

Final Report: DOE DE-SC0007994

Forrest Brewer and Joe Incandela

This project was slated to design and develop Rad-Hard IP components for 1Gb/s links and supporting hardware designs such as PLL, SER/DES, pad drivers and receivers and custom protocol hardware for the 1Gb/s channel. Also included in the proposal was a study of a hardened memory to be used as a packet buffer for channel and data concentrator components to meet the 1 Gb/s specification. Over the course of the proposal, technology change and innovation of hardware designs lead us away from the 1 Gb/s goal to contemplate much higher performance link IP which, we believed better met the goals of physics experiments. Note that CERN microelectronics had managed to create a 4.7 Gb/s link designed to drive optical fibers and containing infrastructure for connecting much lower bandwidth front-end devices. Our own work to that point had shown the possibility of constructing a link with much lower power, lower physical overhead but of equivalent performance that could be designed to integrate directly onto the front-end ASIC (ADC and data encoding) designs. Substantial overall power savings and experimental simplicity could be achieved by eliminating data transmission to data concentrators and data concentrators and related hardened buffering themselves, with conversion to optical media at a removed distance from the experiment core. We had already developed and tested Rad-Hard SER/DES components (1Gb in 130nm standard cells) and redundant Pad Drivers/Receivers (3+ Gb/s designed and measured performance), and had a viable 1Gb/s link design based on redundant a stuttered clock receiver and classical PLL, so the basic goals of the proposal had been achieved. Below, in chronological order, are the products and tools we constructed, as well as our tests and publications.

1. UCFF1 Design

This first version consisted of a soft-IP (before placement) SERDES design and heavily over-designed link components: clock generators, phase detectors and rad-hard current mirrors and related analog components. Additionally, a voltage-controlled oscillator with partially stabilized operation: the pull-up side was left without stabilization, but the pull-down side was fully stabilized. This allowed direct characterization of the total dose effects on the more sensitive P-type devices as well as allowing testing of the new current source designs. (The current source needed total dose stability against shifting device parameters and also needed to be unconditionally stable with very rapid return to operation after a single event transient (SET). Note, conventional tricks such as the use of band-gap devices would not work because of cumulative lattice damage).

The radiation testing provided data which, coupled with results from other published data, enabled a simple process variation Monte-Carlo transistor model. Designed as an extension to the built-in variance model provided by IBM, the new model provides a leakage and threshold addition satisfying TID exposure to ionizing radiation. It has subsequently become a mainstay for design acceptance and optimization, and is used to characterize new design behavior. The model and measurements are shown

in Figure 1 and Figure 2 below. These results and those of others lead us to believe that our radiation hardness problems could indeed be broken into a TID induced variance and leakage model which could be dealt with using variance tolerant design, and a SEE/MEE model which needed to be solved by correct architecture.

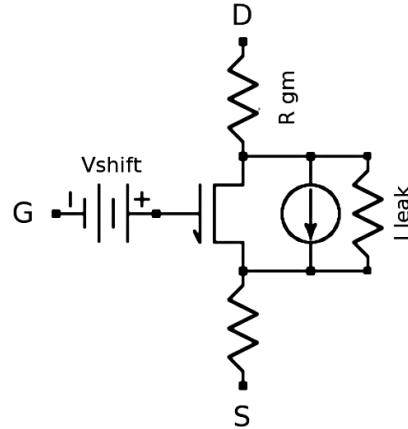


Figure 1: TID Modified transistor model

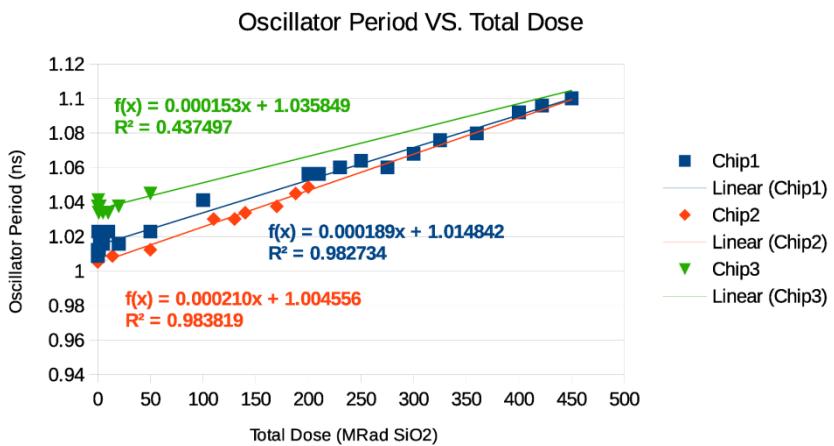


Figure 2: Measured TID response from X-Ray Dose

The SERDES designs were successful, but were at the limit of the standard cell set at 1Gb/s in this technology, with flip-flop latency being the main limitation. On the other hand, the analog design of the pads required an unconventional direct drive of the output with *thin-oxide* transistors. We used a modular-segmented drive with 16 segments so that a SET on a segment would only contribute to the output noise-- not its operation. This design was completed with measured operation performance of better than 3Gb/s, as there was no power savings at lower operation bandwidths. Both the SER/DES and Pad driver designs proved so stable, they were used in several other chips – and scaled into designs

for the 65nm node chips built at INFN. The resulting chip was continually operated and tested to 450MRad in X-ray flux with no errors after the initial startup -- for the 45 continuous hours of test.

2. UCFF2 Design

In constructing the UCFF1 design, we realized that the conventional approach offered only a few avenues to make an efficient hardened design. In particular, hardening an accurate clock oscillator seemed to require very large size -- either inductor or parallel delay element designs. Instead, for the second chip we chose a design based on recent developments in which recent edges are used with a delay locked loop (DLL) to connect the sample to a much shorter time period. This has the effect of increasing the design noise tolerance and reducing the loss period after an SET.

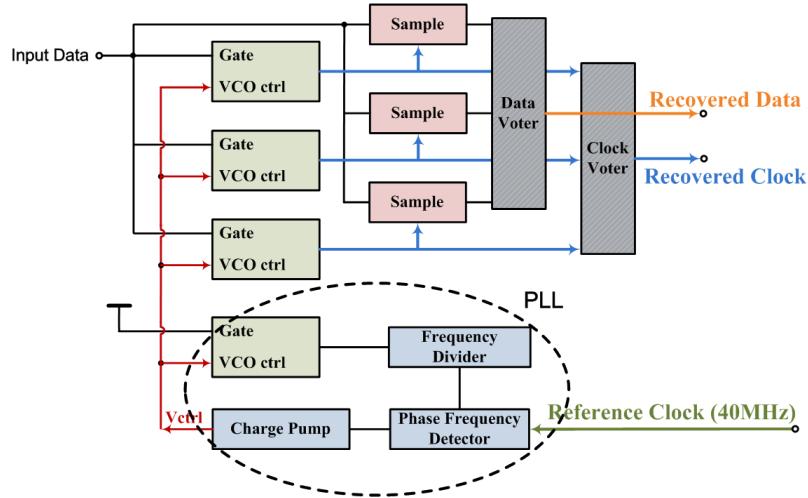


Figure 3: Redundant Clock Recovery CDR

The UCFF2 DLL-based clock and data recovery circuit is realized by gated VCOs as shown in Figure 3. When no data transition is detected at the gate, the VCOs are in oscillation mode and the oscillation loop is closed, creating a clock from a known and stable control voltage established by the separate PLL. In contrast, when a data transition is detected at the gate, the VCOs enter injection mode and the oscillation loop is opened, forcing the VCOs to enter a fixed state in which any frequency/phase errors caused by SETs in the previous data cycle are positively corrected. Shown in Figure 4, this “refresh at data transition” scheme borrows the idea of burst-mode clock recovery put forward by [1][2][3].

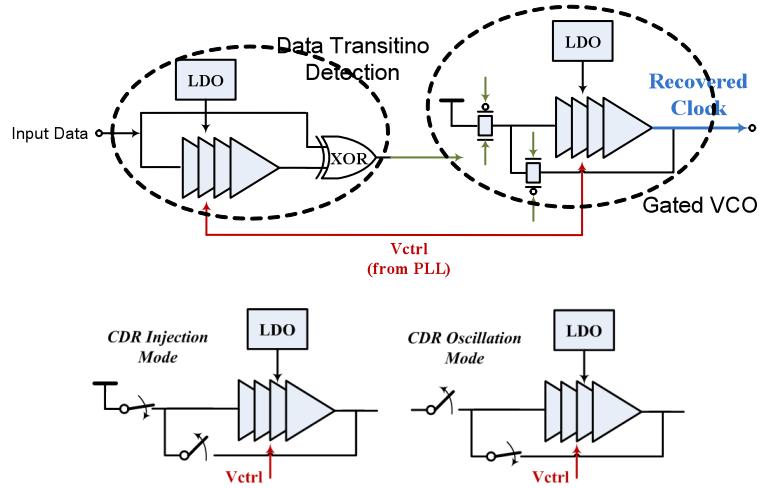


Figure 4: Two operation modes of gated VCO

In operation, the hybrid design exhibits better radiation hardness because the rapid detection and phasing units are triplicated, and the tuning voltage is less crucial to the operation, distributing the error signatures. Board test results of the tuning and locking performance of the PLL are shown in Figure 5. As proved by the test, the PLL is able to lock to any reference frequency between 29 and 51MHz, and provide an accurate 1GHz clock when the reference is 40MHz.

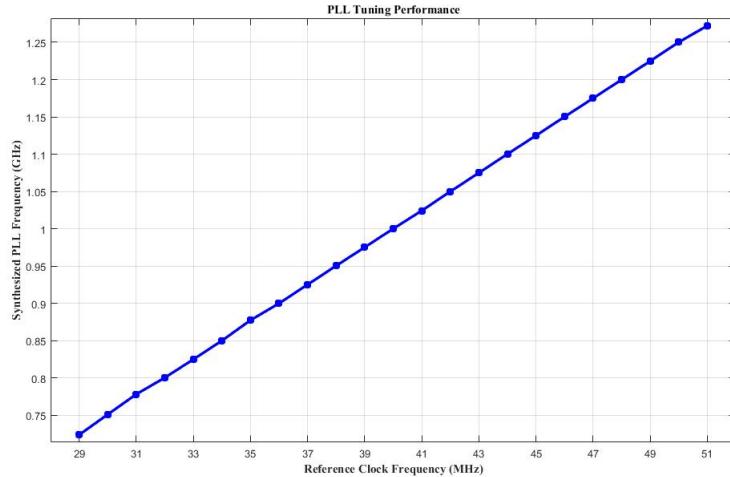


Figure 5: UCFF2 Tuning Performance

The board test result of the SER (clocked by the foregoing PLL clock) is shown in Figure 6. The design in UCFF2 proved to have a +/- 6% global timing tolerance as measured by PLL mismatch to the known source frequency. This is about $\frac{1}{2}$ of jitter tolerance we had hoped for. Further testing revealed 2 issues, 1-substrate mismatch between the master PLL and the slave DLLs and power supply induced jitter increasing the phase noise at the phase detector input. These considerations focused the design on the

issues of non-locality of the detection decisions and lead to the completely new approaches in subsequent chips as we needed rad-hardness and jitter tolerance far in excess to these results for faultless use in accelerator or other physics experiments.

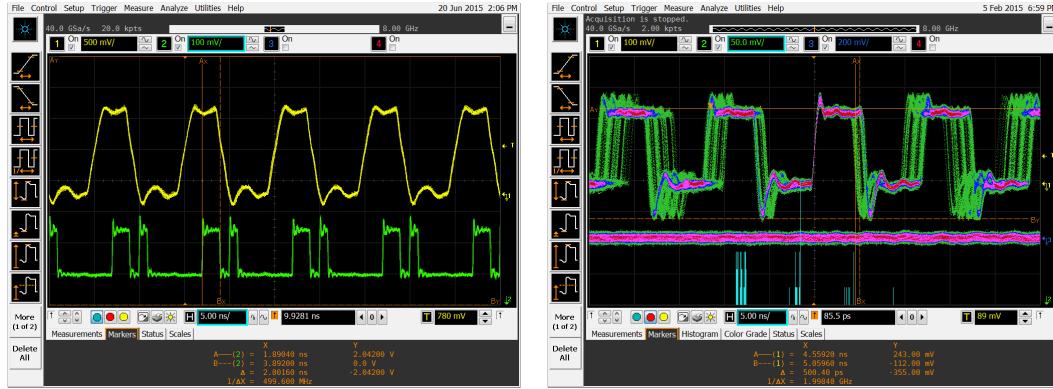


Figure 6: SER input 001101000 (1Gb/s) (right) CDR jitter (left)

The board test results of the jitter the CDR clock is shown in Figure 6. The internal PLL showed good phase continuity and jitter of 22ps, the measured CDR is 135ps, since the VCO oscillation loop keeps being switched on and off owing to its gated structure. This is the normal behavior for this type of detector and indicates correct operation.

- [1] S. Kaeriyama and M. Mizuno, “A 10Gb/s/ch 50mW $120 \times 130 \mu\text{m}^2$ Clock and Data Recovery Circuit,” *Solid-State Circuits Conference, ISSCC 2003 IEEE International*, vol.1, pp.70-478, Feb 2003.
- [2] M. Nogawa, K. Nishimura, S. Kimura, T. Yoshida, T. Kawamura, M. Togashi, K. Kumozaki, and Y. Ohtomo, “A 10 Gb/s Burst-mode CDR IC in $0.13 \mu\text{m}$ CMOS,” *Solid-State Circuits Conference, ISSCC 2005 IEEE International*, vol.1, pp.228-595, Feb 2005.
- [3] Dan Lei Yan, M. Kumarasamy Raja, and Aruna B. Ajikuttira, “A Gated-Oscillator based Burst-Mode Clock and Data Recovery (CDR) Circuit”, *RFIT2007-IEEE International Workshop on Radio-Frequency Integration Technology*, Dec 2007.

3. UCFF3 Design

In parallel with the design of the second chip, we realized that we could drop the analog storage entirely and use asynchronous logic off a new protocol to eliminate all but 2 time scales from the circuits. (The two time scales are the effective pulse width and effective pulse minimum separation in a serial pulse train.) Effectively, this greatly lowers the SET sensitivity of the design by limiting the error length to a few bits at most, even at very-high performance levels. The UCFF3 design makes use of an RZ-pulse encoded 2-wire signaling strategy shown in Figure 7. The strategy is effectively 3-state: transmitting 1, transmitting 0 or non-transmit. Fundamentally, each bit is transmitted along with inherent timing information allowing for fully asynchronous reception. The non-transmit state allows

for sub-mW link off-state engendering very substantial power savings for sporadically used links (which are expected in this application).

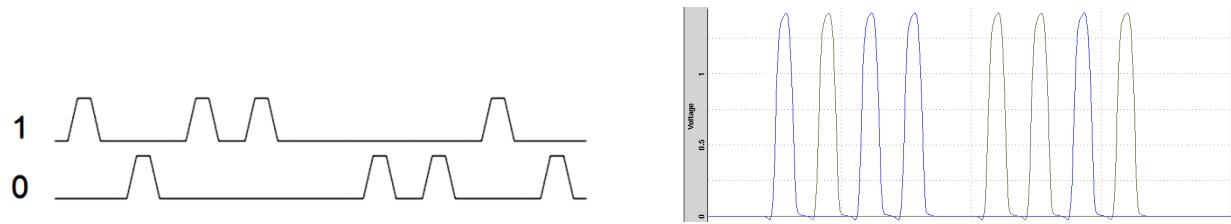


Figure 7: Pulse asynchronous data format

The design of the asynchronous link required several developments: we needed to create a rad-hard high performance standard cell set with extremely low timing variance, we needed a new pad transmitter and receiver design capable of the high speed transits from the pulse signaling model. This model is very robust, but requires effectively 10GB/s signaling to meet 5Gb/s data since both clock and data are sent on each symbol. Abstractly, this seems horrible given conventional communications metrics, however, the new design offers several compensations and we can effectively remove this channel constraint with slightly more complex design. The benefits are enormous jitter tolerance (up to 90% of the symbol period – independently on each bit), greatly reduced complexity to implement radiation tolerance, and vastly reduced power overhead since there are no precision timing components (PLL, DLL or others).

Structurally, the design appears to the logic designer as an asynchronous FIFO, with the far end of the FIFO at the far end of the link (shown in Figure 8). In this way, the link is fire-and-forget and does not need initialization or synchronization at the full data rate of 5.2Gb/s. Both ends of the link are constructed using the new pulse gate standard cells which are feedback stabilized to limit their high speed jitter induced from power and other variance sources. (These cells offer the potential to design very high performance asynchronous designs assembled as standard cell models. We built appropriate CAD models to ensure closure of the designs for practical scale systems.)

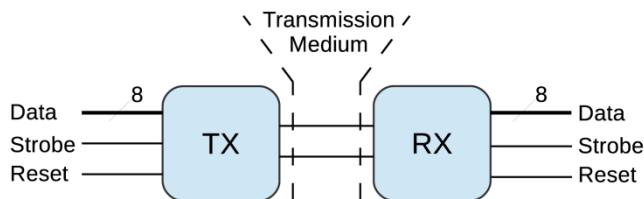


Figure 8: Pulse-link logical model

Remarkably, the new design proved far more robust in simulations as well as being much faster, with much less complexity and layout obligation. It required a redesign of the pads and additionally some high performance low-swing test pads to connect to FPGA compatible ports. The resulting design created links with a 50mW total power budget (both sides, and SERDES) including several stages of modular redundancy. It is not compatible with conventional FPGA I/O but is simple and small enough

that it offers practical savings on analog front-end chips in both pin count and power and can be converted on-the-fly for very cheap IP integration costs.

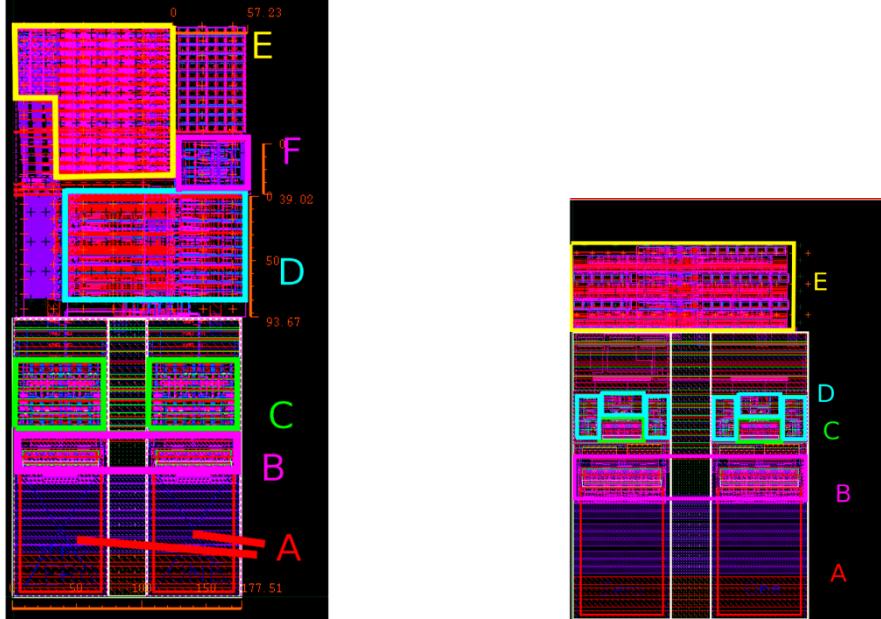


Figure 9: Layout of pulse transmitter and receiver

The layout of the transmitter, test driver and receiver are shown in Figure 9. To get the scale of these figures, the parts labeled “A” are the actual pad locations for wire bonds on the die. The pads in question are 73um wide and a few hundred um tall (smallest of the IBM pad formats). “B” in the figures are the ESD diode arrays for pad protection. “C” in the left figure are the 12x redundant pad drivers while “D” is the complete pulse transmitter. (“E” and “F” in the left figure are high rate PRBS generators for testing). On the right figure, “D” are the triplicated receiver/detector modules while “E” is the complete receiver. (Apparently empty space between “D” and “E” are bus requirements set by the pad frame design for compatibility with the pad ring).

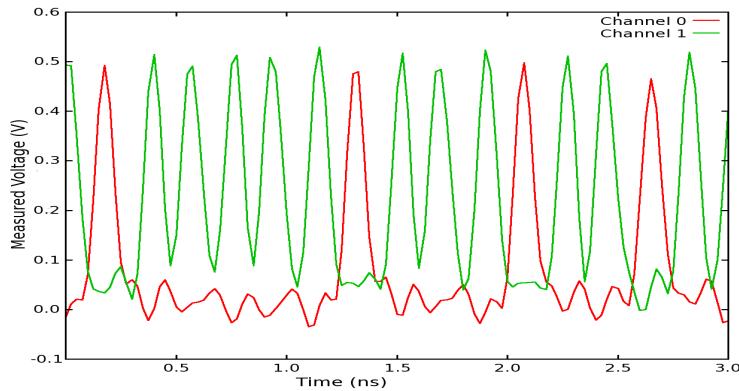


Figure 10: Transmitter output: 1011111011101101 @ 5+Gb/s

Board level testing of this design indicated that the transmitter and receiver worked slightly better than expected. The transmitter output at full speed is shown in Figure 10. The uneven amplitude in the

figure is from sample aliasing in the 8GHz scope used to make the measurement. Since the design is inherently asynchronous, we wished to get a handle on the timing jitter and power/performance trade-off it offers. These are shown in Figure 11 below:

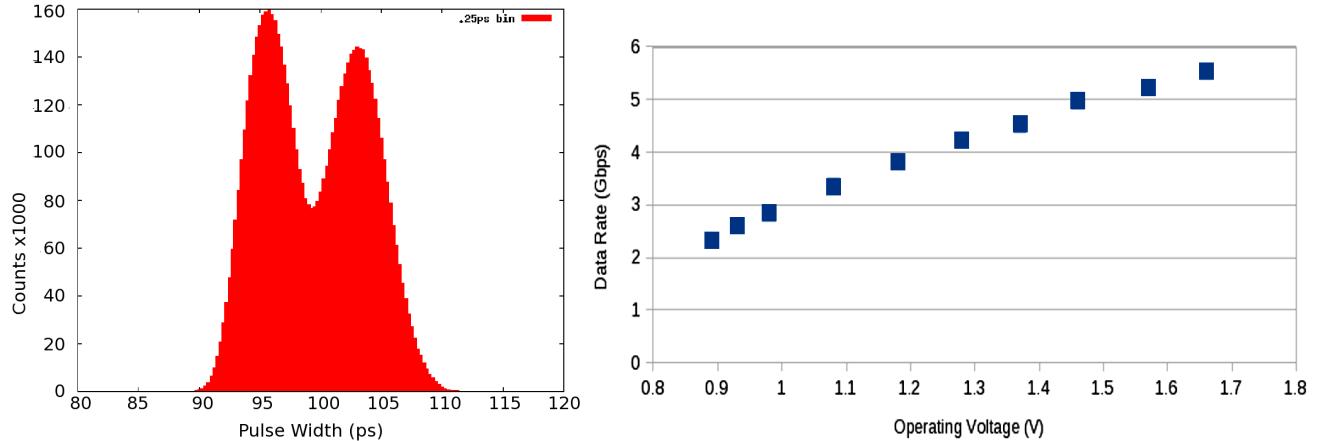


Figure 11: Pulse jitter and data-rate vs. voltage

There are two very interesting observations to be made from these figures: 1. The jitter is bi-modal but occupies a total width of less than 10pS for the measured transmitter output. 2. The link transmission rate is roughly linear with supply voltage from less than 1Gb/s to more than 5Gb/s. These results seem to contradict each other but do not in fact. The pulse gates used in the design have internal feedback to minimize the timing variance in operation, but this compensation operates only at very high rates (e.g. above 2GHz). The construction architecture did not seriously compromise the jitter in delivery to the output, allowing for very good jitter performance in this technology, without use of DLL, PLL or carefully controlled supply voltages. At lower frequencies, the conventional supply vs. performance applies to the link.

Overall, UCFF3 created IP for: 5+Gb/s Pulse Transmitter Pads, 5+Gb/s Pulse Receiver Pads (both redundant and TMR hardened), Asynchronous 5Gb/s SER/DES, 8+Gb/s asynchronous PRBS generator, and 3+Gb/s 1.5V HSTL hardened single-ended (constant current drain) pads. Finally, a possibly most importantly, the new strategy required the development of a pulse-based gate and latch technology that allows relatively simple and safe creation of very high performance asynchronous automata. Such hardened gates were the key to implementation of sub 200pS switching circuits in 130nm CMOS.

4. UCFF4 Design

Nothing is static in the chip business and sale of the IBM fabrication line to Global technologies caused a shift in design technology at CERN. At the same time, uncertainty in the longevity of the former IBM process lead to difficulties in scheduling a 130nm run. A run was finally scheduled and we chose to create a further IP demonstrator based on the pulse design of UCFF3 updated by known issues and setup to appear as a familiar 16-bit asynchronous FIFO interface: 16 bit parallel I/O with a quadrature

timed sampling strobe, both operated in DDR (320 Ms/s = 160MHz). The chip incorporated both a transmitter and receiver section, on the same 16-bit interface (plus reset and strobe) and used separate transmitter and receiver pairs. An additional transmit-receiver pulse pad set was integrated with a few internal pulse gate buffers to test pass-through reliability of the protocol and technology.

As part in parcel, this design was created to allow for radiation testing using both accelerator and x-ray sources with robust I/O, carried to remote test resources. Full testing of this design and radiation testing is still underway, paid for in part by additional funds provided through Fermi-Lab and Prof. Incandella. Test boards were designed with remote interfaces shown in Figure 12.

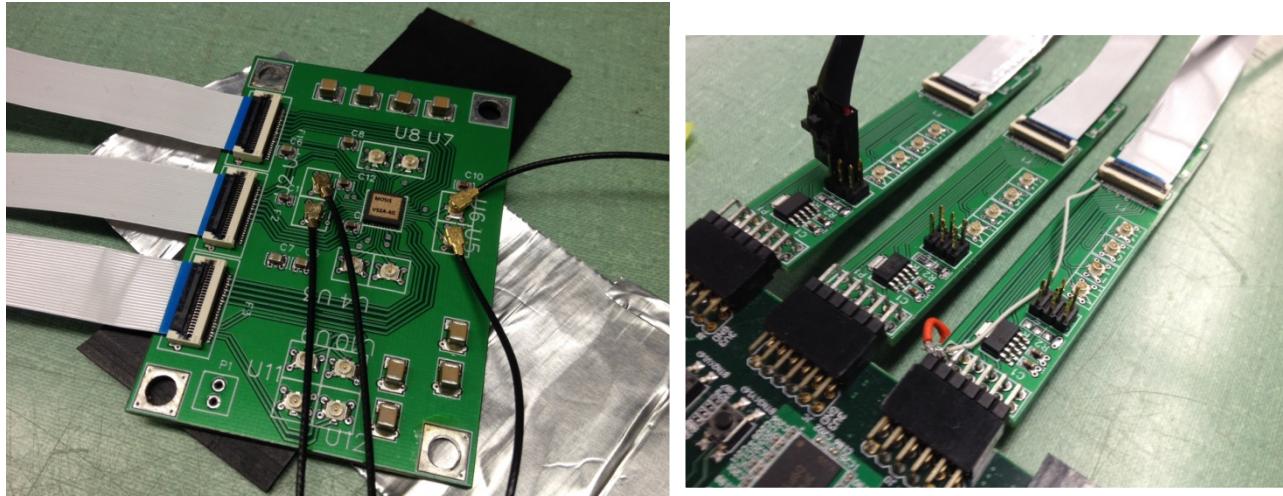


Figure 12: UCFF4 Test boards

The flat cables used in the test setup are commercially available FFC cables which provide low cost digital links (a few \$/cable) and are available in lengths to several feet. Preliminary test of the transmitter output on a long (1.1m RG-178 cable and FEC pigtails) to model cable attenuation is shown in below.

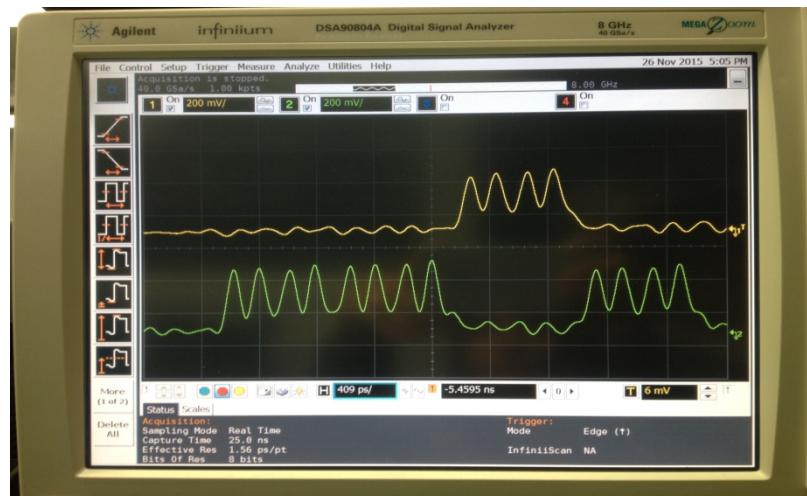


Figure 13: 16-bit pulse transmission @ 1.1m RG-178

Results:

This proposal supported 2 Ph.D. degrees, Merritt Miller (currently a post-doc working on the CMS HGC project) and Di Wang (who expects to graduate by spring 2016. The work resulted in a substantial set of layout and design IP for the 130nm node, with a rad-hard link far exceeding the original requirements of the proposal, in radiation tolerance, performance and power use. These results led to the publication of 4 conference and 2 journal papers so far, with another journal paper in review and 1 conference and 1 journal paper in preparation for submission.

Published papers:

Workshop on Intelligent Trackers (WIT 2014) and journal:

“Multi-Gigabit Low-Power Radiation-Tolerant Data Links and Improved Data Motion in Trackers,” Merritt Miller, Forrest Brewer, Guido Magazzu, and Di Wang, IOP Journal of Instrumentation Vol. 9, No. 12, pp. C12011, 2014.

Real Time Conference (RTC 2014) and associated TNS journal:

“5GB/s Radiation Hard Low Power Point to Point Serial Link,” Merritt Miller, Forrest Brewer, Guido Magazzu, (accepted and in-press IEEE Trans. on Nuclear Science, Nov. 2015).

Hardened Electronics and Radiation Technology (HEART 2015)

“5Gbps Interconnect Links For High Radiation Environments,” M. Miller, D. Wang, G. Magazzu, F. Brewer, Arlington VA, April 2015. (Journal submitted to JRERE and in review).

Nuclear Science Symposium & Medical Imaging Conference (NSS 2015)

“Radiation-Tolerant IP-Cores for 2Gbps Serial Links for the Data Readout in Future LHC Experiments,” G. Magazzu, F. Brewer, M. Miller, D. Wang, IEEE NSS 2015, Sant Diego, CA, Nov. 2015.

Conclusions:

Over the course of this research, the intended aim changed from very conservative rad-hard implementations to quite unconventional ones that have several distinct advantages over the conventional approaches. Although the new signaling strategies have inherent risks – for example, conventional bit-error-rate analysis and tests are useless, they have great simplicity and promise for the application. The asynchronous nature of the link enables a variety of instantiations and performance levels with gigabit level performance achieved at <1.1 volt operation in 130nm. Proponents of conventional links will quickly point out that the described link has substantial common-mode energy,

it needs baseline reset at the end of a long cable and doesn't provide a stable output clock to base local synchronous circuitry on. That being said, the link needs no compensation for up to several feet of cable or PCB at 5+Gb/s, operates end to end at 50mW at peak performance, and at 10% duty cycle has instant wake-up and 5mW total end-to-end power use, allowing for inexpensive redundant links in the experiment.

In the future, we'd like to explore alternative more robust signaling strategies enabled by the low variance design method. In particular, fully differential pulse and high rate asynchronous quadrature show good promise for 5Gb long copper links (3-5m) and 10Gb shielded twisted pair links, both in the same 130nm process likely a mainstay for high performance front-end design.