# FERMI - a new generation of electronic modules for large data acquisition arrays required by high energy physics

A. Dell'Acqua[3], C. Alippi[6], G. Appelquist[8], S. Berglund[8], C. Bohm[4], L. Breveglieri[6], S. Brigati[5], L. Del Buono[9], P. Carlson[10], P. Cattaneo[7], L. Dadda[6], J. David[9], M. Engström[8], G. Fumagalli[7], U. Gatti[5], J.F. Genat[9], G. Goggi[7], M. Hansen[3], H. Hentzell[2, 6], I. Höglund[1], S. Inkinen[3], A. Kerek[10], H. Lebbolo[9], O. LeDortz[9], B. Lofstedt[3], F. Maloberti[5], P. Nayman[9], S.T. Persson[12], V. Piuri[6], F. Salice[9], M. Sami[6], A. Savoy-Navarro[9], R. Stefanelli[6], R. Sundblad[11], C. Svensson[4], G. Torelli[5], J.P. Vanuxem[3], N. Yamdagni[8], J. Yuan[4], R. Zitoun[9]

[1]ABB-Hafo AB, Järfälla, Sweden, [2]Center for Industrial Microelectronics and Materials Technology, Univ. of Linköping, Sweden, [3]CERN, Geneva, Switzerland, [4]Dep. of Physics and Measurement Technology, Univ. of Linköping, Sweden, [5]Dip. di Elettronica dell' Univ. e Sezione INFN, Pavia, Italy, [6]Dip. di Elettronica, Politecnico di Milano, Italy, [7]Dip. di Fisica Nucleare e Teorica dell' Univ. Sezione INFN, Pavia, Italy, [8]Fysikum, Univ. of Stockholm, Sweden, [9]LHPNE, Univ. Paris VI-VII, France, [10]Manne Siegbahn Inst. of Physics, Stockholm, Sweden, [11]SiCon AB, Linköping, Sweden, [12]Swedish MicroSystems AB, Linköping, Sweden

Correspondance to:
C. Bohm
*Fysikum, University of Stockholm,*
*Box 6730, S-113 85 Stockholm, Sweden*

## Abstract

The Front End Readout MIcrosystem, FERMI, is a representative of a new generation of data acquisition modules which utilizes modern design techniques to achieve a high acquisition rate together with intelligent on-line data processing. FERMI is being designed to satisfy the extreme requirements set by calorimeters in the next generation of particle physics detectors. Such detectors are being designed for the future LHC and SSC accelerators at CERN in Switzerland and at the SSC-laboratory in Texas. The calorimeters demand frequent (67 MHz for LHC, 63.5 MHz for SSC) high precision sampling of a large number of input channels (about 5 $10^5$).

Each FERMI module serves 9 channels from which samples are AD-converted, corrected and temporarily stored in a local memory. The data is also merged into a trigger sum processed by digital filters to recover time of incidence and amplitude of incoming pulses. Such data is then fed to a first-level trigger processor which screens irrelevant information. Only data that may contain interesting information is kept for further analysis.

Arrays of 50000 FERMIs constitute formidable processing systems when considering the total computational power and storage capacity.

## Introduction

The new physics content in the data produced by next generation of high energy proton accelerators is deeply buried in large amounts of background information. Provided the particles have been accelerated to a sufficient energy, only about one part in $10^9$ recorded events contain useful information about unanswered questions in physics. In the accelerator rings, counter rotating proton beams intersect in collision regions around which detectors are built to record the debris from the proton collisions. This information is then used to reconstruct the collision itself. It is, however, crucial that the particle fluxes are sufficient to compensate the low yield.

Two such accelerators are presently being considered, LHC at CERN in Switzerland and SSC at the SSC laboratory in Texas. In these, protons will be accelerated in several steps



0-7803-1850-1/94 $3.00 ©1994 IEEE

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

until they reach energies of 8 or 20 TeV in rings of 27 or 87 km circumference, respectively. The protons travel in well defined bunches that collide every 15 ns. Each bunch crossing will cause about 40 inelastic proton collisions at the different collision regions of LHC. The particle flux at SSC is lower, causing only about 1 proton collision each bunch crossing.

The detectors are built in layers with different subdetectors sensitive to different aspects of the reaction products. The volume immediately surrounding the collision region often contain tracking detectors which record ionization paths from charged particles. These subdetectors may contain about 107 channels, which fortunately are rather sparsely activated. If they are situated in a magnetic field it will be possible to estimate momentum and charges from the curvature of the particle trajectories. Outside this there are calorimeters of two types, designed to absorb different types of particles. The first absorbs electrons and photons, while the second absorbs hadrons. Muons are detected in the outermost subdetector. The presence of non-interacting particles can be deduced by summing the momenta of all reaction products from a collision. If all particles are accounted for, the total momenta should vanish.

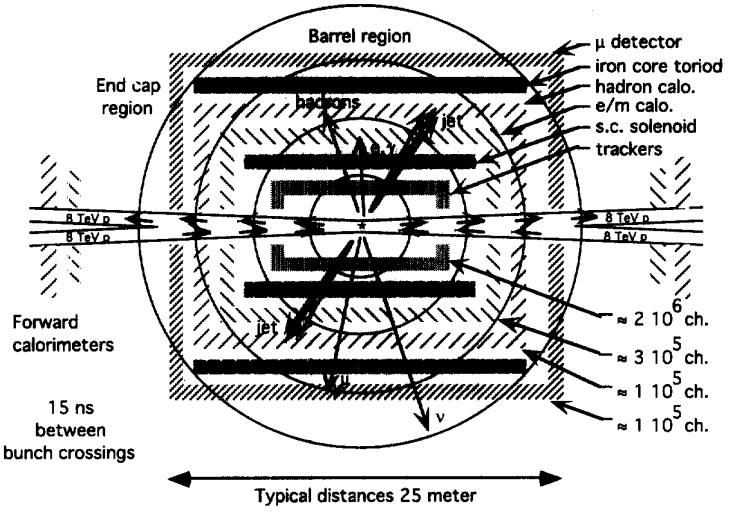


Figure 1. A typical LHC or SSC detector

Recording all information is presently not possible and in any case certainly not meaningful, even if it was possible. The total amount of data is in the order of hundreds of TBytes/sec. The strategy chosen is to filter events in steps, trigger levels, until sufficient data reduction is achieved. The first-level trigger makes a rough calculation based on low resolution data from a subset of the detector. This information is sufficient to reject most data, attaining a reduction rate of about $10^3$ or $10^4$ in this step. Data accepted by the first-level trigger is sent to the second-level trigger, which uses high resolution data within regions of interest from all sub-detectors as a base for its decision. A further reduction of about $10^2$ may be obtained in this way. In the final trigger step, the third-level trigger, all data is considered, at full resolution, to get a final reduction of about $10^2$. The total data reduction achieved in this way is thus in the order of $10^7$-$10^8$. The quality of the trigger algorithms are, however, crucial to the success. They must be as unbiased, efficient and selective as possible.

The values for the trigger processing times are currently estimated to about 2 μs and 2 ms for the first and second-level triggers respectively. This means that all data must be stored in temporary memories during ~2 μs waiting for the first-level trigger decision. The .1% that was not rejected by the first-level-trigger must be kept for ~2 ms.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

## The FERMI data acquisition module

The FERMI data acquisition module is designed to sample, digitize and temporarily store calorimeter data while providing reduced precision data for the first-level trigger. The storage should thus be sufficiently large to keep all data for ~2 μs and some data, i.e. the data approved by the first-level trigger, during ~2 ms. Data approved by the second-level trigger as well is kept in storage until it can be read out to the third-level trigger. The precision and dynamic range of FERMI should match that of the calorimeter used, so that it does not significantly deteriorate the data. It has been found that this can be achieved with a 10-bit AD-converter, provided a suitable analog compression has been applied prior to the conversion. This is done in a non-linear amplifier stage. After conversion, data is expanded to a linear 16-bit representation by applying the inverted non-linearity via a look-up table.

FERMI is implemented as a silicon-on-silicon multichip microsystem, where the microsystem is 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.

Figure 2 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 nine identical channels, while the latter is unique. To achieve a high level of reliability, FERMI is being designed with a high degree of fault tolerance, including implementation in rad-hard technologies. This is important since FERMI will probably be mounted on the calorimeter surface well inside the detector, where the ambient radiation level is fairly high.

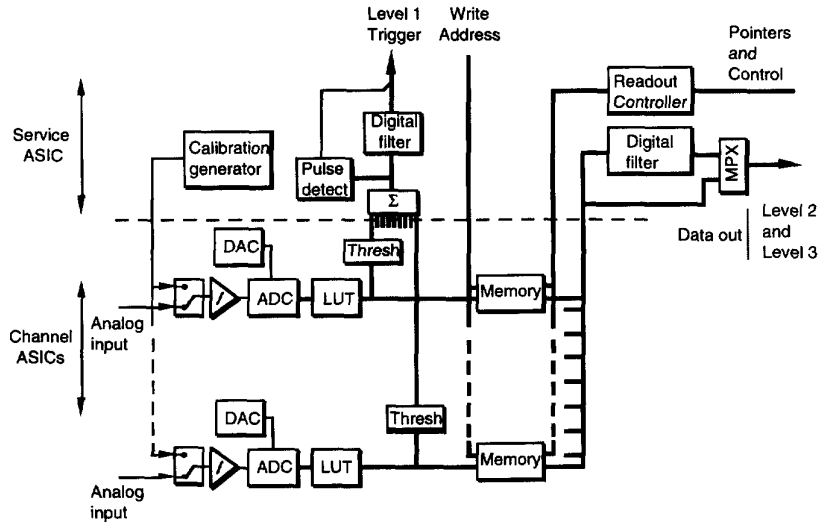


Figure 2. An overview of the FERMI data acquisition module

A FERMI prototype module is being developed within a 3 year research and development collaboration at CERN, RD-16, approved by the CERN DRDC committee in April 1991. The FERMI collaboration, now consisting of nine research institutes and three industrial partners, was thus formed with the aim to explore and develop the technologies required to implement FERMI as a multichannel data acquisition and signal processing module implemented as a silicon-on-silicon multichip microsystem [1], [2], [3].

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

The input signals are compressed and digitized every 15 ns. The resulting 10-bit data are linearized and expanded to the full dynamic range by means of the look-up table. This device can also generate an absolute calibration for each individual channel. The expanded data from all active channels are individually discriminated with independent programmable thresholds to eliminate excessive noise. Sufficiently large channels are then summed and processed in order to provide energy and bunch crossing time information to the first-level trigger processor. A module-level pulse detect signal can set a flag to identify possible pile-up conditions for special processing, i.e. when two pulses appear so close in time that they distort each other's pulse shapes.

The expanded data are also stored in a dual-port memory, anticipating the decisions from the first and second-level triggers. The storage occurs in memory positions given by an external memory management unit, the address generator[4]. A temporal environment, a time frame of programmable length, is associated with each event accepted by the first-level trigger. All memory locations containing sample points not included in a time frame after the first-level decision are returned to the pool of free memory. Time frame memory locations are also returned to the free pool after having been rejected by the second-level trigger or read out to the third-level.

The FERMI I/O facility contains provisions for full or reduced readout of time frames. In the latter case the information contained in the time frame is extracted by means of adaptive non-linear digital filtering techniques. It is envisaged that the reduced readout should be default when pushing data to the second-level trigger processors. These could, however, also return for selective readout of full time frames, which might be necessary in order to resolve specific conditions.

The FERMI microsystem is implemented as a thin-film Multi-Chip Module (MCM-D). This means that several ASIC dies are mounted on top of a thin-film substrate containing interconnections, integral resistors and capacitors. The actual process used for FERMI is based on a silicon substrate. Four metal layers are used, two for interconnections and two for power distribution.

For the FERMI system a number of package options will be available. Results from irradiation tests show that for the final solution a plastic package is preferable, since it contains only small amounts a substances that are susceptible to neutron activation. For the prototypes existing ceramic packages will be used, with the advantage of easier testing and design evaluation. These packages would, however, not be optimized for crosstalk or with respect to radiation effects. The final choice will be based on the results of the evaluation of prototypes as well as the mechanical design of the complete data acquisition system.

### Fault tolerance

FERMI is a very innovative concept, for which basic experimental information concerning fault mechanisms is not available. The lack of information on such basic characteristics as fault distribution in space and time, the dependence of fault rate on accumulated dose, the correlation between faults and the very nature of these faults prevents the development of a complete fault model of the system. As a consequence, the strategy for providing some degree of fault tolerance must take into account survival when exposed to single or multiple faults and a suitable balance between graceful degradation and error correction.

A fundamental point that has to be accounted for concerns the balance between area overhead and actual reliability of the final device. In the specific case of FERMI the most suitable solution for fault tolerance seems to be a combination of techniques dealing with error detection, error correction and dynamic redundancy. It was therefore decided to provide:

- single or multiple fault detection within individual subsystems of each acquisition channel, by means of error correction codes (ECC) (linear, e.g. extended Hamming, or modulo-3, for busses and arithmetic units, respectively)

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

- immunity to single faults within individual subsystems of each acquisition channel, by means of hardware redundancy
- diagnostic information via host-driven actions

Fault tolerance is implemented in FERMI 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 or ordered error checking will alarm the module controller, which will respond with a possible hardware reconfiguration. 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.

In the analog part, radiation damage will introduce a smooth degradation of performance. The reliability of this part is provided by the programmability of the look-up table and the calibration and linearisation functions built into the FERMI architecture. This also includes the analog part of the A/D converter.

Error detection must be capable of identifying the presence of at least one error within the considered function. Single-bit errors should be corrected, while double-bit error detection is adopted to identify invalid data. In addition, the availability of some degree of redundancy will increase system survival in case of independent multiple faults and/or permanent faults. The approaches adopted for the various digital subsystems are briefly described in the "Digital part" section below.

In order to investigate the fault mechanisms inherent to the FERMI design during its normal operating conditions, efforts have been made to estimate the effects of the radiation doses and types encountered in its operating environment. The expected dose of electromagnetic radiation and neutron flux inside future LHC detectors have been calculated using an updated version of the FLUKA program [5]. The electromagnetic radiation will mostly cause single event upsets, while the neutron radiation will cause errors to occur in many places at about the same time. This is due to the fact that defects from a large number of neutrons are required to cause an error. The neutron radiation can also affect the function of the device by activating surrounding material which then will serve as a secondary source of electromagnetic radiation.

## The FERMI System Environment

FERMI relies 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. It will probably 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 calorimeter trigger.
- Local second-level trigger : to read data from several FERMIs and to transmit them to the 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. Different (partial) solutions can be envisaged;

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

- Moving as much as possible of the functionality out of the detector 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 FERMI will also relax the fault tolerance requirement, since the consequences of faults are only local to a given FERMI. Furthermore, the necessary fault tolerance is also easier to realize in the ASICs of the microsystem.
- Grouping as much as possible of the remaining circuitry into a dedicated ASIC, a FERMI board controller, where fault tolerant multiply redundant circuit solutions can be used.

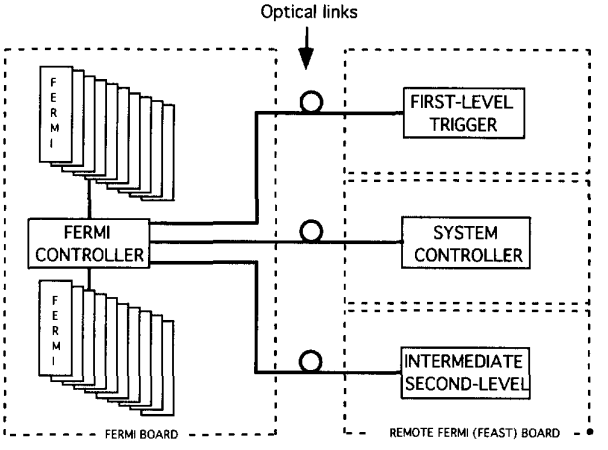


Figure 3. The FERMI system environment

Another important issue is how to solve the communication problems. Coaxial cables take up space and limit transmission speed. Optical fibers, however, are superior in both respects. The coupling between optical receivers and transmitters and the ASICs is easiest realized in a multichip 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, 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.

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 fibers 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 fibers. Fig. 3 shows a possible implementation of these ideas.

### Digital filters

The filter functions of FERMI consist of two different digital filter sections. The first filter is designed to identify a signal providing accurate timing information and relatively coarse energy information for the first-level trigger. It operates on the sum of the nine FERMI

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

channels, while the single channel data are stored in a data memory. Only data accepted by the first-level trigger will be transferred to the second filter.

The second filter is designed to extract an accurate value for the energy provided the sample timing jitter is the dominant source of error. Time frame data can be extracted by bypassing the second filter under special conditions, e.g. for particular classes of interesting events or when a severe pileup condition is detected. Also, two different coefficient banks are available to tailor the filtering to different conditions on an event-by-event basis.

**The first-level trigger filters**

The first-level trigger filter is composed of the summing unit, which adds individually thresholded data from the nine channels of a FERMI, and two FIR filters operating on the channel sum.

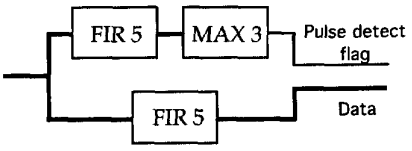


Figure 4. First filter structure for the prototype.

The optimal strategies for extracting time and energy information may be different, and the hardware for the prototype is designed to provide the maximum flexibility for exploring and optimizing both algorithms separately. This is achieved by using two parallel FIR filter branches (see Fig. 4): the timing is determined using a 5-tap FIR filter optimized for time measurements, followed by a three-point maximum finder. The energy is measured using another 5-tap FIR filter optimized for amplitude measurement.

The filter coefficients for both cases can be obtained using either analytical techniques or a least mean square (LMS) error optimization technique with a training set containing signals of different heights and different noise content.

Three different modular architectures have been examined:

i) a conventional architecture based on multiplier modules,

ii) an architecture based on single-bit convolvers [6],

iii) an architecture with parallel barrel shifters and no multipliers.

The third architecture, has been chosen for the prototype because of its inherent simplicity and high speed. It differs from the two other solutions in the fine-grained quantified representation of the FIR filter coefficients. Using only two barrel shifters (both of which can negate the output value) the coefficient values in the range e.g. (-3, 3) can be covered so uniformly and with sufficient precision that the performance requirements of the first-level filters are met without relying on complex multiplier structures.

The resolution of the timing filter, simulated under pessimistic noise and jitter conditions, is such that at 500 MeV, the error rate is 4.7 % for the training method and 1.6 % for the analytical method. For energies above 1 GeV the fraction of events erroneously ascribed to adjacent bunch crossings is less than $5\ 10^{-5}$.

**The second-level trigger filter**

The task of the second-level trigger filter is to measure the pulse amplitude with very high precision using the time frame identified by the first-level trigger. This filter has to provide full accuracy in the presence of sample timing jitter and different noise components. If a pileup condition is detected, a different set of filter coefficients is used. If the pileup is considered too severe, or for special events, the whole time frame can be read out.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

In order to measure the pulse amplitude with very high precision, a non-linear feature extraction method is used. The filter consists of a bank of linear FIR filters , followed by an order statistic operator function. The order statistic operator sorts the input values in ascending order and selects the $r$ - th largest of them as the operator output. This non-linear filter structure is best suited for pulse height measurement when the sampling positions are fluctuating; the FIR subfilters are adapted to different jitter conditions and the order statistic part selects one of them according to the outputs of the FIR filter bank. The procedure to obtain the optimal coefficients for an adaptive FIR/order statistic hybrid (AFOSH) filter is explained in [7]

## FERMI components

## The analog part

### Compression

A FERMI module receives analog signals with a dynamic range of about 15 - 16 bits. Digitization at full LHC speed is performed with a full 10-bit ADC, which requires a dynamic range compression at the analog input level. This adds only a negligible contribution to the total detector resolution, provided an appropriate non-linear transfer function is matched to the resolution of the calorimeter [1].

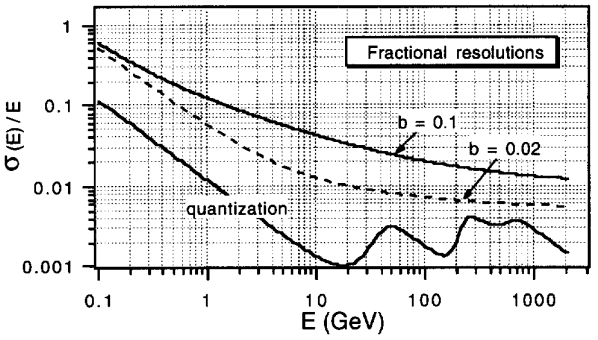


Figure 5. Measured fractional resolutions of the compressor - ADC system compared to that of typical calorimetric detectors.

A piecewise linear transfer function, where the input signal is amplified in stages with different gains and the outputs added linearly, has been chosen. Two different solutions, a bipolar and a CMOS version, have been developed.

Figure 5 shows the measured fractional resolution of a prototype of the bipolar compressor, compared with the fractional resolutions of two typical detectors. The contribution of the quantization noise to the total resolution is minimal, even for a high quality calorimeter.

### Sampling and A/D Conversion

Two different approaches to the analog to digital conversion are being followed, a Two-Stage Pipelined A/D converter and an array of Successive Approximation converters (PSA-ADC). A final choice will be made once prototypes of both types will have been thoroughly tested.

*The Two-Stage Pipelined A/D converter*

Figure 6 shows the topology of the two-stage pipelined A/D converter architecture.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

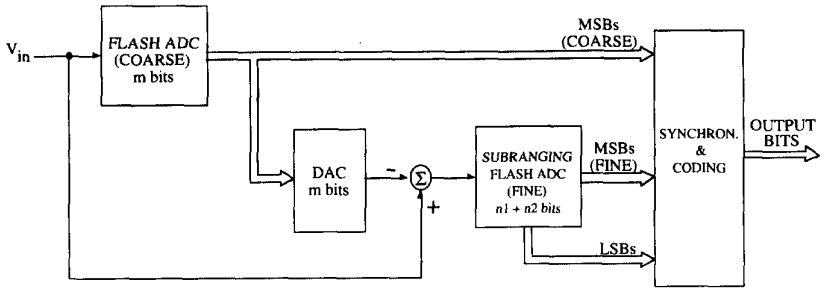


Figure 6. Block diagram of the two-stage pipelined ADC architecture.

A first coarse conversion is carried out by means of a flash ADC, to obtain the $m$ most significant bits (MSBs). The converted bits are then fed to a D/A conversion stage to generate an analog voltage which is subtracted from the input signal. The residue is then converted by a subranging ADC, which gives the remaining $(n_1 + n_2)$ least significant bits (LSBs). Finally, the $(m + n_1 + n_2)$ bits are synchronized to generate the output digital word. In order to reach an accuracy of 10 bits, and to compensate for errors arising in the subtraction and from mismatches between the coarse ADC and the DAC, the resolution of the coarse flash converter is increased by one bit, i.e. $m = 5$, $n_1 + n_2 = 6$. The 11 bits are processed by a digital encoder generating the 10-bit output word.

*The Parallel Successive Approximation ADC*

The principle of the Parallel Successive Approximation ADC (PSA-ADC) is shown in Fig. 7a.

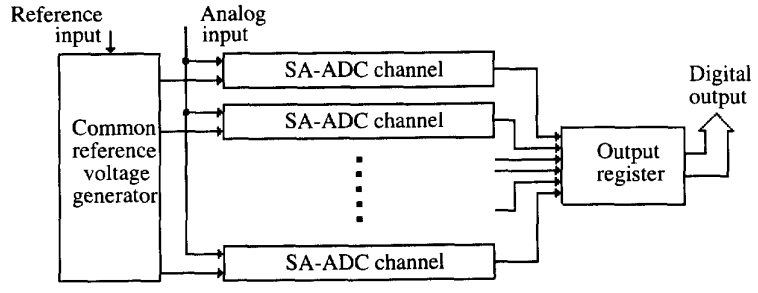


(a) A PSA-ADC.

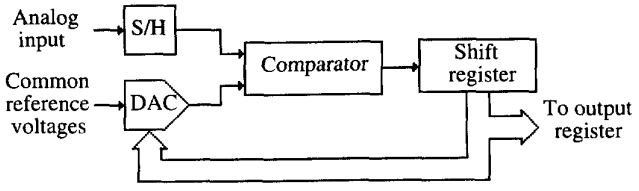


(b) An SA-ADC.

Figure 7. Block diagram of the PSA-ADC

A full converter includes a number of identical successive-approximation ADC (SA-ADC) channels, a common reference voltage generator and an output register. The number of

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

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 comprises a S/H circuit and a successive approximation ADC. Each SA-ADC, shown in Fig. 7b, comprises 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.

The parallel architecture allows high conversion speed by using low speed SA-ADCs. Since each channel is efficient, the cost of such a parallelism is low. Because of its high efficiency, a 10-bit PSA-ADC needs only 14 comparators, in contrast with 62 comparators for a half-flash ADC or 1023 comparators for a full-flash ADC. Therefore, the power consumption is substantially reduced. It consumes one third of the power of a half-flash ADC and much less power than a full-flash ADC. A PSA-ADC prototype has been tested and found to achieve full 10-bit resolution up to 70 MHz.

## The Digital part

The digital part of FERMI is responsible for generating the first-level trigger output, for temporary storage of data and for extraction of data to the second and third-level triggers. It is physically distributed over two types of ASICs the channel and the service ASICs. In each FERMI there are three channel ASICs and one service ASIC (Fig. 8).

The digital electronics in the two ASIC types can be subdivided into blocks. The channel ASICs contain look-up tables, data memory and parts of the first-level trigger functions. The major units in the service ASIC are the local controller, the first-level trigger filter, the output filter and the I/O interface functions.

All programmable registers inside the FERMI system are accessed through a single serial link utilizing a simple protocol. Under normal conditions the internal controller uses this link to configure and test the system. There is also a back-up communication link to allow an external controller to directly access internal registers even if the internal controller has failed.

In order to allow exhaustive self testing of the system, diagnostic registers (fig. 8) have been provided both to insert constant values and to read the current value along the data paths in the system. All diagnostic features are accessed by the a controller through a serial link (see below). As part of the internal test facilities the look-up tables and data memories can be used to store digital test sequences, which can then be sent to the first and second-level triggers.

### The channel ASIC

Each channel ASIC, see Fig. 8, contains three parallel channels accepting data from three corresponding ADCs. These 10-bit data paths are expanded in look-up tables into 16-bit corrected and linearized data. The data are then split into two paths: one going to the local circuitry to form the local first-level trigger sum, while the other merges with data from the other two channels. The result is fed to the data memory at a place chosen by the externally provided insert address pointer. The data selector at the output is used to select which of the three 16-bit data words located at the extract address should be transferred to the service ASIC. Each channel is enabled by a programmable threshold before summing them for the first-level trigger. This facility can be used to discriminate against channel noise and to disable malfunctioning channels entirely, so that they will not be included in the trigger summation.

Each channel ASIC is responsible for storing one of three possible flags, one pulse-detect and two pile-up flags. These flags are generated in the service ASIC and sent to each channel ASIC via a programmable selector. The chosen flag is included into the data stream fed to the data memory.

Fault tolerance is achieved by providing each data word with either an error correcting code (ECC) of Hamming type, allowing single-bit correction and double-bit detection, or a 2-bit modulo-3 residue code. The ECC protection is used when data is stored in the memories and when transferred between the ASICs, while the modulo-3 code is used to check the integrity of data in the arithmetic calculations. Faulty data is marked with an error bit that

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

accompanies data throughout the module and when data is presented at the output ports. All error signals are also recorded in status registers. The status registers then can be read by the controller so that appropriate actions can be taken in case of permanent faults. Further fault protection is achieved by triplicating all fault sensitive logic with a hardware current voting. An additional safety factor is a small associative memory operating in parallel with the main data memory used for patching faulty cells. It can also be used for diagnostic purposes.

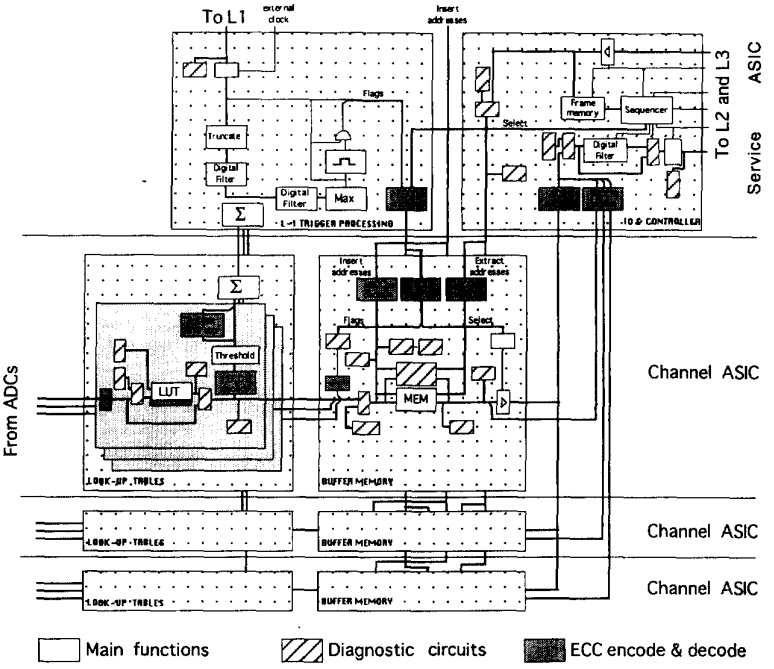


Figure 8. An overview of the digital part of FERMI

### The service ASIC

The service ASIC is common to all channels. It contains data extraction circuits for generating energy and time information to the first-level trigger, digitally filtered data to the second-level trigger and provides a local controller to manage initialization, calibration and monitoring of the FERMI system.

The first-level trigger data, resulting from the local summation on the three channel ASICs, are added together and processed in two different 5-tap digital filters, one for energy determination and one for timing identification (Fig. 8 and 4). The output of the timing filter is applied to a peak finding circuit to identify the bunch crossing and a corresponding pulse-detect flag is merged with the trigger data. 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.

A shift operation scales the filtered energy information in order to optimize data representation for the first-level trigger.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

To control and co-ordinate all the internal functions, as well as to handle their corresponding results, a local FERMI controller is required. This controller will also be the interface of each FERMI module with the external world. The controller should be able to handle locally the necessary computations related e.g. to the calibration and monitoring tasks. During normal data taking, the micro-controller is completely disconnected and its clock is inhibited in order to avoid interferences.

An internal test pulse generation circuitry is capable of generating detector-like pulses that are fed to the analog inputs of FERMI. These pulses are used both as a way to calibrate the system, finding suitable values to be stored in the look-up tables, and to test the functionality of the analog parts. A 12-bit DAC is used to generate a step function which is shaped into an appropriate pulse. Programmable registers are provided to control the amplitude, shape and delay of the pulse.

The clock manager controls the phase relationship between the internal and external clocks by providing a programmable delay. For the ADC, the sampling instant has to be defined with an accuracy of approximately 200 ps with respect to the input signal from the detector, compensating for differences in the time of flight of the particles and differences in the clock distribution delays. A clock enable function is used to disable the internal clock from all parts except for the ADCs. The clock can be enabled both by the internal controller and by an external clock enable input. The latter facility is intended for collective test sequences when a number of FERMIs are activated together by a common test pulse. With the clock manager, the internal controller can also generate bursts of clock pulses, which can be used to sample certain time intervals around a calibration pulse so that it can be examined. The sampling procedure can be investigated by repeating the procedure with different delays of the internal clock relative the external system clock.

## Design methodology

The entire digital part of FERMI is defined using VHDL (VHSIC Hardware Description Language). The connectivity of the system is described using schematics. These schematics can be converted to structural descriptions to obtain a VHDL description of the whole module. In order to allow the use of different design tools appropriate modifications have to be introduced in the code. They have been added in such a way that different versions of the code, eg for simulation and for synthesis, can be generated from the basic source models using a macro preprocessor. The ASIC implementations are verified using testvectors generated from the VHDL simulations.

Before transferring the VHDL code into the first silicon ASIC prototype, a final verification using FPGAs (field programmable gate arrays), will be done. This approach allows to:

- record real signals and compare them with simulated ones
- provide hardware for the test bench to test software routines before the demonstrator is completed
- optimize and iterate the description of the digital part before final implementation.

The breadboard prototype is implemented on VME boards integrated into the FERMI test bench. A common bus from the service IC board connects all the VME boards for serial communication.

The channel IC part is split into two VME boards, one having three independent channels with their LUT memories, counters and multiplexers, the other carrying the main memory and the logic handling pointers from the address generator.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.

The service IC part is on a board housing the filters, the readout sequencer, the controller and the serial communication interface. The output is fed to a fast dual-port memory and read out with LABVIEW™ software. A VME interface allows to load and monitor the hardwired synthesized VHDL FERMI code.

## Conclusions

The design of the FERMI module is well under way. The aim is to fabricate a prototype module in the spring of 94. Large parts of the analog parts have already been produced and tested. The fabrication and test of the entire digital block is foreseen for the second half of 1993 and will be partly implemented in rad-hard technology. The FERMI concept will then most certainly proliferate into a family of modules adapted for different environments.

## 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: "A Digital Front-end and Readout Microsystem for Calorimetry at LHC", *CERN/DRDC/90-74*, December 1990.

[3] The FERMI Collaboration: "Status report on the FERMI project", *CERN/DRDC/93-21, RD-16,* May 1993.

[4] C. Bohm and G. Appelquist, A Memory Controller for FERMI, *CERN RD-16 note #5,* November 1991.

[5] A. Ferrari et al, *CERN RD-3 note #22,* 1991.

[6] L. Dadda, M. G. Sami, "High-speed Parallel Input-output Bit-sliced Fault-tolerant Convolver", *Proc. IEEE 1992 Intl. Workshop on Defect and Fault Tolerance in VLSI Systems*, pp. 287-296, 1992.

[7] S. J. Inkinen, J. Niittylahti, "Trainable FIR-Order Statistic hybrid filters",Submitted to: *IEEE Transaction on Circuits ans Systems -II: Analog and digital signal processing,* April 1993.

Authorized licensed use limited to: UNIVERSITY OF PAVIA. Downloaded on March 10, 2009 at 07:34 from IEEE Xplore. Restrictions apply.