

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

5 Gbps Radiation-Hardened Low-Power Pulse Serial Link

Permalink

<https://escholarship.org/uc/item/71p507v7>

Journal

IEEE Transactions on Nuclear Science, 63(1)

ISSN

0018-9499

Authors

Miller, Merritt P
Magazzu, Guido
Brewer, Forrest D

Publication Date

2016-02-01

DOI

10.1109/tns.2015.2507551

Peer reviewed

5Gbps Radiation-Hardened Low-Power Pulse Serial Link

Merritt P. Miller, *Member, IEEE*, Guido Magazzu, and Forrest D. Brewer, *Member, IEEE*

Abstract—Presented is a high-speed, radiation hardened by design, physical layer data link. The 5Gbps link is implemented in an 130nm process occupying relatively small area and having low power (60mW total). The data rate, radiation hardness, and power consumption are all enabled by asynchronous circuits utilizing a pulse-base wire encoding that, in combination, dramatically simplifies data recovery while increasing timing tolerances in both transmitter and receiver designs. This means that neither the transmitter nor the receiver uses a high-speed clock, instead, small asynchronous state machines handle the transmit and receive functions as well as serialization and deserialization. These small structures contribute to radiation hardness by allowing for efficient modular redundancy at high speeds. The design is also hardened for long environment lifetime (very high total irradiated dose of greater than 300MRad Si). This paper specifically addresses the design implementation trade-offs and high-speed design methodology for the link with simulation results as well as preliminary test data.

Index Terms—Asynchronous circuits, High energy physics instrumentation, Nonlinear circuits, Radiation hardening

I. INTRODUCTION

Data transmission and available bandwidth are major concerns for the development of any electrical system. Sensors in high energy physics experiments, occupants of high-radiation environments by design, benefit greatly from enhanced data transmission rates and power figures. Improvements in the amount of transmitted data in a given amount of time, energy, or number of wires lowers the power and material resources required to achieve a specified data gathering goal. Higher data rates for a given resource allotment potentially enable more complex or thorough data sets [1].

Serial data links have become the key technology for high performance data motion. They use a small number of wires, typically just a single pair, strongly reducing problems of wire-to-wire skew, a key limiting factor for parallel buses; restricting skew requires precise routing, environment analysis and cable choice. Conventional serial links are a common form of data transmission such as USB, SATA, Display Port, to the more specialized Space Wire and similar technologies. While they limit the concerns of wire to wire skew, conventional

serial links have very tight timing constraints and require high performance, low-jitter clocks all of which are sensitive to noise sources and other environmental concerns.

In particular, serial links are difficult to make robust to radiation. Both total dose effects, causing parametric shift, and soft errors, corrupting data, are harder to compensate for when circuitry is pushed to maximum rate operation. Suppressing transients by slow response, as suggested for memory upset hardness in [2], is not a viable option for high rates. Soft errors can manifest as both corrupted data (where the information in the signal is wiped out) and corrupted timing information, typically skew between the clock and data stream. Timing corruption can come through strikes to PLL control logic, resulting in microseconds of phase error, completely disrupting the link [3]. Corrupted data can be handled with error correction codes, and/or protocols that re-send information. The trade-offs between these options depend on link application and available bandwidth. Packet resend and error correction can be implemented in higher levels of the protocol stack, allowing for extra layers of protection. Timing errors, on the other hand, are the sole responsibility of the low-level hardware and upsets to the clock recovery circuitry can cause long periods of link down-time.

In this work, we present a serial link technology specifically targeted for very high radiation environments. This technology is the product of co-design of the serializer, transmitter, receiver, and deserializer as a whole. This approach generated an unconventional line code, as well as a set of self-timed asynchronous state-machines to implement it. To the external user, the interface appears to work like a standard synchronous parallel link, 8 bits wide with a maximum data strobe rate of 625MHz – a rate much relaxed from the 5GHz bit rate. The benefits of the approach are many, including the lack of a high-speed clock, reasonable power budget (~60mW total between transmit and receive), small layout area (normal differential pad size extended $50\mu m$ in 130nm process technology), and the ability to idle the link in a low power (<1mW) mode without a delay for entering or exiting this mode.

II. RADIATION TOLERANCE AND DESIGN PROCESS

Both total dose damage and malfunction due to single event upsets were considered in the design of this transceiver. The design methodology consists of making an estimation of the radiation effects on components and confirming that the transceiver continues to function correctly across the range of expected damage. The 130nm process node is not only widely used and well-studied, it is also well characterized for radiation

Manuscript received June 17, 2014; revised September 30, 2015.

This work was supported in part by the U.S. Department of Energy grant DOE-CERN CMS SLHC "Protocols and IP-Cores for Control and Readout in Future Higher Energy Physics Experiments"

Merritt P. Miller, and Forrest D. Brewer are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara CA 93106. (e-mail: merrittmiller@ece.ucsb.edu, forrest@ece.ucsb.edu)

Guido Magazzu is with Istituto Nazionale di Fisica Nucleare Sezione di Pisa, Edificio C - Polo Fibonacci Largo B. Pontecorvo, 3 - 56127 Pisa (e-mail: guido.magazzu@pi.infn.it)

environments [4], [5], [6], [7], [8]. Radiation tolerance of the proposed technology relies heavily on the characteristics of the 130nm process for total dose hardness, and relies on circuit design for hardness against upset. While this model is becoming more common in deep sub-micron radiation hard circuits in high-energy physics, it is counter to the model used for traditional work in space environments, where the process was also looked at as a source of upset hardness [9].

Verification of the radiation tolerance of the circuits relies on characterization of the expected effects, simulation of the impact(s) of radiation damage, and comparison to the expected behavior of the transceiver as a whole. Circuit hardness evaluation starts with approximating the potential damage mechanisms for the devices. Existing studies were heavily relied on for data, especially the data from [5], [6]. Additionally a test chip was fabricated, from which a further characterization that carried out to 450Mrad to supplement the existing data (See Section II-B as well as Fig. 2 and Fig. 3). Design changes were made to reduce the likelihood that there would be an unusually sensitive component in the system. Finally, using the characterizations already gathered, simulation of different damage cases was carried out to verify the circuit behavior. In the event that the system failed verification, the design was modified with the information gathered in earlier verification attempts, creating an iterative process. A number of architectures were considered in the search for a system-level solution.

A. Total Dose

Core transistors manufactured in the 130nm process are known to have a very high tolerance to total radiation dose [4], [5], [10]. At this process node, the primary effect of damage is a witnessed increase in leakage current due to an activated side-channel in the transistor [5], that is a leakage mechanism related to charges trapped in the Shallow Trench Isolation (STI) oxide at the sides of the transistor. The lack of sensitivity to radiation away from the sides of the transistor stems from the fact that in sufficiently thin oxides trapped charges are known to neutralize via a tunneling mechanism [11]. Because damage is concentrated on the edges of the channel, the transistors are laid out with higher width to length ratios than a non-radiation tolerant design would require. This reduces parametric shift due to radiation damage. It is possible to create fully enclosed source or drain transistors that have no side leakage channel (no shallow trench isolation), however, such transistors have no variance characterization in the technology, they would not fit with the scheme of Total Ionizing Dose (TID) mitigation via variance simulation of augmented models.

Design qualification is based on simulation at the transistor level taking foundry data and augmenting it with available data on total dose damage for the 130nm process at the transistor level. This is similar to existing compact modeling methods such as [12], [13], [14]. The spice-compatible model used is shown in Fig. 1. Foundry variation models are retained with radiation induced damage adding to the expected variation.




Fig. 1. Equivalent circuit model of a transistor damaged by total dose. V_{shift} models the threshold shift due to trapped charge under the gate. The non-ideal current source I_{leak} models the leakage channels at the side of the transistor due to trapped charge. The resistors R_{gm} model reductions in transconductance due to radiation damage

B. Qualification of Total Dose Model

This method of modeling was qualified by creating a test structure and verifying its behavior under X-ray irradiation. The results of this test were first presented as part of a talk at the 32nd Hardened Electronics and Radiation Technology (HEART) Technical Interchange Meeting[15]. The exposure used 40KeV X-rays to a total dose of up to 450Mrad. The test circuit is a ring oscillator that was designed to operate at roughly 1GHz – the oscillator was uncontrolled so that the output frequency could be used to detect changes in delay. Because the existing studies showed a lowering of thresholds for both NMOS and PMOS in this process, the ring oscillator was deliberately constructed so that PMOS timing dominated, meaning slow-downs would represent a measured voltage shift in PMOS. During irradiation both frequency of the oscillator and current draw of the whole chip were monitored. Fig. 2 shows the measured ring oscillator period for three irradiated samples. The highest dose sample was tested to 450Mrad, confirming the gradual degradation model. Measurements of the period suggest that the oscillator slows on average 0.18ps/Mrad; roughly equivalent to .23mV of threshold shift per MRad. Additionally measurements of leakage current for the test chip were made. Due to the nature of the test chip, current demand for a single transistor, or gate could not be isolated, but whole chip demand, shown in Fig. 3, correlated well with the results of [5] allowing that data to be used for single transistor leakage.

C. Single event upset

Single event effects, both single event transients and single event upsets, are a large problem for chips in high-energy physics experiments. These experiments exist to record data relating to ionizing particles; a large number and variety of particles are likely to strike the transmitter while it is operating.

Typically cells can be characterized by the amount of charge required to cause an upset, Q_{crit} ; cells with higher Q_{crit} are harder to upset [2]. Finer scale technology is typically




Fig. 2. Ring oscillator period VS total irradiated dose. The 1ns nominal period slows roughly 0.18ps for each 1MRad of total dose which correlates to about .23mV/MRad voltage shift. The measured shift is close to linear, supporting a linear model of degradation.

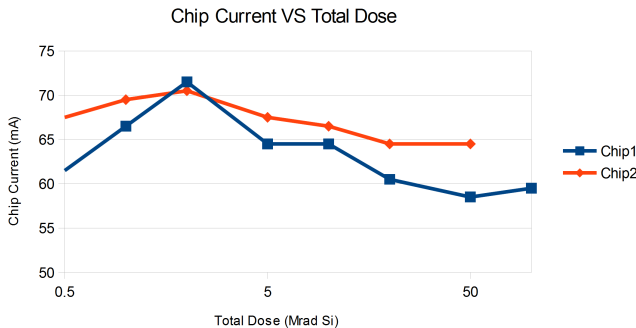


Fig. 3. Measured whole-chip current leakage for a test 130nm chip under X-ray Irradiation. Peak leakage measurements agree with the findings in[5]

easier to upset requiring either error correction or special designs that prevent upset [2], [7]. The high-performance demands of the serializer and deserializer limit cell choice thus system-level mitigation is used instead: modular redundancy, and limiting both the life time and amount of stored state. Triple Modular Redundancy (TMR)[16] in the circuits was chosen as the primary means of error correction. Because the implementation is asynchronous and high-rate most signals are timing critical. All signals are triplicated and voted on at least once per cycle of the data strobe. In the 130nm process, single latches are known to be vulnerable to upset. Multiplicity is one of the few ways of storing information that reduces upset risk [7].

III. ARCHITECTURE

A. System interface

On both the transmit and receive interfaces, the system presents a 8-bit data bus, a data strobe, and a reset line. All of these signals are presented in triplicate to meet the redundancy requirement. On the receive side the data bus and data strobe are both outputs, with the data strobe generated directly from the data stream. On the transmit side the data strobe and data bus are both inputs to the system again triplicated to meet the redundancy requirement. On both sides, the reset signal is an input and also triple redundant. The reset signal is used to put




Fig. 4. Interface of the transmitter and receiver. The interfaces on both sides have a data strobe and data buses (all triplicated) as well as an asynchronous reset.




Fig. 5. 2 of 3 voter cell used extensively in this implementation. This cell has highly predictable timing allowing the use of triple modular redundancy in timing critical paths.


the transmitter and receiver into a known state at boot-up or in the event that a higher-level system catches an error. This system interface is described in Fig. 4.

B. Data strobe and Reset Triplication


In this system, all signals are triplicated including data strobe and the asynchronous reset. This means that voting on data strobe and asynchronous reset occurs within the system. Such voting may give many designers pause, due, in part, to concerns with adding logic to timing critical paths. Key to allowing both a voted high-speed data stream, and enabling triplicated, timing critical systems was the implementation of a glitch-free, fast voter. This voter consists of a 2-2-2 And Or Invert (AOI) gate followed by an inverter wired to be a voter, as seen in Fig. 5. By using this same gate on all external signals, there is little issue of skew, since each data line, as well as each data strobe line, should see similar delay under normal operation. When one of the three lines disagrees with the other two, the forward propagation in the voter is delayed. This increase in delay is projected, via simulation, to be no more than 25ps when only 2 of the three triplicated inputs are in agreement, as compared to all three in agreement. This limited skew was achieved in part by the relative sizing between the AOI gate and the inverter.

IV. ENCODING

The proposed encoding is a two-wire encoding, where transits are allowed on only one wire at a time. A pulse on one wire is one combined data/timing symbol and a pulse on the other wire is another symbol. There is a '1' wire and a '0' wire; a pulse on the '1' wire indicates the arrival of a 1 bit; a pulse on the zero wire indicates the arrival of a 0 bit. This encoding



(a) Idealized pulse stream showing 1-wire and 0-wire



(b) Simulated pulse stream; blue pluses are for the 1 line, gray for the 0 line

Fig. 6. An example pulse-stream. Both 1 and 0 lines shown. These pulses encode 1011 0010 in two 4-bit nibbles. In the constructed circuit, the serializer sends 4 pulses per data strobe edge. This allows timing recovery to be a simple state-machine that toggles once for every 4 pulses.

makes the decoding state-machine very easy to construct, as well as making pre-emphasis an easy task. Fig. 6 shows the encoding example. This encoding has been studied before for high-speed links; a skew compensation circuit was described in [17]. There was also an exploration of such signaling for fast buses in Field Programmable Gate Arrays in [18]. Both works were inspired by the use of pulse surfing based pipelines described in [19]. This encoding has the advantage of easy clock recovery, similar to the Data/Strobe encoding used in SpaceWire [20], [21].

A. Timing constraints for pulse encoding

The pulse system is built using self-resetting domino logic gates. Self resetting domino is the native logic from pulse-timed systems, and enables rapid pulses to be maintained within the system [17], [19], [22]. Rapid pulsing simplifies high-speed design by allowing the extensive use of fast set-reset latches for all jobs that a D-flip-flop would do in traditional clocked logic. The whole transceiver is built out of this logic style, including transmit amplifier and receive amplifier. Signals coming from traditional, level-based systems use edge detectors to create pulses, while set-reset latches are used to change pulses back into levels.

The upper bound on pulsed circuit delay is set by the constraints for properly generating and detecting pulses. Each pulse must be both detectable and distinguishable from the next one, thus placing a limit on pulse width and pulse spacing. In the design style used, the minimum pulse space is identical to the characteristic pulse width; the delay time of the self-reset feedback loop used in the self-resetting logic. During the design phase, an intrinsic reset time on the order of 70ps was deemed the minimum practical duration for the selected 130nm process. This gives a minimum pulse to pulse timing of 140ps internal to the transceiver, and the faster limit in the design margin. The packaging and transmission medium also plays a critical role in setting pulse shape. This pulse system targeted the Quad Flat No-lead (QFN) packaging

system, since it is both inexpensive and high-performance. The draw-back of QFN is bond-wire inductance, limiting the available bandwidth for signaling [23], [24]. Within the confines of this package, rise and fall times faster than 40ps are difficult to achieve within a reasonable power budget and layout. Thus the transmitting pulse width was set to be above 80ps, with an amplitude, .5v nominal relative to ground, set to allow for detection of an attenuated and noisy signal. This extended pulse implies a minimum pulse-pulse timing of 160ps for successful operation. In cases where package-limited speed was not required, longer pulse times could be used, with a voltage vs. detectability trade-off that results in lower operating voltage and thus substantially reduced power.

The relative timing of the two pulse streams is also critical to proper function. Unlike a true differential signal, skew in the two pulse streams is not a detection problem (no cross-over point distortion); the primary concern is to not have an ambiguous ordering of pulses; if both '1' and '0' pulses arrive at the same time it is impossible to re-construct which was supposed to come first. The transceiver is constructed to place a simple arbiter inside the receiver amplifier to help extend the margin of overlapping signals that can be correctly decoded. Such arbiters have shown success at resolving cases of overlap in time to nearly a single picosecond [17].

B. Transmission Medium Considerations

The pulse encoding is not a traditional DC-balanced differential encoding, as information is communicated in both the differential and common mode of two wires, and DC-balance is neither required by the receiver nor provided by the transmitter. The design of the presented detector assumes two independent ground-terminated channels are available. The relative timing of the two transmission channels must meet the timing constraints detailed in Section IV-A. The rationale behind this design is that ground termination allows low-power idle and a simple termination scheme. The original intended medium was to be two close shielded conductors with some coupling such as a shielded twisted pair cable (STP) or two printed circuit strip-lines. All testing was carried out with only 1cm PCB traces and coaxial cables; long runs where each channel conductor does not have its own shield would require separate verification that there would not be substantial cross talk.

V. DESIGN

Both the serializer and deserializer work on the concept of a cascade of cells that activate each-other in sequence. For short bursts, the serializer will operate as fast as the individual cells can fire, slowing down to match the data rate of the transmit data strobe. This slow-down occurs on each strobe edge, giving the characteristic bunching of the transmitted pulses. The deserializer takes advantage of these gaps, fitting the recovered data strobe edges into these spaces. Given the design of the system, the maximum rate at transmit is set not only by the maximum rate that the transmitter can support but also the maximum rate the receiver can support.

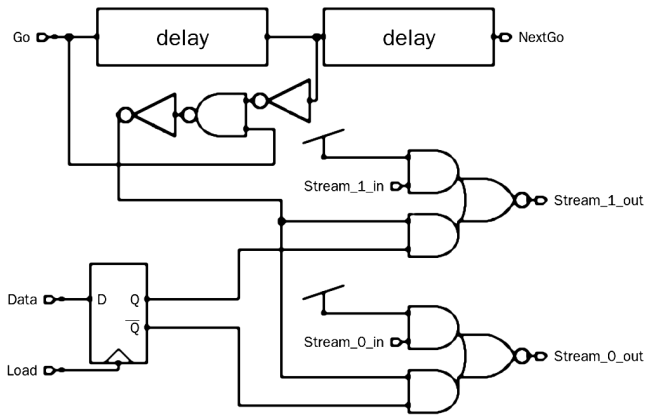


Fig. 7. Serializer Single Cell. The cell allows data from previous stages until the enable signal from the previous cell allows it to transmit. The first cell in the chain is triggered by either the system clock, or a data valid signal in an asynchronous system.

A. Serializer

The serializer is constructed of cells as shown in Fig. 7. The serializer passes data until the rising edge of the enable signal, at which it creates a pulse corresponding to the bit loaded into the cell. The cell creates a delayed copy of its enable. This enable is suitably delayed to create a correct spacing between pulses – that is the time for the pulse width of the serializer cell plus at least a minimum pulse-space worth of time. The relative delay between the delay circuits and the pulse-generating circuits is critical and care must be taken so that the delays of these two systems track with each other. This imposes layout matching constraints on the cell. The serializer action is thus dictated by this chain of cells and events. The serializer can function properly because the forward delays of the enable signal are laid out such that they will be longer than the cell takes to generate a forward going pulse. The pulse shape is derived from a sub-section of this delay line to further reinforce this relationship – that the generated pulses are appropriately sized and spaced. Fig. 8 shows the pattern for the whole system two groups of four cells that have this forward propagating logic. One group of four operates on the rising edge of the data strobe, while the other group works with the falling edge of the strobe. This helps the clock regeneration be more symmetric (since the original pulse pattern is similarly timed from the rising and falling edges of the clock).

B. Deserializer

The deserializer has a cascade of cells, schematic in Fig. 9, similar to the serializer with the exception that instead of a delayed enable triggering a cell, the capturing of a bit triggers the next cell. This again creates a cell with a number of timing-critical details. The detection of capture time plus the enabling time of the cell must be at least as long as the pulse-width of an incoming pulse (to prevent double-capture) but must be no longer than the pulse width plus the pulse spacing (to prevent missing the next pulse). This narrow window in time is easily met, because the serializer system sits after a pulse-shaping stage. The system guarantees that the pulse width is three

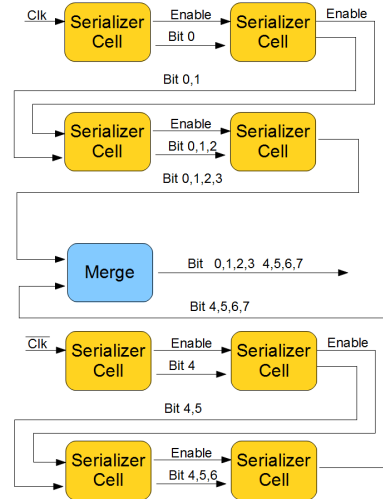


Fig. 8. Serializer Architecture. Showing the two banks of 4-bit chains of serializer cells. The two banks operate off of opposite phases of the clock to evenly spread the pulses between clock phases. This compensates for the self-timed behavior of high-speed operation.

gate delays, the loop between detection and the next detection stage is four gates long (as shown in Fig. 9). The deserializer's function thus requires that the gates in the deserializer track with the pulse-shaping gates in terms of behavior.

The clock recovery behavior of the deserializer comes from separating the deserializer into two equal-length banks. A toggle latch is used to track which bank is currently being filled by the bit-stream and which bank is static. The toggling of this latch generates a signal that can function as a clock signal for the purposes of latching data. This clock signal has all of the inherent jitter of the data system, as there is no PLL to average the jitter across multiple cycles. Though the resulting signal may not be appropriate for all possible clock applications, this is an extremely small state machine for clock recovery.

C. Transmitter and Receive Amplifier

The transmit and receive amplifiers are made out of self-resetting circuits as these circuits are very good at maintaining the characteristics of pulsed systems [17]. The voltage of the pulse at the receiver was selected so that a high-rate, low power receiver could be constructed. Large voltage swings allow a large difference in current between on and off states. The specification of signaling between ground and >450mV was selected for this implementation. Because the system is ground referenced, and a low-current idle was desired, a low-threshold NMOS was selected for the input amplifier stage. With a small-bias current, a fast low power front-end amplifier is created. The ground referenced voltage also allows for a 0 dissipation transmitter without needing an additional termination reference source. The ~500 mV swing height was chosen for two reasons. First 500mV is sufficiently large to operate the receiver as a switch rather than a linear amplifier.




Fig. 9. Deserializer Cell schematic. The deserializer should have its first cell normally enabled, the enable out of the last cell is a data valid signal for the system. Alternately, in a clock recovery system, the receiver is broken into two equal banks. A toggle latch is used to generate the clock. One bank is enabled on each phase of the clock, with the bank done signal causing a toggle. This is the extremely small state machine for clock recovery that makes such a system tractable.

Second the 1.5V power rail could be used to directly drive an output at 750mV through a matching network; allowing for both margin in the signal and a mild form of pre-emphasis. In short, the transmitter and receiver are correctly sized CMOS inverters and have similar characteristics, including limited dissipation when not switching.

The receiver circuit, Fig. 10, is a non-linear amplifier constructed in the self-resetting style. The large pull-down transistor makes the amplifier more sensitive to incoming signals; the lower threshold of which is set by a small, always-on, PMOS transistor. The minimum signal height that the system is sensitive to is thus just over the threshold of the NMOS transistor, which is of low-threshold type in the implementation. The action of the receive amplifier, and all of the self-resetting gates, is to integrate above the detection threshold and then rapidly reset. The receive amplifier thus can be seen as looking for pulses of adequate size over a given threshold. The noise floor is thus set by the threshold of the input, about 250mV in this case, and the detection process is not harmed by excessively high voltage pulses as long as they are suitably short in duration (meet the pulse-spacing requirement).

D. Driver Segmentation

Most of the circuits in the link are built with triple redundancy for hardness against upset. The output driver stage could not easily be made redundant; voting in such a high power circuit would be prohibitively power and area hungry, and still could lead to transients in the output circuitry. To combat transient effects the driver circuit was designed as a set of 12 segments, resistively coupled, and designed such that it would still meet specifications with a single element driving incorrectly. Each segment of the drive circuit is connected resistively to the output, shown in the schematic in Fig. 11. Each segment takes triplicated input, and either drives a shaped pulse or not. In the event that there is an upset in any of the




Fig. 10. Receiver amplifier, including protection diodes. The current through the small, always-on, PMOS transistor sets the lower threshold for the charge integration process. This circuit is specifically sensitive to pulses, but is somewhat agnostic about signaling levels.




Fig. 11. Segmented driver schematic. By splitting the driver into multiple segments and resistively coupling those segments, the worst case of single event upset can be compactly analyzed. A 12-way segmented driver with matching resistors has a worst-case output at 85% of rated voltage in the event of an upset.

12 driver slices 11 of the slices will have the correct behavior (either driving or terminating the line) and the upset slice will have incorrect behavior. In the case that the errant slice is falsely driving a driver with 1/12 of the rated current will drive into both the line and a terminator that is at 12/11 of the line's rated impedance resulting in an erroneous pulse 1/23 (~4.3%) of the full voltage – an acceptable amount of noise. In the case that the slice in error is falsely refusing to drive a pulse then 11/12 of the current is going into both the line and a terminator at 12 times the line's impedance in which case the output pulse has 11/13 (~85%) of the rated voltage – a worse but still acceptable amount of noise for the driver.

VI. LAYOUT

The layouts of these systems are very small ($100\mu m \times 200\mu m$) in addition to the required bond-pads. The number generator occupies much of the extra size in the transmitter layout while the receiver has no extra attached hardware. The layout of the transmit system is shown in Fig. 12. Key elements of the transmit system are highlighted such as (E) The random number generator (F) the logic handling the test-mode (D) the serializer (C) the 12-way segmented transmit amplifier (B) the Electro-Static Discharge (ESD) protection clamps and (A) the

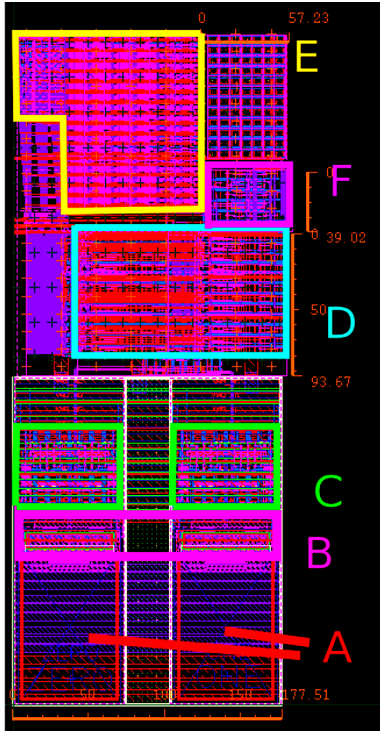


Fig. 12. Transmitter layout. Labeled sections: A marks the bond-pads. B refers to the ESD protection diodes. C is the 12-way segmented drivers, allowing for high-speed and hardness against upset. D is the serializer. E is random number generator used for testing; it implements a 17bit linear feedback shift register in triplicated logic that can generate 6+Gb/s. F is a state machine to approximate an as-fast-as possible clock.

bond pad. This should give reference to the relative sizes of all of the involved components. The entire system is designed to be as compact as feasible, to enable easy integration into other systems. Future work on the layout may see some further integration of the serializer into the pad area, making an even smaller layout. The layout of the receive block is shown in Fig. 13 and similarly consists of a (A) bond pad and (B,C) ESD protection, additionally (D) receiving amplifiers and (E) deserializer. The deserializer has no provision for special test included in the layout, and can also see easy integration due to its small size. Again future developments may see some of the logic block incorporated into the pad region, allowing for a much more dense system.

A. Correlated upset

Voting, rather than intrinsic hardness, is the primary means of protecting against upset in this transceiver, and hence the layout needs to be sensitive to proximity of related circuits. The distribution of charge caused by an ion strike can have an effective track diameter on the scale of $1\mu\text{m}$ [25]. A minimum (center to center) cell spacing of $13.4\mu\text{m}$ between cells that fed the same voter was enforced, and guard rings were placed between all related cells. The exact distance was selected to meet a demand for $10\mu\text{m}$ spacing, need for guard rings, and to meet basic goals of layout self-similarity and room for routing. No cells were flipped so that the center to center space was also the similar device to device space. Transistors were limited to




Fig. 13. Receiver layout. A is the bond pads, B shows the HBM ESD diodes. C is the Termination network and CDM ESD protection. D highlights the triplicated receive amplifiers for each channel. E is the full deserializer, including recovery latches, and data strobe generation.

$5\mu\text{m}$ in any direction. Given these rules, a particle would have to strike at an angle of roughly 6.8° from parallel to the surface in order to create a correlated upset. The sensitive trajectories represents a solid angle of $.09\text{sr} \approx 7\%$ of all possible traversing trajectories, all which force a particle to travel through a large amount of material on the way to the sensitive area.

VII. PROJECTIONS AND RESULTS

A. Projections

Performance estimates were generated using post-layout extraction and the models of expected radiation induced damage. The characterization relied on Monte-carlo simulation to verify that many different stages of damage and process variation were considered. This methodology renders the best available approximation of performance.

The whole link (minus the test generator) is projected to use less than 60mW (40mA @ 1.5V) when operating at 5Gb/s – the average maximum operating frequency was predicted to be 5.1Gb/s. Of the estimated dissipation 35mW is estimated for the transmitter, 4.1mW was estimated for the receiver. The combined power estimate for the serializer, the deserializer and the random number generator is 46mW with 25mW taken by the random number generator and roughly 10mW each for the serializer and deserializer.

B. Test Configuration

A small printed circuit board was fabricated to facilitate test. The circuit board hosts one chip and allows the connection of micro-coaxial cables to test data transmission. This test fixture is shown in Fig. 14. 15cm long pigtails that adapted the micro-coaxial connectors to common SMA type connectors were attached to each of the '+' and '-' labeled ports serving as 1 and 0 bit-streams. 1m RG-174 cables were used to attach the transmit port to either test equipment or another test module for combined TX/RX test. Due to the unique encoding, no




Fig. 14. Test fixture for the first manufactured version of the 5Gbps link. High-speed signals are routed to coaxial U.FL type connectors. All testing is currently conducted with coaxial cables. The central chip under test is 4mm on a side; the whole fixture is 60mm on a side.

transmitter simulator was available to exercise the receiver on its own.

Three key test configurations were considered. The first has a test chip driving a signal through a cable pair into the front-end of a high-speed oscilloscope, an Agilent DSA90804A in this case. This first mode uses the random number generator built into the chip and allows the data to be analyzed in software from a captured scope trace. The second test configuration has one chip driving another, to test the capability of the receiver. In this case a characterized stream is driven into the receive side of a different chip, allowing the receiver to be tested. The third test is to take the transmit side from a chip and loop it back into its receive side for measuring power consumption in a mode that would emulate full-duplex operation.

C. Measurements

To verify correct serializer and transmitter behavior a number of measurements were taken on pulse streams captured by a 8GHz oscilloscope. A measured portion of the pulse stream is shown in Fig. 15. Single data sets consist of 1.28ms of data gathering, roughly 6.8 million bits at a rate of 5.3Gbps. As the trace shows, the noise floor is roughly 100mV and the pulse peaks are measured at over 450mV.

Power: The test chip has an idle current measured at 29mA at 1.5v, full-speed transmit only is measured at 68mA, which goes to 76mA with transmit and receive both active as part of a loop-back test. By subtracting the idle current (power consumed by other devices on the test chip) the full link current is measured at 47mA including the current for the random number generator. This gives an operating power of 71mW, less than the projected 85mW for the random number generator and transceiver combined. If this power savings was spread proportionately amongst the test structures and the transceiver, the actual operating transceiver power is estimated to be 50.1mW while running at 5.3Gbps.

Serializer Speed: The self-timed nature of the communications link means that speed can be changed by altering




Fig. 15. Measured pulse stream captured from a test sample. The green trace represents the '1' data and the red trace represents the '0' data. Trace shows roughly 100mV of noise and at least 450mV of pulse peak height.

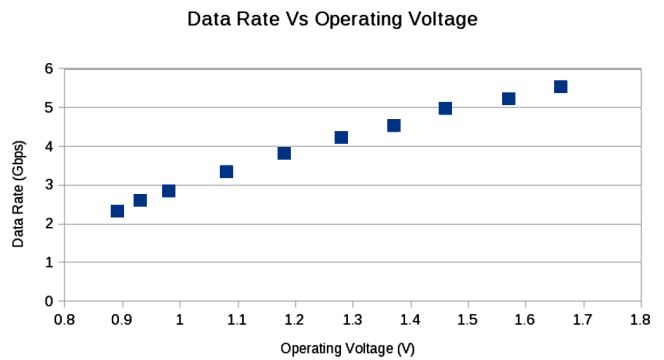


Fig. 16. Measured data rate vs. I/O operating voltage for a single chip. Data gathered while the chip was running without changing chip configuration. The self-timed nature of the logic means that the operating frequency scales with voltage, allowing lower power, lower rate operating modes. The measured sample is capable of operating with a supply of .89V at a rate of 2.32Gbps and at a supply of 1.66V at a rate of 5.53GHz by just changing the supply voltage. 1.5V is the recommended maximum voltage for the process.

the operating voltage of the chip. The chip demonstrated correct functioning between 0.89V and 1.66V, as predicted by simulation. Fig. 16 shows the measured response to voltage scaling, giving a range of 2.3 Gbps to 5.53 Gbps operation. The data for this curve was gathered while the chip was running without changing chip configuration – the same random number generator, serializer, and transmitter function over the voltage range of 0.89V to 1.66V with only changes in timing, which they are designed to tolerate.

Signal Integrity: The traditional eye diagram doesn't capture the readability of a received signal when the receiver is asynchronous and looking for a pulse. Detectivity can be characterized by the pulse amplitude at the receiver, and the pulse duration. For 1.5V core voltage, and 1m of RG174 cable the distribution of pulse voltage is shown in Fig. 17, the pulse widths shown in Fig. 18.




Fig. 17. Measured pulse amplitude after 1m of cable, 1.5V core voltage

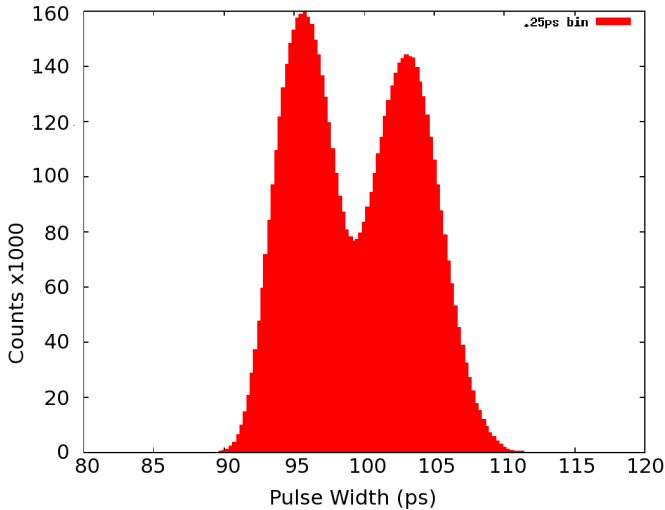


Fig. 18. Pulse width measured after 1m of cable. Note: two different pulse widths were not created by design, but are considered acceptable behavior

Pulse Time Stability: The relative pulse detection time for this transmission scheme demonstrated unexpected stability. As described in Section V-A, two groups of four pulses are used in a single transmit cycle. The additional reset time between words makes the first pulse of a word one that follows the longest of the characteristic breaks. The relative detection time of pulses is plotted in Fig. 19. No pulse in the 6.8 million bit data set is more than 20ps away from the average arrival time of pulses in the same bit position. The transmitter and receiver are neither designed to perform this stability nor are they required to for correct operation, but it is an interesting observation none the less.

Irradiation Testing: While no single event upset testing has been yet completed due to a fabrication bug, two proton irradiation tests have been performed, one to $1.1 \times 10^{14} p/cm^2$ with 36MeV protons, and one to a total dose of $4.95 \times$




Fig. 19. Relative pulse detection time for a whole word. Pulse arrival time is very stable compared to the required no-overlap condition. Timing relative to the last pulse of the previous word.

$10^{14} p/cm^2$ with 800MeV protons; Neither sample showed signs of changes in behavior. The sample exposed to 36MeV protons also had continuous current monitoring and showed no evidence of latch-up.

VIII. CONCLUSIONS

Given the importance of data transmission for high energy physics experiments, any improvement in power and data rate are useful. In the very high radiation environments of cutting edge particle accelerators, large amounts of accumulated radiation damage and charged particles must be tolerated. By changing the line coding and using asynchronous serializers and deserializers, mitigation of radiation effects can come at lower cost to the developer.

The presented pulse link represents a large amount of power savings compared to other efforts at similarly radiation hard serial link technology in 130nm, notably the serializer and deserializer in CERN's GBT consume 330mW and 450mW respectively [26], and run at a similar rate. On an energy per bit level, the pulsed link consumes $\frac{1}{20}$ of the power as this more traditional design.

By making system level changes such as unconventional signaling, asynchronous circuitry, substituting high Q_{crit} circuits with voted circuits, high performance can be had at a lower power cost.

ACKNOWLEDGMENT

The authors would like to thank Prof. J. Incandela for his support of this work.

REFERENCES

- [1] M. Miller, F. Brewer, G. Magazzu, and D. Wang, "Multi-gigabit low-power radiation-tolerant data links and improved data motion in trackers," *Journal of Instrumentation*, vol. 9, no. 12, p. C12011, 2014. [Online]. Available: <http://stacks.iop.org/1748-0221/9/i=12/a=C12011>
- [2] P. E. Dodd and L. W. Massengill, "Basic mechanisms and modeling of single-event upset in digital microelectronics," *Nuclear Science, IEEE Transactions on*, vol. 50, no. 3, pp. 583–602, 2003.

- [3] H. H. Chung, W. Chen, B. Bakkaloglu, H. J. Barnaby, B. Vermeire, and S. Kiaei, "Analysis of single events effects on monolithic PLL frequency synthesizers," *Nuclear Science, IEEE Transactions on*, vol. 53, no. 6, pp. 3539–3543, 2006.
- [4] H. J. Barnaby, M. McLain, and I. S. Esqueda, "Total-ionizing-dose effects on isolation oxides in modern CMOS technologies," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 261, no. 1-2, pp. 1142 – 1145, 2007.
- [5] F. Faccio and G. Cervelli, "Radiation-induced edge effects in deep submicron cmos transistors," *Nuclear Science, IEEE Transactions on*, vol. 52, no. 6, pp. 2413–2420, 2005.
- [6] L. Gonella, F. Faccio, M. Silvestri, S. Gerardin, D. Pantano, V. Re, M. Manghisoni, L. Ratti, and A. Ranieri, "Total ionizing dose effects in 130-nm commercial CMOS technologies for HEP experiments," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 582, no. 3, pp. 750–754, 2007, VERTEX 2006 Proceedings of the 15th International Workshop on Vertex Detectors.
- [7] C. H. et al., "Radiation hardness characterization of a 130nm technology," in *Radiation Effects Data Workshop, 2007 IEEE*, vol. 0, July 2007, pp. 123–130.
- [8] B. Haugerud, S. Venkataraman, A. Sutton, A. Prakash, J. Cressler, G. Niu, P. Marshall, and A. Joseph, "The impact of substrate bias on proton damage in 130nm CMOS technology," in *Radiation Effects Data Workshop, 2005. IEEE*, July 2005, pp. 117–121.
- [9] R. C. Laco, J. V. Osborn, R. Koga, S. Brown, and D. C. Mayer, "Application of hardness-by-design methodology to radiation-tolerant ASIC technologies," *Nuclear Science, IEEE Transactions on*, vol. 47, no. 6, pp. 2334–2341, 2000.
- [10] H. J. Barnaby, "Total-ionizing-dose effects in modern cmos technologies," *Nuclear Science, IEEE Transactions on*, vol. 53, no. 6, pp. 3103–3121, 2006.
- [11] J. Benedetto, H. Boesch, F. McLean, and J. P. Mize, "Hole removal in thin-gate MOSFETs by tunneling," *Nuclear Science, IEEE Transactions on*, vol. 32, no. 6, pp. 3916–3920, Dec 1985.
- [12] X. Huang, A. Francis, A. Lostetter, and H. Mantooth, "Compact modeling of environmentally induced radiation effects on electrical devices," in *Aerospace Conference, 2004. Proceedings. 2004 IEEE*, vol. 4. IEEE, 2004, pp. 2597–2607.
- [13] D. Blackburn, T. Robbins, and K. Galloway, "Vdmos power transistor drain-source resistance radiation dependence," *Nuclear Science, IEEE Transactions on*, vol. 28, no. 6, pp. 4354–4359, Dec 1981.
- [14] P. Pavan, R. Tu, E. Minami, G. Lum, P. Ko, and C. Hu, "A complete radiation reliability software simulator," *Nuclear Science, IEEE Transactions on*, vol. 41, no. 6, pp. 2619–2630, Dec 1994.
- [15] M. Miller, D. Wang, G. Magazzu, and F. Brewer, "5gbps interconnect links for high radiation environments." Presented at the 32nd Hardened Electronics and Radiation Technology (HEART) Technical Interchange Meeting, Chantilly, Virginia, 2015.
- [16] R. E. Lyons and W. Vanderkulk, "The use of triple-modular redundancy to improve computer reliability," *IBM Journal of Research and Development*, vol. 6, no. 2, pp. 200–209, 1962.
- [17] M. Miller, G. Hoover, and F. Brewer, "Pulse-mode link for robust, high speed communications," in *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*. IEEE, 2008, pp. 3073–3077.
- [18] P. Teehan, G. G. F. Lemieux, and M. R. Greenstreet, "Towards reliable 5Gbps wave-pipelined and 3Gbps surfing interconnect in 65nm FPGAs," in *Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays*. ACM, 2009, pp. 43–52.
- [19] B. D. Winters and M. R. Greenstreet, "A negative-overhead, self-timed pipeline," in *Asynchronous Circuits and Systems, 2002. Proceedings. Eighth International Symposium on*. IEEE, 2002, pp. 37–46.
- [20] IEEE, "IEEE Standard for Heterogeneous Interconnect (HIC) (Low-Cost, Low-Latency Scalable Serial Interconnect for Parallel System Construction)," *IEEE Std. 1355-1995*, pp. i–, 1996.
- [21] S. Parkes and P. Armbruster, "Spacewire: a spacecraft onboard network for real-time communications," in *Real Time Conference, 2005. 14th IEEE-NPSS*, June 2005, pp. 6–10.
- [22] I. Sutherland and S. Fairbanks, "Gasp: A minimal FIFO control," in *Asynchronous Circuits and Systems, 2001. ASYNC 2001. Seventh International Symposium on*. IEEE, 2001, pp. 46–53.
- [23] F. Mortan and L. Wright. (2004, February) Quad Flatpack No-Lead Logic Packages. Texas Instruments. [Online]. Available: <http://www.ti.com/lit/an/scba017d/scba017d.pdf>
- [24] Y.-L. Lai and C.-Y. Ho, "Electrical modeling of quad flat no-lead packages for high-frequency ic applications," in *TENCON 2004. 2004 IEEE Region 10 Conference*, vol. 500. IEEE, 2004, pp. 344–347.
- [25] P. E. Dodd, O. Musseau, M. R. Shaneyfelt, F. W. Sexton, C. D'hose, G. L. Hash, M. Martinez, R. A. Loemker, J.-L. Leray, and P. S. Winokur, "Impact of ion energy on single-event upset," *Nuclear Science, IEEE Transactions on*, vol. 45, no. 6, pp. 2483–2491, 1998.
- [26] P. Moreira, "GBT Project Status," Mar 2011, ACES 2011 - Common ATLAS CMS Electronics Workshop for SLHC.