A Digital Front-end and Readout Microsystem for Calorimetry at LHC

The FERMI project

A. Dell'Acqua*, M. Hansen, B. Lofstedt, J.P. Vanuxem
CERN, CH-1211 Geneva 23, Switzerland

C. Svensson, J. Yuan
Department of Physics and Measurement Technology, University of Linköping, S-581 83 Linköping, Sweden

H. Hentzell
Center for Industrial Microelectronics and Materials Technology, University of Linköping, S-583 30 Linköping, Sweden

C. Alippi, L. Breveglieri, L. Dadda, V. Piuri, F. Salice, M. Sami, R. Stefanelli
Dipartimento di Elettronica, Politecnico di Milano, Sezione INFN, Pavia, I-20133 Milano, Italy

P. Cattaneo, F. Fumagalli, V. G. Goggi
Dipartimento di Fisica Nucleare e Teorica dell' Universita' e Sezione INFN, Pavia, I-27100 Pavia, Italy

S. Brigati, U. Gatti, F. Maloberti, G. Torelli
Dipartimento di Elettronica dell' Universita' e Sezione INFN, Pavia, I-27100 Pavia, Italy

P. Carlson, C. Fuglesang, A. Kerek
Manne Siegbahn Institute of Physics, S-104 05 Stockholm, Sweden

G. Appelquist, S. Berglund, C. Bohm, N. Yamdagni
Fysikum, University of Stockholm, S-113 46 Stockholm, Sweden

R. Sundblad
SiCon AB, S-583 30 Linköping, Sweden

* on leave from Dipartimento di Fisica Nucleare e Teorica dell' Universita' e Sezione INFN, Pavia

Abstract

We present a digital solution to the front-end electronics for calorimetric detectors at future supercolliders based on high speed A/D converters, a fully programmable pipeline/digital filter chain and local intelligence. Questions of error correction, fault-tolerance and system redundancy are also considered.

A system integration of a multichannel device in a multi-chip, Silicon-on-Silicon Microsystem hybrid will be used. This solution allows a new level of integration of complex analog and digital functions, with an excellent flexibility in mixing technologies for the different functional blocks.

This type of VLSI multichip integration allows a high degree of programmability at both the function and the system level, and offers the possibility of customizing the microsystem with detector-specific functions.

INTRODUCTION

The electronics for signal detection and triggering for experiments at future proton supercolliders will be one of the most complex systems conceived so far for high energy physics. For calorimeters in particular, complexity, speed and decision (trigger) requirements call for an unprecedented level of data identification, compaction and processing already at detector level. That, and the requirements of massive trigger processing early in the acquisition chain, require the front-end stage of these detectors to include a major portion of the total electronics system.

Front-end signal processing for calorimetric detectors is essential in order to achieve adequate selectivity in the trigger functions, and data identification and compaction before read out required in the high rate environment of a high luminosity hadron machine. Other crucial considerations are the extremely wide dynamic range and bandwidth requirements for the front-end electronics, as well as the volume of data to be transferred to the subsequent stages of the trigger and readout system [1]. These requirements are best met by an early digitization of the detector information, followed by on-chip digital signal processing and buffering functions at both the first-level and second-level decision latencies.

In the following we will describe FERMI [2], a dedicated R&D project for digital readout and processing of calorimeter signals.

PHYSICS REQUIREMENTS AND GLOBAL PARAMETERS

Dynamic range and resolution

The dynamic range necessary for calorimetry in the multi-TeV energy region is dictated by physics considerations and detector characteristics.

0018-9499/93$03.00 © 1993 IEEE

At the higher end of the energy spectrum, the basic parameters are maximum detected particle (e.g. electron) energy and detector granularity. Lateral calorimeter segmentations, which optimize both energy and position resolution, correspond to deposits of a large fraction of electron energy in a single cell. At the lower end of the spectrum, energy isolation criteria for the identification of electrons in a hadronic environment, as well as accurate measurements of hadron jet energies, require adequate sensitivity down to the level of the noise generated by event pileup and electronics. Assuming a maximum detectable mass of about 4 to 6 TeV for an hypothetical Z' boson decaying into an electron pair, and an overall noise level well below 100 MeV per calorimeter cell, the required dynamic range turns out to be 15 to 16 bits. Since the noise level and the fraction of the total energy deposited in a calorimeter cell tend to scale in first approximation with its size, this dynamic range will not depend, to first order, on the details of the geometrical segmentation of the detector.

It is not feasible, with today's technology, to implement A/D converters with more than 10 bit resolution at the sampling speed (60 to 70 MHz) required by the bunch frequency of the proposed supercolliders. Therefore, a compression of the dynamic range must be done before applying the analog signal to the converter. This compression has to obey the criteria outlined below.

The A/D converter resolution and the characteristics of the transfer function of the compression are determined by the energy resolution of the detector and by the requirement that quantization noise does not contribute appreciably to the overall resolution of the system.

The energy resolution $\sigma_d(E)$ of the detector is usually parametrized as a function of energy [3] as:

$$\frac{\sigma_d(E)}{E} = \frac{a}{E} \oplus \frac{b}{\sqrt{E}} \oplus c \qquad (1)$$

where a is the noise term, b a scaling term and c a constant. A quantization process matched to the detector resolution implies:

$$LSB(E) \leq \frac{\sigma_d(E)}{k} \qquad (2)$$

If n(E) is the (non-linear) transfer function of the A/D converter,

$$\frac{dn(E)}{dE} = \frac{1}{LSB(E)} \qquad (3a)$$

then eq.(2) translates into a bound on the derivative of the transfer function,

$$\frac{dn(E)}{dE} \geq \frac{k}{\sigma_d(E)} \qquad (3b)$$

The value of k is set by the allowed contribution of the quantization noise σ_n, to the overall resolution:

$$\sigma_n(E) = \frac{\sigma_d(E)}{k\sqrt{12}} \qquad \sigma_{tot}(E) = \sigma_d(E) \oplus \sigma_n(E) \qquad (4)$$

The combined effects of all noise sources (including conversion) on the theoretical calorimeter performance can then be described by the fractional increase of the detector resolution:

$$R(E) = \frac{\sigma_{tot}(E)}{\sigma_d(E)} - 1 \qquad (5)$$

Any practical implementation of a compression-converter system should produce a function R(E) bounded from above at all values of E.

The normalized integral of expression (3b) can be used to define the minimal transfer function $n_{min}(E)$ which provides the relevant bounds on a wide family of possible transfer functions. The detailed functional form of any transfer function fulfilling these bounds is largely irrelevant. This holds true provided it obeys the following general requirements at all energies; of being above the minimal transfer function, both in value and first derivative.

Transfer functions

Practical implementations of non-linear transfer functions satisfying the bounds given by $n_{min}(E)$ can be obtained with piecewise approximations having linear and/or logarithmic segments. A practical compressor circuit using a 4-fold piecewise linear transfer function, together with a 10-bit A/D converter, has been implemented.

Fig. 1 shows this transfer function, together with the "minimal" ones for two detector families:

A : b=0.10, c=0.01, representative of scintillating fibre or liquid argon calorimeters;
B: b=0.04, c=0.005, representative of high performance scintillating crystal calorimeters.

The "minimal" functions result from integrating expression (3b) in the energy range between 0.1 GeV and 2 TeV, with $k=1/\sqrt{2}$ (corresponding to an upper limit of R(E) = 0.08). All transfer curves are normalized to two counts at E=0.1 GeV.

For both detectors a noise term a=0.05 GeV was assumed.

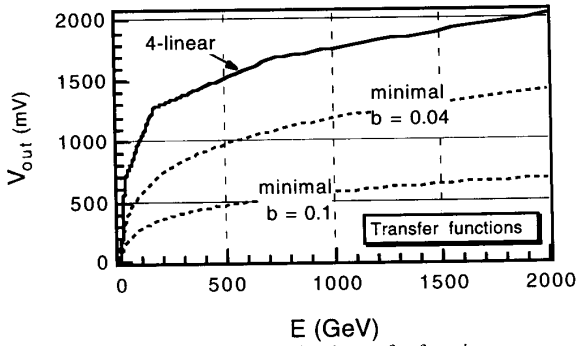

Fig. 1 Actual and minimal transfer functions

Fig 2 shows the measured quantization noise of the combination of the 4-linear compressor and a 10-bit A/D converter as a function of energy. The fractional resolution of the two detector families described above are also shown. It can be seen from the figure that the contribution of the quantization noise to the total error will be small.

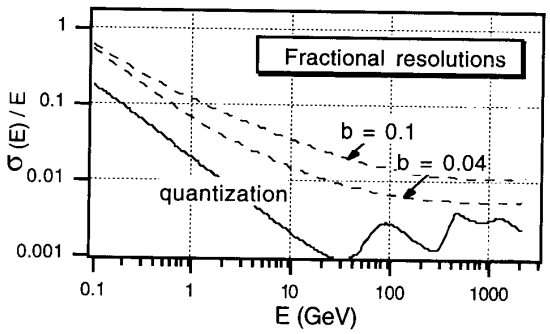

Fig. 2 Quantization noise compared to detector resolutions

The *relative* effects of the composite transfer function-digitzation described above on detector resolutions are shown in fig. 3.

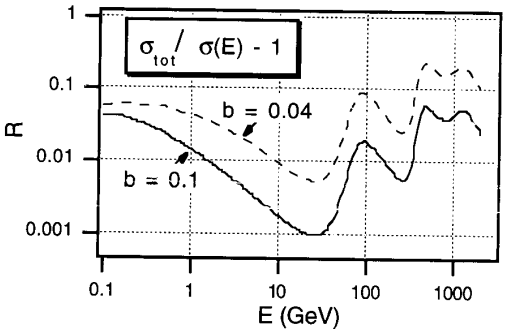

Fig. 3 Relative error contribution from the compressor/converter

On the average, up to 100 GeV the intrinsic resolution of the calorimeter increases by a *relative* factor of the order of a few percent. For higher energies, where the constant term dominates, the overall effect of the resolution is of the order of 0.001.

It can be concluded from the above discussion that the combination of a suitable compression scheme and a practical 10 bit A/D converter can be an adequate solution for early digitization of the data from most types of calorimetric detectors in the LHC/SSC energy range.

DESCRIPTION OF THE FERMI STRUCTURE

FERMI is a multichannel data acquisition and signal processing module implemented as a silicon-on-silicon multichip microsystem. A microsystem implementation means that the module is realized as a chip carrier with a large silicon multi-layer substrate serving as an active circuit board on which integrated circuits are bonded. The substrate contains also diffused and surface-mount components, interconnections and transmission lines [4].

The FERMI microsystem is designed to perform dynamic range compression, 67 MHz digitization, digital signal processing (DSP) and buffering of the data up to and including second level trigger latency. It is intended to be mounted directly on the detector, along with a multi-FERMI controller performing control, calibration and other system functions. Due to the high level of radiation and limited accessibility, the design will incorporate a high degree of redundancy, fault- and radiation- tolerance.

In the current design, nine parallel channels are digitized every 15 ns and the results stored in a dual-port memory until the decisions from the first and second level triggers become available. The storage occurs in memory positions given by a memory management unit, the address generator [5]. A temporal environment (a time frame containing a number of consecutive samples) is associated with each event accepted by the first level trigger. All memory locations containing sample points corresponding to negative first level decision are returned to the pool of free memory. Time frame memory locations are also returned to the free pool as a result of a second level trigger reject or after having been read out to the third level.

The FERMI readout facility contains provisions for full or reduced readout of time frames. In the latter case the time frame is summarized by means of digitally filtered data. It is envisaged that the reduced readout should be default when pushing data to the second level trigger processors. These could, however, return for selective readout (pull) of full time frames, which might be necessary in order to resolve "pile-up" conditions. The data transfer from FERMI can thus be characterised as "push with selective pull".

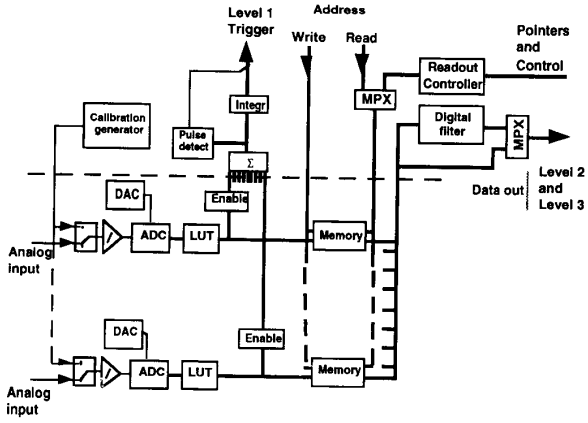

Fig. 4 An overview of FERMI

Fig. 4 contains all major FERMI components in the form of a simplified block diagram. It is divided into two main parts: an acquisition part and a common service part. The former consists of twelve identical channels on three identical chips, while the latter is unique. Only nine channels are used for data acquisition, the others serving as a redundant elements.

To achieve a matching reliability in the service part, it is designed with a high degree of fault tolerance.

The acquisition channel begins with a non-linear amplifier to achieve signal compression. The compressed signal is applied to the A/D converter where it is sampled at the rising edge of the internal 67 MHz clock. This event also starts the digitization process. The resulting 10-bit data are linearized and expanded to the full 16-bit dynamic range linear representation by means of a 1024 word look-up table (LUT) containing the inverse transfer function of the detector - electronic chain. The expanded data are stored in the data memory at locations supplied by external circuitry (pointers from the address generator, embedded in an error correction/detection envelope).

To avoid underflow conditions with bipolar shaped pulses a programmable DC level is added to the analog signal before the A/D conversion.

In order to provide information for the first level trigger the expanded data from all active channels are summed and fed to a digital discriminator for pulse identification. This module-wide data is also summed over time to emulate pulse integration, with the result truncated to 12 bits in order to reduce the complexity of the trigger system. This precision is considered sufficient for the first level trigger algorithms.

The digital discriminator, emulating an analog constant-fraction discriminator, produces a module-level pulse detect flag. This flag is sent, along with the trigger data, to the first level trigger. The discriminator output also sets a "pulse-detect" flip/flop, which resets after a predetermined time. A new pulse detection, occurring while "pulse_detected" is active, will set the pile-up detect flag. This flag is submitted together with the data to warn later processes that a pile-up condition occurred at this point.

External circuitry keeps track of where all samples are stored until the first level trigger has made its decision. If a sample point is accepted as a valid event, a predefined temporal environment (a time frame) is kept in the memory while locations that do not belong to such intervals are released. Released locations are recycled to be used as new input pointers.

External circuitry is also responsible for translating the event identifier, in a read command, into a pointer sequence corresponding to the time frame in question. The pointer sequence, together with readout control information is transferred to the attached FERMIs where it is stored in the readout controller. The readout controller initiates the required sequence i.e. a full time-frame readout or a digitally filtered readout.

A detailed description of the system aspects can be found in [6].

THE COMPRESSOR AMPLIFIER

FERMI receives analog signals with a dynamic range of 15 to 16 bits, but will use a 10-bit A/D to convert the information into digital form. As shown above, this gives only a small contribution to the total detector resolution, provided that the non-linear transfer function is appropriately matched to the resolution of the calorimeter.

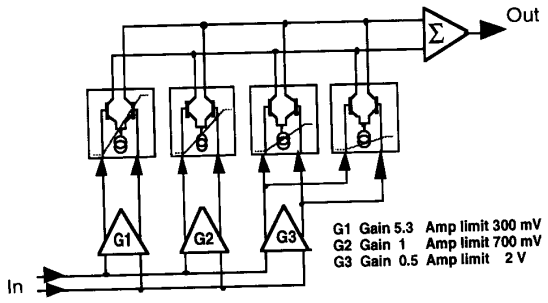

Fig. 5 Block diagram of the compressor amplifier

Two approaches have been studied in more detail, a circuit with a linear-logarithmic response, and another with a piecewise linear transfer characteristic. The preferred solution is provided by the second one, where the input signal is amplified in four stages with different gains and the outputs added linearly (Fig. 5)[7].

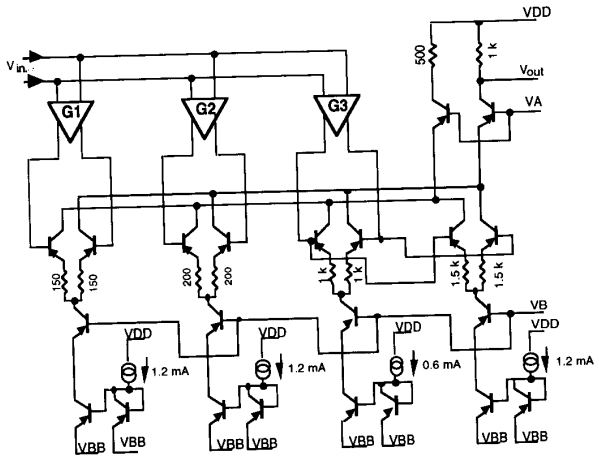

Fig 6 Circuit diagram of the compressor amplifier

The circuit currently under test is shown in detail in Fig. 6. The gains are defined by the ratio of the load resistance of the common output stage to the individual emitter feedback resistances, and the break points are set by the current sources.

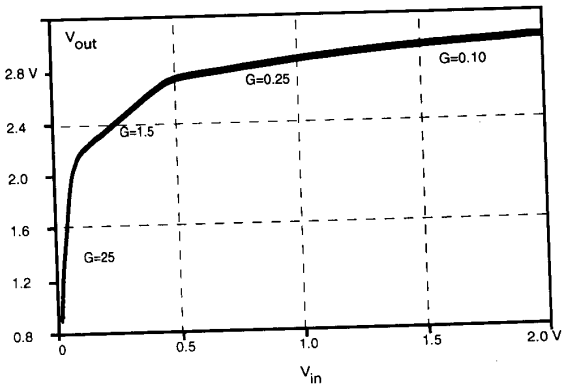

Fig. 7 Effect of component tolerances on the transfer function of the compressor

A SPICE analysis of the circuit of Fig. 6 was performed using parameters for the Gennum GA911 process, a medium performance technology (f_T = 2.5 GHz) available as prediffused analog arrays.

Fig. 7 gives the result of a Monte Carlo run, where the gain-defining passive components were varied within the tolerances given by the manufacturer. The width of the band indicates the expected channel-to-channel variations.

The settling behaviour of the circuit is good, thanks to the fast npn-transistors used. Fig. 8 shows the settling for a set of input signals ranging from 5 mV to 2.0V (2.5 LSB to full ADC range). The input is a ramp with a rise time of 5 ns.

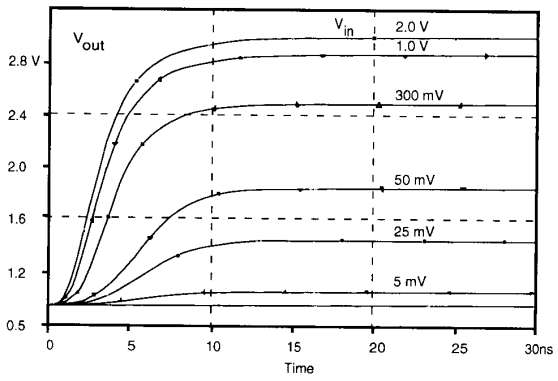

Fig. 8 Settling time of the compressor

For an input signal of 50 mV (i.e. 50 GeV), the settling time of the output to better than 1 LSB (2 mV) is within the sampling interval i. e. 15 ns, as shown in Fig. 9.

The SPICE analysis of the circuit indicates a dissipation of around 60 mW per channel. The circuit has been implemented using the GA911 process and is currently under test.

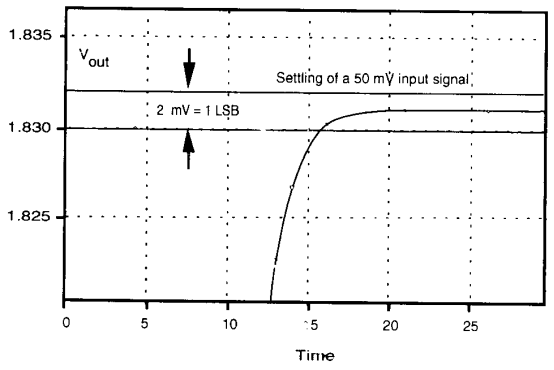

Fig. 9 Detail of the settling time (to within 1 LSB)

SAMPLING AND A/D CONVERSION

We are currently following two different approaches to the Analog to Digital conversion, a "Pipeline flash ADC" and an array of Successive Approximation converters. A final choice can only be made once prototypes have been thoroughly tested.

The pipeline flash ADC

System overview

The block diagram of the converter is shown in fig. 10. A first coarse conversion is carried out by means of a flash ADC, providing the m most significant bits. These bits are fed to a DAC generating a voltage to be subtracted from the analog input voltage.

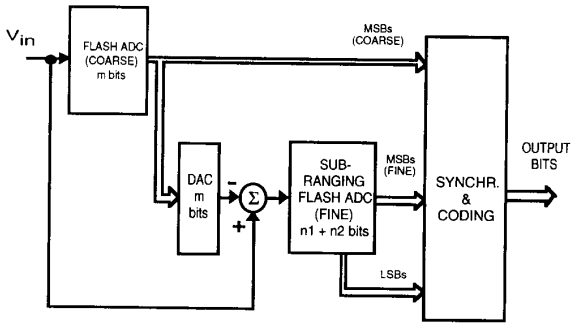

Fig. 10 Block diagram of the pipeline flash A/D converter

The obtained residue is converted with a two-step flash ADC, thus obtaining the remaining ($n_1 + n_2$) bits. Finally, the converted bits ($m + n_1 + n_2$) are combined to form the output digital word. As the target is to have 10 bits of accuracy, $m + n_1 + n_2$ should be equal to 10. In order to compensate for errors in the subtraction and mismatches between the coarse ADC and the DAC, the resolution of the coarse flash converter is augmented by one bit. In the proposed converter, we set $m = 5$ and $n_1 + n_2 = 6$. This turns out to be the best compromise between complexity, resolution and speed.

Flash ADC

The coarse ADC uses a conventional 5 bit flash architecture. The output is applied to a D/A converter based on a resistive string. This generates an analog voltage which is subtracted from the input voltage to obtain a residue which is then fed to the two-step flash ADC to determine the remaining bits. The resistive string of the DAC is matched with the string contained in the coarse flash ADC converter.

Two-step flash ADC

The second ADC operates following a two-step flash subranging approach, where the reference levels are generated by means of an intermeshed resistive network [8] (fig. 11).

A main resistive string (coarse resistor section) defines the reference voltage levels for the n_1 most significant bits (MSBs). 8 levels are generated to obtain the 3 MSBs. These levels are further divided by additional resistive strings (fine resistor section), which define the reference levels for the n_2 least significant bits (LSBs: $n_2 = 3$). The MSBs are evaluated by an array of 8 high-speed comparators. Based on the result of the coarse comparison, a "LSB switch control" selects the

appropriate fine resistive string for the fine comparison which is performed by means of a second array of 8 comparators.

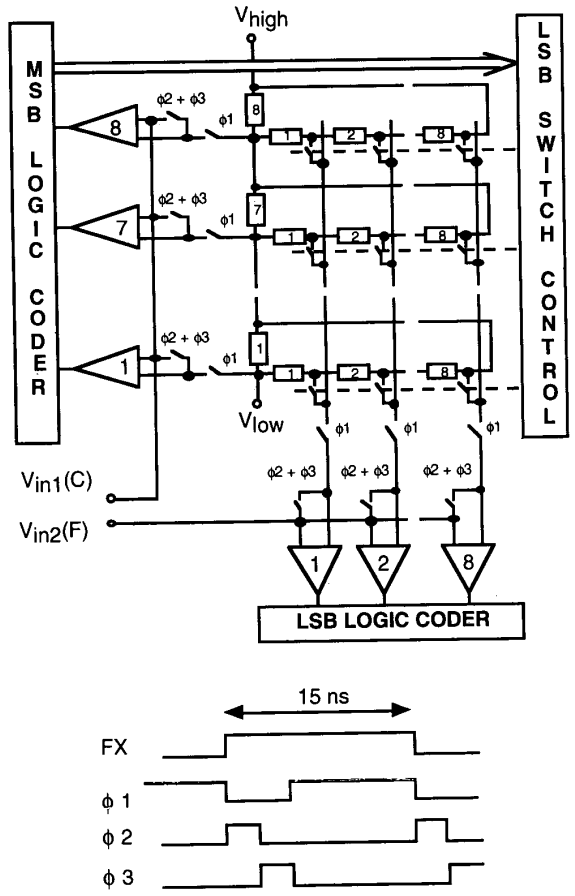

Fig. 11 Architecture of the subranging ADC
C = coarse comparison. F = fine comparison.
FX is the generic operating phase of the SU&H circuits.
ϕ1, ϕ2, ϕ3 are the operating phases of the comparators.

To achieve a very high conversion speed, the converter works in a pipeline fashion. While a residue is being processed for the coarse comparison, the previous residue is processed for the fine comparison. The results of the coarse and fine comparison are then synchronized by means of two logic coder networks. These are, in turn, combined with the MSBs given by the coarse flash converter for the same input sample.

A sample-and-hold stage is needed in front of the subranging converter in order to feed the same residue to the two conversion sections at different times. The same stage is also used to perform the subtraction between the analog input sample and the result of the coarse D/A conversion, thereby giving rise to a "subtract-and-hold" (SU&H) operation. In order to obtain the required accuracy, the operation of the SU&H circuit developed for this converter is divided into four phases, i.e. a sampling phase, a subtract and prediction phase, a first hold phase (for the coarse comparison) and a second hold phase (for the fine comparison).

The length of each operating phase of the SU&H circuits is = 15 ns. Therefore, to achieve the required high-frequency operation (67 MHz conversion rate), the proposed system uses a bank of five interleaved SU&H circuits (fig. 12). For the first SU&H circuit, sampling is performed during phase Φ1, prediction is performed during phase Φ3, and the output value is held during phase Φ4 (for the coarse comparison) and during phase Φ5 (for the fine comparison). For the second SU&H circuit, the four successive operating phases are Φ2 (sampling), Φ4 (prediction), Φ5 (hold for coarse comparison) and Φ1 (hold for fine comparison). A similar shift is present for the operating phases of the other SU&H circuits.

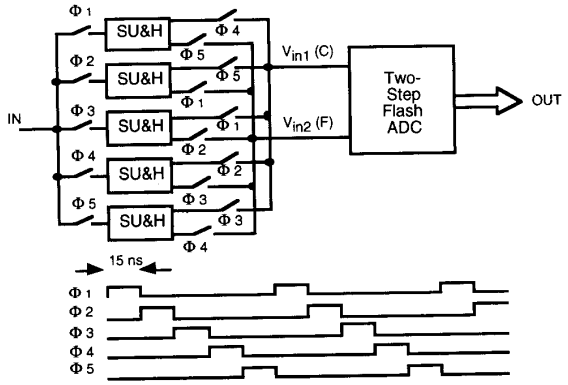

Fig. 12 Block diagram of the subtract-and-hold section.

To obtain a high accuracy, the two-step flash ADC comparators need three operating phases, i.e. a reset phase (for recovery from overdrive conditions), an autozero phase (for offset compensation) and a comparison phase. The comparison phase is the most critical one, therefore this phase has been set twice as long as the other two phases (the total sampling period is 15 ns). Therefore, a 268 MHz master clock is required to generate the three clock phases for the comparators. This clock has to be exactly synchronized to the experiment bunch crossing frequency, as the sampling phase clocks (Φ1, Φ2, Φ3, Φ4, Φ5) are derived from it.

The resistor matrix used in the subranging ADC is not very large. Therefore, the time constants associated with the resistive network and the parasitic resistances and capacitances of the analog switches allow the converter to have full resolution at the required sampling frequency. The overall number of comparators required is small compared to a single two-step flash converter, thereby achieving a reduction in silicon area and power consumption.

Subtract-and-hold circuit

The subtract-and-hold operation is based on the sample-and-hold (S&H) architecture shown in fig. 13 [9]. The S&H is made up of two transconductance operational amplifiers (OTA's), A1 and A2, whose outputs are connected in parallel [10]. This is possible thanks to the large output impedance of

the OTA's. The two OTA's are operated alternatively one in an open-loop configuration and the other in a closed-loop configuration. C_{l1} and C_{l2} represent the load capacitances formed by the two comparator banks performing the coarse and fine comparisons. V_{ref} is a reference voltage, which is assumed as the reference level for all voltages. V_{SS} is the negative supply voltage.

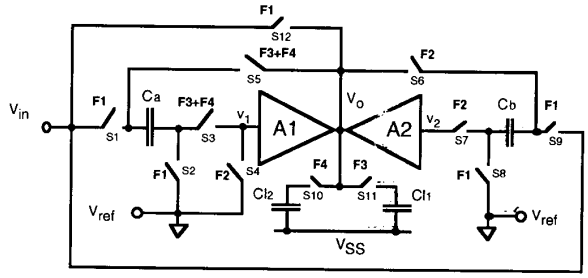

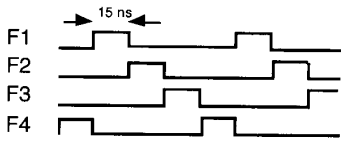

Fig. 13 Basic sample-and-hold architecture.
F1 = sampling; F2 = prediction;
F3 = hold for MSB's; F4 = hold for LSB's.

During phase *F1* (sampling phase), the input signal voltage V_{in} is sampled onto capacitors C_a and C_b. In order to avoid any slewing problem during the next phase, the output terminal is connected to the input through S12.

During phase *F2* (prediction phase), the sampling capacitor C_b is connected in feedback around the gain stage A2, while the gain stage A1 is left in an open-loop configuration, with its input grounded. Therefore, the output node is driven to the value of the input sample, $v_{in}(F1)$. The error in the output voltage due to the amplifier finite gain is that of a non inverting buffer, and is equal to $-\frac{1}{A_{02}} v_{in}(F1)$, where A_{02} is the DC gain of the stage A2.

During phase *F3* (hold phase), the gain stage A2 is connected in an open-loop configuration and, thanks to its parasitic input capacitance, holds the prediction of the input sample. The sampling capacitor C_a is connected in feedback around A1. In this way, the output node of the structure is again driven to the value $v_{in}(F1)$. Thanks to the prediction performed during the previous phase, the variation in the output voltage during phase *F3* is very small ($-\frac{1}{A_{02}} v_{in}(F1)$) and therefore the magnitude of the error E due to the finite gain of the stage A1, A_{01}, is also very small:

$$E = \frac{1}{A_{01}}\left(\frac{1}{A_{02}} v_{in}(F1)\right) = \frac{1}{A_{01} A_{02}} v_{in}(F1) \qquad (6)$$

As shown in equation (6), the error in the output voltage is equal to that of a unity-gain buffer using an operational amplifier with a gain equal to $A_{01}A_{02}$.

Of course, the same error exits during phase *F4*, which performs the fine comparison.

When the input offset voltages of the two stages, v_{os1} and v_{os2}, are taken into account, it is easy to see that the output offset voltage during *F2*, $v_{o,os}(F2)$ is given by

$$v_{o,os}(F2) = v_{os1} \frac{A_{01}}{1 + A_{02}} + v_{o2} \frac{A_{02}}{1 + A_{02}} \qquad (7)$$

When the loop is closed around A1 (phases *F3* and *F4*), the resulting output offset becomes:

$$v_{o,os}(F3) = v_{o,os}(F4) = v_{o,os}(F2) \frac{1}{1 + A_{01}} = -\frac{v_{os1}}{A_{02}} + \frac{v_{os2}}{A_{01}} \qquad (8)$$

Therefore, the proposed architecture also provides for offset compensation.

The architecture exhibits low sensitivity to amplifier finite gain and compensates for offset voltages. Therefore, simple stages with good high-frequency performance can be used while keeping a high static accuracy.

A S&H circuit was designed using a 1.2-μm double-polysilicon n-well CMOS technology. Stages A1 and A2 are implemented by means of single-input cascode amplifiers with a common bias current source. The simulated performance is shown below.

First stage DC gain	44	dB
Second stage DC gain	30	dB
Settling time (99.9%, worst case C_{load})	<14	ns

The subtract-and-hold circuit is derived from the basic diagram in fig. 13. During phase F1, sampling of the analog input signal is performed. An extra phase, F2, was introduced in order to synchronize the SU&H and the two-step flash ADC.

Two-step flash ADC comparator

The design of the comparator for the two-step flash ADC is critical as it must have a very high sensitivity. In order to minimize clock feed-through effects, a DC coupled comparator was developed. Its architecture is shown in fig. 14. Block $A1_C$ is a high-speed amplifier which drives a latch. Block $A2_C$ provides the auto-zero function for offset compensation. The output latch provides a rail-to-rail output.

The operation of the comparator can be divided in three phases, i.e. a comparison phase, a reset phase and an autozero phase (the names of these phases refer to the operation of the amplifying stage).

During $\phi 1$ (comparison phase), the input signal is applied to the input amplifier and the output value is stored on the input capacitors of the latch. The latch is set in a precharge condition.

During $\phi 2$ (reset phase), the amplifier and the latch are disconnected from each other. The latch is driven into one of

its stable conditions, depending on the value stored at its inputs, and the amplifier ($A1_C$) is reset. In this phase the input terminals of the amplifier, as well as its output terminals, are short-circuited in order to recover from a possible overdrive condition.

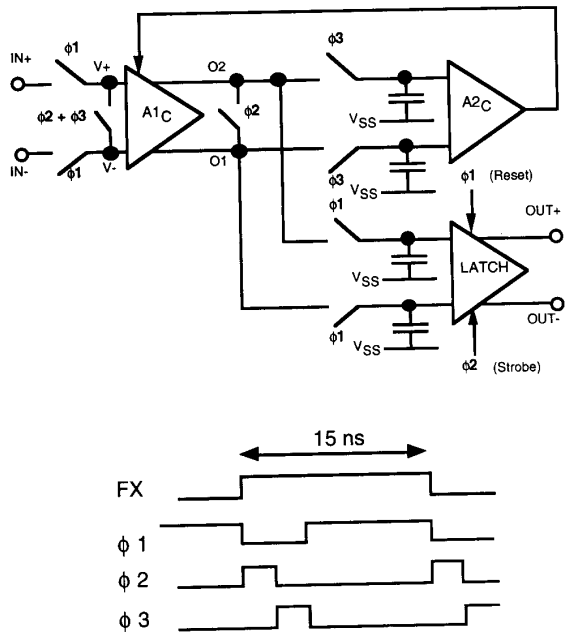

Fig. 14 Architecture of the two-step flash ADC comparator. $\phi 1$ = comparison; $\phi 2$ = reset; $\phi 3$ = autozero.

During $\phi 3$ (autozero phase), only the amplifier inputs are short-circuited and the output voltage is sensed in order to perform the offset compensation.

Output data are valid during phases $\phi 2$ and $\phi 3$.

The amplification stage $A1_C$ is implemented with a fully differential structure, in order to provide high speed and sensitivity with a large input common-mode range. The common-mode feedback required is realized with continuous-time techniques. To obtain high operation speed, a high-transconductance input differential stage is used; moreover, the output capacitance is minimized by picking up the differential output signal by means of two source followers, which drive the output latch. To ensure a very fast recovery from overdrive conditions during the reset operation, the maximum output swing of the stage during the comparison phase is limited by placing two diode-connected MOS transistors between the high-impedance differential output nodes. Offset compensation is performed by means of a negative-feedback loop. During phase $\phi 3$, the input terminals are short-circuited and the output offset voltage is sensed. The measured offset is then fed back to the amplifier after low-pass filtering (obtained with switched-capacitor techniques) and amplification (obtained with the differential stage $A2_C$).

A differential latch is used as the final stage. A very fast reset is provided during phase $\phi 1$ in order to avoid hysteresis. To achieve the maximum operating speed, minimum channel length transistors are used.

The comparator was designed using a 1.2-μm double-polysilicon n-well CMOS technology. The simulated typical performance is shown below.

Sensitivity	1	mV
Comparison phase length	<5	ns
Power dissipation (V_{DD} = 5 V)	6.5	mW

The comparator was integrated in a test chip. Preliminary measurements showed a sensitivity of 4 mV. However, a few phases were generated externally, therefore a better performance is expected for a fully-integrated circuit. The measured propagation delay in the comparison phase was 5 ns with an input voltage of 40 mV.

A new version of the comparator is being designed, in order to reduce silicon area and power consumption.

The parallel Successive Approximation ADC

The principle of the ADC is shown in fig. 15a. It consists of a number of identical successive approximation ADC (SA-ADC) channels, a common reference voltage generator and an output register [11]. The number of channels is given by (k+n), where k is the number of clock cycles needed for comparator auto-zeroing and n is the number of bits. Each channel contains a S/H circuit and a successive approximation ADC. The individual SA-ADC, as shown in fig. 15b, consists an auto-zeroed comparator, a shift register and a digital-to-analog converter (DAC). The output register collects digitized data from each channel and pipelines them to the output. Fig. 15c shows the timing diagram of a complete PSA-ADC, where each channel is successively skewed by one clock cycle. Since each channel needs (k+n) clock cycles to execute both the auto-zero and the binary search procedure and (k+n) channels are used, the PSA-ADC refreshes digitized data at each clock cycle.

The parallel architecture allows high conversion speed by using low speed SA-ADCs. Since each channel is so efficient, the cost of parallelism is low.

Because of the high efficiency a 10-bit PSA-ADC needs only 12 comparators, compared to 62 comparators for a half-flash ADC or 1023 comparators for a full-flash ADC. Therefore, the power consumption is substantially reduced. The high efficiency also results in small area. The core circuit of the first prototype chip occupies an area of 3.5 mm^2 and dissipates about 175 mW.

In a parallel architecture the threshold spread between the individual sub-channels must be kept small. This is realized by using a common reference voltage generator and auto-zeroed comparators. Also, the offset of the analog multiplexing must be kept low. This is simply solved by using passive S/H circuits in each channel, which is possible because of the very low input capacitance of the comparator.

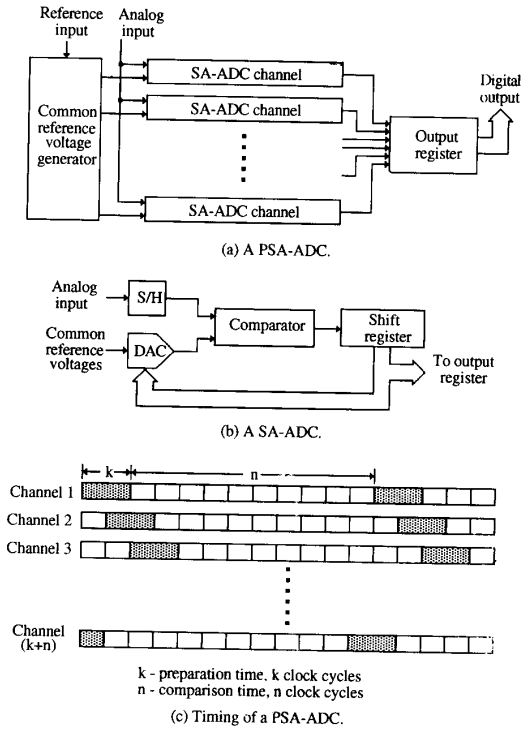

Fig. 15 Principles and timing diagram of a PSA-ADC.

Finally, it is not realistic from either the complexity or the load capacitance point of view to switch 1024 reference voltages to 12 channels. Instead, a two-step comparison method has been implemented, which reduces substantially the number of switches.

The block diagram of a comparator is shown in fig. 16. It combines the advantages of differential and single-ended amplifiers and a combination of AC and DC coupling [12].

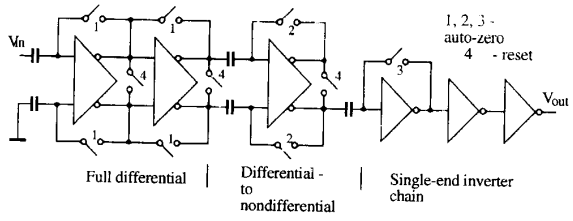

Fig. 16 The 6-stage comparator.

Two directly cascaded fully differential amplifiers are used as the first stages. The third stage transforms from full-differential to single-ended. Stages 4 to 6 are simple inverters. The latch function is performed at stage 6 by adding one more transistor. Capacitive couplings is used where auto-zero operations are performed. A 10-bit resolution needs a total gain of 2500 (1mV to 2.5V). Since 6 stages are used, an average gain of 4 is sufficient for each stage, which is advantageous for both a fast auto-zero and a fast response.

The value of the coupling capacitors is critical. Note that the finite time constant determined by the leakage current of the switch transistor is not a real problem for a fast comparator. Instead, the low limit is set either by the allowable voltage division between the coupling capacitance and the input capacitance of the next stage or by the allowable clock and charge feed-through. In a standard CMOS process, unfortunately, the parasitic capacitance of the coupling capacitor to the substrate is about half of its value. Therefore, the high limit is set by either the time constant $t_0=C_L/g_m$ or the allowable auto-zero time. In order to gain speed and reduce power consumption, the coupling capacitance has been pushed to its lower limit. The chosen value, 0.1pF, is just twice the input capacitance of the next stage. The gain loss of 33% is acceptable, since each stage needs only a gain of 4 and both the response time and the auto-zero time are greatly reduced. The problem caused by clock and charge feed-through will be solved by dummy switches, the differential structure and a three-step auto-zero procedure.

The success of a multi-step auto-zero relies on each stage being kept in its linear region even under the worst residual offset propagation. This demands a large linear region and a large common mode range. The current design meets these demands. The accuracy of the comparator depends mostly on the front stages. The two direct coupled full differential amplifiers give a total gain of more than 25 by introducing positive feedback transistors.

The first test version of the comparator was designed and fabricated in 1991. The statistics of the comparator offset for 62 comparators from 8 chips is summarized in fig. 17. The response time for an input of 1mV is measured to be less than 6ns, including the extra pad driver. The power consumption is less than 10mW.

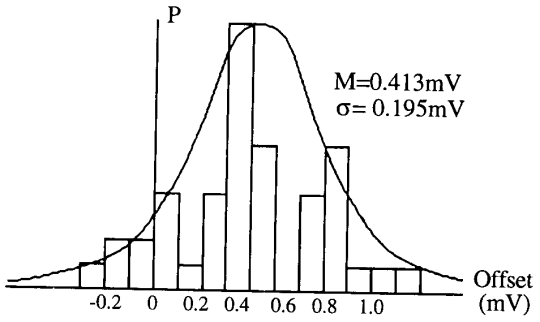

Fig. 17 Measured offset distribution.

The specifications of this comparator are given below.

Gain	> 70 dB
Offset	< 1/2 LSB (1mV)
Worst response time	< 10 ns
Power consumption	< 10 mW
Area (with all sub-circuits)	0.06 mm^2
Technology	1.2 μm standard CMOS

In order to reduce the complexity of the reference switch network, the coarse and fine references are fed into the comparator through two separate capacitors, as shown in fig. 18.

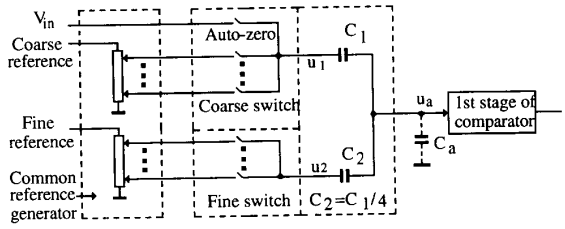

Fig. 18 Dual-capacitor input circuit.

The summation of the two references is done by using the principle of charge-sharing. The charge stored on node "a" during the auto-zero phase is

$$q = u_a + (u_a - u_1)C_1 + (u_a - u_2)C_2 = u_a(C_a + C_1 + C_2) - u_1 C_1 - u_2 C_2 \quad (9)$$

where u_a, u_1 and u_2 are potentials of nodes "a", "1" and "2". C_a is the total input capacitance at node "a". During the auto-zero phase $u_1 = u_{in}$ and $u_2 = 0$. According to the principle of charge conservation, after auto-zero, any new u_1 and u_2 will result in a new u_a and the difference will be

$$\Delta u_a = \frac{(\Delta u_1 C_1 + \Delta u_2 C_2)}{(C_a + C_1 + C_2)} \quad (10)$$

Note that both the analog input and the coarse references are fed to node "1" so there will be no error. However, the fine reference is fed to node "2" and there will be an error if $C_1/C_2 \neq 4$. The relative error is $\Delta = (1 - (C_1/C_2)/4)$. For a two-step ADC, Δ should be less than $1/2^{(1+n/2)}$, i. e. less than 1.5% for n=10. In a CMOS process, the matching of capacitor values can be made better than 0.1% - 0.2% for ratios in the range 1 to 16, well within the required 1.5% limit. In order to increase the effective input voltage Δu_a, the value of C_2 was reduced to $C_1/4$ and the fine reference voltage was increased by a factor 4. This differs from the configuration in [13].

The S/H circuit

In order to avoid the offset differences introduced by the S/H stages in a multi-channel architecture, the buffer normally used in a S/H circuit has been removed. This purely passive S/H circuit which is combined with the input stage of the comparator in our PSA-ADC is shown in Fig. 19.

The sampling time-constant is controlled by the switch transistor S1 and the sampling capacitor C_S. A large C_S is also necessary for reducing the residual clock and charge feed-through from S_1 while the main part of the feed-through will be absorbed by the dummy transistor. During the sampling phase, C_1 is charged by both the input voltage and the auto-zero voltage of the first stage. Switches S_1 and S_3 will be turned off simultaneously to finish the sampling and the first-step auto-zero. The residual clock and charge feed-through will be cancelled by auto-zero steps 2 and 3.

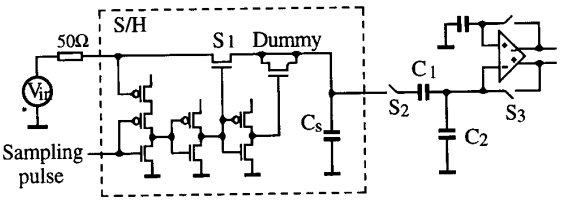

Fig. 19 Clock-slope compensated passive S/H.

The delay of the clock buffer is controlled by injecting a part of the input analog signal into the switch control circuit. This compensates for the sampling point error caused by the finite clock slope, which is about 0.3 ns. According to simulation, an 80 ps time error will occur without compensation, which will cause a 13 mV (~ 3-bit) sampling error for a 50 MHz input signal. The error is reduced to only 4 ps after compensation, corresponding to a 0.6 mV error (< 1/2 LSB). The detailed optimization is being studied with a dedicated multichannel circuit under test.

The complete PSA-ADC

The first test version of the complete PSA-ADC was designed and sent for fabrication early 1992. The layout was done in AMS 1.2μm single-poly and double-metal CMOS technology.

The complete layout was simulated by using the switch level simulator TMODS. In the simulation, except for the analog part which was replaced by equivalent digital logic, all elements were simulated, including the pads. The results indicate a performance conforming to specifications.

The measured performance of the first prototype of the complete ADC are given below.

Resolution	10 bits @ f_s = 50 MHz 9 bits @ f_s = 70 MHz
Diff. non lin.	≤ ± 0.65 LSB
Int. non lin.	≤ ± 0.65 LSB
Number of comparators	12
Area	1.8 x 1.9 mm^2 (core)
Power consumption	
core	175 mW
output drivers	40 mW
Technology	1.2 μm CMOS

A second prototype is now under test, still using 1.2 μm, with the hope of achieving 10 bits resolution at sampling frequencies >70 MHz. Further improvements are possible by scaling the circuit to a finer technology.

DIGITAL SECTION

The digital part of FERMI is distributed over two ASIC types: the service and the channel ASICs. The design of the channel part has now reached a state where the implementation on silicon can begin. This is not yet the case for the service part, since some questions concerning the FERMI interface to the first and second level trigger processes still remain open.

The digital electronics on these two ASICs can be divided into the following functional blocks; On the service ASIC the major units are the controller, the trigger block and the output block. The channel ASIC contains the look-up table and the memory.

In the following we will shortly describe the different circuits in a functional order.

The Look-up table

The look-up table, shown in fig. 20, contains the inverse of the composite transfer function of the preceding, complete, analog chain. It corrects for production tolerances, e.g. differences in the compressor function, the gains of the preceding stages and all other sources of non-linearities and systematic effects. The look-up table will convert the compressed 10 bit data from the A/D converter to a 16 bit linear representation.

The look-up-table function is equipped with the necessary registers and counters to allow easy loading of the data.

As an additional precaution, a multiplexer has been introduced to provide an emergency bypass path around the look-up table. This facility can be used when there is a look-up table failure and the spare channels have been fully used up.

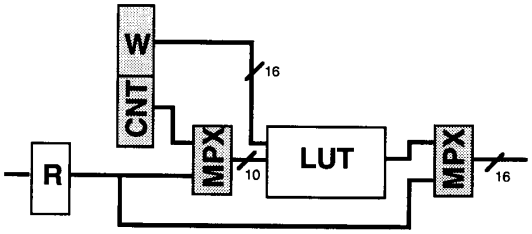

Fig. 20 Look-up table

The counter can also, as is described in the controller section , be used to store a digital test pattern in the look-up table. In this mode the clock input of the counter is connected to the bunch crossing clock, causing the stored information to be outputted to the pipeline and the first level trigger. This mode provides a way to generate digital test patterns for the first and second level triggers.

Event detection and Trigger summation

In the trigger summation the data from all active channels are summed in 2 steps, where the first occurs on the channel IC. This energy sum is applied to a pulse detection circuit (fig. 21), which generates a flag relating each detected pulse to a bunch crossing. Pile-up flags are generated when two such flags are too closely spaced in time. Two different pile-up flags are envisaged to signal severe and mild pile-up conditions.

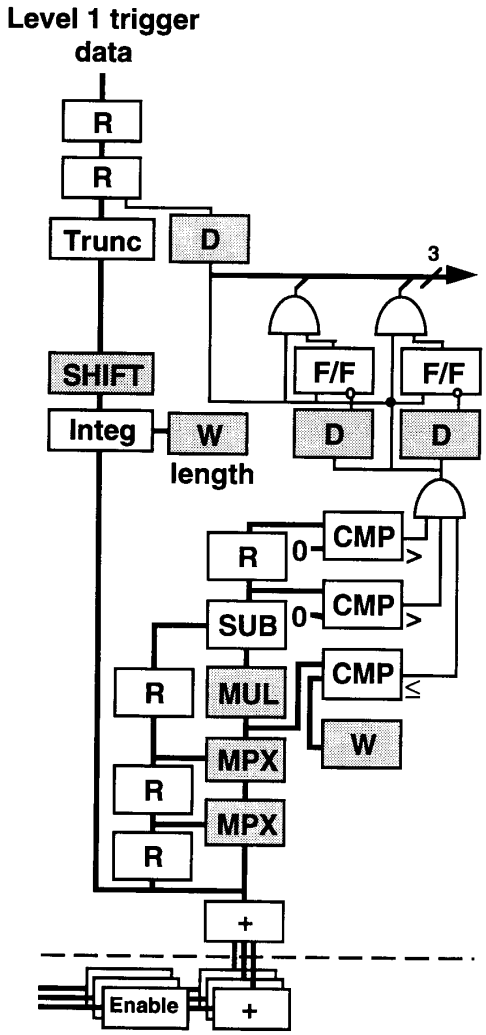

Fig. 21 The digital constant fraction discriminator.

The integration is implemented as a simple digital filter with coefficients 1 or 0. The result of this operation is truncated to a 12-bit word to minimize the number of output pins. A shift operation scales the result to the proper size to allow a more precise data representation. An alternative is to replace the integrator with a more general digital filter (see below).

The output data to the first level trigger is double buffered to adapt to the global clock. This is necessary since there is a phase difference between the internal and the global clock due to a programmable clock delay.

The pulse detection circuit emulates a constant fraction discriminator with pulse amplitude veto. By manipulating the

multiplexers (different delays) and the programmable shift function (different amplifications) it is possible to adjust the discriminator for different amplitudes and rise times. The minimum pulse veto level is a programmed value chosen to optimize the noise rejection. The module level pulse_detect flag is combined with its associated data and transferred to the first level trigger process.

Data memory

The memory block, see fig. 22, contains an option for inserting diagnostic data at its input and a spy register at its output, in order to allow thorough testing of the memory function. The use of error correction codes (ECC) in the memory will provide an on-line error correction and data integrity control by setting error bits in the status register when severe errors are detected. The ECC algorithm is chosen as to correct single errors and to detect double errors. All surrounding control and error correction logic, including the addressing circuitry, are two-fold redundant. The address input is protected with a similar ECC arrangement (i.e. single-error correction and double-error detection).

Additional fault tolerance is provided by a small 10-cell 2-port associative memory which intercepts the main memory access when a write attempt into a faulty memory location is detected.

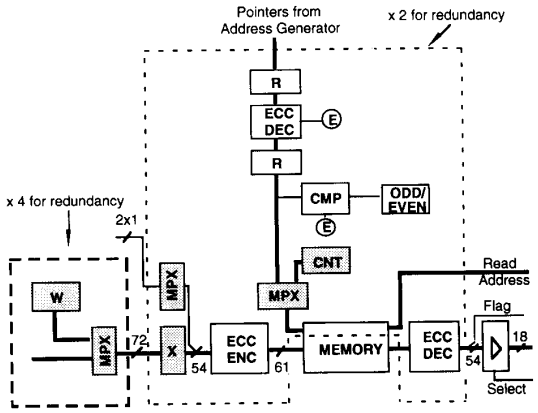

Fig. 22 The Data memory

The memory function is realized by two toggling single-port memory banks. It is therefore required that the addresses alternate odd and even. By comparing the state of the odd/even flip-flop, which toggles at each clock cycle, with the least significant bit of the address, the function of the external addressing circuitry is monitored. A mismatch will cause an error report.

Output stages

This part, see fig. 23, contains the frame pointer buffer memory (FPB), the read-out controller and the digital filter.

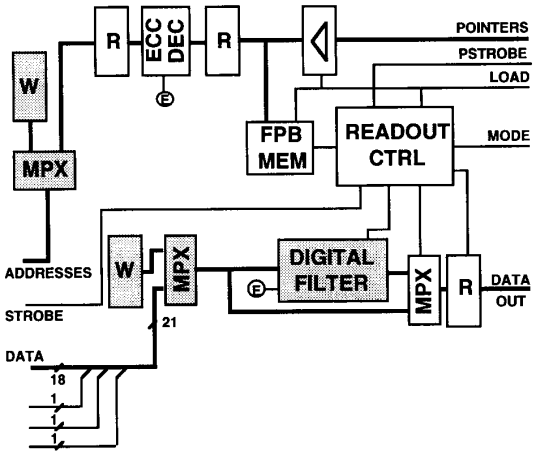

Fig. 23 Output stages

The FPB is loaded with a set of pointers corresponding to the requested time frame. The mode information determines whether a full or a filtered read-out is requested, and in the latter case specify which of two available coefficient banks should be used.

The read-out process starts immediately after loading of the pointers and provides control signals required for producing and storing one datum into the read-out register.

The process is resumed when the datum is read, by the external controller, from the register.

All pointers are subject to an error correction process correcting and reporting all one-bit errors and reporting two-bit errors.

Diagnostic registers and multiplexers are provided to check the pointer stream and the output data with normal or preset data.

Digital filters

In the processing chain of data samples, digital filtering allows to extract compacted information from the time frame associated with each event. Two different implementations, one for the first level trigger, and another for the second/third level are foreseen.

For the latter the filtering hardware is implemented as the inner product of weight and sample vectors. The samples are supplied by the data memory and the weight vectors are read out in skewed form from a dedicated memory, in synchronism with the samples [14] together with their respective modulo-3 residues.

The inner-product unit is composed of a pipelined multiplier generating the products in skew form and of a skew-input accumulator, as shown in fig. 24. In this representation the bits are transmitted individually but at different times, as in the serial form representation (LSB first). The carry informations propagate both in space, as in parallel adders, and in time under clock control, as in serial adders. This skew architecture allows operations to be performed at high speed.

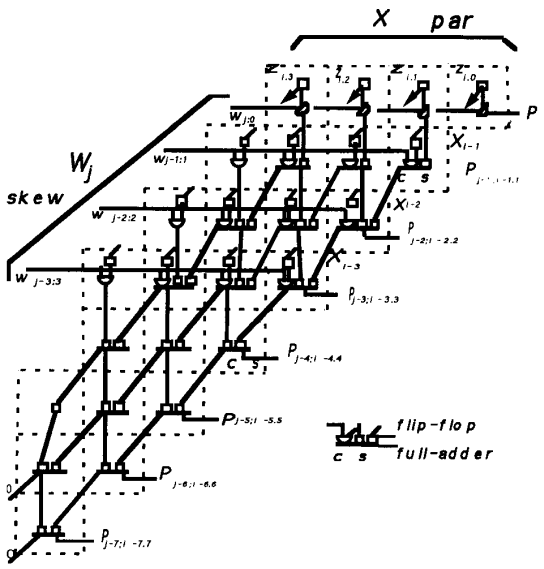

Fig. 24 The inner-product unit

The associated latency (in clock periods) is:

$$L = n_{bits} + \log_2 N_{samples} \qquad (11)$$

An efficient resolution of double pulses depends on the number of convolution terms and on a suitable set of coefficient weights, whose values will be determined in relation to the specific pulse shape of any given calorimeter. These coefficients could be changed at will, even during operation, in order to optimize filter performance under changing conditions.

The study of digital signal processing algorithms for pulse finding and feature extraction is concentrated on efficient operations and new filter design techniques to obtain optimal solutions in the presence of noise and timing jitter effects. The filters considered include also non-linear [15] and linear-non-linear hybrid structures as well as linear filters based on signal interpolation.

As an alternative to the filter for the first level trigger described above, an adder/convolver is under study. Since here the latency is of greater importance than resolution, an upper limit of ten clock cycles is considered, along with a truncation of the data to twelve bits.

The FERMI controller

The service part contains a controller function to initiate and control all programmable FERMI features. The controller will also monitor the FERMI functions and the serial communication links that connect it to the outside world. A back-up communication channel, capable of low-level commands, will also be implemented. This link will provide the same functionality as the internal controller if supported by an external controller. The FERMI controller is also able to generate controlled calibration pulses which are injected into the analog inputs. This will allow the calibration of the look-up tables and other registers. By selecting a special mode the controller can load a digital test pattern into the look-up table memory. This pattern can then be extracted in a system-wide test mode, in order to test the system response to arbitrary pulse distributions.

The clock server function allows the controller to enable or disable the internal clock. It may also generate clock bursts of well defined lengths. This function will also provide means to adjust the internal clock phase to a fraction of a ns.

VHDL simulations

The digital section of FERMI is described in the hardware description language VHDL. This description, apart from providing a compact, well defined and easily interpreted specification, allows to verify the functionality and consistency of the design up to and including the layout phase. It also helps to produce test inputs to the detailed submodule designs so that their conformity to specifications can be verified.

The current description is continuously being updated to include the latest modifications.

The FERMI system environment

The FERMI modules rely on external circuitry to provide control sequences and to read out data. If a number of FERMIs are mounted on a board, this circuitry can be of considerable complexity [16]. It might have to provide the following functions:

-Local first level trigger : to combine data from several FERMIs into one trigger-cell value, which is transmitted to the global first level trigger.

-Local second level trigger : to read data from the FERMIs and to transmit them to the proper second level trigger processor.

-Address generator : to generate new addresses each 15 ns, pointing at free storage locations for the sampled input data, and to generate the pointers required for readout.

-Controller : to initialize, control and supervise the FERMIs.

These functions must be extremely fault tolerant, since any failure will seriously reduce the functionality of the entire system. Such a fault tolerance is very difficult to achieve with standard random logic. There are different solutions solving parts of the problem.

-Moving as much as possible of the functionality either out of the detector or into the FERMI. This will relax the fault tolerance requirement. A good fault detection strategy will identify and localize faults, so that the erroneous part may be replaced.

-Moving functions into the FERMI will also relax the fault tolerance requirement, since the consequences of the faults are only local to a given FERMI. Furthermore, the necessary fault tolerance is also easier to realize in the ASICs of the microsystem.

-A third possibility is to put as much as possible of the remaining circuitry into a dedicated ASIC, a FERMI board controller.

Another important issue is to solve the communication problems. Coaxial cables take up considerable space and limit transmission speed. Optical fibres, however, are better in both respects. The coupling between the optical receivers and transmitters and the ASICs is easiest realized in a microsystem. A reasonable solution is therefore to collect all the functions that cannot be relocated to the outside of the detector into one fault-tolerant microsystem, which will communicate electrically with the FERMIs and optically with the system controller resp. first and second level triggers. One of the relocated functions is the address generator, which in this design is global. This has the advantage of eliminating the loss of synchronism between local address generators as a possible source of errors.

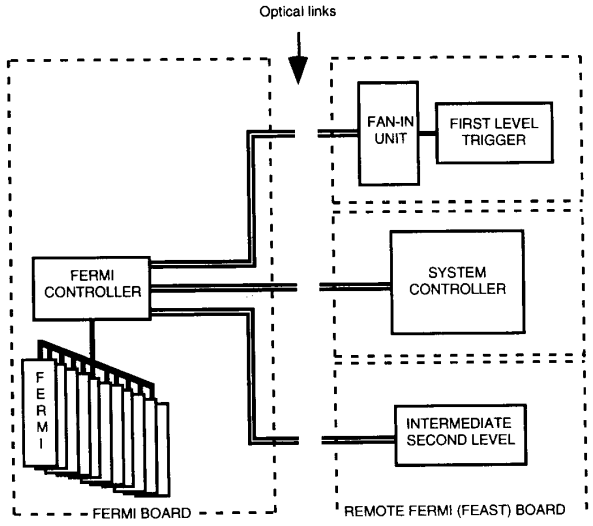

Fig. 25 The FERMI system environment

A 1 Gb/s two-way optical link is sufficient for communication with the second level trigger for a board containing about 36 FERMIs. This connection can also provide a communication channel for control and status reporting. Two similar fibres can transmit addresses from the address generator, the 67 MHz bunch crossing clock, as well as test and global command pulses. The communication with the first level trigger could probably be handled by nine such fibres, unless zero suppression is used to reduce the required bandwidth. Fig. 25 shows a possible implementation of this scheme.

FAULT TOLERANCE

General fault tolerance policies

The definition of fault-tolerance policies relies on the choice of fault models. In the absence of detailed statistical information on failure mechanisms, it was decided to adopt models based on effects of ionizing particles as measured in the context of space exploration. This involves the existence of both transient and permanent failures, which implies the development of (self) diagnostic mechanisms capable of distinguishing between the two failure types and the development of corresponding reconfiguration policies.

Given the complexity of the overall FERMI structure, it is reasonable to adopt a hierarchical approach to fault tolerance, involving both distributed (local) actions and centralized policies. The actual decision on such a distribution will be reached only when the overall architecture will be completely finalized. A preliminary description of these mechanisms is outlined below.

With reference to the filters, distributed self-test mechanisms based on concurrent techniques are foreseen, so as to minimize error latency. Suitable coding techniques, e.g. modulo-3 arithmetic, are being studied, allowing to grant relevant fault coverage while requiring a modest increase of silicon area.

The granularity of self-testing has to be balanced against the corresponding overhead; different solutions are being studied, associated with the various design alternatives. Thus, should the skew filter design be chosen, self-testing will be provided at the level of the individual bit convolver, this being also the elementary subsystem envisioned for reconfiguration. On the other hand, if a totally parallel approach is chosen, self-testing will be related to the individual multiplier, which will also be the basic block for reconfiguration.

Reconfiguration after fault could be implemented either by local self-controlled solutions or by host-driven techniques involving boards or subsystems at higher hierarchical level. Both policies are being explored, the final choice depending not only on global architecture evaluations, but on system considerations as well. A policy to be decided upon concerns the possibility of accepting (or, alternatively, discarding) a number of "faulty" results, i.e. allowing a limited error propagation through a fault-affected system prior to reconfiguring it.

Considering the filter functions discussed above, on-line self-driven reconfiguration can be effective for all alternatives,

with a reasonably low silicon overhead, under the assumption of hard (permanent) faults.

If the various FERMIs will be controlled by remote controllers interconnected via optical fibres, reconfiguration can be remotely controlled, and the local mechanisms distinguishing between permanent and transient faults need be quite simple. On the other hand, local autonomous reconfiguration can be designed for a minimum propagation of errors before reconfiguration takes place.

A more complex analysis is required for the problem of fault-tolerance in the memory section. While for memories simple coding techniques can be adopted (simulations will be required to determine the optimum balance between redundancy due to the possible coding solution and the resulting robustness), the most delicate point is the memory management section. The address generator currently considered, while very efficient in other respects, is quite delicate with respect to faults, and massively redundant solutions do not appear very attractive. An alternative design approach based on the use of associative memories, currently under consideration, would provide robustness with respect to faults. Its design will be developed in detail (and possible implementations will be considered) so as to verify the actual viability of such an alternative.

Fault tolerance implementation

The fault tolerance in FERMI is implemented in different ways depending on the logical function of each stage. It can be automatic, as in memories with error correction code (ECC), or it may require controller intervention. In the latter case, on-line error checking will alarm the controller, which will respond with a hardware reconfiguration, thus eliminating the source of error. In functions with a high level of parallelism, it is usually straightforward to add spare parts. Another method is to use voting procedures on parallel data paths.

Each channel IC, containing four independent channels, will be equipped with a crossover switch dividing each channel into parts. Within these parts it will be possible to disable one of the four channels and reconfigure the remaining ones. The principle is shown in fig. 26.

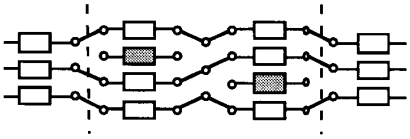

Fig. 26 Reconfiguration of channel IC

In arithmetic parts, such as the circuits for first level summation and digital filtering, results can be checked by applying modulo-3 checking.

STATUS OF THE PROJECT

The non-linear compressor amplifier prototype is designed, implemented and currently under test. Two prototype versions of an A/D converter have been designed and implemented, one of the versions has been tested both at the sub-part level as well as a complete prototype while the other has been partly tested and a complete prototype is currently being processed.

The digital part is modelled using the VHDL language and is being verified using XILINX programmable arrays as test media.

Test and calibration procedures are being worked out, partly as results of the VHDL modelling, partly as results of the continuously ongoing system-wide simulations.

A floor-plan for a complete microsystem is under design.

CONCLUSIONS/SUMMARY

In this paper we have demonstrated the feasibility of producing calorimeter readout electronics, based on early digitization, for the next generation of hadron colliders. The digital representation of the data allows the generation of very sharp thresholds for the trigger processes as well as unlimited latency for the decision processes, flexibility and a very large degree of programmability of the entire system. The problems of handling signals with very high dynamic range has been solved by the use of a non-linear transfer function matched to the resolution of the particular detector. A/D converters fulfilling the requirements, specially what concerns speed, size and power consumption are being developed. Digital circuits providing the necessary storage, processing and communication functions are described. The questions of fault tolerance and reconfigurability have been addressed at each stage of the design. A Silicon-on-Silicon microsystem packaging technology is used to achieve a compact, yet flexible, readout system.

REFERENCES

[1] V.G. Goggi and B. Lofstedt: "Digital Front-end Electronics for Calorimetry at LHC", *Proc. ECFA Large Hadron Collider Workshop, CERN 90-10, ECFA 90-133, vol. 3, pp. 190-200, December 1990.*

[2] The FERMI collaboration: "Status Report on the FERMI Project", *CERN/DRDC92-26, May 1992.*

[3] For a review see: R. Wigmans *Ann. Rev. of Nucl. and Part. Sci. vol 41 (1991) p. 33*

[4] A. Novak *et al.*, "A Standard Microsystem Module", *Presented at the ISHIM Meeting, Nordic Chapter, Åland, Finland*, September 1990.

[5] G. Appelquist and C. Bohm: "A memory controller for FERMI", *FERMI note #5, November 1991.*

[6] G. Appelquist, C. Bohm and B. Lofstedt: "A system specification of FERMI", *FERMI note #12, November 1992.*

[7] R. Chase, private communication.

[8] A. G. F. Dingwall, V. Zazzu: "An 8-MHz CMOS subranging 8-bit A/D converter", *IEEE J. Solid-State Circuits*, vol. SC-20, pp. 1138-1143, Dec. 1985.

[9] F. Maloberti, G. Palmisano, G. Torelli: "A novel approach for high-frequency gain-compensated sample-and-hold circuits", *Intern. Symposium on Circuits and Systems*, San Diego, California, May 1992.

[10] A. Baschirotto, R. Castello, F. Montecchi: "Finite gain compensation techniques for high-Q bandpass SC filters", *1990 IEEE Int. Symp. on Circuits and Systems*, New Orleans, Louisiana, May 1990, pp. 2813-2816.

[11] C. Svensson and J. Yuan: "Report on PSA ADC", *FERMI note #4, June 1991*

[12] C. Svensson, K. Chen and J. Yuan, "Decision making in fast A/D converters", *report LiTH-IFM-IS-154, 1989, unpublished.*

[13] A. Cremonesi, F. Maloberti, G. Torelli, and C. Vacchi, "An 8-bit two-step flash A/D converter for video applications," *IEEE 1989 Custom Integrated Circuit Conference*, pp. 6.3.1 - 6.3.4, 1989.

[14] L. Dadda: "A fast pipelined inner-product unit", *FERMI note #7, April 1992.*

[15] S. Inkinen and Y. Neuvo, "Base Line Normalization of High Energy Physics Detector Signals Using Sparse Median Operations", *Accepted to the IEEE Workshop on Nonlinear Digital Signal Processing*, Tampere, Finland, January 1993.

[16] G. Appelquist and W. Iwanski: "Simulations of a FERMI board using VHDL", *FERMI note #6, November 1991 (also EAST note 91-18).*