Researchers at CERN predicted that commercially produced submicron CMOS would provide a suitable technology base for detector electronics used in particle physics accelerators. Their collaboration with IBM Microelectronics has resulted in the demonstration that CMOS 6SF provides such a base for mixed-signal chips in a harsh environment. The first paper, by Faccio et al., illustrates some of the novel design techniques employed to achieve a measure of radiation tolerance, while the second paper, by French et al., highlights the impressive small-signal analog performance of one of the chips planned for use in the CERN Large Hadron Collider in Geneva. *-Dr. Guy Anastaze, IBM Microelectronics EMEA Marketing, Geneva, Switzerland.*

IBM CMOS 6SF for the Challenging Requirements of the High Energy Physics Experiments at the Large Hadron Collider

Federico Faccio, Giovanni Anelli, Michael Campbell, Giovanni Cervelli, Pierre Jarron, Kostas Kloukinas, Alessandro Marchioro, Gianni Mazza of INFN (Torino, Italy), Paulo Moreira, Pisana Placidi of the U. of Perugia, Angelo Rivetti of the U. of Torino, Walter Snoeys, Thomas Toifl, and Kenneth Wyllie. Unless otherwise indicated, all are with CERN (Geneva, Switzerland).

The readout and control of particle detectors at the CERN Large Hadron Collider (LHC) involves in excess of 10-million channels of electronics, with interactions at 40 MHz producing up to 10^5 interesting events per second. This requires a wide number of ASICs whose specifications are challenging in terms of low power, integration density, and low noise. Additionally, they have to survive the radiation environment generated by the particle beam collision. All these requirements are met with the use of IBM's CMOS 6SF technology, accompanied by a dedicated layout methodology to achieve radiation tolerance; this paper describes several functions successfully implemented into full-custom ASICs using this approach.

Introduction

The Large Hadron Collider (LHC) presently under construction at CERN, the European Laboratory for Particle Physics (Geneva, Switzerland), is scheduled to be operational in 2005. This new particle accelerator, placed in an underground 27-km-long circular tunnel (Figure 1), will generate collisions between two 7-TeV proton beams in the heart of high-energy physics detectors. The beams will cross at a frequency of 40 MHz, with roughly 25 particle pairs colliding at every bunch crossing, each producing a spray of up to 100 particles. The four LHC experiments, ATLAS, CMS, ALICE, and LHC-B, are the *electronic eyes* looking at the collision products, and each consists of tens of millions of sensor channels.

Figure 1. Aerial view of the CERN underground accelerator complex. The larger ring, 27 km long, will house the Large Hadron Collider. On the right, the runaway system of the Geneva International Airport and a portion of Lake Geneva give an idea of the scale.

The integration of all the complex data processing and control functions involves a large number of ASICs. The inner detector layers of each experiment will be placed close to the collision point and aimed at reconstructing the particle tracks. Because the detectors need to operate close to 0 °C, the thermal budget is very limited. To control the heat generated by these high-density ASICs, they must run at the lowest possible power. Low noise is also required for all circuits that read out signals from particle detectors. In some cases, as for the transmission of digital optical signals outside the detectors at 1.6 Gbps, the use of a fast technology is also necessary.

All the above requirements — low power, high density, low noise, and high speed — are satisfied using CMOS 6SF, IBM's 0.25-µm CMOS technology. There is, though, another requirement peculiar to the LHC experiments: radiation tolerance.

The Enclosed Layout Transistors (ELT)

Particles produced by the beam collision generate a radiation environment in which the detectors and associated electronics have to operate. This means that each ASIC must be able to tolerate the radiation level at its location in the experiment; the closer to the collision point, the higher the radiation. In terms of an integrated ionizing dose, radiation tolerance well above the levels achievable with commercial technologies (typically 10–40 Krad) is necessary in the inner detector layers.

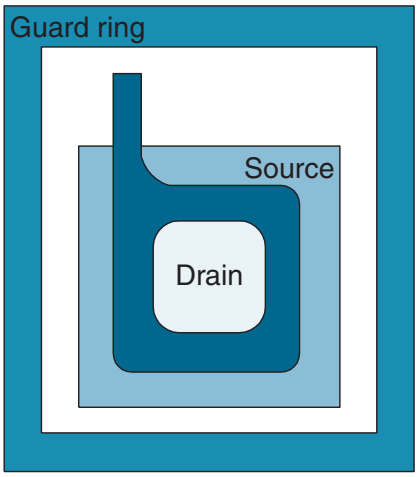

Figure 2. Layout of the ELT, surrounded by a p+ guard ring.

When exposed to an ionizing dose, CMOS technologies are affected by charge buildup and defects are created in the silicon-dioxide layers. When this occurs in the gate oxide, the effects at the transistor level are threshold voltage shift, mobility degradation, and noise increase. Charge buildup in the thicker field oxides opens leakage paths, either between source and drain of the same n-channel transistor or between adjacent n+ diffusions. Source-drain leakage occurs at the edge of the transistor, where a charge trapped in the field oxide can induce an inversion layer and eventually open a conductive path beyond the control of the gate.

The effects in the oxide are inversely proportional to the oxide thickness and sharply decrease below about 10 to12 nm. Since gate-oxide thickness is constantly reduced in each new technology generation, the dominating radiation-induced failure mechanism can be traced to the appearance of leakage currents. CMOS 6SF is no exception to this rule, and the required radiation tolerance can be reliably achieved only with the use of a particular layout technique based on ELTs (also called edgeless transistors and shown in Figure 2) and guard rings. ELTs eliminate radiation-induced leakage currents between source and drain, whilst guard rings prevent the formation of any other possible leakage path. The radiation tolerance requirements of LHC can be met by the combination of using this layout technique systematically in CMOS 6SF and the inherent radiation tolerance of the 5-nm gate oxide.

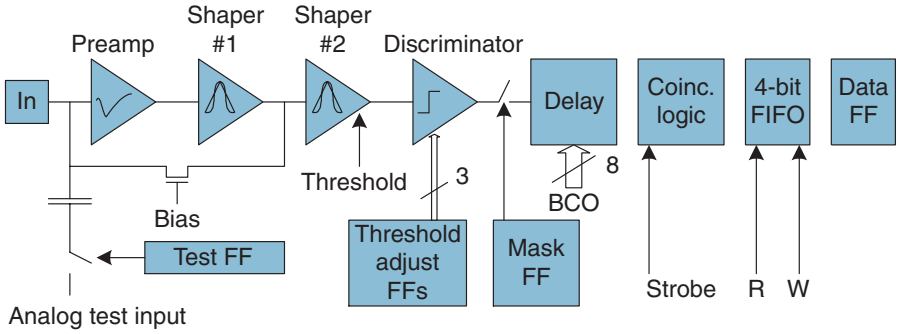

Figure 3. Architecture of one channel of the PIXEL chip.

After the verification of the effectiveness of this approach on elementary devices and basic building blocks, a significant effort was devoted to the development of a small digital library, whose cells were designed with ELTs and guard rings. This opened the way to the development of ASICs for LHC experiments.

ASICs for the LHC Experiments

This section describes a few ASICs representing good examples of some typical functions implemented in the LHC experiments. They have all been developed using the layout techniques described above. The analog blocks were fully custom-designed; the digital blocks were designed with the help of automatic tools and the dedicated library.

Readout of Particle Detectors

The PIXEL circuit is designed to read out signals from the innermost layers of the tracker detector, called the *pixel vertex detector*, of the ALICE and LHC-B experiments. The full detector consists of silicon radiation sensors based on reverse-biased PIN diodes segmented with a 50-µm fine pitch and connected by bump-bonding technology to the readout ASIC. This bonding technique makes it possible to build a hybrid pixel detector in which each sensor element is connected directly to the corresponding readout electronic channel. In the case of ALICE, these sandwiches of readout and sensor chips are assembled in several layers to cover an area of about 0.2 m^2 with ten-million channels.

One PIXEL chip has an area of about 2.1 cm^2 and integrates up to 8000 electronic channels (about 13-million transistors). Each channel comprises a low-noise preamplifier, a shaper filter amplifier, a comparator, a delay line, and readout logic, as shown in Figure 3. The channel fits in an area of 50 µm x 425 µm and consumes less than 100 µW. Testability is essential because flip-chip assembly can only be done using *known good dies* to achieve reasonable yield of a detector assembly where six chips are mounted. Every pixel in the system can therefore be fully tested using the analog input and test flip-flop. To improve the uniformity of all channels, a 3-bit threshold adjust is possible for each channel. Moreover, all biases (current and voltage) can be adjusted because they are generated on-chip by 42 8-bit digital-to-analog converters (DACs) located in the periphery at the bottom of the chip.

PASCAL, whose function is preamplification, analog storage, conversion from analog to digital, is optimized for the silicon drift detectors that will be used in the tracking apparatus of the ALICE experiment. Since only a few out of the many events produced in a typical particle collision are of interest and need to be recorded, the use of a fast analog-to-digital converter (ADC) immediately after the amplifying-filtering stage is not very efficient in terms of power consumption. PASCAL uses a more efficient approach; the output of the amplifier is sampled at the required rate by a low-power circuit implemented as an array of elementary switched-capacitor sample and holds. Only the analog samples corresponding to the interesting events are then digitized by a lower-speed ADC.

The ASIC contains 32 channels, each consisting of an amplifier and a switched capacitor array (SCA) with 256 sampling cells. The SCA allows the amplifier output to be sampled at a rate of 40 MHz with a power consumption as low as 2 mW/channel. Two adjacent channels share one ADC, based on the charge redistribution technique and featuring a resolution of 10 bits and a conversion speed of two-million samples/sec.

Timing Distribution

Correct operation of the LHC detectors is heavily dependent on the synchronization of the several electronic elements and systems that comprise the entire data acquisition chain. In the tracker detector of the CMS experiment, synchronization is achieved by accurate distribution of two signals: the *LHC Clock* (40.08 MHz clock signal) and the *First-Level Trigger Accept* signal. The clock signal indicates the precise instants when the detector signals should be sampled and stored and the trigger signal indicates that an interesting physics event is currently

stored in the detector memory circuits. It also associates, without any ambiguity, the electrical signals that result from the detection of groups of particles with the collision event that has been their origin. This identification makes possible the reconstruction of the physical phenomena being observed.

Optical fibers will be used to broadcast the timing signals. Since maintenance access to the detectors will be highly restricted and expensive, it is necessary to embed redundancy in the timing distribution system. To contain the additional system cost that would result from the duplication of every optical channel, the clock and trigger signals are combined in a single signal. This solution allows the introduction of redundancy in the system while maintaining the same number of optical-fiber links for timing distribution. However, at the receiving ends it requires the presence of ASICs to recover the clock and trigger signals.

The core functions of this circuit are the trigger decoder, a phase-locked loop (to recover the 40-MHz LHC clock signal), and a phase shifter that implements the detector-timing-calibration function (Figure 4).

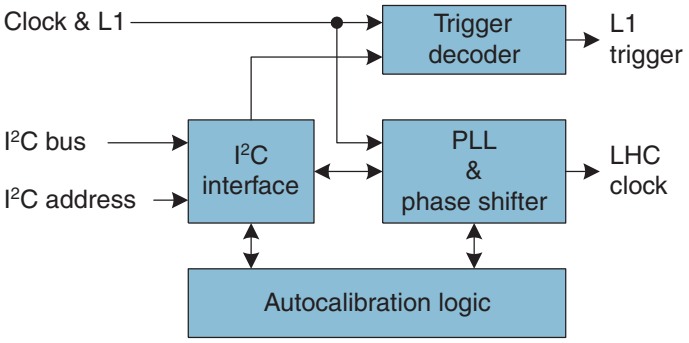

Figure 4. Block diagram of the ASIC recovering the clock and trigger signals.

Digital Data Transmission

To transmit the huge amount of raw data (on the order of several 100 TBps) from inside the particle detector to the outside computation facilities, cost-efficient bandwidth has to be provided through the use of optical links. The transmitter ASIC that has been developed for this purpose is called the gigabit optical link (GOL). The GOL can be easily combined with the existing commercial receivers located in the counting rooms because it implements two different industry-standard encoding schemes: the conditional-invert master transition (CIMT) and the 8B/10B fiber-channel encoding, which is also used in Gigabit Ethernet. Both encoding formats add 20% data overhead to limit the maximum run-length of consecutive zeros or ones.

The transmission speed can be chosen to be either 20X or 40X the 40-MHz reference clock, resulting in data rates of 0.8 or 1.6 Gbps. The interface to the detector is either 16 or 32 bits wide and, in both cases, accepts data at a rate of 40 MHz. All parameters of the circuit (for example, the bias current of the laser driver) can be programmed via either an I²C or a JTAG (IEEE 1149.1) interface. In addition, the JTAG interface is used for boundary-scan test of the printed wiring board connections and internal testing of the digital logic. The block diagram of the GOL is shown in Figure 5.

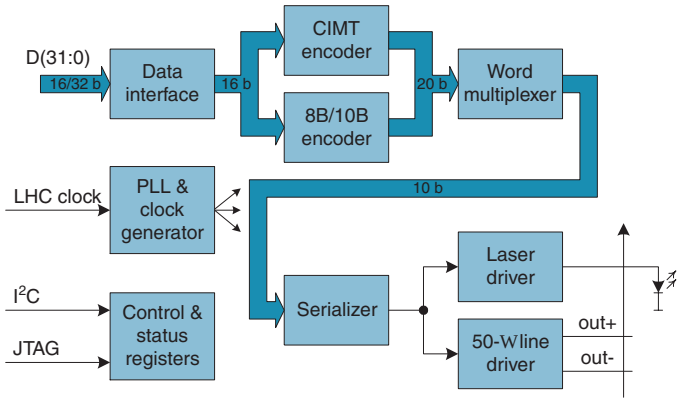

Figure 5. Architecture of the gigabit optical link (GOL) chip.

Analog Data Transmission

The LHC requirements for optical-link drivers and data-transmission receivers include high analog performance (good linearity and low noise) and low power (due to the large number of integrated channels). The three-way linear laser driver (LLD) ASIC is intended for analog data transmission as part of the 50,000 optical-fiber links of the CMS particle-tracking system. A combination of linearization methods achieves 9-bit equivalent dynamic range over a bandwidth of 100 MHz while maintaining wide input common-mode range and limited power dissipation. The linearly amplified signal is added to a dc standing current generated on-chip that can be set over a wide range (-60 mA to +60 mA). The latter capability allows tracking of changes in laser threshold that can occur from aging or radiation damage during operation in the LHC environment. The driver gain can also be adjusted to compensate for the overall gain spread, which is to be expected from the cost-optimized analog components being used for the links. The driver gain and bias current are set via an I²C serial interface.

Conclusion

By using CMOS 6SF for ASIC design, we can produce circuits that satisfy the low-power, high-density, low-noise, and high-speed requirements of the latest generation of high-energy physics experiments. By laying out the transistors as enclosed geometry devices and systematically using guard rings, it is also possible to achieve the necessary level of radiation tolerance.

•••

Dr. Federico Faccio is a staff microelectronics engineer at CERN. He has been working on radiation effects on semiconductor devices and technologies, and on the development of integrated circuits for high-energy physics applications. He is in charge of coordinating the submission of all engineering and production runs from the LHC experiments in the IBM CMOS 6SF technology. His research interests include radiation effects on bulk and silicon-on-insulator MOS and bipolar technologies and circuit design for analog signal processing and optical data transmission.

Dr. Giovanni Anelli is a staff engineer at CERN. He has worked on the characterization of radiation effects on submicron CMOS technologies and on techniques to design radiation-tolerant integrated circuits in commercial deep submicron CMOS technologies. He is currently involved in the design of several analog- and mixed-signal VLSI circuits for high-energy physics applications.

Dr. Michael Campbell worked for a short period at the Battelle Institute, Geneva, Switzerland, before joining CERN where he is a staff engineer. He has worked mainly on the design and characterization of electronics of hybrid silicon pixel detectors for high-energy physics experiments applications. His other research interests include the application

of pixel-detector technology to fields such as medical imaging.

Dr. Giovanni Cervelli is a staff engineer at CERN. His design focus is on full-custom analog drivers and receivers for optical links at the LHC. He is responsible for user-coordination and final design verification in multi-project-wafer engineering runs. He supports the custom radiation-tolerant digital library developed at CERN in CMOS 6SF.

Pierre Jarron is leader of the front-end section of the CERN Microelectronics Group. He has worked on several low-noise electronics for particle physics and on the development of the silicon strip and pixel detectors. He is currently involved in ASIC projects for the four LHC experiments.

Dr. Kostas Kloukinas is a staff microelectronics engineer at CERN. He has worked for a number of years on projects studying the radiation effects in microelectronic circuits where he mainly concentrated on the development of a radiation-tolerant standard-cell library using the IBM CMOS 6SF process. He is currently working on the design of a radiation-tolerant digital ASIC for a readout system of a silicon microstrip detector. He is also responsible for the development of a modular radiation-tolerant SRAM using the IBM CMOS 6SF process.

Alessandro Marchioro is a staff engineer at CERN where he has participated in various generations of high-energy physics experiments. He has worked on the design of hardware and software for fast real-time electronic systems for the readout of detector data from particle physics experiments. In 1992, he took a sabbatical year during which he worked at the IBM T.J. Watson Research Center, Yorktown Heights, NY, on the design of a fast microprocessor. He currently serves as head of the digital microelectronics section at CERN, where he is involved in the design of several dedicated ASICs for the LHC experiments.

Gianni Mazza is a staff engineer at INFN, Torino, Italy. He is involved in the design of mixed-mode ASICs for high-energy physics experiments and hadron therapy.

Dr. Paulo Moreira is an engineer at CERN. He worked for INESC (Portugal) in the development of telecommunication systems and in UCNW (UK) in the development of fiber-optic receivers for multi-Gbps optical communication systems. At CERN, he has been involved in the development of integrated circuits for the future LHC high-energy physics experiments. His current interests include the design of CMOS mixed-signal ASICs and the development of digital and analog data links.

Dr. Pisana Placidi earned her Ph.D. thesis on the design of a radiation-tolerant PLL in submicron technology. She has been a doctoral student at CERN. She is presently working at the University of Perugia. Her main field of interest is CMOS mixed-circuit design.

Dr. Angelo Rivetti is a research assistant at the Department of Experimental Physics of the University of Torino, Italy. He has been working on the design of amplifiers and analog-to-digital converters in submicron CMOS technologies. He now works on the design of mixed-mode integrated circuits for high-energy physics instrumentation.

Dr. Walter Snoeys is an engineer in the CERN Microelectronics Group. He was previously a consultant in the area of integrated sensors, and then worked on CMOS technology integration with Mietec, Alcatel, Belgium. At CERN, he is currently working on the design of radiation-tolerant particle-detector readout chips in commercial deep-submicron CMOS. His research interests include analog circuit design, integrated sensors, technology and device physics, and radiation effects. He has authored or co-authored more than 40 technical papers and holds one US patent.

Dr. Thomas Toifl is an engineer in the Microelectronics Group at CERN. He has been working on radiation-hard integrated circuits for the synchronization of particle physics detectors. His present research interests include the design of radiation-hard phase-locked loops and delay-locked loops, frequency detection in clock recovery circuits, and analog and digital measures for achieving robustness against radiation effects.

Dr. Kenneth Wyllie is an engineer in the CERN Electronic Design Group. He is involved in the design of mixed-mode ASICs for use in high-energy physics experiments. Prior to his current position, he worked in the field of radiation-hardened silicon sensors.

News

www.chips.ibm.com/news

IBM introduces e-business tools to speed custom-chip designs

Earlier this year, IBM announced an Internet-based custom-chip-design tool to accelerate the introduction of complex chips for high-performance servers, networking gear, and pervasive computing products.

Cadence Assura tool supports IBM SiGe technology with complete physical verification solution

Cadence Design Systems, Inc., the world's leading supplier of electronic design products and services, recently announced that its new Assura physical-verification solution has been successfully optimized to support chips designed with IBM's silicon-germanium (SiGe) technology.

IBM, Wind River announce network-processor technology collaboration

IBM and Wind River Systems, Inc., recently announced an agreement to provide network equipment vendors a comprehensive network-processor hardware and software solution for advanced routing and switching applications.

Probing the Nature of Matter with CMOS 6SF: The APV25 Readout Chip for the CMS Tracker

Marcus J. French and Lawrence L. Jones, Rutherford Appleton Laboratory, London, England, and Geoffrey Hall, Blackett Laboratory, Imperial College, London, England

The APV25 is a complex custom integrated circuit designed to read out signals from ionizing particles produced during colliding beam interactions in experiments at the international accelerator laboratory, CERN. It has been constructed using the IBM CMOS 6SF process that takes advantage of both the low-noise analog and fast digital features of the technology. The particular challenge faced in forthcoming colliding beam experiments is to build large detector systems that operate at extremely high interaction rates. These systems must be built with high-performance yet low-power electronics that will be located in inaccessible and hostile detector environments. This article explains how IBM's technology helps us meet these challenges.

Introduction

The IBM 0.25-µm CMOS 6SF technology was selected for use in experiments to answer fundamental questions about our universe at a new particle accelerator called the Large Hadron Collider (LHC) under construction at the European Center for Particle Physics (CERN). Beginning in the year 2005, physicists will probe deeper into the origin of mass with this high-energy machine by colliding two proton beams. Four major experiments are located 100 m underground in a 22-km-circumference tunnel through which the protons are transported in vacuum pipes using radio-frequency electric fields and superconducting dc magnets. Crucial to the experiments are complex detection systems where CMOS 6SF is being used to develop a new generation of instrumentation to analyze particle tracks and to *image* nuclear interactions. The digital and analog features of CMOS 6SF provide excellent system performance.

The CMOS 6SF process was initially investigated because it was known to be a high-volume, state-of-the-art technology that promised good prospects for high yield of a relatively large and complex design like the APV25. The number of chips required for LHC applications is large enough so that quality and uniformity of production are major concerns, best addressed by a well-proven technology in regular use. The thin gate oxides used in a submicron feature size technology like CMOS 6SF were also known to be a prerequisite for several performance improvements, but these types of processes were not available to CERN users in 1998 when this work began. An additional factor was the significant cost savings that could be obtained if a standard commercial process, in which stability, reliability, and process control could be relied on, was used instead of specialized custom foundries.

The interior of the region of the accelerator where experiments are performed contains concentric layers of detectors [1] that measure the trajectories of charged elementary particles that emerge from proton collisions and follow helical paths in the strong magnetic field (Figure 1). The curvature depends on particle momentum, and much of the interior of the experiment is contained within a six-meter diameter superconducting coil whose current creates the high magnetic field required to contain the particle trajectories. A few inner layers are made of pixelated semiconductor ionization detectors, followed by many layers of silicon microstrip detectors. These are followed by dense calorimeters that stop particles and absorb their energy. In this way, both momentum and energy can be measured. The scale of electronic integration in these systems is unprecedented. For example, one of the experiments is the Compact Muon Solenoid (CMS) experiment. S*olenoid* describes the magnet design of the experiment, while *muons* are the only type of charged particles that escape the experiment; their detection is important to signal possible interesting new physics. The silicon microstrip detector of the CMS experiment contains about 10-million strips and has pioneered the application of the IBM 0.25-µm technology [2]. The instrumentation is composed of about 100,000 custom circuits, called the *APV25*, in a 6-m x 2.4-m cylindrical volume in which about 50 kW will be dissipated.

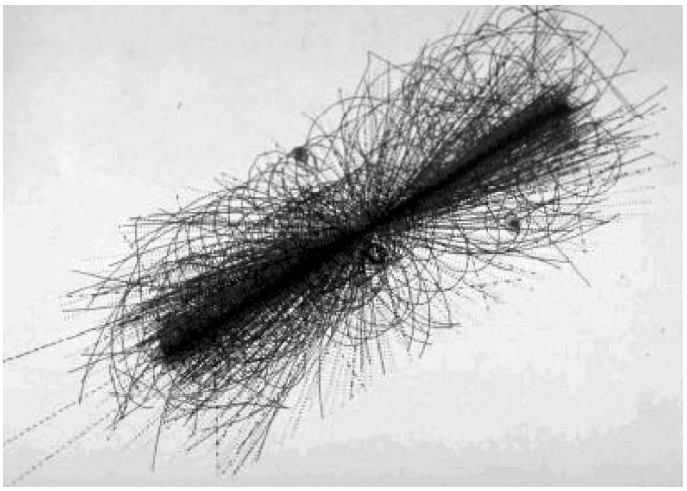

Figure 1. A typical LHC interaction in which about 20 collisions occur during each crossing of the two proton beams.

Silicon microstrip sensors are reverse-biased p-n diodes. When a high-energy charged particle crosses a layer, a column of ionization is generated. The current that then flows in an applied electric field is read out through an ac-coupled connection to the sensor. The signal is very small, around 25,000 electrons (~4 fC). Each strip capacitance is ~20 pF, and the pulse height from every channel is measured, so charge amplification with very low noise is required.

Interesting physics events can happen in any LHC beam crossing, which are spaced at 25-ns intervals. This constrains the timing performance of the system and demands fast signal processing. With 10-million analog channels, 4 x 10^{14} analog samples are collected every second. Fortunately, not all data must be converted and stored. The CMS calorimeter detectors generate triggers at up to 100 KHz for interesting interactions, so only 10^{12} samples are output from the silicon strip system each second.

Each microstrip is read out by a charge amplifier sampled at 40 MHz. Data are stored temporarily in readout circuits while the trigger decision is made, so samples are written to a capacitor-based analog memory that works as a barrel store, providing the 3- to 4-µs delay. Then, following a trigger, samples are

processed by an analog circuit on the APV25. This confines the detector signal to a single beam crossing interval and measures signal amplitude and identifies the bunch crossing.

Analog pulse height data are multiplexed differentially from pairs of front-end chips a short distance to a laser driver. Electrical-to-optical signal conversion employing edge-emitting semiconductor lasers then allows transmission over a 100-m fiber-optic cable to electronics adjacent to the cavern where the remainder of the readout system is located. A powerful array of interconnected computers on the surface above the experiment subsequently processes the data. Most other ASIC components in the system, such as the laser driver, are also constructed using the IBM 0.25-µm process.

Special System Requirements

The CMS tracker system must operate continuously with minimal downtime and high reliability because the electronics are virtually inaccessible once installed in the experiment. The tracking system is made of concentric cylinders tiled with detectors mounted on specially engineered cooling systems that maintain sensor temperatures at about -10 °C to minimize radiation damage effects (Figures 2 and 3). The sensors are wire-bonded on a fine pitch and assembled with lightweight materials to minimize unwanted particle scattering.

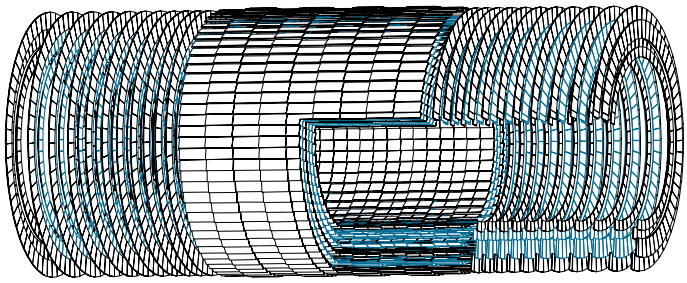

Figure 2. The layout of the microstrip tracking system. Overall dimensions are ~6 m x 2.4 m.

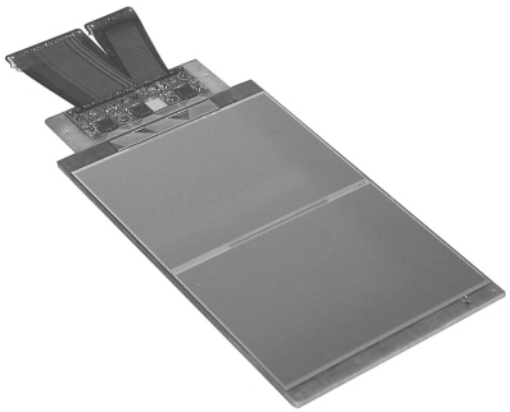

Figure 3. A prototype CMS tracker module with a 170-mm-long silicon detector readout by three APV25s.

To tightly control power dissipation, the APV25 has been designed with a budget of about 0.3 W per chip. This amounts to 25 kW plus similar-size ohmic heat losses in cables whose size is severely constrained by available space between the detectors. The very high magnetic field acts on electric currents to generate strong forces, so the mechanical design of cables and other conductors also requires care. Thus, the control of the power cabling and cooling systems is a crucial factor for having the complete system work safely.

An unusual feature of the circuit is the high channel density compared to commercial chips; 128 detector strips are read out by every APV25, which is bonded on a pitch of 44 µm at the input and arranged in two lines of staggered bond pads. While this arrangement is considerably more dense than commercial parts, it cannot be avoided while minimizing non-detecting area.

Because so many electronic channels are concentrated in a small volume with severe access limitations, complex control and test facilities are built into the APV. On-chip calibration and bias trim allow performance parameters to be measured remotely. Comparison with test data from die testing will allow monitoring of system performance over time to plan maintenance procedures.

The Signal Processing Scheme

The basic principles of signal processing in a system where signal-to-noise is important are well known. In a system with time constant τ, the equivalent noise charge *(ENC)* is simply related to input capacitance:

$$ENC^2 = \frac{\alpha C^2}{\tau} + \beta\tau + \delta C^2 . \quad (1)$$

In this expression, *C* represents the total capacitance at the amplifier, which is usually dominated by the detector. α, β, and δ are constants determined by amplifier technology and characteristic noise sources; δ is usually unimportant provided *1/f* noise is small.

At the LHC, the optimum time constant cannot be used because of the high interaction rate. Fast shaping is essential to reduce signal pile-up, but it worsens noise from a CMOS amplifier, where the constant α usually dominates; but slow pulse shaping leads to higher shot noise after detector radiation damage. Deconvolution offers a way to use either fast or slow shaping, depending on operating conditions [3]. An important consequence of the method is that it can be implemented as an elementary CMOS circuit with low-power consumption.

The Deconvolution Method

In a linear amplifying system with an impulse response *h(t)*, the output *v(t)* from an initial signal *s(t)* can be written as a convolution integral:

$$v(t) = \int_{-\infty}^{t} h(t\text{-}t') \cdot s(t') \cdot dt' . \quad (2)$$

If the amplifier output voltage is sampled at regular intervals so that the measurements of *v(t)*, for example, are *(v1 v2 v3 v4 v5 ...),* it is convenient to write the equation in matrix form as:

$$V_i = \sum_j H_{ij} S_j \text{ or } V = H \cdot S \text{ and } S = W \cdot V = H^{-1} H \cdot S . \quad (3)$$

A faster pulse *S,* which has an effective presence in only one beam-crossing interval, can be formed from a weighted sum of several consecutive time samples. For the CR-RC filter commonly used in amplifier systems, the elements of the weight matrix *W* reduce to a particularly simple form. A suitable switched capacitor filter can be constructed by forming a weighted sum of only three consecutive voltage samples (Figure 4).

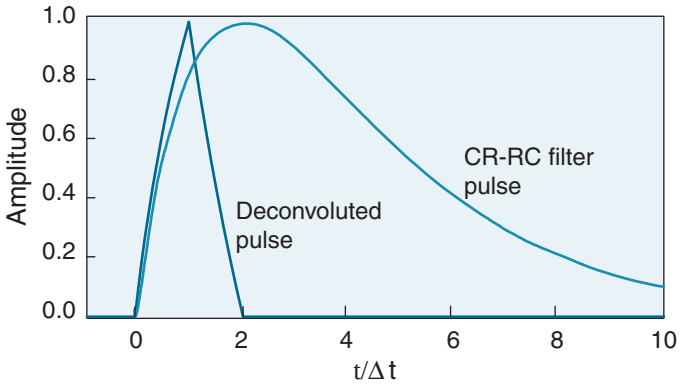

Figure 4. Ideal CR-RC pulse shape and the result after deconvolution. t is the time interval between samples.

APV25 Architecture

Each APV25 channel (Figure 5) contains a preamplifier and shaper followed by a 192-column deep memory into which samples are written at 40 MHz [4,5]. A switchable unity gain inverter is used to permit signals of either polarity. The shaper is an effective 50-ns CR-RC filter with shaping adjustable over a limited range. Up to 30 locations of data awaiting readout are flagged so that they are not overwritten. The chip can be operated in three modes: peak, deconvolution, or multi-sample readout mode. The APV's system features include programmable on-chip analog-bias networks, remotely controllable internal test-pulse generation, and a slow control interface.

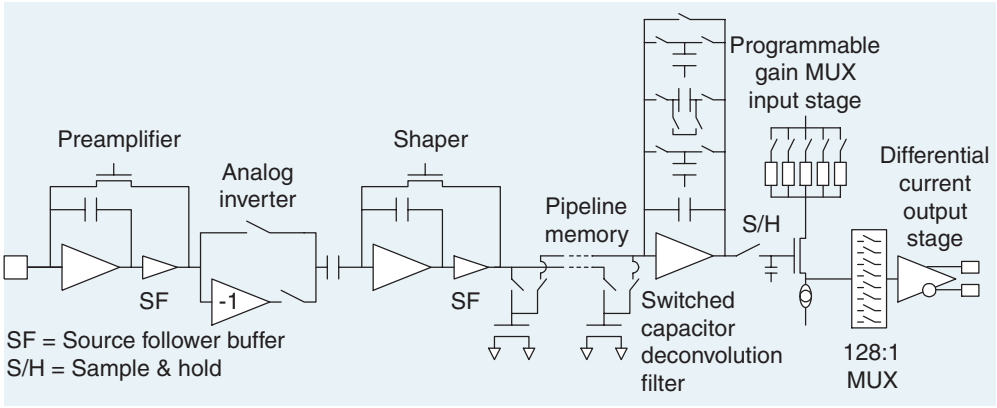

Figure 5. Schematic of the APV25 functionality.

The pipeline is a 128 x 192 array of switched capacitor cells that was designed by using the gate capacitance of NMOS transistors biased in strong inversion. Each cell comprises two transistors to perform the read or write operations and a storage capacitor. The pipeline is read out by the deconvolution processor, which is an amplifier with a capacitor network in the feedback loop. Capacitor ratios define weights used by the deconvolution algorithm.

A 128:1 multiplexer drives analog output from the chip after voltage-to-current conversion by a programmable resistor. Data from two APV chips are interleaved to arrive at the final transmission speed of 40 million samples per second.

The chip is 7.1 mm x 8.1 mm. APV25-s0, the first version, was delivered in October 1999 and demonstrated excellent performance, including radiation tolerance. Minor deficiencies were found in uniformity of the on-chip calibration signal, lower gain than design, and a digital timing error transparent to the user. Noise performance was satisfactory but depended on channel location. This arose from metallization resistance because tracks from input pads at the bottom edge of the chip were significantly longer than those at the top. All these deficiencies were corrected in the final APV25-s delivered in September 2000.

Approximately 500 APV25-s0 chips from four wafers were available for evaluation. An automatic-probe test facility had been developed to screen die during production, and excellent yield was measured.

Analog Design

The preamplifier is a charge-sensitive amplifier with a PMOS input transistor with dimensions of 2000 µm/0.36 µm and current of 400 µA. It consumes 0.9 mW and is the dominant contribution to the total APV25 power budget of 2.3 mW/channel. Test transistors of both NMOS and PMOS types were built on prototyping runs. Figure 6 shows measured noise as a function of frequency. The results indicate why the p-MOSFET was selected. The size and bias conditions were chosen to give optimal performance in the APV25 design. The total front-end gain is approximately 100 mV/MIP (MIP = minimum ionizing particle signal ≈ 25,000 electrons).

Figure 7 shows average amplifier pulse shapes in peak and deconvolution modes [6]. The peak mode pulse is a good approximation to the ideal, and deconvoluted data illustrate the effectiveness of the technique in achieving a short pulse shape. Figure 8 shows typical noise measurements in both modes; the increase in noise after deconvolution is a simple consequence of the fast filter.

The noise target performance for silicon microstrips in CMS is 2000 electrons, which can be achieved (assuming amplifier noise alone) for detectors with capacitance up to 25 pF. Lab measurements using a ^{90}Sr β source are shown in Figures 9a and 9b.

Uniformity of the analog amplifier and storage elements is crucial to system performance. Wide variance in capacitor value or amplifier gain would lead to additional noise in the system. The impact of pipeline cell non-uniformity can be demonstrated by storing and retrieving a signal at every pipeline location. Figure 10 shows the gain for all 192 pipeline cells for one channel. The width of the gain distribution indicates close capacitance matching between cells.

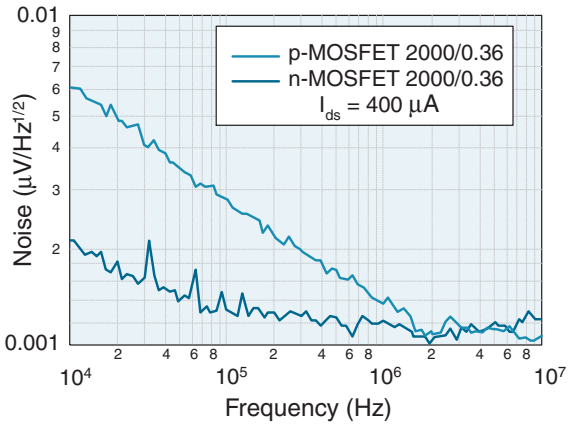

Figure 6. Spectral noise density of NMOS and PMOS transistors.

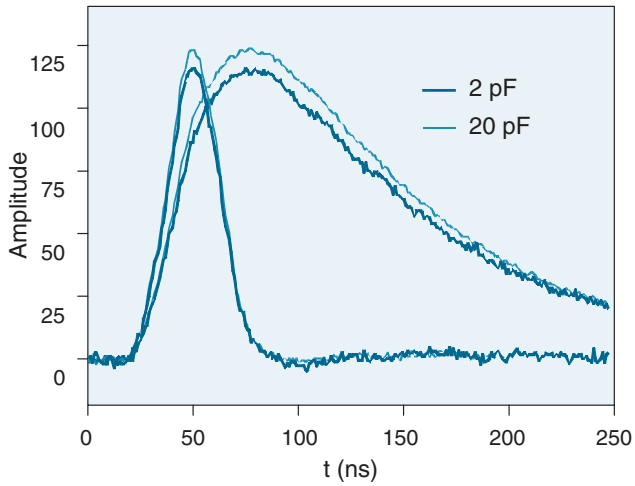

Figure 7. APV25-s1 pulse shapes for various input capacitances.

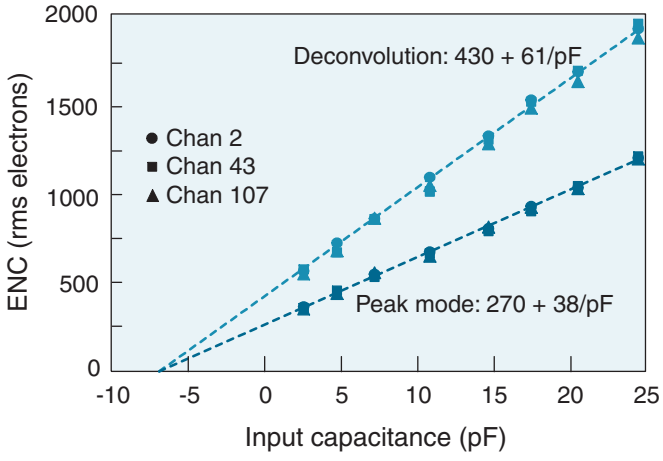

Figure 8. APV25-s1 noise dependence on input capacitance.

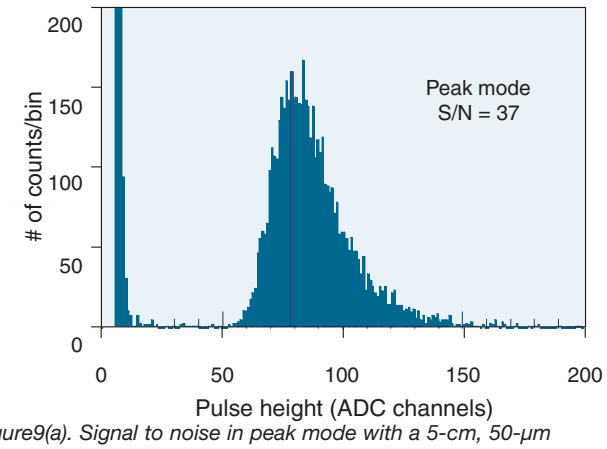

Figure9(a). Signal to noise in peak mode with a 5-cm, 50-µm microstrip.

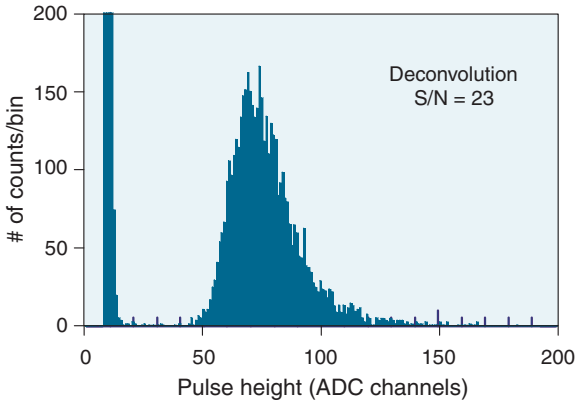

Figure 9(b). Signal to noise in deconvolution mode with a 5-cm, 50-µm microstrip.

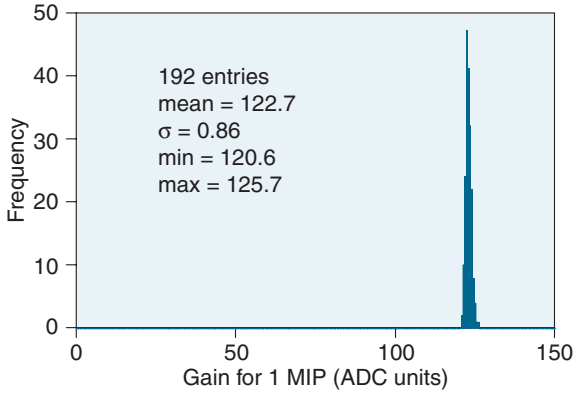

Figure 10. Pipeline gain uniformity.

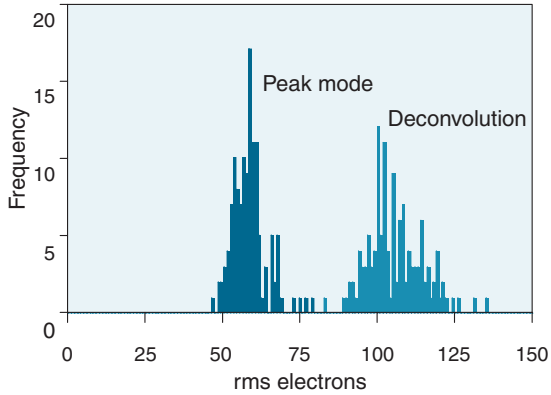

Figure 11. Pipeline baseline level fluctuations for all 128 channels in units of signal electrons.

Another evaluation can be made by studying the effective noise contribution to a single channel from pedestal variation with pipeline location. After converting to equivalent noise charge (Figure 11), it is clear that the variation between channels is negligible in both modes of chip operation.

Conclusion

The APV25-s1 readout chip has shown outstanding performance, benefiting considerably from the CMOS 6SF technology features. Large-scale construction of the CMS microstrip tracker system is due to begin this year.

Acknowledgements

We are grateful for the support of this work by the UK Particle Physics and Astronomy Research Council. We thank Markus Friedl for the photograph of the detector module.

References

[1] *CMS Tracker Project TDR & Addendum, CERN/LHCC 98-6*, 1998, and *CERN/LHCC 2000-016*, 2000.

[2] A. Marchioro, "Deep Submicron Technologies for HEP," *Proceedings of the 4th Workshop on Electronics for LHC Experiments, CERN/LHCC/98-36*, 1998:40-46.

[3] S. Gadomski et al., "The Deconvolution Method of Fast Pulse Shaping at Hadron Colliders," *Nuclear Instruments and Methods in Physics Research*, A320, 1992:217.

[4] APV25-s1 User Manual, *http://www.te.rl.ac.uk/med/*

[5] L. Jones et al., "The APV25 Deep Submicron Readout Chip for CMS Detectors," *Proceedings of the 5th Workshop on Electronics for LHC Experiments, CERN/LHCC/99-09*, 1999:162-166.

[6] M. Raymond et al., "The APV25 Deep Submicron Readout Chip for CMS Detectors," *Proceedings of the 6th Workshop on Electronics for LHC Experiments, CERN/LHCC/2000-41*, 2000:130-134.

•••

Marcus J. French is responsible for the microelectronics design group at the Rutherford Appleton Laboratory where most of the APV25 design work was done. He has worked in the field of ASIC design for 14 years and has worked on numerous design projects for custom instrumentation systems.

Lawrence L. Jones is the lead engineer for the design project. He has worked at the Rutherford Appleton Laboratory for 11 years and designed many mixed-signal ASICs. More recently, he has been focusing on high-resolution data conversion circuits for imaging systems.

Professor Geoff Hall is a particle physicist at Imperial College, London, and is responsible for managing the electronics for the tracking system in CMS. Starting from a Ph.D thesis and several years' work in bubble chamber physics, he began serious development of electronic particle detectors about 20 years ago in the hopes of doing better and faster experiments on newly discovered quarks. Most of his research work for the last decade has focused on LHC experiments, particularly CMS, and the group he leads has done much original work investigating radiation damage in silicon detectors and microelectronics. A part-time activity is x-ray photon detection, with special reference to synchrotron sources and possible biomedical applications.

For information on IBM's CMOS 6SF technology, please visit:
http://www.chips.ibm.com/services/foundry/offerings/cmos/6sf/

For more information on the APV25 or the Microelectronics Design Group Instrumentation Department of the Central Laboratory of the Research Councils, please visit: *http://www.te.rl.ac.uk/med/*

News

www.chips.ibm.com/news

IBM releases Linux-based network-processor software development tools

IBM unveiled an integrated suite of Linux-based software development tools and services that support the IBM PowerNP™ network processor.

IBM and LSI Logic announce technology licensing agreement

IBM and LSI Logic Corporation announced a technology licensing agreement intended to accelerate the integration of high-performance digital signal processor capability into custom chips for next-generation networking equipment, wireless handsets, and other advanced communications products.

TeleHubLink Corporation announces IBM as its microchip technology supplier; THLC's microchip prototype to be developed with IBM's deep submicron semiconductor technology

TeleHubLink Corporation, a secure wireless and broadband communications encryption provider, has announced that the company has signed an agreement with IBM for the design and prototyping of its secure communications ASIC microchips called Hornet.