“CMS News”

M. Hansen, CERN

J. Hegeman, CERN; E. Hazen, BU; T. Gorski, UW

CMS application of VME

- **Semi-custom VME crates from WIENER**
  - Common backplane except ECAL (HP J0)
- **Crate controllers from CAEN**
  - 6U and 9U VME cards
  - PCIexpress controller card serving up to four branches
    - ➔ Formerly PCI controller card serving one branch
  - Register access support by HAL
- **System monitoring and archive of read values through cDCS**
- **No CMS integration support**
  - DAQ, TTC, TTS, etc.

CMS application of microTCA

- Industry standard shelfs with required backplane options
  - Redundant telecom backplane with AMC port 2 and 3 routed to MCH1 and MCH2 respectively
- Industry standard system controller (MCH) from two recommended vendors
- Industry standard power modules from two recommended vendors
- Industry standard bulk power supply from one recommended vendor
- CMS integration support through “AMC13” Located in redundant MCH slot
  - TTC and TTS, DAQ fan-in
- Register access through Ethernet
  - Ethernet switch in MCH
- System monitoring and archive of read values through cDCS
  - Although not entirely implemented yet

CMS application of microTCA
Areas possible to improve

- **System power at its limit**
  - Larger more powerful FPGAs difficult to integrate
- **Cooling capacity at its limit**
  - Boards essentially covered with a heat sink
- **DAQ bandwidth insufficient for some applications**
  - E.g. Pixel readout per AMC; not a technical problem but has lead to endless discussions
- **Slow Control bandwidth is sufficient**
  - Although shared with e.g. local DAQ
- **Several different solutions and implementations of e.g. IP address assignment even though the problem to solve is identical for all systems**
  - Again, the subject lead to endless discussions that ended in divergence

CMS application of microTCA
Areas possible to improve

- **System power at its limit – *ATCA?***
  - Larger more powerful FPGAs difficult to integrate
- **Cooling capacity at its limit – *ATCA?***
  - Boards essentially covered with a heat sink
- **DAQ bandwidth insufficient for some applications – *ATCA?***
  - E.g. Pixel readout per AMC; not a technical problem but has lead to endless discussions
- **Slow Control bandwidth is sufficient**
  - Although shared with e.g. local DAQ
- **Several different solutions and implementations of e.g. IP address assignment even though the problem to solve is identical for all systems – *Common Approach!***
  - Again, the subject lead to endless discussions that ended in divergence

ATCA

- **Size advantage over microTCA**
  - **not as large as one would think**
    - **➔ 50% more area, 100% more front panel**
- **Power and associated Cooling advantage over microTCA**
  - **400% more available power**
- **CMS ATCA back end blade: CBE blade**
  - **May require tuning for larger projects for link count / FPGA size**
- **CMS integration ATCA switch module: "Blade13"**
  - **TTC++, TTS++, DAQ interface**
- **Common IPMC**
  - **E.g. design supported by CERN PH-ESE; including hardware module; scheduled early 2016.**

Possible generic CMS ATCA shelf

continuing the microTCA direction

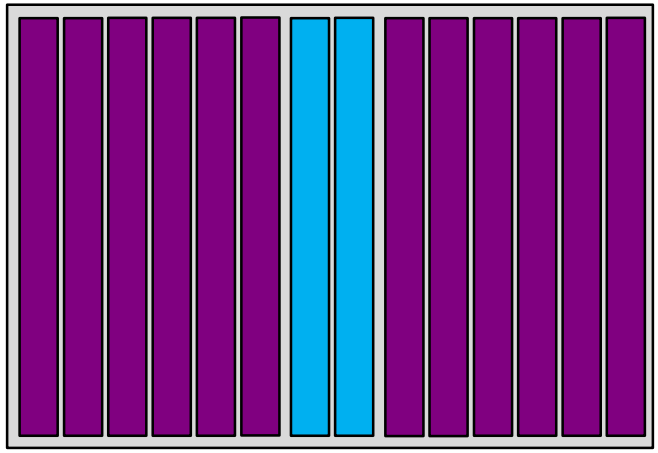

- **Common shelf specification**
- **Generic Blade with, perhaps, custom FPGA and link count but *with standard services***
- **Hub card with CMS interfaces: TCDS++ and DAQ *with standard services***
  - **Allowing for 800 Gbps DAQ per shelf or more without bending standards**

Standard Services wish-list

- **IPMC**
  - ***Well supported*** **open firmware / software as an FRU for forward compatibility**
- **Ethernet End point and Register access**
  - **Either *well supported* e.g. IPbus or equally *well supported* Embedded e.g. linux TBD as an FRU for forward compatibility**
    - ➔ **FRU size to fit a single width AMC card**
    - ➔ **Several commercial SOMs available that meet size requirement**
    - ➔ **If more processing power required – e.g. COM Express mini with an ATOM CPU – not applicable to AMC**
  - **Firmware upload / upgrade mechanism**
- **Power bank**
  - **Predefined main voltages with support for monitoring and customisation**

ATCA in CMS

- *Suggestions* for ATCA applications in CMS:
  - Single shelf connection for TCDS
  - Option for single-point DAQ if bandwidth is enough
    - Could be either COTS switch blade or custom “Blade13”
- Some other ideas:
  - Suggest special use for slot 2 for timing (details next)
  - Plan for at least standard GbE base switch in slot 1 (could be a fancy 40GbE switch blade if desired...)

ATCA Backplane connectivity

- Several clocks sourced from hub slots, all bussed to node slots
  - CLK1
    - ➔ 8 KHz fixed
  - CLK2
    - ➔ 19.44 MHz fixed
  - CLK3
    - ➔ Up to 100 MHz user specified
- As these clocks are bussed and not p-p they may or may not be suitable for any high precision clock distribution, especially since the MLVDS drivers have proven to be sensitive to temperature
- I have omitted any redundancy scheme in order to simplify this talk – in addition, do we need redundancy on this level?

There is already a document

https://www.picmg.org/wp-content/uploads/PDG_0-R1_0-RELEASED-2013-04-231.pdf

Physics Design Guide
for
Clocks, Gates & Triggers
in
Instrumentation

PDG.0 R1.0

23 April 2013

PICMG®

This is an excellent reference

you should read it!

ATCA Timing Options

Here is what I took away from the document....

- Option 0 – PICMG 3.0 "Synchronization Clocks"
  - Bussed, long delays, impedance control not so good
  - Use M-LVDS levels.
    - Suggest we stay away from these
- Option 1 – Base interface (P23 rows 5 and 6)
  - 4 pairs each from slots 1 and 2 to each blade, star-connected
  - Normally used for Ethernet but could be hijacked
  - Assigned as two Tx and two Rx pairs but this is optional
    - This would be incompatible with a COTS switch, but if we do this only for slot 2, why not?
- Option 2 – Fabric interface
  - 8 pairs from each slot to all other slots (full mesh)
    and from slots 1 and 2 (dual-star)
  - For DAQ it is unlikely we need all 8 pairs
    Tx/Rx directions are just suggestions...
    Could use e.g. 2 pairs downstream for TTC and clock and 1 upstream for TTS
  - This would leave 5 pairs for DAQ (4 at 10Gb upstream and 1 for handshake)
    - The clock source could be in any slot of a full-mesh shelf
- These are just ideas

DAQ

- Suggest to use fabric (not base) lanes
- Essentially we can do whatever we want with 8 pairs each to slot 1, 2 (or even another slot)
- If we want to support Ethernet, we need to maintain the Tx/Rx pair assignment in the standard
- Can in principle send 40Gb/s *per slot* to a hub.
  - In the (near) future this could go to DAQ on 100Gb links
- Could also include front or rear DAQ link in blade
- *Suggestion:* decide fairly soon on a preliminary DAQ interface which works conveniently both over fiber and backplane fabric. We can change later.

ATCA Backplane connectivity

- **Base interface: Dual star**
  - **Four pairs**
    - ➔ carrying 10/100/1000base-T (!) Ethernet or two 100base-TX ethernet
      - ⇘ as opposed to our microTCA 1000 base-X Ethernet
- **Common backplane topology: Dual star**
  - **Four bidirectional pairs from each blade to each of the hub slots @ 10 (25) Gbps = eight pairs**
  - **Analogy with CMS microTCA application could suggest**
    - ➔ 4 DAQ pairs @ 10 (25) Gbps; blade -> hub
    - ➔ 1 DAQ flow control @ 10 (25) Gbps; hub -> blade
    - ➔ 1 TCDS++; hub -> blade
    - ➔ 1 TTS++; blade -> hub
      - ⇘ Direction reversed; ~allowed by standard
    - ➔ 1 HP clock; hub -> blade
- **Two more talks today touches on the subject**

J23 Connector (zone 2) Node slot

Table 6-5 J23/P23 connector pin assignments for Node Boards/Slots (Base and Fabric Interfaces)

Row #	Interface Designation	J23/P23 Connector Pairs							
		a b		c d		e f		g h	
1	Fabric Channel 2	Tx2[2]+	Tx2[2]-	Rx2[2]+	Rx2[2]-	Tx3[2]+	Tx3[2]-	Rx3[2]+	Rx3[2]-
2		Tx0[2]+	Tx0[2]-	Rx0[2]+	Rx0[2]-	Tx1[2]+	Tx1[2]-	Rx1[2]+	Rx1[2]-
3	Fabric Channel 1	Tx2[1]+	Tx2[1]-	Rx2[1]+	Rx2[1]-	Tx3[1]+	Tx3[1]-	Rx3[1]+	Rx3[1]-
4		Tx0[1]+	Tx0[1]-	Rx0[1]+	Rx0[1]-	Tx1[1]+	Tx1[1]-	Rx1[1]+	Rx1[1]-
5	Base Channel 1	BI_DA1+ (Tx1+)	BI_DA1- (Tx1-)	BI_DB1+ (Rx1+)	BI_DB1- (Rx1-)	BI_DC1+	BI_DC1-	BI_DD1+	BI_DD1-
6	Base Channel 2	BI_DA2+ (Tx2+)	BI_DA2- (Tx2-)	BI_DB2+ (Rx2+)	BI_DB2- (Rx2-)	BI_DC2+	BI_DC2-	BI_DD2+	BI_DD2-
7	n/a	Unused	Unused	Unused	Unused	Unused	Unused	Unused	Unused
8	n/a	Unused	Unused	Unused	Unused	Unused	Unused	Unused	Unused
9	n/a	Unused	Unused	Unused	Unused	Unused	Unused	Unused	Unused
10	n/a	Unused	Unused	Unused	Unused	Unused	Unused	Unused	Unused

NOTE: Color is used in the table above only to clearly indicate Interface groupings.

Base interface
Four pairs ea

From slot 1

From slot 2

ATCA Front Board and RTM

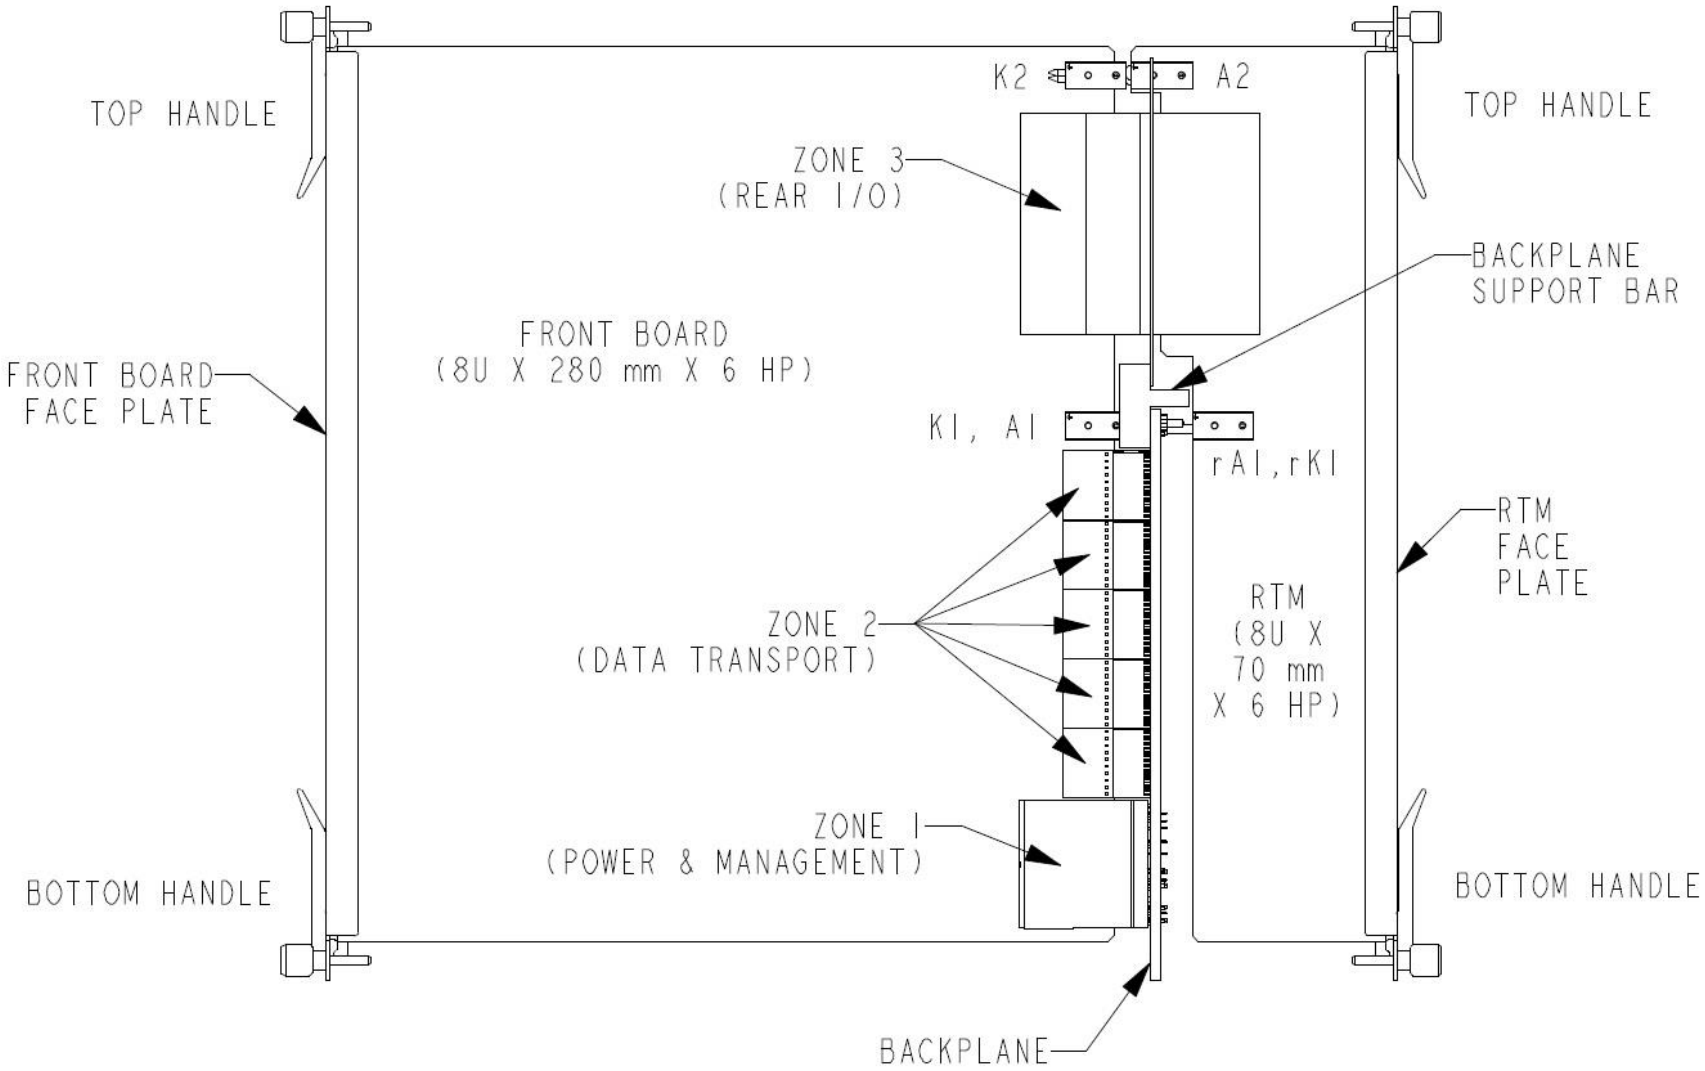

Matching Trigger Needs to ATCA (5)

- Ultra-Fast Trigger Backplane I/O (>10G)
  - Two Zone 2 Issues:
    1. The Full Mesh, consisting of multiple stars, is somewhat awkward for η-φ mapping, which would be better supported by some sort of lattice
    2. Uncertain as to how fast the ADFplus connector can be reliably pushed
  - Zone 3 may be the more promising avenue for ultra-fast intra-crate communications
  - PICMG should eventually address the question of 100G Ethernet on the Fabric Interface, but perhaps not soon
  - Nonetheless, Full Mesh Fabrics have more potential than Dual Star Fabric Interfaces, and seem a very worthwhile target for Phase 2 R&D

Typical ATCA Shelf
(aka crate)