

EPICS Collaboration Meeting
September 20, 2016

ATCA Based Accelerator Controls & Detector Platform

Ryan Herbst
Department Head, Advanced Electronics Systems

(rherbst@slac.stanford.edu)



SLAC TID-AIR
Technology Innovation Directorate
Advanced Instrumentation for Research Division




Defining The Architecture

LCLS-II High Performance Controls Systems (HPS)

- Beam Position Monitor (BPM)
- Bunch Charge Monitor (BCM)
- Bunch Length Monitor (BLMN)
- Machine Protection System (MPS)
- Timing Delivery System

LCLS-I Controls Upgrade Systems

- Beam Position Monitor (BPM)
- Bunch Charge Monitor (BCM)
- Bunch Length Monitor (BLMN)
- Low Level RF (LLRF)



Other Systems & Experiments

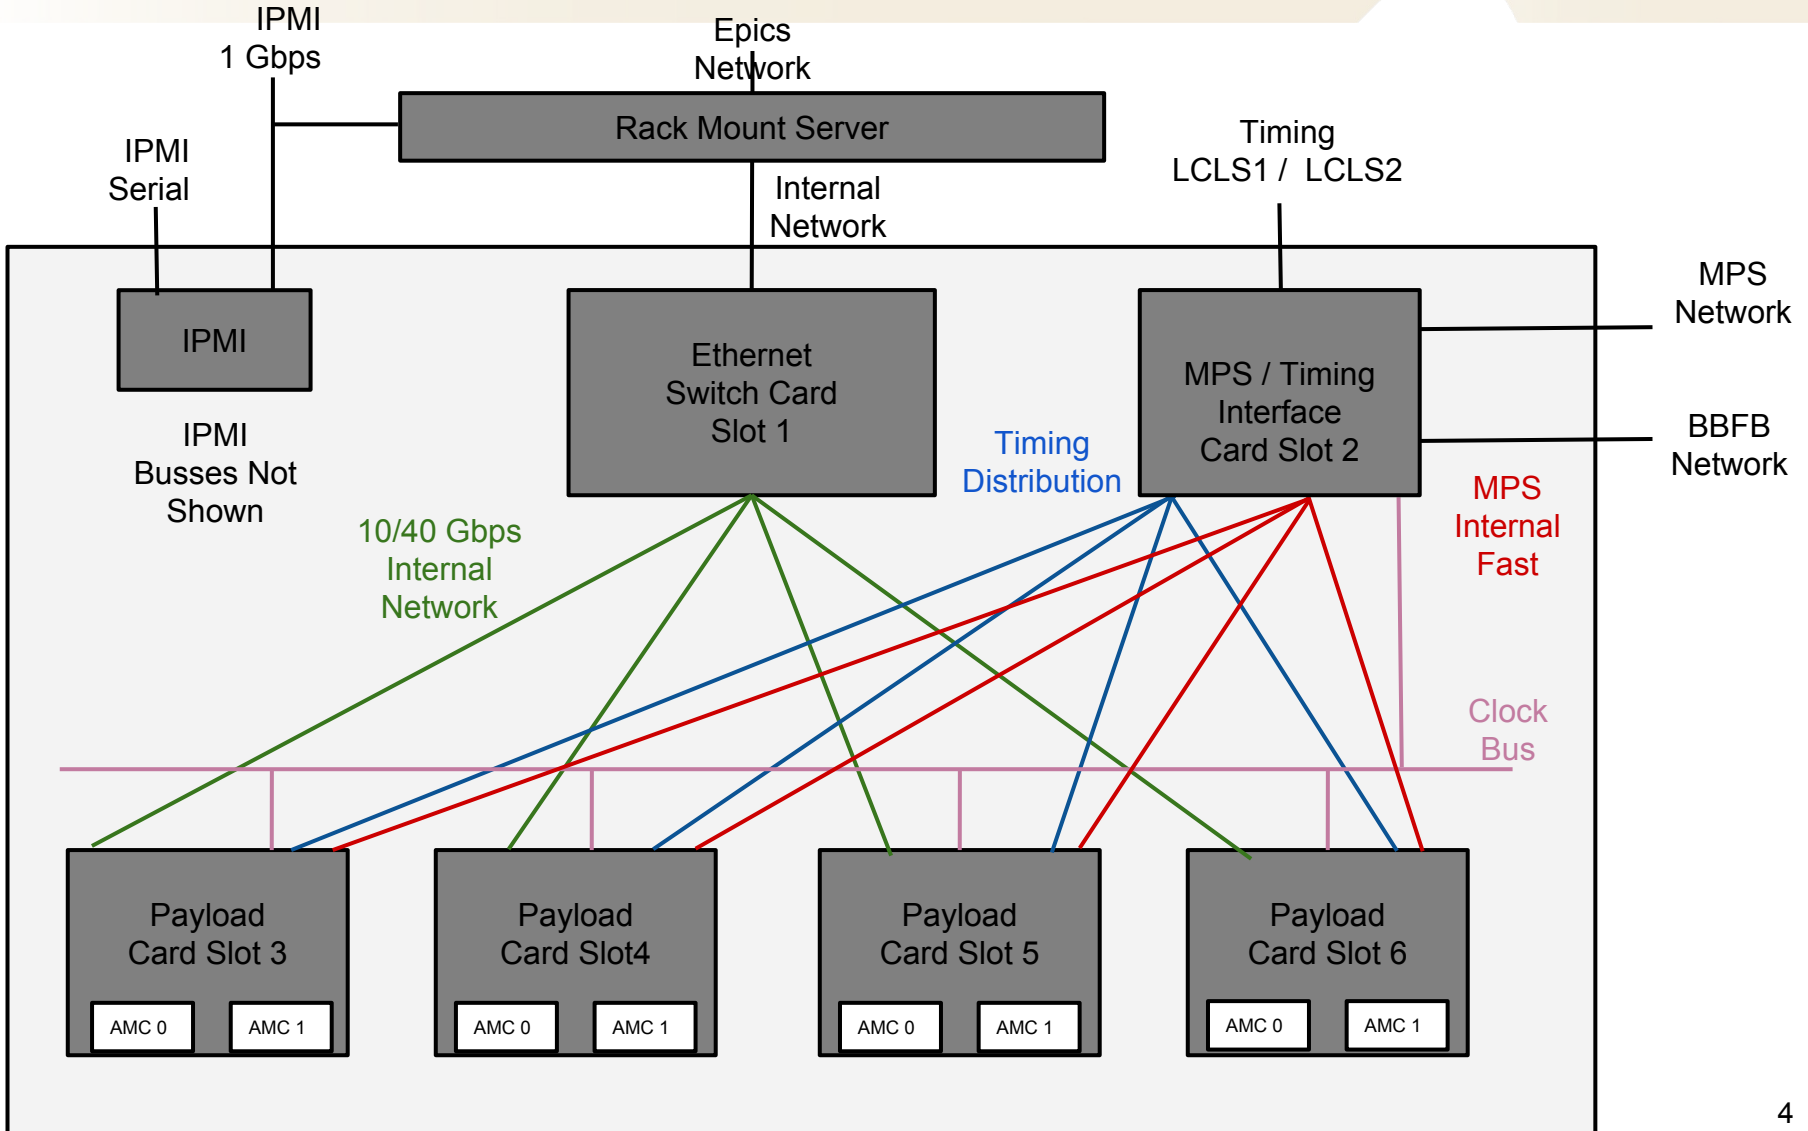
- SSRL Low Level RF Upgrade
- Transition Edge Sensor Detector for Cosmic Microwave Background experiments and X-ray spectrometers

- TID-AIR developed an ATCA based common platform for accelerator controls
 - LCLS-2, LCLS-1 and SSRL
 - Combines high performance digital processing with low noise RF design

LCLS-2 Common Platform Architecture

- The LCLS-2 common platform architecture is formed around a few key concepts
 - ATCA based packaging
 - ATCA = Advanced Telecommunications Computing Architecture
 - 10G Ethernet based interconnections for control, monitoring and inter-communication
 - A common digital board used for all applications
 - A rear transition module to allow for flexible higher level network access
 - An analog / RF application card
 - Maximize analog to digital separation within a crate to minimize cross-talk and interference
- In addition the architecture stresses the following design principles
 - Maximize re-use of hardware, firmware and software components
 - Minimize the total number of distinct hardware units created
 - Minimize cabling
 - Simplify hardware unit debug and replacement
 - Ensure future upgradeability of analog and digital components separately
 - Minimize overall cost
 - Including cost of deployment & long term maintenance
 - Optimize the density of overall system

HPS Interconnects & External Connections



Chassis Example

8 BPMs in one ATCA crate

BPM Carrier 4 – Slot 6

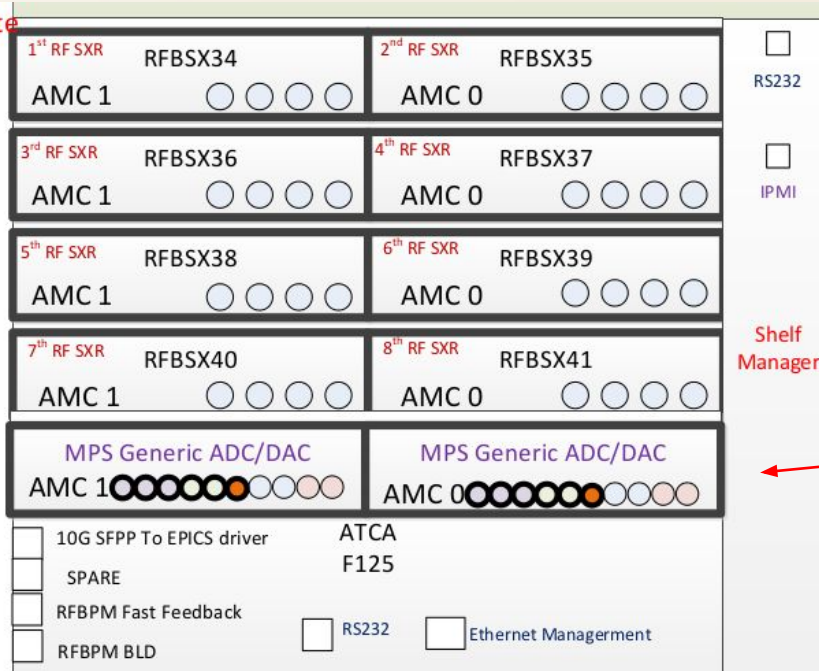
BPM Carrier 3 – Slot 5

BPM Carrier 2 – Slot 4

BPM Carrier 1 – Slot 3

MPS Link Node – Slot 2

10G Switch – Slot 1



Example deployment with 8 BPMs and 6 MPS beam loss monitors

MPS link node card with BLMs

ASIS 7 slot ATCA crate
Possible to support 10 BPMs plus link MPS link node in 6U


Low cost Single slot solution



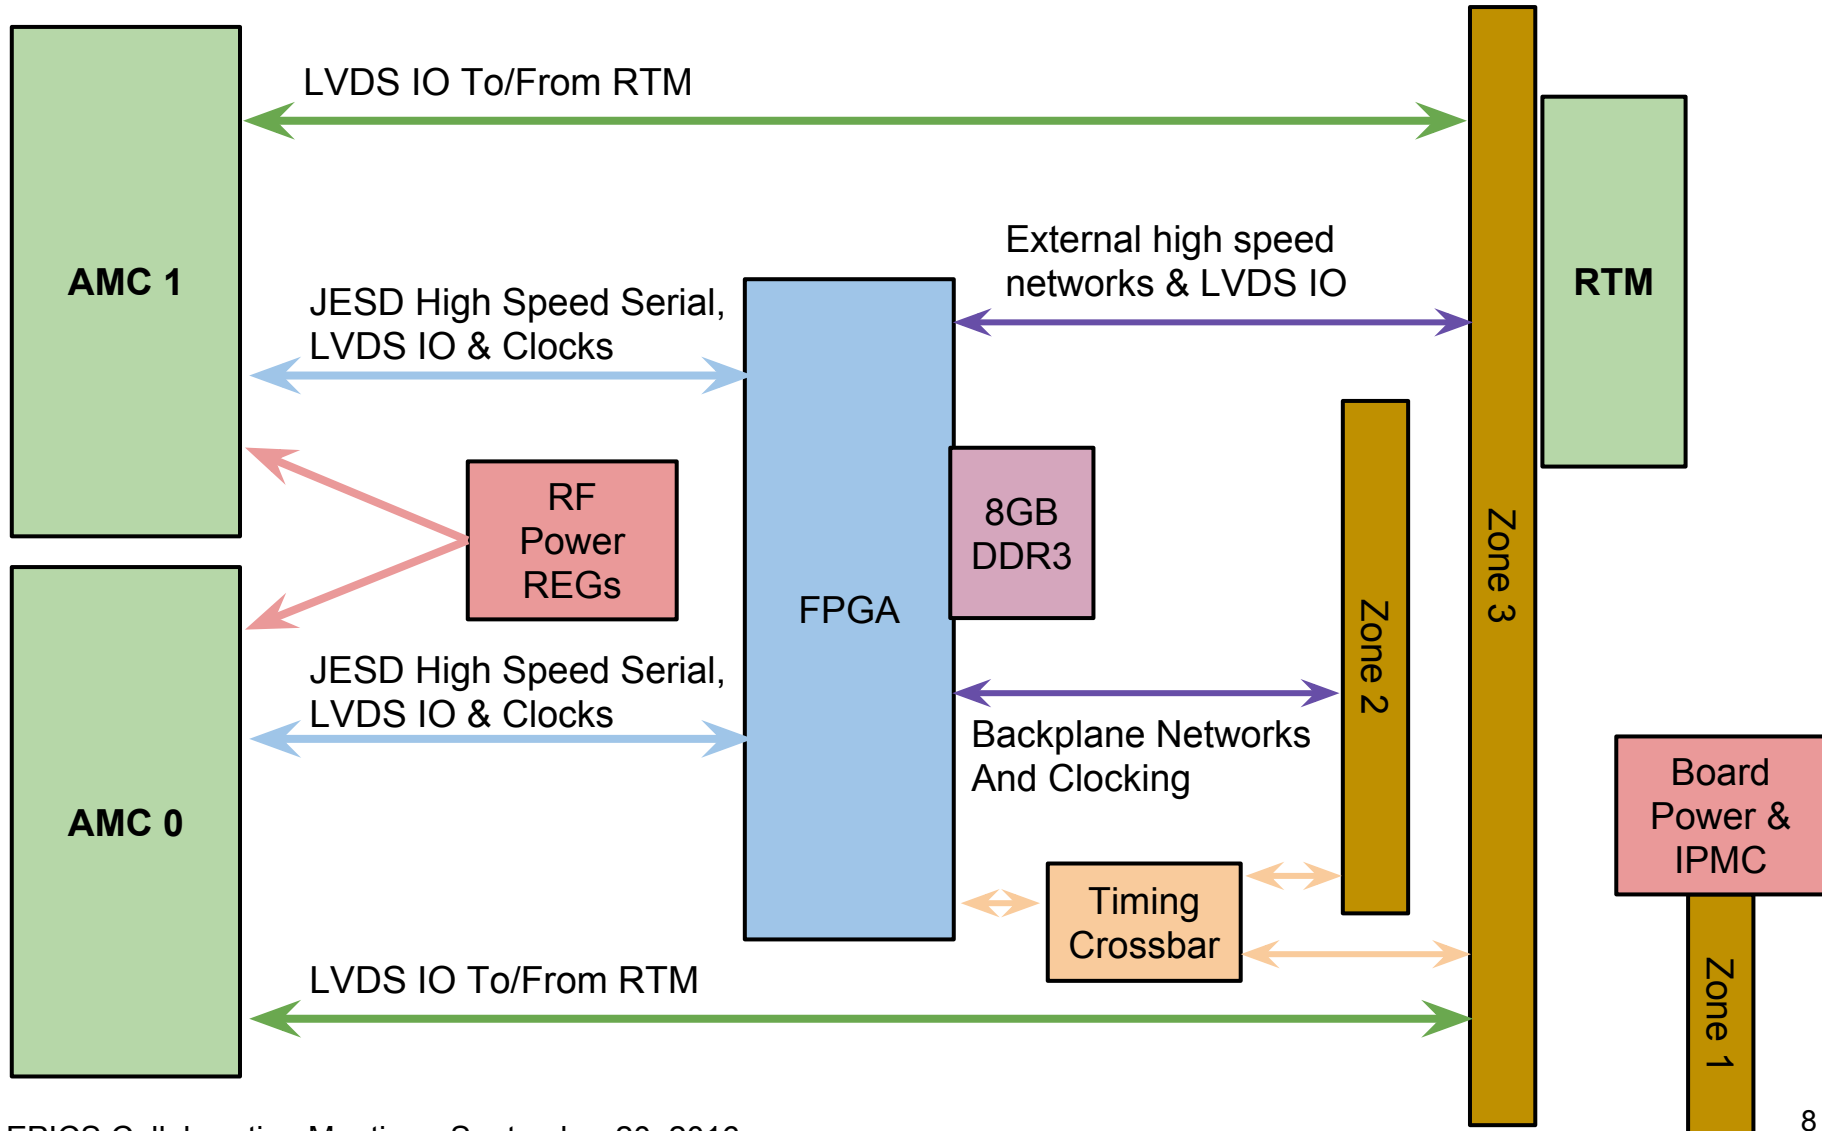
ATCA Benefits

- In ATCA each card operates as an independent unit
 - No need to power cycle or reboot the entire crate
 - The crate is always powered
 - Individual cards are powered on and off as needed
 - Application cards are hot swapped
- Multiple sub-systems can coexist without impacting each other during maintenance
 - IOCs can remain running while card is power cycled or firmware is updated
 - Register access will time out with warnings until hardware comes back online
 - CPU reboot is not required when cards are added or removed
- Crate maintenance can occur while the system is running
 - N+1 power supply redundancy allows power supply hot swap
 - IPMI redundancy allows shelf manager replacement during operation
 - Fan trays can be replaced while system remains running

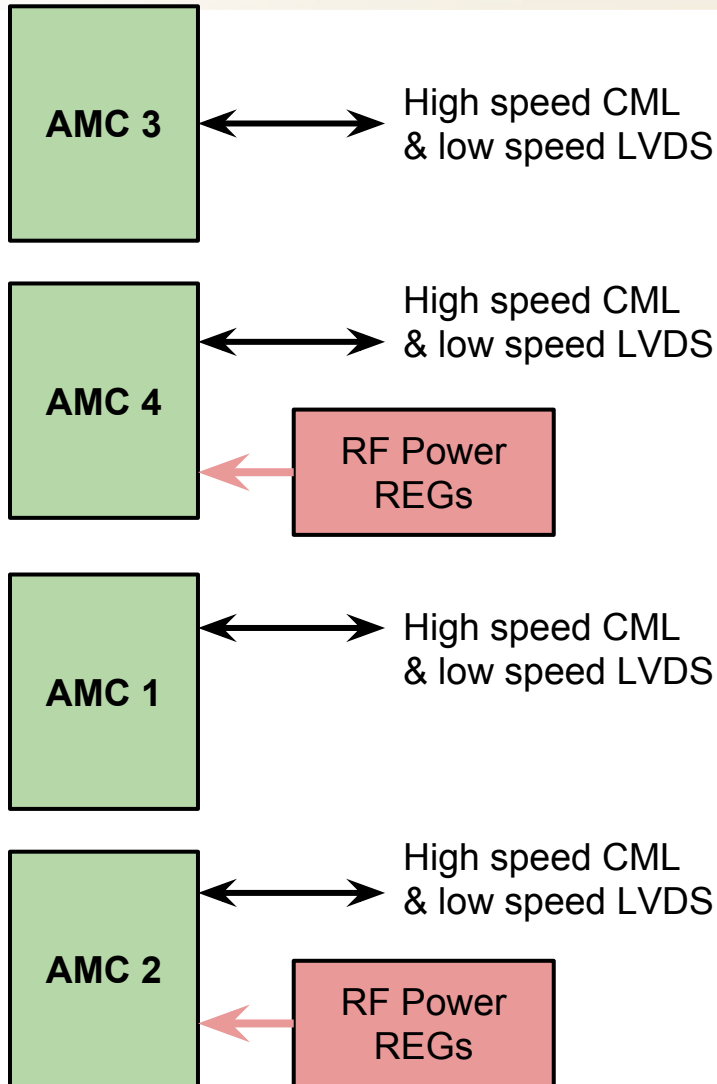
ATCA AMC Carrier Board



ATCA AMC Carrier Interconnects




AMC Connectors



- The AMC carrier card supports 4 standard AMC connectors
- AMC connectors 1 & 3 are fully AMC compliant
- AMC connectors 2 & 4 are mostly AMC compliant
 - Not present in commercial dual wide AMC cards
 - Some LVDS lines mapped to RF supplemental power regulators
 - Power can be turned off when a standard AMC card is inserted
- Default mechanical loading option is for two dual wide AMC cards
 - AMC connectors 2 & 4 not used on standard dual wide AMC cards
- High speed serial lines used as JESD204b lines to support high speed ADCs and DACs
 - 7 pairs per dual wide AMC card
 - Expandable to 10 with larger FPGA
- Connection mapping to FPGA ensures that a PCI-Express compliant AMC card can be supported
- IPMI wiring is fully compliant
 - AMC carrier IPMC needs software update to be fully AMC compliant

AMC Outline




Location of optional second AMC connector

Optional second connector provides access to additional JESD channels & supplemental power supplies
(makes AMC card non-compliant)


From AMC.0 R2.0 specification section 2.2.1.2

AMC Card Examples




LCSL-2 BPM board with secondary AMC connector.

Demo ADC/DAC/BPM/LLRF Board.




AMC Card Examples




Network board with 8 SFP+

MPS Digital Input Board





General purpose ADC/DAC board.
3 x ADC, 2x DAC + trigger I/O
370Mps



AMC Profile

ATCA Chassis




ATCA card is cut out to allow for full height AMC cards

Provides for additional component height

Provides separation between digital and analog sections

Standard configuration for full height AMC card support as opposed to mid height (uTCA)

Rear Transition Module (RTM)




MPS digital inputs
Slow analog inputs
Digital interlock outputs

8 SFP+ modules
2 for timing
6 for high speed networks

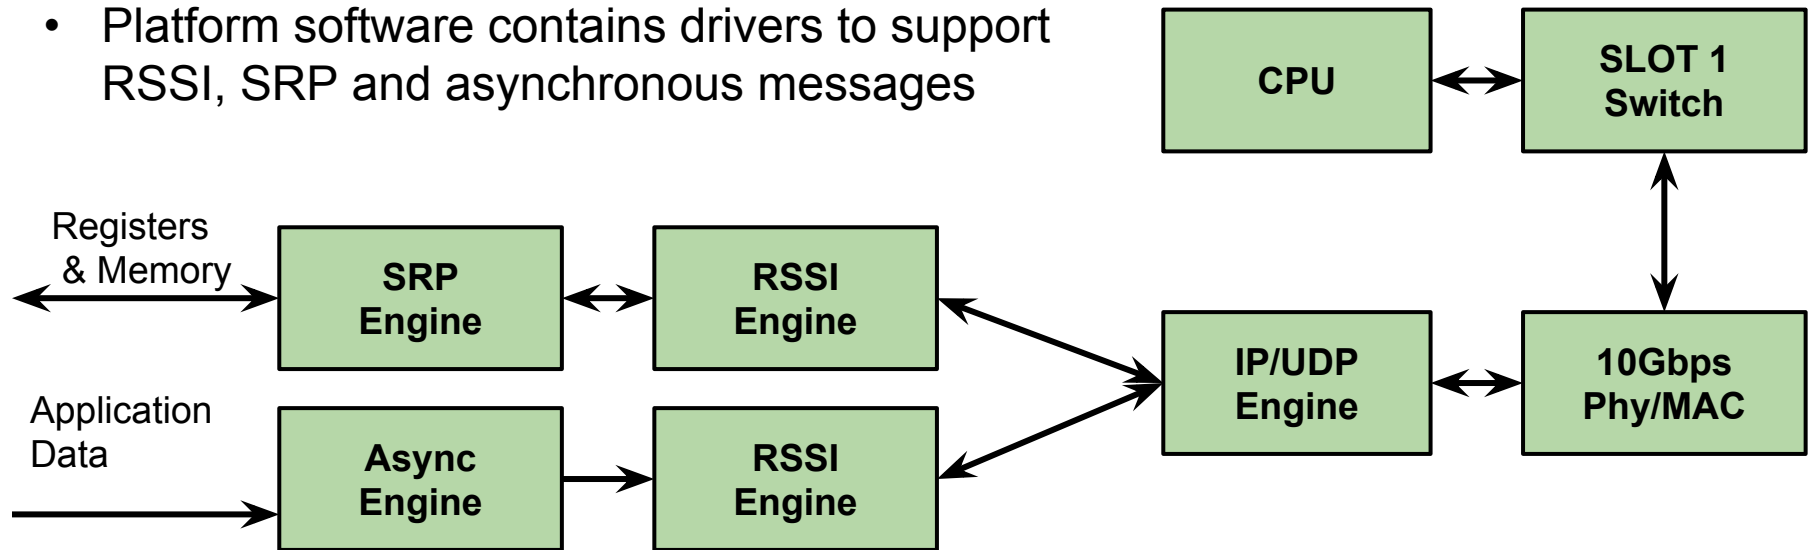
Inter FPGA Communication: 10G Ethernet

- 10G Ethernet is used for flexible communications between application cards utilizing the central switch card
- Communications use UDP protocol for end point and channel identification
 - Each FPGA supports multiple UDP ports to distinguish traffic
- Higher level protocol required for reliable message delivery over UDP: RSSI
 - Reliable SLAC streaming interface, based upon RUDP protocol
 - RFC-908, RFC-1151, draft-ietf-sigtran-reliable-udp-00
 - Additional features added to support flow control
 - Facilitates breakup of large transfers into MTU sized messages

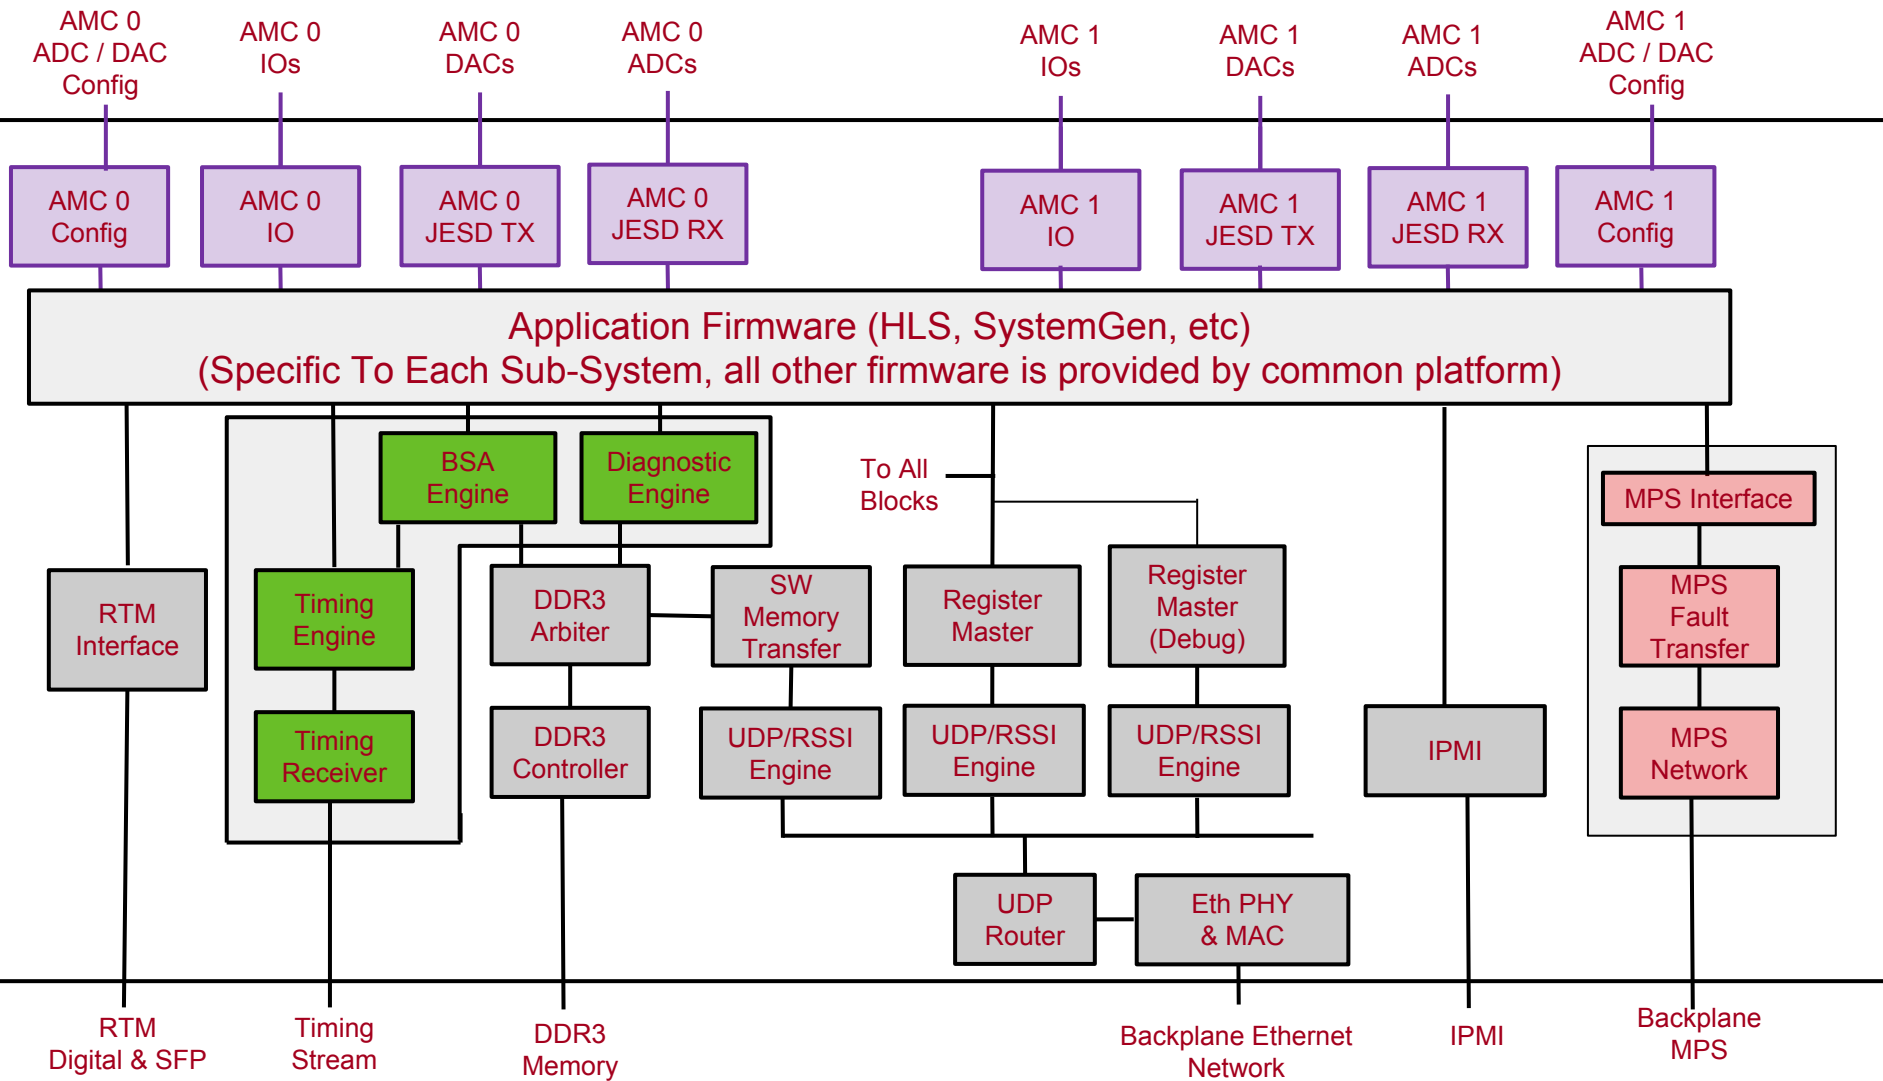


SW To FPGA Communication: 10G Ethernet

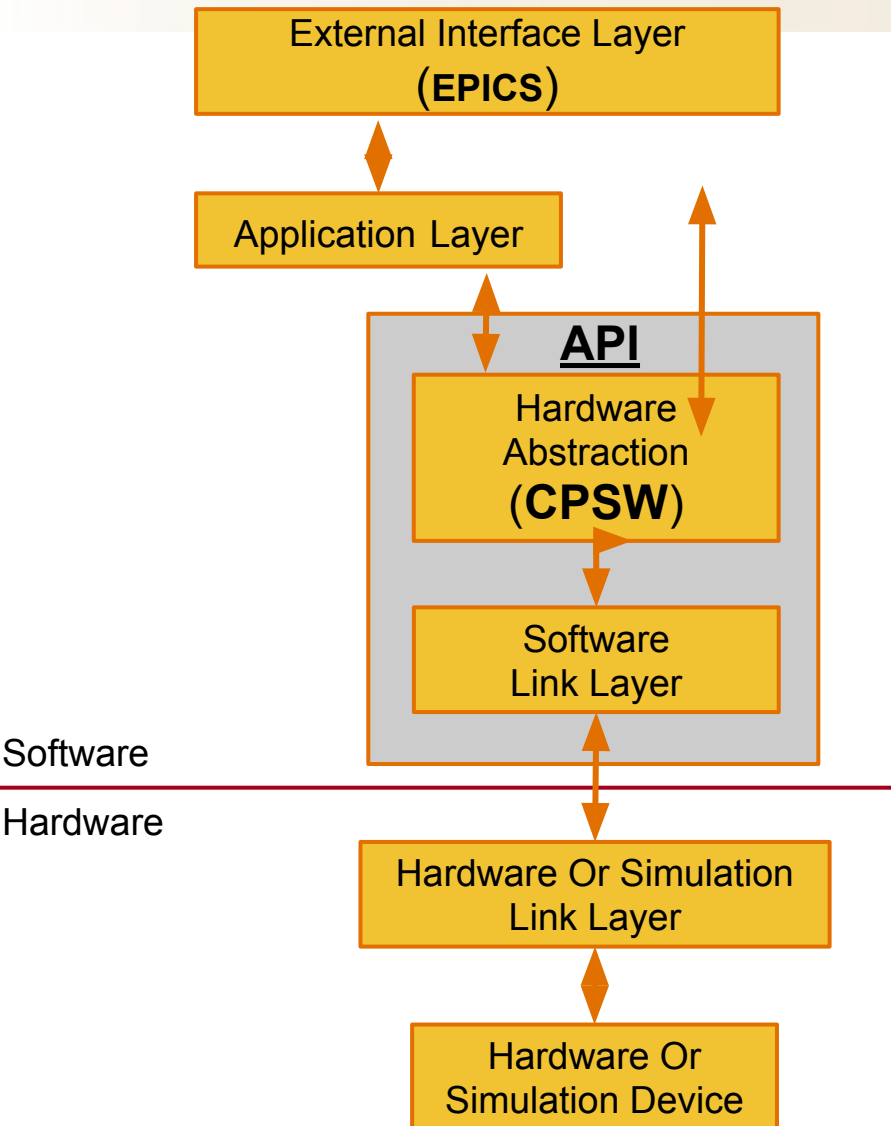
- A CPU will manage one or more HPS crates via direct connection to crate 10Gbps switch card
- Register access protocol to facilitate both individual register access as well as bulk memory access: SRP (SLAC Register Protocol)
 - Register access used to configure and monitor FPGA
 - Memory access used to read BSA and diagnostic buffers
- Asynchronous message support to move raw data from hardware to software
 - Interrupt messages
 - Streaming diagnostic data
- Platform software contains drivers to support RSSI, SRP and asynchronous messages



Common Platform Firmware

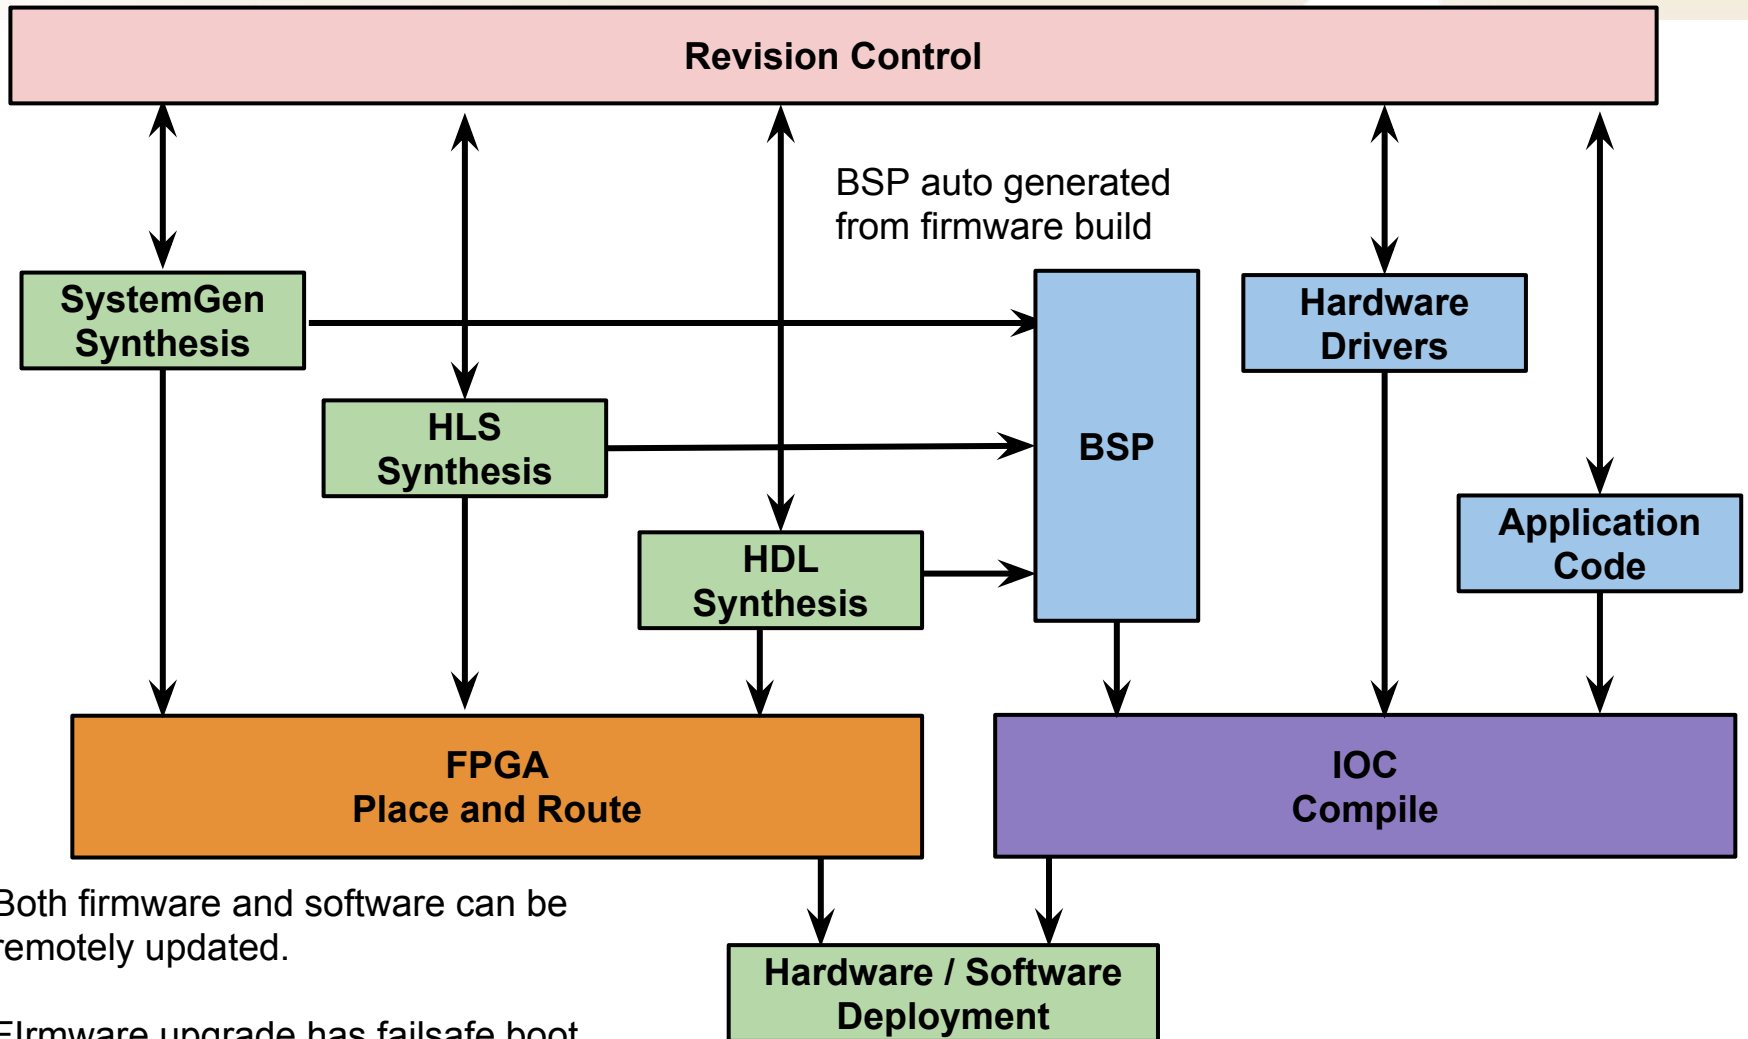


HPS Common Platform Software



- All common layers of software and firmware are developed by the common platform sub-system
 - Encourages re-use and minimizes parallel development
- For application development a suite of commonly libraries and application modules are available as well
 - CPSW (Common Platform SW)
 - i.e. Generic ADC/DAC driver
 - Firmware FIFOs and memories
- All firmware and software is developed with re-use and portability as a goal
- Past projects developments with this design approach have helped in accelerating the LCLS-2 HPS development

Firmware And Software Builds



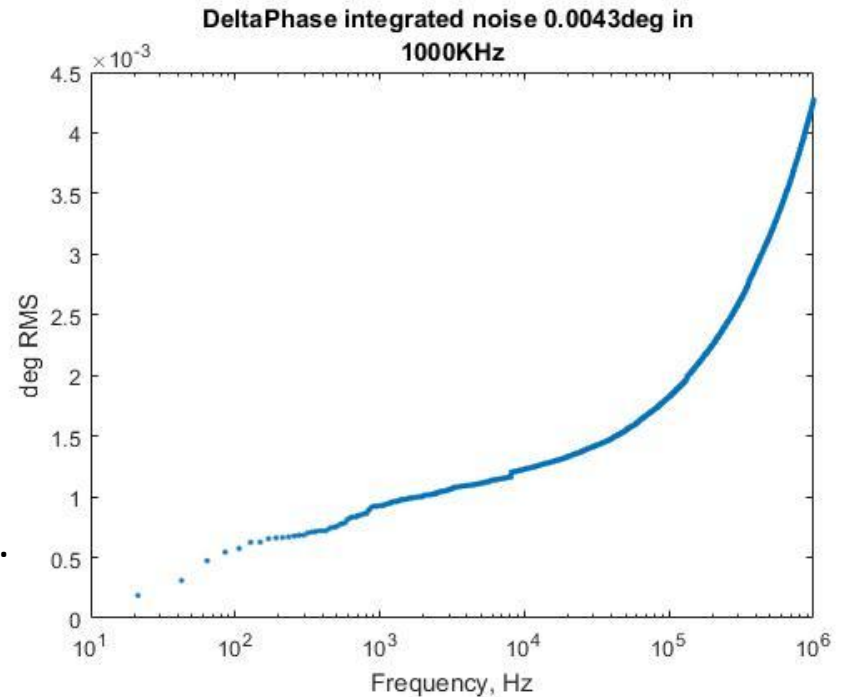
Both firmware and software can be remotely updated.

Firmware upgrade has failsafe boot loader image.

LLRF System performance is very good!

Test Conditions:

- Laboratory environment
- Tested in ATCA crate in normal operating mode
- External 2856 Reference (LCLS spare)
- External 17MHz clock simulates EVR reference
- 2856 split from reference to input channel
- On-board LO and Clock oscillators locked by FPGA to reference and ext. clock
- Raw data transferred to Matlab for processing.



Measured noise 0.0043° much better than 0.01° requirement.
Corresponds to -143dBc/Hz in 1MHz BW

Signal generation noise of 0.019° noise does not quite meet 0.01 degree spec, but this function has just begun testing, expect significant improvements

[Performance details from Dan Van Winkle](#)



The End

Thank You!

Backup Slides

Platform Selection: ATCA vs uTCA

- ATCA and uTCA both have many variants
 - We are comparing SLAC's common platform ATCA and uTCA as used in DESY
 - uTCA assumes PCIe backplane
 - ATCA assumes Ethernet backplane
 - ATCA assumes AMC carrier with JESD interface to FPGA
- ATCA is a mature standard with industry acceptance and support
 - Many sources for crates, power supplies, switch cards and shelf managers
 - Strong competition lowers cost and increases availability
- uTCA.4 is a new standard in a niche market
 - uTCA is a later offshoot standard from ATCA
 - uTCA.4 is a variant on uTCA
 - Many leading xTCA vendors have abandoned both uTCA and uTCA.4
 - Risk for longer term availability after DESY construction bubble dissipates
- Interconnects have consequences to the system architecture
 - PCI-Express in uTCA tightly couples the processor to the payload card
 - Geographic limitations
 - Processor hardware must be informed when cards come and go
 - Ethernet in ATCA decouples the processor and the payload card
 - Geographically unlimited
 - Message based communication
 - NAD like architecture with the advantage of clean mechanical packaging

Platform Selection: ATCA vs uTCA

- Timing delivery is important
 - Encoded timing streams with bit rates ~4Gbps are common in modern systems
 - In uTCA the MCH serves as both the switch and the timing hub
 - Typically you want a commercial switch but a custom timing hub
 - Commercial MCH designs did not support the bit rates required
 - Would require a custom MCH, but PCI-Express vendors have difficult NDA requirements
- Crate based processors are not cost effective
 - Chassis vendors charge a premium for processor blades
 - uTCA requires the processor to be in the crate
 - ATCA allows an external cost effective processor
 - Can be the same model as used in non-crate based portions of the control system
 - Crate based processor is still an option, but not a necessity
- uTCA.4 has limited power capacity
 - AMC connector and RTM use signal pins for power
 - 12V based distribution, vs redundant 48V in ATCA
- ATCA crates support DC power input
 - Useful in very low noise applications
- ATCA main board size allows more flexibility in analog/digital divide
 - uTCA.4 assumes 50/50 space allocation, forcing ADCs onto digital board
 - SLAC's AMC carrier design allows the ADCs and DACs to be closely matched to analog section
 - Small digital section allows independent analog and digital upgrade paths
 - Allows independent analog and digital development
 - ATCA AMC provides better shielding and more space


Platform Selection: ATCA vs uTCA: Our Experience

- Experience with uTCA at SLAC
 - A version of the BPM design was implemented in uTCA
 - The limited height of uTCA did not allow the proper filters to be installed
 - Locating the ADCs on the front card limited designer's choice
 - uTCA does not appear to have a roadmap for PCI-Express 3.0
 - The AMC specification does not yet include PCI-Express 3.0
 - Backplanes have been built but do not meet the PCI-Express length limitations
 - In house testing shows not all payload slots reliably sync at 3.0
 - The RTM connector specified in the uTCA.4 standard is limiting
 - RTM side connector does not support 12Gbps
 - RTM connector is not the highest density available in the family
 - RTM does not have dedicated power pins, instead use signal as power
 - The AMC IPMI standard is still new causing interoperability issues
 - Some cards do not interoperate properly
 - Will be fixed over time
- Experience with ATCA at SLAC
 - TID-AIR has successfully deployed ATCA based designs in multiple experiments
 - All designs tested with fully loaded 14-slot crates
 - In house IPMI expertise
 - Demonstrated successful interoperability with commercial cards
 - Demonstrated low latency high bandwidth inter-card communication
 - TID-AIR has prototyped analog AMC carrier design
 - Demonstrated low noise operation of various analog front end types
 - BPM design fits well in AMC daughter card with proper filters in place
 - LCLS-II ready hardware

- ATCA = Advanced Telecommunications Computing Architecture
 - PICMIG 3.x
- Front board is 280 mm deep by 322 mm high by 30.48 mm on center
 - Lots of usable board space with plenty of room left after placing power supplies and management
- Rear Transition Module (RTM) is 322mm high by 70mm deep
 - Perfect size for network connections
- Backplane architecture
 - Dual star 1Gbps Ethernet for management
 - Generic dual-star or full mesh data inter-connect, 4 bidirectional differential lanes
 - 10Gbps & 40Gbps Ethernet (100Gbps upgrade path)
 - Custom interconnect
 - MLVDS timing bus, consisting of 6 differential pairs
 - Redundant I2C management bus
- Centralized shelf manager with hot swap support
 - IPMI
- Well established multi billion dollar market
- Stable shelf management specification with strong interoperability
 - Commercial IPMI firmware/software available
- SLAC has 15 years of ATCA design experience

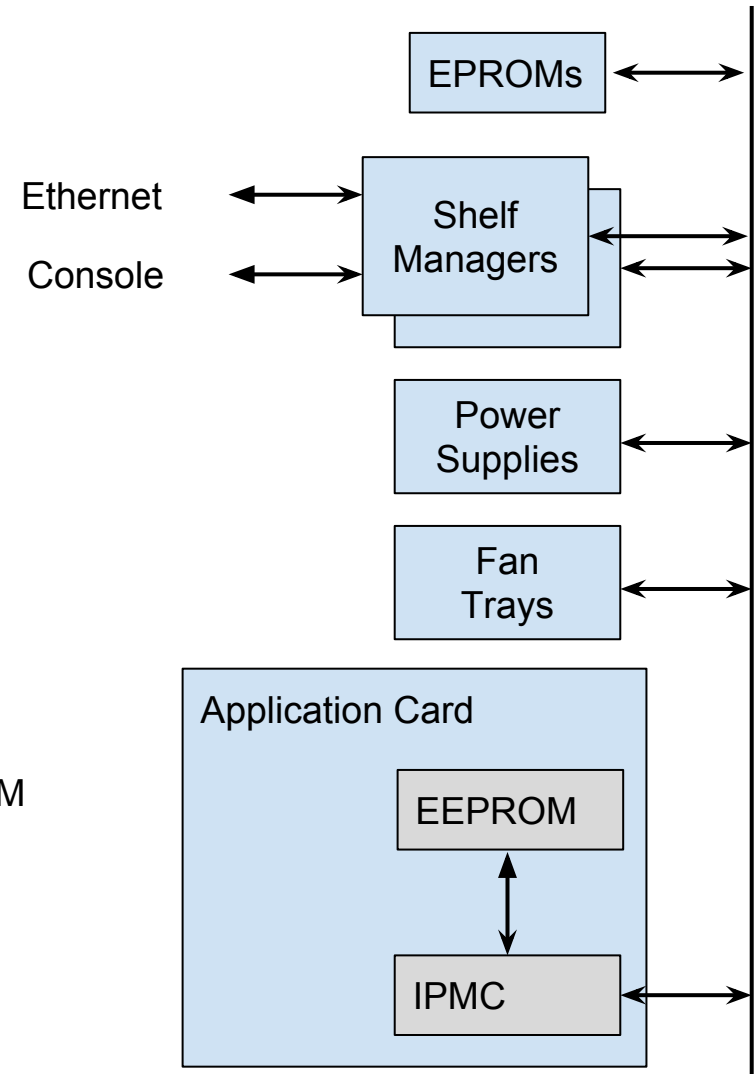


ATCA Shelf Varieties



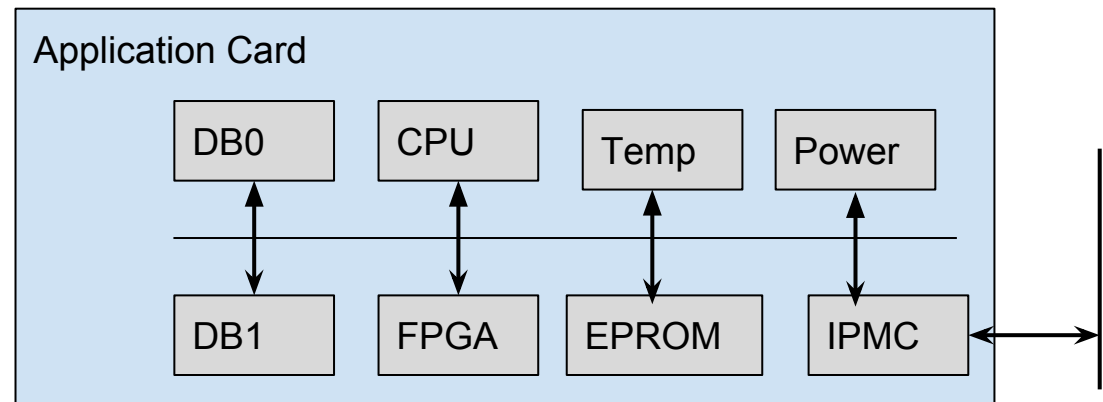
ATCA Shelf Management & IPMI

- ATCA uses IPMI for management purposes
 - Intelligent Platform Management Interface
- Manages and monitors all shelf based components
 - Power supply status and power
 - Shelf inlet and exit temperatures
 - Fan speed control and monitoring
 - Application card control and monitoring
- Redundant EEPROMs contain all shelf information
 - Shelf serial number, location and ID
 - Shelf manager IP/MAC address
- Application card hosts IPMC
 - Intelligent Platform Management Controller
- IPMC hosts all application card information in local EEPROM
 - MAC addresses
 - Serial number, card type & revision




ATCA Shelf Management & IPMI

- IPMC performs the following functions on the application card
 - Relays shelf and card information to local CPU and/or FPGA
 - Enables local power and resets
 - Monitors temperature of board and components
 - Can be configured to power down board when thresholds are exceeded
 - Identified daughter boards and verifies their compatibility with main board
 - Verifies application card is compatible with other application cards
 - E-keying
 - Turns off backplane interfaces to avoid damage
- IPMI handles all of the monitoring and self protection required by a deployed design



ATCA Components

- Almost all components can be replaced in the field
- Redundancy is available if desired
 - N + 1 redundancy for power supplies
 - Redundant shelf managers
- System is designed to handle one fan failure in each fan tray
 - Shelf manager generates alarm to request fan tray replacement




ATCA Components - Backplanes



Zone 1:
Power &
IPMI


Zone 3:
RTM Access

Zone 2:
Base &
Fabric stars



Standard ATCA Backplane

6 Slot Example



Base Interface Details

- Base interface provides connectivity between the two switch cards and the payload cards
- Each switch card slot has connectivity to both shelf managers (redundancy)
- Back plane base traces are treated like Cat6
- ATCA specification does not allow for custom use of the base interface
- Not appropriate for timing distribution or MPS feedback use
- The base interface will not be utilized by the HPS base ATCA carrier board

Base Interface Interconnections

Connector	Base Ch.	Logical Slot					
		1	2	3	4	5	6
P23	1	ShMC	ShMC	1-3	1-4	1-5	1-6
P23	2	2-2	1-2	2-3	2-4	2-5	2-6
P23	3	3-1	3-2				
P23	4	4-1	4-2				
P23	5	5-1	5-2				
P23	6	6-1	6-2				

Fabric Interface Details


- Table on the right shows the dual star fabric connections for a 16-slot backplane
- The SLAC LCLS2 ATCA system will make use of the standard dual star backplane
- Star 1 will be mastered by an Ethernet switch card in slot 1, providing Ethernet to payload slots 2-16
- Star 2 will be mastered by the LCLS-II carrier card in slot 2, providing timing distribution to slots 3-16.
- The ATCA card in slot 2 will also accept MPS network signals from slots 3-16 over star 2
- The ATCA specification allows for custom use of the fabric channels
- None of the ATCA standard specifications are violated in the SLAC ATCA for LCLS-II solution

Table 6-12 Dual Star Backplane routing assignments

Logical Slot #		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Connector	Fabric Channel #																
P20	15	16-1	16-2														
P20	14	15-1	15-2														
P20	13	14-1	14-2														
P21	12	13-1	13-2														
P21	11	12-1	12-2														
P21	10	11-1	11-2														
P21	9	10-1	10-2														
P21	8	9-1	9-2														
P22	7	8-1	8-2														
P22	6	7-1	7-2														
P22	5	6-1	6-2														
P22	4	5-1	5-2														
P22	3	4-1	4-2														
P23	2	3-1	3-2	2-2	2-3	2-4	2-5	2-6	2-7	2-8	2-9	2-10	2-11	2-12	2-13	2-14	2-15
P23	1	2-1	1-1	1-2	1-3	1-4	1-5	1-6	1-7	1-8	1-9	1-10	1-11	1-12	1-13	1-14	1-15

MPS Messages

- PGP = Pretty Good Protocol
- Architecture independent – can be deployed on any 8B/10B SERDES
 - Copper or optical
 - Speed defined by SERDES and internal data path
- Unlimited frame size
- 4 virtual channels, each with separate firmware interface
 - 4 frames in flight at any given time
 - Avoid head of line blocking for configuration messages
- Built in flow control support
- Cell based protocol
 - Large frames segmented into 512byte cells for transport
 - Guaranteed cell ordering
 - CRC protected cells, errors forwarded to end point
 - Per VC arbitration at the cell level
- Low overhead (96% efficient after 8B/10B conversion)
- Can be used to transport timing and trigger
 - Low latency, deterministic trigger transport interface
- Unidirectional & bidirectional protocol
 - Example: 1 downstream link and 4 upstream links
 - Upstream and downstream links can be different line rates
- Supports lane bonding for wider data path
- Standard interface for many experiments
 - LCLS, SID, LSST, ATLAS & Others



SLAC PGPCard
4 lane x 3.125Gbps
Full duplex 9.6Gbps

RSSI Layer (Reliable SSI)

- file:///afs/slac/g/reseng/svn/repos/FirmwareCoreLibrary/trunk/RssiCore
- Reliable communications layer based upon RUDP (Cisco implementation)
 - RFC-908, RFC-1151, draft-ietf-sigtran-reliable-udp-00
 - AxiStream (SSI) based, agnostic to transport layer
 - Works with PGP, UDP, USB, etc
 - May consider modification to unbalance stream directions
 - Each direction of the stream may have differing requirements
 - Max segment size
 - Max outstanding segments
 - May just be a minor change in the handshaking
 - DMA interface is an example
- Parameters configured via generics
 - Max segment size
 - Max outstanding segments
 - Timers, etc
- Configurable client/server modes via generics
- Internal width fixed (wider is better)
 - Interface widths configured via AxiStream configuration
- Application side TREADY de-asserted if client/server connection is not established
- Provides in order segment transmission
- EOF is forwarded

