# CMS MicroTCA crate concepts & AMC card requirements.

Gregory Iles, Magnus Hansen, Tom Gorski and Eric Hazen
31 January 2011  Version 0.9 (Draft)

## 1 General concepts

This document describes a proposed CMS-specific configuration of a MicroTCA crate. It describes requirements for crate mechanics, backplane interconnections and some prescriptions for signal protocols on a few of the interconnections. It does not describe any software requirements, nor prescribe any particular implementation except as examples.

We have chosen to explore MicroTCA as a crate system to replace VME for the next generation of electronics cards inside the CMS experiment at CERN. MicroTCA offers a flexible, high density, high performance backplane that is based on the serial standards in use today (GbE, PCIe, SRIO, SATA, etc). It is relatively inexpensive for both the card manufacture and the customisation of the backplane if required.

MicroTCA is based on the AMC (Advanced Mezzanine Card) standard developed by the PICMG group for ATCA cards. Up to 12 AMC cards can be inserted directly into a MicroTCA backplane. A MCH (MicroTCA Carrier Hub) provides connectivity between slots, although direct connections between slots are also allowed. The system can operate in redundant mode with a second MCH (MCH2) connected to each AMC card and to the primary MCH (MCH1). For CMS we focus exclusively on this redundant type of "dual star" crate.

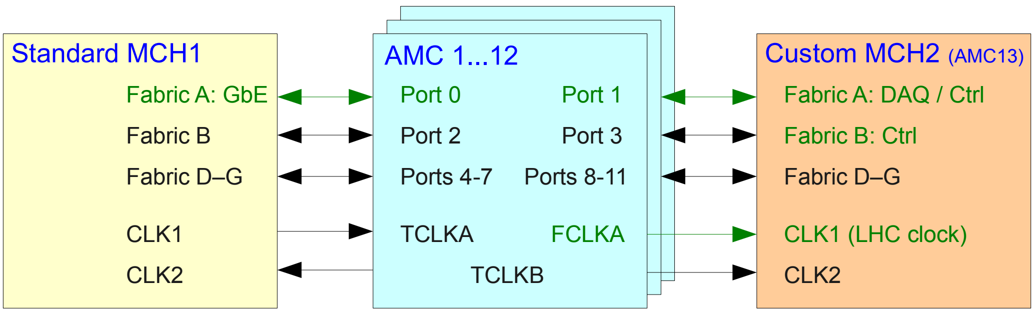


*Illustration 1: Typical "Dual-Star" MicroTCA Backplane.*
*Note that each interconnection shown represents a "star" connection from each MCH to up to 12 AMCs. The Fabric connections are bidirectional (one pair each direction).*

Illustration 1 shows a simplified view of the dual-star configuration. The connections shown in green color represent “required” features of the CMS standard configuration described in this document.

In commercial telecoms applications, a redundant crate provides hot spare capability by providing two redundant MCH modules with duplicate signal routing from each MCH to each AMC. In CMS we use this capability differently: MCH1 is an off-the-shelf module which provides management and communications functions described in the MicroTCA and AMC standards, while MCH2 is customized to provide LHC timing and DAQ functions specific to CMS requirements.

# 2 CMS-Specific Configuration

Now the CMS-specific MicroTCA configuration is described. Essentially we propose to use an off-the-shelf dual-star crate with some key features selected and a custom MCH2 module. The AMCs of course would be customized anyway.

The key features of the proposed configuration are:

- Commercial MCH1 for crate management, GbE communication, and other user features as desired
- Custom MCH2 providing:
  - LHC 40.xx MHz low-jitter clock distribution
  - Fixed-latency controls distribution (aka TTC)
  - DAQ functionality; readout of data from AMCs
  - Buffer management communications for TTS-like functions as well as possible selective readout control
- Approved crates with the following features:
  - 12 full-height double-width AMC slots preferred (but see Section 2.2)
  - Two standard (single-width) MCH slots
  - *JTAG access to all modules via JSM (optional)* *[Omit? Ed]*
  - Approved power modules with xxxV bulk input (voltage T.B.D.)
  - Vertical airflow for cooling
- Backplane with the following interconnections:
  - Dual-star routing of Fabrics A, B, D, E, F, G to MCH1 and MCH2
  - Dual-star routing of CLK1 to TCLKA (MCH1) and CLK1 to FCLKA (MCH2)

## *2.1 Custom MCH*

A MicroTCA MCH is built up from 4 PCB tongues, which each have an AMC-style edge connector. Access to the AMC cards is via “fabrics”. The MCH has 6 fabrics, labelled A to G (excepting C), which each provide up to 12 bidirectional high-speed[1] serial ports. For CMS we plan to use a “dual star” topology redundant backplane with two MCH slots. Table 1 shows the proposed use of the fabric and clock ports for CMS.

| Fabric | AMC Port | MCH1 | MCH2 | Category | MCH Tongue | CMS Use |
|---|---|---|---|---|---|---|
| A | 0 | Yes | | Common Options | 1 | GbE |
| A | 1 | | Yes | | 1 | DAQ |
| B | 2 | Yes | | | 2 | |
| B | 3 | | Yes | | 2 | Timing |
| | TCLKA | Yes | | Clocks | 2 | |
| | FCLKA | | Yes | | 2 | LHC CLK |
| D | 4 | Yes | | Fat Pipes | 3+4 | User |
| E | 5 | Yes | | | 3+4 | User |
| F | 6 | Yes | | | 3+4 | User |
| G | 7 | Yes | | | 3+4 | User |
| D | 8 | | Yes | Extended Fat Pipes | 3+4 | User |
| E | 9 | | Yes | | 3+4 | User |
| F | 10 | | Yes | | 3+4 | User |
| G | 11 | | Yes | | 3+4 | User |
| | 12...20 | | | Unassigned | | |

*Table 1: MicroTCA Dual-Star Backplane Port Use*

For CMS, we place essentially no restrictions on communication between the commercial MCH in the MCH1 site; users may use standard protocols (eSATA, PCIe, etc) on the ports routed to MCH1 (0, 2, 4-7). We reserve only the port 1 and 3 connections to MCH2. However, please see Section 3.1 for some important information on PCIe clocking.

1 Typical bit rates from 1-10Gb/s

The proposed MCH2 board stack is shown in Illustration 2. Two of the four tongues are required. Tongue 1 functions as a "13th AMC", which means that it contains an MMC (module management controller) similar to the one on a standard AMC rather than the hub functions normally associated with an MCH. The concept of using a redundant MCH slot for an AMC card is supported and part of the MicroTCA specification (A.4.2, PICMG® Specification MTCA.0 R1.0).

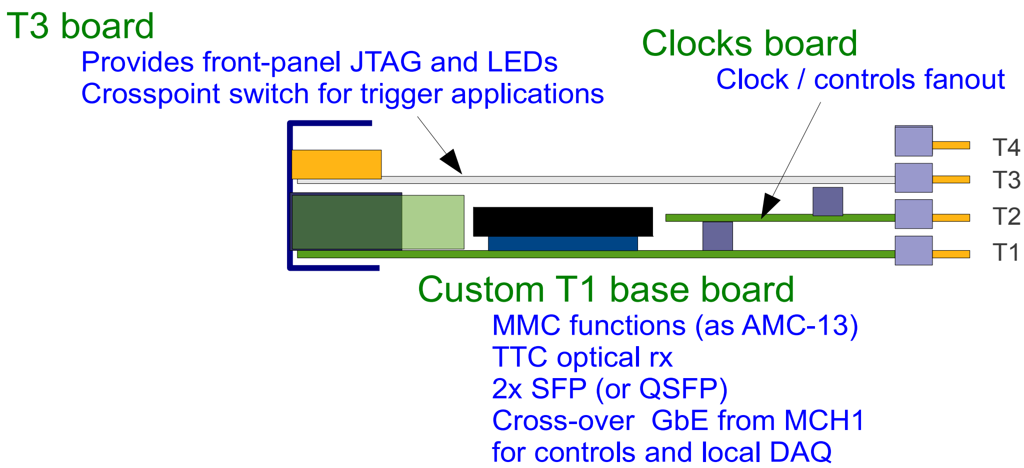


*Illustration 2: CMS MCH2 Proposed Board Stack*

## 2.2 MicroTCA Crate

We propose a crate layout similar to that shown in Illustration 3. Twelve full-height, double-width AMC modules are accommodated, along with two MCH. Note that in MicroTCA terminology, "full-height" refers to the horizontal dimension in Illustration 1.

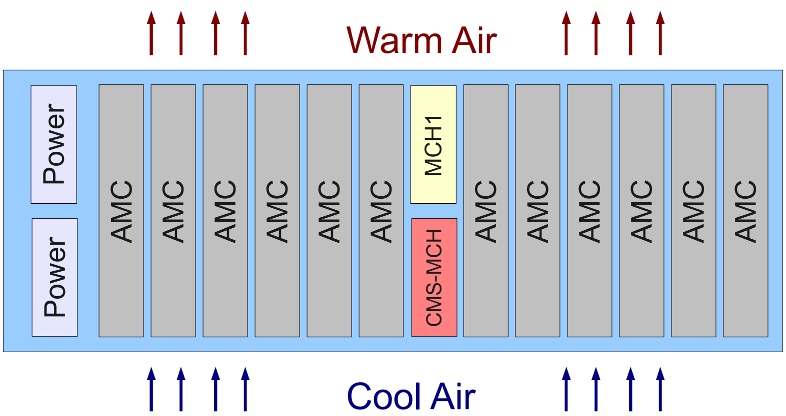


*Illustration 3: Proposed MicroTCA Crate Layout.*

This crate layout is quite similar to some on the market already.

An alternative crate layout is proposed by the xTCA Physics Profile working group, which provides 12 mid-height AMC slots.  It is recommended that AMC designers confine their designs to the mid-height envelope to permit their installation in such crates.

### 2.3 Signal Encoding

The complete details of the protocol transmitted over each backplane port is beyond the current scope of this document, but here is a brief summary.

**Fabric A**

The primary channel for data acquisition is on Fabric A (AMC port 1).  This channel may either operate at a multiple of the LHC clock frequency or based on a 125MHz crystal.  This link should be an AC-coupled CML connection to a SERDES transceiver.  The bit rate should be in the range 1.0Gb/s – 3.125 Gb/s, and the AMC endpoint of this link should be capable of operating correctly when connected to a Xilinx GTP™ SERDES block.  This requires that the SERDES on the AMC be capable of operating from either the FCLKA (LHC clock) or a crystal.

Fabric A will be used as follows:

- DAQ data *will* be carried from AMC to MCH2
- *May* also be used for fixed-latency controls (TTC)
- *May* also be used for buffer status (TTS)
- *May* also be used for auxiliary data in either/both directions (i.e. selective readout)

**Fabric B**

The primary channel for fixed-latency, fixed-frequency control information is Fabric B (AMC port 3).  This is an LVDS, DC-coupled  channel operating at a multiple of the LHC clock frequency in the range 80-400Mb/s.  Fabric B will be used as follows:

- Fixed-latency controls from MCH2 to AMC
- Buffer status from AMC to MCH2 (TTS)
- Auxiliary data in either/both directions (i.e. selective readout)

In addition, the LHC clock (40.079xx MHz) will be distributed on the MicroTCA MCH CLK1 network.

These assignments use only tongues 1 and 2 of the MCH site, and only Fabric A (ports 0/1) and Fabric B (ports 2/3).  Fabrics D-G and the corresponding tongues 3 and 4 of the MCH site are left open for the user.

**CLOCK 1**

The user **should** connect AMC TCLKA (CLK1 in PICMG® AMC.0 R1.0) and FCLKA (CLK3 in PICMG® AMC.0 R1.0) to a M-LVDS receiver.

# 3 Appendix 1: Other options

## *3.1 PCIe Clocking*

MicroTCA is effectively available in two flavours for 12 slot backplanes:

(a) telecom based with redundant clock network plus fabric-B
(b) non-redunant clock network without fabric-B, but with an additional clock for fabric protocols that require it (e.g. PCIe).

This split exists because PCIe required an extra clock. The only spare clock at the time was the redundant MCH clock (i.e. AMC-CLK3 in PICMG® AMC.0 R1.0). This has since been renamed in PICMG® AMC.0 R2.0 to FCLKA to indicate that it is a "Fabric" clock. Ports 7-12 of fabric B on tongue-2 of MCH1 were re-assigned to drive FCLKA, which in the past had been driven from the redundant MCH (MCH2-CLK1). AMC ports 2 and 3, which would normally have had a connection to fabric-B were directly connected without a MCH connection.

We believe that the original telecom standard is more suited to high energy physics because we may wish to have redundancy and fabric-B provides an additional interface to the AMC card when the other fabrics are effectively already assigned (e.g. fabric-A is reserved for communication (Ethernet); fabric C is not present because that region of the MCH tongue 2 is used for MCH-CLK1 and MCH-CLK2 and fabrics D-G are reserved for a fat pipe (x4 lanes).