A Guide to the Benefits of the Lattice Nexus FPGA Platform for Mission-Critical Applications

A Lattice Semiconductor White Paper.

January 2021

Learn more:

www.latticesemi.com

Contact us online:

www.latticesemi.com/contact
www.latticesemi.com/buy

TABLE OF CONTENTS

There is growing demand for mission-critical applications in the industrial, automotive, communications, aerospace, and defense markets. Today, the Lattice Nexus™ Platform provides a true differentiator for FPGAs destined for use in mission-critical systems.

Introduction (MPUs vs. FPGAs)

Today's mission-critical systems can require significant computing power. One well-known computing solution is to use microprocessor units (MPUs), such as those found in PCs and workstations. Although these processors may appear to be powerful, in reality all they are doing is performing simple operations like adding two numbers together or comparing two numbers to see which is larger. Similarly, although they may appear to be fast, this is because their system clocks are running sequentially at 2.4 GHz or higher.

The real issue is that, while MPUs are good for performing decision-making tasks, they can be inefficient when it comes to performing many data processing assignments. As a result, MPUs tend to consume large amounts of power and generate a lot of heat while performing their duties.

A more efficient way to perform signal and data processing in embedded applications is through the use of field-programmable gate arrays (FPGAs). FPGAs benefit from an inherent parallel architecture to run data processing operations in parallel with low latency. As was noted in the column Fundamentals of FPGAs: What Are FPGAs and Why Are They Needed?: "At the heart of any FPGA . . . is its programmable fabric, which is presented as an array of programmable logic blocks. Each of these logic blocks contains a collection of elements – including a look-up table (LUT), a multiplexer, and a register – all of which can be configured (programmed) to act as required."

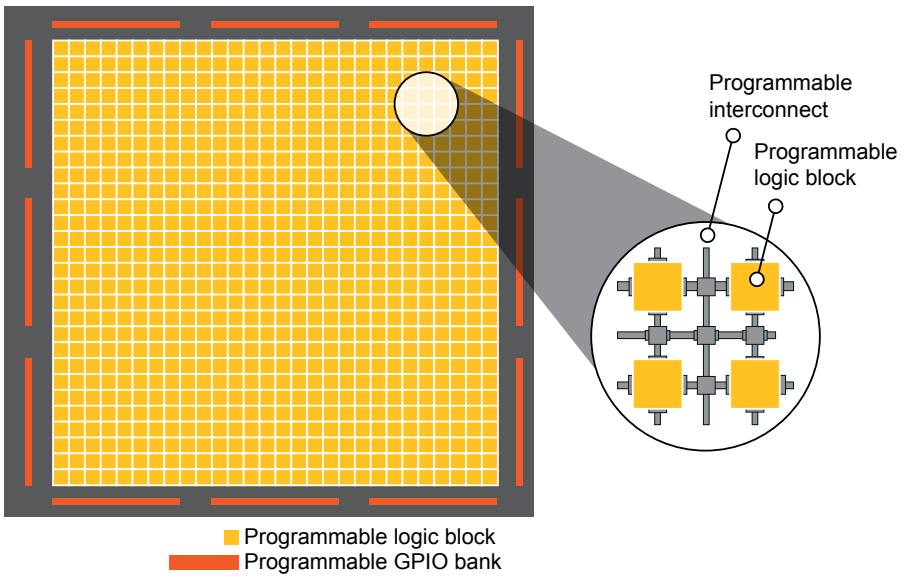

Figure 1. Generic FPGA programmable fabric (Image source: Max Maxfield).

One way to visualize this is as "islands" of programmable logic in a "sea" of programmable interconnect. Groups of programmable logic blocks can be configured to perform desired logical functions, while the programmable interconnect can be used to connect the logic blocks to each other and to the primary general-purpose input/outputs (GPIOs).

It is important to note that the above is an extreme simplification. In addition to their programmable fabric, FPGAs can contain memory blocks equating to megabits of RAM along with thousands of digital signal processing (DSP) functions. Similarly, in addition to their general purpose I/Os, FPGAs can include high-speed SERDES blocks that can support gigabit serial interfaces, along with high-speed interfaces to external memories.

One of the FPGA's main claims-to-fame is that the programmable fabric can be configured to perform appropriate data processing algorithms in a massively parallel fashion, orders of magnitude faster than MPUs while consuming a fraction of the power. Furthermore, unlike a SoC, whose algorithms are "frozen in silicon," the design in an FPGA can be reconfigured "on-the-fly." Unlike SOCs and ASICs, FPGAs don't contain any user IP when they are manufactured as the IP is programmed by the end user. This important value enables FPGAs to exist in unprotected supply chains without the potential of user IP theft or tampering. As a result of all this, FPGAs are ideal for today's mission-critical systems, from industrial robots to radars and guidance systems to communications infrastructure equipment.

Challenges with Radiation

In order to increase capacity, boost performance, reduce power consumption, and lower costs, each new generation of silicon chips features smaller and smaller transistors. Today, the size of the structures created in the silicon is measured in a few tens of nanometers (nm), where one nanometer is a thousandth of a millionth of a meter. These structures are so small that they can be affected by the levels of radiation found on Earth.

Furthermore, systems developed for mission-critical applications are oftentimes deployed in environments that experience significantly higher levels of sustained radiation exposure, including high altitudes and even space.

Two classes of radiation effects that are most critical to systems intended for mission-critical applications are single event effects (SEEs) and total ionizing dose (TID).

The term SEE refers to an effect resulting from a single ionizing particle (electron, proton, ion, photon...) that induces an immediate response in an integrated circuit. Handling SEEs requires resilience to radiation at the current time of the event. By comparison, the TID results in degradation of the semiconductor crystal lattice due to an accumulation of radiation exposure over time. Typical TID effects include shifts in the switching thresholds of transistors, increased current leakage, decreased performance, and -- ultimately -- functional failures. Thus, addressing the TID requires resilience to radiation over longer periods of time.

The term single event upset (SEU) refers to a SEE that strikes a sensitive node in a micro-electric circuit causing a change of state. For example, an SEU could cause a register element or a memory cell to flip from a logic 0 to a logic 1, or vice versa. Unlike problems caused by the TID, an SEU is designated as being a "soft error" because it can be corrected.

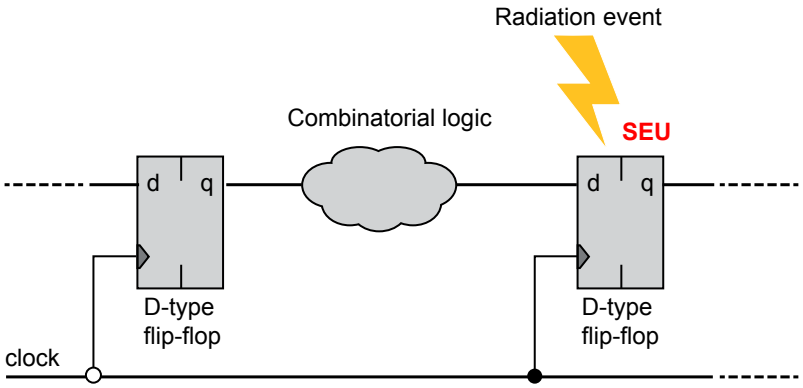

Figure 2. Single event upset (SEU) in the sequential logic (Source: Max Maxfield)

Unfortunately, ever shrinking fabrication processes are also resulting in a trend toward increased multiple cell upsets (MCUs). What this means is that, since the structures in the silicon are now so close together, an SEU may actually upset multiple memory elements. Furthermore, the term multiple bit upset (MBU) refers to a MCU that occurs within the same data word or frame, which may negatively impact the system's ability to correct the error.

Another form of SEE is a single event transient (SET), which refers to a radiation event impacting a portion of combinatorial logic resulting in a pulse (a.k.a. glitch or spike).

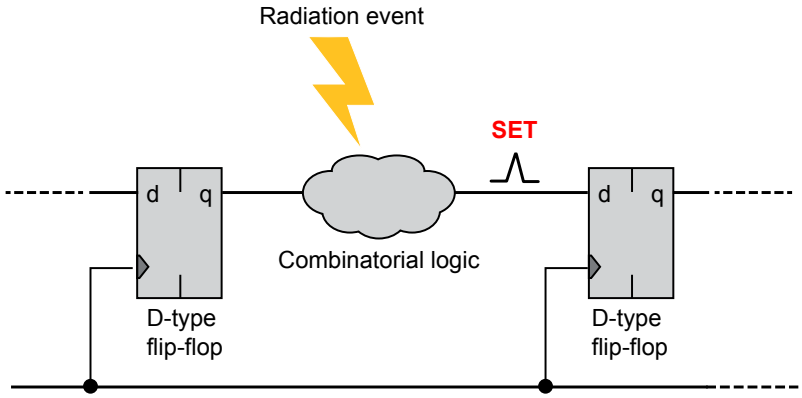

Figure 3. Single event transient (SET) in the combinatorial logic (Source: Max Maxfield)

On its own, a SET isn't too problematic because it will typically have died away before it's seen by the rest of the system. Having said this, if the SET occurs at precisely the wrong time, it could be clocked into a register element or memory cell, at which point it becomes an SEU.

Another potential problem is single event latch-up (SEL), in which a SEE results in the inadvertent creation of a low-impedance path (effectively a short-circuit) between the power and ground rails in a CMOS circuit. If such a state should occur, the device needs to be immediately power cycled (turned off and on again) in order to prevent serious damage. It probably goes without saying that having to power cycle a piece of mission-critical equipment may be problematic in many situations.

There is a further consideration in the case of FPGAs, because -- in addition to their register elements and RAM cells -- they also contain configuration cells, which are used to configure the programmable logic blocks, programmable interconnect, and programmable general-purpose I/O. Different FPGAs employ different configuration cell technologies.

Lattice offers a new class of SRAM-based devices based on the Lattice Nexus FPGA development platform, which provides a true differentiator for FPGAs destined for use in state-of-the-art systems performing mission-critical applications.

Introducing the Nexus Platform

Based on a 28 nm fully depleted silicon-on-insulator (FD-SOI) process, Lattice's Nexus platform provides a huge differentiator in the FPGA market.

The FD-SOI process immediately conveys two significant advantages. First, because it's deployed fully depleted, it's inherently radiation resilient. Quite apart from anything else, FD-SOI is inherently immune to single event latch-up conditions, which means there is no down-time in mission-critical situations that would normally demand a power cycle to exit the latch-up state.
The second noteworthy advantage is flexibility. By varying the biasing of the substrate, users can decide whether they wish to run for high-performance (HP) or low-power (LP). Furthermore, the system can switch back and forth in real-time under program control.

For example, as was previously discussed, an SEU occurs when a radiation event in the form of an energetic particle travels through a register element or a memory cell, ionizing the semiconductor material (charge generation) and creating a brief current pulse (charge collection). This current pulse may be sufficient to disturb the stored value.

Consider the representation of a bulk CMOS process as depicted in Figure 4(a). The radiation event (red arrow) generates an ionized path through the silicon leaving a trail of +/- charges in its wake. These charges subsequently collect at the incident node (blue arrows).

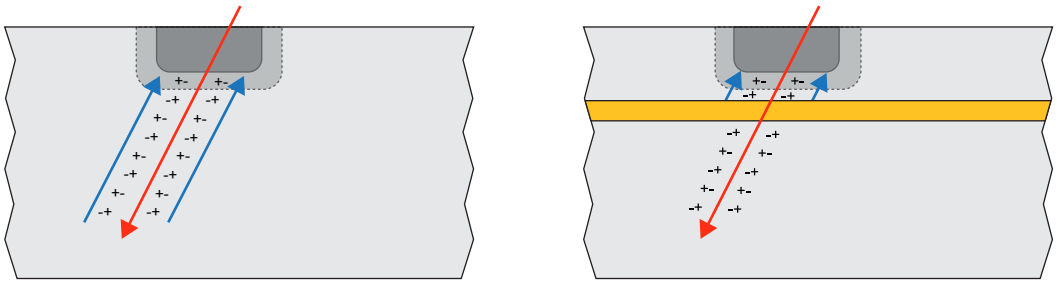

(a) Bulk CMOS Process (b) FDSOI Process

Figure 4. Comparison of the effects of an SEU on a bulk CMOS process(a) versus an FD-SOI process (b).

By comparison, consider the FD-SOI process as depicted in Figure 4(b). In this case, the layer of buried oxide (yellow) isolates the node from the bulk of the generated charge, because any charge below the oxide layer is unable to collect at the sensitive node. Less charge means smaller transient current pulses, which are less likely to upset register elements or memory cells.

Another potential problem noted earlier is that of MCUs and MBUs, in which a single particle is able to upset multiple memory elements. Consider the representation of a bulk CMOS process as depicted in Figure 5(a). Once again, the radiation event (red arrow) generates an ionized path through the silicon leaving a trail of +/- charges in its wake. In addition to collecting at the incident node (blue arrows), these charges may also collect at an adjacent node (purple arrows), possibly resulting in an MCU or MBU.

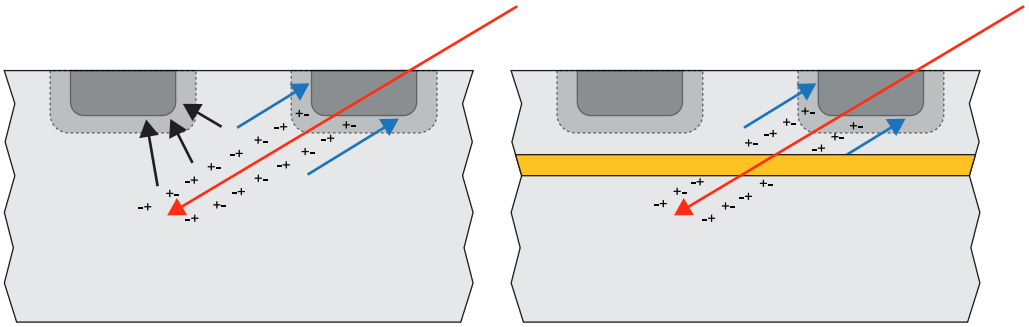

(a) Bulk CMOS Process (b) FDSOI Process

Figure 5. A bulk CMOS process (a) is susceptible to MCUs and MBUs, while an FD-SOI process (b) protects against these upsets.

By comparison, consider the FD-SOI process depicted in Figure 5(b). In addition to isolating the incident node from the bulk of the generated charge, the layer of buried oxide (yellow) dramatically shrinks the sensitive area of each cell, making it much more difficult for a single particle track to affect multiple bits, thereby resulting in a substantial decrease in the generation of MCUs and MBUs.

Designers of mission-critical and safety-critical systems use the concept of the FIT (failures in time) rate. The FIT rate of a device is the number of failures that can be expected in one billion (109) device-hours of operation, (e.g., 1 device for a billion hours, 1000 devices for 1 million hours each, 1 million devices for 1000 hours each, or some other combination thereof).

In the case of a typical FPGA implemented in a bulk CMOS process at the 28 nm technology node, the FIT rate is about 100. By comparison, in the case of a Lattice FPGA implemented using the Nexus Platform's FD-SOI process at the 28 nm technology node, the FIT rate is only 1. This means that, right from the get-go, the Nexus Platform provides an improvement in FIT rate of two orders of magnitude. The Nexus platform FPGAs such as the Lattice CrossLink™-NX provide a detailed SEU characterization report which can be used for estimating failure rates due to radiation effects.

On its own, this makes it possible to truly differentiate FPGAs based on this technology for use in applications such as medical, automotive, defense, and aerospace. But this is only the start, because Lattice's mission is to bring the FIT rate down to virtually zero, and this is achieved by augmenting the FD-SOI process with additional technologies as discussed below.

Lattice Nexus FPGAs

The term error-correcting code (ECC) memory refers to type of data storage implemented in such a way that it can detect and correct any internal data corruption, such as that caused by a radiation event, for example. Remembering that SEUs are referred to as "soft errors," this leads to the concepts of soft error detection (SED) and soft error correction (SEC).

Meanwhile, the term memory scrubbing refers to the process of reading from each memory location, correcting bit errors (if any) with an error-correcting code, and then writing the corrected data back to the same location. Memory scrubbing is often used in the case of mission-critical and safety-critical systems, and also in systems that are subject to high radiation environments.

Designers typically have to implement memory scrubbing functionality themselves, thereby consuming valuable programmable logic resources. By comparison, Nexus FPGAs include dedicated intellectual property (IP) blocks that automatically perform ECC-based memory scrubbing as a background process.

Furthermore, Nexus FPGAs also have an SED/SEC block built into the configuration memory in order to facilitate rapid detection and correction of errors on a frame-by-frame basis without the need for external circuitry. Although such errors are rare, they are theoretically possible, but if a radiation event were to somehow cause a configuration cell to flip state, this dedicated IP would flip it back again.

The end result is that no uncorrectable SEUs has been observed in Lattice Nexus FPGAs. Even though the underlying process provides a theoretical FIT rate of 1, which means one could occasionally experience a bad bit, that corrupted bit will be made good again almost instantaneously.

Having said this, two situations do exist where errors may be uncorrectable by the internal SED/SEC engine. The first case is that of multiple independent SEUs in which two or more particles randomly upset multiple bits within the same data frame. The second is an MBU whereby a single particle manages to upset two or more bits within the same data frame.

With regard to MCUs and MBUs, during radiation characterization of the Nexus platform, special care was taken to monitor for these effects. These tests confirmed the technology advantages described above, showing a very infrequent rate of multiple cells being affected by a single particle. Furthermore, due to Lattice's memory array design, all observed MCUs occurred in different data frames, thereby making them correctable by the SED/SEC engine.

In addition to testing Nexus FPGAs using real radiation sources, these devices include mechanisms to allow system developers to inject their own simulated radiation events. In fact, developers can inject single-bit errors and multi-bit errors, where such errors can be injected synchronously or asynchronously. Using these mechanisms, developers can verify that the memory scrubbing functionality and SED/SEC engine kicks in and the device keeps running and generating correct results, thereby ensuring that the device works as promised and the design works as expected, even in harsh, radiation-intense environments.

Summary

Like all electronic components, FPGAs may be negatively affected by radiation events, whose effects are becoming more pronounced as the structures on silicon chips shrink in size. Lattice's Nexus Platform is based on the 28 nm FD-SOI process, which allows users to decide whether to run for high-performance or low-power consumption, and to make such decisions in real-time under program control. Furthermore, the FD-SOI process is inherently radiation-tolerant, thereby providing a FIT rate of 1, which is an improvement of two orders of magnitude as compared to standard CMOS FPGAs implemented at the same technology node.

For all these reasons, Lattice Nexus FPGAs are ideal for mission-critical and safety-critical applications for the commercial, industrial, communications, defense, aerospace, and automotive applications.