ESA FPGA Task Force: Lessons Learned

A. Fernández-León [1], A. Pouponnot [1] and S. Habinc [2]
[1] European Space Agency/ESTEC, Noordwijk, The Netherlands
[2] Gaisler Research, Gothenburg, Sweden

Abstract

This paper gives an overview of the audits that ESA made on several FPGA designs included in the Rosetta spacecraft. It describes first what motivated the creation of the FPGA Task Force and what procedures were followed. It presents then a summary of the findings at the different subjects inspected: documentation generated, resources employed, overall design methodology, radiation effects countermeasures, verification and validation practices and miscellaneous system aspects. Results of the radiation tests actually carried out on some of the devices are also commented. The paper is closed with a collection of lessons learned and general conclusions.

I. INTRODUCTION

Rosetta, the first European comet orbiter and lander mission is scheduled for launch in January of 2003. Its eight years journey until rendezvous with comet Wirtanen, includes several asteroid flybys and a two year hibernation period. Onboard autonomy is particularly important because the round trip time for radio signals exceeds 90 minutes for a significant and critical part of the mission. During the course of its Critical Design Review, at the end of the year 2000, it became clear that many FPGAs would be flown in critical parts of its avionics. However, not enough information was available at that moment as to assess the quality of these designs and the reliability of the actual FPGA implementations. Several other forthcoming ESA missions were also found to be making extensive use of FPGAs. In particular, and due to lack of proper documentation, there was great concern about the methodologies followed to achieve adequate insensitivity to SEUs and other radiation effect.

II. FPGA TASK FORCE

To address these concerns, an FPGA Task Force group (hereinafter referred to as FTF) was set up at the beginning of 2001, integrated by ESA staff members and representatives from the prime contractor of the Rosetta avionics. The FTF *Terms of Reference* defining the composition of the team and its main objectives, and a *Work Plan* detailing the initial set of tasks to achieve these goals was drawn.

A. The Goals

The main goals were to identify the most critical FPGAs to be flown in Rosetta, to review the overall quality of these designs, their susceptibility to problems due to radiation, and to assess any potential problems and their solutions at spacecraft and/or operational level.

Out of a total number of 45 FPGAs found, 18 different designs were identified, and, after a first screening, a smaller selection of six critical designs was finally chosen to undergo a more in-depth analysis. All of these FPGAs were ACTEL RT14100A parts, responsible for crucial functions in the Control, Data Management and Power units. Neither one of them could be power-cycled in flight.

B. The Auditing Methodology

The auditing methodology evolved and was consolidated along the one year duration of the reviewing activities.

Initially, specific technical documentation on the FPGA designs (functional specifications, architecture, verification and validation documents, source VHDL code and gate level netlist) was requested for analysis. With an irregular response to this demand, the FTF initiated a series of specialized meetings with the industrial parties responsible for the FPGA designs, in an effort to gather sufficient, sometimes critical first-hand information, from the actual design engineers. The vast amount of information retrieved was itemized and grouped into chapters to ensure completeness of the reviews and easy traceability of the potential problems found (by the ESA auditors). These reports would eventually reflect the industry counter-analyses of the problems indicating the foreseen system impact and their proposed solutions.

To complement and reassure some of the theoretical analyses of the source code and its associated documentation, several ad-hoc gate level simulations were performed. These simulations were mainly focused on understanding and discovering potential functional misbehaviors due to SEU.

Lastly, and to add confidence on the theoretical predictions, several heavy ion tests were conducted on five of the six FPGA designs. Additionally, another ACTEL 14100A device carrying an independent test design was subjected to radiation tests to verify SEU sensitivity levels.

One year after, the FTF has put together a Final Report with all the relevant information, analyses, lessons learned, conclusions and recommendations for future FPGA/ASIC developments.

III. THE FINDINGS

One of the outcomes of the FTF activities has been a comprehensive *technical questionnaire*, in the form of an Excel book, that can now be re-utilized in future FPGA (or ASIC) design investigations. It can be used as a checklist tool for the auditor to make sure that all the critical information about the design is available or, if not available, requested. Of course, it will then be the auditor team's responsibility to identify any real or potential problems with respect to each information item in the list, and to decide which actions should be taken in each particular case. Normally, and as seen from this FTF's experience, a typical action will be to raise the problems found to the design and system engineers in order to obtain a counter analysis and arrive to possible solutions.

Although in this paper we will focus more on the negative aspects found, i.e. those which could be improved for future developments, it is fair to say that there were as well many positive aspects which not only met the FTF expectations, but also exceeded them to the point that they can also be categorized as positive *lessons learned*.

Following the structure of the aforementioned questionnaire chapters, we present a summary of the findings here below.

A. *Documentation*

This very important subject can indeed be considered as the element which triggered the concerns and distrust on the FPGA designs. The lack of visibility on what design methodology had been followed to generate these FPGAs along with the sometimes absolute lack of information required to exert a minimal control on the quality of these microcircuits were found unacceptable and sufficient reason to set out the FTF activities.

A common practice observed was that designers did not generate specifically dedicated documentation for the FPGA designs. Instead, they had most of the functional, architectural, verification and validation information scattered across system and board documents. This made it very difficult and time consuming to find and control all relevant information.

For some designs, there was no formal track of what the verification campaign consisted of, and what results were obtained. A simple verbal "everything was simulated successfully, but there was no time to document" was the only answer provided to the FTF. Another weak point found with respect to some requirements specifications was the fact that there was not a precise way to trace and identify individual itemized requirements (that could later on be identified as such for verification/validation purposes). However, a simple number tag attached to each requirement was used in some documents, making it much easier and safer to work with throughout all the design phases.

Other information that was often not available or irregularly distributed among several documents was that concerning radiation requirements and countermeasures adopted.

In many cases, some of the requested documentation was prepared ad-hoc for the purposes of the FTF. This was primarily the case for the questions on design protection against radiation effects, SEU probabilities and their functional impact.

Overall, the documentation made available was formally correct, up to date in most cases, and, sometimes, subjected to a rigorous and precise internal control system.

B. *Resources: Design Team, SW, HW*

Human resources allocated for the management and engineering of the designs were found to be, in general, quite reasonable. Teams were normally composed of one Project Manager, one System Engineer and, most of the times, only one FPGA Design Engineer per device. All designers had a more or less strong background in VLSI design, most of them having designed ASICs for space applications in the past. Some had not had any previous experience doing FPGAs, but managed well to learn and use the specific FPGA tools.

However, the fact that only one person was entirely responsible for all the detailed design tasks proved to be a serious limitation for the verification and validation phases, as well as for the overall production of documentation. The FTF independent reviews spotted several functional issues that had gone unnoticed during the design internal review.

VHDL code was used for design entry in all instances, and typical state of the art simulation and synthesis tools were used: ModelSim, VeriBest (Mentor Graphics) and Synplify (Synplicity). FPGA vendor-specific tools and libraries were used for place and route, back-annotated timing analysis and FPGA programming: ACTEL Designer, ACTEL Silicon Sculptor and ACTEL Activator.

In many cases, the system/board designers did not understand VHDL, which meant they had to rely entirely on the information provided by the FPGA designers.

The principal problems found which directly related to the tools used were: unknown bugs and unexpected synthesis results that could only be discovered by detailed netlist inspection and/or fault injection, as opposed to normal stimuli at inputs. More details are given in following sections E and F. Some design teams lacked the capability to evaluate "code coverage" in their simulation tools.

C. *Physical Device Parameters*

All FPGAs reviewed in this exercise were ACTEL RT14100A devices whose main features are, (according to

ACTEL data books [3]): anti-fuse based, 10k equivalent ASIC gates, 228 I/O available, 5.5V CMOS 0.8µm technology by MEC, radiation tolerant enduring TID 20-70 krads (Si), LET limits of 28 and 110 MeV.cm^{2}mg^{-1} for SEU and latch-up respectively.

The percentage of device utilization in terms of logic cells was always high (above 85%), whereas in terms of pins was low (less than 50%). In several cases, redundancy schemes to counteract SEU effects was ruled out or minimized due to lack of space.

The working frequency of the devices studied was relatively low in all cases and never seemed to pose any problems.

D. *Design Methodology*

The overall design methodology was based on previous ASIC developments, and adapted to the specific FPGA tools. Only one of the three different design houses happened to be putting together their own internal ASIC/FPGA design flow, which, by the way, was targeted to acquire ISO9001 qualification.

Most of the critical issues during the design (coding) phase that would ensure a good quality end result had been taken into account. However, the exercise of reviewing aspects such as the reset approach, metastability, power-up, glitches at outputs, test modes, internal floating nodes, clock domains, clock skew, deadlock states, and SEE design protections revealed that many design tasks should be further improved in future ASIC/FPGA developments.

This should be started by the mere need to document properly how all of these issues have been tackled, not only to improve customer confidence, but also to facilitate potential future modifications in case of bugs or re-utilization.

The general external asynchronous reset was never synchronized internally to minimize possible metastability problems at reset de-assertion (and thus avoid unknown states after reset). ACTEL was inquired about this issue and they have admitted not to have any experimental data on the issue. In any case, the internal state of all storage cells after reset was carefully reviewed by both the FTF and the designers themselves to ensure that the circuits had been designed to be fully deterministic.

Metastability and synchronization of different clock domains had been reasonably tackled in all designs, but not with the same level of analysis of the subject (only a few had gone through the process of calculating Mean Time Between Failures figures).

E. *Radiation Effects Countermeasures I: Protections against bit-flip*

Most of the FTF review efforts was spent on investigating the level of protection of the designs against single event effects, mostly Single Event Upsets (SEUs). The approaches varied quite significantly from one design to the other:

Only one of the six FPGAs audited had Triple Modular Redundancy applied to all its storage cells. This TMR was introduced automatically in the design by means of the synthesis tool.

TMR was completely discarded for two other FPGA designs adducing lack of space and an equally effective alternative method. Instead of TMR, they opted for implementing “combinatorial flip-flops” which exhibit a much lower SEU sensitivity. The rationale behind this decision was carefully documented, concluding that the probability of getting a bit flip in either FPGA during the entire mission life (10.5 years) was 0.6×10^{-3}.

On the other three FPGAs, TMR was applied selectively (to save area too), only for the most critical functions. Only 8%, 13% and 14% of the flip-flops in these three FPGAs respectively were TMR protected.

Even though the probabilities of getting a bit flip were very low, the FTF went through the exercise of reviewing every single storage cell in each design (not TMR protected), and tried to identify the internal functional impacts of getting an SEU there. The goal was to see if an SEU could have a severe or catastrophic impact on the overall mission, no matter how low the probability. Moreover, there was serious concern that multiple SEUs could accumulate because TMR effectiveness could be weakened during the hibernation phase due lack of clock activity. This type of detailed SEU analysis had only already been done by the designers of the two C-FF based FPGA designs, and proved to be very useful in spotting various potential system problems.

After extensive FPGA and system level analysis of the many potential problems found, it was decided not to change any of the designs. Instead, in the event of any of the individually analyzed SEUs, the security and well functioning of the units were to be preserved with SW and operational level countermeasures. Just to give a flavor of what type of potential SEU functional mishaps were found, we enumerate a few: loss or generation of invalid commands, degradation of on-board time, failure of internal reconfiguration commands, spontaneous triggering of internal resets, loss and/or corruption of telemetry frames, entering erroneous functional modes due to SEU accumulation, spurious alarm settings, deadlock in bus arbiters, spurious accesses to external memories, etc.

It has become clear that this kind of SEU-Failure Mode Effect Criticality Analysis should always be carried out in future FPGA/ASIC developments where flip-flops are left unprotected against SEUs.

F. *Radiation Effects Countermeasures II: Protections against deadlock states*

The other big front while investigating potential SEU problems was searching for deadlocks in Finite State Machines (FSM). The FTF did several interesting findings in this area which highlighted the need to better understand and control what the synthesis tools can and cannot do.

In one FPGA the combination of TMR with automatic recovery from illegal states in one-hot coded FSMs, all of it presumably done automatically during synthesis, proved not to work as expected. It was confirmed that TMR was appropriately implemented (at least in the pilot cases which were inspected and simulated), however, when forcing two SEUs in the FSM to enter an illegal state, the FSM not only did not come back as promised to the idle state, but instead, oscillated and sometimes locked indefinitely. The designer had not verified these structures in sufficient depth, and always trusted blindly the results of the synthesis tool.

In another FPGA, a critical FSM had been originally encoded to endure illegal transitions due to single-bit flips, however, the synthesis tool stripped off the redundant flip-flops. Not only the intended protection against SEUs was eliminated, but a new potential problem appeared on the netlist because an internal control signal was left dependent on a single flip-flop value, thereby creating yet another chance of single flip functional havoc.

On a third case, the synthesis tool decided to replicate flip-flops of an FSM to cope with certain high fan-out load requirements. These replicated flip-flops increased the number of SEU-induced "illegal states" of the FSM. This went unnoticed to the designer, who never performed any fault injection simulations nor inspected this section of the netlist nor read these details in the synthesis log files. Once this circumstance was discovered, it was seen that the FSM would oscillate when entering these illegal states. Fortunately enough for this particular case, the oscillations would only last a finite amount of clock cycles since the FSM would eventually be reset by the rest of the control logic in the FPGA.

In some cases, the FSM optimizer of the synthesis tool was switched off and the type of FSM coding and the recovery mechanisms were manually coded by the designer to have full control (or at least, more control) over the resulting netlist. However, this did not prove to be sufficient when it came to preserving redundant protection logic associated with the FSMs.

G. *Verification Practices*

The usual RTL and gate level simulations had both been done in all cases. However not all designs had been inspected and simulated as thoroughly.

Only three FPGAs had been subjected to code coverage test to ensure that every single line of VHDL code had a purpose and was exercised during simulation. This verification exercise was done on the FTF request, prior to the beginning of the audits. In the rest of instances, the code coverage tool was not available or this test was simply not considered of importance.

Most of the times, gate simulations were only performed after place and route, once a Standard Delay Format (SDF) file was available with the timing information of the actual layout.

Only self-standing simulations for individual designs were run, even though, in some cases, the same design house was responsible for several FPGA designs that interacted on the same board/unit. Instead, it was preferred to carefully analyze that their common interfaces worked as expected.

The planning and results of all the verification process was irregularly documented. Sometimes, nice comprehensive records of everything had been generated, including compliance matrices to allow easy tracing of all the requirements and the tests done to ensure their fulfillment. In one case they simply had nothing except their own hand written lab notes.

Practically in all instances the same person who coded the design was the one creating the testbenches and checking the simulation results. The FTF believes that this practice can easily lead to mask problems that often will only surface when an independent, fresh approach is implemented to verify the design.

In one case, the verification of all the simulation results was always done by visually inspecting output and internal signals waveforms after every run, instead of the much preferred way of implementing a self-checking go/no-go test. This approach made impossible, among other things, to compare automatically the results of the early RTL simulations with the final post-layout runs.

In some cases pseudo-random test cases were generated to verify unforeseen functional cases.

As already discussed, designers had done very little verification of the SEU protection mechanisms. The FTF realized the intrinsic difficulties to verify exhaustively, for example, all the TMR structures, but actively got involved and urged the design teams to, at least, script and visually inspect the gate level netlists, and to conduct fault injection simulations for all critical FSMs and other representative pilot cases. These extra efforts uncovered many synthesis-related problems to everyone's surprise (see chapter III, section F).

Finally, it should be noted that one important obstacle that the FTF had to face when working in these independent verification activities was the fact that, for three of the six FPGAs, the FTF could have only restricted access to the source VHDL code, and always at the contractor's premises. This turned out to be a serious handicap and ESA is determined to avoid this problem in the future by contractually ensuring that the source code for any design

funded by the Agency is delivered and fully accessible, at least, for independent review at the customer premises.

H. *Device Programming and Validation*

The devices were programmed with ACTEL specific tools, following strict control and handling protocols for all the blank FPGA devices, the programming files and the burnt FPGAs.

No special problems were encountered neither during the programming process, nor during or after the burn-in tests that were done for three of the devices. These tests were carried out despite the fact that ACTEL firmly discourages to perform burn-in as it would not be of any added value, according to them.

Validation of the devices was performed at board/unit level, sometimes without the participation of the actual FPGA designers. Power consumption calculations were normally postponed until this stage.

I. *System Aspects*

Power-up was one of the principal concerns for the FTF, and it was found to be adequately managed at board level. Pull-up and pull-down resistors had been placed to neutralize the inactive pins, and all the critical control signals coming out of the FPGA were externally blocked until normal FPGA operation was secured after power-up and reset phases.

The power-on reset lines were also adequately driven, and only in one case a strange configuration was found in which a feedback loop from an output of the FPGA itself was used to control the termination of the reset. Additional analysis and temperature tests were conducted to take the reset circuitry to extreme cases, and it was finally concluded that the circuit was reliable. However, and due to the intrinsic uncertain behavior of the FPGA pins during power-up, the FTF would recommend not to make the power-on reset dependent on FPGA outputs.

All unused pins were left floating, and all the special ACTEL configuration pins (i.e. MODE, SDI, DCLK and V_{PP}) were found to be properly biased as per the device datasheet recommendations.

All the FPGAs were used either two or four times within the system, since redundant sub-units and cross-connection were used to achieve fault tolerance at this level.

IV. SEU Radiation Tests

SEU radiation tests were performed on engineering model FPGA devices that were carefully chosen as representative of the actual flying parts. The goals were two: to build up confidence in the predicted SEU sensitivity of the device itself, and also to complement the theoretical analyses on the critical functions and their predicted behavioral changes in the event of SEUs.

The measured SEU sensitivity of the ACTEL RT14100A FPGAs used in Rosetta was found to be well correlated with former tests results obtained in 1997 [4].

Permanent faults induced by heavy ions, due most likely to anti-fuse dielectric rupture, were observed during that last test campaign. The estimated probability of such permanent faults was found to be very low (10^{-7} permanent faults per FPGA during the entire mission) and, therefore, of no significant risk to the Rosetta mission.

During radiation test campaign at board level, i.e. to identify the functional impact of SEUs, many SEU effects were recorded. The FPGA boards were submitted to Ne, Ar, and Kr tests with tilt angles from 0 to 60 deg. covering a LET range from 5.85 to 68 MeV with a fluence of one million ions/cm^2 per test. Most of the findings confirmed the theoretical analyses that had been previously performed, although not all the predicted effects could be measured or observed (some of them due to their extremely low probability of occurrence).

A few unexpected single event functional anomalies were also observed. They had never been theoretically predicted and remain unexplained to date, despite the designers and FTF efforts. The probability of these unforeseen effects was found to be negligible. The problems were found to be either self recovering or controllable from ground.

V. Summary of Lessons Learned

■ Above all, a reliable ASIC/FPGA Development Methodology should be applied for the definition, design, verification, physical implementation and validation phases of any ASIC/FPGA to be flown as part of the spacecraft platform or critical payload. This should be contractually enforced. Here at ESA we will continue to make applicable our own internal standard [1], or any other equivalent methodology proposed by the contractor. As soon as the new forthcoming ECSS standard on this subject [2] is available, ESA will start using it as an applicable document in all projects where ASICs or FPGAs are to be developed.

■ There should always be clear contractual obligation for the design house to provide its customer with the "*essential information*" necessary to conduct an independent review and quality control of the FPGA development (just as it is normally done for ASICs). The documentation package should always include, at least:

- a self-standing and comprehensive FPGA Specification describing in detail all its functions, architecture and interfaces.

- comprehensive Verification and Validation Plans and Result Reports.
- HDL source codes (VHDL, Verilog, etc) and/or schematics of both, design and simulation platforms.

■ The radiation threats to the circuits of the FPGA/ASIC should be assessed and documented separately. The radiation effects countermeasures should be evaluated, justified and properly implemented and verified.

■ A comprehensive assessment of all possible malfunctions due to SEU (flip-flop by flip-flop, FSM by FSM) should be performed and propagated to the system engineers for further evaluation and potential feedback to the design.

■ We reconfirmed once again the old lesson learned which is never well enough learned: never trust completely what the CAD tools will do (or you think they will do) for you.

■ Gate level netlist inspection and simulation (including fault injection in critical cases) should be done to verify not only the nominal functionality, but also the externally-induced-error (e.g. by radiation or noise) correction/mitigation logic such as redundancy schemes and automatic recovery mechanisms out of illegal FSM states. This is particularly important if an "intelligent" CAD tool has automatically implemented this logic.

■ Unwanted extra logic, such as flip-flop replication, just as well as unwanted logic removal, such as redundant logic optimisation, can damage SEU protection logic and go unnoticed with traditional verification practices. The results of automatic synthesis tools should always be scrutinized with respect to these two issues.

■ Verification activities should be done (or complemented) by somebody else than the actual designer of the circuit, in an effort to uncover problems or special cases that might not be apparent unless looked at from an independent point of view.

VI. Conclusions

This reviewing exercise has served to acquire a realistic view of what some of the actual trends are in the design and utilization of FPGAs for space applications. The audits have revealed a serious need to enforce stricter FPGA design and developing methodologies, just as the ones normally applied to ASICs. The lessons learned and conclusions to be applied in future FPGA developments for space projects are many, and cover a wide range of design-related subjects: documentation practices, human resources and tools, design methodology, radiation effects mitigation techniques at different design levels, verification and validation practices, system level design aspects, etc. The intrinsic complexity and vulnerability of FPGA designs to both designer and radiation induced mishaps are, unfortunately, equivalent to those of ASICs, and they should therefore never be underestimated.

ACKNOWLEDGMENTS

We would like to thank Ulrich Gageur, August Knoblauch, Bernard Gillot, Jan Johansson, Karl Engström, Dag Mortensen, Peter Spjuth, Stanley Mattsson, Eric Losman, Kimmo Myllyoja, Juha Kanerva, Jaakko Toivonen, Jyrki Laaksonen, Lars Sejr Krisensen, Kim Plauborg, Benny Ingvardsen, Jan van Casteren, Ralph de Marino, Reno Harboe Sørensen and Richard Creasey as members of the ESA FPGA Task Force and staff of the companies who have helped with the Rosetta FPGAs audit. They all have been very supportive and we thank them very much for their collaboration, their patience, and the high level of quality of their work.

References

[1] *ASIC Design and Manufacturing Requirements*, WDN/PS/700 Issue 2, October 1994, ESA. Available at ftp://ftp.estec.esa.nl/pub/vhdl/doc/DesignReq.pdf

[2] ECSS-Q-60-02, Space product assurance, ASIC Development. Currently under development, as Draft 4. Expected for public review by last quarter of 2002.

[3] www.actel.com

[4] Radiation Evaluation of ACTEL A14100A FPGA, ESA Contract N0 11356/95/NI/Fm, CCN-1WP-1A Final Report, FPGA Summary Report, SE/REP/0084/K Dec 1997