



XAPP864 (v2.0) April 1, 2010

SEU Strategies for Virtex-5 Devices

Author: Ken Chapman

Summary

Xilinx® devices are designed to have an inherently low susceptibility to single event upsets (SEUs). This application note provides a substantial discussion of strategies and representative calculations for handling SEUs with an emphasis on reliability when addressing these low probability events.

This application note introduces an SEU controller macro that can be included in any Virtex®-5 FPGA design to implement an SEU detection and correction scheme. This application note is also a useful tool in evaluating the different methods for addressing SEUs.

Due to the infrequent and unpredictable nature of real SEUs, small scale testing of their effects and system verification is impractical. For this reason, the SEU controller macro can emulate an SEU by deliberately injecting an error into the FPGA configuration to confirm the subsequent detection and correction. Injection of errors can also be used to assess SEU mitigation circuits implemented in a design and to verify the claims made in this application note.

This application note focuses on the Virtex-5 family, however, much of the discussion is also applicable to the Spartan®-6, Virtex-6, and Extended Spartan-3A families.

Introduction

SEUs have the potential to affect most digital electronic circuits. Xilinx takes this issue seriously, and by improving the underlying technology, Xilinx devices experience very low levels of SEUs. Xilinx also recognizes that SEUs are unavoidable within commercial and practical constraints, so the company provides built-in SEU detection in the Virtex-5, Spartan-6, Virtex-6, and Extended Spartan-3A families to simplify and improve system design.

A Note for Experienced Users

The study of SEU behavior and the impact on Xilinx devices (the Rosetta Experiment) is an ongoing project that is providing an improved understanding of the subject. For this reason, readers familiar with SEU issues, especially with regards to Xilinx FPGAs, are advised to study the whole of this application note and review their choice of strategy when dealing with this phenomena. Xilinx continues to take the subject of SEUs very seriously and therefore the dynamics of this subject are changing. The good news is that the situation has improved, but even positive changes should initiate a review of strategy. If you have previously used techniques to handle SEUs in Virtex-4 devices and are now working with Virtex-5 devices, there are valuable differences due to the improved capabilities of the newer devices.

Applications requiring the highest level of reliability should use an appropriate SEU risk mitigation scheme. This application note discusses the factors that should be considered and presents strategies that can be used with Virtex-5 devices.

The Rosetta Experiment

The Rosetta experiment is an ongoing project at Xilinx that collects real measurements of SEUs and applies the knowledge gained when engineering each new product. The test data for Virtex-5 and Extended Spartan-3A families confirms that these devices have significantly lower susceptibility to SEUs than their predecessors.

Details of the Rosetta experiment are available in [WP286](#), *Continuing Experiments of Atmospheric Neutron Effects on Deep Submicron Integrated Circuits*. Calculations in this application note are based on values reported in [UG116](#), *Device Reliability Report*.

Risk Assessment and Specification

As a starting point, a target specification for system reliability should highlight critical sections of a design and provide a guideline value for the required reliability of the FPGA design. This is typically expressed as the failures in time (FIT) rate, which is the number of failures that can be expected in 10^9 hours (approximately 114,155 years) or the mean time between failure (MTBF).

Based on the SEU FIT data reported in [UG116](#), *Device Reliability Report*, a Virtex-5 device shows a nominal 131 FIT/Mb for the configuration cells with a 90% confidence range of -20% to $+26\%$. For example, the XC5VLX50T device on the ML505 development board has approximately 11.37 Mb of configuration cells and therefore has a nominal susceptibility of 1,489 FIT or an MTBF of 77 years ($114,155 \text{ years}/1,489 \text{ FIT} = 77 \text{ years MTBF}$). The device configuration FIT has a 90% confidence range of 1,192–1877 FIT or an MTBF of 61–96 years.

The number of configuration cells, which are of concern, can be determined from [Table 1](#). The values given in the column [Clock Cycles per Readback CRC Scan^{\(1\)}](#) relate very closely to the number of 32-bit words of configuration that are verified during each Readback CRC scan of the device. Therefore, multiplying these values by 32 and dividing by 1,000,000 provides the number of configuration cells in Mb for each device. For example, using a XC5VLX50T device yields $355,190 \times 32/1,000,000 = 11.37\text{Mb}$. These values are smaller than the size of the original configuration image because only the cells that should remain unchanged are relevant and all block memory contents are no longer included.

Table 1: Readback CRC Clock Cycles and Scan Times

Device	Clock Cycles per Readback CRC Scan ⁽¹⁾	Readback CRC Scan Time at 60 MHz ⁽²⁾ (mS)	Longest Readback CRC Scan Time using ConfigRate = 38 ⁽³⁾ (mS)
XC5VLX30	226,122	3.77	11.90
XC5VLX50	339,200	5.65	17.85
XC5VLX85	573,392	9.56	30.18
XC5VLX110	764,534	12.74	40.24
XC5VLX155	1,062,030	17.70	55.90
XC5VLX220	1,450,382	24.17	76.34
XC5VLX330	2,175,590	36.26	114.50
XC5VLX20T	152,609	2.54	8.03
XC5VLX30T	236,782	3.95	12.46
XC5VLX50T	355,190	5.92	18.69
XC5VLX85T	589,382	9.82	31.02
XC5VLX110T	785,854	13.10	41.36
XC5VLX155T	1,083,350	18.06	57.02

Table 1: Readback CRC Clock Cycles and Scan Times (Cont'd)

Device	Clock Cycles per Readback CRC Scan ⁽¹⁾	Readback CRC Scan Time at 60 MHz ⁽²⁾ (mS)	Longest Readback CRC Scan Time using ConfigRate = 38 ⁽³⁾ (mS)
XC5VLX220T	1,471,702	24.53	77.46
XC5VLX330T	2,207,570	36.79	116.19
XC5VSX35T	300,578	5.01	15.82
XC5VSX50T	450,884	7.51	23.73
XC5VSX95T	805,862	13.43	42.41
XC5VSX240T	1,852,838	30.88	97.52
XC5VTX150T	1,069,656	17.83	56.30
XC5VTX240T	1,663,418	27.72	87.55
XC5VFX30T	305,826	5.10	16.10
XC5VFX70T	611,686	10.19	32.19
XC5VFX100T	880,646	14.68	46.35
XC5VFX130T	1,100,816	18.35	57.94
XC5VFX200T	1,570,922	26.18	82.68

Notes:

1. The number of clock cycles does not easily correlate with the configuration image sizes shown in the data sheet. This is due to the highly optimized nature of the readback CRC scan, which avoids variable user data (for example block RAM contents).
2. Maximum external clock rate for readback CRC circuit operation is 60 MHz.
3. ConfigRate = 38 is the highest value guaranteed to remain below 60 MHz. It can be as low as 19 MHz yielding the scan times shown.

These basic calculations indicate the infrequent occurrence of SEUs that affect the device configuration. Even so, some critical applications must consider appropriate precautions and define post-SEU actions.

When more products are deployed, the probability of an SEU affecting any one of them increases proportionately. For example, if the above XC5VLX50T is used in 1,000 products, the nominal FIT across all products is 1,489,000 and represents one SEU every 28 days. This should not be confused with the probability of an individual device being affected. Also, the probability of an individual device incurring a second SEU is determined by the FIT (or MTBF) of the individual device and not the collection. This is an important consideration when assessing suitable strategies for an application. For example, thorough worst-case calculations that consider the largest Virtex-5 device operating at the highest populated locations in the world only indicate one upset per year. The selected strategy should be based on this infrequent potential for disruption even if it was the MTBF for the collection of products to be shipped that motivated the consideration of a requirement to choose one.

The Rosetta experiment demonstrates that in the vast majority of cases, an SEU only changes (flips) a single configuration bit. Multi-bit upsets (MBUs) due to a single ionizing particle almost never occur. Also, there is a high probability that this bit will have little or no effect on the design because less than 20%, and typically less than 10%, of the configuration cells have any significance to a design implementation. An SEU affecting device resources that are not used (for example, unused CLBs, I/O, DCMs, block RAMs, etc.) will have no effect. The percentage of device resources used by a particular design is available in the MAP report, and the device FIT or MTBF calculation can be scaled to reflect the proportion of device used.

Every configuration frame consisting of 1,312 bits contains 12 built-in error correction code (ECC) bits. Any change to the ECC bits caused by an SEU has no effect on the active design.

Each frame also contains 16 unused bits (bits 656 to 671 as shown in the figure captioned “Configuration Words in the Bitstream and Configuration Bits in a Frame” in [UG191](#), *Virtex-5 FPGA Configuration User Guide*). Although less obvious, invariably 8 to 13% of the bits in the configuration memory map can never be flipped by an SEU and can be deducted from the total configuration size when calculating the FIT of a device.

There is no effect when a large number of configuration cells are unused within the area of an otherwise used resource. For example, the programmable interconnect has many possibilities but only a few of those apply to a particular design. Consequently, an SEU that causes the connection of an unused segment of interconnect to another unused segment has no effect on a particular design. Even a connection of an unused segment to a used segment is unlikely to have any noticeable effect.

As a conservative approximation, the device configuration FIT can be divided by ten to provide a design configuration FIT, which for the XC5VLX50T device translates into a design configuration MTBF of 610 to 960 years with a 90% confidence level. This further improves when the design occupies less than 100% of the resources.

If an application needs reliability at this level, then other aspects of the system and design not related to SEUs almost certainly require analysis similar to the basic risk assessment already described. For example, intermittent faults have been traced back to the use of asynchronous resets and cross coupling between traces on PCBs when the fault was initially attributed to SEUs. The use of asynchronous resets in a high reliability design is not recommended and priority should be given to eliminating this common design practice before proceeding with SEU topics.

Comparison of Configuration and Data SEUs

If an SEU occurs in a Virtex-5 device, the state of a bit is flipped. That bit can be associated with the configuration of the device or it can be a change to the operational data of the current design. In this situation, data means anything that is a variable within the design, including the contents of RAM and flip-flops. Although the focus of this application note is the impact of SEUs on the integrity of configuration cells, SEU effects on data must be considered because they also influence the selection of precautions used for configuration.

In most cases, a single bit change within informational data can be tolerated and ignored, for example, a single corrupted ASCII character in an e-mail or one incorrect pixel in a video image. More significantly, any transitory data is overwritten by new information, meaning that the effect of the SEU is short lived. However, when the data takes the form of instructions or a state required in the continuing operation of the design, the effect of an error can be significant and prolonged, for example, a state machine might enter an illegal state or the IP address for communication be made incorrect. In such cases, risk assessment should be performed, and suitable precautions should be considered.

As described previously, the FIT or MTBF associated with the configuration cells of a Virtex-5 device can be calculated using the results from the Rosetta experiment. In a similar way, the FIT or MTBF can be calculated for the contents of the block memories of the device. The XC5VLX50T device has sixty 36 Kb block RAMs, having a total memory capacity of 2.21 Mb. The nominal FIT/Mb value for block RAM contents reported in [UG116](#), *Device Reliability Report* is 643, which yields a device block RAM FIT of 1,421 or an MTBF of 80 years.

Although there are over five times more configuration cells than block RAM memory bits, the FIT for the device configuration (1,489) is similar to the FIT for the device block RAM (1,421). The configuration cells are robust based on their requirement to remain static most of the time while block RAM memory must switch between states quickly for operational reasons. This makes them more susceptible to SEUs.

It is common for designs to use a large proportion of the available block RAM resources, which increases the probability that an SEU will affect the data within a design. In comparison, the typical amount of configuration cell usage is low and the design data FIT will dominate over the

design configuration FIT. Therefore, when block RAM is used extensively in an application, the importance of data should be given consideration before that of configuration.

The [Macro Size and Analysis of Reliability](#) section provides an example of how to perform a risk analysis calculation for any design. It estimates the configuration and data FIT rates that can be compared or combined to determine the total FIT of the macro.

Data contained in flip-flops are least likely to suffer an SEU. Accelerated tests predict the FIT for flip-flops to be as low as 1 to 2 per megabit. Given this low value and the low number of flip-flops in a device, flip-flops can be normally omitted from risk calculations. For example, the XC5VLX50T has approximately 0.03 Mb of flip-flops, which means a maximum device flip-flop FIT of 0.06 or an MTBF of nearly 2 million years, even if all flip-flops in the device are used.

However, if any data value is highly significant, including the data held in flip-flops, the design should contain precautionary logic to detect, correct, ignore, or recover from an SEU disturbance in a way that is appropriate to the application. Only the design has the ability to determine when data is corrupted because data is variable with a meaning specific to the application. Since the block RAM content has the potential to dominate the requirement for precautionary logic, the block RAMs of Virtex-5 devices are provided with an ECC option, which can be exploited. However, logic circuits that require absolute data integrity to operate correctly might have to be implemented with a degree of redundancy, such as employing a triple modular redundancy (TMR) technique.

Strategies for Handling Configuration SEUs

Given the previous risk assessment values, it is understandable why the vast majority of applications can ignore the whole subject of SEUs. Xilinx continues to commit time and resources to the Rosetta experiment to provide the information and data needed to assess the risks. This data should be used as a starting point, and claims not backed by similar data should be viewed cautiously.

There are applications (often small parts of larger designs) for which even the smallest risk is unacceptable, and some precautions and actions must be considered. An SEU disturbance to a configuration cell has the potential to change the definition of the design itself, and unlike informational data, is not overwritten by a new value. However, this possibility should not prevent even high reliability systems from benefiting from Virtex-5 FPGA technology. The strategies described in this application note assume that this is the requirement and that even small risks are unacceptable. It will also focus on the effects of SEUs on configuration cells.

The strategies discussed are scheduled maintenance, emergency maintenance, running repairs, and a combination of all. Because an SEU is a highly unlikely event, the designer must carefully consider the disadvantages of adopting a particular strategy as well as its advantages.

Scheduled Maintenance Strategy

Now consider the possible effects resulting from an SEU that flips a configuration cell that directly impacts an active design. The effect on the design can either be almost instantaneous, or irregularities might not be noticed for a significant time. For example, a disturbance to the main system clock distribution can have a marked effect almost immediately but a change to a circuit used to display the hours on a clock display might not become apparent for literally hours.

The point is that some parts of an application are more critical than others, and it is the critical parts to which precautions should be applied. It is useful to assess the time an SEU takes to affect a function and to consider the consequences that a failure can have. An important question to ask is if the system is required to maintain normal operation after the event or is it only required to fail safely? The answer helps to determine the appropriate precautions and actions to take.

An SEU is a soft error, meaning that its effect can be reversed and has no lasting damage. This is significantly different from a hard error such as a broken wire to a connector, which typically

requires the replacement or physical repair of some part of a system. For this reason, the FIT or MTBF associated with SEUs should not be confused with that of product life expectancy. Whenever an FPGA is configured (for example, following the application of power or cycle of PROGRAM_B), all configuration cells are defined for the required design regardless of any previous state, and any error caused by an SEU is removed.

Although there are some applications that operate continuously for very long periods of time, very few are expected to operate continuously without interruption for the entire product lifetime. Realistically, most applications experience relatively frequent power cycles or periods of inactivity when maintenance can be performed. The scheduled maintenance strategy is intended to fully exploit these opportunities. Obviously, a power cycle inherently results in reconfiguration of the device and requires no further thought. But the best use of the scheduled maintenance strategy exploits every available opportunity to reconfigure the device during normal operation and reconfigure whenever it is not playing an active role in the system. No attempt is made to determine if an SEU has occurred; the reconfiguration simply repairs any corruption, if it exists. This is the same concept as performing regular maintenance on an aircraft, where certain parts are replaced at regular intervals even if they appear to be perfectly serviceable.

Even with the confidence that any errors are corrected by the next scheduled device reconfiguration, what happens for the period of time between an SEU disturbance and the next device reconfiguration? This is when precautions must either maintain normal service or fail safe as the application requires. Some degree of redundancy is required in the design, such as the use of TMR techniques, which accommodate a failure in one circuit because the remaining two circuits continue to provide normal functionality.

By including the redundancy that the application requires, a single bit error in the configuration resulting from an SEU should be acceptable. The probability of the device receiving another SEU before the next scheduled maintenance should also be considered. This is easy to determine because the same MTBF figure applies, and the shorter the time to the scheduled maintenance, the smaller the opportunity for a second SEU to happen. The previous calculations showed that the nominal design configuration MTBF for an XC5VLX50T device was at least 610 years, so if the device experiences an SEU, it is not likely that a second SEU will occur in the lifetime of the device, let alone the hours until the next service.

Even if a second SEU were to occur before the next service reconfiguration, it would have to cause a single bit error in a specific place to have an adverse effect. Given that less than 10% of configuration bits have a direct impact on the operation of a design, there is less than 1% probability that both SEUs would affect the design. If a design also employed TMR techniques, a second error in the module that has already failed does not matter. Likewise, an error occurring in a different part of the application outside of the remaining modules of the previously affected TMR function is covered by its own TMR circuits.

In summary, scheduled reconfiguration definitely corrects any errors that have occurred, and if a design contains preventative measures to cope with an SEU, this scheme should be acceptable.

Emergency Maintenance Strategy

The concept behind the emergency maintenance strategy is the same as a car's brake warning light, which indicates the brake pads might be worn. The appropriate course of action is to take the car to the garage to have the brakes checked and replaced as soon as possible rather than wait until the next scheduled service.

For an FPGA, this means bringing forward the next reconfiguration of the device when an SEU impacts the configuration cells. An analysis must determine if the warning is worthy of immediate reconfiguration or for how long it can be delayed, but the objective is to reconfigure as soon as practically possible.

The key to the emergency maintenance strategy is detecting if an SEU has occurred in the configuration cells. The Virtex-5 and Extended Spartan-3A devices provide a built-in readback CRC facility to make this possible. A comprehensive description is provided in [UG191](#), *Virtex-5 FPGA Configuration User Guide* and [UG332](#), *Spartan-3 Generation Configuration User Guide*. In simple terms, the built-in circuit continuously scans (reads) the configuration cells of the device and computes a 32-bit CRC value. A 32-bit CRC is able to detect any change in as many as 2^{32} bits, which far exceeds the size of the largest devices. If the CRC value computed by a scan differs from the CRC value computed by the first scan immediately following device configuration, a change to the configuration has occurred and the disturbance is indicated by driving the INIT_B pin Low. Virtex-5 devices also provide an internal signal that is driven High, which can be observed on the CRCERROR output of the FRAME_ECC_VIRTEX5 primitive.

Reaction to the disturbance varies according to the application. Two factors must be recognized. First, the error signal indicates a change in the configuration of the device and does not (and cannot) indicate a change to informational data of the application. If parts of the design are data critical, these parts still need precautionary circuits of their own. Second, the detection of an SEU by the readback CRC circuit takes some time. [Table 1](#) shows the scan times for each Virtex-5 device.

The time to report an SEU depends on the relative positions of the scan read location and the location of the SEU when it occurs. When only using the Readback CRC mechanism, an SEU positioned later in the scan is reported in less than one device scan while an SEU positioned earlier in the scan requires completion of the next device scan. Hence, the time to detection is up to two device scans, with an average detection time of one device scan. If the SEU controller macro is also employed, then the maximum time to detection is reduced to one device scan, with the average time to detection being half of one scan along with the capability to perform correction. If the maximum detection time can be tolerated before an otherwise instant reconfiguration of the device is initiated, then the emergency maintenance strategy avoids the requirement for any special circuitry. If the detection time is considered significant (for example, 20 ms equates to 3,000,000 clock cycles in a design employing a 150 MHz clock), then some precautionary circuits should be part of the design. Depending on the balance of precautions included in the design and the requirements of the application, the urgency of reconfiguration can be assessed, and reconfiguration can be delayed until a suitable point in design operation.

Because less than 10% of the configuration cells have a direct effect on a typical design, even if the readback CRC has reported a configuration change, then statistically more than nine out of every ten configuration SEUs will have no effect. Forcing emergency maintenance reconfiguration for every instance could be more of an issue than continuing to operate with precautionary logic, as described in the regular maintenance strategy.

Based on the values for the XC5VLX50T calculated previously, then even the low end of the nominal MTBF range is 61 years for an SEU affecting any of the configuration cells and indicates that emergency maintenance will be a rare event anyway. In comparison, the low-end MTBF for an SEU that impacts the operation of the design is at least 610 years suggesting that emergency reconfiguration can be avoided.

In summary, consider using the emergency maintenance strategy when the design can tolerate a configuration fault for the duration of the readback CRC soft error detection time plus the reaction time needed to initiate the device reconfiguration. If this is the case, the design can avoid all precautionary circuits associated with configuration errors, simplifying the design, and keeping it small enough to fit in a lower density device. This is a significant advantage because a smaller device is statistically less susceptible to SEUs.

Readback CRC Considerations

Figure 1 is a representation of the dedicated readback CRC, which is activated by the POST_CRC constraint. For highest reliability of detection, the internal configuration oscillator more commonly associated with CCLK in master modes should be employed, and the INIT_B pin should be used to observe the configuration status. INIT_B is also driven Low in response to an attempt to program with a corrupted configuration image or in preparation for programming, so it is already important to monitor this signal in any high-reliability system. An external controller requires adequate intelligence to interpret the reason for INIT_B being Low, based on the configuration status of the device.

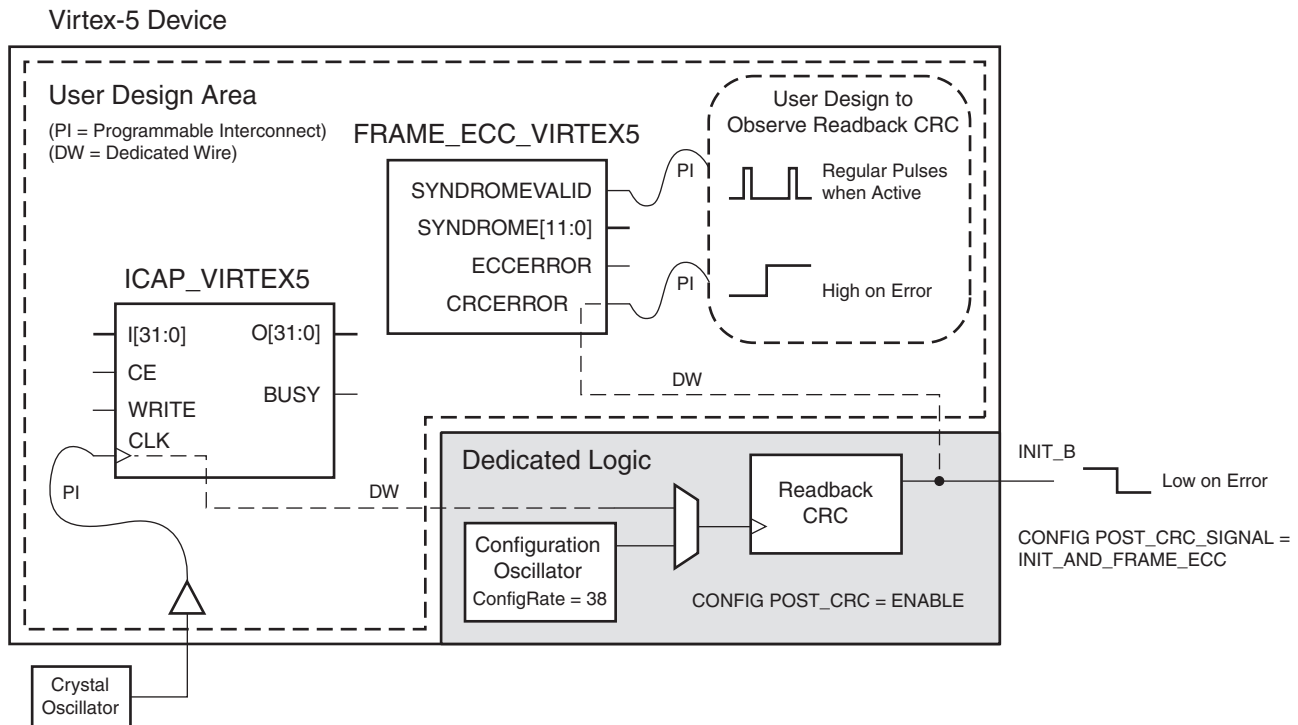


Figure 1: Readback CRC Arrangement in a Virtex-5 FPGA

The optional FRAME_ECC_VIRTEX5 primitive provides internal access to the same (even though inverted) CRC error status as presented externally on the INIT_B pin and can be used to observe a configuration error and take actions appropriate for the application. It is possible to disable the INIT_B pin using a POST_CRC_SIGNAL = FRAME_ECC_ONLY constraint so that only the internal signal is used to report a configuration error. However, it must be appreciated that the programmable interconnect associated with the routing of the internal CRCERROR signal and any user logic that is intended to respond to it is also susceptible to SEUs. Therefore, the INIT_B pin represents the highest reliability reporting point. For the very highest reliability systems, both reporting points can be used with different connections and circuits providing a degree of redundancy.

If an alternative clock is desirable (typically for predictably fast detection times), it can be supplied by a connection to the CLK input of the optional ICAP_VIRTEX5 primitive (or the USRCCLKO input of the STARTUP_VIRTEX5 primitive not shown). Reliable operation of the readback CRC logic is now dependant on the integrity of the alternate clock supply and its connection to the ICAP_VIRTEX5 primitive. The clock should be provided as directly as possible (as illustrated in Figure 1) because any additional logic and programmable interconnect in the path is also susceptible to SEUs. If the external clock stops or if it fails to reach the readback CRC circuit, then device scanning stops and it is not able to report any errors.

The FRAME_ECC_VIRTEX5 primitive also provides access to the SYNDROMEVALID signal. While the primary purpose of this signal is in relation to the 12-bit SYNDROME and ECCERROR outputs, they are not used in this scenario. Instead, the SYNDROMEVALID signal can be used to confirm that the readback CRC circuit is actively scanning, and in so doing, confirms the integrity of the clock being used.

Throughout each device scan, the SYNDROMEVALID signal pulses High for one clock cycle at intervals of 41 readback CRC clock cycles, followed by one 49 clock cycle interval at the end of each complete scan. When the internal master configuration clock is employed, the precise timing of the pulses will be unknown, but a simple watchdog circuit can be implemented to determine the presence or absence of the pulses without being concerned with their exact timing or spacing. The use of the same clock source for the watchdog circuit and the readback CRC is detrimental to system reliability so the use of the internal master configuration clock is also advantageous in this respect. Although a watchdog circuit itself is also susceptible to SEUs, any such occurrence still results in an error report being generated by the readback CRC circuit a short time after the SEU, so this case is naturally covered. If the watchdog circuit fails to detect pulses, action will need to be tailored to the specific application. It might be appropriate to consider a failure of the readback CRC as less of an emergency than an actual readback CRC configuration error. When using an external clock, the failure could be associated with the external oscillator or connections and such a hard error cannot be fixed as easily as reconfiguration repairs a soft error within the FPGA.

Readback CRC Scan Times

The number of clock cycles required for each scan of a particular device is fixed, but the frequency at which the circuit is clocked, and hence the time to complete each scan of the device, is dependant on the user settings and/or the particular device.

The most reliable configuration detection is achieved when the readback CRC circuit is driven by the internal oscillator normally associated with master mode configuration (CCLK). This highest reliability relates to the fact that no external oscillator is required and eliminates the risks posed by physical components and connections as well as avoiding the risks associated with SEU strikes to the programmable interconnect used to route the clock signal as part of the user design. However, the internal oscillator is subject to variations in manufacturing process, operating voltage, and temperature (PVT), and has the potential to deviate up to $\pm 50\%$ from its nominal value, which is defined by the ConfigRate option in the BitGen tool. The nominal ConfigRate values available are 2, 6, 9, 13, 17, 20, 24, 27, 31, 35, 38, 42, 46, 49, 53, 56, and 60 (MHz). The selection of a higher value results in a shorter readback CRC scan time, so it is probably best to specify a value greater than the default value of 2 (MHz). Care should be taken when selecting higher frequencies to ensure that it is compatible with your configuration method as well as ensuring that the readback CRC maximum frequency of operation is not exceeded should the clock be 50% above its nominal value. For this reason, the maximum value of ConfigRate selected should be 38, yielding a frequency in the range 19 to 57 MHz with corresponding range of readback CRC times ([Table 1](#)).

Alternatively, a user clock can be supplied to the readback CRC circuit, which allows the scan time to be predictable and typically faster. However, note that the reliability of the detection is slightly degraded by the dependence on the external clock, the clock's connections to the device, and the risk that SEUs could impact the programmable interconnect routing the clock from the pin to either the ICAP or STARTUP primitives. It is therefore sensible to ensure that the clock source is reliable and that all connections (especially programmable) are kept as simple and as short as practically possible.

As a further precaution and general indication that readback CRC is operational, it is possible to monitor the SYNDROMEVALID output of the FRAME_ECC_VIRTEX5 primitive. This signal pulses High once every 41 clock cycles (49 cycles at the end of each complete device scan) and can therefore be used as a heartbeat indicating all is well. If the clock is lost for any reason (including a hard failure loss of the external clock), then the heartbeat will cease, and the system must take appropriate action.

Running Repairs Strategy

The running repairs strategy is useful in applications where it is desirable to maintain operation following an SEU while carrying out a rapid localized repair of single bit errors. Given the high probability (>90%) that an SEU has no effect on a particular application, this strategy avoids the interruption to service associated with the emergency reconfiguration strategy, which invokes a full device reconfiguration. However, when an SEU does impact the application, the repair is very important, and the limitations of the running repairs strategy must be considered.

As with the emergency maintenance strategy, the built-in readback CRC feature of a Virtex-5 device is instrumental in facilitating the detection of changes to the configuration cells. Using this dedicated circuit ensures that the detection is reliable and rapid. Although a configuration upset is detected and reported, only the behavior of the design and system can deduce if the SEU has any effect on the operation. As described earlier, an SEU has no effect in the majority of cases. However, when an SEU does impact the application, the potential exists for undesirable application behavior from the time the SEU occurs until it is detected and corrected.

A logical error caused by an SEU can also result in prolonged corruption of informational data but this might be tolerable if the repair is made quickly and has no lasting effects. For example, if a configuration bit in the path of a streaming video stream is corrupted, it might cause a disturbance to the displayed images for a few frames until the repair is complete and new video data overwrites the older corrupted images. However, a logical error which adversely affects the control paths or operating states of a design can have lasting effects and precautionary logic might be required in this situation. One potential solution is to issue a reset to all critical circuits upon detection of an error and then remove the reset after the error has been corrected. Although this interrupts normal operation, it will be of significantly shorter duration than the time required for a full reconfiguration. This is especially true when the higher density devices are used.

Having detected an error, it must be located. The exact location can be determined using the configuration ECC and syndrome calculator circuit which is included in every Virtex-5 device. As described in [UG191](#), *Virtex-5 FPGA Configuration User Guide*, there is a 12-bit ECC value embedded in each frame of configuration consisting of 41 words of 32 bits (1,312 bits). As each frame is read over 41 clock cycles, the built-in circuit calculates a 12-bit syndrome for the current content which combines with the embedded 12-bit ECC value to expose any error (even if the error is within the 12 ECC bits). It is then possible to use this 12-bit syndrome, which is presented by the FRAME_ECC_VIRTEX5 primitive, to identify the location of a single bit error in the frame that has been read.

In principle, the correction is straightforward. The corrupted frame of 1,312 bits is read into a buffer, the single bit error is isolated by interpreting the 12-bit syndrome, and then it is corrected. The correction simply inverts the bit to reverse the inversion caused by the SEU. Finally, the corrected frame is written back into the configuration cells. In practice, this task is quite complex, so this application note introduces an SEU controller macro to perform all the necessary steps.

Once an error has been corrected, the readback CRC circuit restarts with the error reporting signals reset. The next complete scan of the device provides confirmation that the device configuration is restored. Although the ECC bits enable any single bit error to be corrected, it is not possible to correct multiple bit errors within the same frame (1,312 bits). In the unlikely event two errors occur in the same frame, the error condition will persist and the system must determine a suitable course of action.

Some multiple errors can be repaired because they take the form of two single bit errors in adjacent frames. Although extremely rare, the SEU controller could be observed to detect and correct each bit separately in rapid succession. Closer examination of the correction reports generated would reveal the adjacent frames and therefore reveal that the errors resulted from the same SEU.

As mentioned before, even multiple bit errors that cannot be repaired will probably have no real effect on the operation of the active design, but extremely high reliability systems must plan for this possibility. If an application must maintain operation even in this extreme case, TMR techniques are probably used extensively throughout the design. In such cases, the level of redundancy almost certainly means that the design can continue operating until the next scheduled reconfiguration of the device.

Ultimately, the SEU controller macro should only be included in a design when the benefits and limitations of the running repairs strategy are understood. If the design contains adequate redundancy, then a regular maintenance strategy or delayed emergency maintenance strategy should be adequate and addition of the SEU controller might add unnecessary complexity.

Combined Strategy

Systems requiring the very highest reliability benefit from using a combination of elements from the scheduled maintenance, emergency maintenance, and running repairs strategies. This multi-layered approach continues the redundancy concept by using each scheme to provide cover for the next.

To provide a good foundation, the design of the system has to be superior in all respects. The susceptibility of Xilinx products to SEUs is low and published values define the level of risk. These contributions should be exploited.

Preventing failures is preferable to having to cope with them during service. Every aspect of the design implementation contributes to the level of reliability (or unreliability), so poor design practices (such as the use of asynchronous reset in HDL code) must be avoided.

When considering the most beneficial combinations for an FPGA-based system, the operation of an aircraft is a useful analogy when reliability is paramount. Once in service, regular maintenance of the aircraft includes the replacement of parts even if there is no visible evidence of a problem. This prevents failures due to wear over time as well as fixing defects missed during in situ inspections. Reconfiguring an FPGA at regular intervals offers similar advantages by correcting possible unseen errors and ensuring a clean start for each period of operation.

Before each flight, checks are made to ensure all critical systems on the aircraft are working properly. If the check reveals a fault, the aircraft is grounded until the fault is corrected. Using the built-in readback CRC circuit of the Virtex-5 device to continuously check the FPGA configuration performs a similar function. If a system is in a standby mode and an error is detected, an emergency reconfiguration can be invoked to repair the fault. This avoids the potential of accumulating errors over time and ensures the device is in perfect condition when it enters an operational mode.

Once an aircraft is in flight, any failure is undesirable. An active warning light alerts the crew to the nature of the fault so that they can take appropriate actions. These actions vary depending on the severity and location of the failure. It might be possible to contain the problem and continue to the planned destination, or a diversion and an emergency landing might be required. The most challenging scenario is when the plane must maintain flight because there is no suitable landing site, for example, in the middle of the Atlantic. In all cases, once the plane has landed, the aircraft is removed from service and repaired before being flown again. This applies even if the fault was considered to be temporary and resolved by the crew during the flight.

In a Virtex-5 device, the warning is provided when the built-in readback CRC scan detects an SEU. In most cases, the SEU does not affect the operational design. However, the cases that do impact operation require that appropriate action is taken, depending on the severity of the impact. If the device can be safely reconfigured immediately, this is the most obvious response to the warning. If continuous operation must be maintained, then redundancy and localized repairs should be considered. These techniques can be used separately or in combination depending on the design's specific requirements.

When redundancy is included in the design, a decision needs to be made as to how much is suitable to cover the potential risk and maintain operation until the system can be repaired. The Virtex-5 FPGA built-in ECC logic exploited by the SEU controller macro introduced in this application note enables localized repairs to be made during operation. The worst case time to detect and repair an error can be estimated using the table of readback CRC scan times (Table 1) because the scan time for detection is the dominant factor. In the worst case scenario, detection would take one complete scan. In comparison, the error would appear to be corrected almost immediately. It is important to remember that following correction, the effects the error had on operational states and data might prolong and a localized reset to circuits might be appropriate. Regardless of all the measures employed to maintain operation, it would still then be wise to carry out a full device reconfiguration at the earliest convenient opportunity as this ensures that all data and states as well as the device's configuration are known to be good.

Introducing the SEU Controller Macro

The SEU controller macro serves two purposes. The principle purpose is to correct configuration errors resulting from an SEU as soon as they are detected and hence to facilitate the [Running Repairs Strategy](#). The secondary, but highly valuable purpose, is to provide a convenient way to emulate SEUs within the Virtex-5 device by injecting errors in a controlled and predictable way into the configuration memory. This tool enables the evaluation of upsets on a particular design and to verify the claims made in this application note even if the decision is taken to exclude the macro from a final system.

The Virtex-5 FPGA SEU controller is provided as fully-supported LogiCORE™ IP with detailed user information. Follow this link to the license and download area:

http://www.xilinx.com/member/sem_core/index.htm

The SEU controller macro occupies less than 4% of the smallest Virtex-5 device and an insignificant amount of the larger devices in the family. The macro is included in a design where the only definite requirement is the application of a clock (50 MHz recommended). Inside the macro, the ICAP_VIRTEX5 and FRAME_ECC_VIRTEX5 primitives are used in the same manner as previously described in [Readback CRC Considerations](#) to clock and observe configuration status as the readback CRC circuit scans the device. When a frame containing an error is scanned, the resulting syndrome ECC error is observed by the SEU controller to trigger the correction procedure immediately.

A small set of simple control and status signals are all optional and can be connected to a design as required. The macro also contains an optional UART interface that is highly recommended for making connections during board and system testing between two I/O pins on the device. Using a simple terminal (e.g., HyperTerminal on a PC), it is possible to control the emulation of an SEU and observe in detail the detection and correction process virtually independent of the main design.

Although many experiments are possible using the SEU controller, one example is to emulate 1,000 SEUs and build a statistical profile of how susceptible a design is to an SEU within a given device and how the device behaves on the rare occasions when it is affected. As suggested previously, it is expected that less than 100 emulated SEUs (typically less than 50) would have any effect on the active design. With correction enabled, the challenge is often being able to detect a disturbance. This hands-on experiment provides invaluable experience to a designer. Xilinx highly recommends obtaining a copy of the SEU controller macro, using it in a design, and connecting the UART to a terminal for this experience.

Macro Size and Analysis of Reliability

This section provides techniques to estimate the FIT and MTBF of macros occupying only a portion of a given device. The example used is the SEU controller macro but the same methodology can be applied to evaluate the SEU susceptibility of each section of a design, or at least those sections that are most critical to operation.

The SEU controller macro occupies approximately 174 logic slices (with some possible variation due to mapping in a complete design) and one block memory of 18 Kb (block RAM in Virtex-5 FPGAs is actually 36 Kb, but these can be divided to form two smaller memories). The macro also includes the ICAP_VIRTEX5 and FRAME_ECC_VIRTEX5 primitives required for the detection, correction, and error injection tasks.

Since the macro is itself part of the FPGA, it is also susceptible to SEUs. For this reason, the same analysis must be applied to the macro as well as the whole device or any part of the design.

Starting with the configuration cells, since the nominal value reported in [UG116](#), *Device Reliability Report* is 131 FIT/Mb. However, a way is needed to convert the 174 logic slices and their associated interconnect together with the configuration and interconnect (not data contents) of the one block RAM into a number of configuration bits. To facilitate this process, [Table 2](#) can be used to generate a reasonably accurate estimate. For the macro, that value is $(174 \times 1,181) + (1 \times 585) = 206,078$ bits, or 0.206 Mb. This results in a nominal macro configuration FIT of 27 or an MTBF of approximately 4,228 years.

Table 2: Approximate Number of Configuration Bits Associated with the Most Common Device Features

Device Feature	Approximate Number of Configuration Bits
1 logic slice	1,181
1 block RAM (36 Kb)	1,170
1 block RAM (18 Kb)	585
1 I/O block	2,657
1 DSP48E slice	4,592

Notes:

1. In all cases, the device feature includes all the programmable interconnects associated with getting signals to and from the feature.
2. These features account for approximately 95% of the total configuration cells of each device.
3. Block RAM does not include the data contents of the memory which must be analyzed separately.

An alternative, somewhat pessimistic way to estimate configuration FIT of a design is simply to evaluate the percentage of a given device that is occupied and use that proportion of the total configuration bits for the device. For the SEU controller macro in an XC5VLX50T device, that equates to 2.42% of the slices and 1.84% of the block RAMs.

[Table 1](#) states that the XC5VLX50T device has 11.37 Mb configuration bits. Therefore, the macro is estimated to be associated with approximately 2.4% of 11.37 Mb (which is 0.27 MB), leading to a nominal configuration FIT of 35 or a MTBF of 3,262 years.

This estimate is 30% more pessimistic than the previous, more accurate estimate, but might be a quicker way to obtain a first approximation, especially when evaluating designs that use a wider selection of features not covered by [Table 2](#).

Regardless of how the nominal configuration FIT is estimated, less than 10% of configuration cells (actually less than 5% in typical designs) would directly impact the active design if an SEU occurred. Therefore, it is reasonable to scale the estimates by a factor of 10. In this case, scaling the more accurate estimate reveals an operational configuration FIT for the SEU macro of approximately 2.7 FIT or an MTBF of 42,280 years.

Next, the susceptibility of the block RAM contents to SEUs must be considered. This is initially a straightforward calculation because there is one block RAM of 18 Kb, which is a total of 18,432 bits (or 0.0184 Mb). The nominal block RAM data FIT from [UG116, Device Reliability Report](#), is 643, resulting in a block RAM data FIT for the SEU controller macro of 11.85 or an MTBF of 9,633 years. However, closer analysis of the macro operation reveals that less than half of the block RAM data contents can be considered critical, and therefore the meaningful data FIT of the macro is approximately 5.9 or an MTBF of 19,348 years. The number of data bits associated with flip-flop and distributed RAM contents are so small that no meaningful values can be generated.

Combining the configuration (FIT = 2.7) and data (FIT = 5.9) values, the complete operational susceptibility of the SEU controller macro to a potentially disruptive SEU is a FIT of approximately 8.6 or an MTBF of 13,274 years.

SEU Controller Macro

Table 3: SEU Controller Macro Matrix

Parameter	Description
Developer Name	Xilinx
Target Devices (stepping level, ES, production, speed grades)	Virtex-5 FPGAs
Source Code Provided	Yes (HDL only)
Source Code Format	VHDL and Verilog
Design Uses Code/IP from an Existing Reference Design/Application Note, Third Party, or CORE Generator™ software	Yes Incorporates PicoBlaze™ processor
Simulation	
Functional Simulation Performed	Verified in hardware (simulation models for configuration unavailable)
Timing Simulation Performed	
Testbench Used for Functional Simulations Provided	
Testbench Format	
Simulator Software Used/Version (for example, ISE® software, Mentor, Cadence, other)	
SPICE/IBIS Simulations	
Implementation	
Synthesis Software Tools Used/Version	XST
Implementation Software Tools Used/Versions	ISE software, version 10.1, service pack 3
Static Timing Analysis Performed	
Hardware Verification	
Hardware Verified	Yes, including proton beam testing to verify real SEU
Hardware Platform Used for Verification	ML505 and ML507 Virtex-5 FPGA evaluation platforms as-well-as test boards used to verify multiple Virtex-5 devices of various densities during beam experiments

Conclusion

The selection of an SEU strategy should begin with a realistic assessment of the risk of occurrence for which Xilinx provides some of the most comprehensive data in the industry. The effect that an SEU can have on your design should be assessed and compared with the requirements for your system to maintain operation should such an event occur.

Virtex-5 devices contain a built-in readback CRC circuit, which automates the detection of SEU errors and enables the system to take appropriate action.

Designs required to maintain operation must evaluate the merits of redundancy within the design, and the ability to perform local error correction using the SEU controller introduced in this application note. The time from SEU detection to correction is rapid, but redundancy might still be considered necessary to bridge this period. Having redundancy might also imply that localized correction is not required for these rare events, but the macro can help keep the redundancy circuits smaller, reducing costs, and actually decreasing the statistical probability of an SEU in the first place.

Revision History

The following table shows the revision history for this document.

Date	Version	Description of Revisions
02/20/09	1.0	Initial Xilinx release.
03/05/09	1.0.1	Typographical edits.
04/01/10	2.0	Added TXT devices to Table 1 . Updated the application note to the latest data. Numerous changes to move from a reference design to a new generation SEU controller macro available as LogiCORE IP.

Notice of Disclaimer

Xilinx is disclosing this Application Note to you “AS-IS” with no warranty of any kind. This Application Note is one possible implementation of this feature, application, or standard, and is subject to change without further notice from Xilinx. You are responsible for obtaining any rights you may require in connection with your use or implementation of this Application Note. XILINX MAKES NO REPRESENTATIONS OR WARRANTIES, WHETHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL XILINX BE LIABLE FOR ANY LOSS OF DATA, LOST PROFITS, OR FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES ARISING FROM YOUR USE OF THIS APPLICATION NOTE.

CRITICAL APPLICATIONS DISCLAIMER

XILINX PRODUCTS (INCLUDING HARDWARE, SOFTWARE AND/OR IP CORES) ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS IN LIFE-SUPPORT OR SAFETY DEVICES OR SYSTEMS, CLASS III MEDICAL DEVICES, NUCLEAR FACILITIES, APPLICATIONS RELATED TO THE DEPLOYMENT OF AIRBAGS, OR ANY OTHER APPLICATIONS THAT COULD LEAD TO DEATH, PERSONAL INJURY OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE (INDIVIDUALLY AND COLLECTIVELY, “CRITICAL APPLICATIONS”). FURTHERMORE, XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED FOR USE IN ANY APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE OR AIRCRAFT, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR. CUSTOMER AGREES, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE XILINX PRODUCTS, TO THOROUGHLY TEST THE SAME FOR SAFETY PURPOSES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN CRITICAL APPLICATIONS.

AUTOMOTIVE APPLICATIONS DISCLAIMER

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS.