Non-GBT links for data transmission

8 March 2016

Merritt Miller (UCSB)

Pulse Link Research

- Research Contributions by: Aditya Dalakoti, Carrie Segal, Di Wang, Guido Magazzu, Forrest Brewer
- Support: Joe Incandela, DOE-CERN CMS SLHC “Protocols and IP-Cores for Control and Readout in Future Higher Energy Physics Experiments.”
- Key Papers:
  - M. P. Miller, G. Magazzu and F. D. Brewer, "5 Gbps Radiation-Hardened Low-Power Pulse Serial Link," in IEEE Transactions on Nuclear Science, vol. 63, no. 1, pp. 203-212, Feb. 2016.
  - Miller, M., F. Brewer, G. Magazzu, and D. Wang. "Multi-gigabit low-power radiation-tolerant data links and improved data motion in trackers." Journal of Instrumentation 9, no. 12 (2014)

This is a presentation about

- Front-end to back-end data transmission for high radiation detector areas
  - High bandwidth, over copper
  - As little power as possible
  - Radiation precludes using non-hard or space-rated components
- Alternate approaches to a dedicated transceiver IC
  - Design details of a GALS approach to the problem

The Data-Power Problem

- Example : a full rate data read-out of a large-channel-count detector

  ~5M Channels

  ~5bits/Channel

  - including compression, abstraction, header information

  => ~10^15 bits/s of data at 40MHz

  - Cisco expects 5x10^14 bits/s for the whole internet in 2019*
  - Requires roughly 100,000 Links @ 10Gbps **~75kW** if using LpGBT**
- Data reduction techniques can help BUT
  - If the detector participates in triggering, some of the full rate data must flow
  - Higher detector occupancy may
    - reduce compression efficacy
    - increase the amount of data that needs to be interpreted with more complex algorithms

*Data rates taken from white paper at:
http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html
** Paulo Moreira “LpGBTX Specifications Summary”

What Does Non-GBT Mean?

- In this context we are talking about the GBT chip and not the project/associated blocks
- The (lp)GBT use case looks something like:

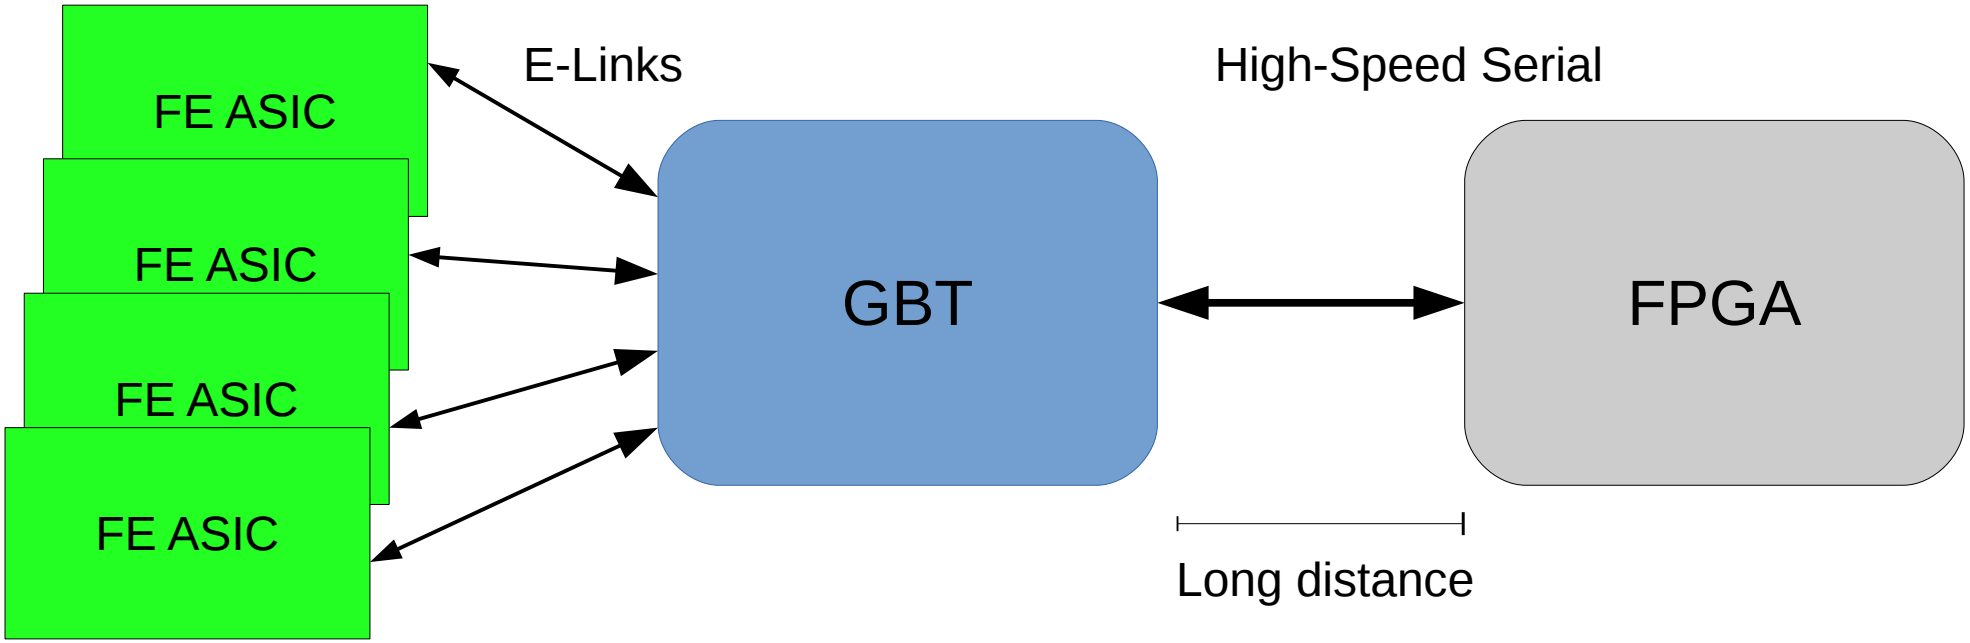

Use-Case Draw-Backs

- E-Links cause power dissipation at both ends
  - ~6mW/ E-link channel for each transmitter*
    - ~40mW to feed the GBT using 1.28G E-Links
    - ~160mW worst-case (28E-Links)
  - E-Links may require serializers and multiple clock domains in the FE if the native speed is slower than E-Link speed
- Adaptable, Multi-clock model on GBT means power consumption to drive clocks, phase adjusters etc.
  - ~140mW of timing power used
  - Long SEU recovery time
- PCB wiring space and complexity to interconnect FE & GBT
- Error Correction applied after 1 round of data transmission and reception

*Power estimates infered from Paulo Moreira “GBT Project: Present & Future” ACES 2014

Alternate Use Model (Non-GBT)

- Embedded Serializer, Delay-line based timing
  - No medium-speed links
  - Less latency (fewer stages of serialization, latching)
  - Possibly faster SEU recovery (Delay lines clear SEUs in 1 word)

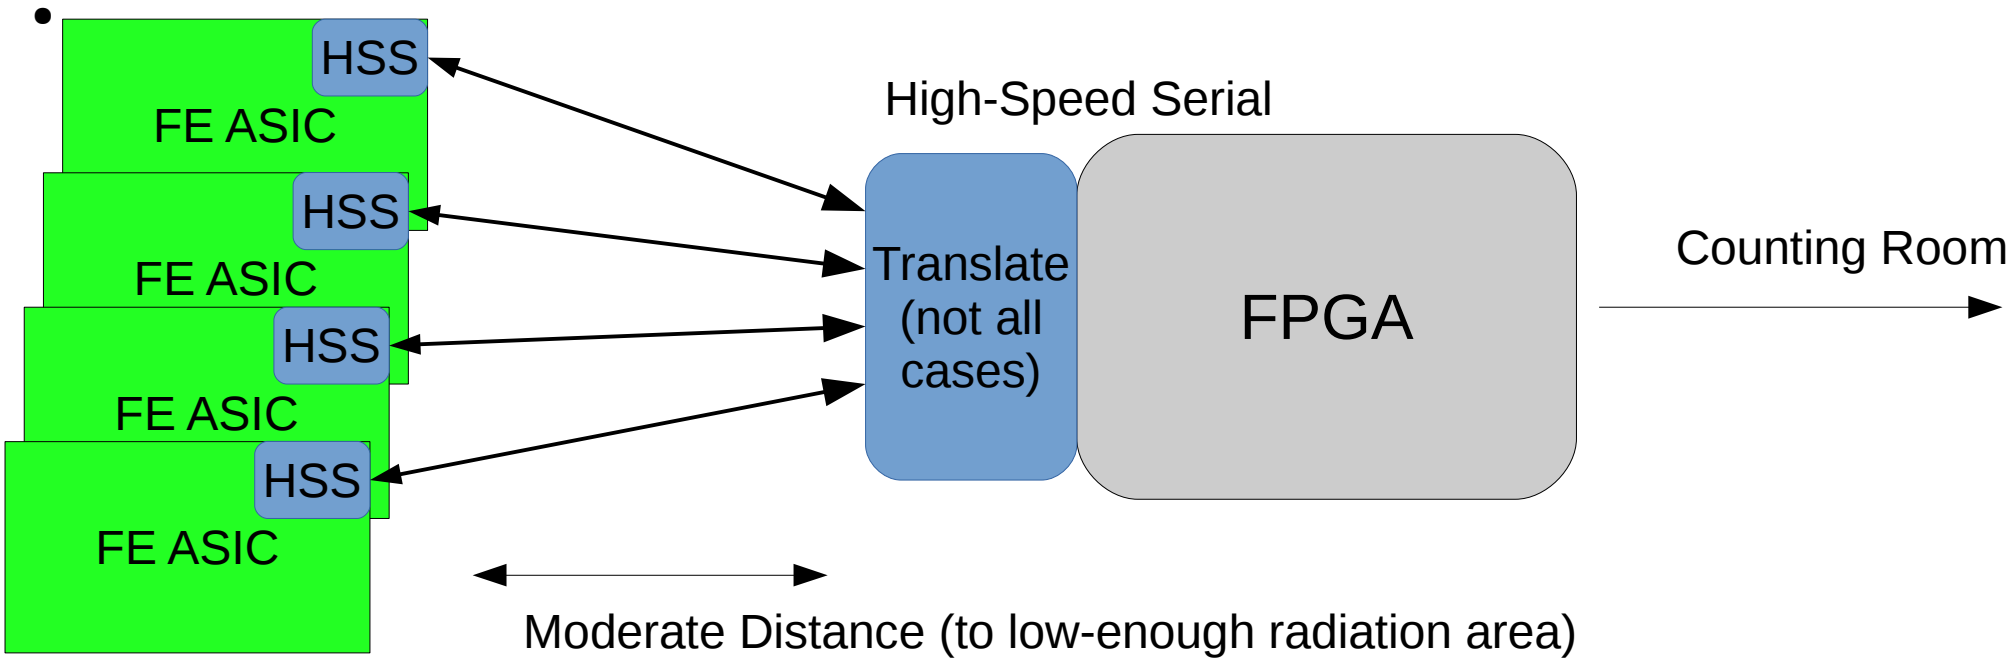

Links with this use-case

- LHCb VeloPix transmitter
  - Sychronous, DLL-Based transmission
  - ~60mW SER + TX*
- UCSB Pulse link (more details presented here)
  - Asychronous logic, Delay-line transmission
  - ~45mW SER + TX, ~15mW for RX
    - Custom RX required
- Note both links give ~ 5Gbps @ ~60mW in 130nm
  - ~5x lower energy per bit transmitted than LPGBT

*Gromov, V., V. Zivkovic, M. van Beuzekom, X. Llopart, T. Poikela, J. Buytaert, M. De Gaspari, M. Campbell, and K. Wyllie. "Development of a low power 5.12 Gbps data serializer and wireline transmitter circuit for the VeloPix chip." Journal of Instrumentation 10, no. 01 (2015): C01054.

Pulse Link

What is Pulse Link and how does it fit into this picture?

- Implements a rapid-wake burst-mode high-speed serial link
- Is a set of IP cores to be dropped into a host ASIC
- Performs serialization, deserialization, and TX/RX functions
- Includes rudimentary pre-emphasis and decision feed-back filters
- Operates with low max power (~60mW @ 5Gbps)
- Idles at very low power ~1mW
  - Wake/sleep happens automatically at word boundary; no extra delay
  - Very good for variable data; Unused bandwidth is unused power
- Is designed for high-radiation environments

Has this link been used before?

- This exact combination is not in service in any application to the best of the author's knowledge
- Pulse Asynchronous circuits have been used
  - Self-resetting logic exists in Pentium 4 for example
- RZ line codes have been used before
  - Common decades ago for electrical links
  - Still has applications in some optical links
- Asynchronous line codes have been used before
  - For example IEEE 1355 DS-DE links (and SpaceWire that take advantage of those style links)

Radiation Hardness Features

- Designed with high-radiation environment in mind
  - No bias voltages used
  - No narrow transistors used
  - Asynchronous circuits broadly tolerant of small changes to performance
- redundancy used through the entire design
  - From receiving amplifiers at each pin to output drivers
  - Output drivers used parallel redundancy (for practicality)
  - Everything else used TMR
- If radiation damage exceeds tolerances circuit slows instead of fails

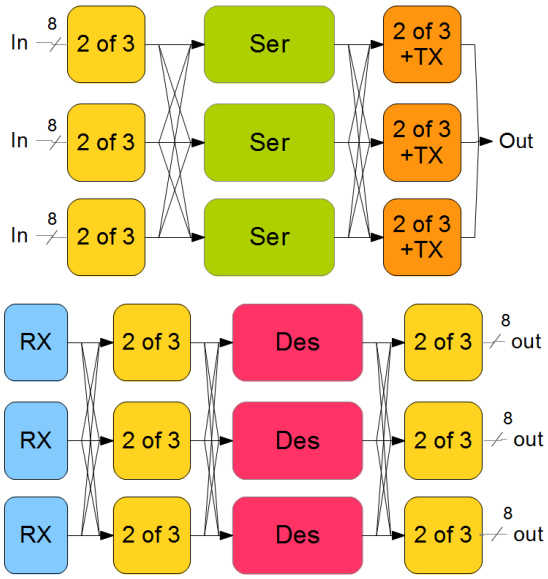

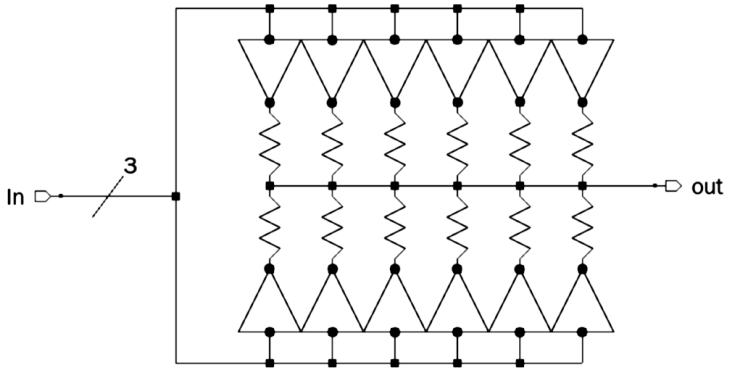

Pulse Logic

- Pulse Link is enabled by a type of Asynchronous Pulse Logic
- This logic family (related to self-resting domino logic) is known for fast operation
- Fast Pulse latches
  - Sampling aperture as low as 80ps
  - Clock->q (equiv) as low as 40ps
- Allows the creation of fast SER & DES without the need for synchronous multiplexing/sampling
  - Both operations become counting & latching (as they would at lower speed with standard cells)

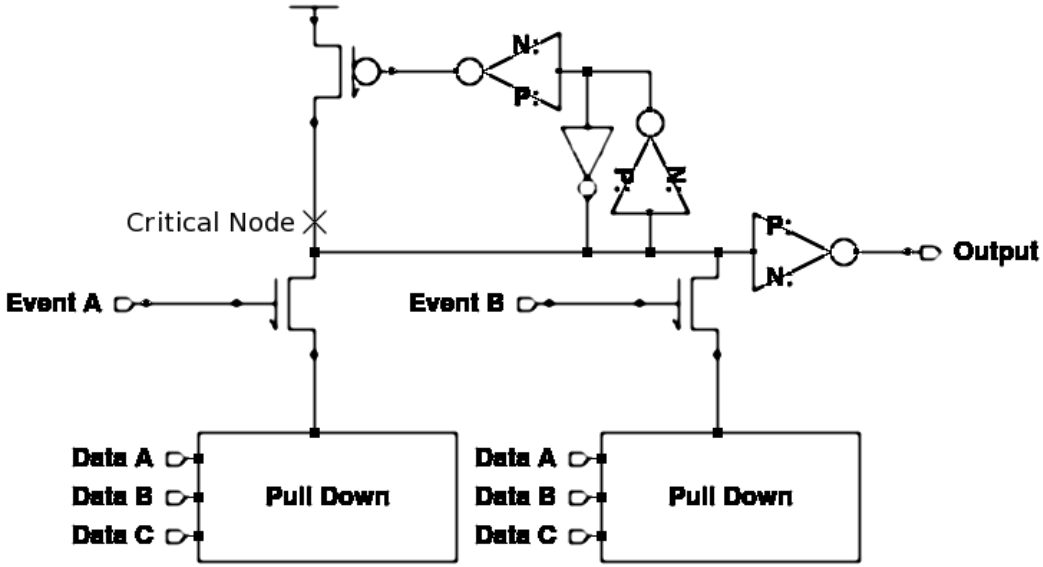

Link Internal Function

- Link logic uses pulses to mark time (events)
- Serialization is translation of a register (electrical levels) to pulses that mark when a 1 or 0 should be transmitted
- A serializer cell is a pair of AND gates and a done signal
- Multiplexing can be done with or gates
- The Serializer is grouped into sets of 4bits
  - 4 bits serailize before voting on the 'go' signal
  - A stage's register (latches) goes opaque when go reaches the 1st bit,
  - latches becomes transparent when go exits the block
- 2 groups of 4 make a group of 8 that runs continuously (one is loading while the other is serializing)
- We have also extended this to a 16 bit DDR version (allows for ~150MHz clock)

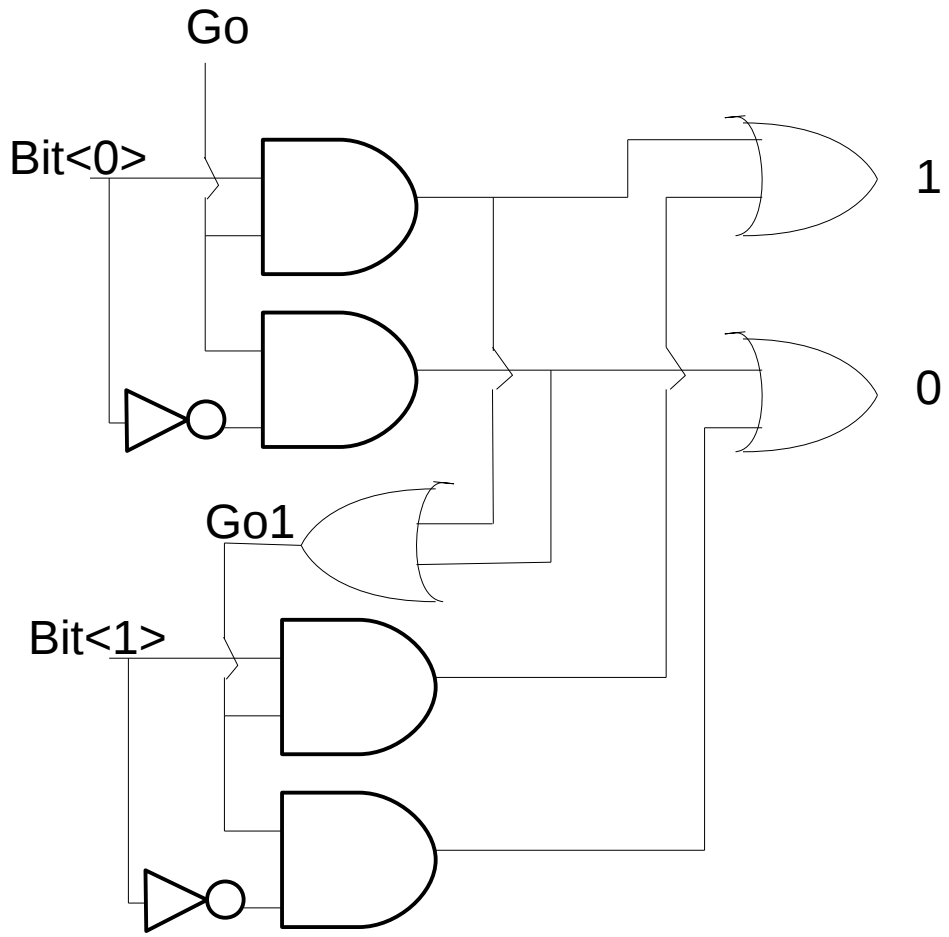

Pulsed Signaling

- Asynchronous Logic Requires an explicit timing signal (cannot imply a clock)
- Pulse Link Must provide a transition for each bit for the signal to be read
- Original link construction uses electrical pulses, which are the native format for the logic family
  - Originally 2 wires carrying pulses
  - Planned revision to true differential (more compatible with twinax cables)
  - Differential Signaling Should improve noise immunity considerably

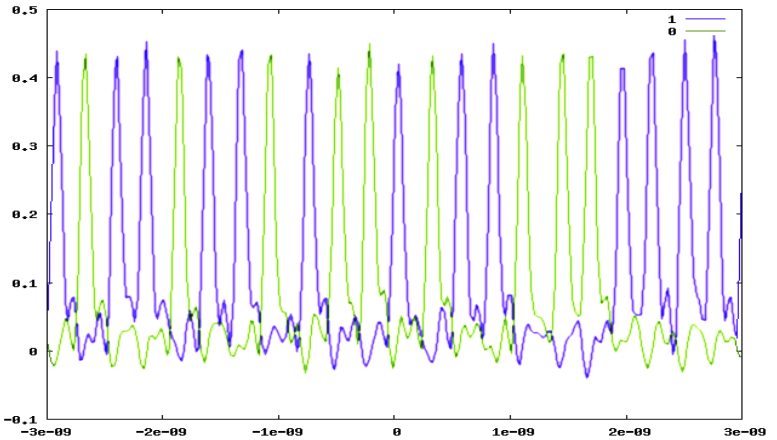

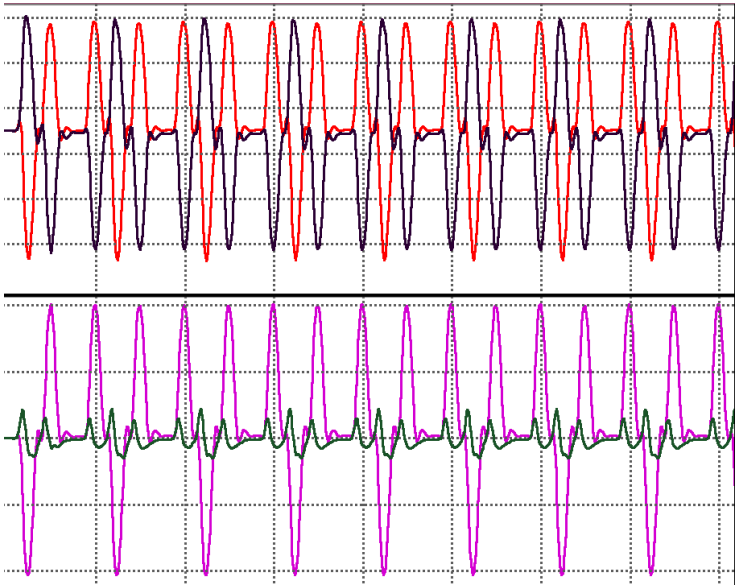

Pulse Detection

- Pulses are detected via a non-linear circuit
  - integrates the area over a given noise threshold
  - Resets integral when pulse is detected
  - Requires reset to be reasonably tuned to the expected pulse shape
- Has some immunity to dispersion and ISI
- Represents a low-order decision feed-back correction circuit
- Detection for UCFF4 prototype RX is roughly 7.5×10^{-12} Vs over threshold
  - Simulation of a 300um PCB trace show TX pulse strength of 9×10^{-12} Vs after 1m
  - Detection improvements for twinax cable planned for next prototype (~6mos.)

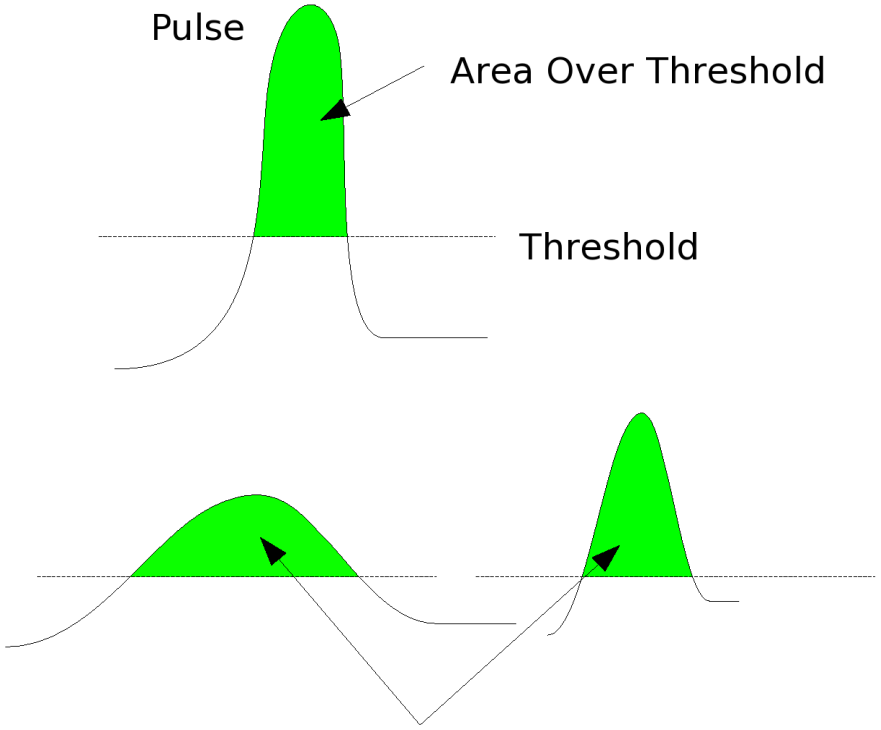

Current Performance

- Demonstrator is currently functional on a test bench
- 1.15m transmissions confirmed above detection threshold
- Transmission rates to 4.8Gbps confirmed
- Comparison of simulations and measurements suggest that lowering the noise threshold will improve transmission distance
  - Threshold is ~80mV too high for expected levels of dispersion

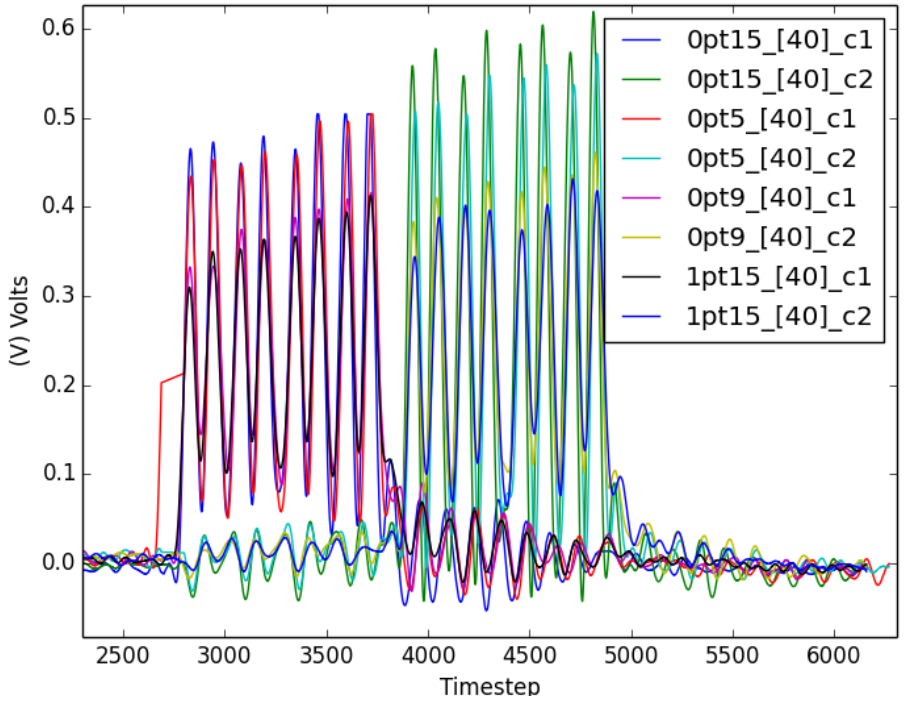

Transmission verified to 1.15m

Inter-cycle Timing Stability

RX strobe matches TX strobe with high fidelity.
Above: 5000 cycles from scope trigger, output strobe shows no significant jitter

Intra-cycle Timing Stability

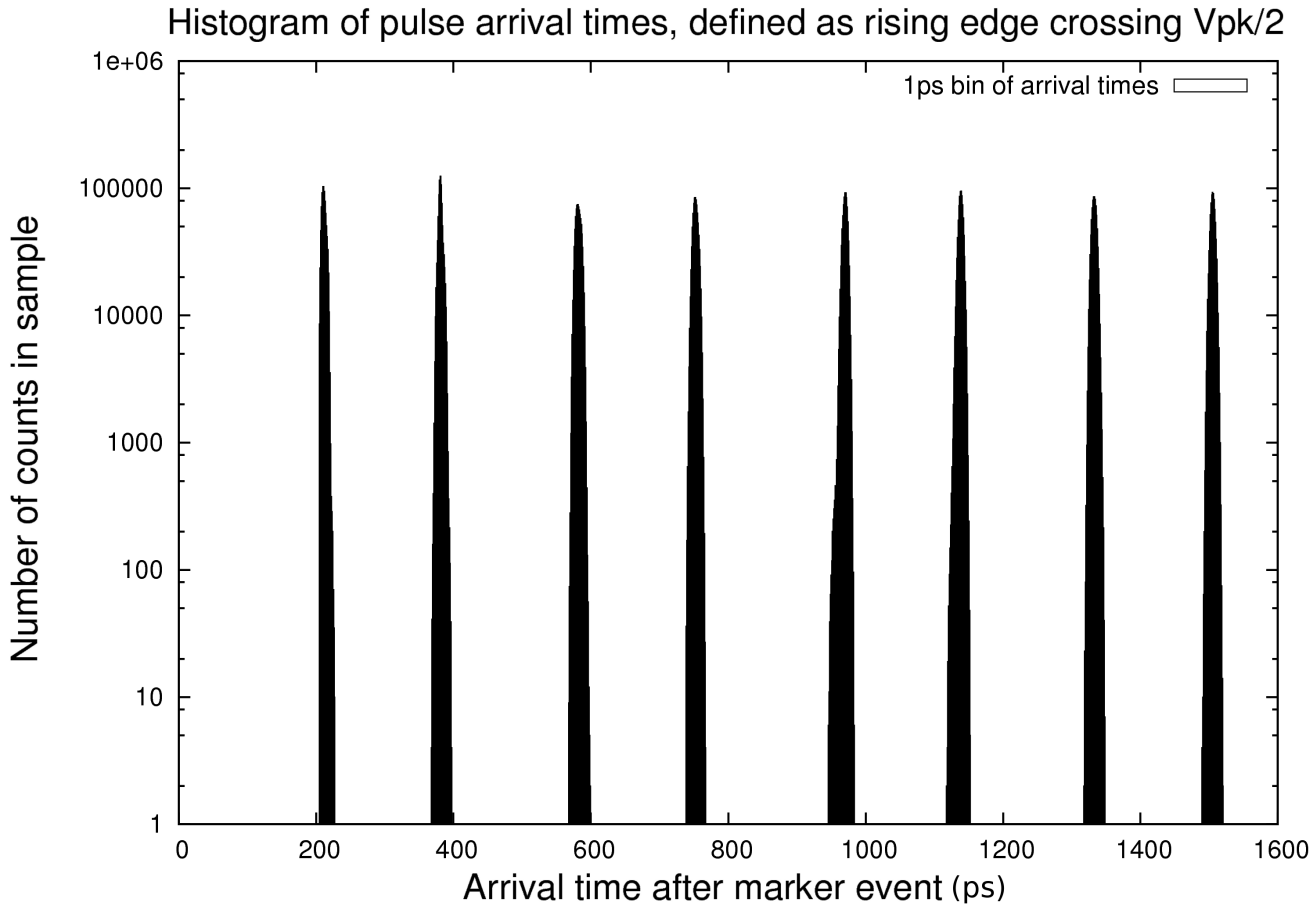

Pulses, while unevenly spaced, do show up in predictable time slots (+/- 10ps jitter)

Image from M. P. Miller, G. Magazzu and F. D. Brewer, "5 Gbps Radiation-Hardened Low-Power Pulse Serial Link," in IEEE Transactions on Nuclear Science, vol. 63, no. 1, pp. 203-212, Feb. 2016

Max Data Rate Vs Vdd

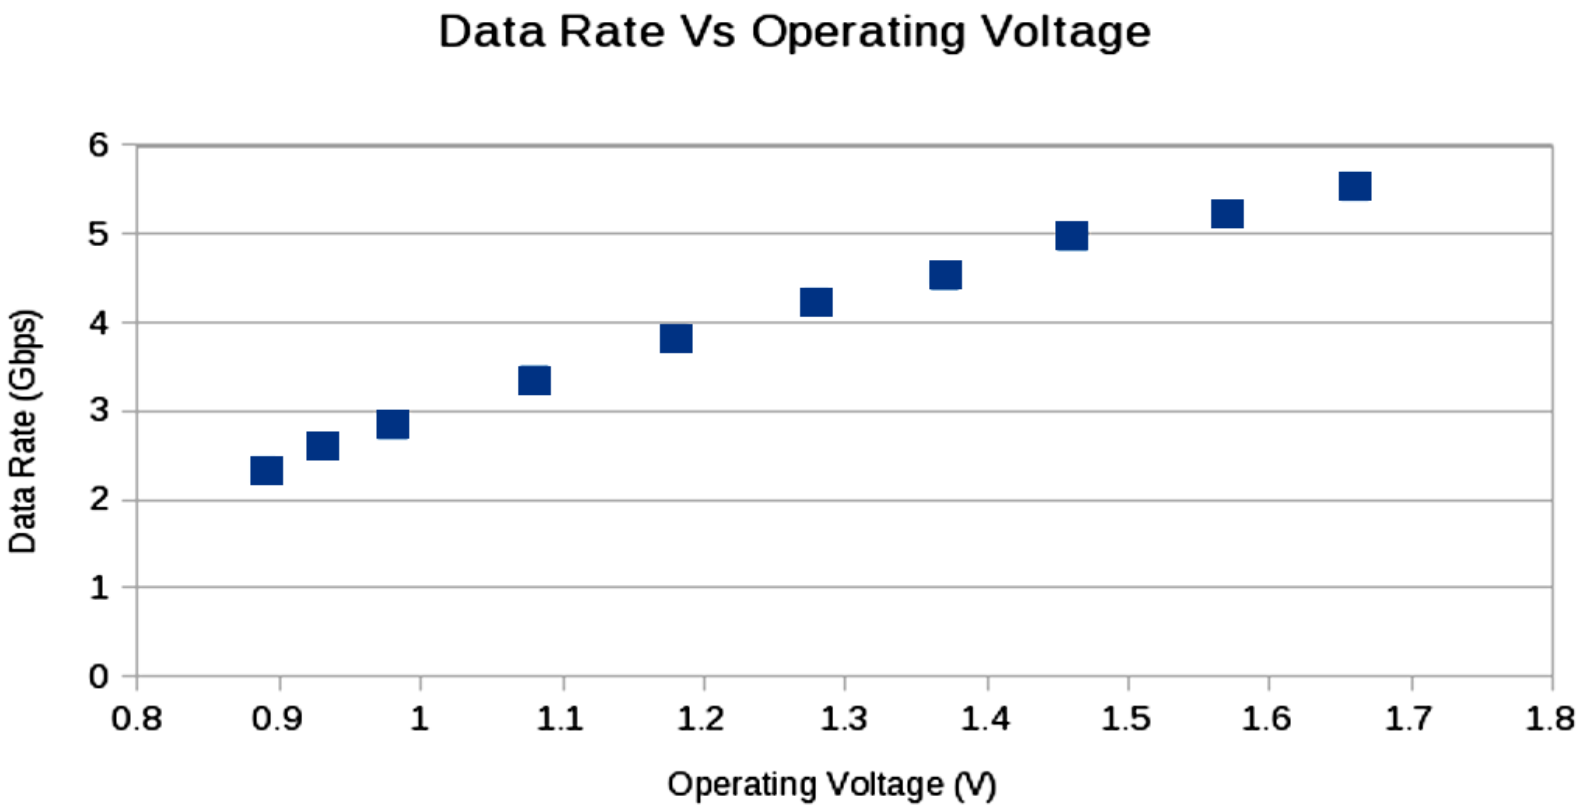

Tx Data rate for current pulse link design.
Link merely slows with loss of power down to ~0.95v

Image from M. P. Miller, G. Magazzu and F. D. Brewer, "5 Gbps Radiation-Hardened Low-Power Pulse Serial Link," in IEEE Transactions on Nuclear Science, vol. 63, no. 1, pp. 203-212, Feb. 2016

Planned improvements (part of the IBM → TSMC transition)

- Transition to true differential signaling
  - Pseudo differential signaling has propigation issues in twinax cable
  - Opportunity to improve drive power
- Investigate reducing Tx power in exchange for increase in Rx power
  - changes in signal levels,
  - active filters
- Investigate lower-power 3-pam
  - Note: operation of 3-pam could enable much higher (1.5-2x) data rates --to be investigated

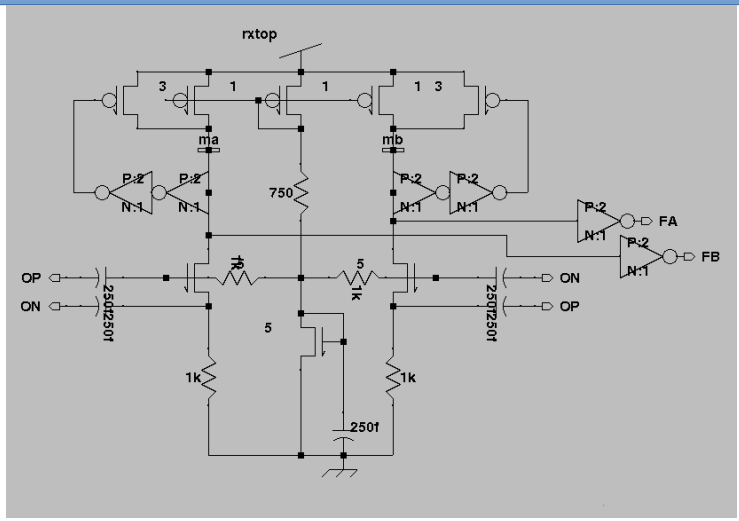

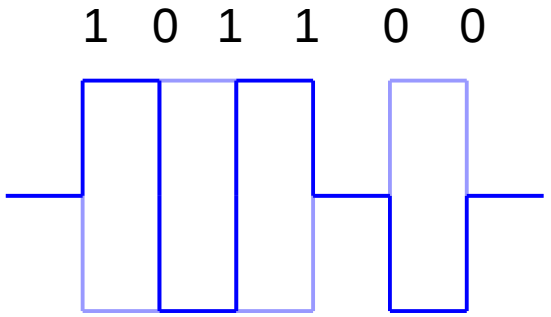

Differential Async 3-PAM

In Conclusion

- There exists the chance to get much more data out of a detector for a given amount of power
  - Integration can have big savings
  - Removing high-speed clocks also has big savings
- Pulse links (and some of the technology they are based on) can further the improvement
  - Instant Idle helps with variable data rates
  - Total energy spent is related to the average rather than peak data rate
- Some detectors may take better advantage esp. full data readout
  - Reduced module complexity
  - Reduced energy per bit
  - Can afford to do data aggregation off-module (FPGA?)