

PROGRAMMABLE LOGIC APPLICATION NOTES

Richard Katz
Microelectronics and Signal Processing Branch
NASA Goddard Space Flight Center
301-286-9705
rich.katz@gsfc.nasa.gov

This column will be provided each quarter as a source for reliability, radiation results, NASA capabilities, and other information on programmable logic devices and related applications. This quarter the focus is on some experimental data on low voltage drop out regulators to support mixed 5 and 3.3 volt systems. A discussion of the Small Explorer WIRE spacecraft will also be given. Lastly, we show take a first look at robust state machines in VHDL and their use in critical systems. If you have information that you would like to submit or an area you would like discussed or researched, please give me a call or e-mail.

1999 MAPLD Conference
September 28-30, 1999
Kossiakoff Conference Center
JHU/Applied Physics Laboratory
Laurel, Maryland

The 2nd annual Military and Aerospace Applications of Programmable Devices and Technologies Conference will address devices, technologies, usage, reliability, fault tolerance, radiation susceptibility, and applications of programmable devices and adaptive computing systems in military and aerospace systems. The program will consist of approximately 60 oral and poster technical presentations and 20 industrial exhibits. The majority of the conference is open to US and foreign participation and is unclassified. There will be one classified session at the secret level, for U.S. citizens only. For conference information, please see the Programmable Technologies Web Site (<http://rk.gsfc.nasa.gov>).

1999 IEEE NSREC and RADECS Papers

A number of papers were given at the 1999 IEEE NSREC on programmable devices with the meeting held in July 1999. Other programmable-related papers will be given at the 1999 RADECS conference during September 1999. This section will list the titles and first author information for each of these articles. E-mail addresses for NSREC first-authors may be found at:

<http://www.nsrec.com/email.htm>

Single Event Upset Immunity of Strontium Bismuth Tantalate Ferroelectric Memories, J.M. Benedetto.

The Impact of Software and CAE Tools on SEU in Field Programmable Gate Arrays, R.B. Katz.

Design Guidelines for COTS in Military and Space Systems, P.S. Winokur.

Reprogrammable FPGA for Space Applications, J-J. Wang.

The Effects of Architecture and Process on the Hardness of Programmable Technologies, R.B. Katz.

Radiation Effects on Advanced Flash Memories, D.N. Nguyen.

SEU and Microdose Measurement Based on FAMOS Transistors, P.J. McNulty.

Total Ionizing Dose Effects in SRAM-Based FPGAs, B.G. Henson.

Total Dose and Dose-rate Effects on Start-up Current in Antifuse FPGA, J. J. Wang. (RADECS)

Total Ionizing Effects in a SRAM-based FPGA, D.M. Gingrich (RADECS).

What's New?

A large amount of data, reports, papers, application notes, and conference information are being stored on our companion Programmables Technology www site, <http://rk.gsfc.nasa.gov>. In order to make it easier to keep readers up to date, all new additions to the site are being listed in chronological order on our "What's New" page. This can be found at: http://rk.gsfc.nasa.gov/What's_New.htm

The site has some new areas including conference information, low voltage dropout regulators, and ferro-electric memories (FRAMs) on the memories page.

Wide Field Infrared Explorer (WIRE)

WIRE was a Small Explorer (SMEX) spacecraft which unfortunately had a failure after launch which prevented the spacecraft from meeting any of its science objectives. A programmable device was at the center of this mishap and has been the subject of much discussion. We will present here the failure review board's Executive summary along with some technical discussion about the failure. The main section of the Board's report is at: <http://rk.gsfc.nasa.gov/richcontent/Reports/wiremishap.htm>. Appendix F, which provides the analysis of the failure mechanism, is on-line at: http://rk.gsfc.nasa.gov/richcontent/Reports/WIRE_Report.PDF.

Executive Summary

The Wide-Field Infrared Explorer Mission objective was to conduct a deep infrared, extra galactic science survey. The Wide-Field Infrared Explorer was launched on March 4, 1999, and was observed to be initially tumbling at a rate higher than expected during its initial pass over the Poker Flat, Alaska, ground station. After significant recovery efforts, WIRE was declared a loss on March 8, 1999.

The WIRE Mishap Review Board has determined that the telescope instrument cover was ejected earlier than planned and at approximately the time the WIRE pyro electronics box was first powered on. The instrument's solid hydrogen cryogen supply started to sublime faster than planned, causing the spacecraft to spin up to a rate of sixty revolutions per minute over the twelve hours following the opening of the secondary cryogen vent. Without any solid hydrogen remaining, the instrument could not perform its observations.

The root cause of the WIRE mission loss is a digital logic design error in the instrument pyro electronics box. The transient performance of components was not adequately considered in the box design. The failure was caused by two distinct mechanisms that, either singly or in concert, result in inadvertent pyrotechnic device firing during the initial pyro electronics box power-up. The control logic design utilized a synchronous reset to force the logic into a safe state. However, the start-up time of the Vectron crystal clock oscillator was not taken into consideration, leaving the circuit in a non-deterministic state for a time sufficient for pyrotechnic actuation. Likewise, the startup characteristics of the Actel A1020 FPGA were not considered. These devices are not guaranteed to follow their "truth table" until an internal charge pump "starts" the part. These uncontrolled outputs were not blocked from the pyrotechnic devices' driver circuitry. There has been no evidence or indication of any component failure although component failures were considered in the investigation.

A significant contributing cause of the anomaly was the failure to identify, understand, and correct the electronic design of the pyro electronics box. Design errors in the circuitry, which controlled pyro functions, were not identified. The pyro electronics box design was not peer reviewed, and other system reviews conducted by the instrument design organization did not focus on the electronics box. At the time the Systems Design Review was conducted for WIRE the design of the pyro electronics box was not completed. It is the assessment of the WIRE Mishap Investigation Board that a peer review held during the design process, by people with knowledge

of and expertise regarding pyro circuit design would have identified the turn-on characteristics that led to failure.

A large number of failure scenarios were evaluated during the investigation to determine the cause of the cover ejection. These included; pre-launch, launch, powered flight, separation, software, operations, design and component reliability faults. Based on comprehensive, systematic review of data, it was determined the cover was most likely ejected at the time the WIRE pyro electronics box was turned on due to a transient condition that exists in the pyro electronics during startup. This transient condition is the direct result of the non-deterministic initialization of a Field-Programmable Gate Array (FPGA) that controls both the arming and firing circuits in the pyro electronics.

Although some design attention was given to the startup behavior of the FPGA, the design contained unidentified idiosyncrasies that triggered the cover ejection. The system design did not contain sufficient start-up lockout protection or independent provisions to prevent the FPGA startup operation from propagating to the firing circuits.


The anomalous characteristics of the pyro electronics unit were not detected during subsystem or system functional testing due to the limited fidelity and detection capabilities of the electrical ground support equipment. Post-flight circuit analyses conducted as part of the failure investigation have predicted the existence of the anomaly and it has been reproduced confidently using engineering model hardware.

Some Technical Details

This section will cover some of the key factors surrounding this failure and discuss the principles behind them. These issues are relatively common, some of which have been discussed here previously. As a result of this investigation, a new application note has been written along with a NASA Parts Advisory. These may be found at the following url's: http://rk.gsfc.nasa.gov/richcontent/General_Application_Notes/StartupNote.pdf and http://rk.gsfc.nasa.gov/maplug/Notices/NASA_Advisory_046_ActelStartup.pdf

The design implemented in the FPGA utilized a synchronous reset circuit. If one would assume a random state of all flip-flops during the power-on period, then the circuitry would have a 1 of 4 chance of failing catastrophically, in the WIRE configuration. This idealized model applies here since the synchronous reset relies on a rising clock edge to put the FPGA's circuits into the reset condition. However, real crystal clock oscillators do

not start instantaneously and have a startup delay that can last for tens of milliseconds or more, depending on the oscillator design, the frequency of the crystal, and other factors. One key "other" factor in the WIRE mishap was the rise time of the power supply. The figure below shows the start time characteristic of a WIRE flight spare oscillator as a function of power supply rise time. For these tests I used a linear ramp for the power supply.



Summary of start time characteristics of a flight spare oscillator at 10°C. Start time is a linear function of power supply rise time using a ramp generator as the power supply.

Note the linear relationship between oscillator startup time and power supply rise time. The time measured here is from the power supply startup until the first edge output from the oscillator. It took additional time for the oscillator to stabilize. These 200 kHz oscillators would either put out pulses of incorrect width or drop pulses until the device stabilized. Clearly, care must be taken in any logic design with respect to the reset topology. Normally an asynchronous clear would be applied with a synchronous removal; this would ensure a quick reset function with synchronous removal to prevent metastable states in sequencers.

Using the idealized model mentioned above of a random flip-flop power-on state, we could then hope to see some evidence of failure if the circuit was tested enough times. This does not necessarily apply and the philosophy of "testing in reliability" is again shown to be false. The power-on state of flip-flops, which are not guaranteed to be in any particular state, were shown to be clearly not random.

In particular, it was shown that in repeated power-on trials, flip-flops in the FPGAs (A1020, A1020B) would consistently power-up in the same state, for stable "conditions." This was demonstrated both on the lab bench and indirectly shown on the WIRE Pyro box engineering model in an effort to


replicate the failure. Bench testing showed that the flip-flop's initial state was also a function of power supply rise time. The mechanism here is the circuit design inside of the FPGA, the effect of asymmetrical load capacitances, and other uncontrolled parameters. After numerous (> 30) trials getting identical results with a power supply rise time of about 1 μ s, a very slow rise time was used and the flip-flops powered on in the opposite state.

Another factor involved in FPGA flip-flop initial state determination for WIRE was the amount of time the flip-flop has been powered off. In this part of the study it was shown, as mentioned above, that repeated trials yielded unchanging results. However, after letting the circuit sit unbiased for an extended period of time, hours, the flip-flops would many times power up in the "opposite" state for just one power-on cycle.

A related case was engineering model testing of GLAS instrument electronics. Here a "working circuit" suddenly ceased to function when the +5V power supply was changed. In this case A14100A devices were used. Analysis showed that the change in the power supply's startup condition changed the power-on state of flip-flops. Based on the symptoms of the failure, it was suspected that the flip-flops which perform the "control function" of the FPGA were not being properly cleared. The MODE pin was tied to +5VDC and the change of the power supply resulted in a change of the power-on state of the flip-flops. This is a good reminder for users of Act 1,2,3,XL, and DX technology parts to always verify that the MODE pin is properly biased to ground during startup. If the Actionprobe is used, it will drive MODE high at the appropriate time. For SX devices which have IEEE 1149.1 test circuits, "Revision 0" parts must have an independent clock drive TCLK with TMS high. For revision 1 parts the TRST* pin should be biased at ground.

Another characteristic of the A1020 FPGA used in the WIRE Pyro Box circuitry was that the outputs of the device were direct inputs to the relay and FET drivers. There was no circuitry utilized to block the outputs of the FPGA during the power-on interval. While not inherently the case, many programmable devices, not just Actels or A1020's, have outputs that are not controlled while the device is powering up or initializing. Each device must be analyzed on a case by case basis. It is noted that some future SX devices, currently in design, will have outputs that are "power-up friendly." The drivers will come up in a tri-state condition and resistors, programmed in either a pull-up or pull-down configuration, will hold the output pin at the appropriate logic level until the device is powered up and stabilized.

Again, testing has shown that a device can not easily be "characterized" for start-up transient performance. Like flip-flop power-on state, the size of the transient, including whether one is seen at all, is a factor of the power supply rise time and the amount of time the device has been powered off. According to Actel documentation, it is also a factor of device temperature. For design/analysis purposes, it should be assumed that an unpredictable transient will occur and that the device powers up with uncontrolled I/O's (except for devices especially designed for safe power-on). As a result, logic that blocks the outputs of the programmable device should be used, in conjunction with a power-on-reset circuit, to ensure that critical signals are under control. Similarly, it should be assumed that device inputs may behave temporarily as outputs. This effects circuits such as power-on-reset circuits where an input may source current during the transient, affecting the amount of time that the reset is active for. The figure below shows the transient response of a flight spare A1020 from the Small Explorer WIRE project.



Output transient on start-up of WIRE flight spare S/N 001 A1020 FPGA observed after 24 hours powered off. The bottom trace is V_{CC} while the top two traces are the ARM and FIRE signals. All signals are at 2 volts/division. Attempts to immediately repeat the transient failed, with both critical outputs, Cover and Arm, maintaining logic low output levels with no glitches detected. The probability of a transient is a function of the rise time of the power supply and the amount of time the device has been off, as a result of a "memory effect". The duration of the transient is also a function of the rise time of the power supply. Results on flight spare S/N 002 as well as 3 non-flight A1020B's and another A1020 were similar. Vertical scale is 2V per division. Horizontal scale is 20 ms per division. Note that under these conditions, both outputs were latched in the logic '1' state.

Low Voltage Dropout (LVDO) Regulators

With the move to mixed-voltage systems, the need for low voltage dropout regulators are increasing. The two devices selected for initial test offer the capability of powering small (LM2931CT) or moderate (LM1117T-3.3) loads. Commercial samples were obtained with both models procured in plastic packages. The devices were subjected to TID testing in a Cobalt-60 cell, proton testing at UC Davis, and for the LM1117T-3.3 only, heavy ion tests. The LM2931CT was not tested for heavy ion SEE because of trouble decapping the samples.


The bias and load circuit for these devices are not reproduced here. They are available for download from the internet in .pdf format from: http://rk.gsfc.nasa.gov/richcontent/LVDO_Regulators/Run1_LM2931_LM1117/regulator_3volt.PDF

Cobalt-60 Test


One device of each type was irradiated at 2.84 rad(Si)/sec. In situ monitoring of the current was performed and each device was biased with a 66 Ω load resistor. Additionally, at periodic intervals, the input voltage was swept and the outputs measured. This permits determination of the device's transfer function and dropout voltage without disturbing the devices under test.

Testing of the devices continued until just over 60 krad(Si) was reached with only minimal changes in the devices' parameters and no failures observed. The test was terminated because of facility availability limitations. Future testing will be done at a higher dose rate.


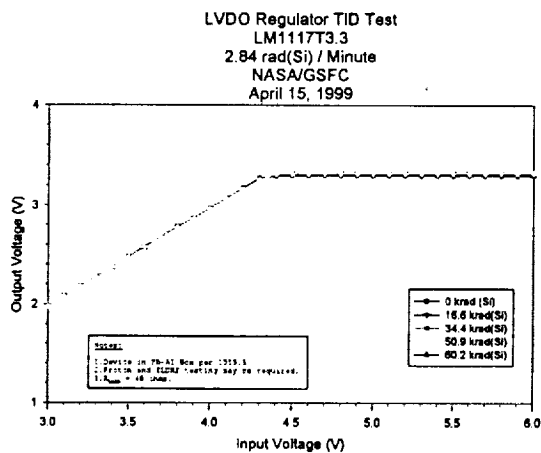
The figure below shows the change in input current over the course of the testing. As can be seen, only small changes were observed. Approximately 50 mA of the current displayed on the graph is from the load on the regulators' 3.3 VDC output.



Similarly, only small changes in output voltage were recorded for each of the devices. In this case, the LM1117T-3.3 did considerably better, showing significantly less than a 50 mV change over the 60+ krad(Si) exposure.



As described earlier, in situ transfer functions were obtained during the irradiation. The data shows that adequate margin exists for this room temperature evaluation for regulation at 3.3 VDC.



Proton Test

The LM1117T-3.3 and the LM2931CT were subjected to proton tests. Two LM1117T and three LM2931CT devices were irradiated with 63 MeV protons. The input voltage for all runs was 5V and output voltages were approximately 3.3 VDC. The initial output voltage of the LM2931CT is adjustable and is set by trim resistors; the LM1117-3.3 comes trimmed to 3.3 VDC. All tests were done at room temperature and annealing effects were not measured.

The chart below summarizes the proton test data (courtesy of Dr. Robert Reed, NASA Goddard Space Flight Center). No significant radiation effects were observed. The following notation is used for the chart:

- I_0 Initial input current
- I_F Input current after irradiation
- OUT_0 Initial output voltage
- OUT_F Output voltage after irradiation

Device	S/N	I_0 mA	I_F mA	OUT_0 V	OUT_F V	Dose krad Si
LM1117T	1	55	55	3.31	3.32	150k
LM1117T	2	55	55	3.31	3.31	150k
LM2931CT	1	49	50	3.18	3.19	50k
LM2931CT	2	49	51	3.18	3.20	100k
LM2931CT	3	50	51	3.21	3.17	150k

* Current includes driving a DC load of 66 Ω .

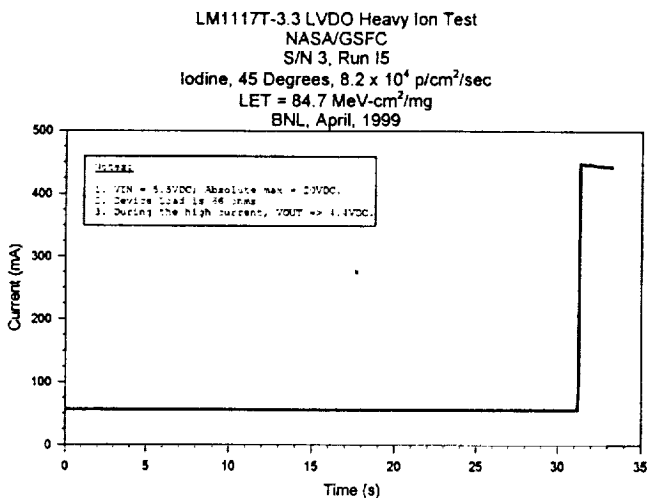
Heavy Ion SEE Test

Three LM1117T-3.3 low-voltage dropout (LVDO) linear regulators were tested with heavy ions at Brookhaven National Labs in April, 1999. The units were procured as commercial parts in plastic packages. This device has a dropout voltage of 1.2V @ $I=800$ mA, making it suitable for producing a 3.3VDC supply from a "standard" 5V logic supply. Most runs were made with a worst-case max logic supply of $V_{in} = 5.5$ VDC, although the device, as specified on the data sheet is capable of tolerating higher input voltages. Some runs were made with a worst-case min logic supply of $V_{in} = 4.5$ VDC.

The devices all showed fluctuations in regulated output voltages during the runs. Start and end values are listed in the table on our www site. It is noted that the changes are small and negligible for standard logic circuits.

All three devices passed at $V_{in} = 5.5VDC$ with Iodine, normal incidence; this is an LET of $59.9 MeV\text{-}cm^2/mg$. All three devices went into a "latchup-like" state at either 30 degrees (LET of $69.1 MeV\text{-}cm^2/mg$) or at 45 degrees (LET of $84.7 MeV\text{-}cm^2/mg$). In this mode, the input current increased by about 400 mA and the output went from 3.3 VDC to approximately 4.4 VDC, until power was removed. S/N LV1 was destroyed.

A typical strip chart of current during a heavy ion irradiation, when the device enters its high current mode is shown in the figure below.



Detailed test heavy ion SEE data can be viewed on-line at:

http://rk.gsfc.nasa.gov/richcontent/LVDO_Regulators/BNL0499/LM1117T-3.3_BNL0499.htm

NASA Lessons Learned

The Lessons Learned Information System (LLIS) is a NASA-wide lessons learned repository. The LLIS offers search capabilities to permit various searches (e.g., NASA Center, date, Project, search string, etc.). Additional categorization capability is under evaluation for future implementation by the LLIS Steering Committee. The NASA Lessons Learned url link, <http://llis.nasa.gov/>, will take you directly to the LLIS Home Page. The Recently Submitted Lessons url link, http://llis.nasa.gov/llis/new_lessons.html, will take you directly to a list of LLIS lessons in time descending order allowing easy access to view the most recently approved lessons.

Is It Safe?

This section will discuss some of the issues involved with designing robust finite state machines (FSMs) in VHDL and some recent developments in a VHDL synthesizer. Additional information can be found in *The Impact of Software and CAE Tools on SEU in Field Programmable Gate Arrays*, to be published in the IEEE Transactions on Nuclear Science, December 1999. Example input and synthesized outputs will be given along with a discussion of the results in the next edition. Time limitations prevent this from being completed here with the proper checking and verification.

Sequencer design can be broken down into several stages. There is the logical design that results in a finite state machine (FSM) which implements the desired function. At this stage logical names are used for each state. In a VHDL implementation, a separate enumerated type is often used, making the code very readable and easily maintainable. A structure of the state machine is then selected. VHDL synthesizers often provide, independent of the HDL code, several options. There are many forms, but a simple register with feedback is commonly used, with the combinational logic providing the next state signals to the state register. The sequence of states is encoded using one of several methods such as a sequential or a gray code. Another popular structure for FSMs is a "one-hot" implementation. The one-hot structure uses one flip-flop per state with exactly one flip-flop in the state register set at any time. The implementation is straightforward and is essentially a shift register initialized such that exactly one of the flip-flops is a 1. This configuration makes decoding of a state trivial and frequently results in a high-performance implementation. The one-hot structure is often used for FPGA designs that are in general register rich; designs implemented in CPLD architectures often use one of the encoded forms.

Independent of the state machine structure, a high-reliability system must not contain any lockup states. These are unused states that can not sequence into a valid state; the state machine is literally locked up. A correctly designed system should never enter one of these unused states. However, a Single Event Upset or other electrical transient or power supply disturbance may cause a soft error and result in an unused state being entered. Since one-hot implementations are often used in FPGAs they will be discussed here in detail. Sequential or gray coded state machines are also a concern, with a detailed discussion of those types of machines discussed in the reference mentioned above.

A simple two-phase, non-overlapping clock generator is used for this example. This machine has four states and can logically be represented in VHDL code by an enumerated type such as:

```
Type StateT Is (Ph1, Ph2, Ph3, Ph4);
```

Using the one-hot encoding, a state assignment is selected by the synthesizer and the states represented in four flip-flops can legally be:

```
0001
0010
0100
1000
```

However, there are 16 possible states of this four flip-flop state vector. Four are used in legal states and 12 are unused. The state machine can transition into any one of 5 illegal states from an SEU; any of the 12 illegal states can occur from a disruption to the power bus or other disturbance or malfunction. The one-hot implementation makes *any* SEU a transition into an illegal state. Since the implementation is essentially a shift register, the fault will never be cleared until the system applies a reset. For example, if the state register, as a result of an SEU goes into state 0101, then we will see the following sequence of states:

```
0101
1010
0101
1010
```

with no hope of recovery. Similarly, if one of the "hot" flip-flops is cleared by an SEU, then the machine will never leave the 0000 state.

There are other structures which help in making a modified one-hot state machine implementation robust. As an example, when a "one-hot" implementation in Actmap is selected, only n-1 flip-flops are used and the all 0's state is a valid state in their implementation. This eliminates the problem of clearing a state bit; the all 0's case is legal and valid. Additionally, a NOR function of all flip-flops' outputs is performed and is input into the D-input of the first flip-flop in the shift register. This tends to clear situations where multiple flip-flops are set by holding off the input of a '1' to the first stage of the shift register. As an example, assume that we have entered, because of an SEU, the state 011 and that the rest of the state machine is well designed. The FSM will transition through the following sequence and then recover:

```
011
001
000
100
```

Similarly, if a state bit is cleared, the NOR function will force the next state to be 100, a valid state.

FSMs using sequential state assignments are also at risk. If the number of used states is not an integral power of 2, then there will be unused states with undefined transitions. Note that use of the VHDL "Others" clause, for any state encoding, will not provide transitions from the unused physical states to a valid logical state. The Others clause operates only on the states defined in the enumeration; it does not operate on physical hardware states. This is disconnect between the abstracted VHDL language and real hardware. There is no mechanism to directly talk about a physical implementation at this level of abstraction; obviously, it can be done using structural coding which eliminates the benefits of the synthesizer and schematics can be used, often a more appropriate tool. Additionally, depending on the tool being used, it's settings, and perhaps even it's revision level, unused states in the state machine that are included in the enumeration may be eliminated by an optimizer that determines that the states are either unreachable or that have no effect on the output.

There is a technique that has been developed, which obviously does not apply to one-hot implementations but can be used, if care is applied, to FSMs using either a sequential or gray code state assignment. This is described in greater detail in the reference but a robust state machine can be coded in VHDL by ensuring that all possible physical states are in the enumeration and that the optimizer can not eliminate them. The preservation of the states and transitions may be possible via synthesizer directives and attributes. In the VHDL domain, a solution would be to force the number of states in the enumeration to be an integral power of two via the introduction of dummy states. Then an "extra" input should force the state machine into a sequence through these states with a dummy output. This will force the states to be reachable and significant.

The problems with robust state machines have been discussed with various vendors. One has added a "safe" mode option to the FSM encodings, since the hardware is not easily and efficiently controlled at the VHDL level, as shown briefly above. This safe encoding feature is controlled via attributes placed into the HDL code.

The synthesizer's algorithm in this mode will add extra overhead since circuitry is needed for the detection of an illegal state and recovery. For this study, I have used Synplify Lite version 5.1.5a. An overview of their algorithms and effects will be given here. Detailed examples of input at the VHDL level and output at the netlist level in the form of a schematic will be in the next edition, as the EEE production deadline is now here. The examples, used here as a framework for the discussion, was a two-phase, non-overlapping clock generator targeted to SX technology. In SX, an "R-Cell" is used as the flip-flop element.

It is obvious that there will be extra combinational logic to detect entry into an illegal state that will assert an error signal. In the implementation examined here, there are two additional R-Cells in the "safe" implementation. These are used for forcing the state machine back into a legal state when an illegal state is detected. The two R-Cells form a simple shift register, with the first R-Cell clocked on the same edge as the FSM and the second R-Cell clocked on the opposite edge. The recovery of this circuit uses the first R-Cell to latch in the signal indicating an error. This is passed to the second R-Cell in the pair, clocked on the opposite edge. This second R-Cell drives the asynchronous inputs to the other R-Cells through 1 stage (in the simple test case used) of combinational logic.

There are two impacts to this implementation. The first, obviously, is that the flip-flop count has increased which will slightly increase the SEU cross-section of the design, since an error in the recovery flip-flops will force the system to change its state erroneously.

The second impact of this recovery mechanism is for timing analysis and margin. When analyzing this circuit, which at the VHDL code level appears to only use the rising edge of the clock, the designer/analyst must also analyze the path from the negative edge-triggered flip-flop to the other devices clocked on the positive edge. This signal must be removed in a half clock cycle. Of course, the worst-case half cycle time period will be less than one-half of the clock period as a result of asymmetry in the clock signal at the R-Cell's inputs. This may be the critical timing path in the design.