Developing LHC Detectors

H.J. Hilke and P.G. Innocenti

CERN, Geneva, Switzerland

Detectors for the LHC call for novel approaches and new technologies to handle the unprecedented high luminosities.
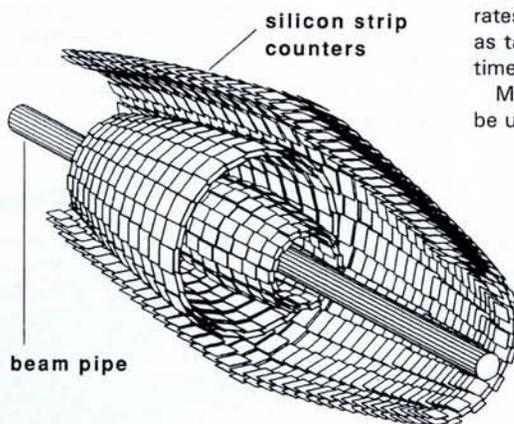
LHC places stringent demands on detector performance which present designs cannot fulfill. The principal reason for this is the very low cross-section for events of interest. While the background events have a large cross-section of 0.1 barn, the most interesting reaction channel pp → H + χ; H → ZZ + four muons has a cross-section on the order of 10^{-15} barns, depending on the mass of the Higgs particle H, *i.e.*, one such reaction will occur in only about 10^{14} events. This leads to the requirement of unprecedented interaction rates: to register just 100 such events in 10^7 s (about one year with 100% operating efficiency) a luminosity of $10^{34}/cm^2s$ is required, equivalent to more then 10^9 interactions per second or 10^3-10^4 times the rates for previous colliders. For the LHC scheme presently proposed, a maximum luminosity of $1.7 \times 10^{34}/cm^2s$ with beam bunch crossings every 15 ns, the H event will on average occur simultaneously with 19 background events. Experimentalists will have to disentangle these 20 events by determining their respective production vertices scattered over a distance of about 15 cm in the direction of the beam. Moreover, the high beam energies lead to high beam multiplicities. Finally, precise determination of energy and momenta of energetic particles up to the TeV range requires large detectors.

Today's technology is in many areas inadequate for realising systems able to meet future detector requirements for high energy physics. The preparation of new experiments has traditionally concentrated on intense development activities for detector techniques. Thus, a R. & D. effort was started at CERN as long ago as 1987 within the specially funded LAA Project to tackle the requirements for large precision detectors with high granularity and capable of handling many particles at very high rates. A major programme involving the entire CERN community was launched in 1990 under the framework of CERN's Detector R. & D. Committee (DRDC). Some 26 of more than 40 proposals for the development of detectors, electronics and data acquisition have so far been approved. They receive financial support from participating institutes and from CERN in the ratio of roughly 2:1.

We shall summarise the status of these developments by illustrating how they meet the physics and technical requirements for

Hans Hilke is the Deputy Division Leader of the Particle Physics Experiments Division at CERN.

Pier Giorg Innocenti is the Division Leader of the Electronics and Computing for Physics Division at CERN.

Fig. 1 — Conceptional layout of a silicon vertex detector. The detector contains some 5×10^6 read-out elements at a pitch of 50–100 μm providing position resolution of < 20 μm in two coordinates. The three layers around the beam pipe reach from a radius of 9 cm to 35 cm and cover a length of 1.7 m, thus requiring some 13 m² of silicon detectors.

the main features (tracking, calorimetry, data acquisition, triggering, and magnets) of detector systems for the LHC.

Particle Tracking

Outer tracking

The tracking of particles comprises two distinct sections, namely the outer section beyond the calorimeters, and the inner section from the beam vertex up to the calorimeters. Outer tracking identifies muons and measures their momenta, with triggering on transverse momentum. Thousands of m² of detector are required, usually arranged in at least three radial layers, so low-cost construction is essential. On the other hand, a spatial resolution of 0.1–0.2 mm in one coordinate is also necessary to achieve the high momentum resolution of $dp/p < 0.1\,p$ (where p is in TeV) — five to ten times better than for LEP detectors — demanded in the LHC's TeV range. Good segmentation in the second coordinate and a very fast response are required to avoid ambiguities and pile-up.

Besides honeycomb strip chambers, most devices under investigation are variants of proven designs, namely drift chambers with small drift cells. Resistive plate counters, parallel plate chambers, limited streamer tubes, and scintillators are being studied for fast triggering and local timing of the drift cells. However, considerable effort is required to arrive at viable solutions for huge systems having high precision and acceptable costs. The **honeycomb strip chamber** is a novel device offering economic, large-scale construction: honeycomb structures are formed from corrugated Mylar or Kapton foils with 5 mm wide metallised strips running perpendicular to the anode wires inside the hexagons of the honeycomb. A resolution of 80 μm has been obtained from the strips.

Inner tracking

Inner tracking between the beam pipe and the calorimeters has several tasks: to determine the sign of the e± up to energies of about 1 TeV; to improve electron and photon identification (by isolation cuts in conjunction with the calorimeters); to provide a check on the calorimeter calibration; to improve the muon momentum measurement, especially for lower momenta; and, when LHC is operated at lower interaction rates, to identify decays from particles such as taus and B mesons with very short life times of the order of a picosecond.

Most of conventional techniques cannot be used in the inner region because of the extremely high particle rates and densities. At approximately 10 cm from the beam, the intensity is about 2×10^7 charged particles/cm²s, and is falling off as the square of the distance.

Silicon strip detectors (or possibly GaAs equivalents) are the only devices that can be considered for the innermost "vertex detector" layers. Large systems at LEP with 10^5 readout elements ($\approx$ 6 cm long strips at a pitch of 50 μm) have given excellent resolutions (down to 7 μm). For LHC, however, systems with about 100 times more elements are required (Fig. 1). These demand ingenious solutions for precise but lightweight mounting structures, efficient cooling and rapid but cost-effective read out. In addition, both the detectors and the electronics must withstand unprecedented radiation levels (up to 10 Mrad from charged particles and 10^{14} neutrons/cm² for one year of operation). This is two to three orders of magnitude more than what conventional devices tolerate. Many groups are studying the causes of radiation damage and possible solutions (including silicon-on-insulator processes and GaAs) in close collaboration with industry. Substantial progress has been achieved and it seems likely that the required radiation hardness can be reached within a few years.

Other exciting developments in **silicon detectors** are appearing, such as silicon drift chambers and pixel detectors for two-dimensional read-out with extreme granularity. In a silicon drift chamber, the information from the ionization caused by the passage of a charged particle is drifted over centimetres inside the silicon, thus providing excellent precision with a strongly reduced number of read-out channels. The read-out elements in pixel detectors can be as small as 75×500 μm² with front-end electronics placed on top of each channel.

Chips with 1000 elements have recently been operated very successfully. Owing to the enormous number of elements in larger devices, "intelligent" read out is mandatory: interesting information announces where it sits and empty cells are skipped. This development is still underway, but existing devices have already found applications in the less-demanding environments found in X-ray spectroscopy and in the observation of faint stars (where an extraordinarily low noise of 5 rms electrons has been achieved).

Scintillating fibres and microstrip gas counters are two other novel detectors which are most promising for the next layers (radii $\approx$ 0.5–1.5 m). The development of fibres with new wavelength shifting dopants giving a large Stokes shift combined with high quality extrusion and coating mean that light can be caught inside 30 μm diameter fibres and transmitted for more than a metre with good efficiency. Track localization to 35 mm has been achieved with 20 cm long targets. For LHC, a system is proposed with some 160 million fibres, 60 μm in diameter and up to 2.3 m long, read *via* some 400 novel optical imaging delay tubes gated by the trigger. These delay tubes with their silicon pixel readout remain at present the element requiring the most development effort. The fibre tracker promises excellent momentum resolution, good vertex reconstruction both normal and parallel to the beam axis, and low occupancy (less than 1%) even at the highest luminosities.

Microstrip gas counters have already shown good results: a resolution in space of 30 μm and the capability to handle particle intensities above 10^6 Hz/mm^2. However, many details need to be studied before large systems with more than a million channels may be contemplated. The basic design principle is simple: it is copied from drift chambers, but without wires thus permitting a reduced spacing of read-out anodes (Fig. 2). Alternating narrow (10 μm) anode and wider (50 μm) cathode strips, with a pitch of 200 μm between the anodes, are deposited on a high-resistivity (10^{11} Ωcm) substrate, *e.g.*, special glasses or plastics such as Tedlar. The approximately 3 mm thick gaseous layer is closed by a cathode plate. Careful tuning of many parameters is necessary to obtain a stable gas amplification (up to 10^4) and good efficiency. The design allows many variants: even curved constructions are relatively simple with plastic substrates and the absence of wires. Microstrip gas counters will undoubtedly find many applications well before the operation of the LHC, especially when high precision is required at high particle rates.

Time projection chambers (TPC) are proven devices but an exciting new application is in sight for heavy ion experiments. A TPC is the only purely electronic detector delivering three-dimensional tracking information, an absolutely essential property if one wishes to identify and measure separately the very large number of particles developed in heavy ion collisions at high energies.
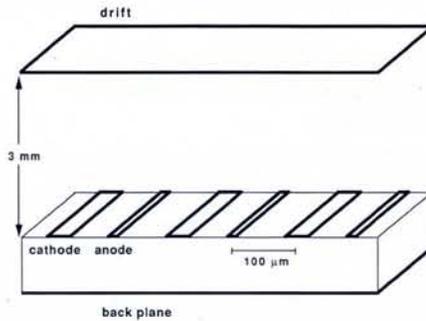

Fig. 2 — Schematic view of a microstrip gas counter. Thin metallic anodes and wider cathodes are deposited on a support with high resistivity, e.g. glass or plastic. The narrow gas gap is closed by the drift electrode.

For Au on Au central collisions at LHC, some 10 000 charged particles traverse a 5 m long detector with inner radius of 1 m, so statistical analysis should be possible on single events. Roughly 100 such events would be produced per second at a luminosity of 2×10^{26}/cm^{2}s. The principle modification with respect to the standard TPC design consists in completely subdividing the read-out planes into small read-out pixels of < 1 cm^2. This leads to > 3×10^5 electronic channels for a TPC with an outer radius of 2.5 m, so economic designs with a low power consumption are needed. The aim is to achieve a two-track separation of less than 1 cm in space and the identification of most particles.

Calorimeters

Calorimeters measure the energy of particles by total absorption. They usually consist of two sections: an electromagnetic part to identify and measure electrons and photons followed by a hadronic section to measure protons, neutrons, pions, and kaons. Rapid response, high granularity, hermeticity, and radiation hardness are mandatory in the LHC environment. A dose of up to 1 Mrad and a few 10^{13} neutrons/cm^2 will be deposited in the calorimeters during one year of operation.

Hadron calorimeter

The requirements for the energy resolution of the hadron calorimeter are not particularly stringent. Besides liquid argon and fibre calorimeters of similar construction as those described below, development groups are also investigating read-out techniques based on flat, thin liquid argon gaps, radially arranged scintillator tiles or silicon pads.

Electromagnetic calorimeter

For the electromagnetic calorimeter, an energy resolution of about 10%/$\sqrt{E}$ with E in GeV and an energy independent error of 0.5–1% are considered the minimum requirements: a much higher resolution is essential for some specific reaction channels, while the need for a constant term of only about 0.5% places severe demands on homogeneity and calibration.

A novel design of the classic **liquid argon** calorimeter and a novel technique, the fibre calorimeter, have been investigated for the "standard" resolution. The former is inherently radiation hard and easy to calibrate precisely. An accordion structure has been developed for the lead and read-out plates to improve the response time, homogeneity and hermeticity (Fig. 3). Radial strips are read out in two sections towards the inner and outer radii. Excellent homogeneity and space resolution and the expected energy resolution have been obtained for a prototype device.

In the **fibre** calorimeter, scintillating fibres pointing with a small angle towards the vertex are embedded in lead. This calorimeter has a rapid response and is relatively simple to construct, even as a combined calorimeter with a hadronic section for excellent energy resolution and linearity up to the highest energies. However, precise calibration and the long-term monitoring of stability (especially under high irradiation) to the 0.5% level have to be demonstrated.

Exciting new avenues are in sight for electromagnetic calorimeters having the highest resolutions. An intensive search for new materials followed the LEP L3 detector collaboration's pioneering development of bismuth germanate (so-called BGO) crystals leading to 11 t of BGO being used to give resolutions below 1% for particle energies above 2 GeV. A promising candidate is **CeF$_3$**: it responds much more rapidly than BGO (decay times of 5-20 ns instead of 300 ns), shows no temperature dependence and withstands more than 1 Mrad of radiation in small samples. Collaboration with industry has started to develop cheap but reliable mass-production techniques. If successful, there is no doubt that CeF$_3$ or
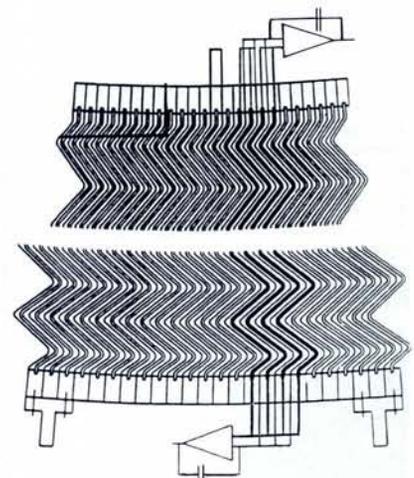

Fig. 3 — Schematic view of a liquid argon calorimeter of the accordion type. Alternating lead converter plates and thin read-out foils are folded with high precision to form liquid argon gaps 2 μm in width. Strips on double-sided copper-clad Kapton foils are read out to amplifiers at the inner and outer radii. The construction offers excellent homogeneity and hermeticity as well as a fast response.

similar crystals will find multiple applications in different fields, *e.g.*, in high precision tomography.

Data Acquisition

In order to cope with the high particle rate and maintain the required identification power, LHC detectors will have high granularities. A general-purpose detector with central tracking could have as many as $10^7 - 10^8$ electronic channels. Present estimates indicate initial trigger rates of the order of $10^4 - 10^5$ Hz, implying data acquisition bandwidths of the order of Gbytes/s and distributed CPU (central processing unit) power amounting to several hundreds of thousands of machine instructions per second.

A **multi-level** trigger and data acquisition system is necessary to accommodate the unprecedented event rates, trigger rejections and data volumes. A detailed study of the characteristics of this architecture will only be possible once the full design of the detectors is known. Nevertheless, by using a model of a general-purpose detector combined with the accelerator parameters and predicted particle cross-sections, evaluations of the system requirements have been made to show that the data acquisition system must be designed for the highest possible performance. Two features are essential, namely, pipelined buffering and hierarchical data collection. The system must also have by a high degree of parallelism for easy scaling and adaptation to the evolution of the required performance.

The electronics must resist radiation damage with little performance degradation over an expected lifetime of around 10 years. The expected doses over this period (10 Mrad and 10^{14} neutrons/cm²s) are well above what is tolerated by industrial components, particularly those used for analog circuits. Understanding radiation damage of electronic components at these high rates and finding manufacturing processes for custom VLSI circuits that show tolerable degradation challenges designers of LHC detector circuits.

Software aspects cannot be underestimated. The organization of the multiplicity of tasks in the several-thousand processor environment of the trigger system, not to speak of the control and on-line monitoring of the experiment and the management of thousands of CPU's distributed at the various levels of the data acquisition process, will require complex software. Modern software engineering methods and tools must be applied for its development and maintenance.

Multi-Level Trigger Systems

Much of the electronics for LHC detectors must be installed on or near the detector since moving the huge amount of data generated by the detector would require unmanageably large signal bandwidths and numbers of cables. Sophisticated trigger systems based on several levels are then needed to control the data flow. Relatively crude decisions made quickly at the first level can be refined at the second and third
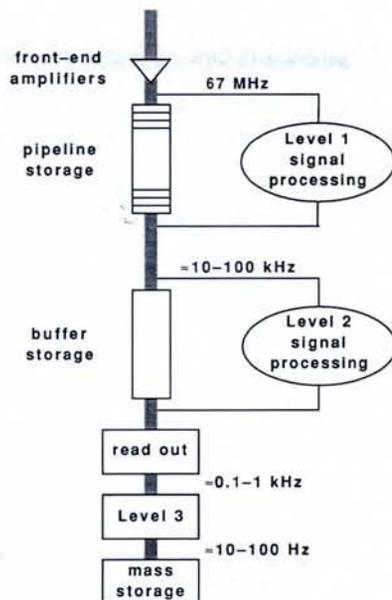
Fig. 4 — Conceptual model of a multi-level trigger system.

levels using more detailed information from the detectors and more complicated algorithms. The function of each trigger level is to reduce the trigger rate to the point where it can be accepted by the next level.

First-level processing should be very rapid because information from all the detector's channels ($10^7 - 10^8$ for a detector with central tracking) must be stored until the decision is available. It is expected that the first-level storage will be carried out by electronics placed at, or near, the detector. The size of a typical LHC detector is 30 m by 20 m so only allowing for cable delays, a first-level trigger decision needs at least 400 ns. Given that this time will be much longer than the 15 ns intervals between bunch crossings and that each event may contain several interactions, information for many events must be stored for each channel during first-level processing. Storage devices known as **pipelines** are expected to fulfill this rôle. While the trigger signal is being processed, pipelines reproduce the function of a delay cable as follows: pipeline memories, either analogue or digital, store new data every 15 ns until the first-level decision is available, after which the data are either discarded or transferred to the second-level.

The first-level trigger may take up to 2 μs to deliver a verdict to the detector electronics, this latency being made up of a combination of the response times of detectors and their associated electronics, cable delays and propagation delays through trigger electronics. However, the trigger must be able to start analyzing a new event every 15 ns. We therefore imagine a trigger system will concurrently processing many events, each separated in time by 15 ns. Pipelined processing where synchronised programs are allocated to each processor, a technique commonly employed in high-performance parallel computers, can be adopted for this.

It is important to combine data from different detectors in a single first-level central trigger for, say, electron-muon coincidences. With thresholds for electron and muon triggers chosen to have good efficiency for interesting physics reactions, the rate after the first-level trigger will thus be $10^4 - 10^5$ Hz. A reduction by a factor of at least 10^3 relative to the beam crossing rate, and by more than 10^4 compared to the interaction rate for luminosities in excess of 10^{34}/cm²s, is required.

When an event is accepted by the first-level trigger, data from the whole detector are stored in a second-level **buffer memory** during second-level trigger processing. The second-level system must also access detailed information from the detector so there is an intimate relationship between trigger processing and data storage at this level. One can buffer the data on the detector and only move those data which are needed by the second-level trigger processors. Alternatively, one can move all the data from the detector into the trigger processor system. Since trigger rates reach 10^5 Hz, moving, after each first-level trigger, the very large volumes of data ($> 10^5$ bytes per event) generated by some detectors will be difficult. Buffering the data locally until after the second-level trigger has accepted an event may therefore be preferred.

The second-level trigger should be able to increase rejection with respect to the first-level trigger decision by a factor of 10^2 so the rate after the second-level trigger will be $10^2 - 10^3$ Hz. In the model shown in the Fig. 4, events selected by the second-level trigger are fully read out before third-level processing. This trigger will already have made sophisticated decisions based on information from several (maybe all) detectors. So third-level processing may then have to perform full-event reconstruction and make physics analysis cuts at rates of some hundreds per second in the case of W → eν and W → μν decays. This could be carried using **clusters** of commercial processor systems: each incoming event is allocated to a processor which performs the third-level processing for that event. The limiting factor is then not so much computational power but the need to move data from read-out electronics to the processors.

Detector Magnets

To permit a fair comparison of choices for the LHC detector magnets from the technical point of view, a small study group has established common rules for comparing technical risks, construction problems and cost estimates, including scaling laws for cost variation with size and magnetic field strength. This led to a reduction in the number of options and in the size of some of them.

Nevertheless, the size, the mechanical stresses and the stored energy (up to 3 GJ) of some of the options for superconducting magnets are an order of magnitude larger than those of the last generation built for the LEP and HERA experiments. It became clear that intensive R. & D. work was neces-

sary before construction could be envisaged. The study group was charged with the task of defining the basic problems and establishing the necessary R. & D. programme. Three main areas of concern have been identified: stability of the magnet with respect to perturbations, fabrication of the conductor and impregnation of the coils.

A detailed plan is being worked out for comparing proven techniques and possible extensions. The conductor should carry 20 kA in fields up to 4 T. With up to 70 km of conductor required for one magnet, its cost represents a large fraction of the total cost. It is anticipated that some two years of intensive R. & D. work in close collaboration with industry will be required before the final specifications can be established.
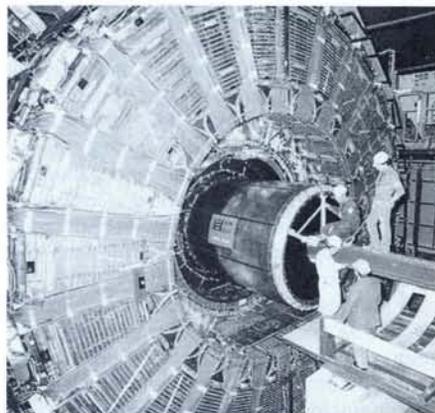
Particle Detectors and Industry

The experimental challenge at LHC calling for a close collaboration between the particle physics community and industry promises important spin-offs.

New frontiers in particle physics could historically only be reached by combining progress in accelerator technology with that of experimental techniques, including detectors, electronics, data acquisition, and analysis. Experiments had to become not only larger and more rapid but also much more complex and sophisticated: they could only be constructed and operated by large international collaborations.

Given the need to develop international cooperation and to handle enormous amounts of data, it is no surprise that the particle physics community vigorously promotes broadband data communication systems as well as open (interconnectable) multi-vendor computer networks. The same philosophy of open systems and networking applies to distributed computer-aided design and engineering whereby design teams in several locations can contribute simultanously to projects. The community also remains a test-bench for frontier developments in computing as it requires groups of computers operating in parallel ("farms") and machines based on massively parallel processing for large simulations.

A LHC experiment will contain many millions of read-out channels, each corresponding to the requirements of a high-definition television receiver. Both will require similar high-quality components so the incentive for a close interaction with industry leading to low-cost production is clear. High energy physics will again offer a thorough evaluation of performance limitations.

Owing to the need for rapid data acquisition, particle physics initiated new international communication standards such as NIM, Camac and FASTBUS, where the first two subsequently became widely used in industry. The community is presently participating in the development of a new generation of even faster bus systems for data. Trigger formation will call for a very flexible cross-linking of large data volumes at high rates based on integrating several industrial techniques only now entering development.

The DELPHI detector at LEP.

For the detector magnets, after initial problems, European industry managed to meet the tight tolerances and to be competitive in producing the superconducting cables required for two LEP experiments. The higher performance requirements for the LHC experiments magnets clearly represent a new challenge.

Regarding detector components, many new techniques have been invented by particle physics groups during the past decades, often relying on special industrial products (composites, high purity quartz, conductive plastics) and some finding various other applications (wire chambers of all kinds, new scintillators, *etc.*). This development activity continues and has indeed been stimulated immensely by LHC's stringent demands. Take the case of scintillating crystals: emerging requirements have brought together a large group of research institutes and companies involved in crystal growth — the former mainly evaluate performance, while the latter contribute their experience in mass production. These crystals find important applications outside particle physics owing to their compactness, fast response and high resolution.

So too for silicon detectors. The particle physics community initiated a broad line of developments in industry, both for the detectors and the associated electronics which requires an ever-increasing degree of large-scale integration. Many applications are on the horizon, including simple versions of recent pixel devices that have already entered X-ray research and astronomy. The next generation will permit readouts several orders of magnitude faster than today's CCD's used in video cameras.

These few examples illustrate that the interaction with industry proceeds on three levels. Particle physics often acts as a normal consumer of more-or-less standard products. In many cases, however, it participates actively in the developments of new products with exceptional performance. Finally, it sometimes invents techniques which find extensive use in other fields.