RAD-HARD/HI-REL FPGA

Jih-Jong Wang, Brian E. Cronquist, and John E. McGowan
Actel Corporation, Sunnyvale, CA 94086, USA

Richard B. Katz
NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

Abstract

This paper describes the attributes and goals for a radiation-hard and high-reliability Field Programmable Gate Array (FPGA). The first Qualified Manufacturer List (QML) radiation-hardened antifuse FPGA, RH1280, is characterized. Its total dose and Single Event Effects (SEEs) are tested and the results are reported. Trade-offs and limitations in Single Event Upset (SEU) hardening are also discussed.

I. INTRODUCTION

In recent years, FPGAs have gained popularity in aerospace applications. The gate array architecture and high integration level (from thousands to tens of thousands of gates) makes a potent replacement of large numbers of discrete logic devices. The major advantages over more traditional mask-programmed gate arrays are the cost per design and turnaround time. This is exploited in the implementation of changes during system development. The electronic system designs for aerospace applications are full of such examples.

For military applications emphasizing high reliability, there are several choices of FPGAs that were adapted from the commercial marketplace, employing a number of different architectures and technologies. For space applications, military grade FPGAs have been screened by total dose and Single Event Effects (SEE) testing, with some success and a strong lot-to-lot radiation performance variation for some parameters. Actel's 1020s (ACT 1 family) and 1280s (ACT 2 family) are widely accepted in this way by spacecraft electronic designers.

However, the lot-to-lot screening approach cannot keep up with the shrinking of the device geometries and process changes. Proven devices with good radiation tolerance could be obsolete in a few years, with new, more advanced devices not meeting the required radiation tolerance. The evolution of Actel 1020 series shows such a trend. As shown in Table 1, the present device with the smaller geometries, A1020B, has lower total dose tolerance than the older and larger siblings, A1020 and A1020A. The production lifetime of each generation is just four years. Also note, for many applications and environments, shielding is not practical or feasible.

Table 1. Actel 1020 devices with total dose tolerance

Device	Technology	Total Dose	Lifetime
A1020	2.0 μm	>100 krad(Si)	1988-92
A1020A	1.2 μm	100 krad(Si)	1991-95
A1020B	1.0/0.9 μm	<20 krad(Si)	since 93
A1020DX	0.5 μm	N/A	N/A

To fulfill the continuing demand for 1280 and 1020, especially for applications where high radiation survivability is a must, QML-qualified radiation-hard RH1280 and RH1020 devices were developed. These RH devices are fabricated in an Integrated-Circuit (IC) foundry using a 0.8 μm radiation-hardened CMOS process technology.

II. FPGAS IN RADIATION ENVIRONMENT

A brief discussion about the radiation effects on various architectures and technologies of FPGAs is presented in this section. Basically, there are three types of FPGAs: SRAM, EEPROM, and antifuse. Each has a unique technology and architecture, offering trade-offs in functionality, performance, reliability, and radiation-hardness. To compare the radiation effects of FPGAs built on these three technologies, we should examine the radiation response of the most important aspect of them, the 'switching' mechanism built into the interconnect matrices.

The memory switch element is usually composed of a memory bit, such as an SRAM or EEPROM, and a MOSFET pass transistor. The memory bit switches on and off the pass transistor to configure an interconnection node to accomplish the programmability of the FPGA device. Since the memory switch is re-writable, the corresponding FPGA is called re-configurable, or re-programmable. The

antifuse approach is more direct, in that the switch is physically built right on the interconnection node. The switching node is 'open' before programming due to a dielectric layer separating the interconnections. After programming, a highly conductive link is permanently formed at the node. The architecture for antifuse is one time programmable (OTP).

The re-programmability in the memory switch of volatile FPGAs is a major liability in a severe radiation environment. Since the memory-controlled switch is configuring the interconnect of logic circuitry, any error, such as an SEU, a failure from total dose, or an incorrect load, may produce catastrophic hard failures in either the FPGA or other devices. Note that the required number of programmable memory locations for an SRAM-based FPGA is quite high and far exceeds 'user' memory bits; this results in an increase of the SEU cross-section by approximately two orders of magnitude. In contrast, the hard-wired antifuse switch does not suffer any consequence from SEUs, total dose effects, or loading errors.

A recently discovered effect, Single Event Dielectric Rupture (SEDR) [1], may cause failure for both memory and antifuse FPGAs. The failure is due to the electric rupture of a thin dielectric of the antifuse or a gate oxide. The rupture is strongly dependent on the electric field strength in the dielectric and gate oxide, the Linear Energy Transfer (LET), and the incidence angle of heavy ions.

Although the phenomenon has been carefully studied, the detailed mechanism of SEDR has not yet been fully understood. One probable picture starts with a current being initiated by the electrons and holes generated by heavy ions. Under a high electric field (as high as 5×10^6 V/cm in some antifuse dielectrics), the impact ionization effect amplifies the initial heavy ion induced current to the critical value necessary to start a thermal runaway. Experiments and data analyses are still ongoing to validate this picture.

It should be noted that, so far, SEDR is only found in ground tests using very heavy ions generated by accelerators. The calculated failure rate for SEDR is low; an instrumented "K-Fuse" A1280A will be flown on a satellite to attempt to detect a rupture in the space radiation environment.

III. ACTEL ANTIFUSE FPGA

Because of the sensitivity to total dose and SEE effects, available SRAM and EEPROM FPGAs are considered non-radiation-hard and usually are not recommended for applications in severe radiation environments. The dominant FPGA for space flight applications is Actel's antifuse device.

Actel's antifuse device uses a channeled architecture in which rows of logic modules are interspersed with routing channels [2]. There are two basic logic modules, combinatorial (C) and sequential (S) modules. The C-module can implement an 8-input function. The S-module consists of a C-module's functionality followed by a sequential block to implement high-speed flip-flop or latch functions within a single module. ACT 2 and Act 3 both have these modules while ACT 1 has only the C module. Additionally, each family has I/O modules. In Act 1, these consist of simple buffers. There is storage in the Act 2/3200DX and Act 3 devices; latches in Act 2/3200DX and registers in Act 3.

The programmable interconnect switch is formed using a proprietary Programmable Low-Impedance Circuit Element (PLICE) technology, in which an oxide-nitride-oxide (ONO) dielectric is sandwiched between heavily n-doped polysilicon and diffusion interconnects. The programming is accomplished by applying high voltage (about 20 volt) pulses to the ONO. The antifuse switch is then fused by an electro-thermal mechanism. Once ruptured, the antifuse will be soaked by a few hundreds of current pulses to reduce the resistance to as low as 200 Ω.

An on-chip DC-to-DC converter, implemented by a charge pump circuit, is designed to provide a 10 volt power supply for an isolation circuit. These isolation circuits are used in every logic module to isolate the 5 volt logic from the 20 volt programming path while the device is programmed. Once programmed, the isolation circuit is always turned on by the charge pump, which opens the data paths into and out of the logic module.

The total dose effects on the more recent, smaller geometry ACT 1 and ACT 2 products, including A1020B, A1280, A1280A and 1280XL, showed a distinctive failure mode related to the charge pump circuit. Figure 1 shows the static leakage current I_{CC} has a sudden surge at a certain dosage. When the voltage generated by the charge pump degrades as a result of total dose effects, the isolation circuits will fail to pass full logic levels, and the inverters (tens of thousands) in the logic module array will start to pass totem pole currents from V_{CC} to ground.

In the following sections, the first ever QML/RHA certified FPGA device, RH1280, will be introduced by its performance, reliability, and radiation tolerance. The device is assembled in a 172-pin Ceramic Quad Flatpack (CQFP) package.

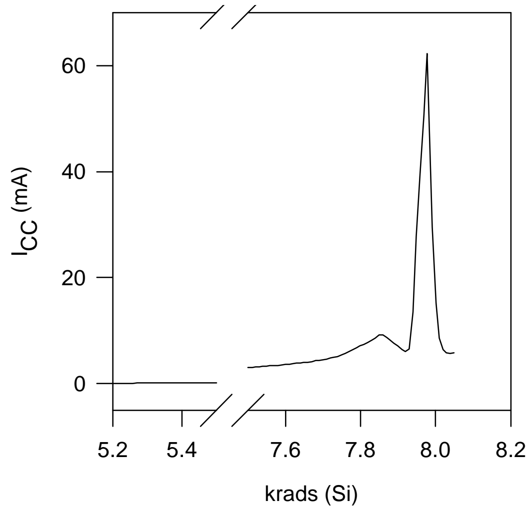

Figure 1 TID Test of A1020Z

IV. RH1280 CHARACTERISTICS

The speed and power consumption of the RH1280 are characterized by measuring various AC and DC parameters. Actel publishes the complete document for customer reference [3]. Here, it is of interest to compare the performance of the RH1280 with its commercial counterpart, the 0.8 µm technology A1280XL, to see if the radiation-hardened wafer fabrication process compromises the performance.

Table 2. Speed and power of RH1280 and A1280XL

	Binning Delay	Standby I_{CC}	Dynamic I_{CC}
A1280XL	90 ns	0.5 mA	24.4 mA
RH1280	71 ns	0.5 mA	19.7 mA

As shown in Table 2, the speed is compared by evaluating a 'binning circuit delay', which basically is a collective measure of the logic module and routing delays. Usually this parameter is used to 'bin' the product to different grades of speed (see reference [2]). At room temperature, the speed grade of RH1280 is 20% better than a standard graded A1280XL. The dynamic current consumption for RH1280 is also 20% better. Table 3 shows the overall speed difference when the timing of a typical design is simulated using Designer™ 3.0.1. Since the parameters were measured from the silicon, the simulation provides a faithful comparison of the register to register and register to I/O delays.

Table 3. The register and I/O delays of a typical design in RH1280 and A1280XL. The min/max values are shown

	Inpad=>Reg	Reg=>Reg	Reg=>Outpad
A1280XL	6.2/15.5 ns	4.1/30.4 ns	4.3/64.1 ns
RH1280	5.0/13.6 ns	3.6/24.3 ns	3.6/51.9 ns

The only significant difference between the RH1280 and the A1280XL (0.8 µm) is the wafer fabrication process. The RH1280 uses a proprietary total dose hardened process provided by Lockheed-Martin Federal Systems' foundry. The front and back end metallization is very advanced. It includes Ti-silicided diffusion and polysilicon, chemical mechanical polish (CMP) planarization, W-plug, and TiN sandwiched Al-Si-Cu metallic interconnections. The better speed and power consumption is probably due to lower parasitic interconnect resistance and capacitance resulting from these advanced metallization technologies.

The reliability of the RH1280 is enhanced by its QML process flow. The production has V-class screening and Q-class technology conformance inspection (TCI) [4]. The electromigration margin is also significantly improved. The traditional contact/via process in the A1280XL often left step coverage as low as 30%. In the RH1280, CMP planarization and W-plug processed contact/via have no steps at all, or in the traditional sense, offer 100% step coverage.

V. RH1280 RADIATION PERFORMANCE

The RH1280 has radiation hard assurance (RHA). Its radiation survivability is extensively tested by total dose effects, dose rate effects, and single event effects.

A. Total Dose Testing

Pre- and Post-irradiation tests were performed on an Advantest Logic Test System at power supply voltages of 4.5 V, 5 V, and 5.5 V at case temperatures of -55 °C, 25 °C, and 125 °C. Total dose irradiation was done using a gamma cell source with dose rate of 152 rad(Si)/sec at room temperature. The bias condition during irradiation has all power pins at 5 V and all input pins grounded. Rebound annealing to simulate the space environment was accomplished by 5 V biased annealing at 100 °C. The testing flow is: pre-irradiation measurement, measure after 300 krad(Si), measure after annealing for 168 hours, measure after another 300 krad(Si), and the final measurement was done after another 168 hours of annealing.

A test design pattern was programmed into the part. DC and AC parameters, as defined in the Actel data book, [2] are measured. A failure is defined as any functional failure or parametric shift outside of specification limits. The results show no parametric or functional failures at any temperature or voltage following the 600 krad(Si) anneal.

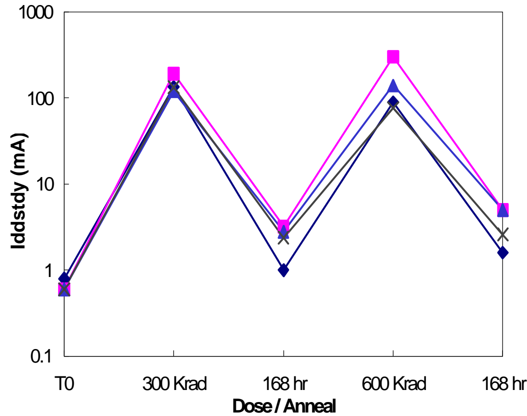

Figure 2 RH1280 Standby I_{CC} at pre-irradiation (T0), and after series of irradiation and biased anneal.

Two key parameters which usually determine the total dose tolerance of Actel's devices, standby I_{CC} and binning circuit delay, are plotted in Figures 2 and 3 respectively. The results of four samples are displayed. Figure 2 shows that I_{CC} increased out of the specification (30 mA) immediately after radiation but always decayed well below the specification after each annealing. Degradation of the binning delay is < 10% at 300 krad (Si), which is the guaranteed tolerance for the RH1280.

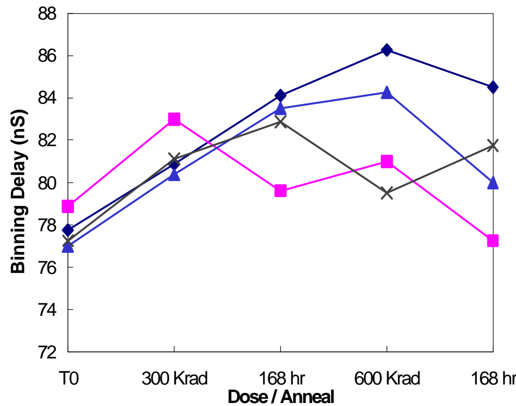

Figure 3 RH1280 binning circuit delay at pre-irradiation, and after series of irradiation and biased anneal

The total dose induced charge pump failure in 1280s was never encountered in the RH1280. This might be explained by the threshold voltage shift (ΔV_t) of the high voltage NMOSFETs in the charge pump circuit (details will be published separately). The hardened gate oxide in the RH1280 reduced ΔV_t drastically. That keeps the output voltage from the charge pump well within the functional range even with very high total dose accumulated.

B. Dose Rate Testing

Dose rate testing was performed using the LINAC at Rensselaer Polytechnic Institute (RPI). The irradiation source was a 17 MeV electron beam with 50 ns pulse width. V_{CC} is biased at 5.5 V for the latch up test and 4.5 V for the upset test. Measurements were done at two temperatures, 40 °C and 125 °C. The sample size was six units.

The DUT was programmed with FIFO strings designed using only the S-modules and FIFO strings using only the C-modules. Both FIFO strings could be exercised simultaneously during testing. There are 270 flip-flops in the C-module FIFO and 592 for the S-module FIFO. The coverage of this design is 90% of the total available modules. Testing was performed by shifting high/low ('1' and '0') states through both FIFO types and recording errors using an event counter on both outputs. The input pattern was an alternating '1' and '0' checkerboard (Zebra) pattern with a frequency of 4 MHz.

The onset (threshold) dose rate for upset was measured at 1.03 x 10^9 rad (Si)/sec. Latch up was never detected up to a dose rate of 3.5 x 10^{10} rad (Si)/sec. No temperature dependence was detected.

C. SEU Testing

Heavy ion tests were performed using the Tandem facility at Brookhaven National Laboratory. The DUT was biased at the worst-case condition, V_{CC} = 4.5 V. The temperature of the device was either nominal (25 °C), 40 °C, 80 °C, or 125 °C. Clock frequencies used varied from static to 4 MHz. The DUT was programmed with the same FIFO strings as those for the dose rate upset test. The three patterns were '1's, '0's, or the zebra pattern.

The test results were very similar to A1280As that were tested previously by Aerospace and NASA/GSFC [5]. For S-modules, a '1' pattern is much more vulnerable than the '0' pattern, while for C-modules, a '0' pattern is slightly softer than a '1' pattern. The difference in upset susceptibility between '1' and '0' is very significant in that the worst-case of either '1' upset or '0' upset will dominate the error rate. Nevertheless, the zebra pattern test is more representative of real applications. Figure 4 shows the cross-section versus LET data of the zebra pattern, together with the aforementioned A1280A data for comparison. For the S-module, LET_{th} = 4 MeV-cm^2/mg, and saturated cross section = 320 μm^2/bit. For the C-module, LET_{th} = 17 MeV-cm^2/mg and saturated cross section = 110 μm^2/bit.

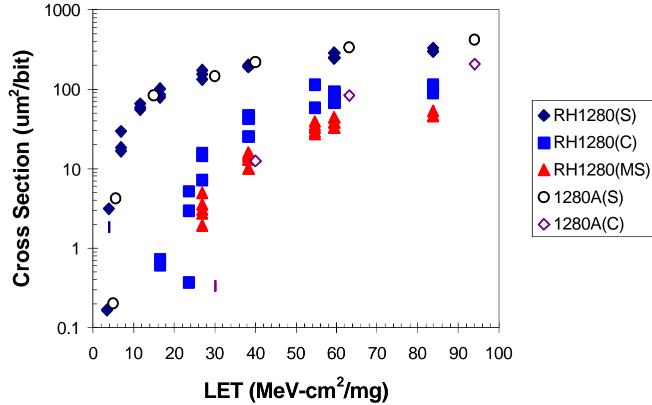

Figure 4 SEU data of RH1280 and A1280A FIFOs made of S-, or C-, or MS (modified S)-modules were tested.

The low LET_{th} of the S-module prompted two actions: the proton SEU testing, and the search for hardening techniques. The proton test data are shown in Figure 5, where the threshold energy is identified as 30 MeV, and the saturated cross section is 6 x 10^{-5} μm^2/bit. The hardening techniques will be discussed in section VI. One of the techniques recommends using a modified, or by-passed S-module, in which the dedicated flip-flop in the S-module is by-passed by the design software. This modified S-module should have the similar SEU susceptibility as that of the C-module. The test results are shown together with those of regular S- and C-modules in Figure 4. The MS-module has LET_{th} = 26 MeV-cm^2/mg and saturated cross section = 51 μm^2/bit.

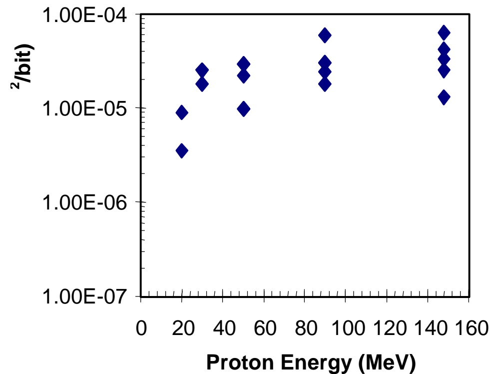

Figure 5 RH1280 proton SEU data, which were measured on FIFOs made of S-modules only.

There are two I/O latches associated with each input/output pin. This latch was also subjected to SEU effects testing. Its test circuit was designed by wiring an I/O latch to an adjacent one to form a FIFO string with twenty (20) I/O latches. Using the same test methodology, the LET_{th} was determined to be 19 MeV-cm^2/mg, and the cross section 100 μm^2/bit. Since the population of the I/O latches in a particular design is always low, the device upset rate contributed by it is also correspondingly low.

All the SEU tests showed no measurable temperature or frequency dependence. The upset was always due to the logic modules with clock upset not detected.

D. SEL Testing

This testing uses the same setup as that of the SEU testing. The device was immune to SEL up to an LET of 177 MeV-cm^2/mg and a fluence up to 10^6 ions/cm^{-2}. Although not tested to the facility limit, no SEL was found in the commercial A1280XL either. The reason is three-fold: first, all the logic modules are protected by double guard-rings in both n- and p-wells; a maximum number of well taps are also implemented to reduce the parasitic resistance in the latch-up path; and epitaxial wafers are used to further reduce the parasitic resistance. RH1280 inherited all those plus two more improvements: it has a thinner epitaxial layer (5 μm) and its retrograde n-well has reduced parasitic resistance.

E. SEDR Testing

SEDR effects were tested using two different test methodologies. The first method was employed in the RH1280's QML certification. No functional failure was detected after the device was bombarded with Au (LET = 80 MeV-cm^2/mg) at normal incidence to a fluence of 10^6 ions/cm^{-2}. The worst case bias, V_{CC} = 5.5 V, was applied.

NASA/Goddard did independent tests using a slightly different setup with a more sensitive monitor. As described in detail in a previous publication [6], the current of the power supply is monitored and I_{DDQ} techniques are used as a diagnostic. The rupture of an antifuse will give an increase in I_{CC} when a bias is placed across it. When each antifuse is ruptured, I_{CC} will increase step by step as shown in Figure 6. This demonstrates SEDR and classifies the failure type (i.e., short, bridging fault). Usually each ruptured antifuse contributes several mA of extra I_{CC}; typically, most ruptures do not cause a hard failure.

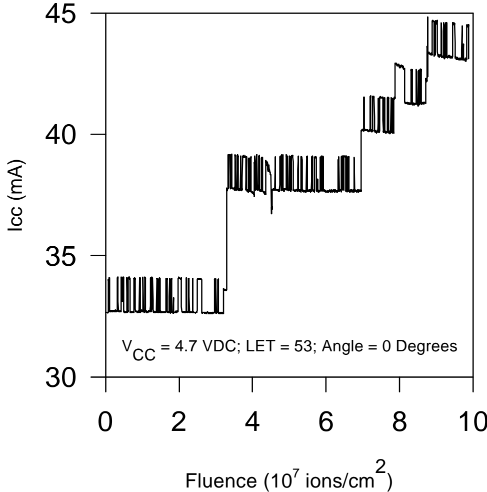

Figure 6 RH1280 S/N 063 Antifuse Rupture

The SEDR testing is very expensive because very costly parts are destructively evaluated with long beam times. Efforts were concentrated on measuring the threshold LET based on the criterion of the first observable antifuse ruptured. Extensive testing determined that LET_{th} > 37 MeV-cm^2/mg at V_{CC} = 5.5 V with fluence > 10^7 ions/cm^2. Note that the test has to be done at 90° incidence; this effect does not follow the cosine law with failures decreasing at increasing angle [1].

VI. SEU-HARD DESIGN

FPGAs provide the flexibility of hardening flip-flops at the logic design level. Following a SEU hardening design guide that has been published [7], a user can implement his design with required hardness. There are trade-offs between the hardness, performance, and number of usable modules. For example, to combat proton SEU for Low Earth Orbit (LEO) missions, the user can restrict flip-flop selection to those made of either C-modules or the combinatorial portion of S-modules. For hardness levels exceeding 37 MeV-cm^2/mg, Triple Module Redundant (TMR) techniques can be applied. Storage, D-type, and J-K-type flip-flops can all be constructed using TMR techniques. However, users need to be familiar with the trade-offs between hardness and the number of usable modules to implement a hardened design effectively.

A. Moderate SEU Hardness

A variety of flip-flops can be created using two C-modules (one module functions as a Master and the other as a Slave) and will achieve SEU tolerance of approximately 1 x 10^{-7} upset/bit-day. Testing showed that the C-modules, the combinatorial part of the S-module (or modified S-module, MS), and the I/O latches have a high enough threshold LET to be immune to protons. To achieve the above moderate levels of SEU tolerance in the RH1280, it is thus necessary to avoid use of all S-module flip-flops (SFFs). Additionally, it is anticipated that single module transparent latches will also show moderate SEU hardness; these blocks are in fact used to construct the master-slave flops.

B. Triple Modular Redundancy

Triple modular redundant (TMR) flip-flops achieve superior SEU performance in the RH1280 family, even if the SFF is used. TMR has been used in the past in many applications and is ideal for effectively SEU-hardening SEU-soft flip-flops in gate arrays. D-type, register, and J-K flip-flops can be constructed. The MUX based architecture of the RH1280 allows an efficient implementation of TMR.

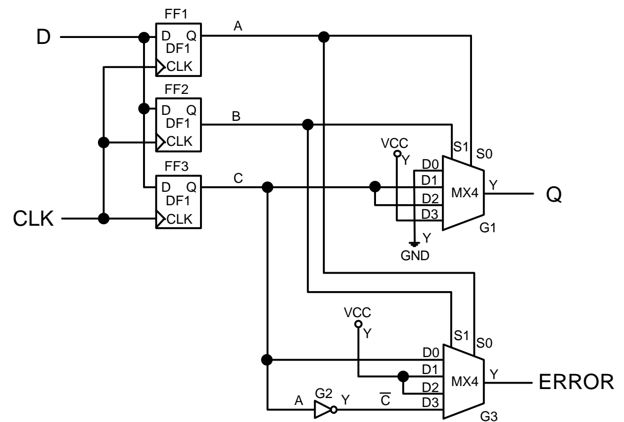

Figure 7 D-type flip-flop implemented with TMR and optional error detector.

D flip-flop: As shown in Figure 7, three D-type flip-flops are connected in parallel to the clock and data inputs. A voter (or majority circuit) is used to create a “hardened” output. The outputs of two flip-flops, A and B, go to the selects of the voter MUX. Should both A and B read logic zero, MUX input D0 is selected and since it is tied to GND, the output of the MUX will read logic zero. Similarly, if A and B read logic one, the output of the MUX will read logic one. Should A and B disagree due to an SEU (or other possible cause), the MUX will select flip-flop C via inputs D1 or D2. We know C agrees with either A or B and thus the MUX has “voted” to produce data agreed on by two of the three flip-flops. An inverter and MUX using similar logic structures can implement an

optional error signal. The output of this MUX will be low so long as all three flip-flops agree. If any flip-flop disagrees, the output of the error MUX will go high. This is useful for testing the redundant logic.

The high degree of data protection using TMR allows the use of SFFs for the three parallel flip-flops while still achieving superior SEU performance. Reference [7] shows how to calculate effective SEU hardness. If the D-type TMR circuit were to feed another similar circuit, such as in a shift register, then an optimal use of the S-module can be achieved by using three separate voter MUXes. Although the number of MUXes has increased, the place-and-route software will "combine" each MUX into an S-module and use the SFF for the flip-flop, thus allowing implementation of the entire TMR circuit into an average of just three S-modules per bit. Additionally, the speed penalty of a separate voter module is eliminated. Lastly, since the voters are redundant, a combinatorial upset on a voter will affect just one flip-flop and be voted out.

Care should be taken when TMR circuits is used. First, the output of the voter may be susceptible to a "glitch" due to a logic hazard. This is not a problem if the TMR is feeding the data input of another synchronous function. However, the TMR output should never feed asynchronous inputs such as flip-flop clocks, clears, sets, read/write inputs, etc. A second consideration for the use of TMR is if gated clocks are used; if the flip-flops are not continuously clocked, then errors can "accumulate" or be "integrated" in the TMR triplets over time. Calculations show that even low to moderate speed clocks will perform well. Ground testing using artificially high SEU rates showed no difference in performance for clock rates ranging from 4 kHz to 10 MHz. A third consideration with TMR is that the proper operation of the parallel flip-flops is not fully testable after the FPGA is programmed because if one flip-flop is not functioning, the TMR circuit will "repair" it. While it is highly unlikely this will occur, 100% fault coverage testing of the TMR circuit's flip-flops can be achieved in a number of ways. First, the optional error monitoring circuits can be used and tied into an OR-tree and brought off chip for monitoring. Other methods use the Action Probe™, which allows the outputs of any two internal logic modules to be monitored on special external pins. Thus all internal parts of the TMR circuit can be read, either on ground or in flight.

C. TMR Circuits with Refresh

For memory elements such as loadable registers, a modified TMR circuit is used (Figure 8). This circuit will constantly "refresh" itself by feeding corrected data back into the inputs of the flip-flops when the enable (E) input is low, permitting error-free data to be held indefinitely; this circuit topology works by placing the error correction function in the feedback loop. When the enable is high, new data is loaded into the TMR triplet. Because of the MUX-based Act 2 architecture, this topology has a very efficient implementation.

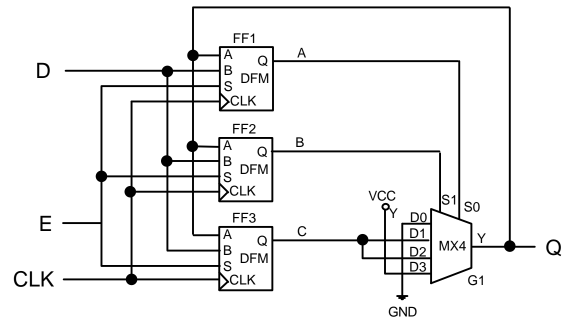

Figure 8 Register element with TMR

A J-K flip-flop TMR circuit (with refresh) can also be implemented (Figure 9). It operates in a similar principle to the register circuit described above, again with the voter circuit inside the feedback loop. The topology also takes advantage of the MUX-based architecture and *combinability*, but since the voter MUX and inverter (for toggling) cannot be combined, this J-K flip-flop is less area efficient than other types of flip-flops.

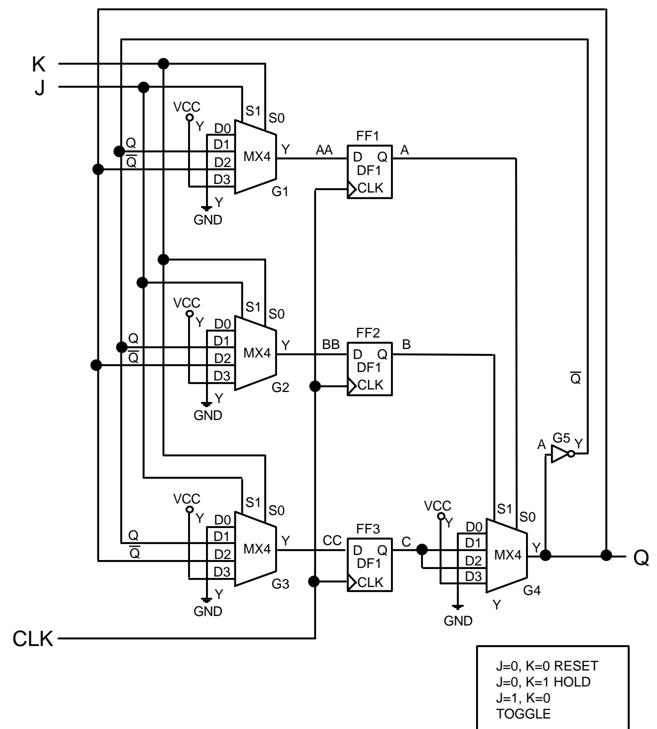

Figure 9 J-K flip-flop with TMR

VII. SUMMARY AND FUTURE TRENDS

In summary, a QML/RHA certified antifuse FPGA was developed by implementing an existing commercial design on a rad-hard process. This device is suitable for all aerospace radiation environments, including LEO, GEO, and deep-space. Logic level hardening methodology to combat SEU can be captured in the user software. However, trade offs between SEU hardness and gate density are inevitable.

A sea-of-modules antifuse FPGA architecture, soon to be introduced by Actel, will improve density, speed, and cost. Since a proven design philosophy was implemented, SEL immunity should be preserved. At the same time, the in-operation electric field in its antifuse is one order of magnitude less than in the ONO antifuse. This should eliminate SEDR. Preliminary SEE and total dose data will be published shortly.

REFERENCES

[1] R. Katz, R. Barto, and G. Swift, "Anti-fuse/FPGA Reliability in the Heavy Ion Environment," NASA Tech Briefs, 1995.

[2] Actel Corporation, FPGA Data Book and Design Guide, 1996.

[3] R. Feldman, "RH1280 Radiation-Hardened FPGA Qualification Book," Feb. 1997.

[4] MIL-STD-38535, MIL-STD-883D.

[5] R. Katz, R. Barto, P. McKerracher, B. Carkhuff, and R. Koga, "SEU Hardening of FPGAs for Space Applications and Device Characterization," IEEE Transactions on Nuclear Science, Dec., 1994.

[6] R. Katz and G. Swift, "Diagnosis of Multiple Faults Using I_{DDQ} Techniques," International Test Conference I_{DDQ} Workshop, 1995.

[7] Actel Corporation, "Design Techniques for Radiation-Hardened Field Programmable Gate Arrays, Rev. 1," 1996, Edited by K. Hayes and R. Katz.

ACKNOWLEDGEMENTS

The authors would like to express their appreciation to Gary M. Swift of the NASA Jet Propulsion Laboratory and to Martha O'Bryan of Hughes/STX for their assistance.