

# A Radiation Hardened Reconfigurable FPGA

Shankarnarayanan Ramaswamy<sup>1</sup>, Leonard Rockett<sup>1</sup>, Dinu Patel<sup>1</sup>, Steven Danziger<sup>1</sup>  
Rajit Manohar<sup>2</sup>, Clinton W. Kelly, IV<sup>2</sup>, John Lofton Holt<sup>2</sup>, Virantha Ekanayake<sup>2</sup>, Dan Elftmann<sup>2</sup>  
<sup>1</sup>BAE Systems, 9300 Wellington Road, Manassas, VA, 20110, USA

<sup>2</sup>Achronix Semiconductor Corporation, 333 W. San Carlos Street, San Jose, CA, 95110, USA  
703-367-4611

ram.ramaswamy@baesystems.com

*Abstract*—A new high density, high performance radiation hardened, reconfigurable Field Programmable Gate Array (FPGA) is being developed by Achronix Semiconductor and BAE Systems for use in space and other radiation hardened applications.<sup>12</sup> The reconfigurable FPGA fabric architecture utilizes Achronix Semiconductor novel picoPIPE technology and it is being manufactured at BAE Systems using their strategically radiation hardened 150 nm epitaxial bulk CMOS technology, called RH15. Circuits built in RH15 consistently demonstrate megarad total dose hardness and the picoPIPE asynchronous technology has been adapted for use in space with a Redundancy Voting Circuit (RVC) methodology to protect the user circuits from single event effects.

As described previously BAE Systems has produced, completed radiation testing, and qualification of a radiation hardened 16M Static Random Access Memory (SRAM) built using the RH15 process [2]. The BAE-Achronix Proof of Concept (POC) RH FPGA device, to be named RadRunner, currently being fabricated uses this same memory cell technology for storing the device FPGA fabric configuration data [3].

## TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. ACHRONIX PICOPIPE TECHNOLOGY.....	2
3. SEE REDUNDANCY VOTING CIRCUIT.....	4
4. RH15 FPGA WITH RVC TEST CHIP RESULTS...5	
5. RESULTS.....	7
6. CONCLUSION.....	7
ACKNOWLEDGEMENTS.....	7
REFERNCES.....	8
BIOGRAPHIES.....	8

## 1. INTRODUCTION

BAE Systems in Manassas VA has been developing and producing radiation tolerant technologies and products for the nearly thirty years, playing a vital role as a merchant supplier of space-qualified microelectronic components and sub-systems for military, civil, and commercial space applications. In partnership with the DoD community, BAE Systems recently completed the modernization of their microelectronic fabrication facility and are currently in production of prototype and flight hardware at the strategic radiation hardened 150 nm CMOS technology node [1]. The key features of their fully scaled, rad hard 150 nm CMOS technology (RH15) are shown in Figure 1.

### RH15 Technology Features

#### Rad Hard 150nm CMOS Technology (RH15) Features

Minimum Feature Size	150nm
Isolation	Shallow Trench Isolation
Device Options	26 A / 70 A
S/D Engineering	Halo with S/D Extensions
Supply Voltages	1.5V / 3.3V
Gate Electrodes	N+ Poly (NFET) / P+ Poly (PFET)
Metal Levels	7 Levels (Planarized BEOL)
Poly & Diffusion Silicide	CoSi <sub>2</sub>

#### Radiation Hardness Assurance Levels

<u>Environment</u>	
Total Ionizing Dose (rd(Si))	>1M
Single Event Upset (errors/bit-day)	<1E-10
Single Event Latchup (MeV-cm <sup>2</sup> /mg)	>120
Neutron Fluence (n/cm <sup>2</sup> )	>1E14
Prompt Dose Upset (rd(Si)/s)	>2E9
Prompt Dose Survival (rd(Si)/s)	>1E12

Figure 1 – Key Features of RH15

This paper describes the new asynchronous high density, high performance, radiation hardened FPGA being built in RH15. The chip has been designed and manufactured utilizing this 150 nm radiation hardened CMOS process employing seven levels of metal. The chip size measures at 21.4 mm x 21.4 mm.

Fabrication, test and packaging will all be performed at the BAE System Manassas facility. Conventional wire-bonding and flip-chip packages are planned. The initial package will be a Column Grid Array 624 pin package.

## 2. ACHRONIX PICOPIPE TECHNOLOGY

Achronix high-performance FPGAs achieve, on average, three times the throughput of existing FPGAs implemented at the same process node. For example, Achronix Speedster FPGAs implemented in a commercial 65 nm process achieve a peak performance of up to 1.5 GHz. This paper


<sup>1</sup> 978-1-4244-2622-5/09/\$25.00 ©2009 IEEE

<sup>2</sup> IEEEAC paper #1243, Version 8, Updated December 10, 2008

introduces the unique Achronix technology which makes this performance possible.

Achronix FPGAs consist of a conventional I/O Frame surrounding a picoPIPE logic fabric (Figure 1). The frame is similar to the periphery of other high-end programmable devices and includes configurable I/Os, clocks, PLL's, etc. The frame provides the off-chip interfaces and forms the boundary between the picoPIPE core and these interfaces. All data entering and exiting the core must pass through the frame.

At the architectural level, the internal picoPIPE fabric is virtually indistinguishable from a conventional FPGA fabric. The only distinction is that the data throughput is substantially increased.



**Figure 1 - Achronix FPGA Architecture**

The picoPIPE fabric is formed from an array of Reconfigurable Logic Blocks (RLB), connected through a programmable fabric. Each RLB is surrounded by Switch Boxes (SB) (Figure 1). The Switch Boxes route global signals across the picoPIPE fabric.

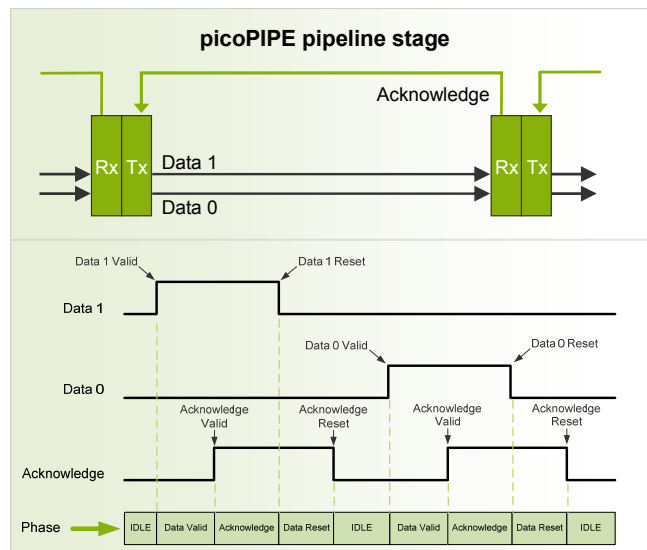
Each RLB contains eight, 4-input Look Up Tables (LUT). In addition to RLB's and programmable routing, the picoPIPE fabric also contains Block Random Access Memory (BRAM). A BRAM is a dedicated two-port memory containing 18,432 bits of RAM. One port provides write access to the memory, while the second port allows for the reading of memory contents. The RadRunner contains 10 of these blocks. Users may configure the BRAM blocks for different data width and depth (36 x 512, 18 x 1024, 9 x 2048, etc.) to meet the needs for their design. The RLB's and routing are implemented using picoPIPE technology, enabling much higher throughput than non-picoPIPE-based FPGA devices. The memories in the RadRunner device are designed to operate at 350 MHz to match the peak device performance.

*Data Tokens*

An important concept to highlight is the 'Data Token'. In conventional logic, a Data Token is a logic value at a clock edge. With traditional logic implementations, data is always

present, but is only valid (and therefore propagated) when a clock edge is received at a storage element. Hence every time data is propagated from one storage element to the next, only a distinct, valid data value or 'Data Token' is propagated.


picoPIPE stages are the atomic elements of the Achronix FPGA fabric. The fabric is capable of implementing any logic function. To implement a logic function, picoPIPE's use explicit Data Tokens, rather than propagating data in response to a global clock edge. Rather than using a global clock to distinguish individual data tokens, picoPIPE stages use explicit local handshakes to pass data. The picoPIPE stages in our FPGA use three wires, two data wires and one enable wire, to implement a four-phase handshake protocol (Figure 2). The key innovation that enables the picoPIPE fabric to operate at high frequency is this new representation of Data Tokens.



**Figure 2 - picoPIPE Four-phase Handshake Protocol**

*picoPIPE Elements*


The basic elements of a picoPIPE are the Connection Element (CE), the Functional Element (FE), the Boundary Elements (BE) and the Link (Figure 3).



**Figure 3 - picoPIPE Building Blocks**

Links connect CE's, FE's and BE's to form pipeline networks (Figure 4). Once combined into networks the picoPIPE implementation exactly matches the functionality

of conventional logic circuit, but is capable of much higher throughput.



**Figure 4 - picoPIPE Pipeline Stages**

Each pipeline stage is capable of holding a Data Token, meaning that picoPIPE technology is highly pipelined by design. In traditional logic designs, adding pipeline stages will change the logic function computed. With picoPIPE technology this is not the case. picoPIPE pipeline stages, such as Connection Elements, can be added without automatically adding a new Data Token into the circuit. This is possible because a picoPIPE stage can contain 'No Data'; the three valid states for a picoPIPE stage are data-1, data-0, or no-data. This is one of the reasons that adding a pipeline stage does not change the logic function computed.

A new Data Token is therefore not introduced when simply adding a new pipeline stage. This means pipeline stages (or Connection Elements) can be inserted anywhere in a circuit without changing the logic function that the circuit computes — in effect, pipelining is transparent to the user.

In synchronous logic, when a pipeline stage (register/storage element) is added, it must still be clocked from a global clock. This automatically inserts a new Data Token, changing the functionality of the logic. With picoPIPE technology, inserting a connection element does not automatically cause a new Data Token to be inserted. Thus the logic is unchanged. Adding pipeline stages does add a small latency delay, but nothing else.

CE's are unlike any existing logic structure. They are capable of being initialized into one of two states, either state holding, or not state holding. When initialized as state-holding, a connection element performs a similar function as a register in a conventional design, and a Data Token is initially created for that pipeline stage. When initialized as not state-holding, it behaves like a repeater. The main difference between a series of un-initialized CE's and a wire is that each pipeline stage is still capable of containing a Data Token, even if it doesn't start with one initially. This property plays a large role in enabling the throughput of Achronix FPGAs to be increased, while maintaining exact logical equivalence to a conventional circuit.

The FE's have functionality equivalent to combinatorial logic. The only difference relates to how ingress and egress data is handled. The local handshaking within a picoPIPE network means the FE's must also handshake data in and out. This handshaking ensures only valid, settled data is propagated.

The BE's get deployed at the boundary where the picoPIPE fabric meets the FPGA Frame. These elements are responsible for converting Data Tokens in the Frame into Data Tokens in the picoPIPE fabric (ingress). They are also used for converting Data Tokens in the fabric back into Data Tokens in the Frame (egress). Therefore every signal entering and exiting the picoPIPE fabric will pass through Ingress Boundary Elements and Egress Boundary Elements, respectively.

When comparing picoPIPE stages to traditional logic several factors need to be considered. The picoPIPE Fabric does not require global clock networks resulting in an area savings when compared to traditional logic. In traditional logic, the global clock networks require silicon area for the large clock buffer circuits and matched length routing tracks to achieve the required low skew timing for high speed operation. The picoPIPE Acknowledge signal wiring uses shorter point to point wiring between transmitter and receiver stages versus the larger area required for implementation of a chip wide global clock network. The two data wires in the picoPIPE implementation increasing the wiring density compared to the single wire needed for traditional logic. The picoPIPE technology derives the increased performance by adding picoPIPE stages into the programmable routing interconnect. While traditional FPGA logic devices selectively use active repeaters for long lines in the programmable routing interconnect, the picoPIPE technology deploys active regeneration of signals at each programmable routing interconnection point.

### *Implementing Synchronous Logic*

The picoPIPE fabric implements synchronous logic by maintaining exact functional equivalence. Several features of Achronix FPGAs and picoPIPE technology enable this equivalence.

FE's provide the combinatorial computation, implementing logic described by the input RTL. CE's provide both the connectivity (local and global routing) and storage (registers) that enable synchronous computation. In the FPGA architecture, LUT's contain FE's and CE's while the Switch Blocks contain only CE's.

The conventional I/O Frame ensures that every Data Token enters the picoPIPE core at a clock edge, and every Data Token leaving the picoPIPE core is clocked out at a clock edge. The functional relationship of inputs and outputs is maintained between the synchronous design specified by the RTL, and the implemented functionality as observed on the boundary of the Frame. Boundary elements ensure that every valid data token in the Frame (data on an input, at an input clock edge) becomes a Data Token within the picoPIPE core. Likewise on the output, every Data Token leaving the picoPIPE core becomes a valid data signal as it is clocked into the Frame (a data value is clocked out of the core).

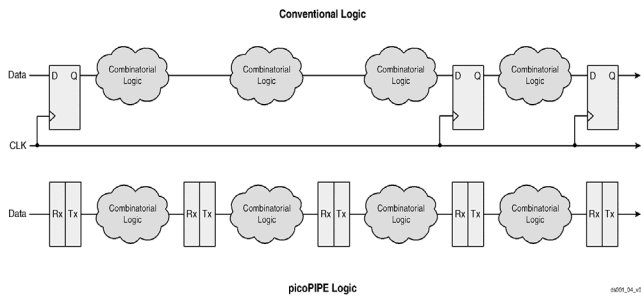
The number of Data Tokens entering, and leaving the picoPIPE core is exactly the same as it would be if the core were implemented with conventional logic.

Another consideration is the number of storage elements in the original design. In traditional, clocked logic each storage element is implemented with a register, which creates its own internal Data Token. As CE's can be configured to be state holding, for every storage element in the original design, a CE is initialized to add an initial Data Token. Thus the number of internal Data Tokens specified in the picoPIPE implementation is consistent with the original design.

*Increased Throughput*

Achronix FPGAs achieve higher throughput compared with existing FPGAs because of the fine-grained pipeline stages. Unlike existing FPGA implementations, these pipeline stages can be automatically inserted anywhere in a design without changing the logic functionality.

As Figure 5 shows, there are often many levels of logic between storage elements. It takes time for data to propagate from the 'Q' register output, through the combinatorial logic and settle at a stable state on the next register's 'D' input. As the clock cannot occur until all data is settled, the clock speed must run no faster than the propagation delay of the longest path in the entire clock domain. Data in every path that is shorter than the longest path (by definition, all paths but one) must wait for the longest path.



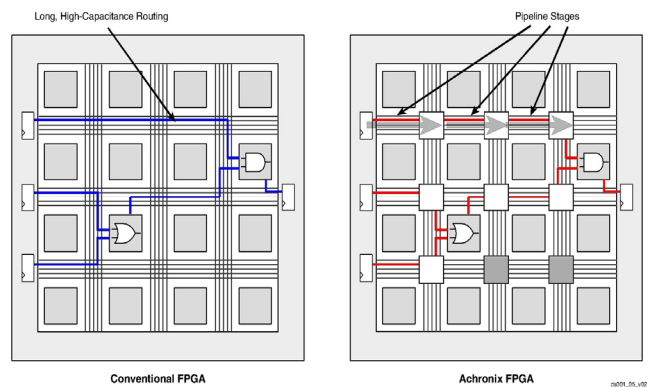
**Figure 5 - picoPIPE vs. Existing FPGA Implementation**

In stark contrast, picoPIPE technology allows optimum pipelining without changing the logic functionality (Figure 5). Each pipeline stage has less logic depth and therefore completes its operation very quickly. This allows the rate of Data Tokens through the logic to be increased, which increases the effective clock rate.

Let's consider the benefits of picoPIPE technology in terms of the physical layout within an Achronix FPGA. Within a traditional FPGA, signals travel on long routing tracks and pass through many routing components (Figure 6). These signals suffer from a high capacitive load; and the larger the FPGA, the longer the paths that need to be traversed.

Additionally, there are many levels of logic between state holding elements (registers).

Within Achronix FPGAs, the built-in pipelining ensures that signals only need to travel on short routing tracks. This reduces the capacitance of the signal at each stage. For larger devices signals still may need to propagate from one corner of the device to the other. While larger devices may have slightly increased latency, unlike other FPGAs, they do not have decreased throughput, as each pipeline stage is capable of holding a new Data Token. Thus the inherent pipelining of picoPIPE technology allows maximum throughput to be maintained, regardless of how large the FPGA becomes. Pipelining also ensures there is only one logic level per pipeline stage, allowing a much faster rate of Data Tokens to be used.



**Figure 6 - Conventional vs. picoPIPE Implementation**

The picoPIPE elements provide the infrastructure for adding transparent pipeline stages. Additional circuits in the periphery of the picoPIPE fabric supply the mechanism for increasing the number of Data Tokens processed by the fabric. During startup of the device this circuitry disables the exit of Data Tokens on the egress side, but Data Tokens are allowed to enter on ingress side. This circuitry can be programmed to a software determined optimum number of additional Data Tokens. The increase in the number of additional Data Tokens depends on the design implemented by the end user of the FPGA. Increasing the number of Data Tokens being processed results in a higher effective processing capacity and external clock rate.

**3. SEE REDUNDANCY VOTING CIRCUIT**

Using the same picoPIPE circuits described above, Achronix has used a circuit design technique to make picoPIPE FPGA circuits tolerate transient faults that can be introduced due to a variety of effects, including radiation effects such as single-event upsets (SEU's) or more broadly single-event effects (SEE's). This circuit has been given the name Redundancy Voting Circuits (RVC).

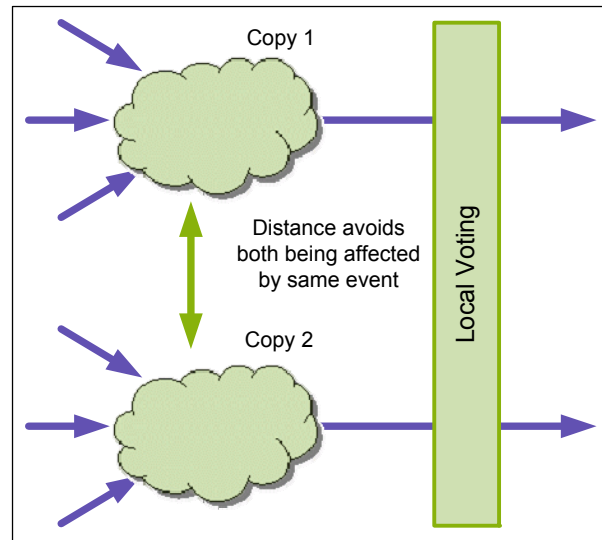
picoPIPE circuits do not use a clock signal for their operation. Since the clock is not present, it cannot be used to filter "glitches" or data hazards. Therefore picoPIPE

circuits and in particular picoPIPE control circuits cannot have any switching hazards on their signals. A transient fault can be thought of as a temporary change in a signal value—a change that causes a “glitch” in the circuit. This error can propagate and create functionality issues, including deadlocks or incorrect data computation.

Radiation effects complicate matters, because the impact of ionizing radiation can disrupt circuit operation at more than one electrical node simultaneously. This is because the radius of impact of the particle might be significantly larger than the smallest features in a modern electronic circuit. Therefore, a “single event” can cause multiple signals that are physically proximate to change simultaneously. Various rad hard investigators have experimentally established this phenomenon, the most direct observation being made in memories where a single particle strike can change the state of multiple adjacent memory bits.

Our SEE immunity is built in through redundancy (Figure 7). Basically two copies of all circuits exist in the logic fabric. The copies are placed non-adjacent to each other with physical spacing of the circuits to avoid a single upset affecting both. Each stage has a local voting mechanism to compare the value of the Data Tokens from each copy. When no SEE has occurred, both copies agree and the local voting mechanism allows the Data Tokens to propagate normally. In the case of an SEE which upsets one half of the redundant circuit, the value of the two copies will disagree and the local voting mechanism will block Data Token propagation. After the SEE energy has been dissipated the circuit will be corrected and the voting mechanism will allow the Data Tokens to propagate after a minimal delay.

Generally it is true that adding more circuitry to implement a logic function increases the total critical area for the overall circuit and thus may potentially make it more vulnerable to radiation effects. However, our approach includes voting and correction on all signals, not just the flip-flops in the design. Hence, vulnerability to radiation requires that the copies of the same signal must be upset simultaneously. If two unrelated signals are upset on different regions of the chip, both errors can be corrected in parallel in most cases. For the implementation to fail, the double error must occur in a localized region of the chip. This improves the vulnerability to radiation effects even though the approach requires more circuitry.



**Figure 7 - Redundancy Voting Circuit**

Redundancy voting happens continuously, not just during the forward propagation of the Data Token. For any signal protected by RVC circuits, the two copies of that signal are monitored constantly for agreement. If any one copy changes due to an SEE, the RVC circuits will prevent that transition from propagating to the next stage. If the Data Token has passed to the next stage and been acknowledged, the circuits are still immune to an SEE changing the data value of one copy in this stage. Because the RVC circuits are continuously active (as opposed to only voting on agreement on the forward propagation of the Data Token), they will observe that only one copy has changed, and block that erroneous change from propagating downstream.

By building the RadRunner FPGAs with RVC, a logic fabric with immunity to SEE’s is realized. The RVC logic fabric combined with the proven BAE Systems radiation hardened (for both SEE and total dose) memory cell will deliver a device that will be extremely useful in a wide range of applications in space flight, military and high radiation industrial applications.

#### **4. RH15 FPGA WITH RVC TEST CHIP**

An effort was undertaken to characterize the single event latchup (SEL) single event functional interrupt (SEFI) and single event upset (SEU) sensitivity levels for the BAE-Achronix RadHard FPGA fabric utilizing the RVC circuits. Initial tests in August 2007 were preliminary and used as a proof of purpose (tester operation). During the preliminary tests, the device under test (DUT) was monitored for single event functional interrupt (SEFI) and for potentially destructive events induced by exposing it to a heavy ion

beam at the Texas A&M University Cyclotron Single Event Effects Test Facility. Future tests will be expanded to monitor single event upsets (SEU). The defined goal of initial radiation testing was to validate that data tokens continue to propagate in the event of heavy ion impact.

The test chip contained a 4x4 array of FPGA tiles implemented with RVC circuits (Figure 8). One dedicated tile had an eight-bit counter (implemented in RVC) that incremented for each output acknowledge for that tile. This allowed for observation of the functioning of the various FPGA configurations at a reasonable frequency with standard low voltage CMOS I/O. Five different FPGA patterns were utilized during the testing. The sample size for the devices under test (DUT) was three including one control that was not exposed to the radiation source. The test chip used during preliminary testing was a test structure and did not contain the actual configuration memory technology that will be implemented in the RadRunner product.




Figure 8 - Test Chip Die Image

*Single Tile Pattern Configuration*

The Single Tile FPGA pattern utilized a single RLB Tile with dedicated picoPIPE “source” and “sink” elements (Figure 10). A “source” element generates a constant Data Token when the receiving picoPIPE element to which it is connected can receive one. A “sink” element generates the output acknowledge and destroys the Data Token. This configuration pattern utilizes six picoPIPE stages:

Total exercised area = 23.5k μm<sup>2</sup>

Total area = 23.5k μm<sup>2</sup>

The single tile toggles between a logical 0 and 1 data token running internally at 290 MHz.




Figure 9 - Single Tile Pattern

*Ring Pattern Configuration*

The Ring Pattern FPGA pattern had the 12 outer tiles configured in a loop of alternating logical 1 and 0 tokens running internally at 37 MHz (Figure 11). This pattern utilized 16 RLB tiles. The inner ring had no observable output, therefore it could not be verified that this token loop was functioning during testing.

- Each tile used:
  - RLB = 23.5k μm<sup>2</sup>
  - Two connection boxes = 2 x 9.4k μm<sup>2</sup>
  - Switch point = 4.2k μm<sup>2</sup> (each is one picoPIPE stage)
- Total area = 744k μm<sup>2</sup>

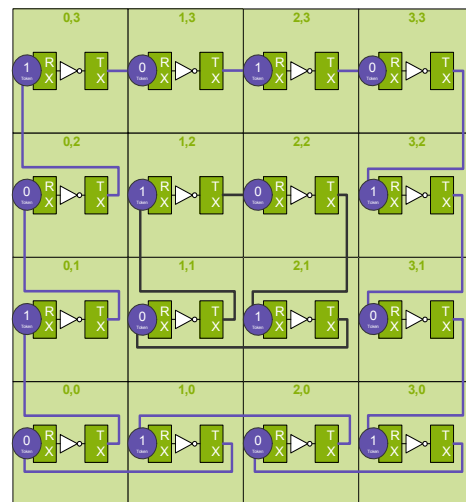


Figure 10 Ring Pattern

### All Tile Pattern Configuration

The All Tile Pattern FPGA pattern utilized all 16 tiles configured in a weave pattern to create a loop of alternating logical 1 and 0 tokens running internally at 112 MHz (Figure 12).

- Each tile used:
  - RLB =  $23.5k \mu m^2$
  - Two connection boxes =  $2 \times 9.4k \mu m^2$
  - Switch point =  $4.2k \mu m^2$  (each is one picoPIPE stage)
- Total area =  $744k \mu m^2$

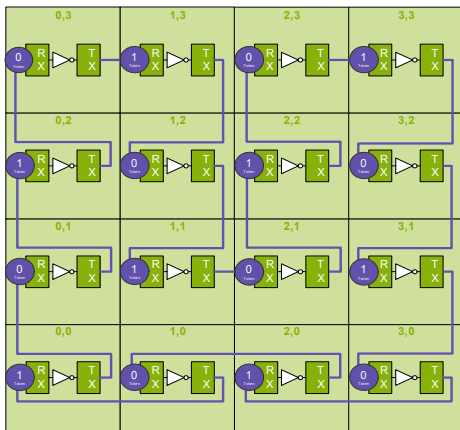


Figure 11 - All Tile Pattern

### Flop1 Pattern Configuration

The Flop1 pattern utilized all 16 tiles configured in a loop with single alternating logical 1 and 0 data token running internally at 135 MHz (Figure 13).

- Each tile used:
  - RLB =  $23.5k \mu m^2$
  - Two connection boxes =  $2 \times 9.4k \mu m^2$
- Total area =  $678.6k \mu m^2$

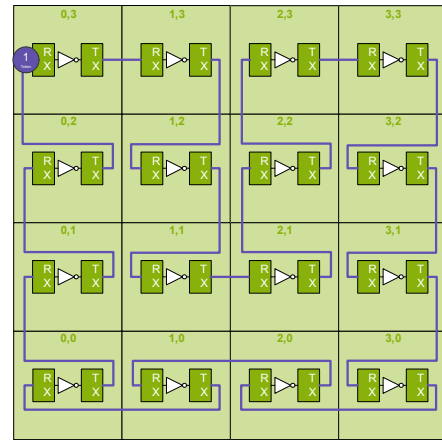


Figure 12 - Flop1 Tile Pattern

### Flop2 Pattern Configuration

The Flop2 pattern utilized all 16 tiles configured in a loop with two alternating logical 1 and 0 data tokens running internally at 68 MHz (Figure 14).

- Each tile used:
  - RLB =  $23.5k \mu m^2$
  - Two connection boxes =  $2 \times 9.4k \mu m^2$
- Total area =  $678.6k \mu m^2$




Figure 13 - Flop2 Tile Pattern

## 5. RESULTS

The ions used for the tests are listed in Table 1. The energetic intensity of each ion type is expressed in terms of its Linear Energy Transfer (LET).

Table 1 - Ions Used for TAMU SEE tests

Ion	Initial LET (MeV·cm <sup>2</sup> /mg)	Angle
Ar	5.7	0
Kr	20.1	0
Kr	29.4	45
Xe	58.1	45

### Test Methods

Devices were electrically tested with the NASA-GSFC REAG (Radiation Effects and Analysis Group) Low Cost Digital Tester (LCDT). The LCDT is a reusable universal digital device tester.

### SEL Analysis

No SEL was observed through the tests with temperatures up to 74 °C and LET of 55 MeV·cm<sup>2</sup>/mg.

### SEE Data and Analysis

For all tests run no SEFI's were observed. With this test vehicle SEU's were not able to be monitored and will be analyzed during the next phase of testing with the POC device. The goals of the testing were achieved as the data tokens were observed to continue propagating for all test configurations.

## 6. CONCLUSIONS

Digital logic designers have long desired a reliable high performance reconfigurable FPGA solution for use in radiation environments. The proven BAE Systems radiation hardened SRAM design and strategically radiation hardened 150 nm epitaxial bulk CMOS technology provide a solid foundation for the POC RadRunner device configuration memory. Combined with the Achronix picoPIPE

technology and proven Redundancy Voting Circuit methodology to provide built in SEE protection a new high performance reconfigurable solution will soon be made available to designers. Digital designers will be able to continue describing their design with familiar VHDL or Verilog Register Transfer Level coding methods and synthesis tools.

## ACKNOWLEDGEMENTS

The SEE testing was completed by Melanie Berg and Hak Kim of MEI Technologies, Inc. with beam time provide by Ken LaBel of the NASA Electronic Parts and Packaging (NEPP) program.

## REFERENCES

- [1] L. Rockett et. al., "The Foundry Modernization at BAE SYSTEMS in Manassas VA – A Progress Report", Government Microcircuit Applications Conference (GOMAC) Digest of Technical Papers, Vol. XXIX, March 2004.
- [2] Tri Hoang; Ross, J.; Doyle, S.; Rea, D.; Chan, E.; Neiderer, W.; Bumgarner, A., "A Radiation Hardened 16-Mb SRAM for Space Applications", Aerospace Conference, 2007 IEEE Date: 3-10 March 2007.
- [3] R. Manohar et. al., "Development of Reprogrammable, Low Power, High Density, High Speed RADHARD FPGAs using Proven Commercial Technology and RADHARD-by-Process Techniques", Government Microcircuit Applications Conference (GOMAC) Digest of Technical Papers, Vol. XXXII, March 2007.

## BIOGRAPHIES

**Shankarnarayanan Ramaswamy** is a senior engineer in the VLSI Customs Product Design and Development area at BAE Systems in Manassas VA. He has over 20 years of semiconductor experience ranging from circuit design, technology development, technology transfer, product engineering, test/characterization and project management. He earned his M.S. degree in electrical engineering and his MBA from University of Kentucky. He is a Senior Member





of the IEEE.

**Leonard Rockett** is the manager of the microelectronics business area at BAE Systems in Manassas VA. He has over 30 years of both engineering and management experience primarily in the areas of radiation hardened technology development and microcircuit design. He has authored or co-authored over 60 published technical reports and papers. He currently holds 15 U.S. patents plus he has other patents pending. He earned his B.S. degree in electrical engineering from Howard University in Washington DC, and his M.S. and Ph.D. degrees in electrical engineering with a specialization in microelectronics from Columbia University in the City of New York. He received his professional certification as a Program Manager from Lockheed Martin. Dr. Rockett is a Senior Member of the IEEE.



**Dinu Patel** is the program engineering manager responsible for process enhancement and yield improvement at BAE Systems in Manassas VA. He is the technical lead for the RHAX250-S installation and qualification program.



He has over 35 years of both engineering and management experience primarily in the areas of process development and product engineering of commercial and radiation hardened products. He has authored or co-authored several published technical reports and papers. He earned his M.S. degree in electrical engineering from N.C. State University. Mr. Patel is a Senior Member of the IEEE.

**Steven Danziger** is the semiconductor technology diversification Program Manager with BAE Systems. He is the program manager for the RHAX250-S installation and qualification program. He earned a B.S. in Physics from Queens College City University of New York and a M.S. in Physics from The Pennsylvania State University. He has over 23 years of experience as an engineer and a manager in semiconductor development and production for various defense and commercial semiconductor




organizations. He has over 30 published technical reports and papers and holds one patent.

**Rajit Manohar** is a Founder and CTO of Achronix Semiconductor Corporation. As an internationally recognized pioneer in asynchronous semiconductor design and implementation, Rajit Manohar has spent over ten years working in the field of asynchronous VLSI. Prior to founding Achronix as the CTO, Dr. Manohar was an Associate Professor of Electrical and Computer Engineering at Cornell University. He was the recipient of an NSF CAREER award as well as numerous teaching awards at Cornell. He was one of the principal designers of MiniMIPS, the first high-performance asynchronous microprocessor. While at Cornell, Dr. Manohar's group successfully demonstrated several asynchronous VLSI systems in Silicon, including the first microprocessor for sensor networks, an event-based asynchronous chip-multiprocessor, and a pipeline-configurable asynchronous FPGA. He holds numerous patents and is widely published in the area of asynchronous design. Dr. Manohar has received numerous awards, but most recently was selected by MIT's Technology Review magazine as one of the world's "Top Innovators Under 35 in Science and Engineering". Dr. Manohar received his B.S., M.S., and Ph.D. in Computer Science from the California Institute of Technology.




**Clinton W. Kelly, IV** is a Founder and the Vice President of Advanced Research for Achronix Semiconductor Corporation and leads hardware research and development activities at the Company. For the last six years, Dr. Kelly has led numerous research and product development programs related to ultra-low-power asynchronous microprocessors and FPGAs. To date, Dr. Kelly has successfully led a variety of processor and FPGA projects in 180nm and 90nm CMOS. Before founding Achronix, Dr. Kelly designed and implemented an ultra-low-power asynchronous microprocessor called SNAP, which is used in sensor networks. This processor was successfully fabricated in 180nm CMOS and demonstrated orders of magnitude less energy than conventional processors currently used in sensor networks. Dr. Kelly also designed and implemented an asynchronous multiprocessor




called the Network on a Chip. Dr. Kelly holds numerous patents and is widely published. Dr. Kelly earned a Ph.D. and an M.S. in Electrical and Computer Engineering and a B.S. in Computer Science, all from Cornell University.

**John Lofton Holt** is a Founder of Achronix Semiconductor Corporation. He is the Chairman and CEO and has worked as a successful consultant and entrepreneur for over 10 years, and has worked in the field of Electrical Engineering for over 15 years. Before co-founding Achronix Semiconductor in 2004 Mr. Holt led his company, Integrated Strategies Group (Formerly Saber Security Solutions), a management technology and venture consulting firm in Washington DC. Under Mr. Holt's leadership, the firm built a multi-million dollar business with Fortune 100 and startup clients in the financial services, manufacturing, aerospace, semiconductor, telecommunications, and pharmaceutical industries. Prior to co-founding Achronix Semiconductor, Mr. Holt was a manager of several strategy and technology consulting practices at PricewaterhouseCoopers and Booz Allen Hamilton that advised Fortune 100 companies as well as startups on their execution strategy and technology investments. Mr. Holt started his career at NASA's Goddard Space Flight Center in 1989 and worked as a design engineer until 1997, focusing on the design, integration, and testing of LIDAR-based flight instruments for remote sensing applications. Mr. Holt is educated as an Electrical Engineer and Physicist with a Bachelor of Science in Electrical Engineering from Princeton University and a Masters of Science in Electrical Engineering from Johns Hopkins University.



**Virantha N. Ekanayake** is one of the original founders of Achronix, but his passion for research led him to the Johns Hopkins University where, until rejoining Achronix in 2006, he served as an Assistant Professor in the Department of Electrical and Computer Engineering. Dr. Ekanayake's research group focused on Asynchronous VLSI and Architecture. Dr. Ekanayake has numerous publications arising out of his work on high-performance asynchronous pipelined memory systems, and ultra-low-power



asynchronous processors optimized for sensor networks. From high-level system specification to back-end custom and semi-custom circuit design, he has been involved in the full design flow for several 180nm and 90nm processors and FPGAs. Outside of academia, Dr. Ekanayake has extensive experience in industry, hailing from Netscape Communications Corporation, Telcordia (formerly Bell Communications Research), and, most recently, Intel Corporation. Prior to earning a Ph.D. in Electrical Engineering from Cornell University, Dr. Ekanayake graduated magna cum laude from Dartmouth College with a B.E. in Electrical Engineering, and received a Masters of Science in Electrical and Computer Engineering from Cornell University.

**Daniel K. Elftmann** is the Director of Strategic Marketing for Industrial, Military, and Aerospace at Achronix Semiconductor. He provides strategic market direction on product features and tools by working closely with customers to define specific market solutions. A veteran of the semiconductor industry, Elftmann held various technical management and marketing positions throughout the industry for nearly two decades. Dan spent six years as the director of product engineering for Actel Corporation after working in various other positions such as world-wide technical manager and field applications engineer. Prior to Actel he served as a principal hardware design engineer for Computer Network Technology where he led the effort to create Verilog HDL Design flows for CPLD and FPGA devices. He also held various other positions at Bull Micral, Control Data Corp., ETA Systems Inc., and FSI Corp. Elftmann holds a bachelors of science in electrical engineering from the University of Minnesota.

