NATURE|Vol 448|19 July 2007|doi:10.1038/nature06078

Detector challenges at the LHC

Steinar Stapnes[1,2]

The best way to study the existence of the Higgs boson, supersymmetry and grand unified theories, and perhaps the physics of dark matter and dark energy, is at the TeV scale. This is the energy scale that will be explored at the Large Hadron Collider. This machine will generate the energy and rate of collisions that might provide evidence of new fundamental physics. It also brings with it the formidable challenge of building detectors that can record a large variety of detailed measurements in the inhospitable environment close to the collisions points of the machine.

Four main experiments have been designed and constructed for the Large Hadron Collider (LHC) machine: ATLAS, CMS, LHCb and ALICE. ATLAS and CMS are large general-purpose experiments. LHCb will study *b*-quark systems, produced predominantly in the forward direction, and ALICE is designed specifically for studies of heavy-ion collisions (see page 302).

This review focuses on the challenges — related to tracking, calorimetry, muon detection, triggering and data acquisition — faced by the designers and builders of the general-purpose detectors ATLAS and CMS, as well as some of the particular issues for the more specialized detector LHCb.

Experimental measurements at the LHC

Inside the 27-km ring of the LHC, bunches of 10^{11} protons will collide 40 million times per second to provide 14-TeV proton–proton collisions at the LHC design luminosity of 10^{34} cm^{-2}s^{-1} (see page 285). With an inelastic proton–proton cross-section of about 100 mb (~10^{-25} cm^2), this gives 25 events per bunch crossing, or a total rate of 10^9 inelastic events per second. This means that around 1,000 charged particles will emerge from the collision points every 25 ns, within a volume defined by $|\eta| < 2.5$, where pseudorapidity, η, is related to the polar angle relative to the beam axis, θ, by $\eta = -\ln[\tan(\theta/2)]$ (Fig. 1a).

This formidable luminosity and interaction rate are necessary, because the expected cross-sections are small for many of the LHC benchmark processes (such as Higgs production and decay, and some of the processes needed to search for and explore new physics scenarios such as supersymmetry and extra dimensions). They also raise a serious experimental difficulty; every candidate event for new physics will, on average, be accompanied by 25 inelastic events occurring simultaneously in the detector.

The very nature of proton–proton collisions creates a further difficulty. The cross-sections for producing jets of particles, through quark interactions governed by quantum chromodynamics (QCD), are large compared with the rare processes being sought — several orders of magnitude larger, even for jet production above 500 GeV. Therefore, one has to look for characteristic experimental signatures, such as final states involving one or more leptons, or photons, or with missing transverse energy or secondary vertices, to avoid being drowned by QCD background processes. Searching for such final states among already rare events imposes further demands on the luminosity needed and on the detectors' particle identification capabilities.

Specific requirements for the LHC detector systems[1–8] have been defined using a set of benchmark processes that covers most of the new phenomena that one might hope to observe at the TeV scale. The first such process is the production of the standard-model Higgs boson, which is particularly important because there is a wide range of decay modes possible, depending on the mass, m_H, of the Higgs boson. If m_H is low (less than 180 GeV, which is twice the mass of the *Z* boson), the natural width is only a few MeV, and the observed width will be defined by the instrumental resolution. The dominant decay mode into hadrons is difficult to isolate, because of the QCD background. Therefore, the two-photon decay channel will be important, as will other channels, including associated productions such as *ttH*, *WH*, *ZH* (see page 270), for which a lepton from the decay of the accompanying particle will be used for triggering and background rejection.

Above 130 GeV, Higgs decay into *ZZ* (one Z being virtual when m_H is below the *ZZ* threshold), with its four-lepton final state, will be the most interesting channel. Above 600 GeV or so, *WW* or *ZZ* decays into jets or into states involving neutrinos (leading to missing transverse energy because the neutrinos are undetected) are needed to extract a signal; for m_H close to 1 TeV, it becomes necessary to tag 'forward' jets, in the region $2 > |\eta| > 5$, from the *WW* or *ZZ* fusion production mechanism. The Higgs might not even be of the standard-model variety; detection of some of the Higgs particles of the 'minimal supersymmetric extension of the standard model' (MSSM) would require very good sensitivity to processes involving tau leptons and *b* quarks.

If supersymmetric particles such as squarks and gluinos are produced at the LHC, their decays would involve cascades that always contain a lightest stable supersymmetric particle, or LSP (if R-parity is conserved). Because the LSP interacts very weakly with the detector, the experiments would measure a significant missing transverse energy in the final state. The rest of the cascade results in a number of leptons and jets.

Several new models, motivated by theories of quantum gravity, propose the existence of extra dimensions[1–4]. In terms of experimental signatures, the emission of gravitons that escape into extra dimensions would result in missing transverse energy; furthermore, Regge-like excitations could manifest themselves as *Z*-like resonances with ~TeV separations in mass. Other experimental signatures could be anomalous, high-mass dijet production and mini-black-hole production with very spectacular decays involving democratic production of jets, leptons, photons, neutrinos, and *W* and *Z* bosons.

The LHC will also allow studies of QCD, electroweak and flavour physics. For example, *t* quarks will be produced at the LHC at a rate measurable in hertz. New, heavy gauge bosons (W' and Z') could be accessible at masses up to 5–6 TeV. To study their leptonic decays, high-resolution lepton measurements and charge identification are needed

[1]Department of Physics, University of Oslo, 0316 Blindern, Oslo, Norway. [2]Department of Physics, CERN, CH-1211 Geneva, Switzerland.

©2007 Nature Publishing Group

up to transverse momenta of a few TeV. Another new-physics signature could be jets produced with very high transverse momenta; if quarks are composite rather than fundamental particles, deviations from QCD expectations in the jet cross-sections could result.

Detector concepts

The necessary detection capabilities for all of these experimental signatures lead to a stringent set of design requirements. First of all, owing to the experimental conditions at the LHC, the detectors need fast, radiation-tolerant electronics and sensor elements. In addition, high granularity of the detectors is needed to be able to handle the particle fluxes and to reduce the influence of overlapping events. Good charged-particle momentum resolution and reconstruction efficiency in the tracking system, and in the inner tracker specifically, are essential. For efficient high-level triggering and offline tagging of taus and *b* quarks (which decay a short distance from the primary interaction vertex at which they are produced), pixel detectors close to the interaction region are needed to observe the distinctive secondary vertices.

Two key requirements are good electromagnetic calorimetry for electron and photon identification and measurements, and full-coverage hadronic calorimetry for accurate jet and missing-transverse-energy measurements. Likewise, good muon identification and momentum resolution over a wide range of momenta, and the ability to determine unambiguously the charge of muons with high transverse momentum, are essential. Finally, triggering the event readout on the presence of leptons, jets, photons or missing transverse energy — and at low transverse-momentum thresholds to ensure high efficiencies for most of the physics processes of interest at LHC — is an absolute requirement to reduce the data rate (a few hundred collisions out of the 40 million taking place every second are finally kept) to a level that can be handled offline.

The layout of the ATLAS detector[1,2] is shown in Fig. 1c. It has an inner, thin, superconducting solenoid surrounding the inner detector cavity, and large, superconducting, air-core toroids, consisting of independent coils arranged with an eight-fold symmetry, outside the calorimeters. The inner detector comprises a large silicon system (pixels and strips) and a gas-based transition-radiation 'straw' tracker. The calorimeters use liquid-argon technology for the electromagnetic measurements and also for hadronic measurements in the endcaps of the detector. An iron/scintillator system provides hadronic calorimetery in the central part of the detector. The muon system is based on gas detectors and has precise tracking chambers and trigger chambers for a robust and efficient muon trigger. The ATLAS detector has a radius of 13 m and is 46 m long, with a weight of 7,000 tonnes.

The design of CMS[3,4] is shown in Fig. 1b . The main distinguishing features of CMS are a high-field solenoid housing a full silicon-based inner tracking system (pixels and strips), a fully active, scintillating crystal electromagnetic calorimeter, and a compact scintillator/brass hadronic calorimeter. Outside the solenoid, there is a hadronic 'tail-catcher' in the central region, and an iron-core muon spectrometer sitting in the return field of the powerful solenoid, with tracking chambers and trigger

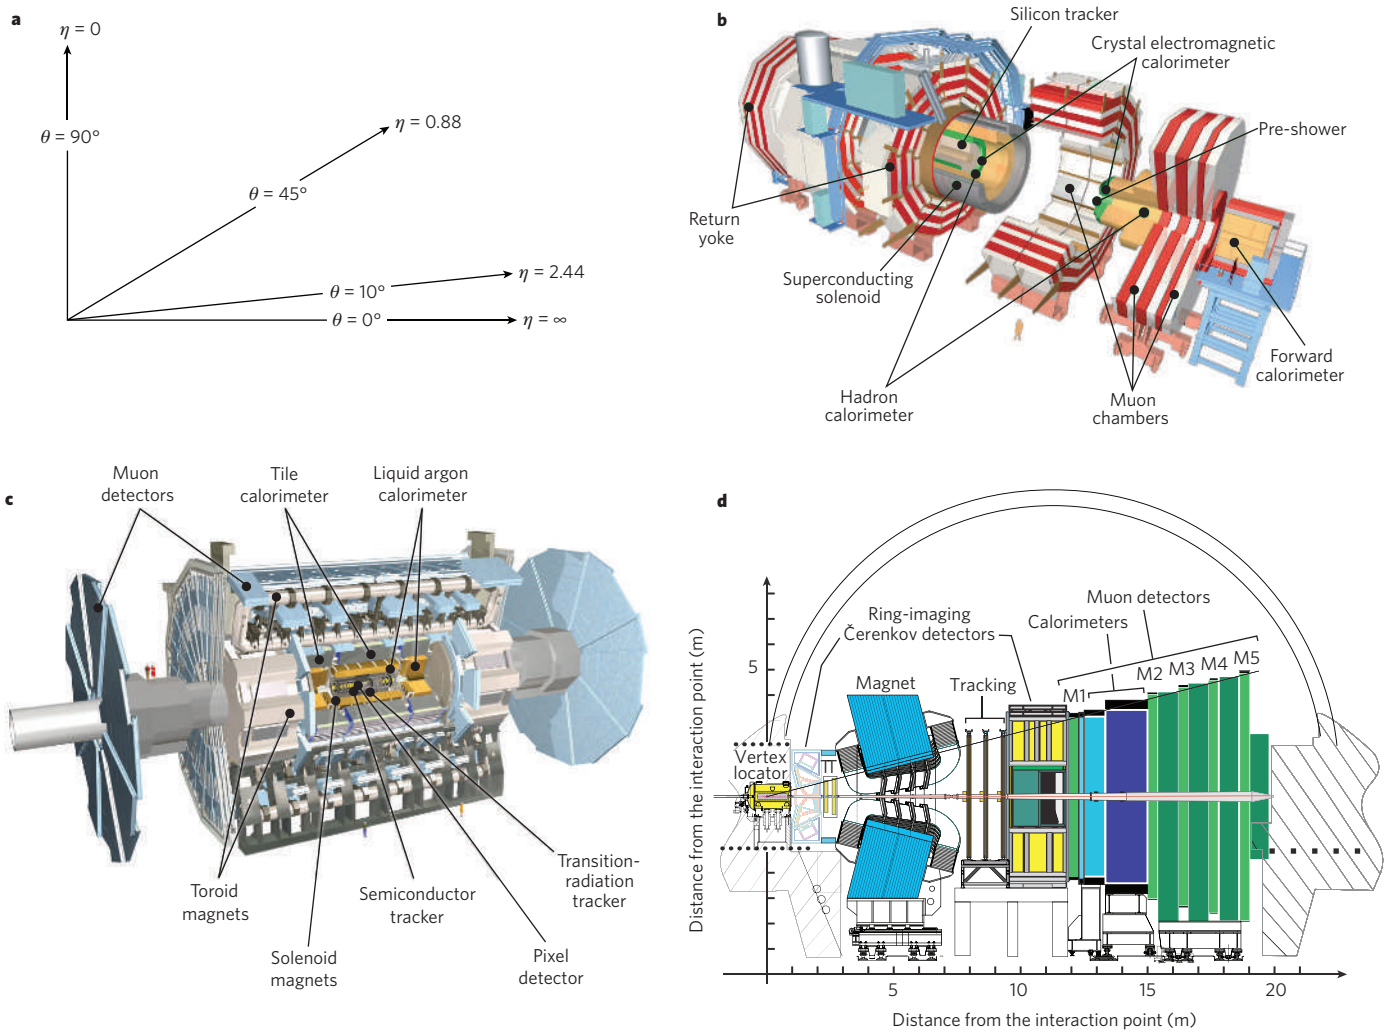

Figure 1 | Detector design. The complex experimental apparatus comprises different detector components, each optimized for a particular task. Regions of the detector volume are commonly described using the variable pseudorapidity, η, which is related to the polar angle, θ, as shown in **a**. The geometry and basic elements of the general-purpose LHC detectors, CMS (**b**) and ATLAS (**c**), are similar, but the layout of the more specialized detector LHCb (**d**) is optimized for detecting the production of *b* quarks in the forward direction. Images b–d reproduced with permission from CERN.

©2007 Nature Publishing Group

chambers. The CMS system is more compact than ATLAS, and has a radius of 7.5 m and length of 24 m, but weighs 12,000 tonnes.

LHCb[8] is designed to operate at a luminosity of 2×10^{32} cm^{-2}s^{-1}, and its instrumentation is concentrated in the forward direction, between 10 and 300 mrad (Fig. 1d), because this is the region in which the pairs of b and $\bar{b}$ quarks that it aims to study are predominantly produced. The LHCb detector has a silicon vertex detector around the interaction region; then a tracking system consisting of silicon microstrip detectors and a straw tracker, and it includes a dipole magnet. It also has two ring-imaging Čerenkov detectors, positioned in front of and after the tracking system, for charged-hadron identification; a calorimeter system and finally a muon system. LHCb, in particular, has to trigger efficiently on the secondary vertices that are the signature of b quarks, and so its vertex and tracking detectors are factored into an early stage of its trigger scheme.

To construct these large detectors requires substantial resources. The ATLAS and CMS communities each consist of more than 150 universities and institutions from about 35 countries with about 2,000 collaborators per experiment. (LHCb is a factor of three smaller.) Research and development for the LHC detectors began around 1990; the construction projects were approved in 1996 and started in earnest around 1998. The effort to build all of the detector components has involved physicists all over the world, with groups of geographically distributed institutes taking responsibility for the construction of various parts according to their specific expertise and capabilities, as well as involving a large network of industrial partners.

Inner detectors

The ATLAS and CMS inner detectors (Fig. 2) are contained in central solenoid fields of 2 T and 4 T, respectively. They provide efficient tracking of charged particles within the pseudorapidity range $|\eta| < 2.5$, allowing momentum measurement and the reconstruction of primary and secondary vertices. Both systems are largely based on silicon, with high-granularity pixel systems at the smallest radii, and silicon-strip detectors at larger ones. ATLAS has a 'straw' tracker at the largest radius.

Silicon detectors are p–n junction diodes that are operated at reverse bias[9]. This forms a sensitive region depleted of mobile charge and sets up an electric field that sweeps charge (electron–hole pairs) liberated by radiation towards the electrodes. Detectors typically use an asymmetric structure, for example, a highly doped p electrode and a lightly doped n region (p-i-n), so that the depletion region extends predominantly into the lightly doped volume. By adding highly doped n electrodes (n-i-n) at the back, the back side can also be read out. Integrated circuit technology allows the formation of high-density micrometre-scale electrodes on large (10–15 cm in diameter) wafers, providing excellent position resolution. Furthermore, the density of silicon and its small ionization energy (the energy needed to created an electron–hole pair) result in adequate signals with active layers only 200–300 μm thick, and the charge mobility is such that the signals are also fast (typically tens of nanoseconds).

The main challenges for the inner detector parts are the high particle rates, the radiation tolerance needed and the control of ageing effects. The ATLAS and CMS trackers had to be designed to withstand high radiation doses (500–1,000 kGy) for the innermost pixel layers, and up to 100 kGy for the systems farther away from the interaction point, after 10 years of operation). As a result, the development of the integrated front-end electronics for these systems has been a major problem to solve over several years and design iterations. These circuits must be fast, radiation tolerant and low power, and are integrated on low-mass modules where cooling and material limitations are severe. Several rounds of testbeam measurements and rigorous irradiation programmes have been necessary to prove that the circuits will function in their final assemblies, as well as after high irradiation.

A similarly stringent research and development programme was needed for the silicon sensors themselves[10,11], for which the major difficulty is bulk radiation damage. The relevant parameter is the accumulated dose in the volume of the inner detector, which varies between 10^{15} and 10^{14} cm^{-2} from the innermost layers to the outer ones (where n_{eq} means the number of equivalent particles, normalized using non-ionizing energy loss cross-sections to the damage expected to be caused by 1 MeV neutrons). These radiation doses have severe consequences for the silicon sensors (as they do for all other module components and for the thermal design of the system). They cause increased leakage current and changes in effective doping, and therefore changes in depletion and operation voltages, leading to type inversion for n-type sensors. After type inversion, the effective doping also shows an increase with time following irradiation (reverse annealing) that is temperature dependent. To maintain the operation voltage within reasonable limits, the sensors are therefore kept cold (–10 °C to 0 °C) throughout their lifetime, which has the added benefit of reducing the leakage current.

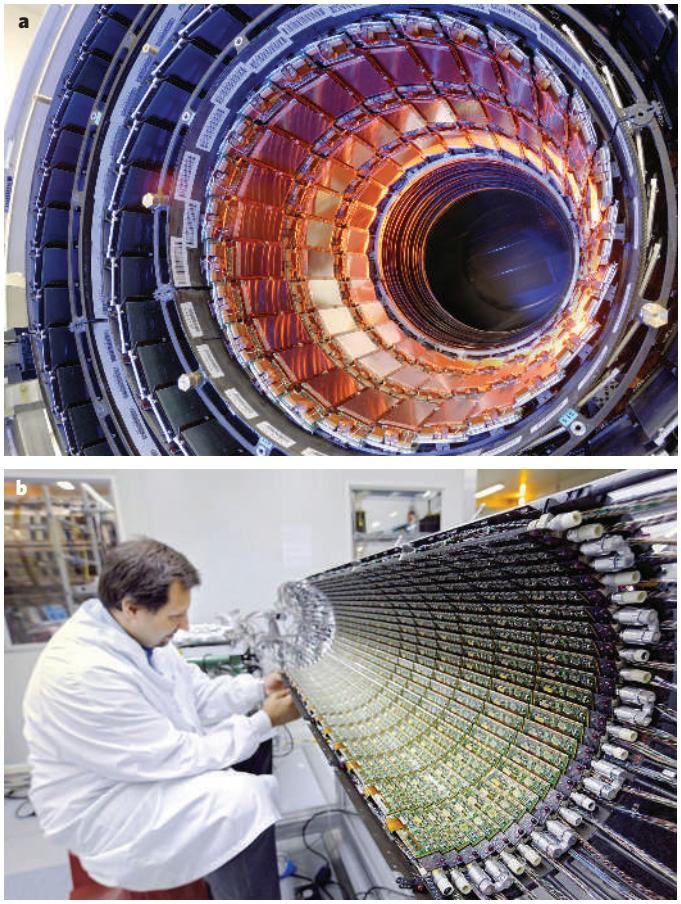

Figure 2 | Tracking systems. a, Silicon microstrip detectors in the CMS barrel region. **b**, The ATLAS pixel detector during the final assembly stage. Images reproduced with permission from CERN.

All of these effects have been carefully mapped out, and various design options have been evaluated in prototypes. Of particular interest are n-i-n silicon sensors, in which the charge-collection region grows with bias voltage from the n-implant side after type inversion following irradiation. Therefore, high efficiency can be obtained from an under-depleted detector. This allows a system to be specified with a lower maximum operating voltage. The ATLAS and CMS pixel systems and the LHCb vertex detector use such sensors. These require double-sided processing and are relatively complex and costly, so for the large-area silicon-strip systems, simpler, single-sided p-i-n designs have been adopted.

Considering the flux of charged particles at increasing radii around the LHC beams, three detector regions are defined in ATLAS and CMS. In the first of these, closest to the interaction point where the particle flux is highest, there are silicon pixel detectors (Fig. 2b), whose cell sizes of 50×400 μm^2 and 100×150 μm^2 in ATLAS and CMS, respectively, give an occupancy of about 10^{-4} per pixel per bunch crossing. To improve the measurement of secondary vertices (typically from b-quark decays) an innermost layer of pixels has been introduced as close to the beam as is practical, at a radius of about 4.5 cm. The lifetime of this layer will

©2007 Nature Publishing Group

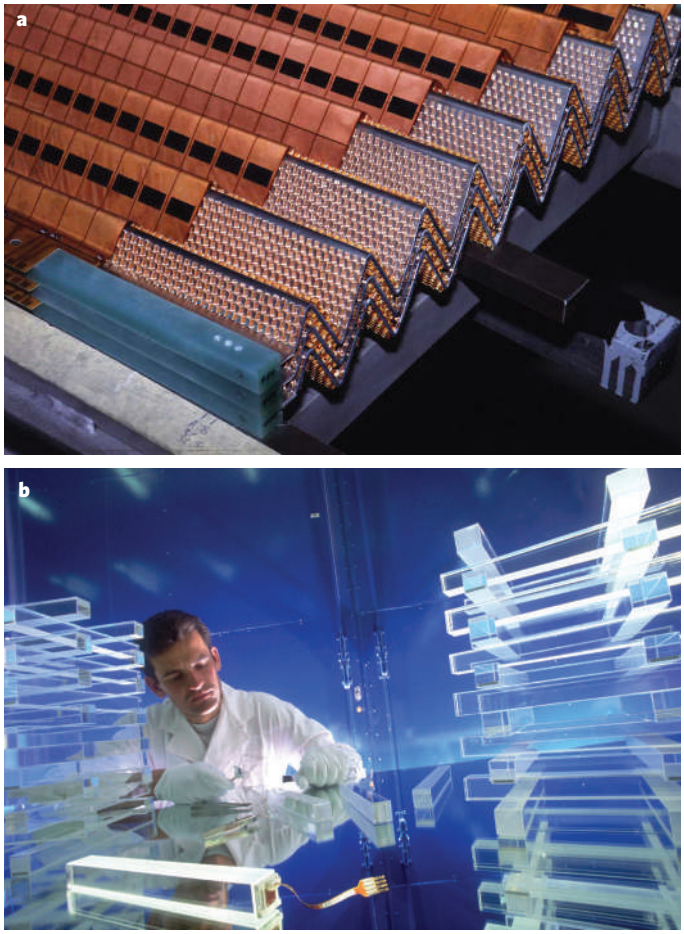

Figure 3 | Calorimetry: different approaches. a, The layers of the ATLAS electromagnetic calorimeter have an 'accordion' geometry. **b**, Tens of thousands of lead tungstate crystals have been prepared and tested, before being assembled into the electromagnetic calorimeter of the CMS detector. Images reproduced with permission from CERN.

be limited, owing to radiation damage, and it may need to be replaced after a few years. The pixel systems in ATLAS and CMS are very much larger than any comparable existing system. The ATLAS pixel system covers about 2 m^2 and has 80 million channels; the CMS pixel system is only slightly smaller.

Pixel detectors are expensive and have high power density, so at a certain radius and system size, silicon microstrip systems become the preferred technology. In the intermediate tracking region of ATLAS and CMS, at a radius of 20–55 cm, the particle flux becomes low enough to use silicon microstrip detectors. Barrel cylinders and endcap discs provide coverage out to about $|\eta| = 2.5$. Strip dimensions of 10–12 cm × 80–120 μm lead to an occupancy of 1–3% per bunch crossing. Both trackers use stereo angle in some of the strip layers (that is, strips placed at a small angle with respect to the z axis, 40 mrad and 100 mrad for ATLAS and CMS, respectively) to improve the resolution in z. In these microstrip systems, it has been essential to find a good balance between the pitch of the cells (determining resolution and occupancy), radiation effects, capacitive load (noise), material length and costs.

Finally, in the outermost region (beyond about 55 cm), the particle flux has dropped sufficiently to allow the use of larger-pitch silicon microstrips in the CMS tracker, with a maximum cell size of 25 cm × 180 μm while keeping the occupancy to about 1%. There are six layers of such microstrip modules in the barrel, accompanied by nine endcap discs providing coverage out to about $|\eta| = 2.5$, amounting to 15,400 modules and 9.6 million channels, and spanning a total detector area of more than 200 m^2.

For ATLAS, at radii greater than 56 cm, a large number of tracking points (typically 36 per track) is provided by the 'straw' tracker — 300,000 straw-tubes embedded in fibre or foil radiators and filled with a xenon-based gas mixture. This detector allows continuous track following with less material per point, and also has electron identification capabilities. X-ray photons are produced through transition radiation because highly relativistic particles such as electrons traverse the straw tracker's multiple interfaces.

Another silicon microstrip detector at LHC, the Vertex Locator[12] of LHCb, has some special features that present significant challenges. The 42 double-sided, half-moon-shaped detector modules are placed at a radial distance from the beam (8 mm) that is smaller than the aperture required by the LHC during injection and must therefore be retractable. For minimizing the material between the interaction region and the detectors, the silicon sensors are inside a thin aluminum box at a pressure of less than 10^{-4} mbar (10^{-2} Pa). The n-i-n sensors used have rφ geometry with pitch (40–140 μm) depending on the radius.

Calorimeters

The calorimeters absorb and measure the energies of electrons, photons and hadrons. In the design of the electromagnetic calorimeters for both ATLAS and CMS, the emphasis is on good resolution for photon and electron energy, position and direction measurements, and wide geometric coverage (up to $|\eta|$ close to 3.0). In the QCD-dominated environment of the LHC, the ability to reject neutral pions is crucial for photon and electron identification. It is also important to have efficient photon and lepton isolation measurements at high luminosities. For the hadronic calorimeters, the emphasis is on good jet-energy measurements, and full coverage (to $|\eta| = 5$) to be able to ascribe the observation of significant missing transverse energy to non-interacting particles (such as neutrinos, or light neutralinos from supersymmetric-particle cascade decays) rather than to losses in the forward regions. Last but not least, the quantities measured in the calorimeters play a crucial part in the trigger of the experiment as signatures of significant parts of the new physics sought at the LHC.

These considerations bring stringent requirements for high granularity and low noise in the calorimeters. The major technical difficulties for the calorimeters are related to the radiation doses (reaching 200 kGy for the electromagnetic part and 1,000 kGy for the hadronic part at the highest η), the sampling speed and the dynamic range needed to measure with low noise and good resolution over a wide energy range.

The ATLAS calorimetry consists of an electromagnetic calorimeter covering the pseudorapidity region $|\eta| < 3.2$, a hadronic barrel calorimeter covering $|\eta| < 1.7$, hadronic endcap calorimeters covering $1.5 < |\eta| < 3.2$, and forward calorimeters covering $3.1 < |\eta| < 4.9$, as shown in Fig. 1c. Over the pseudorapidity range $|\eta| < 1.8$, a presampler is installed in front of the electromagnetic calorimeter to correct for energy loss upstream.

The electromagnetic calorimeter system consists of layers of lead (creating an electromagnetic shower and absorbing particles' energy), interleaved with liquid argon (providing a sampling measurement of the energy-deposition) at a temperature of 89 K. The system's 'accordion' geometry[1,2] provides complete azimuthal symmetry, without cracks, and has been optimized for the high sampling rate environment of the LHC (Fig. 3a). The barrel section is sealed within a barrel cryostat, which also contains the central solenoid, surrounding the inner detector. The endcap modules are contained in two endcap cryostats that also contain the endcap hadronic and forward calorimeters.

The hadronic barrel calorimeter is a cylinder divided into three sections: the central barrel and two identical extended barrels. It is again based on a sampling technique, but uses plastic scintillator tiles embedded in an iron absorber. The vertical tile geometry makes it easier to transfer the light out of the scintillator to photomultipliers and achieves good longitudinal segmentation.

At larger pseudorapidities, closer to the beam pipe where higher radiation resistance is needed, liquid-argon technology is chosen for all calorimetry, for its intrinsic radiation tolerance. The hadronic endcap calorimeter is a copper/liquid-argon detector with parallel-plate

©2007 Nature Publishing Group

geometry, and the forward calorimeter is a dense liquid-argon calorimeter with rod-shaped electrodes in a tungsten matrix.

The approximately 200,000 signals from all of the liquid-argon calorimeters leave the cryostats through cold-to-warm feedthroughs located between the barrel and the extended barrel tile calorimeters, and at the back of each endcap. The barrel and extended barrel-tile calorimeters both support the liquid-argon cryostats and act as the flux return for the solenoid.

The CMS calorimeter system contrasts with that of ATLAS, because of its compactness[3,4]. In CMS, the solenoid is positioned outside the calorimeter, reducing the material in front of it, but also limiting the thickness of the calorimeter itself, and in particular the number of interaction lengths available to absorb hadronic showers.

The electromagnetic calorimeter, with coverage in pseudorapidity up to $|\eta| < 3.0$, comprises around 80,000 crystals of lead tungstate (Fig. 3b). These crystals have a high density and short radiation length, which make for a compact and high-resolution calorimeter. The main challenges are related to ageing/radiation effects and to temperature control (to the level of a tenth of a degree) to make full use the excellent intrinsic resolution of the system. The scintillation light is detected by silicon avalanche photodiodes in the barrel region and by vacuum phototriodes in the endcap region. A 'pre-shower' system, of silicon strip and lead layers, is installed in front of the endcaps to aid rejection of neutral pion signatures.

Surrounding the electromagnetic calorimeter is a brass/scintillator sampling hadron calorimeter, with coverage up to $|\eta| < 3.0$. Brass has a relatively short interaction length, is easy to machine and is non-magnetic. The scintillation light is converted by wavelength-shifting fibres embedded in the scintillator tiles and channelled to novel photodetectors known as hybrid photodiodes, which can operate in high axial magnetic fields.

Even with such compact electromagnetic and hadronic calorimeters in the barrel region, the total interaction length is limited to 7.2 λ (where λ is the interaction length or mean free path of a particle in the material) at $\eta = 0$ inside the solenoid coil. For this reason, a 'tail catcher' has been added around the coil to complement the hadronic calorimetry and to provide better protection against the escape (or 'punch-through') of hadronic energy into the muon system beyond.

The CMS forward calorimeter, constructed from steel and quartz fibres, is situated 11 m from the interaction point, thereby minimizing the amount of radiation and charge density in the detector during operation. The Čerenkov light emitted in the quartz fibres is detected by photomultipliers. The forward calorimeters ensure full geometric coverage up to $|\eta| = 5.0$, for the measurement of the transverse energy in the event and forward jet measurements.

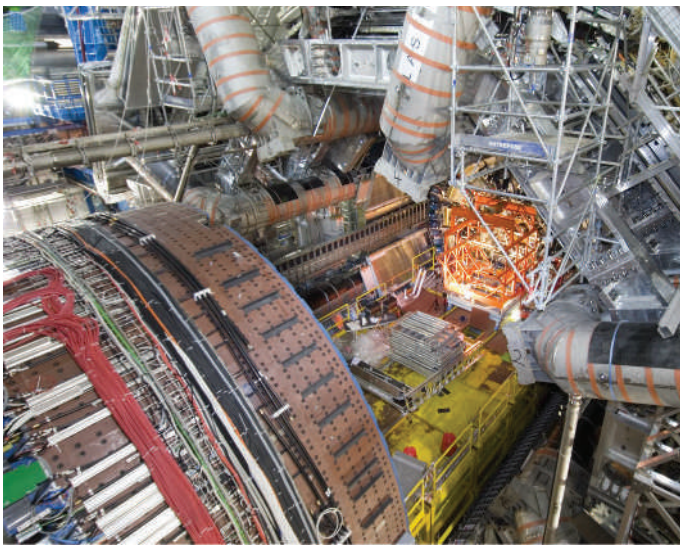

Figure 4 | Final integration. The components of the ATLAS detector are installed in the experiment's underground cavern. Here, part of the inner detector has just been moved inside the barrel calorimeter and toroid systems, while the endcap calorimeters (in the foreground) are kept in an open position to allow access. Image reproduced with permission from CERN.

Muon systems

The outermost detector layers of ATLAS and CMS, and the farthermost layers of LHCb, are dedicated to the measurement of the directions and momenta of high-energy muons, which escape from the calorimeters. Muons form a robust, clean and unambiguous signature of much of the physics of interest at the LHC.

Both ATLAS and CMS have had their overall detector designs optimized and adapted to trigger on and reconstruct muons at the highest luminosities of the LHC. (LHCb will operate, deliberately, at lower luminosity.) Muons must be measured with high efficiency and momentum resolution at low energies (such as in *B*-physics studies), at intermediate energies (for example, in the search for a Higgs decay into four muons), and at very high energies (to identify multi-TeV resonances such as a *Z*′). Wide pseudorapidity coverage is also important, and the ability to trigger on muons with energies of 5–10 GeV is crucial for several of the key physics goals. Finally, good timing resolution and the ability to identify in which proton-bunch crossing the muons were produced are absolute requirements, putting important constraints on the technological solutions chosen for the muon systems.

The muon systems in all three experiments are large-area gas-based detectors (several thousand square metres of multilayer chambers each in ATLAS and CMS). The chambers are divided into two sets, one intended for precise measurements of muon tracks and the other dedicated to triggering on muons. The sheer size of the systems means that there are significant technical challenges related to the stability and alignment of the chambers and to the careful mapping of the detectors' magnetic fields over large volumes. The radiation levels for the muon chambers are much less severe than for the inner detectors or calorimeters, but there are still concerns about ageing of the systems and also the neutron radiation environment of the experimental halls in which the detectors sit. The designs of the beam pipe and the shielding elements in the forward direction have been carefully optimized to reduce the neutron-induced background rates in the muon chambers.

Although the muon-chamber technologies chosen for ATLAS and CMS have many similarities, the magnet configuration in the two experiments is quite different. The ATLAS air-core toroid system, with a long barrel and two inserted endcap magnets, generates a large-volume magnetic field with strong bending power within a light and open structure. Multiple-scattering effects are thereby minimized, and excellent muon momentum resolution is achieved with three stations of high-precision muon-tracking chambers, covering up to $|\eta| = 2.7$. Over most of the range in pseudorapidity, the measurement of track coordinates (in the principal bending direction of the magnetic field) is performed by monitored drift tubes[1,2]. This technology provides robust and reliable operation, thanks to the mechanical isolation of each sense wire from its neighbours in the gas-filled drift volumes of the individual tubes. At large pseudorapidities and close to the interaction point, where the rate and background conditions are more difficult, cathode-strip chambers with higher granularity strip readout are used. The muon trigger system, with a fast time response and covering $|\eta| < 2.4$, comprises resistive-plate chambers in the barrel and thin gap chambers in the endcap regions. As well as triggering, these chambers provide a measurement of a second track coordinate orthogonal to the one measured by the high-precision chambers. In addition to the muon-chamber measurements, the inner detector measurements in the central solenoid of ATLAS contribute to the combined muon momentum resolution of the experiment.

The superconducting solenoid inside CMS is, at 13 m long with a 5.9-m inner diameter, the largest of its kind[3,4]. To achieve good momentum resolution without making overly stringent demands on muon-chamber resolution and alignment, the solenoid will operate at a high magnetic field of 4 T. In CMS, centrally produced muons are measured three times: in the inner tracker, after the coil, and in the return flux, into which four muon 'stations' are integrated, achieving robustness

©2007 Nature Publishing Group

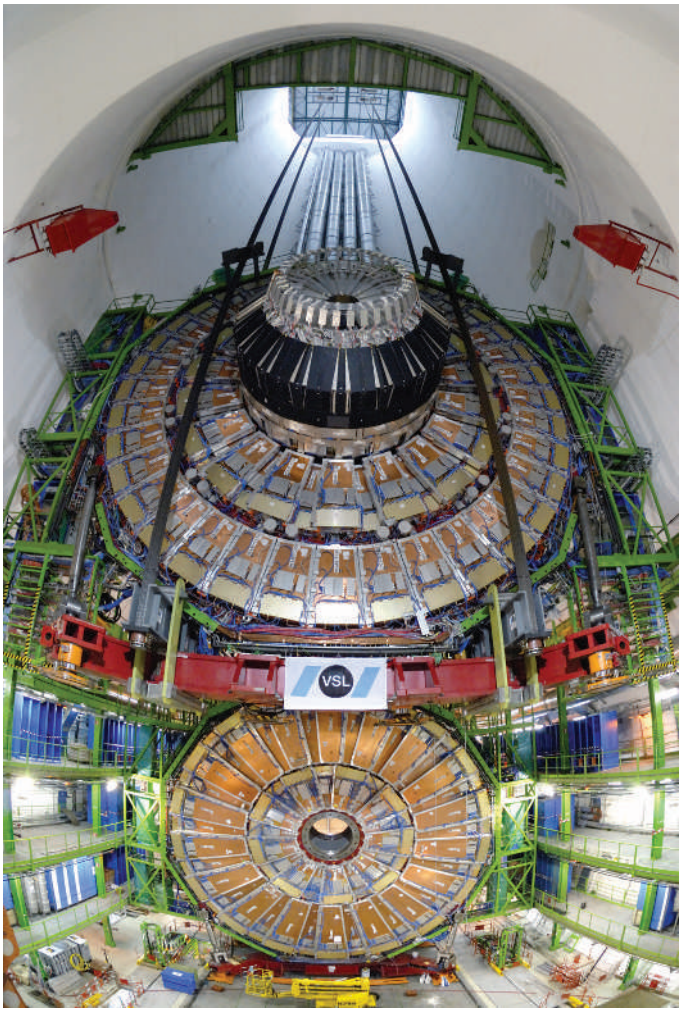

Figure 5 | Going underground. A large unit of the CMS detector, an endcap disc with muon chambers and part of the hadronic calorimeter, is lowered 100 m into its final position in the cavern. Image reproduced with permission from CERN.

and full geometric coverage. Three types of gaseous detector are used: drift tubes in the barrel region ($|\eta| < 1.2$), where the neutron-induced background is small, the muon rate is low and the residual magnetic field in the chambers is low; cathode-strip chambers in the two endcaps ($|\eta| < 2.4$), where the muon rate, the neutron-induced background and the magnetic field are all high; and resistive-plate chambers, in both the barrel and the endcap regions, to provide a fast response with good time resolution and to identify unambiguously the correct bunch crossing (albeit with coarser position resolution).

The intrinsic resolution of the high-precision chambers is in the range 60–150 μm (ref. 6), but the overall performance over the large areas involved (particularly at the highest momenta) depends on how well the muon chambers are aligned with respect to each other and with respect to the overall detector. The high accuracy of the ATLAS stand-alone muon measurement necessitates a precision of 30 μm on the alignment; in CMS, the different muon chambers need to be aligned with respect to each other and to the central tracking system to within 100–500 μm.

Both experiments have intricate hardware systems installed that are designed to measure the relative positions of chambers that contribute to the measurement of the same tracks, but also to monitor any displacements during the detector operation. For instance, in ATLAS, about 5,000 optical alignment sensors and 1,800 magnetic field sensors will track the movements of the chambers and will map and track the magnetic field to an accuracy of approximately 20 G (2 mT) throughout the detector volume. In the CMS, the solenoid magnetic field is more uniform but nevertheless the field is carefully monitored by around 80 sensors. Around 1,400 alignment sensors will provide independent monitoring of the tracking detector geometry with respect to an internal light-based reference system. The final alignment values will be obtained with the large statistics of muon tracks traversing the muon chambers.

The LHCb muon system consists of five stations. Multiwire proportional chambers are used throughout, except for the innermost region closest to the beamline of the first station. This station is placed in front of the calorimeters and represents a significant challenge in terms of material budget, space constraints, rate capability and radiation tolerance. The innermost region of this station, where the particle rates are highest, is equipped with triple-GEM (gas electron multiplier) detectors[13] with pad readout that are particularly suited for tracking in a high particle rate environment.

Triggering and readout

At design luminosity, the LHC will create 10^9 proton–proton events per second, but data storage and processing capabilities are such that data from only about 100–200 carefully selected events per second (each of these interesting events is accompanied by an average of 25 overlapping proton–proton events in the same bunch crossing) can be recorded offline for complete analysis. Hence, there is a need for a trigger system to select only the most important physics signatures and achieve a rejection factor of nearly 10^7.

The trigger systems for the LHC experiments have distinct levels. The first level, based on custom-built processors, uses a limited amount of the total detector information to make a decision in 2.5/3.2 μs (ATLAS/CMS) on whether to continue the processing of an event or not, reducing the data rate to around 100 kHz. Higher levels, using a network of several thousand commercial processors and fast switches and networks, access gradually more information and run algorithms that resemble offline data analysis to achieve the final reduction. The total amount of data recorded for each event will be roughly 1.5 megabytes, at a final rate of 150–200 Hz. This adds up to an annual data volume of the order of 10 petabytes for the LHC experiments.

The challenges to be faced in real-time data collection and reduction are many. The synchronization of the individual parts of the detector — and there are several thousand units to time in — must be accurate to better than a nanosecond, taking into account the flight times of particles to the individual sensor elements. At later stages, there is the second synchronization challenge of assembling all of the data for a particular bunch crossing from various parts of the detector into a complete event.

There is no chance of processing and selecting events within the 25 ns available between successive proton-bunch crossings. Furthermore, the sizes of the detectors and of the underground caverns in which they sit impose a minimum transit time between the detector electronics and trigger electronics. The first-level trigger calculations themselves need to be sufficiently sophisticated to identify clear physics signatures; the decision is based on the presence in the calorimeters or muon detectors of 'trigger primitive' objects, such as photons, electrons, muons and jets above pre-set transverse-energy or transverse-momentum thresholds. It also employs global sums of transverse energy and missing transverse energy. During the transit and processing time — less than 2.5/3.2 μs for ATLAS/CMS — the detector data must be time-stamped and held in buffers.

After an event is accepted by the first-level trigger system, the data from the pipelines are transferred from the detector electronics into readout buffers. The further processing involves signal processing, zero suppression and data compression while the events are examined by a farm of commodity processors consisting of several thousands of central processing units. The design and implementation of the processor farm, switching network, control software and trigger application software are major challenges. The event fragments must be directed from the readout buffers to a single processor and buffer node, using fast switches and networks, in order to perform more detailed calculations of the critical parameters of the event and to reduce the final rate further.

Even after such a large online reduction, huge amounts of data will be recorded. It was soon clear that the required level of computing resources

©2007 Nature Publishing Group

could be provided only by a significant number of computing centres working in unison with the CERN on-site computing facilities. Off-site facilities will be vital to the operation of the experiments to an unprecedented extent. Hence, over the past five years, GRID infrastructure for data processing and storage has been developed[14,15].

The GRID solution is geographically distributed and relies on three tiers[6]. The various tiers have clear responsibilities. The raw data output from the high-level trigger is processed and reconstructed in a Tier 0 computing facility at CERN, producing reconstructed data. Most detector and physics studies, with the exception of calibration and alignment procedures, will rely on this format. A copy of the raw data plus the reconstructed data is then sent to the Tier 1 centres around the world. These share the archiving of a second copy of the raw data, provide the reprocessing capacity and access to the various streams of reconstructed data (corresponding to the major trigger signatures) and allow data access and processing by the experiments' physics-analysis groups. The analysis data produced at Tier 0 and 1 are derived from the reconstructed data and are a reduced-event representation, intended to be sufficient for most physics analyses.

The Tier 2 facilities, each linked to a specific Tier 1 centre, are smaller but more numerous and are used for analysis, data-calibration activities and Monte Carlo simulations. Furthermore, the Tier 1 analysis data are copied to Tier 2 to improve access to them for physics analysis; by contrast, the Tier 1 centres provide safe storage of the large data sets produced at Tier 2 (for example, simulation data). A final level in the hierarchy is provided by individual group clusters and computers used for analysis.

This machinery for processing and analysis is being set up, and tests already indicate that the transfer speed between CERN and Tier 1 that is essential for initial running can be achieved, and that large-scale data production can be carried out in the GRID framework.

Ready to start

Installation of the LHC detectors in their underground caverns began in 2004. The components of ATLAS have been assembled and tested in the experiment's cavern, at 'Point 1' on the LHC ring (Fig. 4), and LHCb in the cavern at 'Point 8'. By contrast, the bulk of the CMS system was assembled and tested on the surface, before being lowered 100 metres into its cavern at 'Point 5' in 15 large lift operations (Fig. 5).

All of the LHC experiments have been operating their detector elements as much as possible on the surface and in 'test beams', and also, following installation, in the underground caverns. The flux of muons from cosmic rays provides a useful test of systems such as the calorimeters, inner detectors and muon systems to check alignment, calibration and the integration of data collection. In parallel, the GRID computing infrastructure and organisation are being planned, implemented and tested. For ATLAS, CMS and LHCb, major exercises of their data processing, software and computing infrastructure have been performed, and more are planned in the run-up to the introduction of beams into the LHC in 2008.

The first collisions at a centre-of-mass energy of 14 TeV are expected by mid-2008. As soon as a luminosity of 10^{32} $cm^{-2}s^{-1}$ is reached, within days the LHC can produce data sets — of W and Z bosons, t quarks, high-transverse-momentum jets, and even supersymmetric particles — that will surpass those of any previous or existing accelerator. At that point, the main physics goals of the LHC will be in full focus, and the aim will be to collect as much data in 2008 as possible.

The decade-long period of detector development and construction is coming to an end. Many of the fundamental challenges addressed by the experiment builders at LHC have been solved successfully, most notably in the development of fast, radiation-tolerant and sufficiently granular detector systems and electronics, integrated in turn in large systems that exceed any existing detector in specification (size, speed and channel count, in particular). Technology advances in computing, switches, networks and software have allowed the development of sophisticated trigger, data acquisition and GRID systems to handle the LHC data rates and volumes — it was far from obvious that this could be achieved when the building of the experiments was initially approved. The accelerator, detectors and off-line systems now need to be completed in their underground areas, commissioned fully and operated efficiently.

A further challenge overcome in the LHC project is the successful collaboration in each experimental team of as many as 2,000 scientists from all over the world, working together for a decade and using their resources and skills efficiently. Thanks to their efforts, and those of the teams building the LHC itself, a new era of research in experimental particle physics is finally within reach. The community can now look forward to the new challenges posed in interpreting the data from the LHC — challenges as great as those that have been faced in the building of the detectors. There is good reason to be optimistic, and the potential rewards, in terms of physics discoveries, make it well worth the effort. ■

1. ATLAS Collaboration. *ATLAS Detector and Physics Performance: Technical Design Report*. Report No. CERN/LHCC/99-15 (CERN, Geneva, 1999).
2. ATLAS Collaboration. *ATLAS: Technical Proposal for a General-Purpose p p Experiment at the Large Hadron Collider at CERN*. Report No. CERN/LHCC/94-43 (CERN, Geneva, 1994).
3. CMS Collaboration. *CMS Physics Technical Design Report* Vol. I. Report No. CERN/LHCC/2006-01 (CERN, Geneva, 2006).
4. CMS Collaboration. *CMS Technical Proposal*. Report No. CERN/LHCC/94-38 (CERN, Geneva, 1994).
5. Ellis, N. & Virdee, T. S. Experimental challenges in high luminosity collider physics. *Annu. Rev. Nucl. Part. Sci.* **44,** 609–653 (1994).
6. Froidevaux, D. & Sphicas, P. General purpose detectors for the Large Hadron Collider. *Annu. Rev. Nucl. Part. Sci.* **56,** 375–440 (2006).
7. Gianotti, F. *European School of High-Energy Physics, CERN Yellow Report*. Report No. CERN-2000-007 219–244 (CERN, Geneva, 1999).
8. LHCb Collaboration. *LHCb Technical Proposal*. Report No. CERN/LHCC/98-104 (CERN, Geneva, 1998).
9. Yao, W.-M. *et al*. Review of particle physics. *J. Phys. G* **33,** 1 (2006).
10. Doležal, Z. *et al*. The silicon microstrip sensors of the ATLAS semiconductor tracker. *Nucl. Instrum. Methods* **578,** 98–118 (2007).
11. CMS Collaboration. *Addendum to the CMS Tracker TDR*. Report No. CERN/LHCC 2000-016 (CERN, Geneva, 2000).
12. LHCb Collaboration. *LHCb TDR 5* Report No. CERN/LHCC 2001-011 (CERN, Geneva, 2001).
13. LHCb Collaboration. *Second Addendum to the Muon System Technical Design Report*. Report No. CERN/LHCC/2005-0012 (CERN, Geneva, 2005).
14. Foster, I. & Kesselman, C. *The Grid: Blueprint for a New Computing Infrastructure* (Morgan & Kaufmann, San Francisco, 1998).
15. LCG project. *LHC Computing Grid Technical Design Report*. Report No. CERN/LHCC/2005-024 (CERN, Geneva, 2005).

Acknowledgements This review article is largely based on the very complete documentation already existing for the ATLAS, CMS and LHCb experiments, and on refs 1–8 in particular. Furthermore, with around 5,000 scientists involved in the construction of the experiments described, I can only cover a small part of the challenges, excitement and difficulties involved in making them a reality. The best I can do is therefore to acknowledge all of the members of these collaborations, and in particular those who have helped me with corrections and comments, and apologize for all the dedicated sub-projects and work I had to leave out of this review.

Author information Reprints and permissions information is available at npg.nature.com/reprintsandpermissions. The author declares no competing financial interests. Correspondence should be addressed to the author (steinar.stapnes@cern.ch).

©2007 Nature Publishing Group