


FPGAs on Mars

Xilinx FPGAs have transitioned from a flight ASIC prototyping platform to playing integral roles in the Mars Exploration Rover Mission.

by David Ratter
Field Applications Engineer
Nu Horizons Electronics
dratter@nuhorizons.com

Selecting the correct components for any engineering project can be a critical and difficult choice. This is clearly true for engineers at the Jet Propulsion Laboratory (JPL) when they must select components used on high-stakes flight projects, and especially important on high-profile missions like the Mars Exploration Rover Mission.

On Mars there are no tow trucks or auto clubs to call if something stops working. So how did Xilinx FPGAs go from performing a predominantly flight ASIC prototyping role to being designed into and flown in projects like the Mars lander and rovers? And what does the future hold for Xilinx and JPL space missions?



The needs of JPL's design engineers were the main driving force behind the paradigm shift from ASIC prototype to flight-qualified part, as were the Xilinx testing, processing, and manufacturing flows for its radiation-tolerant FPGAs. Engineering needs demand meeting mission requirements, both from a functional/performance viewpoint and a time-to-working-product viewpoint. Xilinx FPGAs provide inherent design advantages to meet those needs: high gate densities, rich on-board architectural features, large I/O counts with multiple I/O standards, and the ability to be reprogrammed at any time.

The second reason for the transition was JPL's qualification of Xilinx radiation-tolerant FPGAs into more and more flight situations.

The net results of these efforts can easily be seen on the Mars Exploration Rovers. Inside of each rover (named Discovery and Spirit), two Virtex™ XQVR1000s served as the main brains that controlled the motors. Four Xilinx XQR4062XL devices in the 4000XL family controlled the Mars lander pyrotechnics, crucial to the successful multi-phase descent and landing procedure. Also evident, although not as visible, are the increasing number of future flight missions that JPL engineers are designing with Xilinx radiation-tolerant FPGAs.

In this article, we'll briefly document the parts qualification and design steps, along with the past, present, and future of FPGA-based flight opportunities at JPL.

Flight Considerations


There are three steps that must be accomplished before any part can be used in a flight application for JPL. The first is a general flight approval for a part. The second can be referred to as mission-specific approval. The third is additional design requirements for flight-based semiconductors.

General Flight Approvals

JPL must give general flight approval before any part can be used for a flight application. This process requires that the manufacturer perform numerous additional processing, testing, and quality steps

over and above the normal commercial processing steps.

JPL also meticulously examines device statistical and quality data that is constantly updated by the semiconductor manufacturers. They also examine additional parameters in this phase, including temperature considerations and packaging materials as well as semiconductor characteristics (for example, are there voids or contaminants that in zero gravity could migrate and cause failures?). Only after JPL engineers have conducted exhaustive research and analysis do they approve parts for flight use.



JPL Radiation Effects Group scientists work directly with Xilinx to unify and continually improve the testing, processing, and manufacturing steps used for Xilinx radiation-tolerant parts. This close customer/manufacturer relationship has resulted in a much superior radiation-tolerant product from Xilinx, and for JPL, a high-reliability manufacturing process gives them the utmost confidence.

Specific Flight Approvals

Even after a part or parts has general flight approval, it still must receive mission-specific flight approval. Mission-specific approval is exactly what the term implies: JPL scientists and engineers review the mission-specific environments that the parts will encounter. This includes a detailed risk assessment.

JPL takes into account all aspects of the flight to predict what the part(s) will face during the mission's lifetime. Some of these

parameters might include the number of temperature cycles, total ionizing dose, and predicted rate of radiation exposure. It is possible that parts with general flight approval will not get mission-specific flight approval.

Specific Flight Design Considerations

Only after a part has met the above two criteria can it be used in a flight mission. The design process entails incorporating space-specific flight design requirements that include, but are not limited to, the following single-event phenomena:

- Single-event latch-up (SEL)
- Single-event upsets (SEU)
- Single-event transients (SET)
- Single-event functional interrupts (SEFIs)

In the case of Xilinx radiation-tolerant FPGAs, all single-event phenomena are taken into account either through the radiation-tolerant manufacturing and processing steps or through well-documented design practices:

- The epitaxial layer of Xilinx radiation-tolerant FPGAs eliminates SELs.
- Triple-mode redundancy (commonly referred to as TMR) mitigates SEUs and SETs.
- "Scrubbing," or reprogramming the FPGA, takes care of SEFIs and the accumulation of SEUs.

Let's discuss the latter two design features in more detail.

In simple terms, a design with TMR requires three sets of key logic elements, with a voting structure that allows only the majority decision to propagate through the circuit. The theory is that statistically you are going to get an SEU over some time period. When this upset occurs, it will disrupt a single element (net, route, or bit). When this happens and the element it disrupts is being used, the other two "correct" elements will have the correct value; thus, the correct value will be passed out of the circuit. This is especially important in circuits that use feedback, such as counters and state machines.

Numerous approaches can be taken with respect to scrubbing, from simply reprogramming the FPGA to partial reconfiguration. The simplest method of scrubbing is to completely reprogram the FPGA at some periodic rate (typically 1/10 the calculated upset rate).

For example, if an SEU will occur once every 10 days, then you would reprogram the FPGA every day. However, when you reprogram the FPGA it is not operational during that reprogram time (on the order of micro to milliseconds). For situations that cannot tolerate that type of interruption, partial reconfiguration is available. This technique allows the FPGA to be reprogrammed while still operational.

The Past

Although the both the Discovery and Spirit rovers are still on the surface of Mars and active, they were launched in June and July 2003 and landed on Mars in January 2004. This means, obviously, that the engineering was completed in the past.

As stated earlier, the Mars Exploration Rovers used XQVR1000 devices and the lander used XQR4062XLs. The XQR4062XLs were used during the descent and landing of the rovers on the surface of Mars, while the XQVR1000s were used to control all of the brushed DC and stepper motors for the wheels, steering, antennas, camera, and other instruments on the rovers themselves.

Both the XQR4062XL and XQVR1000 designs used TMR for SEU and SET mitigation as well as scrubbing. JPL engineers achieved TMR in their designs by analyzing the design for feedback nets and other low-level design details (as detailed in Xilinx application note XAPP197), and then inserted or replaced logic with TMR library elements. After the designs were functionally implemented, the engineers went back through the design and inserted the TMR logic where necessary and made other space-specific design changes.

Due to the critical nature of the XQR4062XLs role, FPGA redundancy was utilized to mitigate all single-event phenomena. The scrubbing technique employed was a complete reconfiguration.



Animation by Dan Maas, Maas Digital LLC (c) 2002 Cornell University. All rights reserved. This work was performed for the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the United States Government under Prime Contract # NAS7-1407 between the California Institute of Technology and NASA. Copyright and other rights in the design drawings of the Mars Exploration Rover are held by the California Institute of Technology (Caltech)/ Jet Propulsion Laboratory (JPL). Use of the MER design has been provided to Cornell courtesy of NASA, JPL, and Caltech.

When an upset condition was detected, the FPGA in question was completely reprogrammed, while the redundant FPGAs remained functional. On the other hand, because of its non-time critical nature, when a fault condition was detected on one of the XQVR1000s, the rover was temporarily halted, the FPGA was reprogrammed, and the rover went back on its merry way.

All of the mitigation techniques on the Mars Exploration Rovers worked exactly as designed by JPL engineers. During the seven-month voyage from Earth to Mars, the Xilinx radiation-tolerant FPGAs on the lander were “left on.” During that time JPL collected data of interest, including the upset rate. The upset rates predicted for the FPGAs by JPL matched almost exactly the actual upset rates observed. Also, the upset detection and mitigation techniques implemented on the FPGAs performed their functions flawlessly and allowed for robust and reliable operation.

The Present

JPL is currently working on flight designs with both Virtex-II™, the latest Xilinx radiation-tolerant family, as well as Virtex radiation-tolerant FPGAs. These new missions will fly in the next two to five years and are becoming more and more sophisticated in both mission electronic requirements and design implementation. These current projects more fully utilize the FPGA's inherent benefits.

JPL is particularly interested in the ability to update or revise designs while the spacecraft is either in flight to its final destination or already there. This will allow engineers to implement algorithm enhancements after the spacecraft has left Earth, enabling them to constantly improve design performance during a long mission timeline. This, coupled with partial reconfiguration, will allow an FPGA-based design to have one portion of its design upgraded while the rest of the design remains completely operational.

There have been some breakthroughs in the area of single event mitigation and correction. Xilinx, in partnership with Sandia Labs, recently produced a TMR tool that

will XTMR a design. The XTMR tool takes in a synthesized netlist, analyzes it, and applies the appropriate TMR measures and space-specific design modifications to produce a final XTMR netlist. Not only does this guarantee a much more robust TMR design, it also takes what used to take an engineer days or weeks to perform and reduces the process to a matter of minutes.

The Future

Who knows that the future has in store? The new missions have more demanding requirements: more speed and more integration, with an ongoing goal of less space and less weight. On the Xilinx front, each subsequent family of radiation-tolerant FPGAs (like Virtex-II Pro™ devices) will provide more integration, more architectural features, and more capabilities.

The need for more integration, speed, and reduced space and weight goes hand in hand with continually improving radiation-tolerant FPGAs. Xilinx advances in FPGA technology and the improvements in radiation testing and processing made possible by their relationship with JPL come together to spell success.

Conclusion

The collaborative efforts between JPL, the Xilinx aerospace and defense team, and local support (provided by the manufacturer's representative, Norcomp SC, and Nu Horizons Electronics Distribution) helped pave the way for Xilinx to go from a predominantly ASIC prototyping role to become key components in the successful design and implementation of the Mars Exploration Rovers.

Current JPL flight projects that will launch in the years to come are already being designed using the latest Virtex and Virtex-II Xilinx high-reliability FPGAs. And the continual release of bigger, faster, and better Xilinx radiation-tolerant families means that with Xilinx and JPL, not even the sky is the limit.

For more information, please visit <http://marsrovers.jpl.nasa.gov/home/>, www.xilinx.com/esp/mil_aero/index.htm, www.norcompsc.com, and www.nuhorizons.com. ❧