



The SLAC ATCA Platforms

Ryan Herbst, SLAC National Accelerator Laboratory


Current ATCA Development

- SLAC has been focused on ATCA based data acquisition and control systems
- RCE (Reconfigurable Cluster Element) platform is a full meshed distributed architecture, based upon network based “system on chip” elements
 - “Plug in” architecture for applications
 - Firmware and software development kits
 - Based upon Xilinx Zynq platform
 - Full mesh 10G network
 - 96 high speed back end links






- ATCA based general purpose analog & RF board
- Digital back end is based on Xilinx Ultrascale FPGA
- Supports two double wide dual-height AMC cards for analog and RF processing
- Designed targeted for mixed analog/digital applications such as: LLRF, BPMs, MPS, CMB, TES readout

COB (Cluster On Board)



RCE @ ATLAS CSC MUON Sub-System

- Replaced previous RODs which limited trigger rate $< 70\text{Khz}$
- Integrated with ATLAS timing and trigger system
- Integrated into ATLAS data acquisition
- Successful demonstration of 100Khz trigger rate @ 13% occupancy
- Meets all specifications



RCE @ Heavy Photon Search

SLAC

Power supplies
Low voltage
Sensor bias

FE boards:
Amplification
Analog to digital
Hybrid control/power

Flex cables:
Impedance controlled, low
mass signal/bias/control to
hybrids

Hybrids:
Pulse shaping
Pulse sampling
Buffering


JLab DAQ



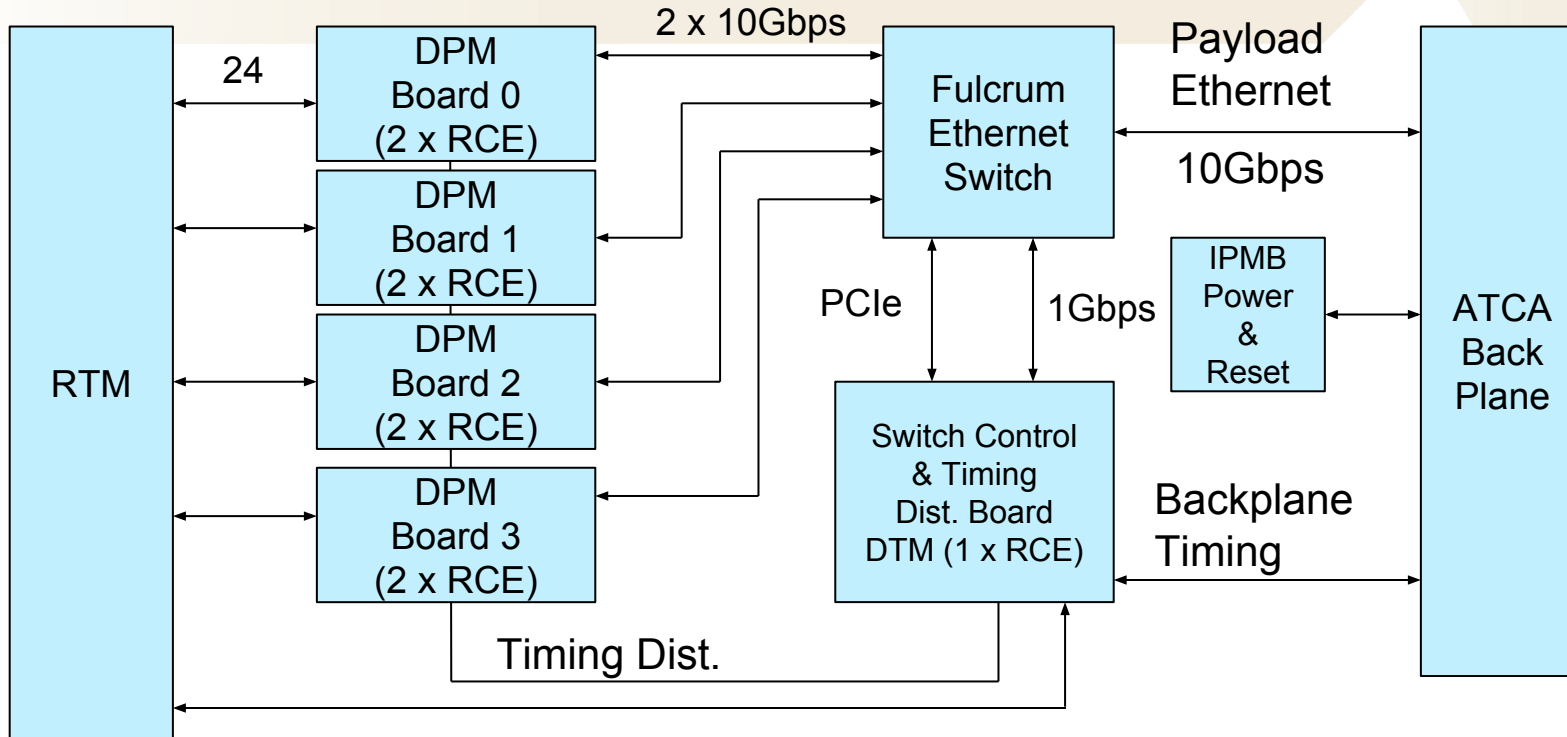
- Running experiment at Jefferson Laboratory Hall B
- Integrated with JLAB's timing and back end DAQ system (CODA)
- Took data at beginning of 2015
- Expect more data runs in 2015/2016

Upcoming Experiments Using SLAC ATCA

- LSST
 - Data acquisition and data cache
- LCLS-1 accelerator controls upgrade
 - Beam Position Monitor (BPM) upgrade
 - Low Level RF (LLRF) upgrade
- LCLS-2 high performance accelerator controls
 - Timing distribution
 - Beam position monitoring (BPM)
 - Bunch charge and bunch length monitoring
 - Machine protection system
- LCLS-2 detectors and data acquisition
- KOTO Experiment
 - Collaborating with University of Michigan
- ATLAS Inner Tracker (ITK) upgrade development (proposed)
- nEXO (baseline)
 - 2nd generation to EXO 200



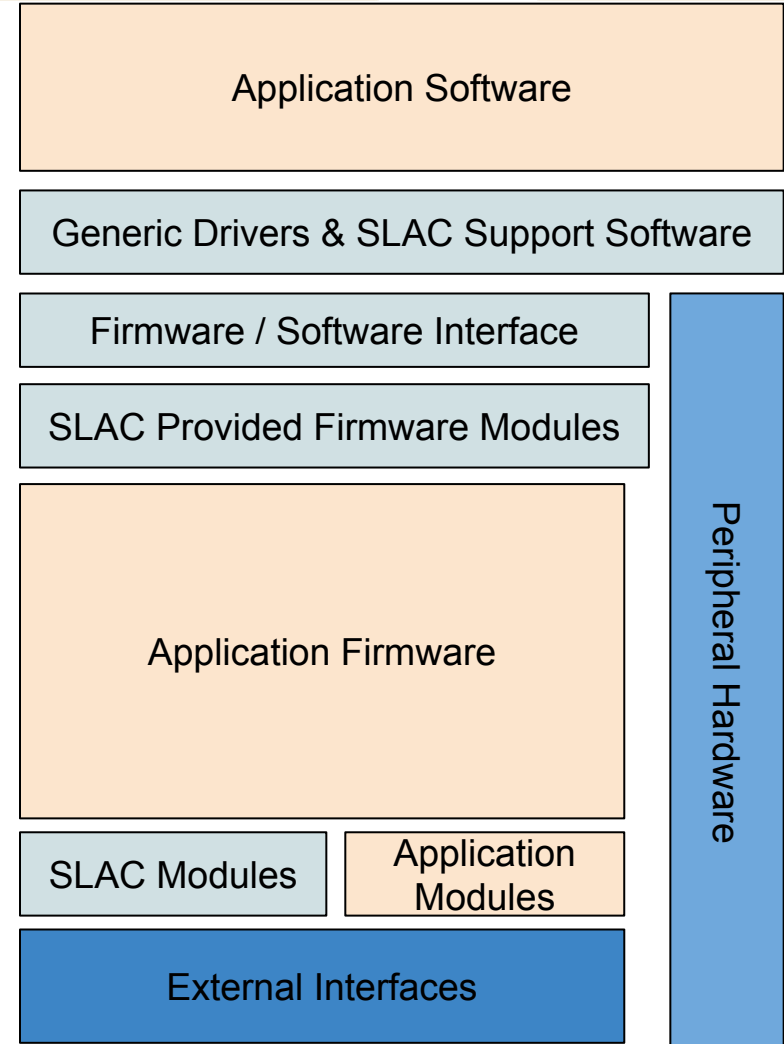
COB (Cluster On Board)





- On board 10Gbps Ethernet switch
 - Supports full mesh backplane interconnect
- 12 high speed links between RTM and each RCE
 - 96 total channels
- On board timing, trigger code and trigger data distribution
- Modular, allowing staged upgrade of individual pieces
 - System can be updated with latest technologies

RCE (Reconfigurable Cluster Element)

- Based on Xilinx ZYNQ 7000 series FPGA
 - ARM (dual-core) A-9 @ 900 MHZ
 - 1 Gbyte DDR3 memory
- Tight coupling between firmware and software
- SLAC provided utilities simplify initial bring up and development
 - Example designs with common interfaces
 - Simple build system with modular libraries
 - Well defined 'sandbox' for application development
 - Extensive library of commonly used modules
 - No commercial cores used
 - Extensive experience supporting outside collaborators who develop application specific firmware and software
- Flexible external interfaces
- Suite of management tools which provide central monitoring, reboot and firmware upgrade





RTM (Rear Transition Module)



- RTM allows the platform to be customized for applications
 - Number and type of external high speed links
 - Targeted timing interface for site timing system
- RTM can be as simple or complex as needed by the experiment
 - Low risk layout for most experiments
 - Well defined interfaces and common power/IPMI blocks
- RTMs can be analog or digital



RTM (Rear Transition Module)



- RTMs have been built with a wide variety of interface types
 - PPOD (12 channel RX or TX)
 - SFTP+
 - QSFP (4 x 10Gbps bi-directional)
 - CXP (12 x 10Gbps bi-directional)
- Space allows for simple or complex circuits
- Enough room to support FPGA(s) for complex interfaces
 - NOVA timing interface for DUNE 35-ton
- 34 IN² of usable board space
 - Compared to 25 IN² on standard PCI-Express card

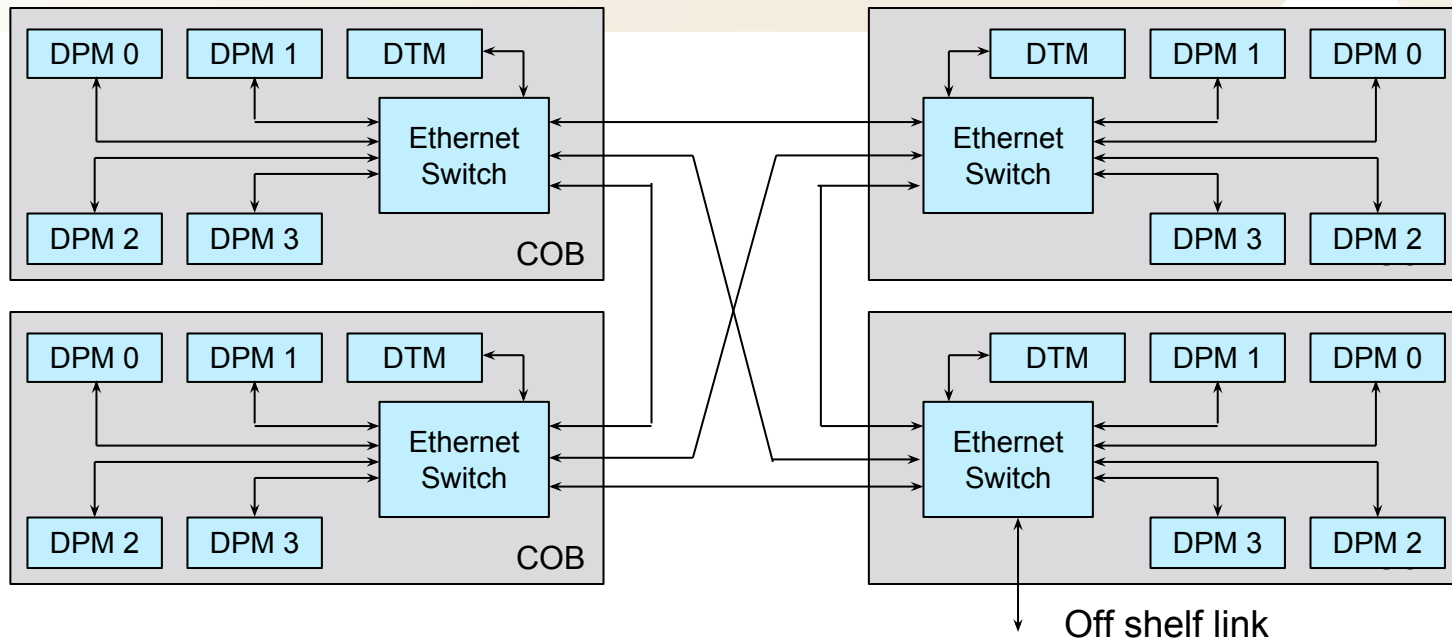
ATLAS CSC RTM

RTM (Rear Transition Module)



- RTMs are extremely flexible and can host a number of exotic implementations
 - LSST hosts RTMs for front end interfacing as well as local data storage for 2 days of camera data
 - Complex crossbars & co-processors can be supported as well
- Isolating application specific logic to the RTM allows experiment specific customization without touching critical routing and power layout associated with FPGAs

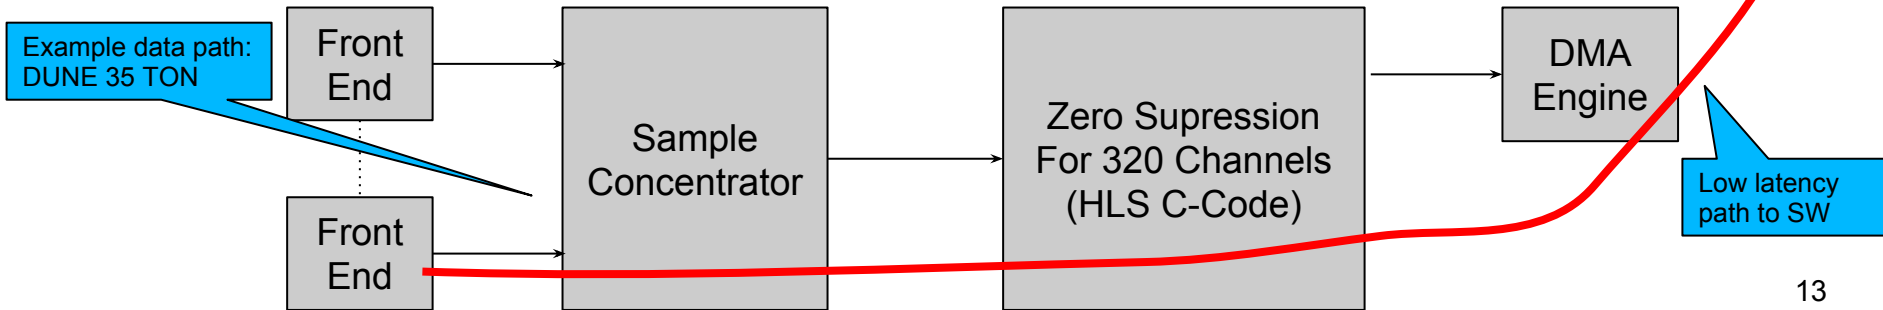
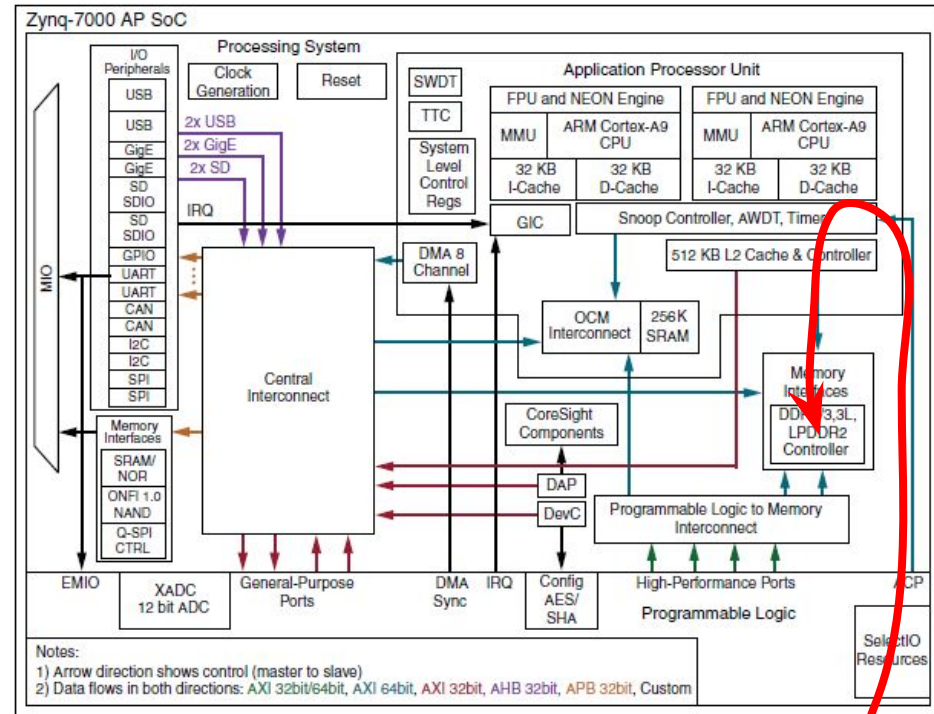
RCE Platform Clustering



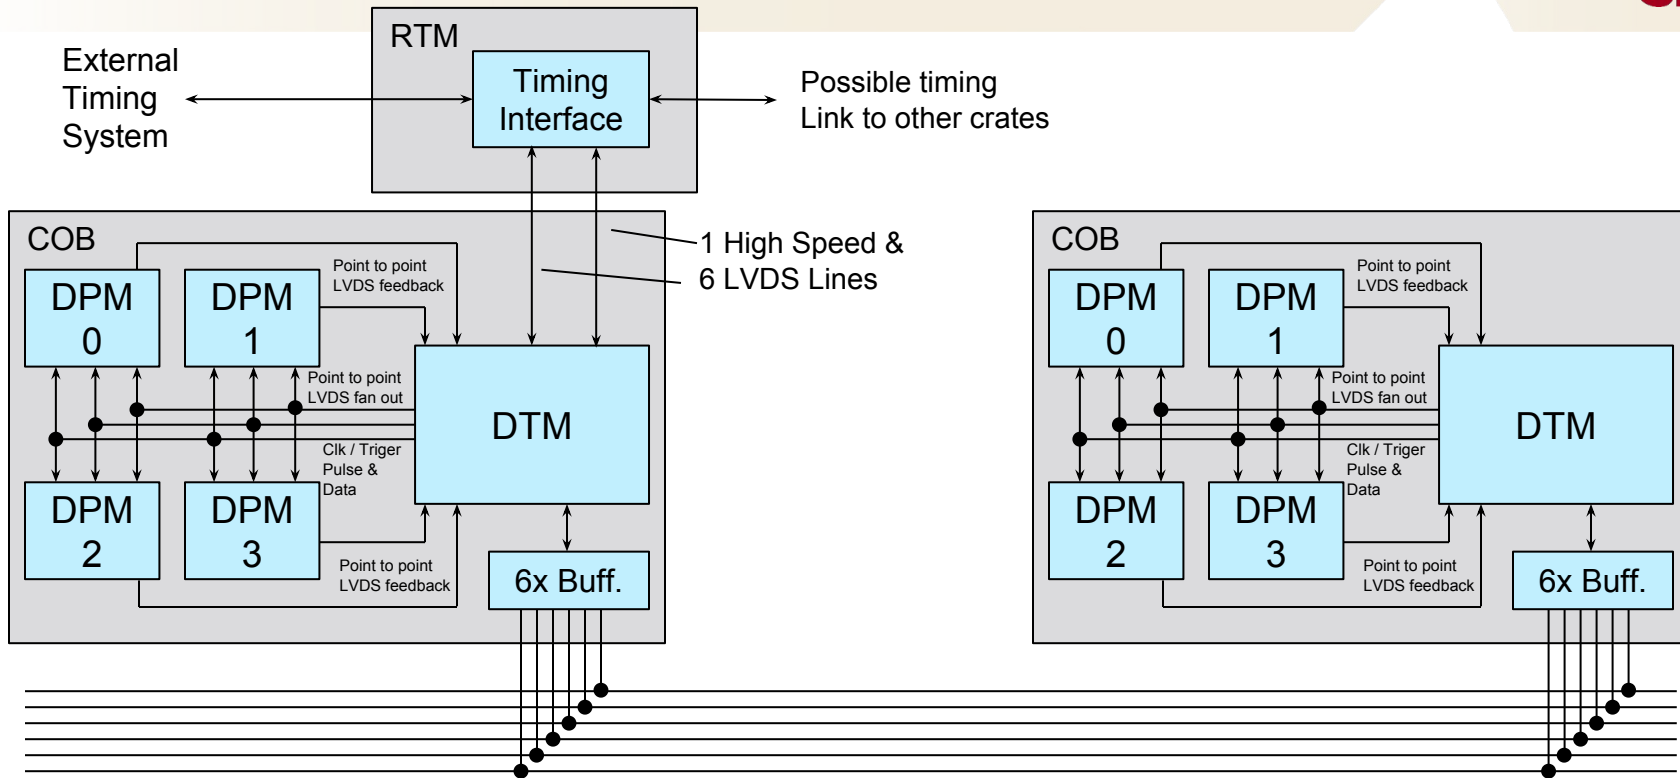
- Tightly coupled 10Gbps mesh network
 - High performance cut through latency of < 300ns
- Software APIs to facilitate the cluster configuration and communication
 - Inter-application messaging
- Firmware APIs to facilitate inter-application messaging
 - Hit maps
 - Edge channel charge
 - Veto
- Hosting FPGAs and daq code on tightly coupled nodes minimizes the number of elements in the data chain, simplifies cabling and minimizes rack space
 - Front end -> RCE -> back end event builder -> storage farm
 - Art-DAQ can be hosted on RCE similar to how CODA was hosted for HPS
- In shelf processor and switch not required, distributed processing & switch exists within each blade!

CPU To Firmware Interface

- The ARM based Zynq architecture allows a tight coupling between application firmware and application software
 - Cache aware interface for firmware data path into processor memory
 - Allows for DMA into cacheable memory (DMA directly to user space)
 - Avoid expensive cache line flushes
- FPGA register access and DMA handshaking performed over general purpose AXI busses
 - Minimal elements between the processor and the firmware
 - Does not involve a complex bus interface and multiple protocol bridges
- Application firmware within the FPGA operates as an extension to the processor
 - Experiment specific co-processor!
 - Does not suffer latency penalties of complex Ethernet or PCI-Express interconnects!



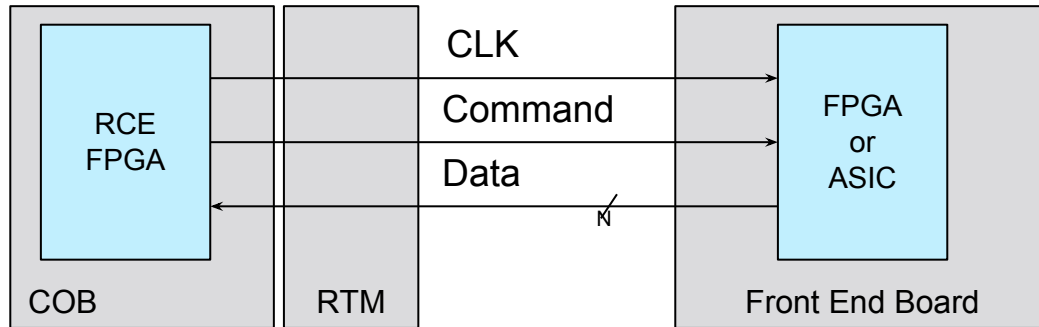
RCE Platform Timing Distribution



ATCA Backplane
6 Pairs of MLVDS
(Up to 14 slots)

- The RCE platform supports internal timing and trigger distribution
 - Lengths are matches to consistent and predictable latencies
- Eliminates errors and uncertainties with complex external cabling
- Flexible timing system interfaces through RTM and DTM firmware
 - Proven track record of interfacing to complex exotic systems

Front End Links



- RCE connects directly to front end board
 - FPGA or ASIC
- Clock is extension of RCE timing distribution
- Command stream supports register access commands and side-band trigger, reset and alignment codes
- Return data streams are 8B10B encoded
- 8B10B based protocols (such as PGP) have lower overhead compared to Ethernet as well as unlimited transfer sizes
 - Cores exists for both FPGAs and ASICs
 - Cold operation is supported
- Round trip opcodes allow system to measure latency to front end
 - Latencies differences can be compensated by adjusting phase of delivered clocks
- Direct connection to front end allows for link errors to be easily identified instead of dropped within a commercial network
 - Point to point electrical or with optical breaks

The top diagram shows a detailed FPGA layout with a callout box stating: "VHDL synthesized 8B10B data transmitter (ASIC core)".

The middle diagram shows another FPGA layout with a callout box stating: "Synthesized ASIC cores can support cold operation for DUNE".

The bottom diagram shows a third FPGA layout with a callout box stating: "SACI: VHDL synthesized generic ASIC control core (ASIC core)".

DUNE 35t Prototype

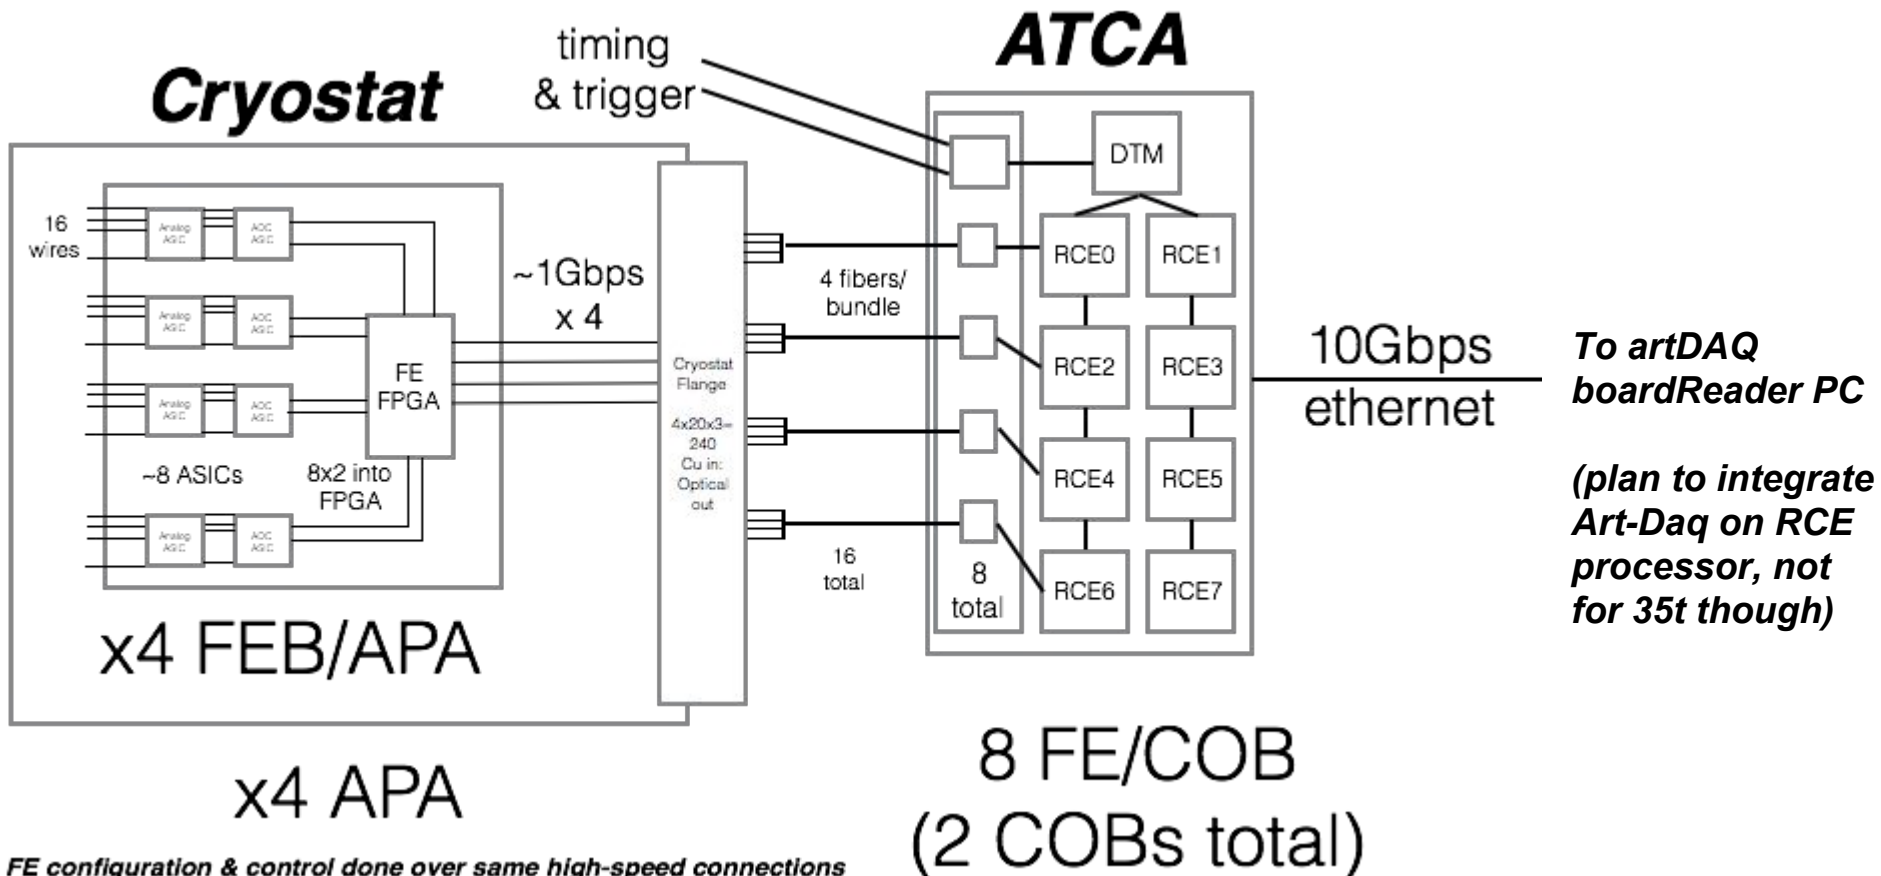


Picture by Bo Yu standing outside of the cathode plane

35t Particulars

- Four (mini) DUNE-style APAs
 - 4 FEBs each → 2,048 channels
- Use the BNL cold electronics
 - cold analog & digital ASICs, FPGA & SERDES
 - very similar to SBND/ProtoDUNE proposed design
- Read out 4 1-Gbps signal channels per-FEB (64 total)
 - copper (Gore) to the flange PCB; convert to optical on flange (QSFP)
 - each set of 4 fibers read into a single RCE→2 COBs total for 35t
 - clock & control also flows through the RCEs
- Run mostly in 2 TPC readout modes:
 - triggered, from scintillating paddles or photon system
 - continuous, with waveform extraction (almost)
- Backend DAQ uses artDAQ to read RCEs, build & aggregate
 - we have plans to integrate some parts of artDAQ into RCE software in the future
 - note that the backend computing network/storage was, for various reasons, designed for high throughput


35t TPC Readout



*To artDAQ boardReader PC
(plan to integrate Art-Daq on RCE processor, not for 35t though)*


FE configuration & control done over same high-speed connections as signal, with dedicated I2C links as a backup

The 35t ATCA shelf & boards @ PC4



One of the RTMs is removed for clarity .

Readout for 2048 TPC channels



RCE System Control, Monitoring & Diagnostics GUI

System Control

System State

Device Name	Run State	Status	Reg Rx	Timeouts	Data/Event	Errors
Data DPM 0	Enable	Warning	7921504	26	561088	2
Data DPM 1	Enable	Ready	14340623	35	564209	0
Data DPM 2	Enable	Ready	16598887	265	562103	0
Data DPM 3	Enable	Ready	16450799	215	564201	0
Data DPM 4	Enable	Ready	16604479	309	566298	0
Data DPM 5	Enable	Ready	16423138	204	564202	0
Data DPM 6	Enable	Ready	16453725	229	566298	0
Data DPM 7	Enable	Ready	415974	0	568404	0

Configuration And Control

HardReset SoftReset RefreshState

Set Defaults Load Settings Reset Counters Pgp Link Status

Run Control

Run State:

Data DPM 0

System Commands Status Configuration

Variable	Value	Decode	Info
DataBuffer(0)			
BufferCount	0x320	800	Info
BurstCount	0x0	0	Info
BurstSize	0x0	0 frames	Info
BusyCount	0x31	49	Info
DropCount	0x0	0	Info
FrameMissing	0x1ae	430	Info
RxBw	4223.7	4223.7 Mbps	Info
RxCount	0x9e360	648032	Info
RxErrors	0x2	2	Info
RxRate	1999.8	1999.8 Hz	Info
RxSize	0x40748	264008 Bytes	Info
TriggerRate	0.9	0.9 Hz	Info
Triggers	0x135	309	Info
TxBw	0.0	0 Mbps	Info
TxCount	0x9baa2	637602	Info
TxErrors	0x0	0	Info
TxPend	0x64	100	Info
TxRate	0.0	0 Hz	Info
TxSize	0x4075c	264028	Info
WorkCount	0x2710	10000	Info
DataPath(0)			
DmaBusyCnt	0x0		Info
ExtTrig	False		Info
ExtTrigCnt	0x14e	334	Info
ExtTrigRate	0x1	1 Hz	Info
LbneFrameDrop	0x0		Info
LbneFrameDrop0	0x0	0	Info
LbneFrameDrop1	0x0	0	Info
LbneFrameDrop2	0x0	0	Info

Read Status

- The GUI is a nice tool to control & monitor the FEBs, RCEs, and data path.
- Communication is via XMLRPC, so is easy to script & program controls as well→this is how we do run control via artDAQ
- Hooks available for EPICs integration

35t Zero Suppression (Waveform Extraction)

- Waveform extraction algorithm implemented for 128 channels using ~25% of FPGA resources
- Adjustable parameters to optimize performance depending on S/N conditions
- Handling of data quality problems (stuck bit, etc.)
- Future plans: nearest neighbors, correlated noise, signal processing

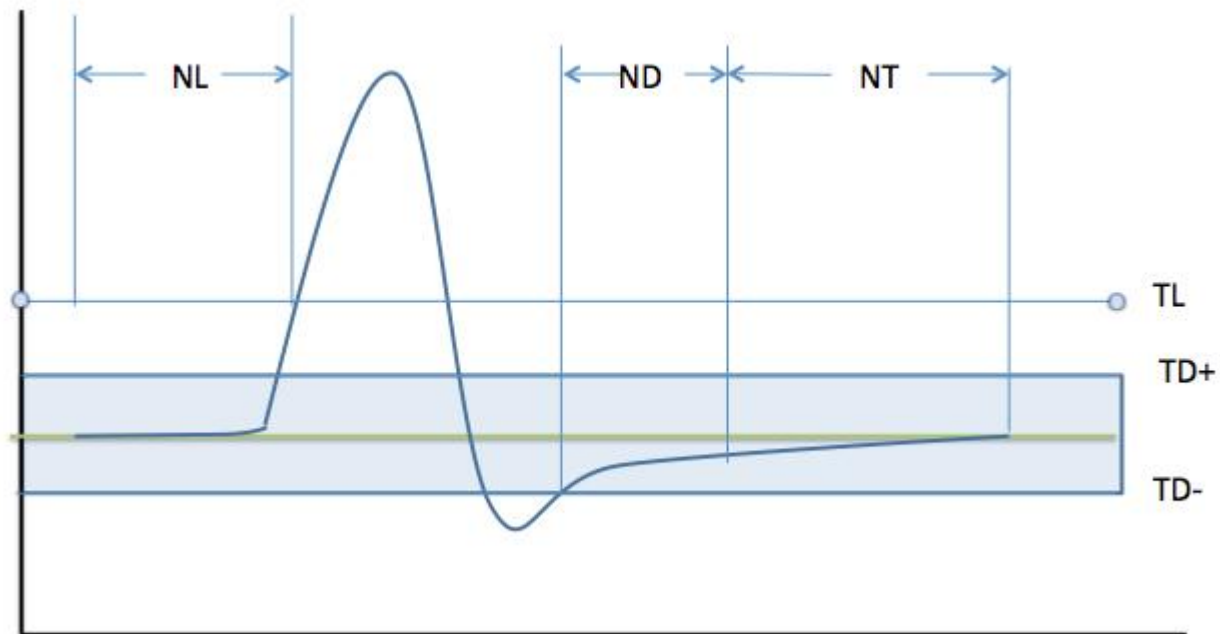
NL: # of Front Porch samples

TL: Threshold

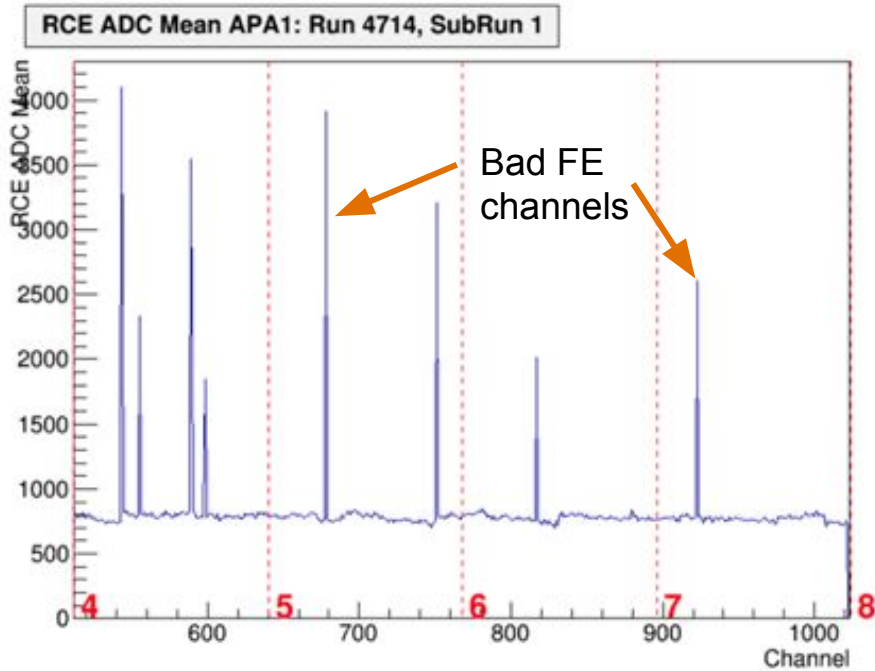
TD+/TD-: Dead band limits

ND: # of required Dead Band samples

NT: Back Porch



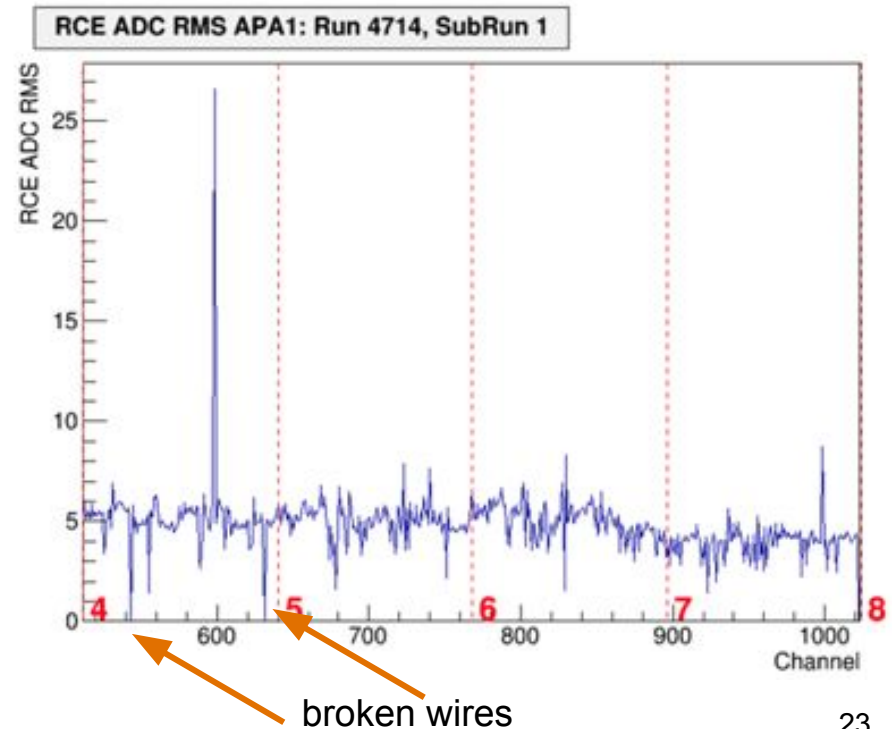
Online Monitoring



These plots show the ADC mean & RMS from online monitoring of a run taken November 18.

We use artDAQ to perform online monitoring of the TPC, SSP, and trigger board readouts, along with general DAQ performance.

In addition, we use ganglia in real time to monitor data throughput and artDAQ behavior.



35t DAQ Progress, Status, and Plans

- High speed links between FEB & RCEs are solid
 - data transmission and clock & control communications are working as designed
- We can run for ~ hours* without failures with all 2,048 channels
 - *our bottlenecks are disk writing (50 MBps) and then the backend network (~90MBps)...35t DAQ not designed to store a bunch of data!
- Currently taking lots of noise data and generally trying to explore limits of system, particularly sensible back-pressure handling across all DAQ systems
- RCE-based waveform extraction firm/software is in development and on schedule to be deployed in December
- Cool down should happen in December, plan to start taking cosmic data then


The SLAC Team



- SLAC Technology And Innovation Directorate (TID)
 - Mike Fazio Director
- SLAC Advanced Innovation For Research (AIR) Division
 - Gunther Haller Division Director
- Closely linked departments encompassing all areas of detector and controls development
 - Sensors: Chris Kenney
 - Integrated Circuits: Angelo Dragonne
 - Electronics Systems: Ryan Herbst
 - RF Systems: Joe Frisch
 - Advanced Controls Systems: Ernest Williams
 - Advanced Data Systems: Matt Weaver
- Professional engineers and scientists from a variety of backgrounds
 - Emphasis on system design with open communication and involvement from all departments
 - Experiences working on many projects
 - Open discussions allowing all members of the division to contribute and learn
- Very close relationship with the Science Directorate
 - Partnership approach to experiment design
 - Cross discipline contribution with engineers and scientists involved in all levels of design
 - Engineers actively support data analysis and scientists contribute to firmware and software design
- Extensive experience building complex systems and handing off to the experiment operators for application development
 - Teaching junior engineers and scientists how to contribute to firmware and software
 - Providing well defined sandboxes to simplify system development
 - Support for higher level tools for application design
 - SystemGen design flow from Matlab for engineers more experience in analog and RF design
 - HLS for scientists and engineers more experience in C-programming

Third Party Integration


- Very skilled at interfacing to foreign protocols
 - Timing systems from JLAB, LCLS-1, ATLAS, NOVA, KOTO
 - Often re-implemented at the protocol layer or adaption of outside firmware
 - JLAB required import Virtex-4 schematic based firmware into Zynq FPGA in Vivado
 - ATLAS SLINK protocol to ROS nodes
 - Camera Link protocol
- Modern detector systems do not have arbitrary lines between sections
 - Typically the demarcation point is within the FPGA and software
 - Systems must be designed with this in mind
 - SLAC has extensive experience with incorporating third party and outside firmware/software into its systems
 - Successful integration of JLAB CODA ROC software
 - Plans for integrating ART-DAQ
 - ATLAS run control software
 - EVG/EVR protocols from Micro-research
 - We have also integrated our protocols into outside designs



SLAC 'PGP-Card'

Generic timing system interface

8 high speed serial links (6.6Gbps)
PGP / CameraLink /
SLINK / SATA &
others....



LCLS-1 / LCLS-2
Timing Receiver

- The SLAC RCE platform has a proven track record of providing reliable data acquisition, detector control and timing distribution in a number of experiments
 - Projects involved collaboration and shared development amongst a number of institutions
- ATCA is the platform of the future
 - Billion dollar industry
 - Adopted by telecom, datacom and large physics experiments
 - Compact design which minimizes rack space
- Mature technology which exists and is being used
 - Plans to transfer RCE to become a product at commercial companies
- SLAC has a long track record collaborating with outside laboratories and institutions
 - Production RCEs are being distributed to external collaborators:
 - Harvard, Brookhaven, CERN, Fermilab, Stanford, JLAB, Oxford, Stony Brook, Berkeley, University of Geneva, Glasgow, Liverpool, Goettingen, UCL, RAS, Adelaide, KEK, MPI-Munich, Bern, Ljubljana, UCL, IFAE Barcelona, LBNL, NYU, Oklahoma State, U Illinois Urbana, U Washington, ANL
 - Collaborators develop application software and firmware based upon RCE base hardware and framework
 - Running and/or baselined for many experiments

- The RCE is Powerful and flexible while allowing application specific customization
 - Includes all parts needed for a detector system
 - Timing, Control & Data acquisition
 - Minimizes rack space & cabling
 - Integrated control, status monitoring & debugging
 - No need for separate systems for each function
 - Tight coupling of firmware and software for high performance data processing, while simplifying design debug
- Modular & scalable design with support for both large and small systems
 - Well defined upgrade path to support next generation components & interconnects
 - Low risk when dealing with product lifespans and component end of life
 - Flexible architecture allowing integration of third party interfaces, timing systems and DAQ systems
 - Allows development of experiment specific interfaces without modifying core complex board design
- Fits well within the DUNE detector design
 - One APA per COB
 - Flexible architecture allowing interfaces to be adjusted to evolving detector design
 - Available FPGA and CPU resources allow for complex data processing