

Exploring Design Dimensions in Flash-based Mass-memory Devices

Original

Exploring Design Dimensions in Flash-based Mass-memory Devices / Caramia, M.; DI CARLO, Stefano; Fabiano, Michele; Prinetto, Paolo Ernesto. - STAMPA. - (2009), pp. 43-48. (Intervento presentato al convegno ACM 4th International Workshop on Software Support for Portable Storage (IWSSPS) tenutosi a Grenoble, F nel 15 Oct. 2009).

Availability:

This version is available at: 11583/2296447 since:

Publisher:

ACM

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Politecnico di Torino

Exploring Design Dimensions in Flash-based Mass-memory Devices

Authors: Caramia M., Di Carlo S., Fabiano M., Prinetto P.,

© ACM, 2009. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM 4th International Workshop on Software Support for Portable Storage (IWSSPS), 15 Oct., 2009, Grenoble, FR and is available at:

URL: <http://140.118.127.40/iwssps2009/iwssps09-design.pdf>

Exploring Design Dimensions in Flash-based Mass-memory Devices

Maurizio CARAMIA

Thales Alenia Space
Command Control and Data Handling
Strada Antica di Collegno 253
10146, Torino, Italy

Maurizio.Caramia
@thalesalieniaspace.com

Stefano DI CARLO, Michele FABIANO,
Paolo PRINETTO

Politecnico di Torino
Dipartimento di Automatica e Informatica
Corso Duca degli Abruzzi, 24
10129, Torino, Italy

{Stefano.Dicarlo, Michele.Fabiano,
Paolo.Prinetto}@polito.it

ABSTRACT

Mission-critical space system applications present several issues: a typical one is the design of a mass-memory device (i.e., a solid-state recorder). This goal could be accomplished by using flash-memories: the exploration of a huge number of parameters and trade-offs is needed. On the one hand flash-memories are nonvolatile, shock-resistant and power-economic, but on the other hand their cost is higher than normal hard disk, the number of erasure cycles is bounded and other different drawbacks have to be considered. In addition space environment presents various issues especially because of radiations. The design of a flash-memory based solid-state recorder implies the exploration of different and quite often contrasting dimensions. No systematic approach has so far been proposed to consider them all as a whole: as a consequence the design of flash-based mass-memory device for space applications is intended to be supported by a novel design environment currently under development and refinement.

Categories and Subject Descriptors

B.8.1 [Reliability, Testing, and Fault-Tolerance]

Keywords

Flash-memory, Space applications, Design environment

1. INTRODUCTION

Flash-memory based systems are gaining acceptance and usage not only in the consumer market but in space applications, as well, where they could play the role of high-capacity storage devices: in fact flash-memory guarantees both the non-volatility in case of power loss and a highest storage density, being at the same time shock-resistant and power-economic [1].

Mission-critical space applications require proper mass-memory device (i.e. solid state recorder). Designing a solid state recorder for space is basically different from developing a typical mass-memory device for common use, mainly due to the peculiar operational environments [1].

This goal could be accomplished by using flash-memories: the exploration of a huge number of design dimensions and evaluating a vast amount of trade-offs among all such dimensions. The most relevant dimensions include flash-memory technology, flash-memory architecture, file management system, dependability enhancement strategies, power consumption, weight, physical size.

Unfortunately in the literature each paper is typically tackling just one specific issue in just one design dimensions, without proposing any systematic approach in order to consider them all as a whole. Such a concurrent exploration capability is mandatory to provide the designers a powerful design environment, capable of supporting them through all the steps of the design cycle, including Architectural Exploration, Design Validation & Verification, (Automatic) Test insertion, Dependability evaluation and so on.

A complex and powerful design environment is thus needed to properly evaluate the impact of the choices in each dimension and the related trade-offs. A novel design environment currently under development and refinement is aimed to support the design of flash-based mass memories, especially for space applications, and is exploited to highlight the effectiveness of the concurrent exploration discussed above.

The paper is organized as follows: Section 2 is a brief overview of the main issue related to the space environment, Section 3 resume the major characteristics and issues associated to flash-memories, while Section 4 addresses a possible architecture for a design environment to support the design of flash-based hard disks for space applications.

2. THE SPACE ENVIRONMENT: AN OVERVIEW

The processing power available for embedded technology and boards just a few years ago is absolutely not comparable to the present one. However in space applications the design choices

have often driven toward older and/or lower-performing radiation-tolerant electronics because of the very strict requirements.

Although each new space application has its own story and increasing requirements [2], a typical mission-critical space system application includes, among the several aspects, the design of a solid state recorder(s). No loss of mass memory data and the guaranteed availability of storage capability at End-Of-Life (EOL) are only some of the many different constraints a solid state recorder for critical space missions needs to satisfy. The requirements of interplanetary missions have to be met and a well-designed flash-based memory system can do it. However its design has to compensate for flash's shortcomings in speed, radiation tolerance, noise, and read/write cycle life: this compensation obviously leverage the costs.

On the one hand vendors should absolutely provide flash-memories physically qualified to survive in the space environment with the help of proper strategies [3] – [6], while on the other hand data integrity, reliability, simplicity, modularity, and autonomy are just some of the key features to fulfill (e.g., reliable storage of context data, also during spacecraft power outage).

Moreover ECC algorithms for error checking and correction of NAND flash-memory should be evaluated by designers in order to accomplish the level of dependability requested by their design. A more detailed survey about the most peculiar flash-memory design dimensions and trade-offs to tackle during the design of flash-based hard disks for space applications could be found in [7].

3. FLASH-MEMORY: CHARACTERISTICS AND ISSUES

Several interesting features of flash-memories properly fit with the requirements of mass-memories for space applications, whereas feasible alternatives need to be evaluated. A relative high capacity would characterize our solid-state recorder: a suitable solution could be DRAMs. On the one hand DRAMs are very fast, reliable and provide a very high data rate, but on the other hand they need a battery pack-up to not lose data and this issue generate an intricate balance between battery mass and data retention time: data retention over years is not feasible and count of battery charge cycles are limited. DRAMs are not discussed anymore in this paper.

3.1 Characteristics

A more attractive solution is the use of flash-memories. There are two major types of flash-memory in the current market: NOR and NAND flash-memory. NOR flash-memory is for EEPROM replacement and is more suitable for program execution, while NAND flash-memory is more suitable for storage systems [8],[9]: [7] briefly sums up the main characteristic of these types of flash-memory.

This paper addresses only NAND flash-memories: in fact they are the most suitable choice for HD replacement. On the one hand, flash-memories are nonvolatile, shock-resistant, and power-economic: a pure flash-memory based solution could guarantee unlimited data retention time and no need of battery backup. On the other hand flash-memories are still much more expensive than hard disk drive memory. Only rather large data structures could be accessed and, in addition, DDR2 SDRAMs provide higher

write/read rate (e.g., 6 Gbit/s) compared to the moderate one accomplished by flash-memories (e.g., up to 80/200 Mbit/s) [1].

Moreover one of the main challenging aspects of flash-memories is that the space already written (i.e., programmed) with data usually cannot be overwritten unless it is erased from the flash-memory device. A NAND flash-memory is usually partitioned into blocks: each block has a fixed number of pages and each page has a fixed size. A block is the smallest unit for erase operations, while read and write operations are done in terms of pages, i.e., a page can be erased only if its whole corresponding block is erased, whereas a page can be read/written independently.

In addition flash-memory wears out after a certain number of erasure cycles (i.e., actually 10^6 for NAND flash-memory): if the erasure cycles of a block exceed this number, it becomes a “bad block” and is not reliable for storing data anymore.

Finally another peculiar aspect of flash-memories is that technology provides the possibility of storing more than one bit of information per cell: in fact newer devices (Multi-Level Cell or MLC) are able to store typically two bits per cell, while traditional flash-memory devices (Single-Level Cell or SLC) store only one bit per cell.

3.2 Main Issues

The investigation of a vast quantity of design parameters needs to be defined during the design of a flash-based system for space application. A possible partial taxonomy could involve *Flash-memory Technology*, *Flash-memory Architecture*, *Flash-memory Testing*, *Dependability and Wearing* and, finally, *Using flash-memory as Hard-Disk*.

Designer should discriminate among which technology to choose (*Flash-memory Technology*) and they may also have to select the most appropriate flash-memory chipset (*Flash-memory Architecture*). [7]

Flash-memory Testing is quite different from testing other kinds of memory: disturbances or faults not conforming to any of the traditionally known fault models used in testing RAMs could be experienced and specific fault models are needed to properly represent the most frequent physical defects are needed. Then efficient test algorithms are needed and Built-In Self Test (BIST) and Built-In Self Diagnosis (BISD) circuits have to been taken in consideration. [7]

In the problem of *Wearing*, a significant role is played by the so called wear leveling techniques, which are aiming at distributing data evenly across each memory block of the entire flash-memory to avoid single block to wear out. Several interesting wear leveling techniques have been proposed [10] – [13] and could be considered or higher capacity flash-memory devices could be used, then especially taking care of the resulting drawbacks in terms of weight and volume [1]. An effective comparative analysis of some wear leveling algorithms could be found in [11].

However wearing, testing and dependability are strictly linked among each other: designers should work hard on exploring all their possible alternatives and combinations, discriminating among the most suitable ones and evaluating the right trade-off among them.

Several challenging aspects need to be addressed when *using a flash-memory as a mass-memory device*. A possible incomplete taxonomy of these aspects could involve:

- implementing an efficient logical to physical address translation process for fast operations (i.e., *Address Translation*) [7]
- looking for proper solutions in order to let OS successfully communicate with NAND flash-memory devices (i.e., *Operating System Management*) [14] – [19]
- proper strategies for reclaiming invalidated space to be erased in order to free some space (i.e., *Garbage Collection*) [7]
- proper techniques to handle blocks exceeding the maximum number of erase cycles (i.e., *Bad Block Management*) [20]

Finally designers should address Error Detection And Correction (EDAC) techniques, evaluating the most proper choice for their design. [7]

[7] is a more detailed investigation about these main flash-memory issues discussed above.

4. FLASH-BASED HARD-DISKS DESIGN SUPPORT FOR SPACE ENVIRONMENT

Several steps are needed in order to provide an efficient support for designing flash-based mass-memory devices that have to survive into a mission-critical space environment. First of all a preliminary study is needed in order to better understand how to decompose a flash-based hard-disk into composing functional blocks. Then a design environment is fundamental to support the design of flash-based hard disks for space applications, following the functional decomposition discussed before.

These two steps are briefly discussed in the sequel of this section and are actually in continuous refinement.

4.1 The Design of Flash-based Hard-disks

A first possible high view of a flash-based mass-memory device is shown in Figure 1. There are three main functional blocks: a *Non-Volatile Memory* is needed to provide integrity of data, a *Volatile Memory* is used for performance reasons, while a *Memory Controller* is managing and controlling the overall system, providing several features. The mass-memory device is interacting with the requests of the external world, e.g., the operations of the OS in use.

It is possible to refine this first architectural view: in fact, during the design of a flash-based HD, the OS and the applications want to successfully communicate with the bare flash-memory chip. On the one hand OS usually would like to exploit its typical system calls (e.g., *open*, *read*, *write*) to work with the mass-memory device, without taking care of anything else. On the other hand, the flash-memory chip would like to receive the most proper commands to accomplish the operations previously requested from the OS. E.g., [13] briefly presents a typical layered system architecture of popular flash-memory-based file systems.

So designers should develop a sort of *managing part* to tackle all the typical issues of flash-memories, presented also in the previous paragraphs and in [7]. This is the most important and challenging part of designing a flash-based HD: many often

contrasting issues, parameters and dimensions are involved in this part, which has to address them in the best way as possible.

This *managing part* could be named “flash-memory manager” and could be split into its composing functional blocks as Figure 2 shows: these blocks represent the main issues a flash-memory based system has to tackle and to solve in the more possible efficient way. Designers have to manage how the logical to physical *Address Translation* is accomplished: reliability and efficiency are only two of the parameters of quality of this aspect. They need also to focus on *Wear Leveling* and *Garbage Collection* techniques and strategies. At the same time, designers need to distinguish among several *EDAC* strategies: a trade-off between needed reliability and related costs leads their choice over a particular code rather than another one. In addition designers have to manage *Bad Blocks*.

4.2 The Evaluation Design Environment

Actually the development of a flash-based hard disk qualified for space applications is provided with no systematic support. Designers have always to think about the most suitable choice for the specific space applications they are dealing with: the huge number of variables and parameters could easily lead to unverified scenarios and to delayed product release.

In fact the level of confidence with these parameters is directly linked with the designers’ skill, cleverness and experience. As a result, a systematic tool to support the design of flash-based hard disks for space applications is needed.

4.2.1 Assessment Criteria

Several aspects of the design of a flash-based system are intended to be evaluated with the help of FLARE tool. Designers have to tackle many critical issues: FLARE could help them to distinguish and identify the peculiar features of these aspects and to evaluate the most suitable solution for them.

Designing a flash-based HD means dealing with NAND flash-memories, which are always partitioned in blocks and in turn each block is divided in pages: once *capacity* is set, designers have to address the *dimension* of each block and of each block or, in the same way, the *number* of blocks and pages for each block. Obviously, designers could decide the *capacity* from the dimension and the number of blocks and pages, but the issue is practically the same.

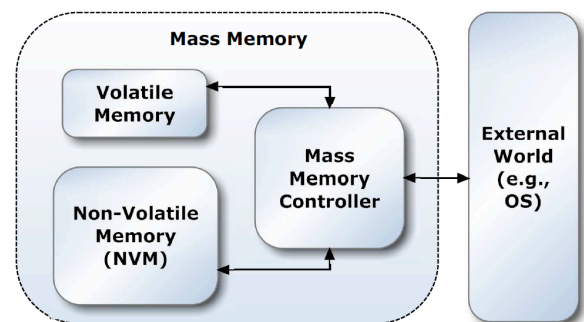


Figure 1 – A Flash-based Hard Disk Architecture

FLARE could help the designers to do this decision, in order to understand which level of granularity would more properly fit with the current design and to decide the most suitable flash-memory chipset.

The *capacity* of the flash-memory is the other fundamental parameter to set: designers should discuss about the physical quantity of flash-memory required by the design. This is a typical issue of space applications: in fact space critical missions require minimizing all the costs as much as possible and the dimension of the flash-memory is the first significant parameter that designers and their companies have to face.

Designers could have to discriminate among different flash-memories of different capacities during the design of their system: FLARE could provide them with an overall evaluation of which *capacity* is more suitable for their specific design.

Designers have to provide a well-defined level of dependability according to their specific design. A fundamental role could be played by the so called *Out-Of-Bound (OOB)* data [16]: they can be exploited also to store some kind of ECC/EDAC codes, in order to accomplish the required dependability. The smart reader could get the unavoidable trade-off between spare data and user data: in fact it is true that a bigger *OOB* area could provide higher level of dependability, but at the same time would provide poor service in term of user data storage. Designers have to tackle this issue and find the most suitable solution for their particular design. Moreover, designers have to evaluate among the *Built-In Self Test (BIST)* functionalities, evaluating at the same time the *percentage of errors detected/corrected*.

An essential parameter to evaluate is the *percentage of wearing of each block*: especially in mission-critical space applications resources are always a key-point of the mission and it is desirable or, usually, mandatory that the percentage of wasted resources is as low as possible. E.g., on the one hand it could be enough 2GByte NAND flash-memory with some kind of wear-leveling techniques or on the other hand a bigger NAND flash-memory device could be requested in order to accomplish mission requirements. At the same time, designers could need to explore several kinds of solutions, in order to find that one with the most fitting percentage of wearing. As a consequence, this parameter is strictly linked to the adopted wear-leveling strategies: with the help of FLARE, designers could evaluate how this percentage varies as the wear-leveling techniques change.

As a consequence designers need to calculate the percentage of flash-memory which is not “dead”, i.e., the percentage of blocks which did not become *bad blocks* at the End-Of-Life (EOL). Mission-critical space applications sometimes could explicitly require a fixed amount of flash-memory still alive at the EOL: designers have to evaluate the possible alternatives and to find the most affordable solution at the minimum possible cost for their design.

This is only a possible incomplete taxonomy of what is needed to be evaluated during the design of a flash-based mass-memory device for space applications. Moreover all these parameters are strictly linked together and they affect each other in a complex way: so an exploration of these different and quite often contrasting dimensions is needed and no systematic approach has so far been proposed to consider them all as a whole.

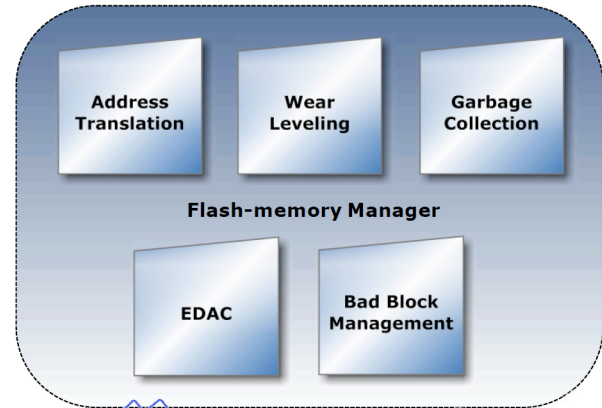


Figure 2 – Main functional blocks of the *Flash-memory manager*

4.2.2 The Evaluation Tool Architecture

The proposed FLash ARchitecture Evaluation (FLARE) design environment is aimed at supporting designers through all the steps of the design cycle flash-based hard disk for space applications, including Architectural Exploration, Design Validation & Verification, (Automatic) Test insertion, Dependability evaluation and so on. FLARE is currently under development.

Figure 3 shows the architecture of the system.

4.2.2.1 The Configuration Management Block

The exploration of all the possible alternatives and design dimensions is allowed with the help of the *System Configuration Management*: designers are able to easily modify the memory configuration block (*Architecture configuration*), the Test infrastructure (*Test configuration*), and all the architectural solutions aimed at tackling Flash aging (*Bad block, Garbage Collection, Wear Leveling Configuration*).

The *Architecture Configuration* block is intended to contain all the details about the architecture of the flash-memory to emulate: capacity, number of blocks and number of pages are only some of the main architectural parameters that the designers are able to set.

With the *Test Configuration* block, the designer can set all the parameters for correctly testing the proposed flash-memory: all the issues addressed in the previous paragraphs are taken into account and the proper fault-models and the specific testing strategies can be specified, always according to the particular application flash-memory is used for.

As clearly showed previously, some wear-leveling strategies are needed to spread writes over the flash-memory: designers are capable to exploit the *Wear Leveling Configuration* block to specify all the details about the wear-leveling strategies to adopt during the emulation campaign. The range of these details can be variable: designers could choose a “simple” less/more aggressive wear-leveling strategy among the ones just provided with FLARE tool or developing their own wear-leveling algorithm could be a valid alternative, in order to evaluate it.

If wear leveling strategies aim to spread write operations over the flash-memory device, at a certain point invalidated space should be reclaimed: in the *Garbage Collection* module, designers are able to specify the strategies to identify a block, to collect its

good pages and to erase it. It is usually strictly connected with wear-leveling strategies: it could even be considered that GC preferences are managed by wear-leveling strategies, but these two issues are kept separated now.

However, blocks exceeding the maximum number of erasure cycles are marked as bad: in *Bad Block Configuration* module designers can set the proper parameters to mark, identify and exclude bad blocks from active space memory. Simple well-known strategies could be used (e.g., *Skip Block Method*) as well as new approaches can be experimented and evaluated by designers.

Dependability of flash-memory need to be guaranteed: designers are able to specify in the *EDAC Configuration* block all the parameters needed to accomplish the required level of data integrity and reliability. E.g., a reasonable question could be if a CRC code is enough to accomplish the requested level of reliability or something more is needed. Maybe some ECC would be absolutely necessary to accomplish the required level, e.g., Orthogonal Reed-Solomon Error Correction Code might be the EDAC strategy designers were looking for. In *EDAC Configuration* block, designers are capable of defining, exploring and evaluating all possible EDAC strategies for their particular design.

The designers' configuration choices feed the so called *Configuration Manager* block: this layer is thought to take care of managing the "static" data coming from the various dimensions of the design of a flash-memory device (i.e. the "note" modules on the right) and to get it across the core functional blocks. This layer is essential for dispatching the updated configuration modules discussed above to the appropriate managing blocks. The architectural choice of having this kind of layer is strictly linked to flexibility: in fact on the one hand if some changes to parameters and algorithms are needed, designers can simply modify the proper module(s) not interesting in the rest, because the *Configuration Manager* layer will take care of dispatching the updated configuration(s) to the appropriate blocks. On the other hand, designers are capable of developing new (compatible) configuration modules, in case they felt like the existing modules were not enough for their needs: adding new configuration modules to the whole architecture would turn in very few efforts, thanks to this way of partitioning, and would result in high modularity and flexibility.

4.2.2.2 The Core Simulator

Flash Memory Simulator is the newly developed system kernel, charged of providing the designer the possibility of simulating and evaluating all the parameters of interest.

The *Flash-memory Simulator* block is one of the most important functional blocks of the FLARE tool: in fact it is thought to emulate the behavior of the configured flash-memory. The desired architecture is specified in the *Architecture Configuration* module discussed before: a "customized" architectural configuration for the flash-memory device could be identified or a ready for use configuration could be chosen from a developed library. Then the *Configuration Manager* takes this information and advertises the *Flash-memory Simulator* block about the architectural details of the flash-memory to emulate.

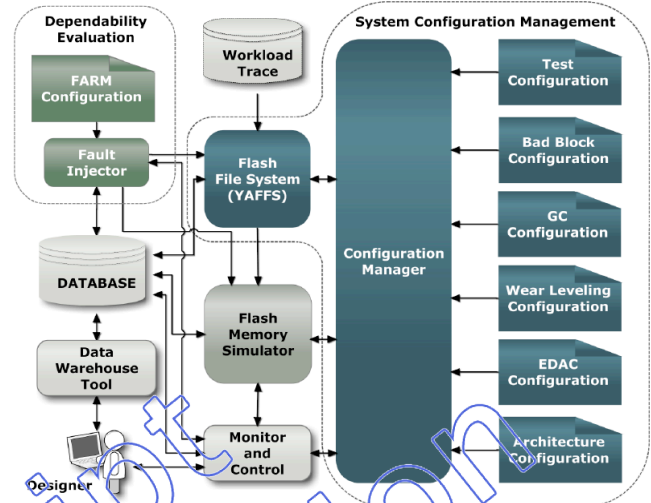


Figure 3 – A detailed view of the FLARE architecture

4.2.2.3 Dependability Assessment

In addition to the overall architecture, some fault injection techniques could be considered: a *Fault Injector* functional block is added for this purpose. It is fed by a *Fault Activation Readout Measure (FARM) Configuration* block [21]: it sets all the needed parameters for the fault injector to make it work as requested. In this way a fault can be injected in the system to evaluate its effect in the emulated flash-memory. Fault injection is an additional function of the FLARE tool: in fact it is represented surrounded by a dotted rounded rectangle in order to highlight this point, i.e., it is not essential to the correctness of the FLARE tool, but at the same time it could be very useful for experimenting various fault injection techniques and configurations.

A *Fault injection* environment provides the designer to assess the target system dependability via a powerful manager of fault injection campaigns in all the part of the system itself. [21]

4.2.2.4 The Monitor and Control Blocks

As the name intuitively suggests, the *Monitor and Control* block is monitoring and controlling the output of the previous *Flash-memory Emulator* block. Designers can have under control all the events of the core blocks in order to get a more comprehensive knowledge about the countermeasure to take in some specific cases. The *Monitor and Control* block is peculiarly different from the *Data Warehouse Tool* block: in fact the last one is a mean with which the user can extract information about the emulated flash-memory at the EOL timeline, whereas the first one is a sort of automatic tool informing the user about the most significant events of the actual emulation campaign.

The use of a *Database* is fundamental to gather all the information needed at the EOL timeline: its role is simply to store data. The user is able to access the data with the help of a *Data Warehouse Tool*: data and metadata can be extracted, transformed and loaded, to easily accomplish all the designers' requests.

5. CONCLUSIONS AND FUTURE WORKS

Flash-memories are gaining acceptance in several fields of application and they could be adopted also in space applications, where they could be exploited as a mass-memory device.

However several and quite often contrasting dimensions need to be explored during the design of a flash-memory based hard-disks and no systematic approach has so far been proposed to consider them all as a whole: FLARE is intended to be the one.

The composing blocks of the proposed architecture highlight the high-level of modularity and flexibility that this tool will be able to provide to designers: in fact each block is intended to be a sort of plug-in block, which can simply be plugged-out and replaced by another block when necessary, without taking care of the rest. At the same time, when little configuration changes are needed they can be easily accomplished without involving the way in which they will be updated in the managing blocks.

FLARE tool is currently under development and refinement: the first implementation data of the tool are intended to be provided soon.

6. REFERENCES

- [1] Cassel M., Walter D., Schmidt H., Gliem F., Michalik H., Stähle M., Vögele K., Roos P. Casel.: "NAND-Flash-memory Technology in Mass Memory Systems for Space Applications", *Proceedings Data Systems In Aerospace (DASIA) 2008, Palma de Mallorca, Spain, 2008*
- [2] Anthony Lai: "Space-ready, radiation-tolerant processor modules: A COTS technology strategy", *Military Embedded Systems Resource Guide, May 2005*
- [3] Brüggemann M., Schmidt H., Walter D., Gliem F., Michalik H.: "Further Heavy Ion and Proton SEE Evaluation of High Capacity NAND-Flash-memory Devices for Safeguard Data Recorder", *8th ESA/ESTEC/TEC-QCA Final Presentation Day, February 2007*
- [4] Schmidt H., Walter D., Brüggemann M., Gliem F., Harboe-Sørensen R., Virtanen A.: "Heavy Ion SEE Studies on 4-Gbit NAND-Flash-memories", *Radiation Effects on Components and Systems (RADECS) 2007, DWL-14, September 2007*
- [5] Schmidt H., Walter D., Gliem F., Nickson B., Harboe-Sørensen R., Virtanen A.: "TID and SEE Tests of an Advanced 8 Gbit NAND-Flash-memory", *Proc. IEEE Radiation Effects Data Workshop, 2008, 38-41*
- [6] Brüggemann M., Schmidt H., Walter D., Gliem F., Harboe-Sørensen R., Roos P., Stähle M.: "SEE Tests of NAND Flash-memory Devices for Use in a Safeguard Data Recorder", *Radiation Effects on Components and Systems (RADECS) 2006, A-3, Volume A-3, 2006*
- [7] Caramia M., Di Carlo S., Fabiano M., Prinetto P.: "Flash-memories in Space Applications: Trends and Challenges", *To appear on East-West Design & Test Symposium (EWDTS) 2009, Moscow, Russia, September 18-21*
- [8] M-Systems: "Flash-memory Translation Layer for NAND Flash (NFTL)", 1998
- [9] Intel Corporation, Flash File System, US Patent 540, 448
- [10] SanDisk Corporation, White Paper: "SanDisk Flash-memory Cards Wear Leveling", *Doc. No. 80-36-00278, October 2003*
- [11] Chang Li-Pin: "On Efficient Wear Leveling for Large-Scale Flash-Memory Storage Systems", *Proceedings of the 22nd ACM Symposium on Applied Computing, 2007*
- [12] M. L. Chiang, Paul C. H. Lee, R. C. Chang. "Using Data Clustering To Improve Cleaning Performance For Flash Memory", *Software - Practice and Experience, 1999*
- [13] Chang Y.-H., Hsieh J.-W., Kuo T.-W.: "Endurance Enhancement of Flash-Memory Storage Systems: An Efficient Static Wear Leveling Design" *Proc. 44th ACM/IEEE Design Automation Conference (DAC) '07, 212-217, 2007*
- [14] Intel Corporation, Technical Report: "Understanding the Flash Translation Layer (FTL) Specification", *December 1998*
- [15] Hsieh Jen-Wei, Tsai Yi-Lin, Kuo Tei-Wei, Lee Tzao-Lin: "Configurable Flash-Memory Management: Performance versus Overheads" *IEEE Transactions on Computers, Vol. 57, no. 11, 2008*
- [16] Chang L. P., Kuo T. W.: "An efficient management scheme for large-scale flash-memory storage systems", *Proceedings of the 2004 ACM Symposium on Applied Computing, Nicosia, Cyprus, 862-868, 2004*
- [17] Woodhouse D., Red Hat, Inc.: "JFFS : The Journaling Flash File System", <http://sources.redhat.com/jffs2/jffs2.pdf>, 2001
- [18] JFFS2, <http://sourceware.org/jffs2/>
- [19] Aleph One Company, Cambridge, UK: "Yet Another Flash File System", <http://www.aleph1.co.uk/yaffs/index.html>, 2002
- [20] Samsung, Application Note: "XSR1.5 Bad Block Management", 2007
- [21] Benso A., Prinetto P.: "Fault Injection Techniques and Tools for Embedded Systems Reliability Evaluation" – *Kluwer Academic Publishers, ISBN: 1-4020-7589-8, 2003*