

Trigger and Data Acquisition

Mike Huffer, SLAC and Wesley Smith, U. Wisconsin
CPAD Instrumentation Frontier Meeting
October 7, 2015

- Thanks for an excellent series of presentations:
 - Darin Acosta, U. Florida
 - Elliot Lipeles, U. Pennsylvania
 - Remi Mommsen, FNAL
 - Jinlong Zhang, ANL
 - Ted Liu, Fermilab
 - Anders Ryd, Cornell
 - Markus Joos, CERN
 - Ryan Herbst , SLAC
 - Kurt Biery, FNAL
 - Hucheng Chen, BNL

- Able to preserve Run 1 (2012) level thresholds
 - This is key because HL-LHC is not just about very high p_T
 - Weak coupling physics at moderate electroweak scale is a key part of the program and a trigger challenge
- Elimination of L2 in favor of split hardware levels
- Split Level Hardware System
 - Allows tracking and fine granularity calorimeter information before full detector readout
- General Migration of Offline-like algorithms to High Level Trigger then to L1

- L1 Track
 - Only input information from region of interest around L0 object
 - Requirements ~ 10 μ s
 - AM chip-based pattern recognition
 - Content addressable memory searches predefined list of paths="patterns"
 - Duplicate patterns to minimize latency
 - Linearized track fits in FPGA
 - Linearized = treat track as a perturbation for a set of precalculated paths
- L1 "Global"
 - Fine granularity offline-like calorimeter processing
 - Track-shower matching, Jet-vertex association
 - Time-multiplexed: multiple events (order 10) processed in parallel
 - Implementation to be studied (FPGA, CPU, GPU, ...)
- FTK++ similar to L1Track and phase-1 FTK
 - Hardware track preprocessor for Event Filter
 - Full detector track reconstruction at ~100 KHz
 - target mainly low energy hadronic activity

- CMS plans to upgrade its detectors and the Level-1 Trigger electronics for HL-LHC
 - For the Level-1 (HLT) Trigger, the goal is to maintain sensitivity for electroweak and TeV scale physics with rate < 750 (7.5) kHz and 12.5 μ s latency
 - Bring more of what was done at HLT to the Level-1 Trigger
- Initial algorithms demonstrate meeting goal, but much more R&D required for implemented algorithms and system design
 - Likely performance can improve further (e.g. particle flow at Level-1)
- Many Level-1 trigger subsystems will absorb many tens of Tbps data input
 - Order of magnitude higher than current Phase-1 upgrade
- Trigger logic similarly expands by a large factor
 - e.g. tracking logic for the silicon detector trigger – full coverage
- Need to achieve this during the next decade within roughly the same budget and overall system size as for the current Level-1 trigger system

- I/O bandwidth and system/PCB design
 - Electronic systems designed to handle O(50) Tbps input
 - High bandwidth optical data links at O(25) Gbps
 - High bandwidth backplane communication
- Logic implementation
 - Ultrascale FPGAs
 - And associated high-level synthesis tools to improve productivity
 - Associative Memory ASIC, 3D technology for high density
- Large, fast memory access
 - Access to >GB memory *near logic resources* for quick, complex calculations via large lookup table (P_T assignment in tracking triggers)

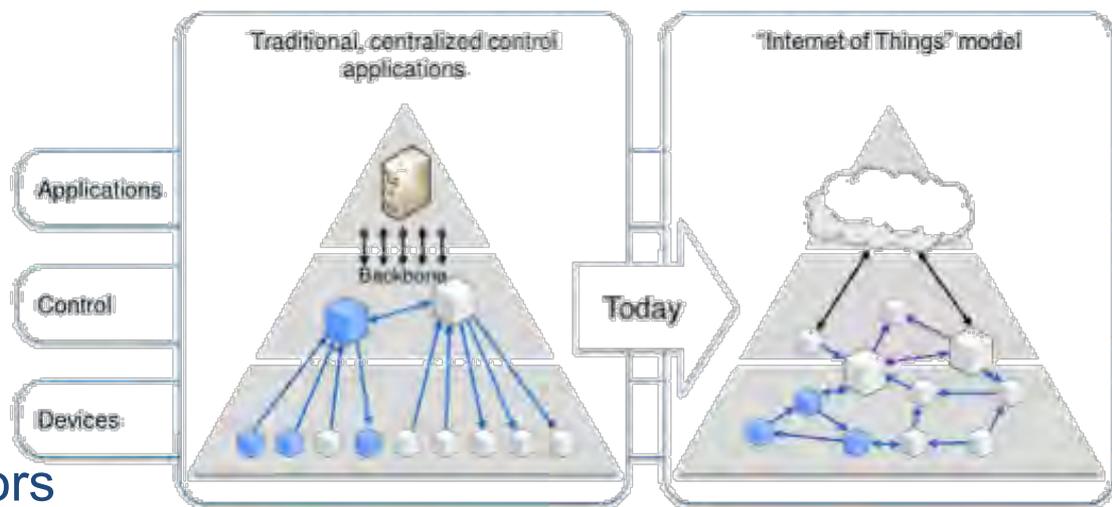
- Electronic controls
 - Clock and trigger control signal distribution to systems
 - On-board processors (e.g. Xilinx ZYNQ) for board control and monitoring

- Platform technologies
 - Evaluate platforms for backplane bandwidth, card real estate, power density ...
 - xTCA? (Telecommunications Architecture)
 - CMS Phase-1 upgrade electronics are based on μ TCA standard
 - Advanced TCA for some Phase 2 implementations?

Note: important to maintain ties with industry as well as ECE colleagues at universities on technology and tool developments

- CMS has a complete new DAQ system for LHC run 2
 - State-of-the-art technology
 - Order of magnitude smaller than old DAQ system
 - Achieves run 1 performance
 - More work needed to use full potential
 - New sub-system readouts are being integrated
- The changes for run 3 (2019) will be less radical
 - But still a lot of planning and work will be needed
- Open field for phase 2 DAQ (2025)
 - Evolution of current DAQ as baseline
 - Opportunities for radical new ideas (next slide)
 - Opportunities for R&D on modest budget
 - High-speed, low-mass, rad-hard readout links (mostly unidirectional)
 - Data-reduction schemes on- or near-detector w/o affecting physics
 - Tie event building and selection into a mesh

- Traditional event building uses hardware inefficiently
 - Needs a lot of resources to transport data which is mostly unused
 - Used only for L1 trigger
 - Not processed by HLT and then discarded
 - Network b/w is used only in one direction
- Think more of a mesh (or Internet of Things)
 - Leave data as close as possible to the detector
 - Pre-process it locally
 - Specialized processors (custom or commercial)
 - Generic CPUs
 - Access it remotely
 - Event-building on demand
 - Continuous calibrations with feedback to processors
 - Allows near offline-quality selection to reduce the event rate
 - Blurs boundary between online and offline reconstruction



- Higher level of commonality between detectors
 - A common object providing functionalities today implemented in detector-specific back-end custom electronics (ROD)
- Increased use of COTS components
 - All ROD-like functionality (including data processing) could most likely be implemented in standard computers by Phase-II
- Performance scalability built-in
 - Programmable connectivity between detector FE and DAQ
- Capability to disentangle ROD-like functions from hardware implementation
 - Different granularity for monitoring, control, data handling ...
 - DCS and DAQ traffic separation

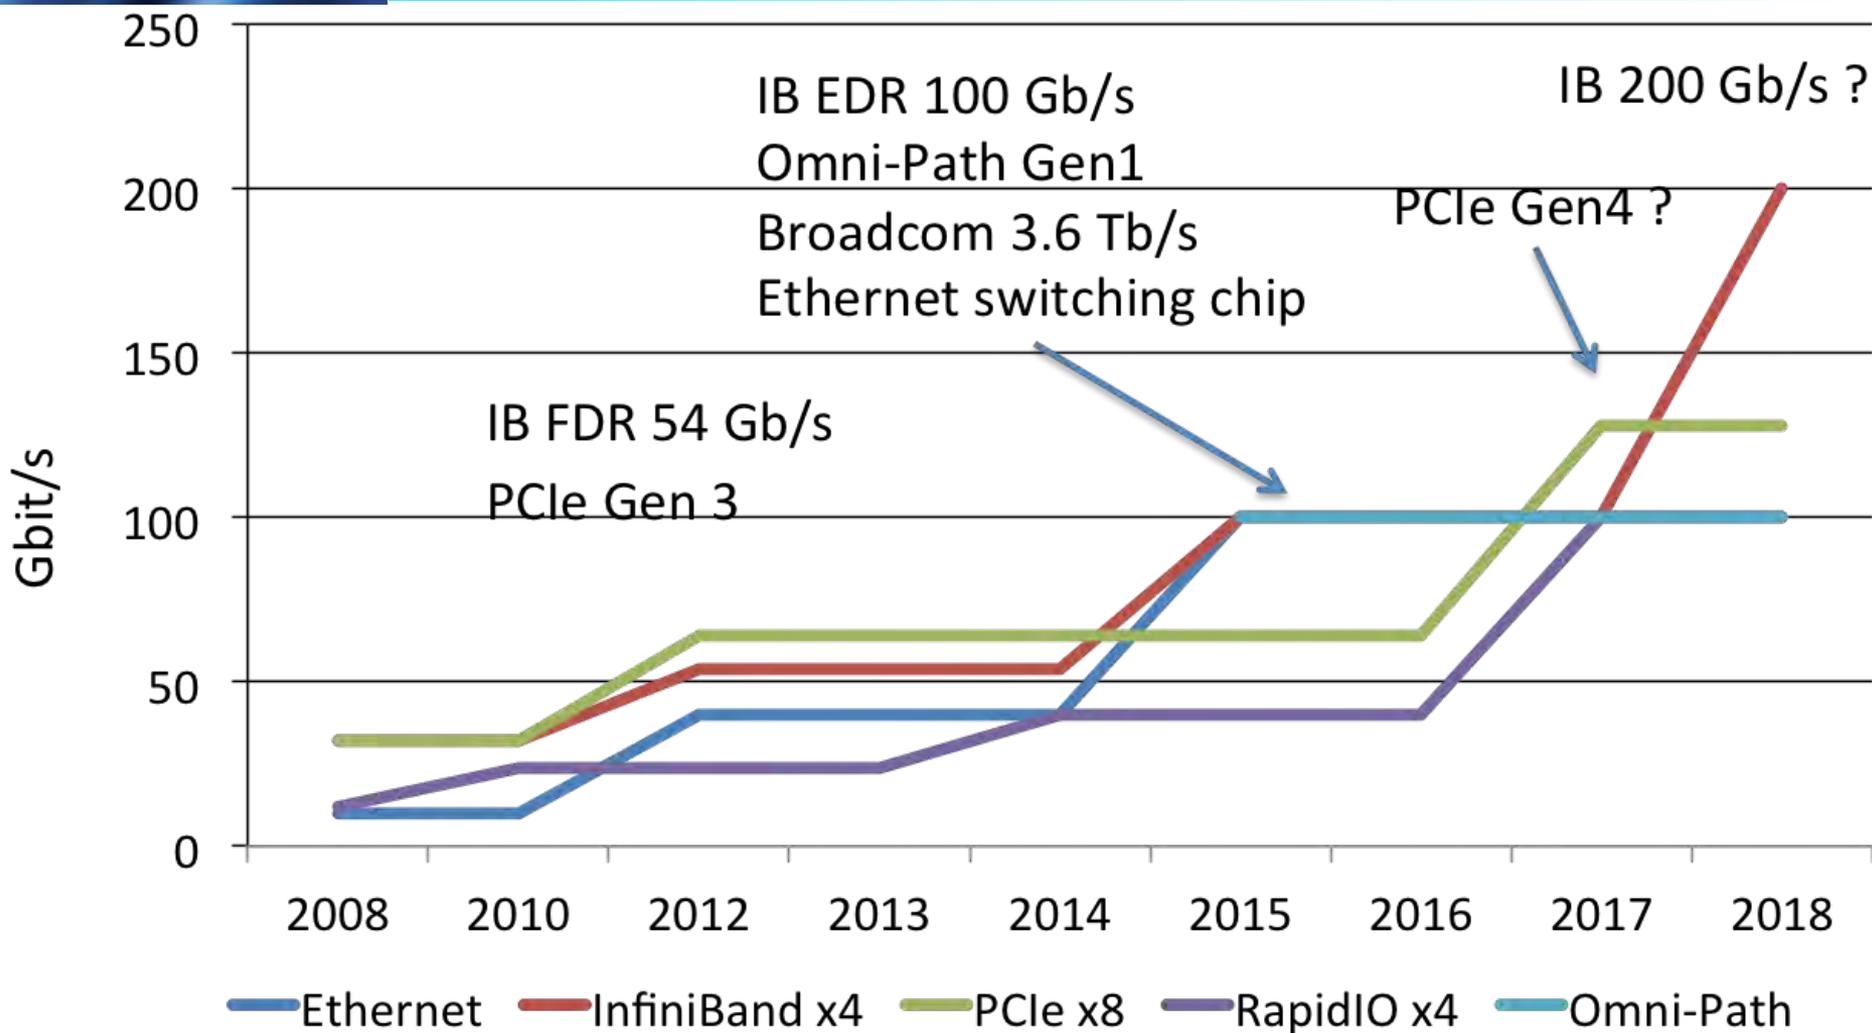
- Higher and higher trigger rates
 - Triggerless not yet possible
- PC-based single-stage data aggregation
 - Ethernet or InfiniBand
 - PCIe4
- Network bandwidth becoming very affordable
 - Changing from the philosophy of “move minimal amount of data”
 - Capability for full event building @ L1A rate (even decouple from HLT)
- Heterogeneous HLT computing (ASIC/FPGAs, GPGPUs, ...)
 - A factor of $O(50)$ in HLT computing power needed wrt to Run 1
 - Full event tracking @ 100 kHz with special hardware (FTK++, GPU, etc)
 - General purpose PCs (performance increase of a factor 10 in ~10 years)
 - Integrated solutions (e.g., Open Compute Project)?
- Tight integration with offline
 - From the blur boundary to the full fusion?
 - Utilization of online resources during non-beam time (e.g., Run Simulation on experiment HLT nodes)

	Type	Max Performance ¹	Max Transceivers	Peak Bandwidth ²
Virtex UltraScale+	GTY	32.75	128	8,384 Gb/s
Kintex UltraScale+	GTH/GTY	16.3/32.75	44/32	3,268 Gb/s
Virtex UltraScale	GTH/GTY	16.3/30.5	60/60	5,616 Gb/s
Kintex UltraScale	GTH/GTY	16.3/16.3	64	2,086 Gb/s
Virtex-7	GTX/GTH/GTZ	12.5/13.1/28.05	56/96/16 ³	2,784 Gb/s
Kintex-7	GTX	12.5	32	800 Gb/s
Artix-7	GTP	6.6	16	211 Gb/s
Zynq UltraScale+	GTR/GTH/GTY	6.0/16.3/32.75	4/44/28	3,268 Gb/s
Zynq-7000	GTX	12.5	16	400 Gb/s

FPGA	Memories*	DSPs*
Virtex 7	~50 Mbits	3,600
Ultrascale	~130 Mbits	5,500
Ultrascale+	~500 Mbits	11,900

- *Trigger uses these

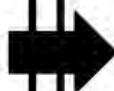
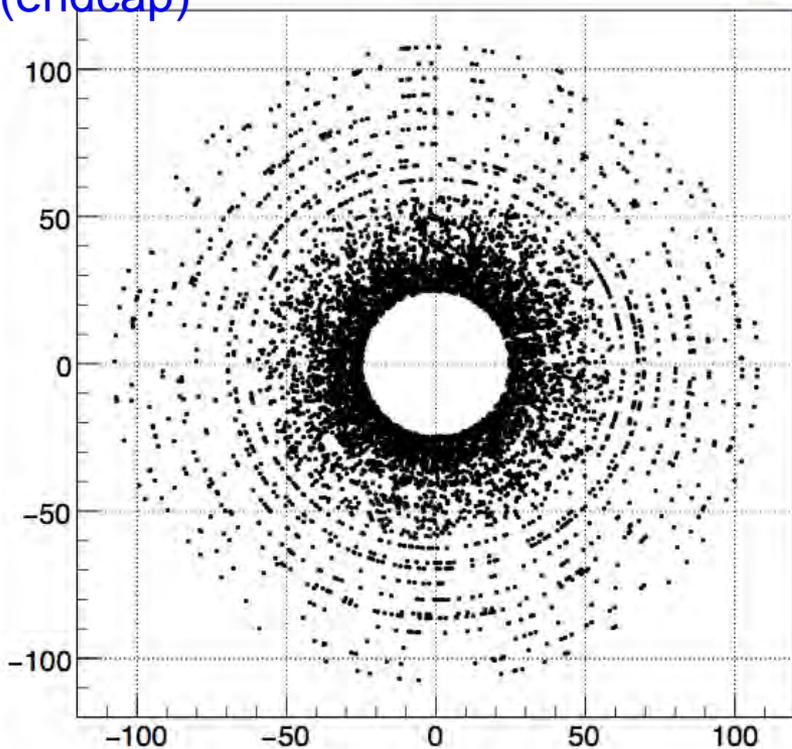
- Readout system will utilize these serDes speeds or faster, so GBT, even IpGBT(to be used for Phase-II) appear modest
 - Lightweight protocol being considered in some cases



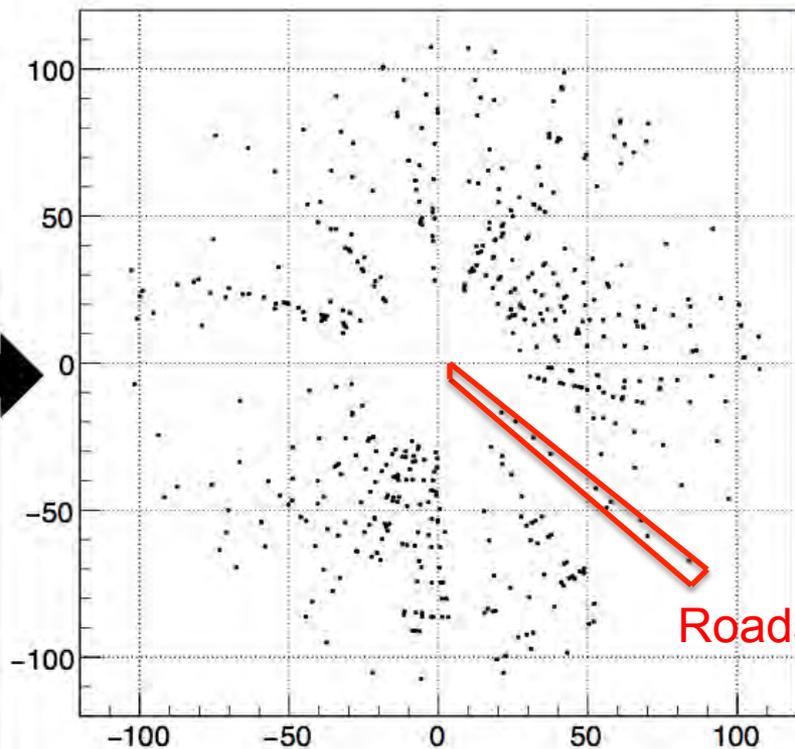
- Network for ~500 of 100 GBE links not a problem in 2024 (Phase-II)
- PCIeGen4 expected in later 2017

- Before and after AM stage: Associative Memory for CMS case at HL-LHC
~100M patterns, or ~2M pattern per trigger tower, or few x 100K per chip
- A factor of > 10 reduction of occupancy, and more importantly, stubs/hits are organized in roads (“hits of interest”), making the track fitting task inside FPGA much easier.

PU200+4 Top event simulation
(endcap)



After AM

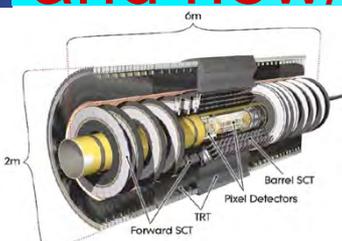




Module design: modern packaging technologies

Frontier Physics reach at high luminosity LHC require the most advanced Real Time processing technology.

LHC developments in this respect are also a "playground" for developing and experiencing novel Real Time field.



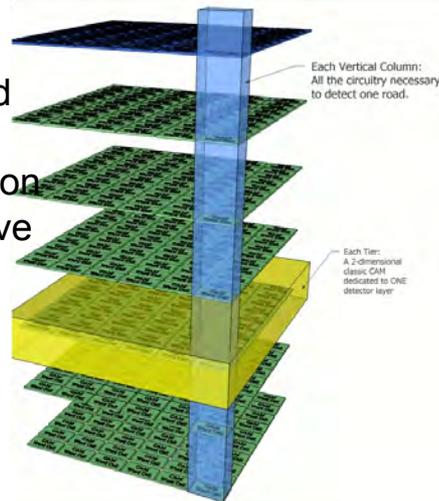
Data transfer:

Data formatting

Telecommunication technology natural fit

VIPRAM in 3D (DOE CDRD)

VIPRAM:
Vertically
Integrated
Pattern
Recognition
Associative
Memory

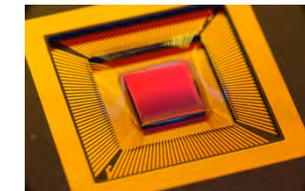
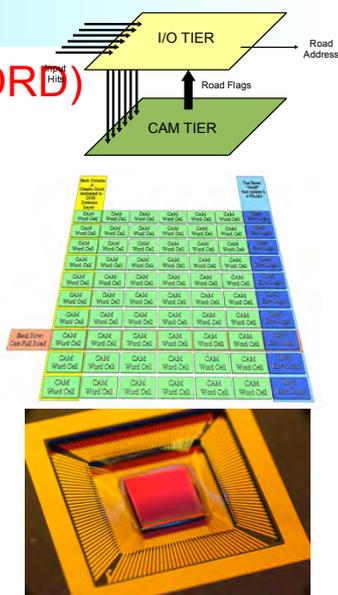


Pattern Recognition

Associative Memories: 3D technology may allow major breakthrough here

Track Fitting

FPGA technology



tracks

- Approximate number of DSP operations required for the different steps in the algorithm

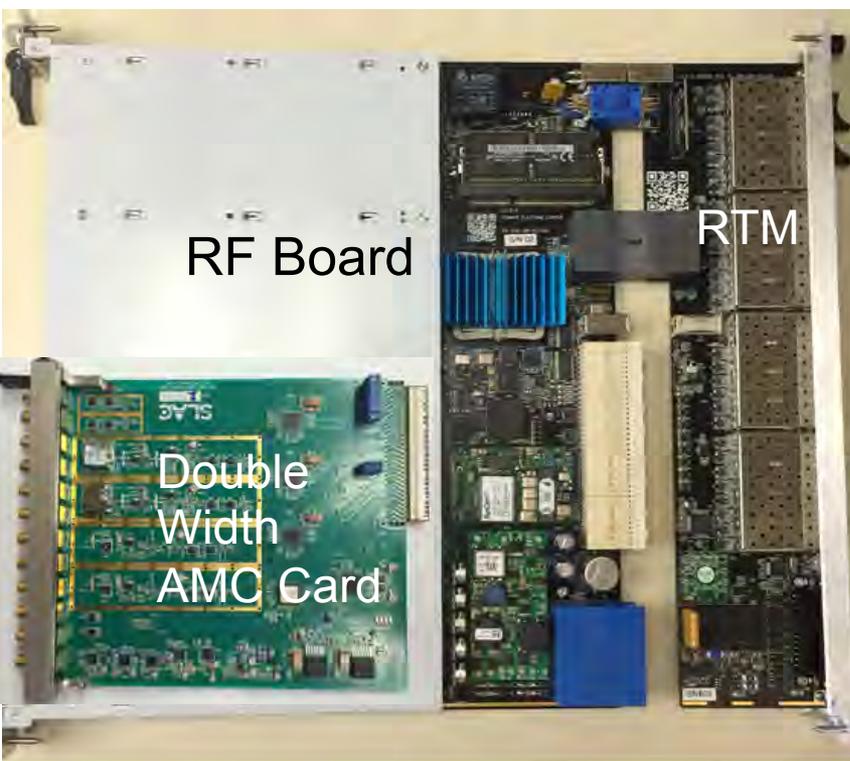
Task	# objects	DSP operations per object	Total DSP operations
Tracklet Parameters	60	20	1200
Tracklet Projections (4 per tracklet)	60	4×10	2400
Matching	200	4	800
Track fit	10	20	200
Total			4600

- Assuming a factor of 4 time multiplexing (100 ns between events) and a 300 MHz project clock each DSP in the FPGA can perform 30 operations. Today’s Virtex 7 (690T) has about 3600 DSPs and can perform ~100,000 operations per event
 - **Need about 5% of the resources**
- The newly announced Ultrascale+ FPGAs will have up to 12,000 DSPs and O(1%) would be needed for the L1 tracking
 - **Challenge is to handle combinatorics and tails**

- How to give xTCA a future in HEP?
- ATCA and uTCA share many concepts
 - Exploit this for controls and DAQ S/W
- In terms of age and performance there are no competing standards at the moment
 - also not in terms of complexity
- To master this complexity our community should:
 - Invest in product evaluations and interoperability testing
 - Develop and share designs
 - Some (well document and supported) open S/W and H/W solutions are already available
 - Share information (e.g. via the xTCA interest group)
 - Avoid using “exotic” technologies
 - E.g.: Prefer PCIe for AMC communication over SRIO
 - Consider contributing (more) to the standardization process

CPAD SLAC RCE (ATCA) Development

- 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
- LCLS-1 LLRF upgrade
- LCLS-2 BPM, MPS and timing system
- SSRL RF booster upgrade
- CMB
- TES (transition edge sensor) photon detector

- Continue development of RCE platform
- - Improved SDK and support tools
 - Research available real time options for Linux
 - 40G upgrade for COB
 - Ultrascale Zynq upgrade
- • Continued R&D into RCE based applications
 - Low latency trigger architectures and algorithms
 - Online data reduction with software/firmware
 - High density flash storage coupled to RCE processors (LSST)
- • R&D into Gigachip (GCI) memory systems
 - Multiple processors/FPGAs accessing central shared memory
- • Demonstrate low noise analog designs in ATCA
 - Linear Coherent Light Source (LCLS)-II controls platform
 - High density photon readout with tight integration to RCE platform
- • Continue R&D into FPGA based RF platforms
 - ATCA based carrier board at the core
 - Build IP library for RF & DSP cores
 - Improved design flow and revision tracking
 - LCLS-1 low level RF upgrade
 - TES based photon detectors with RF
- • Improved hardware abstraction
 - “Device Tree” like description file to define the hardware registers and overall structure

- Direct Involvement: Trigger-less Time-Slice systems
 - MKIDs – 10K or 20K pixel instrument at SOAR in 2016; possible use in CMB; Phase 2 100K pixel or larger system
 - CONNIE/DAMIC – DAQ for 1kg detectors
 - Optical Links – system specification and testing for Versatile Link+
 - Rad-Hard Sensors – continued testing of candidate detectors
 - CAPTAN+X – deployment; reusable firmware blocks
 - Test beam Detectors and DAQ – ready-to-use system
- Indirect and Support Activities:
 - Off-the-Shelf DAQ – work w/test beam, experiment and university users
 - *artdaq*: DAQ Online Software Framework – continue to partner with experiments; enhance functionality with Real-Time Innovations, Inc. (SBIR) Data Distribution Service (DDS) messaging middleware; expand core functionality; software triggering.
 - Possible involvement to integrate DDS in slow controls (EPICS)
 - PREP – prototype NIM coincidence module; wider use of PREP model
 - Test stands – refine model; work with experiments and universities

- Design Activities: Focus on cold electronics
 - Front-end ASIC designs optimized for detector technologies
 - Front-end board designs tailored for experimental environment (radiation, cryogenic temperature etc.)
 - Trigger/DAQ designs explore the ultimate use of COTS solutions (high speed optical link, DSP in FPGA etc.)

- Expertise has been developed in R&D to provide an integrated solution of readout and Trigger/DAQ system for future experiments, examples:
 - Proposal of readout system for SBND in Short Baseline Neutrino program: short-term high priority project
 - R&D in FESOC of ATLAS Phase-II Upgrade and Full Integrated Cold ASIC for TPC readout: long-term R&D with high risk and high impact

CPAD TDAQ Comments

- GPU v. FPGA v. CPU – different capabilities and constraints – need to understand how each will be used.
- CMS and ATLAS rely greatly on FPGAs
 - Will performance scaling continue to follow processors?
 - What about the trends for FPGA to be “not” FPGA?
 - Hardened silicon purpose built blocks, processors communication links etc...
 - Will Intel's acquisition of Altera be a disruption in the FPGA community?
- Keep aware of developments to provide direct optical interconnects to FPGAs.
 - Potentially transformative
- Level-1 Tracking Triggers are essential for LHC Phase-2
 - Transformative and enabling!
- Use of precise timing information in the L1 Trigger should be investigated.
- High Speed Optical Links are a key enabling technology
 - HEP Leadership provided by European groups
- Need to understand Communications Architecture Issues
 - uTCA remains a work in progress in CMS
 - ATCA will most likely be used in CMS
 - ATLAS has adopted ATCA

- Encourage increased coordination and cooperation between National Labs and Universities in TDAQ.
 - Technologies, architectures....
- Encourage collaboration on High Level Synthesis Tools – establish a venue for collaboration on this and other FPGA tools (e.g. Linux on ZYNQ)
- Establish a forum within the US to facilitate the above
- Support R&D in FPGAs, GPUs, Associative Memories, Communication Architectures and Link Technologies