

# Taking Designs to New Heights with Space-Grade Virtex-4QV FPGAs

Special processing, packaging, and SEU mitigation techniques enable designers to perform system-level integration with reconfigurable FPGAs in space-based systems.

by Greg Lara  
Virtex Marketing Manager  
Xilinx, Inc.  
[greg.lara@xilinx.com](mailto:greg.lara@xilinx.com)

Carl Carmichael  
Manager, Space Product Solutions  
Xilinx, Inc.  
[carl.carmichael@xilinx.com](mailto:carl.carmichael@xilinx.com)

Gary Swift  
Radiation Effects Specialist  
Xilinx, Inc.  
[gary.swift@xilinx.com](mailto:gary.swift@xilinx.com)

Designing high-performance electronic systems for space presents designers with a slew of unique challenges. Designers must meet tight power and weight budgets while under enormous pressure to get designs completed on time and with zero flaws.

Companies and government agencies set rocket launch schedules years in advance of actual launch dates. Missing a launch date can delay programs until, literally, planets align; some delays can even lead to project cancellations.

But while space applications have power, weight, and time-to-launch challenges, by far the most stringent requirement is ultra-high reliability. If a component fails in space, you can't send the Geek Squad out to fix it (unless you have astronauts on board).

To make things worse, as spacecraft rise to higher altitudes – and especially as they move beyond the protective blanket of Earth's atmosphere – they are bombarded by increasing levels of radiation that can upset or even disable delicate electronic components.

As a viable solution to these challenges, you can try to integrate functions into an ASIC, which can meet your performance, capacity, and power requirements. But designing an ASIC is increasingly complex, and if the design requires new functionality or if you find a bug late in the design cycle, you may have to re-spin the design or even create a new mask, causing you to miss launch windows.

Some silicon vendors offer radiation-hardened anti-fuse FPGAs, but those devices have limited capacity and performance, which means most of them are unsuitable for achieving SOC-type integration goals.

Increasingly, designers are turning to space-qualified Xilinx® Virtex™ FPGAs, which allow them to achieve performance, power, and integration targets while enjoying flexibility and productivity benefits previously available only in the commercial arena.

Let's examine our latest space-grade offering, the Virtex-4QV FPGA family, and some design techniques and tools that you can use in your design projects to tailor the devices for your space deployment requirements.

**The Xilinx Space Heritage**

Xilinx has provided uninterrupted support to the aerospace and defense (A&D) community since 1989. Our company pioneered the application of SRAM-based FPGAs in high radiation environments, beginning with the XQR4036XL FPGA, which first flew in FedSat, an Australian satellite.

Quad-redundant XQR4062XL FPGAs performed mission-critical landing duties (pyro control) on the Mars Exploration Rover (MER) 2003 landers, which delivered the Spirit and Opportunity MERs to the surface of Mars. Virtex XQRV1000 FPGAs handle motor control functions on both MERs as they explore the Martian landscape.

To accelerate the adoption of reconfigurable FPGAs for aerospace applications, in 2002 Xilinx joined forces with the Jet Propulsion Laboratory, Los Alamos National Labs, Sandia, Aerospace Corp., and SEAKR Engineering to found the Single-Event Effects (SEE) Consortium. Fostering an environment of openness and collaboration, the SEE Consortium leverages the collective talents and financial resources of its member companies to perform exhaustive analysis of complex FPGAs.

The SEE Consortium subjects FPGAs to in-beam radiation testing and then reports peer-reviewed results to the industry. As a result of those efforts, Venus Express, the Mars Reconnaissance Orbiter, GRACE, OPTUS, TACSAT2, CIBOLA, and a number of classified programs have included Xilinx FPGAs. Many more Virtex FPGAs will soon be

orbiting overhead, as our customers have more than 50 new programs currently under development.

**Virtex-4QV FPGA Family Overview**

To bring the benefits of advanced FPGAs to designers of in-orbit, space-based, or extraterrestrial systems, Xilinx recently launched a space-grade variant of its Virtex-4 family. The new Virtex-4QV family of radiation-tolerant FPGAs meets the military's MIL-PRF-38535 Class V screening requirements and offers design groups a compelling alternative to ASIC and one-

time programmable logic technologies. The lineup comprises four of the configurations from the commercial Virtex-4 device family (see Table 1).

High-performance logic, processing, and DSP capabilities make Virtex-4QV FPGAs ideal for demanding applications such as video processing, radar, and encryption. In addition, their packet processing capabilities are enabling a new generation of networked space-based communication systems, such as the IRIS program, to put reconfigurable network routing in orbit.

		Virtex-4QV FPGAs Optimized for:			
		Logic	DSP	Embedded Processing	
Part Number		XQR4VLX200	XQR4VSX55	XQR4VFX60	XQR4VFX140
Logic Cells		200,448	55,296	56,880	142,128
CLB Flip-Flops		178,176	49,152	50,560	126,336
Maximum Distributed RAM (Kb)		1,392	384	395	987
Block RAM/FIFO w/ECC (18 Kb each)		336	320	232	552
Total Block RAM (Kb)		6,048	5,760	4,176	9,936
Digital Clock Managers (DCMs)		12	8	12	20
Maximum Single-Ended I/Os		960	640	576	896
DSP Slices		96	512	128	192
10/100/1000 Ethernet MAC Blocks		–	–	4	4
PowerPC Processor Blocks		–	–	2	2
Total Ionizing Dose (krad)		300	300	300	300
SEL Immunity (MeV-cm <sup>2</sup> /mg)		>125	>125	>125	>125
Package	Area	Available User I/Os			
<b>CFA Packages (CF): Flip-Chip Ceramic Column Grid Array (1.0-mm Column Spacing)</b>					
CF1140	35 x 35 mm		640		
CF1144	35 x 35 mm			576	
CF1509	40 x 40 mm	960			768

Table 1 – Xilinx 90-nm Virtex-4QV FPGA family

## The combination of the XTMR methodology and configuration memory scrubbing provides an elegant and cost-effective method to ensure reliable operation for almost all critical space applications.

To enable design teams to use reconfigurable FPGAs in space, Xilinx tackled multiple technical challenges:

- Make the CMOS silicon immune to latchup induced by strikes of heavy ions (a.k.a., “cosmic rays”)
- Create a high pin-count package that meets the unique requirements of space deployment
- Screen all devices to MIL-PRF-38535 Class V requirements
- Develop a range of upset mitigation techniques backed up by tools and IP

Ionizing particles in space can strike the electronic systems of a spacecraft and produce currents in a device’s silicon junctions that can induce latchup of the parasitic transistors inherent in all CMOS integrated circuits. Latchup is a significant concern for designers because it results in high current that can kill the device.

To provide immunity against destructive latchup, Xilinx fabricates radiation-tolerant Virtex-4QV FPGAs on special wafers that incorporate a thin, lightly doped epitaxial layer. In the event of a particle strike, the doped substrate effectively limits the length of the charge trail that can contribute to a current spike in active nodes (Figure 1).

For each Virtex-4QV device type, the SEE Consortium verifies latchup immunity at maximum  $V_{CC}$  and operating temperature, subjected to a heavy ion fluence exceeding  $1 \times 10^7$  particles/cm<sup>2</sup>, with an effective linear energy transfer (LET) exceeding 125 MeV-cm<sup>2</sup>/mg.

### Packaging Considerations

Space applications impose special requirements on device packaging with regard to thermal cycling and hermeticity. Xilinx has developed a column-grid array (CGA) surface-mount package that meets these requirements while supporting the high pin count of Virtex-4 FPGAs.

For reliable operation in the vacuum of space, semiconductor devices require packaging that can withstand the extreme temperature ranges and thermal cycles encountered outside the thermal conduction of the Earth’s atmosphere.

It isn’t uncommon for one side of a satellite to broil under the searing radiation of the sun, while the other side suffers the deepest chill of space. Components in an orbiting and spinning satellite can find themselves rapidly cycling between these temperature extremes. In such an environment, differences in the coefficients of thermal expansion between PCB and device pack-

ing point than the solder paste that companies typically use to attach the device to the circuit board. The solder columns flex to absorb the difference in expansion rates of the package and board to greatly extend the operating life of the device.

Hermetic packages encapsulate die, which electrically connect to the package pins through wire bonds. Virtex-4 FPGAs are high pin-count devices and thus can’t employ wire-bond packaging. Instead, they use flip-chip packaging, in which solder bumps on the die make direct contact with the package substrate.

For Virtex-4QV FPGAs, we reengineered the device’s flip-chip package to




Figure 1 – A mobile charge produced by a heavy ion strike can induce SEU in an FPGA.

age materials can lead to failures in the solder joints that bind the packaged device to its PCB.

To avoid this problem, we package our Virtex-4QV FPGAs with CGA packaging, which addresses the thermal cycling shortcomings of traditional surface-mount technologies. An enhancement of BGA technology, CGA packaging uses solder columns, rather than balls, made of a more flexible low-tin solder with a higher melt-

address the thermal cycle-induced stress where the die attaches to the package, as well as to meet the reliability requirements of a hermetic package. Through judicious choice of solder bump and underfill materials, lid material, and lid-attach adhesive, we made sure that Virtex-4QV devices satisfy JEDEC moisture sensitivity level 1, NASA outgassing requirements, and corrosion specs, plus first- and second-level thermal cycle reliability requirements.

To ensure that our Virtex-4QV FPGAs meet reliability standards for space deployment, Xilinx performs conformance certification of each manufacturing lot. Following military performance specifications, we perform a full test suite across the full military temperature range, including dynamic and static burn-in.

In addition, we perform full TCI/QCI screening to MIL-STD-883 test method 5004 (screening 100% of units) and 5005 (destructive stress testing) on a lot-by-lot basis. We subject every assembly lot to destructive physical analysis (DPA) to comply with MIL-STD-1580 requirement 16, which involves the construction analysis of die fabrication, package fabrication, and assembly.

**Design Techniques for Upset Mitigation**

In addition to making the silicon more resilient to radiation, Xilinx has developed a range of mitigation techniques that designers can employ to further reduce the likelihood that a charged particle strike will affect the operation of our FPGAs in high-radiation environments.

When developing systems that will operate in hostile radiation environments, you must consider the effects of charged particles such as heavy ions or reaction products from neutrons or protons. As these charged particles travel through the FPGA (Figure 1), they can alter the logic state of any static memory element, resulting in single-event upsets (SEUs). An SEU in the configuration memory array can have adverse effects on expected FPGA functionality.

By thoroughly analyzing your system’s operating environment and radiation-tolerance requirements, you can choose an approach that economically delivers the required reliability. For example, for a mission-critical application in a harsh environment, you could use a mitigation scheme that combines redundant FPGAs with radiation-hardened configuration management hardware.

On the other hand, consider an orbiting satellite carrying an image processing circuit that powers up only periodically to capture scenes of interest. If the statistical risk of an upset during that short window of operation is acceptably small, you might choose

to forego any special mitigation techniques, especially if the worst-case outcome is the corruption of single frame of image data. In this case, the standard configuration process of loading fresh configuration data on each power-up cycle could be a sufficient approach to fixing upsets.

Answering the following questions and understanding the trade-offs between these application requirements will guide you in selecting the appropriate mitigation scheme:

- Is the application error-tolerant?
- What is the expected operating time window?
- What is the FPGA performance requirement?
- What is the FPGA power requirement?
- How cost-sensitive is the application?
- Can the application withstand any possible FPGA downtime? If so, does the single-event functionality interrupt (SEFI) cross section specified in the data sheet meet the application requirement?

Based on your answers to these questions, you can employ a series of mitigation techniques and IP in your design to reduce the risks that a particle strike will affect the operation of Xilinx FPGAs when they are in high-radiation environments. Figure 2 provides an overview for a mitigation scheme selection based on application requirements and environmental constraints.

**Reconfiguration and Memory Scrubbing**

A complete FPGA reconfiguration (or, alternatively, power-cycling the FPGA) is the simplest form of mitigation, as it reloads the configuration memory contents. You can use the FPGA’s configuration memory readback function to determine whether the contents have been corrupted by an ion strike and then initiate reconfiguration. Although readback does not interfere with the device’s operation, a complete reconfiguration interrupts design function for anywhere from a fraction of a second to a few seconds, depending on clock rates and device size.

Because a given design uses only a fraction of the configuration memory cells in an FPGA, an SEU will likely have no effect on system operation. To avoid unnecessary interruptions caused by complete reconfiguration, you can employ the partial reconfiguration capabilities built into Virtex-4QV FPGAs to reset corrupted bits. The FPGA performs this process (called memory scrubbing) automatically, without hindering the operation of the system.

**Hardware Redundancy**

Designers can also fall back on a well-known technique called hardware redundancy, which is critical for high-availability systems – those that can only have minimal unplanned downtime. By employing hardware redundancy, you create two copies of a system so that one copy can keep running

Data Criticality		Low		High	
		No		Yes	
SEU Rate	System Duty Cycle	Low	Hours	No Mitigation	XTMR
		High	Weeks	Configuration Management (Scrubbing)	Configuration Management Plus XTMR
			Years		
			Continuous		

Figure 2 – Mitigation scheme matrix

if the other fails. However, the overall system may not be able to distinguish which of the two systems has failed.

To resolve this dilemma, you can add a third copy of the system – known as triplication – to act as a tie breaker. Logic triplication, plus the addition of special voter circuits at key locations in the design, enables your system to continue running even in the presence of an upset.

You'll find that the ample logic and routing resources available in Virtex FPGAs enable you to implement logic triplication within a single device. This capability is the cornerstone of the Xilinx Triple Modular Redundancy (XTMR) methodology. We offer software called the TMRTool that automates the process of building triple redundancy into your designs to simplify this tedious and poten-

design to periodically scrub the circuits and prevent errors from accumulating. Indeed, the combination of the XTMR methodology and configuration memory scrubbing provides an elegant and cost-effective method to ensure reliable operation for almost all critical space applications.

There are extremely rare cases where the combination of XTMR and scrubbing is not enough. For example, if several ions simultaneously (during a single scrub cycle) strike the device's configuration memory and change the logic function on two signals going into the same voter circuit, the voter circuit will misinterpret the false signal as a valid signal and propagate the false signal through the rest of the circuit.

However, for a reasonable scrub rate, even during the most severe space radia-

### Redundant Devices

Depending on the mission criticality of your design and the orbit in which it will be operating, you might choose to use an even more robust mitigation scheme. For example, you can expect that a triplicated design with full-blown configuration management in geosynchronous orbit will experience only one SEFI per device-century. Thus, if the FPGA's SEFI detection and device reconfiguration requires only one second, the result is one second of device outage per 3 billion seconds of operation.

In the unlikely event that your system requires even higher availability (must have extremely short or no downtime), you should employ a design that involves redundant FPGAs plus dedicated configuration management, as well as voting on

Mitigation Scheme	Mitigation Strength	Board Layout Complexity	Ease in Meeting Timing Constraints	Power Consumption	Component Cost
No Mitigation (Power Cycling)	Weak	Low	Normal	Typical	Low
Configuration Management (Scrubbing)	Medium	Low	Normal	Typical	Low
XTMR	Medium	High	Reduced	~3x Typical	Medium
XTMR + Configuration Management	Strong	High	Reduced	~3x Typical	Medium
Redundant Devices + Configuration Management	Strongest	Medium	Normal	~2-4x Typical	High

Table 2 – Performance overview of mitigation schemes

tially error-prone task. Used in conjunction with the ISE™ software design tool suite, the TMRTool will partially or fully triplicate a design, insert voters, and synchronize feedback path loops. TMRTool also enables the insertion of custom-triplicated design modules.

XTMR by itself is not completely fool-proof. As systems fly through space, they will accumulate errors that will adversely affect their operation. During system operation, the probability of system error increases exponentially as upsets accumulate in the system and defeat the device's voter circuits.

To address error accumulation before it becomes a problem, use a combination of XTMR and memory scrubbing in the

tion conditions (a 90% worst-case solar flare), the probability of this happening is lower than the likelihood of a SEFI – the upset of a control bit that results in the persistent interference of the normal operation of the circuit.

A SEFI typically results in the loss of one or more of the following: configuration capability, power on reset ability, JTAG functionality, a region of configuration memory, or the entire configuration. Recovering from a SEFI requires reconfiguring your FPGA.

You can improve scrubbing by adding the ability to detect SEFIs. Thus, full-blown configuration management combines transparent readback and scrubbing plus SEFI detection.

the outputs of the FPGAs. Table 2 shows the trade-offs that can guide your choice of mitigation schemes.

Xilinx takes great pride in introducing space-grade variants of our commercial devices to bring advanced functionality to designers in the aerospace and defense industries. This functionality allows you to meet payload and power requirements and integrate more advanced functionality into your designs. It also gives you the unique flexibility to make changes to your designs late in the development process and even in space, if required. And as you plan to expand the possibilities of space exploration, Xilinx continues its commitment to helping you make those possibilities a reality. 