

A Modular Fault-Tolerant Approach to Design a Front-End Microsystem for Calorimetry at LHC

Luca Breveglieri, Vincenzo Piuri

Department of Electronics and Information, Politecnico di Milano
piazza Leonardo da Vinci 32
I-20133 Milano, Italy

Abstract - The FERMI system has been designed to perform acquisition and DSP of calorimeter data in high energy collision experiments, planned at the LHC collider (CERN, Geneva, CH). The system relies mainly upon the FERMI chip, a dedicated VLSI multi-chip device: it requires rad-hard features due to the critical application, to the difficulties in maintenance, and to the operating environment. In this paper, a modular approach to support policies for self diagnosis and fault tolerance is discussed.

INTRODUCTION

The electronics for signal detection and triggering for experiments at the future proton supercolliders will be one of the most complex systems conceived so far for high energy physics. Calorimeters are one of the basic measurement systems to observe and evaluate physical events. For such systems, the requirements of complexity, speed and resolution (triggering of events) call for an unprecedented capability of identification, compaction and processing of data. This requires an intelligent front-end for the detectors placed in the calorimeter, in order to select and process adequately the experimental data.

On the other hand, massively fault tolerant approaches to design the front-end for such advanced calorimetric systems must be considered, in particular for the electronic components, in order to provide a high degree of reliability for the measurements, a high availability of the whole calorimeter and possible survival even in the presence of a limited number of faults.

In [1-4], a prototype architecture for a digital front-end (the FERMI system) for a liquid argon calorimeter [1, 2] and a preliminary approach to fault tolerance have been proposed, to support nuclear experiments at the Large Hadron Collider (LHC) of CERN, Geneva, CH. The detection system embodies a very large number - around 5×10^5 - of acquisition channels. It is composed by two main parts:

- a local subsystem, consisting of about 10^5 basic VLSI devices (called the FERMI chips), which provides the first level of data selection, compression and processing, by removing the meaningless samples;
- a remote subsystem, implemented by means of a set of host computers, which processes the outputs of the local subsystem, to perform the final recognition and classification of the relevant events.

The local subsystem and its main component, the FERMI chip, is up to now in the most advanced development phase [1].

This paper presents an advanced, modular approach to design the digital, local subsystem, which is directly connected to the calorimetric sensors. In [3], a preliminary analysis of the fault tolerance features and some possible solutions have been discussed, but a large circuit complexity was adopted to grant system survival through error correction with a limited reconfiguration capability. In the present paper, we consider a different approach: the novelty of the design solution presented consists in adopting a highly regular and simple structure for the FERMI chips, suited to support effectively the error detection capabilities and the system reconfiguration in the actual specific fault and error models. Different levels of protection are envisioned to deal with different classes of faults according to the specific rad-hard environment. Solutions are analysed from the point of view of circuit complexity, (i.e., silicon area), throughput, computational time, detection probability and survival probability, with respect to the expected fault model for the novel operating environment at LHC.

In the second section, the general structure of the FERMI system is envisioned, with particular regard to the architecture of the FERMI chip. In the third section, the fault and error models are introduced and discussed, relating it to the environmental conditions of the system. In the fourth section, the policies for introducing self diagnosis and fault tolerance are presented and discussed. The conclusion lists directions for future research.

STRUCTURE OF THE FERMI SYSTEM

The FERMI system is composed by a local and a remote subsystem. Together they perform three levels of triggering of the input data from the calorimeter: the local subsystem performs the first triggering, while the remote subsystem the second and the third ones. At each triggering level the amount of data forwarded to the subsequent level is greatly reduced. The FERMI chip is the basic module of the local subsystem; its architecture is shown in Figure 1. The present prototype version is a 5×5 cm silicon wafer, implementing three channel chips and one service chip, mounted with the flip-chip technology. Each channel chip contains three acquisition channels, which perform compression of the analog signal from the sensors in the calorimeter, digitisation to 10 bits, expansion to 16 bits by means of a look-up table (a RAM, off-line loaded with a suited transfer function) and merging of the significant samples of three acquisition channels into a single data memory. The addresses of this data memory are supplied by an external address generator, which manages a whole set of FERMI chips.

Significant samples are triggered by circuitry in the service chip. Nine acquisition channels are thresholded, in order to filter out low-level noise, and then added up. The accumulated samples are processed by a non-linear trigger filter, which detects the pulses corresponding to significant events. The stored samples are then processed by a second non-linear filter; the outputs of the second filter are delivered to the second and third level triggerings (performed remotely), together with the first level triggering strobe. The service chips contains a controller, that coordinates all activities of the service chips itself and of the three channel chips of the FERMI chip.

Figure 2 shows the architecture of the two non-linear filters implemented in the service chip. The trigger filter is a programmable 5-taps parallel convolver that processes the accumulated outputs of 9 acquisition channels, followed by a maximum extractor that computes the peak value and therefore triggers the event time. The output filter is more complex: it is formed by a bank of three programmable parallel convolvers, of increasing length, that process data for the second and the third level triggering functions (performed remotely), followed by a votation computed by a programmable rank ordering operator. The three convolvers process data with increasing temporal resolution and the rank ordering operator sorts their outputs, selecting the most significant one. This allows to achieve a compromise between a high temporal resolution and a low latency [6].

The remote subsystem is implemented by means of a group of host computers, connected to the local subsystem with optic fibres, and is placed in a separate and protected environment, while the local subsystem is placed in the immediate neighbourhood of the liquid argon calorimeter it is connected to, in a rad-hard environment.

The whole system is synchronous and works at a sampling frequency of 64.7 MHz, but it is supposed to be speed-up to a sampling frequency of 80 MHz in the next future. All these temporal requirements and the range of the digital signals to be processed (16 bits, two's complement binary representation) are dictated by physical requirements [1, 2, 4]. The FERMI chip is intended to be a prototype, to be tested in the operating conditions and redesigned in the future since the novelty of the operating environment does not guarantee correctness and completeness of the fault model assumptions.

THE FAULT AND ERROR MODELS

The quantity of data to be processed by the FERMI system is extremely large, yet data contain a very large degree of redundancy. However, from physical considerations [1, 2], significant events are expected to occur at a very low frequency with respect to the intrinsic frequency of the




Figure 1: Structure of the FERMI chip.




Figure 2: The two non-linear filters in the service chip.

calorimeter circuitry, and to be mixed with a large number of irrelevant events as well as to be masked by background noise. This implies that, in spite of the large degree of redundancy, both the *completeness* (no loss of data) and the *credibility* (no corruption, masking or aliasing) of data need absolutely be granted.

The project was started in the absence of definite knowledge about the behaviour of microelectronic digital devices in rad-hard environments, in particular the absence of characterisations of the nature and of the spatial/temporal distributions of faults. This has prevented the possibility of formulating realistic complete fault and error models.

Preliminary testing results [5] show that both transient and permanent faults (in general multiple) are to be expected for a set of VLSI technologies, including CMOS and Silicon on Insulator (SOI), in environmental conditions which reproduce those of the LHC. Analog devices exhibit a tendency to graceful degradation, combinatorial (logic/arithmetic) and sequential (memories) circuits suffer from both transient and permanent faults; moreover memories tend to die completely after massive radiation exposure. However, different technologies exhibit a wide range of behaviours; in particular, sequential circuits built in rad-hard technologies (SOI) stand radiation quotas similar to the maximum foreseen for the experiments at LHC without fatal failure. This entitles rad-hard technologies for the design and fabrication of the final FERMI chips.

Given such premises, it is advisable to refer to a functional fault model [10]: the FERMI chip is partitioned into modules and faults are defined as the malfunctioning of a module. In this way, a definite dependency on the physical nature of the faults is eliminated. Multiple faults are modelled as several modules simultaneously faulty, independently of the number of physical faults of the chip.

There exist, however, parts of the FERMI chip for which lower level fault models are applicable: busses and memories. For such modules, a stuck-at fault model [10] is advisable, since it favours the estimation of the performances of error correcting codes, by means of well-known mathematical techniques [8].

As for the multiplicity of the faults, the above mentioned studies [5] indicate that both single and multiple faults are to be expected, for both the functional and stuck-at fault models. Therefore it is not possible to confine to traditional techniques for single faults only, but it is necessary to consider explicitly multiple faults. It must be noted that the functional fault model naturally reduces some multiple faults to a single one, whenever the faults occur inside a single module, thus partially simplifying the problem of handling multiple faults. Moreover, some error correcting codes can be easily extended to partially handle also multiple stuck-at faults [8, 9].

To match the above physical requirements and environmental constraints it has been decided to design the FERMI system giving priority to *self*

diagnosis capabilities, in order to grant that data delivered from one level of the system to the next one are *correct*. Should this be no longer true, the related piece of information is better discarded. *Self correction* is implemented whenever possible, compatibly with costs, measured in terms of the overhead of silicon area, of time and of circuital complexity. The indeterminacy of the fault models has suggested to fulfil the following priority list:

- handling both transient and permanent faults;
- handling single and possibly also some multiple faults;
- providing graceful degradation.

As a result, the designed FERMI system and, in particular, the FERMI chip provides the following functionalities:

- information is complete and credible (up to physical limitations, of course);
- the system is fault tolerant at least to single faults and to some multiple faults, occurring in separate devices;
- the system gracefully degrades in the presence of multiple faults.

Due to the locality of the extremely large number of FERMI chips and to the amount of data to be processed, self diagnosis is performed for the FERMI chip in a highly concurrent and distributed way. This implies that each FERMI chip need autonomously perform self diagnosis. On the contrary, not all fault tolerance policies can be performed locally, but need be deferred to a remote monitoring, in particular the fault tolerance functions requiring a large reconfiguration of the system, by excluding/including a number of acquisition channels. Remote control is necessary for higher level processing of the triggered information, since they are too complex to be performed locally.

THE FAULT TOLERANCE ISSUES

The high number of acquisition channels and the fact that the sensors are embedded in the calorimeter with high redundancy imply that the sampled data are highly redundant as well. Therefore, it is always possible to achieve fault tolerance by programming the second and third level triggerings to discard the samples supplied by the faulty acquisition channels, without compromising the overall performances of the system or, at least, only gracefully degrading them. Therefore it is possible to envision a modular approach to the introduction of diagnosis and fault tolerance capabilities in the FERMI system. This modular approach is a trade-off of diagnosis and fault tolerance capabilities against cost, measured in terms of

the overhead of silicon area and latency. We can define three policies for the protection of the FERMI chip:

- 1) Introduce only diagnostic capabilities, relying upon the intrinsic redundancy of the system and software reconfiguration (of the more advanced triggering levels) for fault tolerance and graceful degrading.
- 2) Introduce diagnostic capabilities and fault tolerance, by using more powerful coding techniques and modular redundancy to protect the core elements of the FERMI chip, in particular the overhead modules required by the introduction of the enhanced codes themselves..
- 3) Introduce diagnostic capabilities and more fault tolerance, by using coding techniques and modular redundancy at the level of the core elements of the FERMI chip and also at the level of the entire acquisition channels.

Clearly, software reconfiguration and graceful degradation can still be used when adopting policies 2 and 3; the diagnosis tools implemented in policy 1 are reusable for policies 2 and 3, with possible enhancement.

Given the lack of definite information on the fault model (in particular, on the distribution of faults), for the implementation of policy 1 in the digital part of the FERMI chip, it is advisable to adopt techniques capable of:

- detecting single or multiple faults within individual subsystems of each acquisition channel, by means of ECC. These are linear cyclic redundancy codes [8] (e.g., parity code) for busses and memory elements, and arithmetic codes [9] (e.g., 3N code) for parallel adders, threshold comparators and parallel convolvers (see Figures 1 and 2);
- providing diagnostic information and for the host-driven reconfiguration of acquisition channels [7].

Following the guidelines presented in the previous section, self diagnosis and has been implemented in the digital functional units composing the FERMI chip, as follows.

Look-up table (LUT): two levels of fault management are provided for this device:

1. **Concurrent testing:** it is split into two further levels.

1a. **Memory word:** the Hamming code is adopted; since no correction is implemented, this code can be used for double detection, therefore covering also some multiple faults. The encoders/decoders are core elements and can be protected for single faults by means of double modular redundancy with a current voting scheme (supported by CMOS technologies), which allows to design voter circuits with a very limited silicon area overhead [4].

1b. **Address decoder:** double modular redundancy with a current voting scheme is adopted, for single fault detection, in the address decoding section. To provide detection of more complex fault distributions each word is associated with the signature of its address and a signature checking is performed at readout, as it is shown in Figure 3.

2. **Periodical testing:** it can be performed through a memory checksum operation, involving not only the data part of the word, but also the signature part. The structure shown in Figure 4 supports such a scheme, by means of the adoption of an address counter and of a functional unit for the evaluation of the checksum. Some degree of host-driven testing is required, to handle the testing procedure.

Data memory: the same technique adopted to protect the look-up table (concurrent and periodical testing) is suited to the protection of the data part of the memory buffer, with some modifications to take writing into account. The data memory address decoder is protected by means of double modular redundancy with a current voting scheme, similarly to the LUT.

Busses: they are protected by means of ECC codes, which are arithmetic codes for the busses connecting the arithmetic functional units and cyclic redundancy codes, in particular the Hamming code, (which can be used for double error detection) for the busses connecting the remaining functional units. All encoders/decoders are protected for single faults by means of double modular redundancy with a current-voting scheme, as they constitute core elements for diagnosis.

Digital filters: arithmetic codes are used for fault diagnosis in digital filters [9]; the encoders/decoders are themselves protected for single faults by means of double modular redundancy with a current-voting scheme, as they constitute core elements for diagnosis and fault tolerance. Figure 5 shows the structure of this detection scheme; the filter is modelled as a convolver followed by a non-linear function, which models attains to both the triggering and output filter (see Figure 1).




Figure 3: Concurrent testing for the look-up table.




Figure 4: Periodical testing of the look-up table.

Samples are encoded, then convolved and processed through the non-linear part of the filter and are eventually decoded. All encoders and decoders are duplicated and their outputs are voted by means of a current checker (not shown in Figure 5), which is very small compared to the size of the duplicated devices. The same duplication and votation is applied to the non-linear function, as it is not protected by the arithmetic code (most arithmetic codes are linear, see [9]); such functional units (max. and rank ordering) are however small compared to the size of the parallel convolvers.




Figure 5: Arithmetic testing for the digital filters.

A similar protection scheme attains also to the parallel adders and the threshold comparators (see Figure 1). Possible arithmetic codes to use are 3N codes and residue codes [9]. 3N codes are very simple, requiring only multiplication and division by 3, but are not suited to correction [9]. It is possible to convert all data to a suitable residue representation, perform all linear and non-linear functions and then restore the ordinary two's complement representation. Encoders and decoders need be duplicated and voted as well. In addition to diagnosis, it is easy to obtain also correction, for a suitable residue representation [9].

The next step is the implementation of policy 2, that is the introduction of fault tolerance at the level of the individual modules forming the FERMI chip, wherever it is feasible, depending on the silicon area overhead and the strict frequency and latency constraints (15 ns sampling period and no more than 2 μ s latency for the first level triggering function [1, 2]). The

protection schemes for policy (1) are retained, but codes are enhanced and some degree of triple modular redundancy is introduced, instead of duplication only. To do this, it is necessary to use coding techniques capable of performing single (and possibly multiple) fault correction and to provide survival to at least single faults within the individual subsystems of each acquisition channel, by means of hardware redundancy [7].

A possible scheme is triple modular redundancy for the most critical and otherwise unrecoverable functional units of each acquisition channel. The shaded modules shown in Figures 1, 2 and 3 are the core elements of the FERMI chip to triplicate: essentially, they are the encoders/decoders themselves, since a fault in these modules would be fatal.

Since the samples registered in the data memory are already validated by the first level triggering function, it is important to grant their credibility. For this reason, the addressing of the data memory is further protected: an associative memory keeps track of the faulty locations in the data memory, as the scheme of Figure 6 shows. The associative memory itself is protected similarly to the LUT; the checker (shaded area of Figure 6) is triplicated, as it is a core element.




Figure 6: Fault tolerance for the data memory.

The last step is the implementation of policy 3, that is the introduction of fault tolerance at the level of acquisition channels; the amount of redundancy to introduce depends on the silicon area overhead. Figure 7 shows a possible scheme: switches are introduced to include/exclude the working/faulty acquisition channels. The diagnosis of each channel is provided by means of the implemented diagnosis tools of policies 1 and 2, which are retained, and by means of the supervision provided by the second and third level triggering functions. The switches are core elements and therefore they need be implemented by means of robust technologies, if not by means of modular redundancy as well.




Figure 7: Fault tolerance for the acquisition channels.

As a general remark on the whole approach (policies 1, 2 and 3), the basic tools used for achieving diagnosis are error detection codes, of logic or arithmetic type, depending on the functions to protect. Fault tolerance is obtained by using error correction codes, wherever the overhead with respect to detection only is feasible. Core elements, i.e. the encoders and decoders are duplicated or triplicated, as the silicon area overhead is limited with respect to the nominal area [4]. In fact, the encoders/decoders are associated to functional units which are mostly memories and parallel convolvers; clearly the sizes of these functional units are very large in comparison to those of the encoders/decoders [4, 8, 9]. The voting, which is itself a core function, is performed by means of a very small and

simple current voting circuit, which does not increase sensibly the core silicon area. Since address decoding for memories is not protected by the ECC's and therefore is a core function, address decoders are themselves duplicated or triplicated and then voted. All diagnostic pieces of information which are not used directly by the FERMI chip are forwarded to the second and third level triggerings.

CONCLUDING REMARKS

The main choices concerning fault detection and fault tolerance have been devised and a prototype of the FERMI chip is currently being designed, which will undergo fabrication and testing by 1994. It implements all envisioned protection policies, but the modular redundancy of acquisition channels. For the prototype a mix of CMOS and Silicon-On-Insulator (SOI) 0.8 μ technologies are used.

The FERMI chip will be placed in a hostile environment, exposed to severe radiation stimulation. The results of the testing will allow to evaluate the effective performances of the self diagnosis and fault tolerance functions of the FERMI chip, and possibly to redesign them. The FERMI chip should be capable of processing information according to the following specifications:

- corrupted data, provided by faulty elements, are *immediately* discarded;
- correct data, provided by working elements, are *safely broadcast* to the remote hosts, which are placed in a less hostile environment;
- the nature of the fault, i.e. whether it is transient or permanent, can be determined;
- the structure can be reconfigured, wherever it is possible; this can be partially done at least at the level of acquisition channels, by discarding the faulty ones and by relying upon redundancy thereof.

The remote part of the FERMI system, i.e. the hosts that perform the second and third level triggering functions and provide a remote control for the FERMI chips, is currently in a preliminary phase of design. However, the design of this part of the system is fairly less critical, since it is placed in a controlled environment. Moreover, the first level triggering function and the DSP functions provided by the FERMI modules themselves greatly reduce the amount of data which are broadcast to the remote hosts, relaxing their performance requirements.

REFERENCES

- [1] The FERMI Collaboration, "Status Report on the FERMI Project - A digital Front-End and Readout Microsystem for Calorimetry at LHC", *CERN/DRDC Report n. 93-21 RD-16*, May 1993
- [2] The FERMI Collaboration, "A digital Front-End and Readout Microsystem for Calorimetry at LHC - The FERMI Project", *IEEE Transactions on Nuclear Science*, vol. 40, n. 4, IEEE Nuclear and Plasma Society Press, August 1993
- [3] The FERMI Collaboration, "System Level Policies for Fault Tolerance Issues in the FERMI Project", *Proc. of the IEEE International Workshop on Defect and Fault Tolerance in VLSI Systems*, Venice, Italy, October 1993
- [4] The FERMI Collaboration, "FERMI - A new Generation of electronic Modules for large Data Acquisition Arrays required by high Energy Physics", *Proc. of the IEEE Int'l Conference on Wafer Scale Integration*, San Francisco, California, U.S.A., January 1994
- [5] A. Kerek, "Front-End Electronics in the Radiation Environment of LHC", *Proc. of the IEEE International Workshop on Defect and Fault Tolerance in VLSI Systems*, Venice, Italy, October 1993
- [6] S. Inkinen, J. Niitylathi, "Adaptive FIR / Order Statistic Hybrid (AFOSH) Filters", *submitted to the IEEE Transactions on Circuits and Systems II: analog and digital Signal Processing*, April 1993
- [7] R. Negrini, M. Sami, R. Stefanelli, *Fault Tolerance through Reconfiguration in VLSI and WSI Arrays*, The MIT Press, Cambridge, U.S.A., 1989
- [8] W. W. Peterson, *Error Correcting Codes*, John Wiley & Son, New York, 1961
- [9] T. R. N. Rao, *Error Coding in Arithmetic Processors*, Academic Press, New York, U.S.A., 1974
- [10] J.A. Abraham, W.K. Fuchs, "Fault and error models for VLSI", *Proc. IEEE*, May 1986