressure, followed by room-temperature screw compressors operating at sub-atmospheric suction pressure. This arrangement combines good efficiency, limited capital expenditure and compliance to the strongly variable demand resulting from the dynamic heat loads [29].

The issue of cryogen supply and inventory management [30] completes this brief overview. The system contains 135 tons of helium, of which about 60% are in the magnets when the machine is in operation, the rest being shared between the distribution pipework, the cryogenic plants and the minimum reserve in the buffer storage vessels. Helium is procured from the market and delivered to CERN in standard liquid transport containers. Upon warm-up of the machine, the helium must be stored and its purity preserved. Long-term storage is done at room temperature in 250 m³ gas vessels at 2 GPa (20 bar), which can only accept about half of the inventory. Part of the helium can also be stored, for limited amounts of time, in 120 000 litre vacuum-insulated liquid tanks at atmospheric pressure. The rest is re-injected in the market via “virtual storage” contracts with the gas vendors. This strategy allows receiving the amounts needed for operation in due time, while limiting the capital expenditure. Liquid nitrogen is used for precooling the 37 500 t mass of the magnets, a task incommensurable with the installed capacity of the cryogenic plants. For this purpose, 10 000 t of liquid nitrogen supplied by 50 m³ semi-trailers are vaporized at rates of up to 40 t per hour, to produce 4.8 MW of precooling power down to about 100 K, thus permitting to bring the whole machine from room to cryogenic temperature in about three weeks.



Fig. 8.12. Cryogenic plant providing 18 kW @ 4.5 K. Left: compressor station. Right: cold box.

8.4 Current Leads: High Temperature Superconductors to the Fore

Amalia Ballarino

Powering the superconducting magnet circuits requires the transfer of about 3 MA of current from the power converters (at room temperature) to the liquid helium environment. Good electrical conductors are also good conductors of heat, so the devices (current leads) used to transfer the current must be designed to minimize the heat in-leak yet safely carry the current. Conventional current leads achieve this by using the helium vapour that is boiled off at the foot of the lead to cool a combined conductor/heat exchanger. Such “self-cooled” leads transfer about 1 W of heat to the liquid helium bath per kA of current (compared with about 48 W/kA for a lead without vapour cooling). The interest of including a High Temperature Superconducting (HTS) section in the leads is that resistive heating in that section is absent and thermal conduction can be low, so heat in-leak is reduced [31].

Already in 1992, just six years after the discovery of High Temperature Superconductivity [32], conceptual studies were underway at CERN on using HTS in the current leads for the LHC, current leads being an evident application [Box 8.4]. Targeted tests revealed that it should be possible to make reliable HTS leads [33], and in 1999 the decision was taken to use such leads at the LHC. The aim was to obtain a saving with respect to conventional leads of a factor 10 on the required liquefaction power. Over a thousand such current leads would be needed, 64 rated at 13 kA for the main dipoles and quadrupoles, 474 rated at 6 kA for the quadrupoles of the matching sections and 520 rated at 600 A for correctors.

High temperature superconductors (HTS)

Box 8.4

Up to 1986, practical low temperature superconductors (LTS), which first appeared in the 1960s, had been studied in depth [Boxes 4.3 and 8.1], and had been adopted for use in detector and accelerator magnets — and crucially a commodity for the magnetic resonance imaging (MRI) industry. But these materials only work at very low temperature; most magnets are cooled using liquid helium (4.2 K), which is expensive to buy and to keep cold. With this background the unforeseen discovery of ceramic materials (cuprates) that are superconducting at higher temperatures, even beyond that of abundant and cheaply available liquid nitrogen (boiling at 77 K, and a good electrical insulator), caused great excitement in the applied science community. But for accelerator scientists, the superior *high field characteristics* of these so-called “high temperature superconductors” (HTS) at *low temperatures* are far more important, as they offer the promise of magnets operating at 20 T and even higher.

The crystal structure of an HTS compound is highly anisotropic, with a stack of layers of oxides, interspersed with atoms of copper oxide. The superconducting current passes preferentially in the planes between the layers rather than perpendicular to it and the intensity depends on the number of layers. HTS behaviour is not explained by BCS theory. Many types have been found, but only two developed for practical use:

BSCCO, which has two main phases, Bi-2212 and Bi-2223 (the numbers refer to atoms of Bi, Sr, Ca, Cu respectively), have critical temperatures of 85 K and 110 K. These were developed rather quickly, and are referred to as first-generation HTS. The multi-filamentary conductors are made via powder-in-tube (PIT) technology; the matrix is silver to avoid reaction with the precursor powders and let oxygen permeate during the heat treatment process. Bi-2212 can be made in the form of wires that can be cabled (before heat treatment at about 800°C, under pressure). Bi-2223 is produced in the form of tape, about 4 mm wide. The silver matrix provides good electrical stabilization. A variant using a matrix of silver-gold alloy to reduce thermal conductivity is used for HTS current leads (to limit the heat conduction into the cold environment). The cost of silver limits the likelihood of widespread large scale application.

REBCO, which has one main phase, RE-123, where RE refers to rare earth (often replaced by Y for Yttrium) and the numbers to atoms of RE, Ba, and Cu respectively, has a critical temperature (with Y) of 93 K. This second generation conductor is produced in the form of 12 mm wide tape that can be split into the “standard” 4 mm wide type. The superconducting layer is only 1 to 3 μm thick, deposited on a strong metallic substrate; a layer of copper provides electrical stabilization. At 4.5 K and 20 T current densities of 400 A/mm² are currently achieved. Flat cables can be assembled from shaped ribbons cut from wide tape (Roebel cables [Highlight 12.3]). This material is presently very expensive to produce, but is favoured by industry due to a potential for cost reduction.

Magnesium di-boride (MgB_2 , $T_c = 39$ K) is not a true HTS, but is easier to produce in the form of multi-filamentary wire, can be cabled, and is very affordable. Although brittle and more sensitive to magnetic fields than Nb-Ti [Box 4.3], it works efficiently at 20–25 K making it well-suited for power transmission [Highlight 12.3].

The LHC HTS current lead programme was the first large-scale, mission-oriented, application of HTS materials [Box 8.4]. The work started with in-depth studies of various types of superconductor, including bulk BSCCO and YBCO materials [33], before converging on the choice of multi-filamentary BSCCO 2223 embedded in silver-gold alloy (Ag5.3wt%Au) tape. The tape was being actively developed for other projects such as power transmission, for which the BSCCO is embedded in pure silver. Silver is necessary for texturing the superconductor, but pure silver is an excellent conductor of heat. Its thermal conductivity is reduced by doping with gold. Specific methods were developed for optimizing leads combining an HTS section and a gas-cooled section. Taking advantage of the availability in the LHC of helium gas at 20 K, it was decided to operate the HTS section in the range 4.2 K to 50 K, and the gas-cooled section between 50 K and ambient temperature [31]. Prototypes were assembled and tested at CERN, where all the technological processes were developed and qualified. The design was based on the application of standard workshop practice and assembly tolerances, to enable normal engineering firms to tender. A 13 kA lead is shown in Fig. 8.13.

The procurement of the series relied on build-to-print CERN drawings and detailed assembly procedures. The 64 HTS current leads rated at 13 kA were manufactured in industry, while the large series of 6 kA and 600 A leads were made at the Budker Institute of Nuclear Physics (BINP), Novosibirsk. Close technical follow-up, with transfer of know-how acquired during the development activity, assured the quality of the production of the current leads, which was completed in a period of about two years. The required 31 km of BSCCO 2223 tape was specified and procured by CERN, where it was also soldered into stacks. These were measured in cryogenic conditions via an external contract prior to delivery to the lead manufacturers. Contracts were also placed for cold acceptance testing of all leads before integration in the LHC.

Since installation the current leads have undergone several thermal cycles and thousands of electrical cycles, as well as suffering some untoward events (e.g. stopping the flow of helium) generating the resistive transition of the HTS elements [34]. The leads were designed to withstand such events: the HTS material is shunted with sufficient normal material to carry the current during a fast ramp down (time constant 105 s in a dipole circuit) and quench protection thresholds were specifically defined to cope with the requirements of the HTS material.

In view of the success of this project, other users have moved to designing HTS leads for manufacture following standard workshop practice. For example, major aspects of the LHC design have been adopted by ITER for their current leads [35].

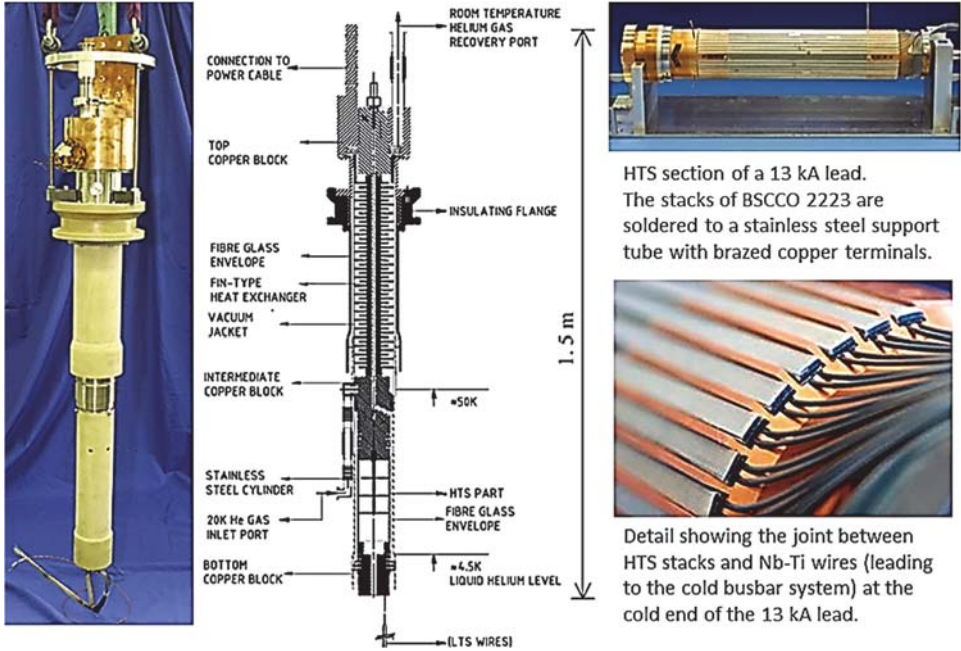


Fig. 8.13. 13 kA HTS current lead for the LHC

8.5 A Pumping Vacuum Chamber: Ultimate Simplicity

Cristoforo Benvenuti

The successful performance of Non Evaporable Getter (NEG) pumping for LEP (Chapter 7) was a strong motivation to envisage also applying a similar solution to the LHC warm sectors. But the LHC warm chambers were too small to host the LEP getter pump. The obvious solution would be to coat the chamber walls with a thin NEG film [36], all the more because such a solution could also be adopted for the chambers in the experimental areas.

It was clear that a thin film NEG coating, if this could be done, would be more effective than the LEP NEG pump for the following reasons: (i) the NEG surface would be much larger, resulting in a lower pressure, and (ii) the coating would inhibit the degassing of the underlying chamber, the major gas source in the accelerator. In addition, the NEG coating pump would be intrinsically safe (no moving parts, no risk of electric short circuits) and simpler (no power supplies, no electric feedthroughs required). To cope with the machine gas load a coating thickness of 1–2 μm would be sufficient using typical getters. However, heating the film for NEG activation implies heating the chamber, therefore activation should be feasible at a temperature that also the Aluminium chambers in the

experimental areas can withstand ($\leq 200^\circ\text{C}$). This temperature was too low for the activation of the available getters, so a suitable NEG alloy had to be found.

The coating process envisaged (cylindrical magnetron sputtering) brought another problem. The coating is obtained by ion bombardment of a cathode. These ions are produced by triggering a discharge in argon or krypton gas (at a pressure of about 10^{-1} Pa), obtained by applying a negative bias voltage of 400 V to the cathode, a rod a few mm in diameter, placed at the centre of the chamber (Fig. 8.14). Producing such cathodes of the different alloys to be explored has a cost and takes time. Furthermore, some of the alloys of the usual getter elements, such as Ti or Zr, are not ductile and the production of the cathode is problematic. To overcome this difficulty, cathodes were made by twisting together wires of pure metals, readily available from industry, as shown in the figure. The atoms extracted by sputtering from the different wires recombine on the chamber walls to provide the desired alloy. In this way the elemental composition of the alloy may be varied and, in addition, for a given choice of metals the relative composition may be changed by changing the diameter of the wires. A smaller wire results in a lower content of the corresponding metal in the final alloy. In this way about 20 different alloys were tested until finally a Ti-Zr-V alloy was found to offer an activation temperature of 180°C , low enough for LHC. For this final choice wires of equal diameter were used, resulting in the same atomic content of these metals [37]. As this technology was not available in industry, the ~ 6 km of room temperature LHC chambers were coated at CERN. Some coated chambers were tested under different conditions showing that extremely low pressures could be achieved (below 10^{-10} Pa). Furthermore, the degassing induced by synchrotron radiation was much lower than for uncoated chambers. Finally, the secondary emission of electrons was reduced by the coating to such a level that the electron cloud was no longer a problem.



Fig. 8.14. Sputtering configuration with a cathode consisting of twisted wires.

Thanks to the NEG coating the vacuum performance has been very satisfactory since the beginning of the LHC operation. This innovation is protected by a CERN patent which has been licensed to several companies and institutes and has since been adopted in accelerators worldwide.

8.6 Vertex Detectors at LHC: In Search of Beauty

Leonardo Rossi

Ordinary matter is made of stable elementary charged particles: up and down quarks and electrons. Heavier replicas of these elementary particles are produced in high energy particle collisions, but their fate is to decay to their stable partners after a lifetime ranging from 5×10^{-25} to 2×10^{-6} s. Two of the heavy replicas, the beauty quark and the tau lepton, are especially interesting, because, belonging together with the top quark to the third family, they may open a window on physics beyond the Standard Model [Box 6.4]. At LHC relativistic beauty quarks and tau leptons are produced abundantly and, having a lifetime of about one picosecond (10^{-12} s), they decay after travelling a few millimetres.

Extrapolating with high precision to the collision points of protons at the LHC to identify and measure decay vertices is therefore of paramount importance. This requires spatial accuracies of ~ 0.01 mm, which are routinely obtained with highly segmented silicon detectors [Highlight 7.7]. The task becomes difficult because of the environment of the LHC proton–proton collisions. At the nominal collision rate of 10^9 interactions per second about 1000 charged particles traverse the vertex detector each time the proton bunches collide at the centre of the experiment. This happens at a frequency of 40 MHz and gives both a large instantaneous hit rate and a large accumulated irradiation dose, which must be considered for long-term operation.

Vertex detectors at LHC must be fast (signal collected in 25 ns), highly segmented (able to cope with > 1000 particles), precise (about 0.01 mm spatial accuracy), radiation resistant (up to 500 kGy) and as thin as practical to minimize loss of information for the downstream detectors. Both ATLAS and CMS chose hybrid silicon pixel detectors [38] as their vertex detector. This technique, nurtured by the RD19 collaboration at CERN in the '90s, had been first employed, on a much smaller scale and with less demanding specifications, in the experiments WA97 and DELPHI. It consists of a sensor directly coupled to its mating electronics read-out chip. The sensor is a matrix of silicon diodes implanted on high resistivity wafers; the read-out chip is a corresponding matrix of front-end electronics with amplifiers and discriminators followed by a coding circuit. The mating between the two parts is obtained through high-density “bump-bonding”,

an operation consisting in the deposition on both matrices of tiny metallic droplets, which are then fused to establish electrical contact between each diode and its front-end electronics channel. The silicon diodes are reverse polarized with a field of $\sim 3 \times 10^3$ V/cm, allowing complete charge collection within 10 ns, to provide a current pulse of ~ 0.5 μ A to the front-end amplifier. This arrangement provides true 3-D hit information and robust pattern recognition capability even in a dense track environment. The very small sensitive volume ($\sim 5 \times 10^{-6}$ cm³) of each pixel minimizes the dark current and provides resistance to large radiation doses.

Figure 8.15(a) shows an individual pixel bump-bonded to its mating front-end channel, while Fig. 8.15(b) shows the exploded view of a module, the basic building block of a pixel detector. ATLAS and CMS were forced to push the hybrid pixel technique far beyond previously established characteristics. Many challenges had to be overcome; the most important being related to instantaneous and integral track density, and to the need to minimize the amount of material. Once installed as the innermost part of each experiment access is difficult, so some degree of redundancy is necessary to minimize the risks of losing decay vertex information.

About 5000 pixel/cm² are required to deal with the space resolution and the doses expected. This required pushing the bump-bonding technology far beyond standard specifications used for the consumer electronics. ATLAS chose 50 μ m \times 400 μ m pixels and digital read-out while CMS decided on 100 μ m \times 150 μ m pixels

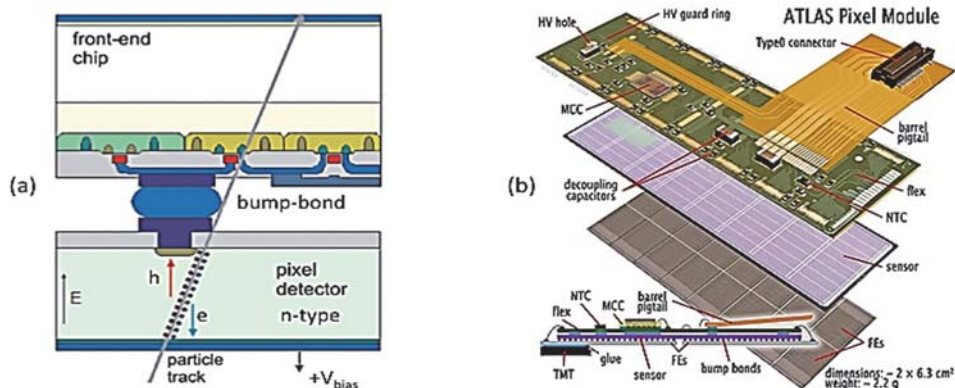


Fig. 8.15. (a) Sketch of one pixel mated to its electronic channel through bump-bonding; (b) exploded view of an ATLAS pixel module where one sensor is bump-bonded to 16 front-end chips. The signals from the 47532 individual pixels are coded in the front-end chips and routed to the read-out electronics via a flexible circuit.

with analogue read-out for charge interpolation, the sensor thickness being ~ 0.3 mm in both experiments. The optimization of charge interpolation pushed CMS to spread the ionization charge between pixels, while ATLAS chose to concentrate most of the charge on one pixel. The different charge spread is obtained through proper orientation of the modules in the magnetic field. The ATLAS design privileges robustness while that of CMS privileges accuracy.

The amplifiers must have a fast rise time, to pass the discriminator threshold (typically set at $\sim 15\%$ of the average signal) within the bunch-crossing period of 25 ns. But the time to return to baseline can be 100 times longer as the probability for a pixel to be hit in a bunch crossing is very low ($< 10^{-4}$). If the return-to-baseline signal slope is constant (as in the case of ATLAS) the signal length can be used to measure the pulse height, to give the specific energy loss in silicon. The front-end electronics has to meet two main specifications: fit all the required functionalities within $< 2 \times 10^{-4}$ cm² and guarantee that the circuit behaviour remains constant (or tuneable) up to doses of 500 kGy. Fortunately, the two requirements could be satisfied by adopting deep-submicron CMOS electronics, made available by IBM through a framework contract with CERN. With gate length of 0.25 μm , not only the designers can arrange ~ 1000 transistors per pixel cell, but the oxide layers become so thin that electrons can tunnel out. The dominant effect of the accumulated dose on the electronics (threshold shifts due to accumulation of charges in the gate and field oxide of CMOS transistors) can therefore be avoided.

The ATLAS and CMS vertex detectors operate in a solenoidal magnetic field of 2 T and 4 T, respectively. They provide three space point measurements over the first ~ 12 cm of the particle trajectories, with each coordinate measured to an accuracy of ~ 0.01 mm in the track bending plane. Each measurement layer represents $\sim 2.5\%$ of a radiation length (X_0), mostly because the high-speed front-end electronics is power-hungry (~ 0.5 W/cm²), and the detector must operate at below -5°C to minimize the radiation damage to the sensors. These requirements determine most of the material in a layer, i.e. conductors to bring the power in and cooling system to take it out, the active part of a layer (sensor and electronics) being $< 20\%$ of the overall X_0 . ATLAS has 1744 identical modules covering ~ 1.7 m² with $\sim 8 \times 10^7$ channels, CMS has 1776 modules of seven different sizes covering ~ 1 m² with $\sim 6.5 \times 10^7$ channels. Both pixel systems meet the demanding LHC specifications with very similar spatial resolution and efficiency, greatly contributing to tagging of events involving beauty quarks and tau leptons.

The ALICE experiment is also equipped with a hybrid pixel vertex detector specifically designed for complex topologies emerging from lead-lead collisions. The environment is quite different: the interaction frequency is only ~ 4 kHz, but

for each collision there may be up to 50 tracks/cm² (mostly low-momentum). The radiation damage is therefore less of an issue, but detector granularity and the amount of material per layer are primary concerns. The pixel size ($50 \times 425 \mu\text{m}^2$) could not be shrunk further, as it is limited both by the bump density and the electronics gate length. ALICE made then a special effort to minimize the amount of material per layer, which was reduced to $< 1\% X_0$, using very thin mechanical supports and distributing signals with flexible circuits made of aluminium tracks on Kapton supports, specially realized at CERN. The ALICE vertex detector consists of a barrel having two layers, at ~ 4 and ~ 8 cm radius, housed in a common carbon fibre support; 120 modules cover $\sim 0.2 \text{ m}^2$ with $\sim 10^7$ channels.

In LHCb the vertex detector plays a key role, since the main physics objective of the experiment is the study of beauty hadrons. Within the acceptance region of the experiment, this study is possible with a luminosity of $\sim 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. LHCb has therefore to sustain a rate that is about halfway between ATLAS/CMS and ALICE. As it is not confronted with the high multiplicity of the Pb-Pb interactions this opens the possibility to use established silicon microstrip detector technology — but requires a special arrangement to optimize vertex resolution. The proton–proton interactions in LHC take place within a cylinder of ~ 0.1 mm radius and ~ 15 cm length. The vertex detector must cover this region and must be as close as possible to the beams to minimize errors due to track extrapolation. The solution adopted has been named VELO (Vertex Locator) and is made of 42 half-moon-shaped silicon microstrip detectors with pitch varying from $\sim 40 \mu\text{m}$ (at the inner radius of 8 mm) to $\sim 100 \mu\text{m}$ (at the outer radius of 42 mm). Each detector is a doublet of 0.3 mm thick sensors, one measuring the radial coordinate and the other the azimuthal coordinate of the tracks. VELO extends for ~ 1 m and a small cut-out in the centre of sensors allows the main LHC beam to pass through unimpeded. The total number of channels is $\sim 1.7 \times 10^5$, all with analogue read-out to optimize spatial resolution through charge interpolation.

The main challenge of the VELO project was to come as close as possible to the interaction region and to minimize the amount of material from the interaction to the first measurement. The VELO system is placed inside the LHC vacuum pipe, but, to protect the integrity of the primary LHC vacuum system, the sensors are separated from the beam volume by a 0.3 mm thick aluminium shield. The silicon detectors must be protected from damage caused by the LHC beam (especially during injection and tuning). The VELO system is therefore remotely retractable by 35 mm and is placed close to the beam only when safe operation is assured.

Vertex resolutions of ~ 0.01 mm in the transverse plane and < 0.1 mm in the beam direction are routinely obtained — performance that contributes to LHCb as a being successful competitor to experiments at e^+e^- beauty factories.

8.7 Large Silicon Trackers: Fast, Precise, Efficient

Luigi Rolandi

Precise tracking is indispensable for any collider experiment. Lepton signatures, based on the combined information of the central tracker and the e-m calorimeter, or the muon system, are privileged signals for potential discoveries allowing to separate new phenomena from the huge background. At the LHC, Z and W bosons are identified via their leptonic decays and are used both to calibrate the detectors and to search for new physics phenomena. About 60% of the energy of jets is due to charged particles: by combining measurement of their momenta with the calorimetric information the jet energy can be determined with good resolution.

The capability to reconstruct detached vertices, combining the information of the central tracker and the vertex detector, allows the identification of hadrons with a b quark. This is an important tool for the precise study of the top quark and for discovery physics in all cases when the new particles have a preferential decay to heavy quarks, such as Higgs bosons or supersymmetric particles [Box 7.2].

The transverse momentum p_t of a charged particle curling in a magnetic field B is derived from the sagitta s of the projection of length L on the plane transverse to the magnetic field

$$s = 750 \mu\text{m} \left(\frac{B}{1 \text{ T}} \right) \left(\frac{L}{1 \text{ m}} \right)^2 \left(\frac{50 \text{ GeV}}{p_t} \right).$$

The transverse momentum resolution, i.e. the momentum measurement precision, is proportional to $1/(BL^2)$. The trajectory is measured at N points with spatial resolution σ , and the momentum resolution scales as σ/\sqrt{N} . With a magnetic field of 4 T and some 15 points having a resolution of $\sigma = 40 \mu\text{m}$ each the momentum resolution is 1% at $p_t = 50$ GeV. With this resolution the width of the peak of a 125 GeV mass Higgs boson decaying to four muons is about 1% and the signal to background ratio in the interval 110–130 GeV is near 2.

The high luminosity and the bunch crossing frequency of the LHC pose severe constraints on the response time of the tracking devices [Box 8.3]. Typically, some 800 particles are produced together with the rare particles of physics interest in the many collisions at each bunch crossing and traverse the detector simultaneously. The pattern recognition algorithm must be able to reconstruct all these tracks belonging to the same collision event. Tracks from different collisions in a bunch crossing can be identified because they originate from a different primary collision

vertex. Since the collision area of LHC has a Gaussian distribution along the beam direction with a σ of about 6 cm, the vertices of the different inelastic collisions are separated by about 1 cm on average. Figure 8.16 shows a 4-muon event selected by the ATLAS analysis searching for the Higgs boson.




Fig. 8.16. Event display of a 4μ Higgs boson decay candidate reconstructed by the ATLAS detector. The muon tracks are shown in red and cross the muon chambers located in the external part of the detector. The insert on the right is a projection perpendicular to the beams showing all charged tracks reconstructed in the same bunch crossing (yellow lines). The insert on the top left shows a region of about 10 cm length where the beams cross with more than 10 reconstructed vertices and their associated tracks.

The complexity of pattern recognition increases as a function of “occupancy”, the average number of hits per event in one elementary detection element. High efficiency pattern recognition requires occupancy below 1% and correspondingly small detection elements. A 10 mm^2 surface covered by a 10 cm long and $100\text{ }\mu\text{m}$ wide silicon detector strip placed at 20 cm from the beam line is hit by charged particles only once every 100 bunch crossings at nominal LHC luminosity.

ATLAS and CMS use 300 micron thick silicon detectors for the innermost part of their tracking systems. In these detectors charge collection time is comparable to bunch crossing period and using fast read-out electronics the pile-up is limited to one bunch crossing. This calls for very large surfaces of the silicon detector. The ATLAS SCT [9] has 15 912 silicon sensors mounted in 4088 modules for a total silicon surface of 63 m^2 and 6 million electronic channels. CMS also uses silicon for the outermost part of detector: the silicon strip tracker [10] has 24 244 sensors mounted in 15 148 modules for a total silicon surface of 198 m^2 and 9.3 million electronic channels. The concept of these large silicon trackers was revolutionary; their development and construction represented an enormous

engineering and logistic advance, considering that the largest silicon trackers existing at that time had a surface of some 6 m^2 [Highlight 5.9].

The silicon sensors for CMS were manufactured for the first time on 6-inch wafers in a standard planar process. This led to significant cost reduction per unit area compared to standard 4-inch wafers, making possible the construction of large ($\sim 85 \text{ cm}^2$) sensors for the outer layers, and greatly reducing their total number.

The precise ($\sim 10 \text{ }\mu\text{m}$) assembly of the sensors and the readout electronics into modules is challenging, and was traditionally performed by trained technicians. This was not possible for the CMS tracker because of the huge number of modules, so the experiment commissioned an automated mounting system based on a precision X–Y coordinate machine with large two-dimensional coverage, video pattern recognition capability and innovative glue-dispensing and pick-and-place tools. These robotic machines could assemble three modules precisely in less than an hour. Several were transferred to collaborating institutes where over 15000 modules were assembled in one year. The total of 25 million electrical connections between strips of two silicon sensors mounted on the same module and between sensors and readout electronics were made using fast (1 Hz) bonding machines.

Particles produced by collisions and secondary interactions in the detectors irradiate the detectors. All elements of the detector, including read-out electronics and cables must therefore be appropriately radiation resistant. The radiation dose (especially high close to the interaction point) degrades the detector response during the 10-year design lifetime. Radiation damage is mitigated by housing the detector in a dry nitrogen gas environment and cooling the sensors to -10°C .

The central trackers of ATLAS and CMS are similar. About 2 m in diameter and 6 meter in length, they enclose the central part of the detector and cover the polar angular range from 10° to 170° . Both operate in a solenoidal magnetic field, 4 T for CMS and 2 T ATLAS. In both trackers the innermost detector layers are built with silicon pixels, with a pixel size of about 0.02 mm^2 to provide space points with $\sim 10 \text{ }\mu\text{m}$ resolution and occupancies smaller than 0.06%. The intermediate layers use silicon strip detectors with high strip density; they provide 6–8 measurements with about $25 \text{ }\mu\text{m}$ resolution of the coordinate perpendicular to the strip and have occupancies below 2%. Some detector layers are mounted with strips rotated by a small angle to provide spatial resolution along the strips. The technology used for the outer layers is different: CMS uses $500 \text{ }\mu\text{m}$ thick silicon strip detectors with coarse pitch providing eight measurement points with resolution of $< 50 \text{ }\mu\text{m}$ while ATLAS uses a Transition Radiation Tracker with 4 mm diameter gas straw tubes providing 35 points with about $130 \text{ }\mu\text{m}$ resolution.

The huge number of front-end electronics channels located on the detectors consumes about 30 kW of electrical power that requires cooling and results in

considerable additional material. The total amount of material varies from about 40% of a radiation length to near 2 radiation lengths at places where the services (cables, pipes) are concentrated. The large amount of material spread along the trajectory of electrons affects the measurement of their energy in the calorimeter. Hadrons are also affected: depending on the angle, some 5%–20% are not reconstructed due to interaction within the tracker volume. The tracker material limits the momentum resolution at low momentum due to multiple scattering.

The intrinsic resolution of the tracking detectors is usually better than that of the detector assembly. The position of the detectors varies with time due to changing environmental conditions (temperature, magnetic field strength, etc. that mainly occur during LHC stops). Alignment procedures are used to measure and monitor their position over time to recover the intrinsic resolution. The corrections from the nominal to the real positions of all detector elements is computed using a large set of reconstructed trajectories of particles from pp collisions and cosmic rays [39]. The statistical error of this method is small, provided a sufficiently large set of tracks is used. A subset of silicon detectors is constantly monitored with optical alignment systems. As stability is imperative, the silicon detectors are mounted on carbon fibre supports with thermal expansion coefficients smaller than 10^{-6} . Depending on the number of independent detectors to be aligned, the simultaneous fit of a large number of parameters can be computationally challenging. The most complex case is the CMS silicon tracker with about 2×10^5 parameters.

The ATLAS and CMS trackers have been operational since the beginning of LHC data taking and perform at design specifications [40, 41]. Less than 3% of the channels are defective. The track quality agrees with simulation to within 2%. Medians of the track-hit residual distributions after alignment indicate that the individual modules are aligned to substantially better than 10 μm . The muon track reconstruction efficiency is at the 99% level. The muon momentum resolution measured with di-muon events agrees with simulation to better than 5%, and the muon momentum scale is known to better than 0.1% [9, 10].

The success of the development and operation of these silicon trackers is the basis for the next generation at the LHC, capable of handling the increase in LHC luminosity with collision rates in excess of 5×10^9 collisions per second. This will require Si trackers with increased radiation resistance, readout electronics with ten times the data transmission capacity and novel approaches for electrical powering.

8.8 Two Approaches to High Resolution Electromagnetic Calorimetry

Tejinder Virdee

The motivation for, and the importance of, high-resolution electromagnetic calorimeters at the LHC arose from the desire to cover the full range of allowed masses for a putative standard model (SM) Higgs boson. In the early 1990s, it was believed that a SM Higgs below a mass of about 100 GeV should be detectable at LEP II, where eventually a limit of 114.4 GeV was obtained.

In the (low) mass interval 100 GeV–130 GeV the dominant decay mode of the SM Higgs boson is into a $b\bar{b}$ pair. It is particularly difficult to cover this region at the LHC, because this decay mode is masked by copious production of $b\bar{b}$ pairs in pp collisions. In the studies performed in the early 1990s the two-photon decay was considered to be the final state most likely to give an early significant signal of a low mass SM Higgs boson at the LHC, provided the energies of these photons could be measured with instruments of an unsurpassed quality. This mode became the benchmark to evaluate candidate devices for electromagnetic calorimetry. Furthermore, in the range 130 GeV–500 GeV, the SM Higgs should be detectable via its decay into two Z bosons that subsequently decay, leading to a final state with either $4e$, or $2e2\mu$, or 4μ , where electrons, as well as muons, play a central role.

In July 2012 the clearest signals for the discovery of the Higgs boson at $m_H \simeq 125$ GeV at the LHC appeared in the same two channels mentioned above.

The design of ATLAS and CMS electromagnetic calorimeters

In the early 1990s, there was no candidate technology for electromagnetic calorimeters that could meet at the same time the diverse requirements such as:

- Very good energy resolution ($\sigma/E \leq 1\%$ at 100 GeV), and accurate enough determination of the angle between the two photons so as not to degrade much the instrumental di-photon mass resolution and hence clearly observe a small and narrow signal on a large background.
- Fine lateral and longitudinal segmentation.
- Fast response: at design operation bunches of protons cross every 25 ns and some 20 pairs of protons would be colliding in every bunch crossing, in addition to the interaction of interest. Thus, the products of the collision of interest may be confused with those from other collisions in the same bunch crossing. This confusion can be reduced by using highly granular calorimeters with good time resolution ($\ll 25$ ns), giving low occupancy (fraction of detector elements that contain useful information), at the price of a large numbers of detector channels.

- Radiation resistance to ensure stable response throughout the lifetime of the experiment. At the smallest angles ($\sim 6^\circ$), the radiation levels reach a dose of 100 kGy during a ten year operation of the detector. This level of radiation severely limited the choice of possible detection media.

When a high energy photon or electron traverses a dense (high-Z) medium, they give up their energy by creating a shower of secondary photons and electrons and positrons [Box 6.3]. The parameters that determine the thickness needed and the lateral extent of the electromagnetic showers are the radiation length (X_0) and Molière radius (R_M), respectively. Typically, $25X_0$ are needed to contain high-energy showers and R_M should be as small as possible. The electrons and positrons either excite the atoms (whereupon scintillation light is emitted through de-excitation) or ionize the atoms (creating electron-ion pairs which separate rapidly in an externally applied high electric field). The amount of scintillation light or electron-ion pairs generated is proportional to the energy of the original particles.

Several R&D programmes were launched, particularly on scintillating inorganic crystals (in view of improving their radiation resistance) and on noble liquids, working as ionization detectors (in view of improving their speed and granularity).

The ATLAS and CMS Collaborations chose two very different instruments, the sampling “accordion” liquid argon calorimeter of ATLAS [9] and the homogeneous PbWO_4 crystal calorimeter of CMS [10] yet meeting overall the same physics goal. Technically they differ in many ways as shown in Table 8.2.

Table 8.2. Some parameters of interest for the e-m calorimeters of the ATLAS and CMS experiments.

Parameter	ATLAS	CMS
Intrinsic energy resolution		
Stochastic term (a/\sqrt{E} , E in GeV)	10%	3%
Constant term (c)	0.2%	0.5%
e-m cluster energy equivalent of noise GeV (b/E , E in GeV)	0.25	0.20
Effective radiation length (cm)	2.3	0.9
Effective Molière radius (cm)	4.5	2.2
Segmentation		
Lateral (major sampling) $\Delta\eta \times \Delta\phi$	0.025×0.025	0.017×0.017
Number of Longitudinal Samplings	4	1
Speed of response	Max. drift time = 400 ns	Time constant = 20 ns
Electronics peaking time	50 ns	45 ns
Material in front of ECAL ($ \eta = 0$ (1.7)) X_0 (excluding solenoid in the case of ATLAS)	0.35 (1.35)	0.35 (1.5)
Sensitivity to radiation	Relatively insensitive	Sensitive (see text)

It is clear that a compromise had to be struck by both ATLAS and CMS in terms of the technical characteristics — some which lead to a “better” performance in one experiment or the other. Nevertheless, both experiments have achieved a similar overall physics performance as evident from the estimated values of the *expected* $H \rightarrow \gamma\gamma$ signal significance from the Run 1 data of 4.6σ and 5.1σ for ATLAS and CMS, respectively.

A further challenge was the design of front-end electronics that had to have low energy equivalent of electronics noise, low-power consumption, be radiation resistant, and capable of covering a dynamic range (~ 16 bits) from a few tens of MeV to several TeV (energy of the largest anticipated signals, for example from a putative Z' decaying in an electron-positron pair).

As opposed to momentum measurements in a magnetic field, the relative resolution of calorimeters improves with energy, typically as its square root. The energy resolution is usually parameterized as: $\sigma/E = a/\sqrt{E} + b/E + c$, where a is the “stochastic” term, b the “noise” term and c the “constant” term. The challenge is to simultaneously keep all three terms small. For homogenous electromagnetic calorimeters it is possible to achieve a value of a few percent whereas in sampling electromagnetic calorimeters $a \sim 10\%$ is normally achievable. However, at high energies, a major limitation arises from the size of the constant term that is related to uniformity and time-stability of response, and importantly the accuracy of the inter-cell calibration. In practice, for large systems it is challenging to reach overall accuracies significantly below 1% at high energies (~ 100 GeV or higher), even for the most performant electromagnetic calorimeters.

ATLAS accordion liquid argon sampling electromagnetic calorimeter

Liquid argon calorimetry offers several attractive performance features, including a response that is intrinsically linear, stable over time and tolerant to high levels of radiation. In liquid argon the secondary charged particles in showers produce electron–ion pairs. The signal develops on the electrodes as soon as the freed charges start drifting in the liquid, i.e. instantaneously. The charges drift for a time of ~ 400 ns in a 2 mm liquid argon gap at a field of ~ 2 kV/cm. A shorter effective response time can be achieved by electronic pulse shaping [Highlight 4.10].

The challenge of a liquid argon calorimeter for use at the LHC lay in the implementation of fine segmentation and fast speed of response. Conventionally, in sampling calorimeters, the absorber and active planes had been oriented perpendicular to the incident particles. But in such a geometry, it is difficult to realise fine lateral segmentation, with small size cells and towers that should point towards the interaction region. Longitudinal subdivision should be implemented without introducing insensitive regions and a large number of penetrating interconnections which degrade the performance. To overcome these shortcomings a novel absorber-

electrode configuration, known as the “accordion” (Fig. 8.17(a) and (b)), was introduced, in which the particles traverse the chambers at angles of around 45°. This accordion geometry enables implementation of high granularity and seamless coverage. The electrodes can easily be grouped into towers at the front or at the rear of the calorimeters, such that adjacent modules can be tightly packed, without introducing cracks between modules. The absorber is made of lead plates, clad with thin stainless steel sheets for structural rigidity as shown in Fig. 8.17(a) and (b), sandwiching copper-clad Kapton electrodes positioned centrally by precision-cut honeycomb. Details of the sampling structure can also be seen in Fig. 8.17(b). Longitudinal subdivision and pointing towards the interaction region is achieved by the way the electrodes are etched (see Fig. 8.17(b)).

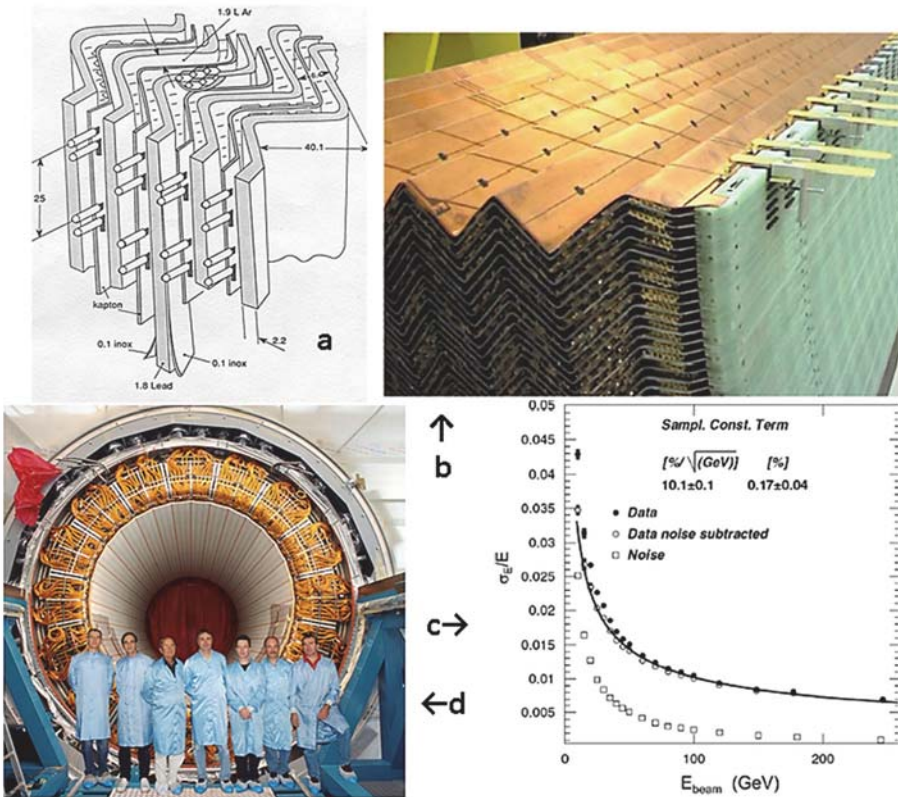


Fig. 8.17. Clockwise from top left: (a) The conceptual drawing of the ATLAS “accordion” liquid argon calorimeter, (b) partial view of a calorimeter module, (c) the energy resolution measured using a prototype for electrons in test beam, (d) the completed barrel part of the liquid argon calorimeter before lowering into the underground experiment cavern.

Considerable engineering challenges were overcome as the several innovative features resulted in major production issues. The most difficult was the fabrication of large electrodes, ~ 2 m in length containing nearly 1000 resistive pads each. This problem was solved by careful monitoring of the production on-site by experts from the collaboration. Approximately 20 000 m² of honeycomb spacers have been used to maintain the flexible electrodes in the centre of the gap between absorbers. To avoid major problems with the high-voltage behaviour of assembled modules, a rigorous and careful cleaning procedure for all parts, especially the honeycomb, had to be implemented. Radiation-tolerant electronics had to be produced for all components in the cavern, including the front-end for nearly 200 000 channels, housed near the cold-to-warm signal feed-throughs.

The performance of the liquid argon calorimeter in test beams is illustrated in Fig. 8.17(c). The completed barrel part of the liquid argon calorimeter before lowering into the ATLAS underground experiment cavern is illustrated in Fig. 8.17(d). The result related to the physics benchmark reaction — that of the diphoton mass spectrum, clearly showing the presence of the Higgs boson, is illustrated in [13].

CMS lead tungstate scintillating crystal calorimeter

Lead tungstate scintillating crystals offer several advantageous performance features. Compact size — a length of 23 cm is sufficient to absorb the energy of high energy electrons and photons, good tolerance to high levels of radiation, and fast response — 90% of the scintillation light is emitted within 20 ns. But it also has a number of drawbacks. The light yield is low and strongly temperature dependent ($-2\%/^{\circ}\text{C}$). The low yield brings the need for photo-devices that can provide gain in high magnetic field — ruling out photomultipliers for the barrel region. *Silicon avalanche photodiodes* (APDs) were eventually chosen because they can provide sufficient gain in magnetic fields. The amount of light collected is affected by loss of transparency due to colour centres created by radiation.

In 1994, tests on a small number of lead tungstate crystals coupled to APDs yielded very encouraging results, convincing the CMS collaboration to adopt this technique. Some examples of the technical challenge faced by CMS are given here. Between 1993 and 1998 much research and development work was carried out to improve the transparency and the radiation hardness of the crystals. The relative amounts of lead oxide and tungsten oxide in the crystals had to be optimized; an affordable means of ensuring adequate purity of the raw chemicals had to be set up; and any remaining defects in the crystals structure had to be compensated by identifying and introducing specific dopants. Although the intrinsic scintillation mechanism is not affected by radiation, tests in particle beams revealed a loss of transparency dependent on the absorbed dose rate. It was also found that the

crystals show recovery of transparency when the irradiation ceased, revealing a dynamic behaviour of the yield of collected light. CMS is able to maintain the good performance by deploying a powerful laser-light monitoring system to precisely track, and correct for, changes in the transparency of each and every crystal. The mass manufacture of crystals took place between 1998 and 2008. The crystals were manufactured largely in Russia, but with some in China. In Russia a high-volume crystal growing (Fig. 8.18(a) and (b)) and processing capacity had to be put in place in a factory previously in the military-industrial sector. The last of $\sim 75\,000$ crystals were delivered in March 2008, and the full electromagnetic calorimeter was installed by mid-August 2008.




Fig. 8.18. Clockwise from top left: (a) An ingot of lead tungstate crystals just after the 3-day growing cycle, (b) a set of cut and polished crystals, and one in the foreground with two APDs mounted, (c) the energy resolution measured using a prototype for electrons in test beam (d) the completed barrel part of the calorimeter just after installation in the CMS underground experiment cavern.

Engineering challenges included the stabilization of the temperature of about 100 t of crystals and 130 000 APDs to 0.1°C (achieved by water cooling) and stabilizing the bias voltage of the APDs to some tens of mV. This was necessary because the APD gain has large temperature ($-2.4\%/^{\circ}\text{C}$), and bias voltage dependence ($3\%/V$). The goal was to keep these contributions to the energy resolution at the ‰ level. The crystals were supported by specially developed, lightweight carbon-fibre or glass-fibre structures and integrated with high performance electronics. The performance of the lead tungstate calorimeter in test beams is illustrated in Fig. 8.18(c). Figure 8.18(d) shows the installed barrel part of the crystals calorimeter in the CMS underground cavern. The result related to the physics benchmark reaction (that of the di-photon mass spectrum) clearly showing the presence of the Higgs boson, is illustrated in [13].

Concluding remarks

It is remarkable, if not surprising, that the ATLAS and CMS Collaborations have chosen two very different techniques for their electromagnetic calorimeters, yet addressing and meeting overall the same physics goals. This did not happen by design or by a committee decision. It happened, because even in these 3000-plus strong collaborations the physicists and engineers use their experience, familiarity, conviction and “taste” in selecting an experimental approach and shaping the decisions. In each of the two collaborations members were familiar with the respective technology and were able to convince their collaborations that the technology could be honed to achieve the required performance. “Required” means achieving a measurement accuracy which approaches the intrinsic performance limited only by the fundamental aspects. There was little room for even the slightest degradation, lest the discovery of the Higgs boson be jeopardized. That both collaborations reported, on the same day of July 4, 2012, the discovery of the Higgs boson with the same experimental significance is a tribute to the foresight, intuition, perseverance and dedication of many hundreds of students, technicians, engineers and physicists worldwide.

8.9 Multigap Resistive Plate Chamber: Chronometry of Particles

Crispin Williams

The ALICE experiment at the CERN LHC is designed to study nuclear matter under extreme conditions of temperature and density produced in collisions of heavy ions. Knowledge of the identity of the particles emerging from such a collision is crucial for this research. That this is not a trivial task can be appreciated by observing Fig. 8.19, which displays a typical heavy ion collision. Among the thousands of particles, which are, for example, the protons, pions or kaons?

Most of these particles are produced with momenta less than a few GeV/c. For these particles a frequently used technique is the “Time-of-Flight” (ToF) velocity measurement. The curvature of the tracks in a magnetic field determines the particle momentum (i.e. “relativistic” mass \times velocity). If the velocity is also measured, its mass can be determined; hence its identity is inferred. All that is needed is to surround the central tracker with a detector capable of precisely measuring the arrival time of the particles. This, together with the time the collision occurred and the distance to the point of collision gives the time-of-flight (and its velocity) and thus, with the momentum measurement, its mass. The detector must be highly segmented, so that each of the thousands of particle tracks has its own individual time stamp. This detector must also work in a strong magnetic field.

In the late 1990s ALICE started to develop its experimental concept. ToF-identification was an essential part, but none of the existing techniques were remotely close to meeting the novel and stringent specifications: a new detector was needed. The breakthrough was the “The Multigap Resistive Plate Chamber” [42].

The principle of this detector is as simple as it is elegant. It is a derivative of the time-honoured gaseous parallel plate detector. A single gap Resistive Plate Chamber (RPC) is shown in cross section in Fig. 8.20. A strong electric field is applied across this gap. A charged particle traversing the gap ionizes the gas molecules producing positive ions and electrons. Typically, a few primary ionisation collisions occur in the 2 mm gap. The electrons are accelerated in the strong electric field, collide with and further ionize the gas molecules leading to




Fig. 8.19. A lead-lead collision observed in the ALICE experiment. Each of the lines represents a particle track. (Source: ALICE).

an avalanche [Highlight 4.8]. The movement of these charges induces signals at the pick-up electrodes. Registering the induced signal provides a measurement of the position of the particle and the arrival time. The time resolution is quite respectable being typically 2×10^{-9} sec (2 ns). But this was not good enough for ALICE where a time resolution of < 0.1 ns (100 ps) was required.

Progress came through a detailed understanding of the processes influencing the time resolution. Only avalanches that are initiated close to the cathode can grow to a sufficient size able to produce a detectable signal. Fluctuations in its initial growth give rise to time jitter (i.e. uncertainty) of the observed signal. Decreasing the width of the gap decreases the time jitter; but this also decreases the efficiency of the detector. The *Multigap Resistive Plate Chamber (MRPC)* is a solution invented and developed by the LAA project at CERN. The gas is divided into several small gas gaps (five in this example) with a single pair of readout electrodes, as shown in Fig. 8.21. The voltage is applied to the outer surfaces of the two external plates. The internal plates are allowed to float electrically, which considerably eases construction. The MRPC is made by stacking sheets of resistive material (window glass in the case of the ALICE TOF barrel). The gap size is controlled by running a fishing line (!) across the surface. This is the short version of a long story of heated debates between experts, on whether such a device would work and how exactly it should be constructed. Details are revealed in [42].

A crucial aspect of precise timing is the readout technique for a highly segmented large-area detector. Suppression of electronic noise is mandatory in order not to spoil the timing resolution. This is best achieved with a ‘differential’ readout, where both pickup electrodes are separated from the electrical ground of the system, and connected to a differential amplifier. An ultra-fast differential

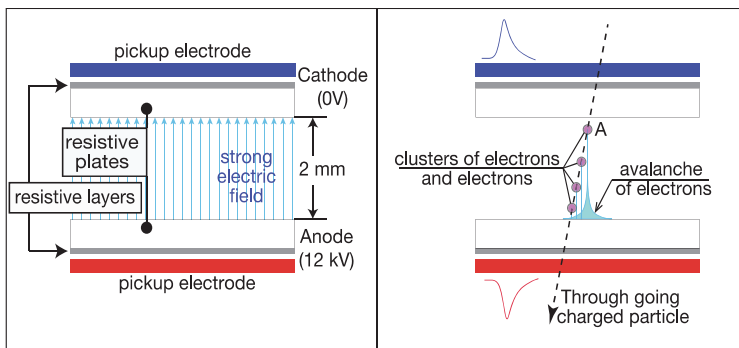


Fig. 8.20. Cross section of single 2 mm gap RPC. A traversing charged particle creates clusters of electrons and ions in the gas. The electrons avalanche across the gap due to the high electric field.

amplifier-discriminator, known as the NINO, was developed in 250 nm CMOS technology and has become the standard for precise timing. Together with these developments, MRPCs have been shown to be excellent high performance TOF detectors meeting the requirement of a systems time resolution below 100 ps [42].

For ALICE a TOF cylindrical detector of radius 3.7 m with a length of 7 m was built surrounding the Time Projection Chamber (TPC). The MRPCs have ten 250 micron gaps; the detector planes are segmented into $2.5 \times 3.5 \text{ cm}^2$ pads, resulting in a total of 160,000 signal channels. All MRPCs are operated at the same applied voltage and the same discriminator threshold. Figure 8.22 shows a typical plot of $\beta = v/c$ (measured by the TOF) versus the momenta measured by the TPC. At low momenta, the dE/dx measurements from the TPC provide additional identification power [Box 5.2]. A similar detector based on these MRPCs has been built for the STAR heavy ion experiment at RHIC. A TOF barrel built from MRPCs will be installed for the BES III experiment at IHEP, Beijing. The CBM experiment at GSI/FAIR is planning to build a 120 m^2 wall of MRPCs for TOF measurements [43]. There are also proposals to build a neutron detector with precise timing with MRPCs embedded between iron plates. MRPCs are also being considered as detectors of 511 keV gamma photons for Positron-Emission-Tomography (PET), where its excellent time resolution would improve the imaging quality.




Fig. 8.21. Schematic cross-section of a Multigap Resistive Plate Chamber (MRPC). Voltage is only applied to the outer surface of the two external plates. All internal plates are electrically floating.




Fig. 8.22. Distribution of particles with given velocity $\beta = v/c$ (measured by the TOF barrel) versus their momenta. (Source: ALICE).

8.10 The LHCb RICH: Lord of the Cherenkov Rings

Carmelo D'Ambrosio and Olav Ullaland

The LHCb experiment is dedicated to the study of CP violation [Box 3.4] and the rare decays of heavy flavours [Box 6.4]. A key to these studies is the identification of the many particles produced in the complex decay chains of particles containing the b-quark. The Ring Imaging Cherenkov (RICH) detector system provides charged particle identification over a wide momentum range, from 2 to 100 GeV/c.

A major requirement for charged hadron identification in a flavour-physics experiment is the reduction of combinatorial background. Many of the interesting decay modes of b- and c-flavoured hadrons involve hadronic multibody final states. At hadron colliders like the LHC, the most abundantly produced charged particle is the pion. The heavy flavour decays of interest typically contain a number of kaons, pions and protons. It is therefore important when reconstructing the invariant mass of the decaying particle to select the charged hadrons of interest in order to reduce the combinatorial background. Knowledge of the particle identity also helps to distinguish between final states of otherwise identical topology. An example is the two-body hadronic decays, $B \rightarrow h^+h^-$, where h indicates a charged hadron. In this case there are many contributions, as illustrated in Fig. 8.23, including $B^0 \rightarrow \pi^+\pi^-$, $B_s^0 \rightarrow K^+K^-$ and other decay modes of the B^0 , B_s^0 and Λ_b [44].




Fig. 8.23. Invariant mass distribution for $B \rightarrow h^+ h^-$ decays in the LHCb data before the use of the RICH information (left), and after applying RICH particle identification (right). The signal under study is the decay $B_0 \rightarrow \pi^+ \pi^-$, represented by the turquoise dotted line. The contributions from different b-hadron decay modes ($B_0 \rightarrow K\pi$ red dashed-dotted line, $B_0 \rightarrow 3\text{-body}$ orange dashed-dashed line, $B_s \rightarrow KK$ yellow line, $B_s \rightarrow K\pi$ brown line, $\Lambda_b \rightarrow pK$ purple line, $\Lambda_b \rightarrow p\pi$ green line), are eliminated by positive identification of pions, kaons and protons and only the signal and two background contributions remain visible in the plot on the right. The grey solid line is the combinatorial background [44].

A measurement of the Cherenkov photon emission angle (θ_C) of a charged particle, together with the knowledge of the refractive index of the Cherenkov medium ($n(\lambda)$), determines its velocity (v), where $v/c = 1/(n \cos\theta_C)$. A Cherenkov detector is not a stand-alone detector. An accurate determination of the mass is obtained by coupling this information with the knowledge of the momentum and the trajectory of the particle [Boxes 3.2 and 6.3]. As the Cherenkov light is produced along a cone at a polar angle θ_C relative to the charged particle trajectory, a spherical or parabolic mirror is used to project these photons in a well-defined, ring-like pattern onto a photon detector plane. This is the basic principle behind a Ring Imaging Cherenkov detector [Highlight 7.8].

The LHCb RICH Particle Identifier comprises two separate detectors, RICH 1 and RICH 2, (Fig. 8.24). The first, RICH 1, is upstream of the magnet and covers the momentum range from ~ 2 to 60 GeV/c . RICH 2 is placed just in front of the calorimeters and covers momenta up to 100 GeV/c . The main task of the two detectors is to efficiently identify charged hadrons, as shown in Fig. 8.23.

Some first generation RICH detectors had problems with ease and reliability of operation. Common to these devices was detection of photons in the far-ultraviolet (VUV) range, with wavelength between 160 nm and 200 nm. The refractive index increases rapidly in this region and uncertainty in determination

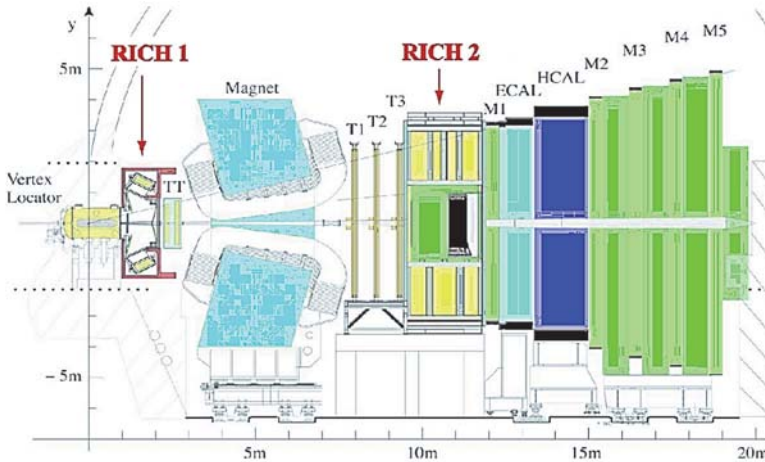


Fig. 8.24. Side view of the LHCb experiment.

of the Cherenkov angle is dominated by the chromaticity of the radiator gas.^a Operating in the VUV region also meant that impurities, such as O₂, H₂O or CO₂, at the ppm level in the gas would absorb the photons. A very high purity gas was therefore required.

The reason for the detection of photons in the VUV region was simply that no single detector for visible photons with a spatial resolution better than a centimetre had yet been developed. Experiments had to rely on specially adapted gas chamber detectors [Highlight 4.8], which required a photon energy in the VUV region to ionize the gas and generate the signal. This was acceptable, as the number of Cherenkov photons grows quadratically with the photon energy so that sufficient numbers could be detected. However, calculations showed that detecting photons in the visible range (VIS) would decrease the chromatic effects on the overall error and thus pave the way to much more precise detectors. Specifically, the error on the Cherenkov angle is given by σ_θ/\sqrt{N} where σ_θ is the single photon Cherenkov angle error and N the number of detected photons. When working in the VUV, experiments tried to maximize N . In the visible range, giving up photons while simultaneously decreasing σ_θ results in a much improved Cherenkov ring resolution. Such a change of perspective was not easy, as it required a complete re-appraisal of the available tools, instruments, hardware and software, involving all RICH system components, from particle-sensitive hardware to electronics and data acquisition and processing. For the hardware the following steps were taken:

- Choice of gases having intrinsically low chromaticity with good sensitivity to particle momenta between 2 GeV/c and 100 GeV/c. The gases chosen were

^a RICH detectors can also use liquid or solid radiators. Here only gas radiators are considered.

C_4F_{10} in RICH 1 and CF_4 in RICH 2. These gases are relatively easy to handle and, by working in the visible range, acceptable gas impurities are in the percent range rather than the few parts per million (or billion) for VUV.

- Development of efficient single-photon detectors working in the VIS with high spatial resolution (pixel sizes from 0.25 to 2.5 mm), called pixel-HPD (HPD for Hybrid Photon Detector). Their development had been started in the 1990s and was pursued at CERN, triggered by the need for precise photon detectors for scintillating fibres. The basic idea was to insert in a vacuum envelope a silicon pixel detector and electronic chip to directly detect and digitize photo-electron signals produced by the Cherenkov photons. This is not an easy task and it took more than 10 years to realize the first mature example of this completely new class of photon detector. They feature high spatial resolution, an impressive signal to noise ratio and can be assembled in arrays. There are 484 pixel-HPDs in the RICHs covering a surface of $\sim 3.6 \text{ m}^2$.
- Development of thin, light, precise and highly reflective mirrors, together with their support and alignment structures, to focus the photons on the pixel HPDs. Mirrors had to be thin to minimize material traversed by the particles. A mirror array contains hundreds of components totalling a reflective surface of $\sim 19 \text{ m}^2$. They are aligned to a precision of 0.05 mrad to form an almost perfect sphere and minimize the error on σ_θ of the single photon Cherenkov angle.

RICH 1 features carbon fibre composite mirrors and RICH 2 has 6 mm thin glass mirror arrays. They are held in place either by carbon fibre (RICH 1) or by aluminium honeycomb (RICH 2) structures with Polycarbonate micrometric adjustments. With these technologies the mechanical structures for these large RICH detectors represent only $\sim 2.5\%$ and $\sim 10\%$ of a radiation length, respectively. To give an idea of its size, the RICH 2 volume is larger than 100 m^3 . After about ten years of development, all the mirrors show excellent alignment stability and tilting errors lower than 0.1 mrad after online corrections. The result of this effort is best illustrated with an example: Fig. 8.25 shows the “world record” distribution of σ_θ , obtained with RICH 2 in 2012 [45].

The ability to detect sharp Cherenkov rings in the LHCb RICH system allows to detect and identify several rings simultaneously. This process is referred to as pattern recognition. A new “global search algorithm” based on statistics and probability was developed. Instead of searching for rings, the probability for each detected photon to belong to a possible ring is evaluated. By constructing and minimizing a likelihood function, the mass and hence the identity of the particles is assigned. The RICH system can handle thousands of photons produced by hundreds of particles up to a million times per second and deliver information about the particles crossing the detector and their underlying physics processes.




Fig. 8.25. Single-photon Cherenkov angle precision of 0.68 mrad for LHCb RICH 2 [45].

For 2020, LHCb is planning to further improve the RICH detectors aiming for operation at 10 times the present collision rate of $3 \times 10^7 \text{ s}^{-1}$ at LHCb and 40 times the readout rate. Beyond that, one can anticipate a new “*quantum leap*” producing the RICH detectors of the future, when new components, ideas and a good dose of human ingenuity will all synergize, thoughtfully and skilfully in the conducive laboratory environment of CERN.

8.11 Signal Processing: Taming the LHC Data Avalanche

Philippe Farthouat

When the concept of the LHC program began to take shape, physicists and engineers were confronted with a seemingly unconquerable world of experimentation: proton bunches colliding 40 million times a second, entailing a billion collisions, 10^{11} particles, and data volume approaching 100 TB (10^{14} Bytes) per second [Box 8.3]. It would take a worldwide-coordinated R&D programme to possibly find solutions to the design and construction of the detectors matching such requirements. Particularly worrisome was the question whether the signal and data processing systems for the hundred million detector channels could be developed and built. It took courage, sweat, stamina, ingenuity, collaboration with industry and a fair dose of luck to write a happy end to this story, as told below.

Each of the several large sub-detectors of an LHC experiment has its custom-designed signal processing system, taking care of the necessary analogue signal conditioning (amplification, filtering, etc.), analogue-to-digital conversion, digital signal processing (data compression, digital filtering, etc.), and interface to the data acquisition system of the experiment. The LHC characteristics, the trigger and data acquisition functional models, the harsh environment and large number of channels are factors that drove the definition and design of the readout electronics and led to some spectacular technical developments [46].

Each experiment contains a large number of readout channels: up to 100 million for the inner trackers, several hundred thousand for the calorimeters and about one million for the muon spectrometers. These astronomical numbers could not be handled without integrating a large fraction of the readout electronics directly into the detector (front-end electronics, FEE), and extensive multiplexing before transmission to the off-detector electronics (back-end electronics, BEE) located in service rooms at distances of up to 100 m away from the detectors. The high density of signal channels demanded the FEE to be implemented with Application Specific Integrated Circuits (ASICs). Located on the detector, they are subject to extreme environmental conditions, in particular radiation and high magnetic fields.

In nominal conditions, the LHC delivers bunches of protons colliding in the centre of the experiments every 25 ns and ideally, one would record all the activity in the detectors every 25 ns. However, with an event data size of a 1 MB or more, a readout and storage bandwidth of several tens of terabytes per second would be required. In addition, most of the events are of little interest for physics studies. It took — a famous example — about 10^{15} collisions before the Higgs boson discovery was announced, based on a few hundred particles identified as Higgs boson decays. Clearly, registering all the information is not the way to proceed.

In the strategy adopted only the events of potential physics interest are selected out of the billion collisions per second. This selection is probably the most critical part of the data processing system. The experiment teams developed a multi-level selection (“trigger”) approach. The level-1 trigger uses hardware processors, implemented off-detector in state-of-the-art Field Programmable Gate arrays (FPGAs) and exploits only a subset of information, e.g. from the calorimeters and the muon spectrometers in the case of ATLAS and CMS. Selection criteria are energetic muons, electrons, photons or events with a high multiplicity of particles, signalled by a large energy deposit in the calorimeter, and provide the basis for a decision to select or reject an event. It delivers a YES (Level-1 accept signal [L1A]) or NO decision every 25 ns within a fixed decision time interval

(“latency”) in the range $2\ \mu\text{s}$ – $3.2\ \mu\text{s}$ and at an average rate of up to 100 kHz for ATLAS and CMS, and a latency of $4\ \mu\text{s}$ and 1 MHz trigger rate for LHCb.

During this latency, about 100 new bunch crossings occur and the signals corresponding to these bunch crossings must be stored in so-called level-1 buffers. The correct data retrieval from the level-1 buffers relies on the fact that the L1A has a fixed latency time of N bunch crossings with respect to the interaction time. When a L1A signal is received by the front-end electronics, the data of the event giving rise to this L1A were therefore stored in the level-1 buffer location corresponding to N bunch crossings earlier. These data are then selectively retrieved and transmitted to the back-end electronics. However, the complete information of an event is spread over several thousand locations and must subsequently be correctly assembled. For this purpose, the entire readout system is synchronized with the 40 MHz bunch crossing (BC) clock of the LHC machine.

Following this transmission to the off-detector electronics the subsequent selection proceeds in the high-level trigger. This is software-based and implemented in arrays (“farms”) of processors (PCs) accessing the data through industrial-standard networks. In the high-level trigger either a subset of data of potential interest as flagged by level-1 trigger information or complete data of an event are analysed in more detail (Event Filter), ultimately reducing the flow of events to a few hundred Hz. These events are permanently stored for further physics analysis. The high-level trigger needs up to 1 s computation time per event. Farms of ~ 2000 PCs with more than 30 000 cores working in parallel perform the processing tasks. As an example, the data flow and selection for ATLAS is shown in Fig. 8.26.

There is a large radiation background when the LHC operates and electronics is strongly affected by this radiation. As an example, threshold voltages in CMOS ASICs are sensitive to the total ionizing dose, the current gains of bipolar transistors are sensitive to displacement damages caused by non-ionizing energy loss, and energetic particles may produce single-event upsets in the digital electronics. Very nasty destructive events such as latch-up or gate burnout may also occur. The radiation levels in the centre of the experiments reach up to 1 MGy and several 10^{14} 1-MeV-neutron-equivalent/cm² for a 10 years of operation.

The front-end ASICs were designed either in radiation-hard processes or in deep sub-micron CMOS technologies with special design rules, such as enclosed layout transistors (ELT), a technique invented years ago, but revived and extensively used for the LHC experiments. All sensitive digital elements were implemented in triplicated logic with procedures such that an upset in one circuit element would not disturb the circuit functionality. Commercial Off-The-Shelf Components (COTS) could be used for detectors receiving lower radiation after thorough radiation qualification.

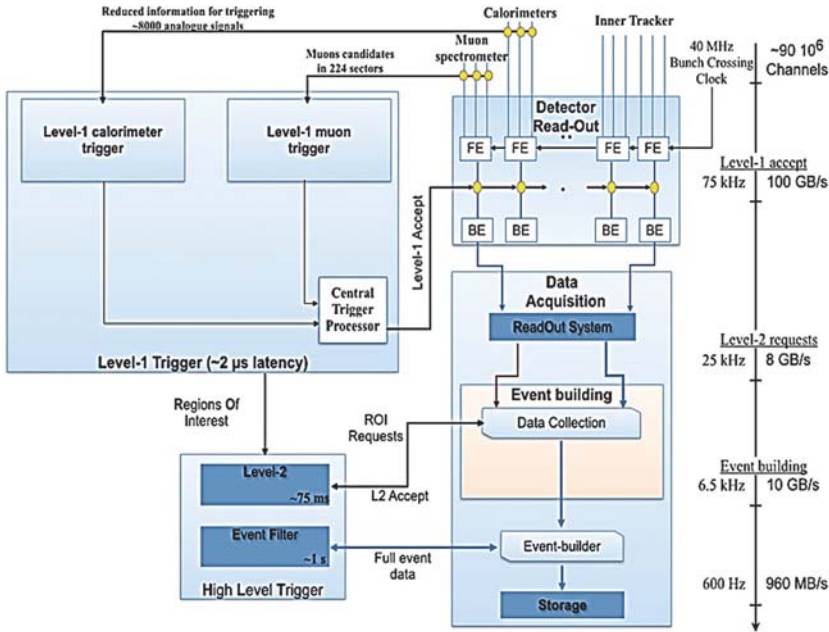


Fig. 8.26. Example of readout, trigger and data acquisition architecture (ATLAS). The front-end electronics (FEE) records the detector signals at the bunch crossing frequency (40 MHz). The information is retained during the Level-1 trigger latency. The data corresponding to a Level-1 Accept are sent to the back-end electronics (BEE) through dedicated front-end links and then, after data formatting and processing, to the data acquisition system. A selected subset of data is sent to the level-2 trigger which makes a decision in about 75 ms. The data fragments of the events to be kept are assembled in the event builder. The complete events are then processed in the Event Filter (making a decision in about 1 s) and the fraction of accepted events are stored for later physics analysis. The required data bandwidths at each level as well as the rates involved are shown.

Developing this radiation-tolerant electronics remained for many years the biggest hurdle towards functioning signal processing systems. Ultimately, solutions were adopted in collaboration with semiconductor foundries, exploiting in new ways the features that were introduced for standard commercial products.

Besides the complex analogue processing, which is different for the various detectors of the experiment and which processes input signals as low as a few thousand electrons, the front-end electronics contains the level-1 buffer. Usually it also handles the analogue to digital conversion (effective resolution up to 10 bits and dynamic range up to 64 000) and the data compression, reducing the data volume to be transmitted and hence the number of readout links.

In most cases digital optical links were selected because of their high bandwidth and very low material mass (optical fibres are by far lighter than copper links). It required the radiation hardness qualification of optical devices (edge emitting lasers and PIN diodes), the development of low-mass packages and for the CMS tracker the development of an analogue optical link. A high-speed radiation-hard serializer able to feed an optical link at 1.6 gigabit per second has also been developed and extensively used. The installed density of optical links is unique: 20 000 are installed in ATLAS and more than 100 000 in CMS.

The power dissipated by the front-end electronics required special attention because it occurs in confined places; bringing power inside the detector requires thick copper cables. The total power consumption of the front-end electronics is ~ 350 kW for each ATLAS and CMS. Operated at low voltages, tens of kA are needed. In order to reduce the size of the cables penetrating the detector volume, DC-DC converters located at places with low magnetic field in the experimental caverns, followed by linear voltage regulators in the detector volume are used. This again required the development, often in collaboration with industry, of radiation tolerant devices that could be installed in the experimental caverns.

The development and deployment of this signal processing system with its large number of radiation-hard ASICs (1.5 million ASICs have been installed in ATLAS and CMS) represents a unique achievement, even more so, considering the radiation levels, to which the readout electronics of the inner detectors are exposed.

This success, unique in so many ways, is a tribute to a world-wide collaboration of dedicated physics, engineering institutes and industry, with CERN assuming a substantial load and the coordination of these activities.

8.12 Giant Magnets for Giant Detectors

Alain Hervé, Peter Jenni and Thomas Taylor

All four major experiments, ALICE, ATLAS, CMS, and LHCb, feature large spectrometer magnets [47]. These devices are the costliest single components of the detectors, and many of the other components rely on their structure for support. In contrast with LEP, where 3 of the 4 detectors were assembled in a “garage position”, the LHC experiments were installed in their final position on the beam. When the detectors are fully assembled the magnets are difficult to access; it follows that any significant repair would require a long shutdown, affecting the entire machine. This gave rise to two important constraints: (i) all the magnets had to be delivered and installed on time, and (ii) the magnet systems had to be designed for utmost reliability. Although the magnets were to be designed,

manufactured, tested and installed under the technical and financial responsibility of the collaborations, they also had to satisfy constraints imposed by CERN: powering, cooling, cryogenics and controls would be supplied and maintained by the laboratory. For reasons of economy this infrastructure had to be standardized. Common safety standards also had to be applied. CERN therefore allocated experienced technical staff to the collaborations to help with the magnet projects, to ensure compatibility with the interfaces and to provide the essential leadership for ensuring success (Chapter 11).

Following the experience with LEP, CERN was already associated with studies for the LHC detector magnets before the collaborations gelled and the magnet designs were frozen. It was realized that the major detectors would require larger superconducting magnets than had been made before, and that they would have to run at higher current, probably around 20 kA. The architecture of the conductors for magnets for experiments — cabled Nb-Ti stabilized with pure aluminium — was well established in previous work (at CEA, Saclay, CERN, INFN, KEK and RAL), but only for currents of up to about 5 kA. The superconducting cable could be similar to that developed for the LHC dipole magnets [Highlight 8.2], but co-extrusion of the cable with the aluminium stabilizer would call for gaining access to very large extrusion presses, and extensive testing to verify that the superconductor was not degraded by overheating during the process. It was also essential to ensure good bonding between the cable and the high purity aluminium, and to be able to control the quality of the bond on production lengths. Collaborations were set up with ETH Zürich and CEA Saclay to develop and validate such conductor for use in windings for a practical magnet [48].

For the winding being envisaged for the powerful CMS magnet, designed for operation at 4 T, the high-purity Al used to stabilize the superconductor was not sufficiently strong to withstand the stress on the conductor when the magnet was excited to full field (4 T). One solution proposed was to surround the pure Al/Nb-Ti insert described above with aluminium alloy, to be co-extruded simultaneously [49]. This was found to be impractical as it required heating to $> 500^{\circ}\text{C}$ causing massive degradation to the superconductor. The alternative was to solder or weld Al alloy flanges to the pre-extruded insert — a process that became the subject of another development in collaboration with ETHZ. Working together with industry it was found that by using electron-beam welding the process was both feasible and practical and resulted in a superconductor degradation of less than 5%.

The strong conductor issue was approached in a different way for the ATLAS central solenoid (CS). A doped pure aluminium [50], sufficiently strong but with low resistivity was developed for this magnet (an in-kind contribution of Japan).

The conductors for the CMS and ATLAS magnets are shown in Fig. 8.27.




Fig. 8.27. Comparison of the superconductors used for ATLAS and CMS.

CMS magnet

The CMS magnet (Fig. 8.28) consists of a 6 m diameter, 12.5 m long solenoid coil surrounded by a steel yoke [51]. The coil is special in that unlike previous magnets of this type that featured single layer coils, it has four layers. For ease of manufacture (and transport) it is divided into five sub-units; the coils are wound on the inside of the support cylinder [52]. The 12 000 t yoke is made of plates between which detectors are installed to track muons emanating from the interaction. The steel, which serves as a general supporting structure for the detector, carries the return flux and is magnetized, so the tracks are curved and momentum can be determined. The coils are cooled by liquid helium in parallel tubes between manifolds top and bottom, forming a thermo-siphon configuration, following that of the ALEPH magnet at LEP. The magnet was designed for 4 T, up to which field it was tested, but it is run at 3.8 T to prolong its lifetime. Besides being presently the single magnet having the largest stored energy (2.6 GJ) in the world, it also features a record ratio of stored energy to cold mass (about 12 kJ/kg) for this type of magnet. This is tolerable due the inductive heating of the external mandrel if a quench should occur, rapidly spreading the quench to the whole winding. In normal circumstances half the energy is extracted in case of a quench.

The magnet was assembled and tested in the surface building above the cavern at LHC point 5, and lowered into position in the cavern in eleven pieces of mass between 600 and 2000 t. This procedure was based on the experience gained with lowering the support tube for the L3 experiment at LEP [Highlight 7.10]. The magnet itself has been working flawlessly since its installation in 2008, but some auxiliary equipment has had to be repaired.




Fig. 8.28. Schematic exploded view of the CMS solenoid magnet.

ATLAS magnet system

ATLAS features a unique configuration of four large superconducting magnets [53]. Most imposing is the 26 m long air-cored barrel toroid (BT), presently the largest magnet in the world, consisting of eight rectangular superconducting coils, each enclosed in its own cryostat. The unusual architecture of this experiment makes it possible to decouple precise muon momentum measurements, made beyond the calorimeters, from momentum measurements in the inner tracking system. The toroidal configuration (Fig. 8.29), completed with two endcap toroids (ECT) [54], provides increasing field integral with decreasing forward angle, corresponding to nearly constant sensitivity in the measurement of transverse momentum over the angular momentum of particles from a proton collider. The low-mass superconducting air-cored solution provides high muon resolution thanks to minimal multiple scattering. A thin, 2 T superconducting solenoid (CS) [55], 5.8 m long, diameter 2.4 m, surrounds the inner tracking detector. Its insulating vacuum enclosure also serves the LAr electromagnetic calorimeter [Highlight 8.8] — a novel approach that minimizes material thickness. For the muons the two independent momentum measurements are combined to get the best resolution. The configuration and size of the system made this project challenging with regard to cost containment and installation. Major CERN contributions were the modification of the original design to simplify (to make more affordable) the BT coil casing of the barrel toroid, and changes to enable powering the three toroidal magnets in series, rather than separately. The stored energy of the BT is 1.1 GJ, and that of the ECTs is 0.5 GJ.




Fig. 8.29. Schematic view of the toroidal magnet system including services connections. The eight BT coils are visible with the two castellated ECTs inserted on both ends.

The three toroids are cooled by pumping supercritical helium through pipes attached to the coil casing, requiring a powerful pump. The dedicated refrigeration plant is located in a side-cavern.

The complete magnet system was commissioned in the cavern in 2008, and has worked reliably during all data taking periods. It is the largest and most complex superconducting magnet system ever built for a particle physics experiment.

LHCb magnet

Like ALICE, the LHCb detector also features a forward dipole. The first design of the magnet was based on the use of superconducting coils. However, at that stage the design was conceptual and the LHCC magnet advisory group (MAG) judged that it would be difficult to flesh out the design and build and test the magnet in the allotted time. Moreover, a quick study showed that sufficient performance could be obtained with resistive coils, avoiding the necessity of local cryogenics, by making a judicious adjustment to the steel yoke [56]. The idea was to maximize the integrated bending field by having sloping poles (Fig. 8.30), in a similar way to a previous magnet that had been designed at CERN for use at the ISR. The LHCC strongly urged the collaboration to pass to the resistive version, which they did. The magnet was designed at CERN and the coils were specified and purchased via competitive tender, with a strong technical follow-up by CERN. The total mass of the magnet is 1450 t; it consumes 4.2 MW when run at an excitation current of 5.85 kA to deliver an integrated field of 3.6 T m.




Fig. 8.30. The LHCb magnet (left) and the Alice dipole magnet (right) assembled, ready for testing.

ALICE solenoid and dipole magnet

This experiment uses the large solenoid designed and constructed at CERN for the LEP L3 experiment [57], complemented with a new large dipole magnet installed downstream for analysing particle (in particular di-muon) trajectories emerging at small angles [58]. The solenoid was overhauled and fitted with an improved water cooling system. The dipole magnet was designed by JINR (Dubna) in collaboration with CERN; steel for the yoke was an in-kind contribution from Russia. The coils were finish-designed at CERN and purchased via competitive tender. The delivery schedule was held thanks to strong technical follow-up and support by CERN (Fig. 8.30). The total mass of the magnet is 800 t; it consumes 5 MW at an excitation current of 6 kA, to deliver a bending strength of 3 T m.

References

1. L. Evans (ed.), *The Large Hadron Collider: a Marvel of Technology* (EPFL Press, 2009).
2. G. Brianti, The Large Hadron Collider project: historical account, *Phys. Rep.* **403–404** (2004).
3. M. Jacob (ed.), *ECFA-CERN Workshop on a Large Hadron Collider in the LEP Tunnel, Lausanne and Geneva*, CERN 84-10-V-1, ECFA-84-085-V-1 (CERN, Geneva, 1984). http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/16/034/16018264.pdf;
M. Jacob (ed.), *ECFA-CERN Workshop on a Large Hadron Collider in the LEP Tunnel, Lausanne and Geneva*, CERN 84-10-V-2, ECFA-84-085-V-2 (CERN, Geneva, 1984). http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/16/034/16034963.pdf.
4. J.H. Mulvey (ed.), *Workshop on Physics at Future Accelerators, La Thuile, Italy*, CERN-1987-007 (CERN, Geneva, 1987). <https://cds.cern.ch/record/179654?ln=en>.
5. Report of the long-range planning committee, CERN/1658. August 1987, see CERN/1670 (1987).

6. *LHC Design Report. Vol. 1: The Main Ring*, CERN-2004-003 (CERN, Geneva, 2004). <http://dx.doi.org/10.5170/CERN-2004-003-V-1>.
7. P. Lebrun, Superfluid helium as technical coolant, *LHC Project Report 125*, (1997).
8. J.H. Mulvey (ed.), *The Feasibility of Experiments at High Luminosity at the Large Hadron Collider: report of the High-Luminosity Study Group to the CERN Long-Range Planning Committee*, CERN-1988-002 (CERN, Geneva, 1988). <http://dx.doi.org/10.5170/CERN-1988-002>.
9. ATLAS Collaboration, The ATLAS Experiment at the LHC, *JINST 3*, S08003 (2008).
10. CMS Collaboration, The CMS Experiment at the LHC, *JINST 3*, S08004 (2008).
11. LHCb Collaboration, The LHCb Detector at the LHC, *JINST 3*, S08005 (2008).
12. ALICE Collaboration, The ALICE Experiment at the LHC, *JINST 3*, S08002 (2008).
13. P. Jenni and T. S. Virdee, The discovery of the Higgs Boson at the LHC, in *60 Years of CERN Experiments*, ed. H. Schopper and L. Di Lella, (World Scientific, 2015) ISBN 978-9814644143.
14. R. Perin, Status of LHC programme and magnet development, *IEEE Trans. Appl. Supercond.* **5**, 189–195 (1995).
15. D. Leroy, R. Perin, G. de Rijk and W. Thomi, Design of a high-field twin-aperture superconducting dipole model, *IEEE Trans. Magnetics*, **24**, 1373–76 (1987).
16. R. Perin, pp. 71-92 in *Proc. CAS-CERN Accelerator School: Superconductivity in Particle Accelerators, Hamburg, Germany, May 1995*, S. Turner (ed.), CERN-1996-003 (CERN, Geneva, 1996). <http://dx.doi.org/10.5170/CERN-1996-03.71>.
17. P. Fessia *et al.*, Selection of the cross-section design for the LHC main dipole, *IEEE Trans. Appl. Supercond.* **10** (2000).
18. A. Siemko, *Safeguarding superconducting magnets*, pp. 35–37, CERN Courier, September 2013.
19. F. Savary *et al.*, Description of the main features of the series production of the LHC main dipole magnets, *IEEE Trans. Appl. Supercond.* **18**, No. 2, (2008).
20. V. Chohan, Operation for LHC cryomagnet tests: concerns, challenges & successful collaboration, *2007 Asian Particle Accelerator Conference (APAC)*, RRCAT, Indore, India, pp. 84–88 (2008).
21. J. Billan *et al.*, Twin Rotating Coils for Cold Magnetic Measurements of 15 m Long LHC Dipoles, *IEEE Trans. Appl. Supercond.* **10**, 1422–26 (2000).
22. D. Leroy, Review of the R&D and supply of the LHC superconducting cables, *IEEE Trans. Appl. Supercond.* **15**, 1454–60 (2006).
23. P. Lebrun, Cryogenics for the Large Hadron Collider, *IEEE Trans. Appl. Supercond.* **10**, 1500–06 (2000).
24. G. Claudet and R. Aymar, Tore Supra and helium II cooling of large high-field magnets, *Adv. Cryo. Eng.* **35A**, 55–57 (1990).
25. P. Lebrun and L. Taviani, pp. 453-476 in *Proc. CAS-CERN Accelerator School: Superconductivity for Accelerators, Erice, Italy, April 2013*, R. Bailey (ed.), CERN-2014-005 (CERN, Geneva, 2014). <http://dx.doi.org/10.5170/CERN-2014-005.453>.
26. P. Lebrun, L. Serio, L. Taviani and R. van Weelderren, Cooling strings of superconducting devices below 2 K: the helium II bayonet heat exchanger, *Adv. Cryo. Eng.* **43A**, 419–42 (1998).
27. V. Baglin, P. Lebrun, L. Taviani and R. van Weelderren, Cryogenic beam screens for high-energy particle accelerators, *Proc. ICEC24 – ICMC 2012*, Fukuoka (Cryogenics and Superconductivity Society of Japan, 2013) pp. 629–634.

28. S. Claudet, P. Gayet and U. Wagner, Specification of four new large 4.5 K refrigerators for the LHC, *Adv. Cryo. Eng.* **45B**, 1269–1276 (2000).
29. S. Claudet *et al.*, Specification of eight 2400 W at 1.8 K refrigeration units for the LHC, Proc. ICEC 18 Mumbai (Narosa Publishing, 2000) pp. 215–221.
30. D. Delikaris *et al.*, Cryogenic operation methodology and cryogen management at CERN over the past 15 years, *Proc. ICEC 22 – ICMC 2008* (KIASC, Seoul 2009) pp. 33–38.
31. A. Ballarino, Large capacity current leads, *Physica C* **468**, 2143–2148 (2008).
32. J.G. Bednorz and K.A. Müller, Possible high T_c superconductivity in the Ba–La–Cu–O system, *Zeitschrift für Physik B Condensed Matter*, **64** (2), 189–193 (June 1986).
33. A. Ballarino, HTS current leads for the LHC magnet powering system, *Physica C* **372–376**, 1413–1418 (2002).
34. A. Ballarino, Operation of 1074 HTS current leads at the LHC: overview of three years of powering the accelerator, *IEEE Trans. Appl. Supercond.* **23** (3), 4801904 (2013).
35. P. Bauer *et al.*, R&D towards HTS current leads for ITER, *IEEE Trans. Appl. Supercond.* **19** (3), 1500–1503 (2009).
36. C. Benvenuti, Non-Evaporable getters: from pumping strips to thin film coatings, *Proc. EPAC 98* (1998) pp. 200–204.
37. C. Benvenuti, P. Chiggiato, F. Ciccoira and Y. L’Aminot, Non-evaporable getter films for UHV applications, *J. Vac. Sci. Technol. A* **16** (1), 148 (1998).
38. L. Rossi, P. Fischer, T. Rohe and N. Wermes, *Pixel Detectors, From Fundamentals to Applications* (Springer, 2006).
39. CMS Collaboration, Alignment of the CMS tracker with LHC and cosmic ray data, *JINST* **9**, P06009 arXiv:1403.2286 (2004).
40. ATLAS Collaboration, Performance of the ATLAS inner detector track and vertex reconstruction in the high pile-up LHC environment, *ATLAS-conf-2012-042*, (2012). <http://cdsweb.cern/record/1435196>
41. CMS Collaboration, Description and performance of track and primary-vertex reconstruction with the CMS tracker, *JINST* **9**, P10009 (2014).
42. E. Cerrón Zeballos *et al.*, A new type of resistive plate chamber: the multigap RPC, *Nucl. Instr. & Meth. A* **374**, 132 (1996).
43. M.C.S. Williams, Particle identification using time of flight (2012), *J. Phys. G: Nucl. Part. Phys.* **39** 123001 (2012).
44. R. Aaij *et al.*, (LHCb collaboration), Measurement of b-hadron branching fractions for two-body decays into charmless charged hadrons, *JHEP* **10**, 037 (2012).
45. M. Adinolfi *et al.*, The LHCb RICH Collaboration, *Eur Phys. J. C* **73**, 2431 (2013).
46. D. Green, *At the Leading Edge: The ATLAS and CMS LHC Experiments* (World Scientific, 2010).
47. T.M. Taylor, The magnets for the LHC experiments, *IEEE Trans. Appl. Supercond.* **10** 1, 342–346 (2000).
48. I.L. Horvarth *et al.*, The CMS conductor, *IEEE Trans. Appl. Supercond.* **10** 1, 395–398 (2000).
49. J.C. Lottin *et al.*, Conceptual design of the CMS 4 T solenoid, *Adv. in Cryo. Eng.* **41** 819–825 (1996).
50. K. Wada *et al.*, Development of high-strength and high RRR aluminium–stabilized superconductor for the ATLAS solenoid, *IEEE Trans. Appl. Supercond.* **10** 1, 373–376 (2000).
51. A. Herve, The CMS detector solenoid, *IEEE Trans. Appl. Supercond.* **10** 1, 389–394 (2000).

52. P. Fabbriatore *et al.*, The construction of CMS superconducting coil, *IEEE Trans Appl. Supercond.* **14** 2, 552–555 (2004).
53. H.H.J. ten Kate, The superconducting magnet system for the ATLAS detector at CERN, *IEEE Trans. Appl. Supercond.* **10** 1, 347–352 (2000).
54. D.E. Baynham *et al.*, ATLAS end cap toroid magnets cryostat design, manufacture and installation at CERN, *IEEE Trans. Appl. Supercond.* **14** 2, 522–525 (2004).
55. A. Yamamoto *et al.*, The ATLAS central solenoid, *Nucl. Instr. & Meth. A* **584** 53–74 (2008).
56. M. Losasso *et al.*, Test and field map of LHCb dipole magnet, *IEEE Trans. Appl. Supercond.* **16** 2, 1700–1703 (2006).
57. F. Wittgenstein *et al.*, Construction of the L3 magnet, *Proc. 11th Int. Conf. on Magnet Tech. (MT-11)*, 130–135 (1989).
58. D. Swoboda *et al.*, Status of the ALICE magnet system, *IEEE Trans. Appl. Supercond.* **12** 1, 432–437 (2002).