TWEPP 2010 – Aachen

Development of a MicroTCA Carrier Hub for CMS and SLHC

E. Hazen - Boston University
for the CMS Collaboration

Outline

- CMS Introduction
  - Detector, Trigger / DAQ system
- MicroTCA introduction
- CMS readout crate – typical functions
- MicroTCA Carrier Hub (MCH) features
- Custom MCH Details
- Test Results
- Plans

The CMS Detector

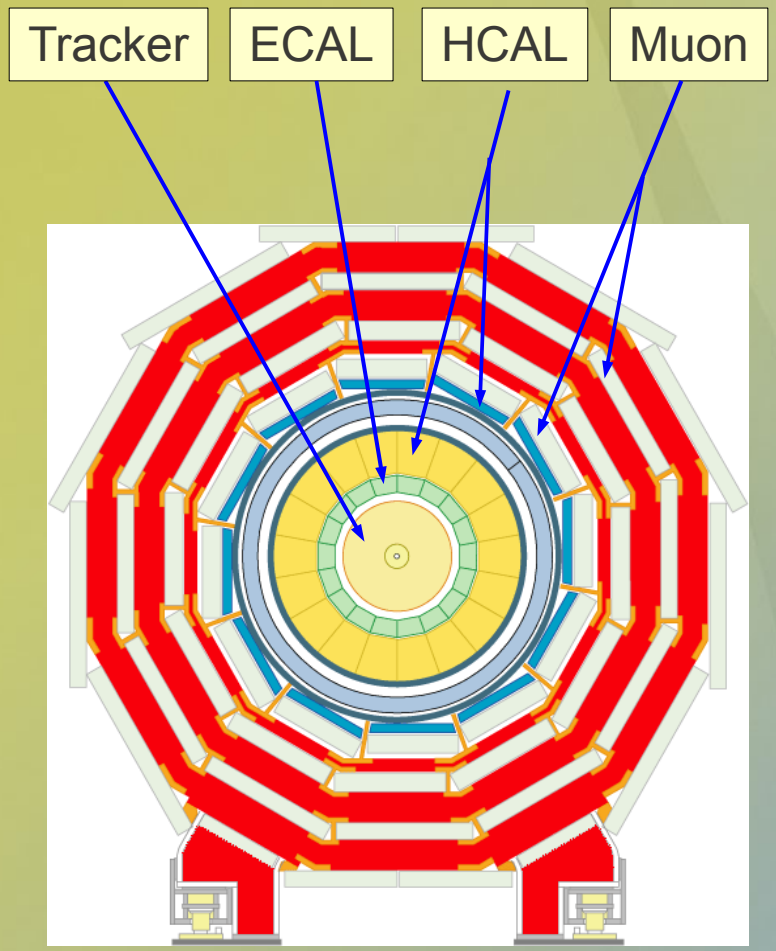

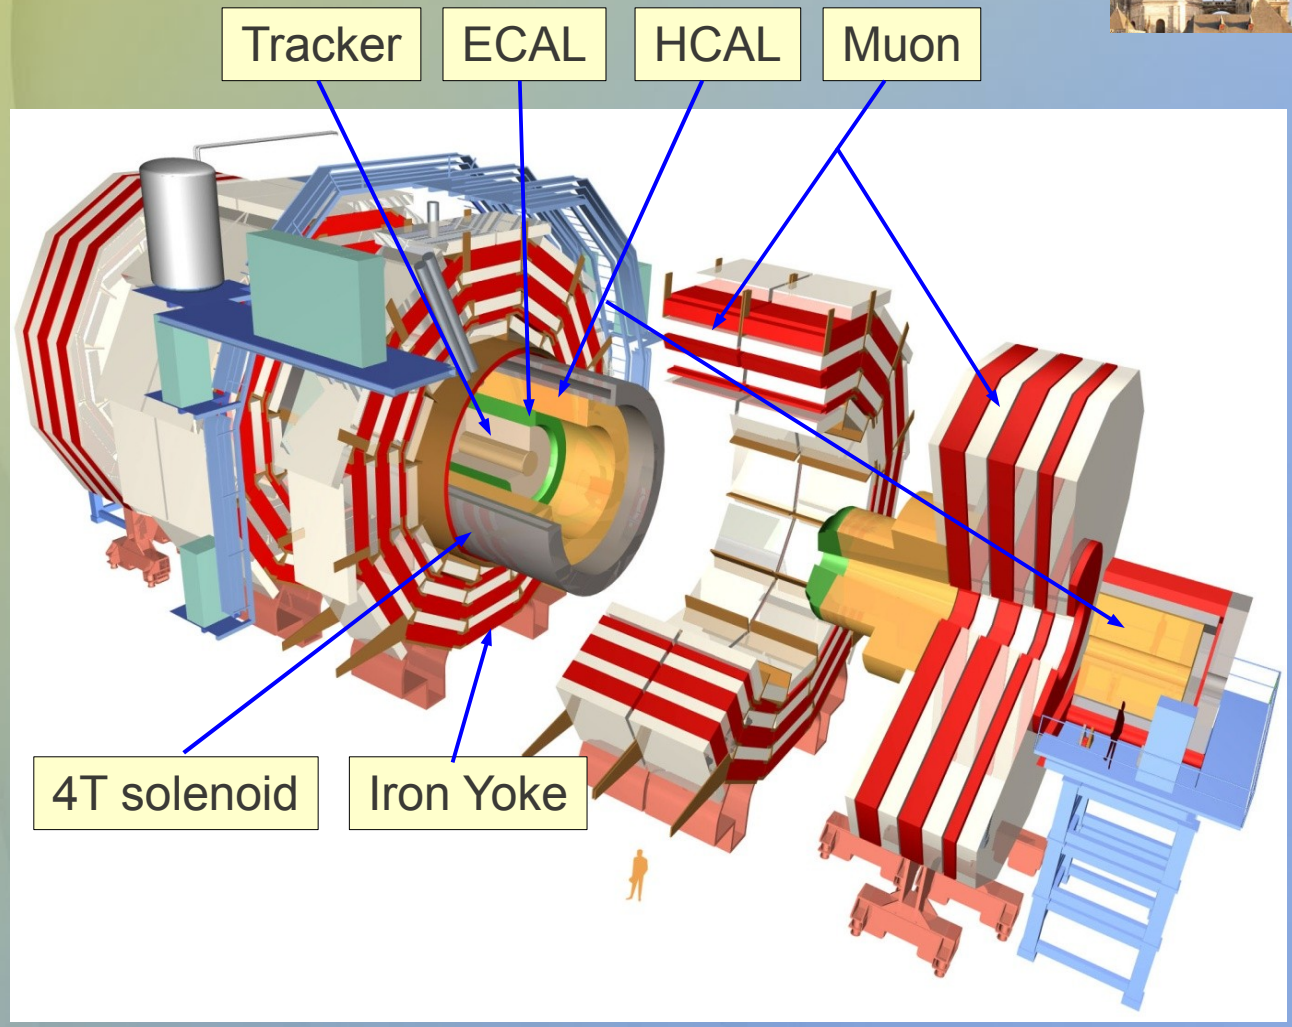

Physics Objectives:
- Higgs search
- SUSY
- Extra dimensions
- Matter/antimatter
- Quark gluon plasma

Total weight:	12,500 t
Diameter:	15m
Length:	22m
Magnetic field:	4T

CMS Trigger/DAQ

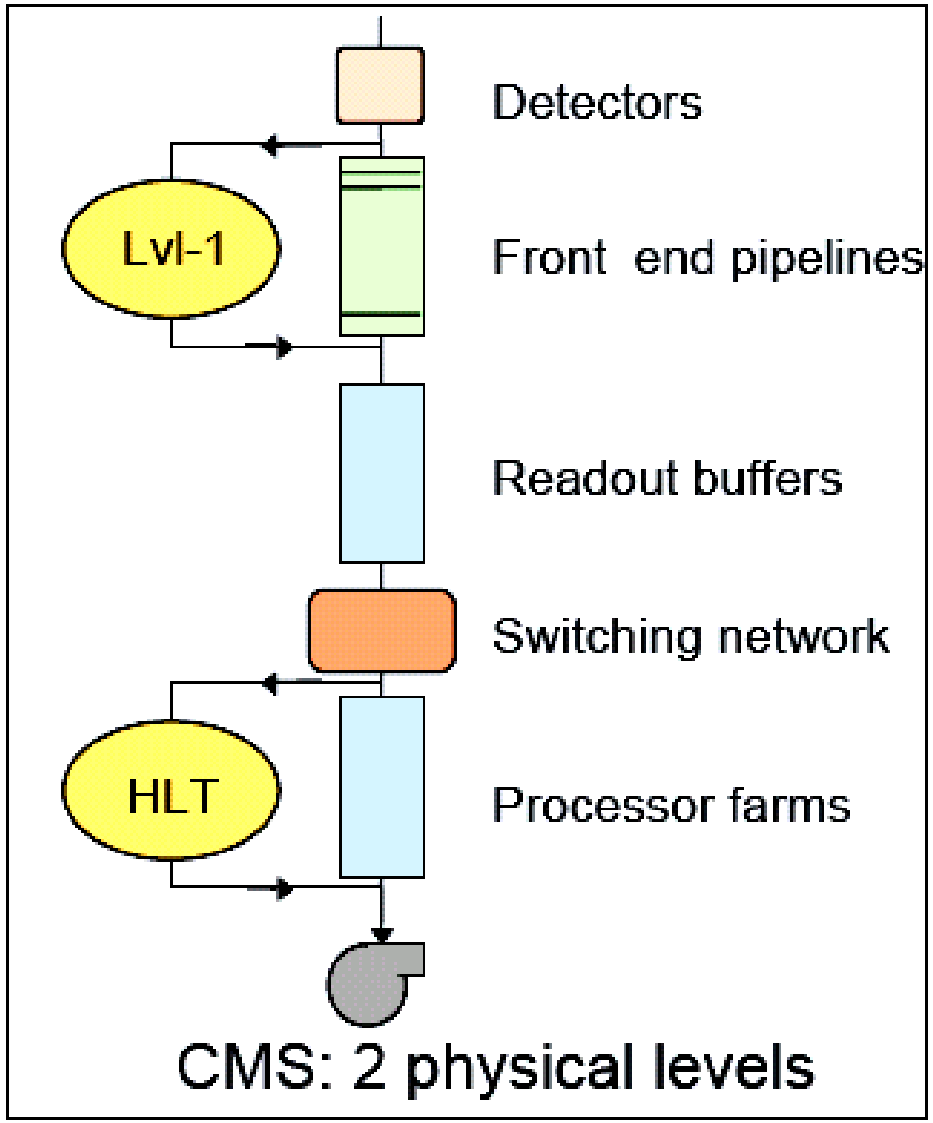

- 40MHz input (one measurement per BX)
- Triggers pipeline in Level 1 at 40MHz
  - L1A rate 100kHz
- Data stored in pipelines for 3.2 uS
- Data copied to readout buffers at 100kHz
- Readout buffer data sent from underground to surface at 200Mbytes/sec * 512[1] links
- High-level trigger in software

Items in red specified not to change for upgrades until ~ 2020

[1] Up to 512. About 446 currently in use.

CMS Trigger/DAQ Hardware

- Currently a mix of VME / cPCI / proprietary standards
- Common functions provided (typically VME)
  - Control / monitoring / firmware updates (VMEbus)
  - Low/fixed latency paths – out-of-band signals
    - Trigger and fast timing
      - cables or custom backplanes
    - Data Concentration / DAQ output
      - cables or custom backplanes
    - Board-to-board communication for trigger
      - Custom backplanes
  - Different subsystems use different, incompatible extensions to the VME standard :(

Existing VME Crates

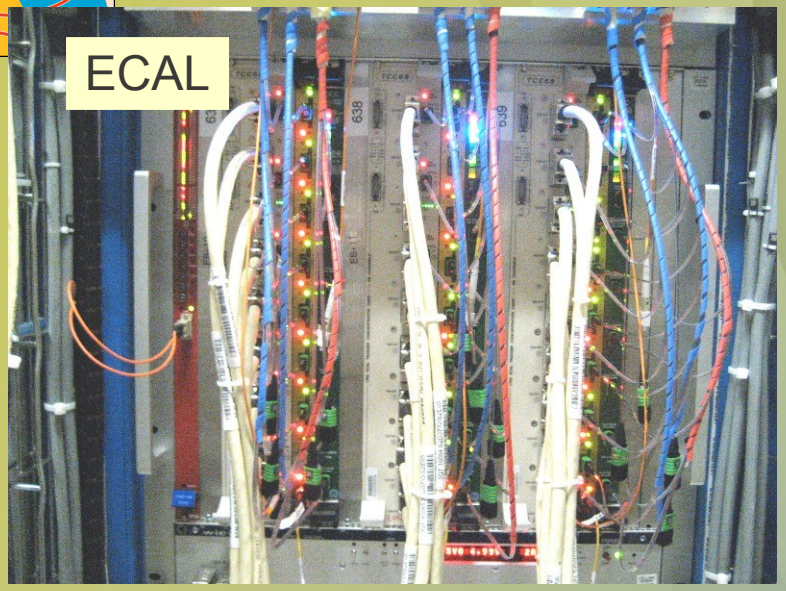

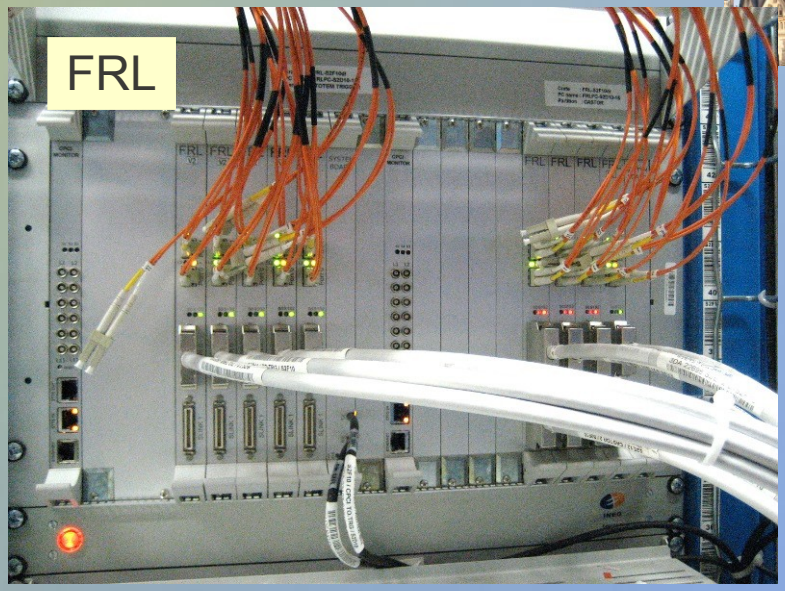

Wide variety of front/rear I/O driven partly by requirements but largely by the taste of the designer! We hope to minimize front panel interconnects in upgrade

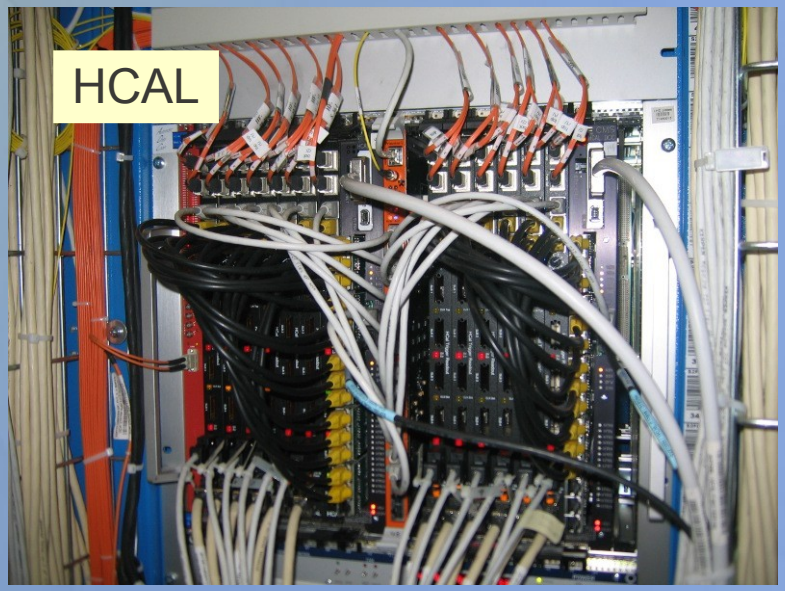

CMS Level 1 Trigger

A significant portion of the CMS trigger / DAQ will migrate to MicroTCA

(my opinion)

Expected to migrate From VME (or similar) to MicroTCA

Possible future migration to MicroTCA

On detector electronics

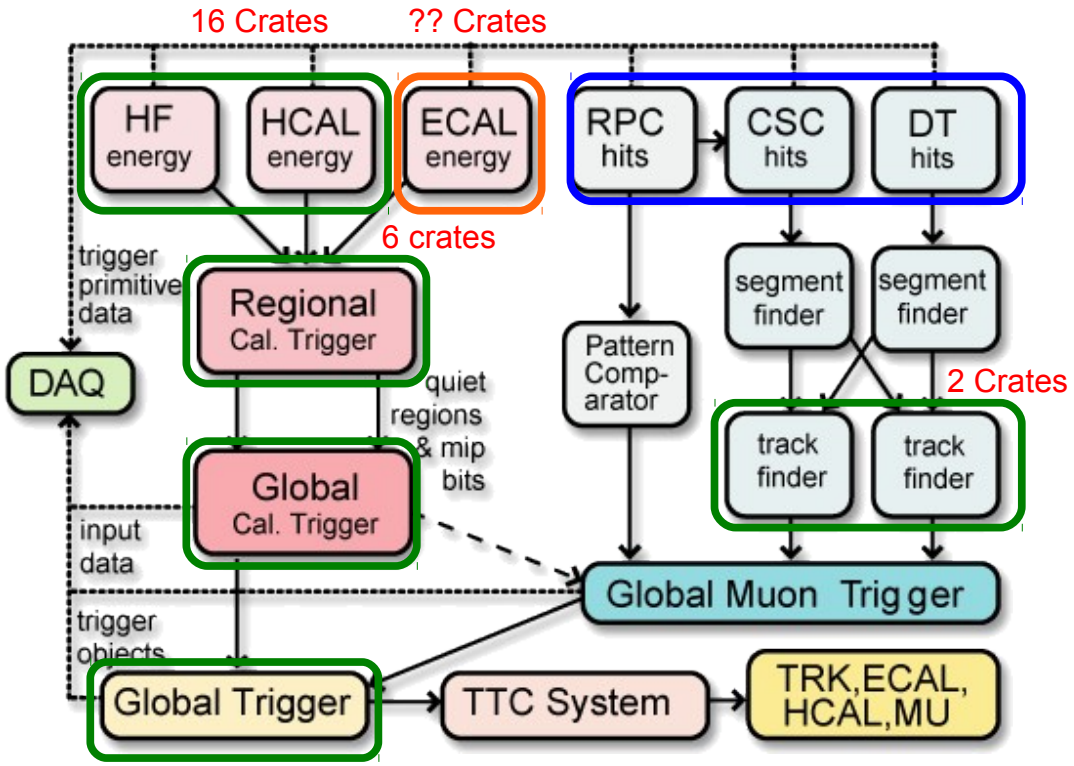

Common Crate Effort

- We are trying to standardize early…

  Draft document on CMS TWiki
  topic: MicroTCA

- Effort organized by G. Iles (see later talk)

CMS MicroTCA crate concepts & AMC card requirements.

Gregory Iles, Magnus Hansen and Eric Hazen
14 June 2010
Version 0.6 (Draft)

1 General concepts

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).

The MCH is built up from 4 tongues, which each have an AMC edge connector. Access to the AMC cards is via "fabrics". The MCH has up to 7 fabrics, labelled A to G, which each provide up to 12 bidirectional serdes ports. Fabric C is often not implemented because the region on that tongue is used to provide clock capability to the MCH. The infrastructure setup is done via IMPI over I2C. Communication is via GbE over fabric-A on tongue-1 to AMC port 0 or 1 (redundant MCH). SAS/SATA is typically through fabric B on tongue-2 to port 2 or 3 (redundant MCH). Alternatively SATA/SAS is simply routed directly between AMC cards. Fabrics D to G on tongues 3 and 4 provide a "FatPipe" that connect to ports 4-7 or 8-11 (redundant MCH). These are used for 4xPCIe, 4xSRIO, 10GbE, etc.

µTCA Introduction

- Derived from AMC
  Advanced Mezzanine Card std
  - Up to 12 AMC slots
    Processing modules
  - 1 or 2 MCH slots
    Controller Modules
- 6 standard 10Gb/s point-to-point links from each slot to hub slots (more available)
- Redundant power, controls, clocks

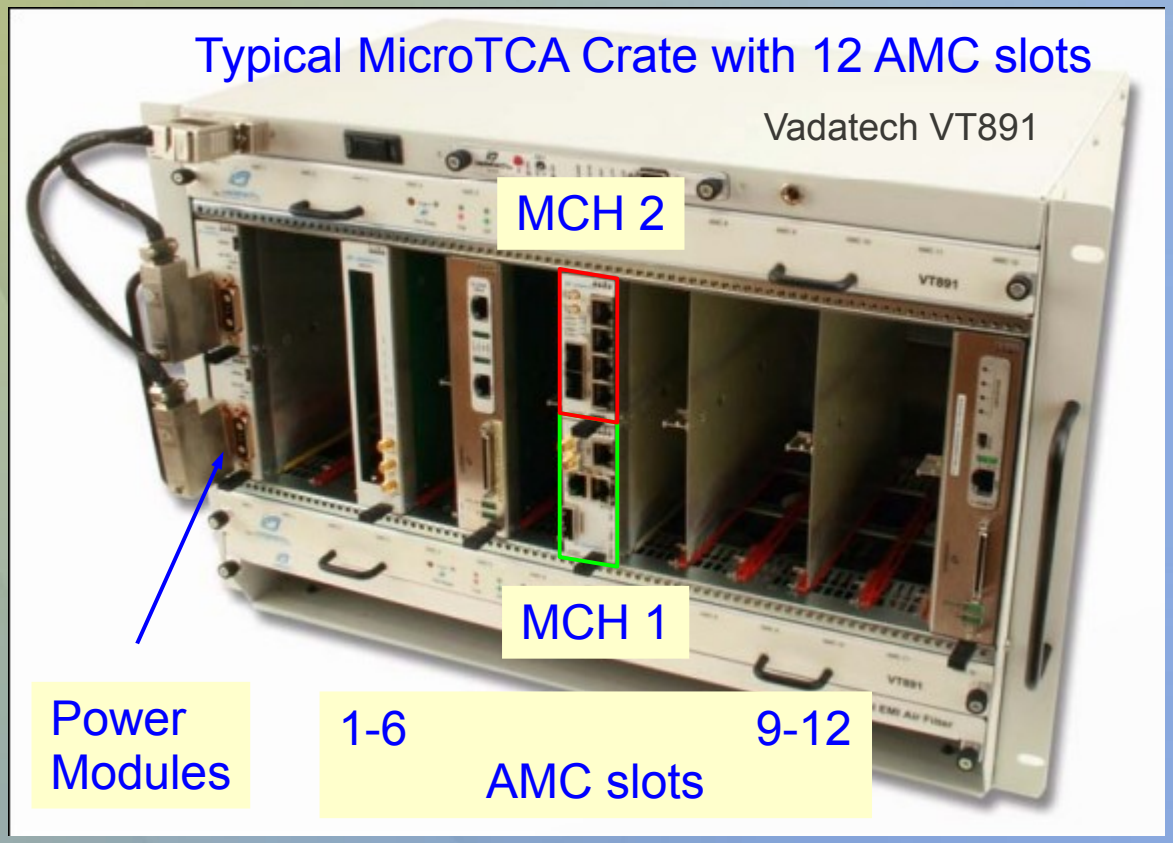

Single Module (shown):
75 x 180 mm
Double Module
150 x 180mm

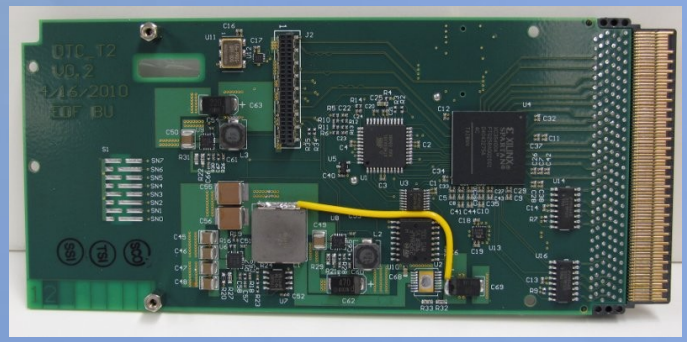

More about uTCA

- Single and double “width” modules (75, 150mm)
- Each AMC can have in principle (20) 10 Gb/sec ports
- Backplane customization is routine and inexpensive
- Redundant hub slots can be co-opted for CMS use:
  - Primary slot for standard hub: management features
  - Redundant slot for CMS-specific features:
    - Fast timing and controls
    - Data concentration and DAQ

µTCA Dual-Star Backplane

Note: Interconnections can be customized by the backplane manufacturer inexpensively.

Bi-directional serial (up to 10Gb/sec) point-to-point links from each AMC to MCH (redundant links to each MCH)

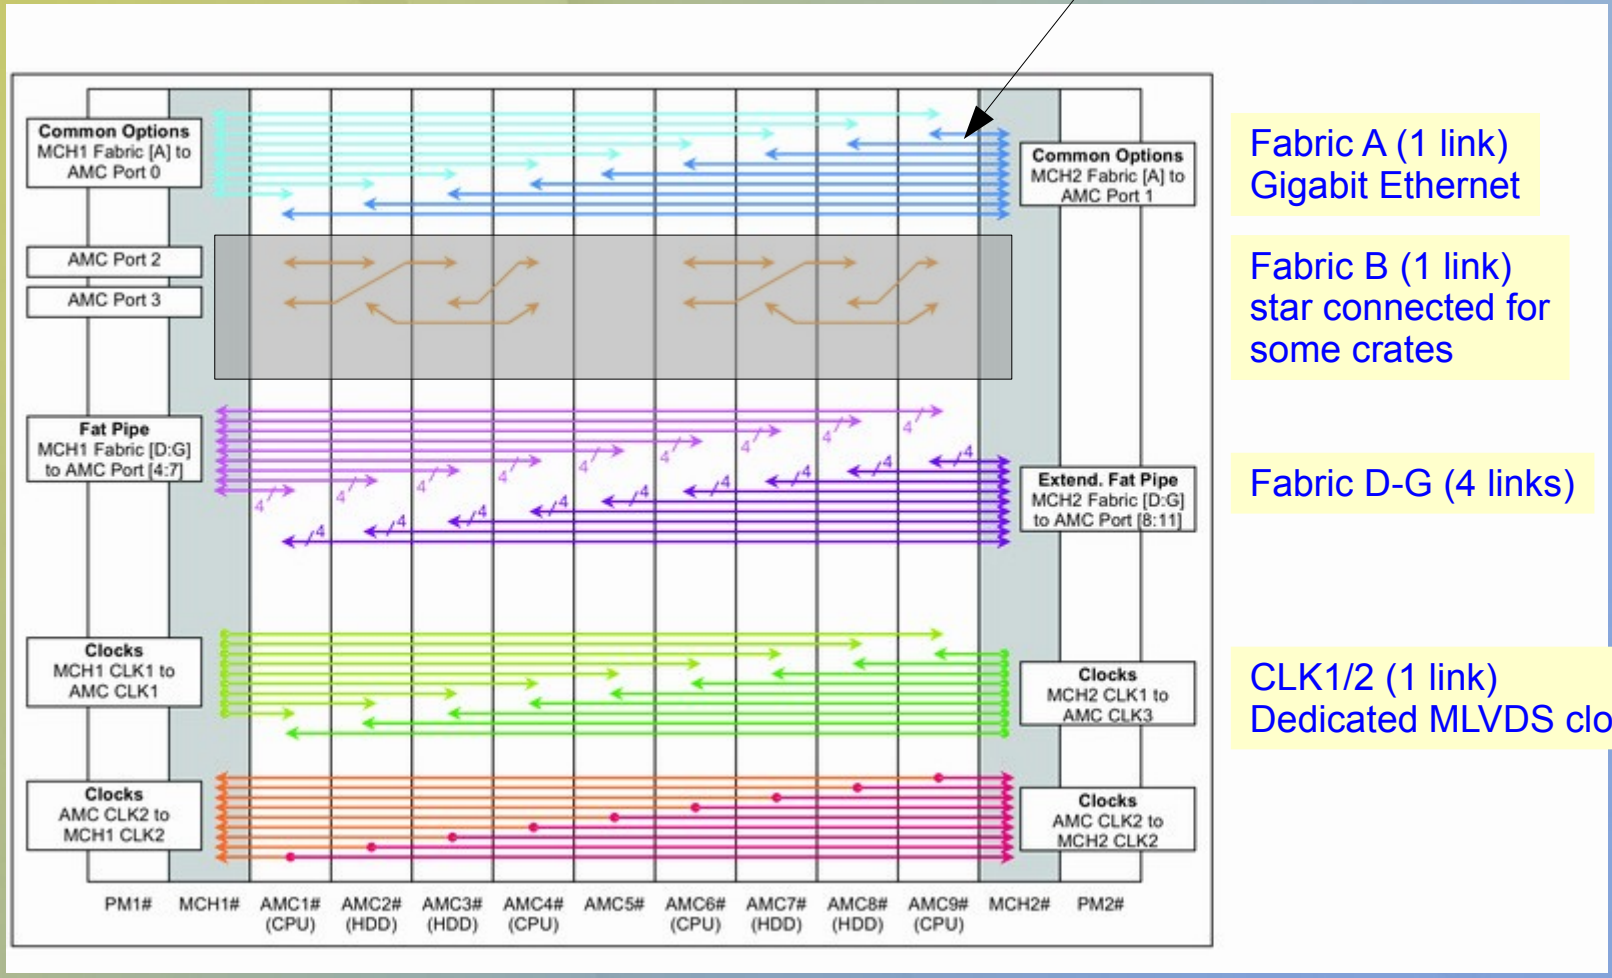

Fabric A (1 link) Gigabit Ethernet

Fabric B (1 link) star connected for some crates

Fabric D-G (4 links)

CLK1/2 (1 link) Dedicated MLVDS clock

MCH 1

MCH 2

CMS Readout Crate (i.e. HCAL)

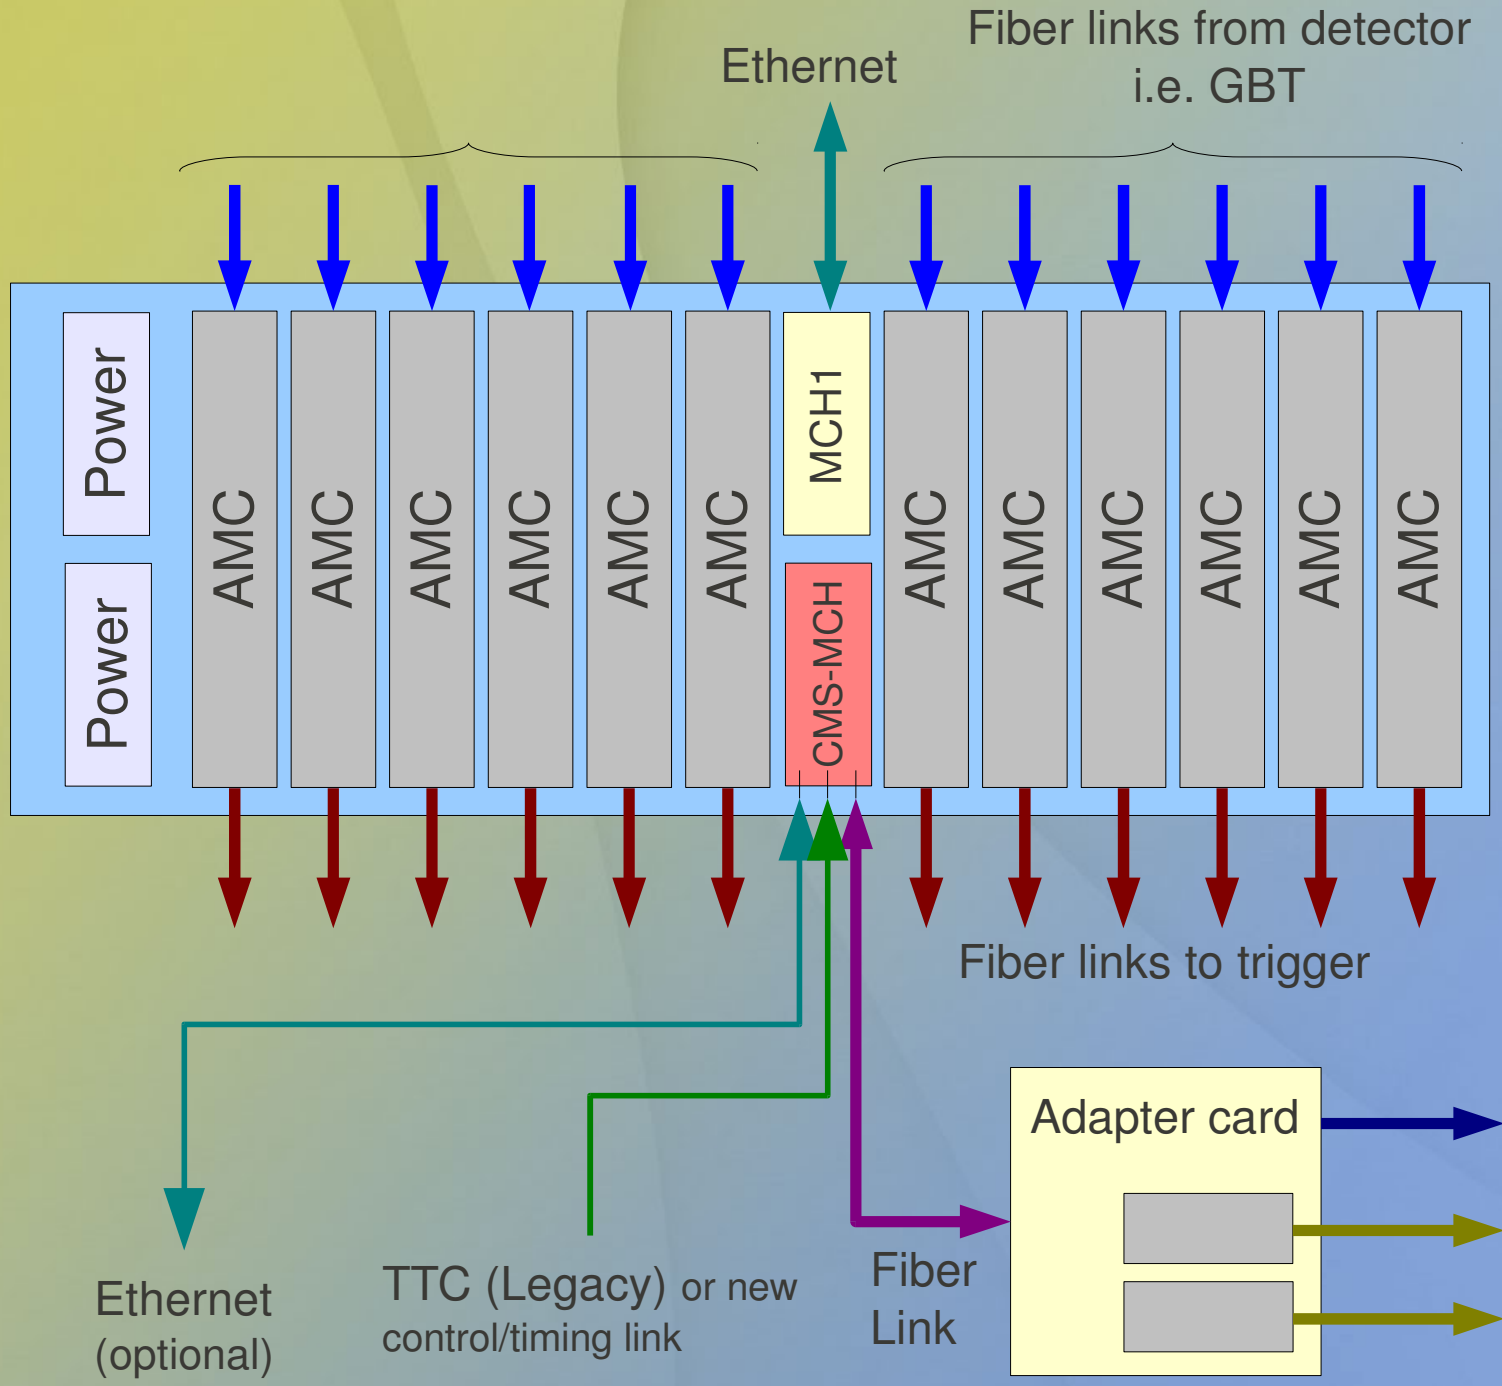

MCH Construction

Four stacked PCBs, 8mm board-to-board spacing!
Mechanically a real challenge.

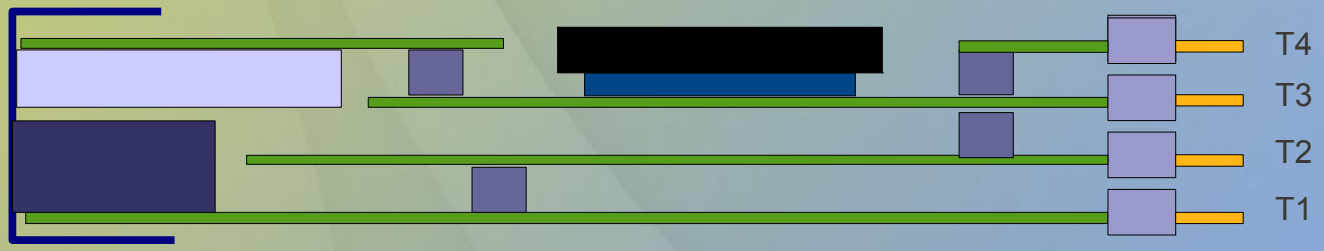

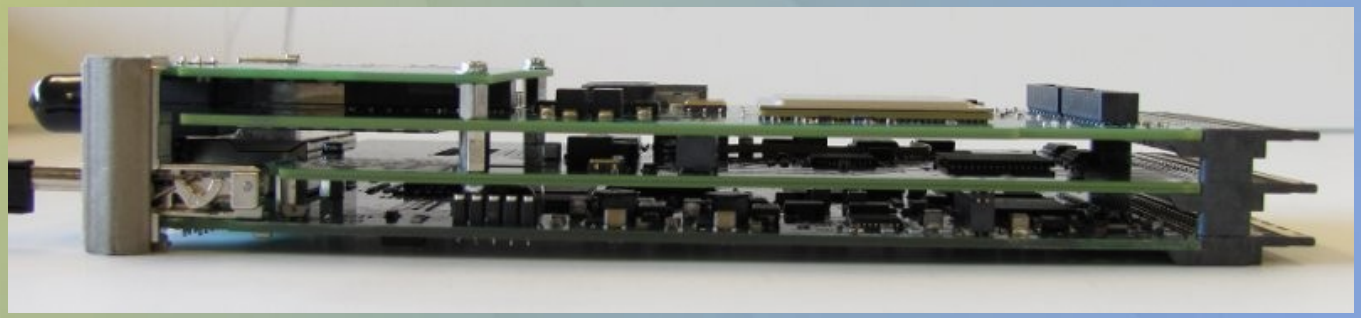

T1 board – Management functions; Fabric A (Ethernet) port
T2 board – Clocks and Fabric B port
T3 and T4 – Fabric D-G ports

HCAL DTC (aka CMS-MCH)

Built on NAT-MCH base board

- Single-width MCH module
- Provides controls, clock distribution (TTC initially)
- Provides DAQ interface via fat pipes (external SLink/TTS adapter)
- Font panel I/O:
  - 2x SFP+ optical transceiver DAQ + Spare
  - TTCrx fiber receiver

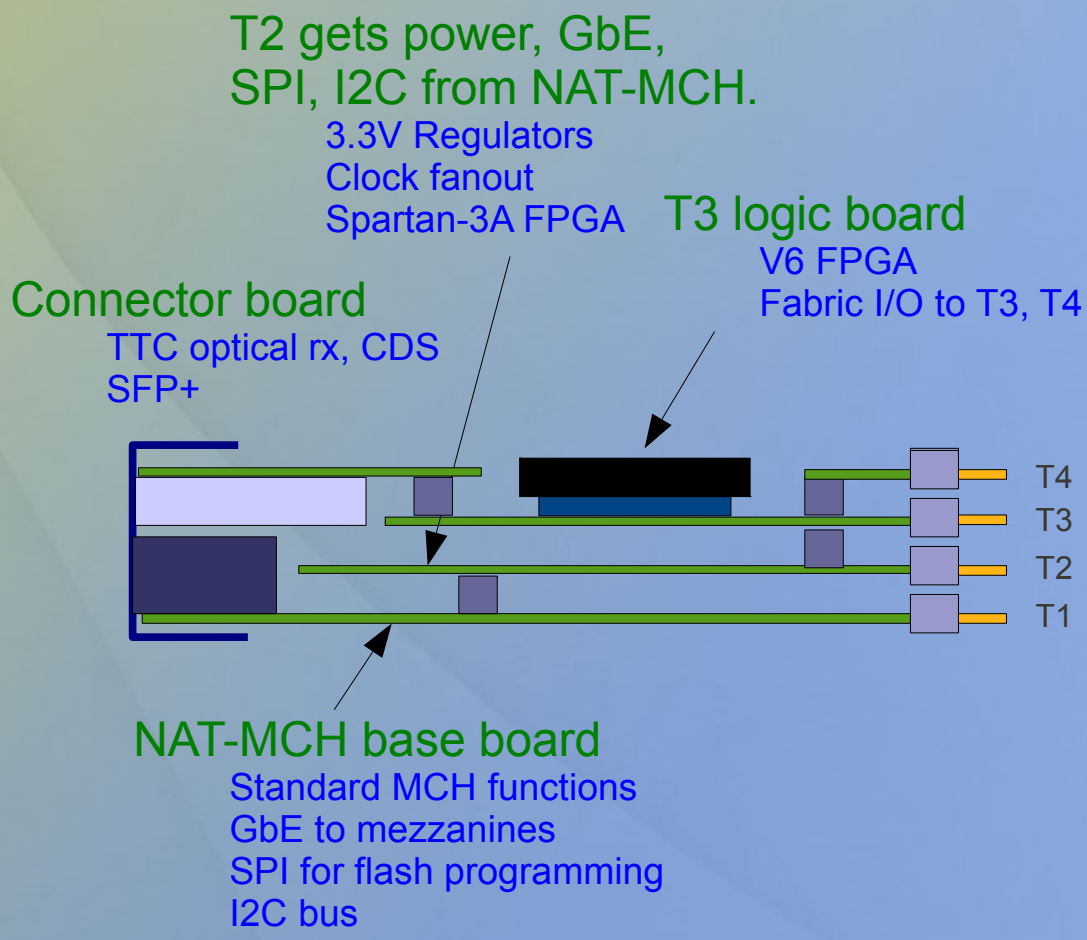

Ports Use

- TTC-derived clock on MCH CLK1 (MLVDS per std)
- Fabric A (port 0/1)
  - 1000BaseX via NAT-MCH
- Fabric B (port 2/3) or Fabric E (port 5/9)
  - Encoded Timing @ 80Mb/sec (currently DC-LVDS)
    May increase speed to ~ 400Mb/s in future
- Fabric D (port 4/8)
  - DAQ (GTX on V6 FPGA)
    - *Port (a/b) represents (MCH1/MCH2) slot*

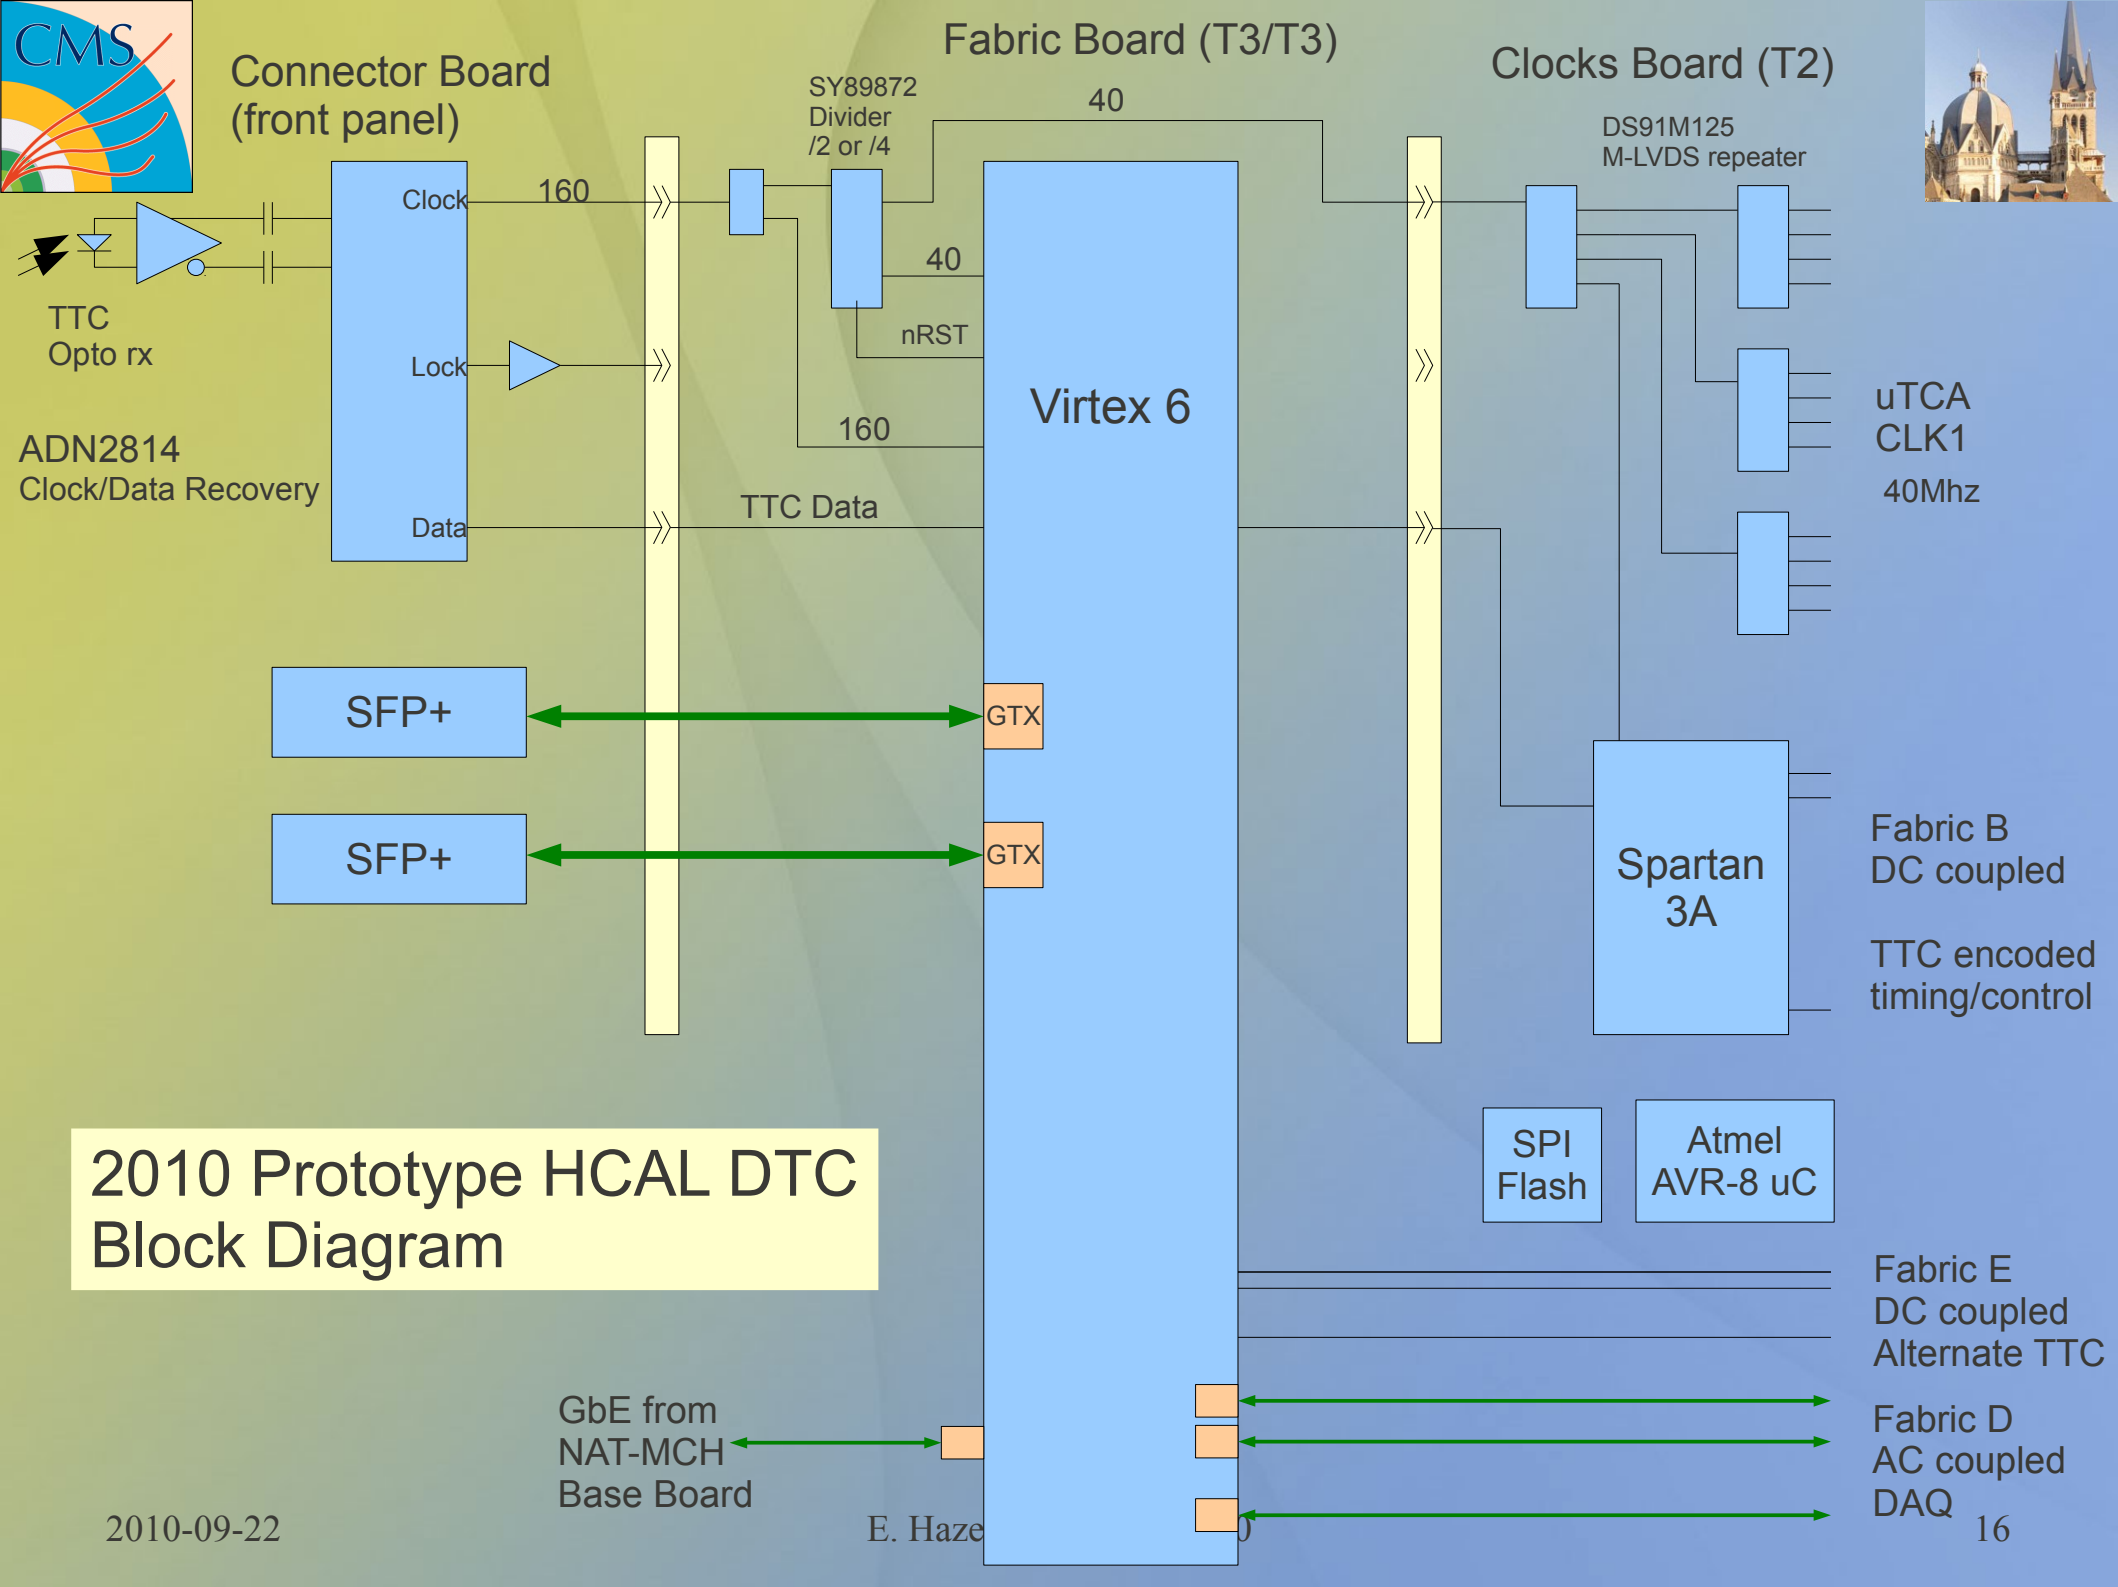

2010 Prototype HCAL DTC Block Diagram

MCH Base Board from NAT Europe

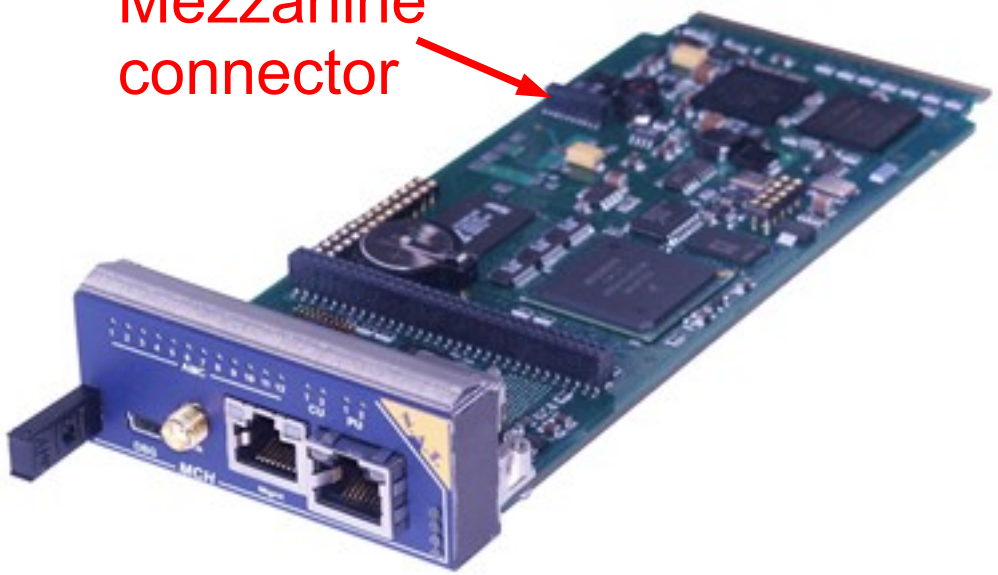

MicroTCA Carrier Hub (MCH)

The NAT-MCH the powerful management and data switching entity for all microTCA (uTCA) systems.

Key Features:

* Support for 12 AMCs, 2 Cooling Units, multiple Power Units
* GigaBit Ethernet switching (Fabric A)
* PCI Express switching (Fabrics D-G)
* SRIO switching (Fabrics D-G)
* XAUI (10GbE) switching (Fabrics D-G)
* Telecom, non-Telecom and Fabric clocks
* Management Controller (MCMC)
* Java based GUI for Linux and Windows Hosts

Provides required MCH management functions
  Simplified initial design
  Allows operation in non-redundant crate (one MCH)

Mezzanine connector supports NAT proprietary stack (clocks, fabric mezzanines) or our own home-made ones

Considerable support and documentation provided by NAT

Services Provided by NAT-MCH

- GbE via switch
  - Not yet used, but plan to use for DTC control and local DAQ
- SPI bus
  - FPGA Flash programming (very convenient)
- IPMI
  - Management (temperature readout, versions etc)
  - DTC control (backup path) via Atmel AVR-8 uC
- Switched payload power

DTC Custom Clocks Board

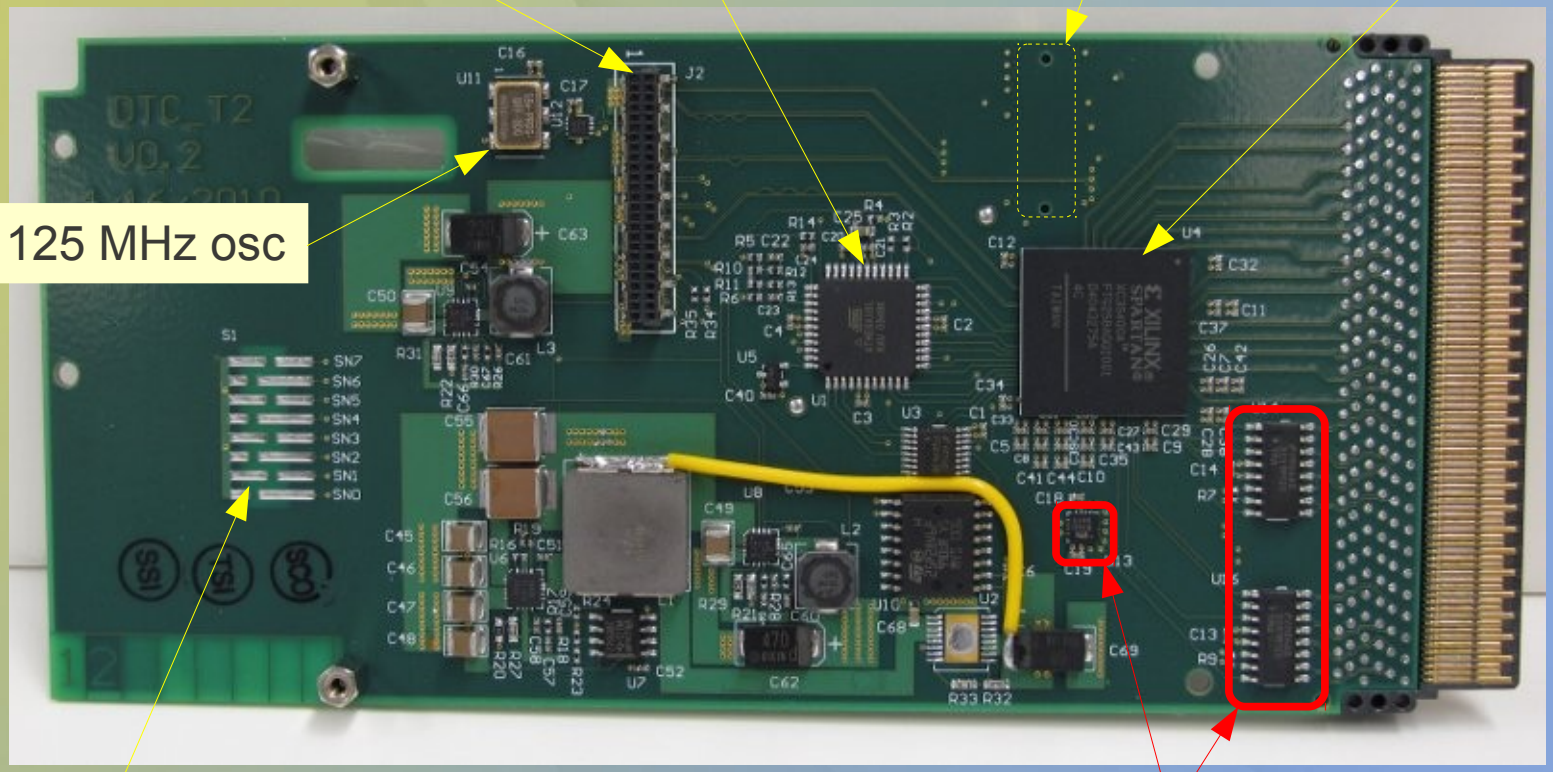

DTC Custom T3 Board

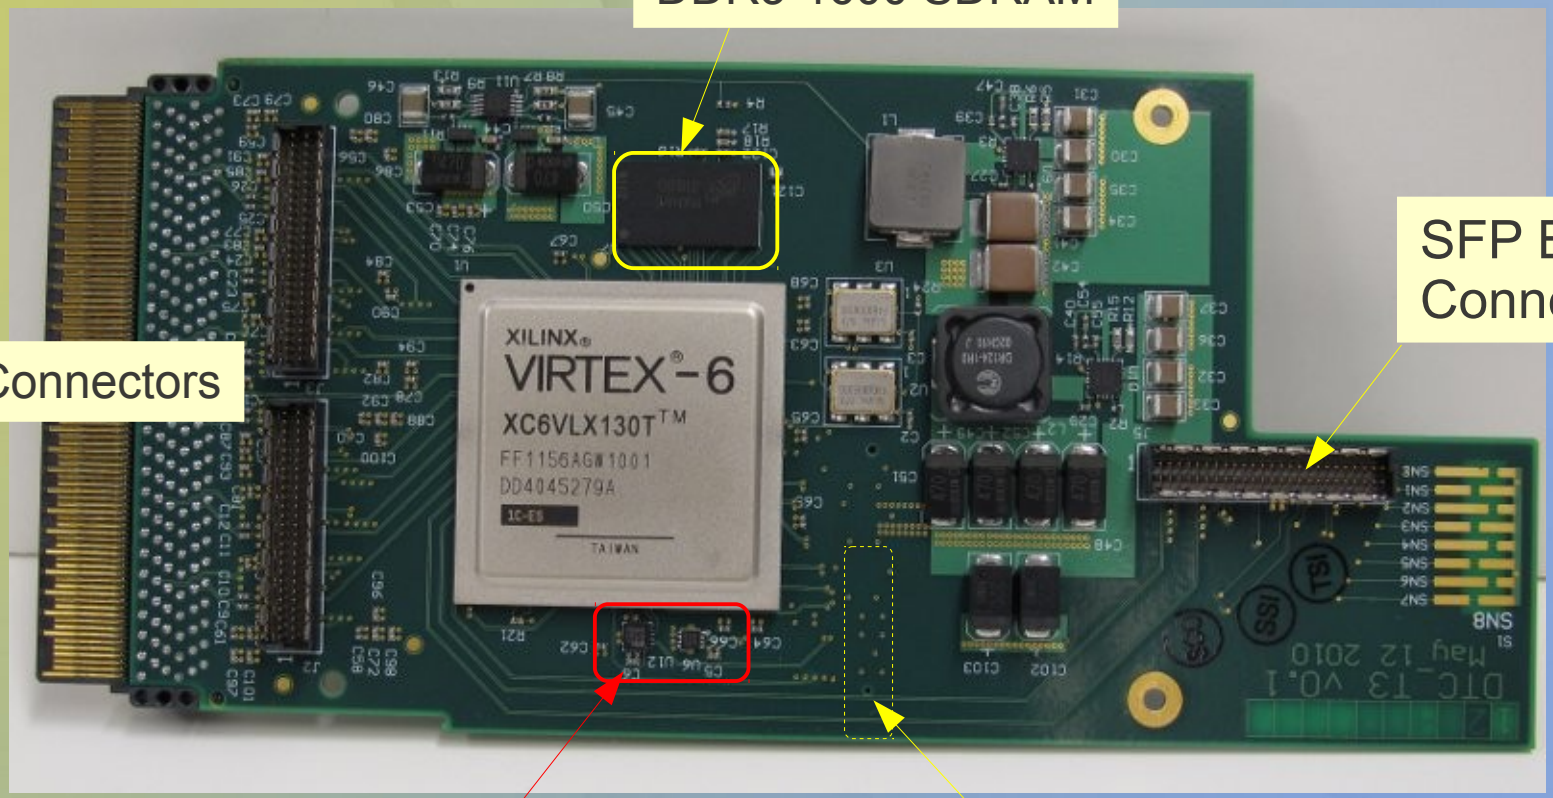

DTC Custom SFP/TTC Board

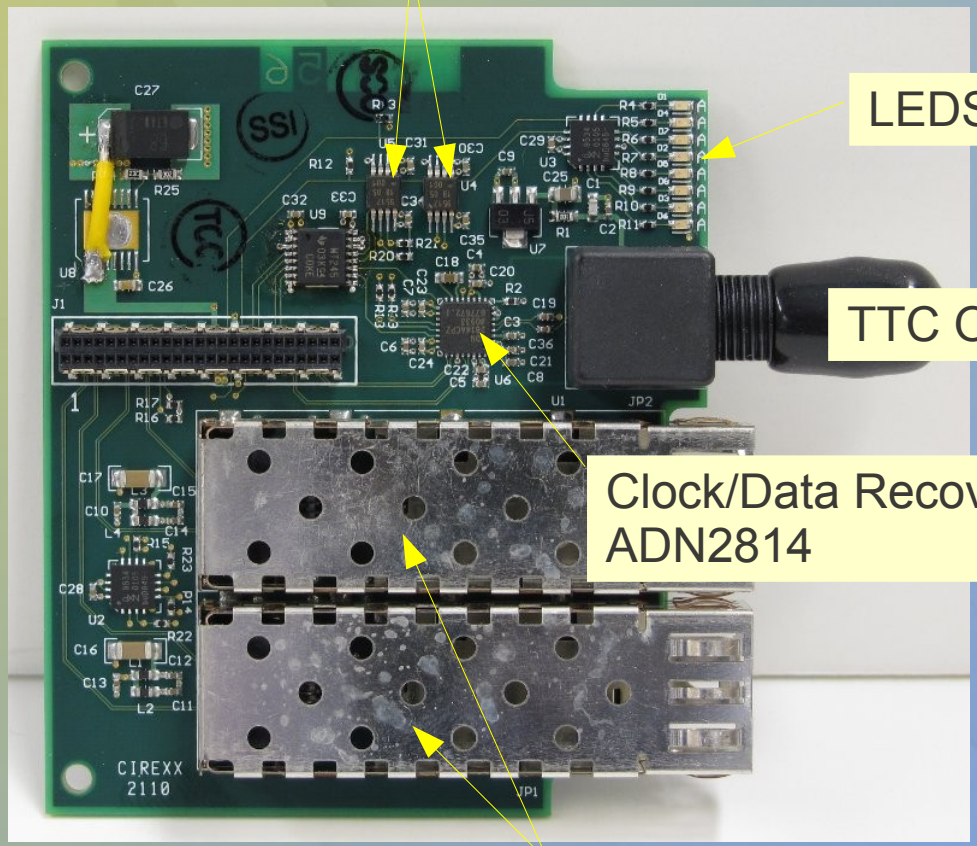

Front Panel

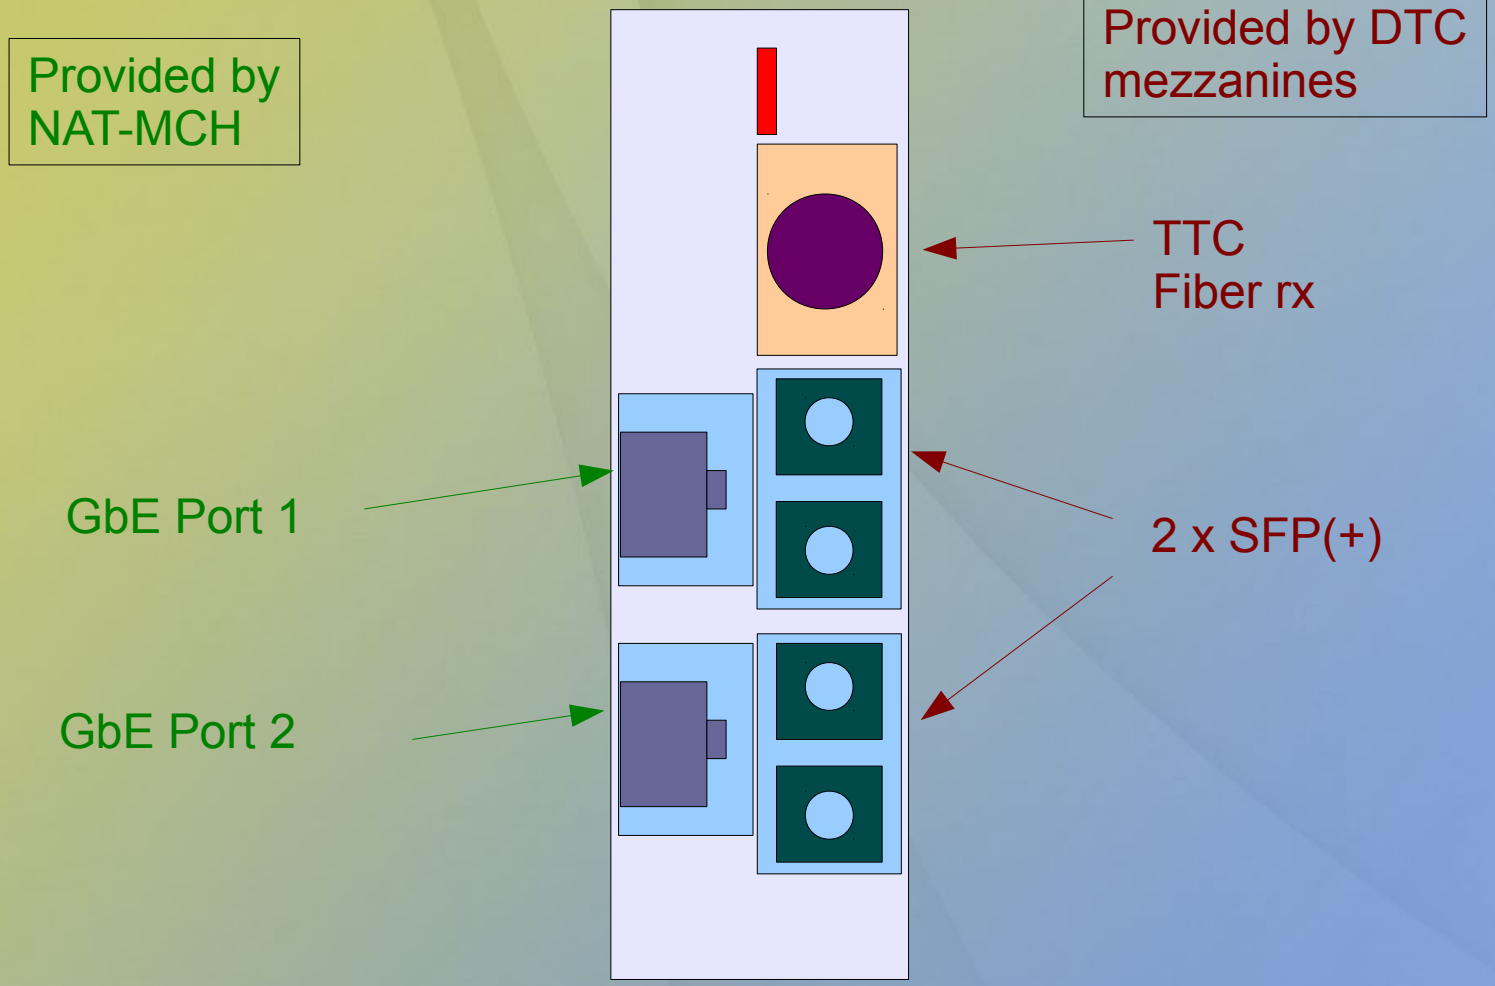
Provided by NAT-MCH
Provided by DTC mezzanines
TTC Fiber rx
GbE Port 1
2 x SFP(+)
GbE Port 2

Test Results

- Firmware and software developed for TB2010
  - Control / monitoring via IPMB
  - CMS xDAQ software support including HyperDAQ
  - TTC clock/data recovery, error-check, distribution
- BER test
  - SFP fiber and backplane link loop-back @ 4.25
  - Pseudo-random bit pattern (2^{32}-1 period)
    - No errors in ~ $5*10^{15}$ bits
  - TTC operated with extensive error monitoring for several days with 0 errors
    - TTC signal interrupted repeatedly to check synchronization – completely reliable

Planned DTC Design Updates

- Eliminate commercial MCH with unneeded features
- Eliminate extra connector board
- Allow for reduced version with timing / controls only
- Move SFP to T3 board – no fast signals on connector
- *Must operate in redundant crate with another MCH*

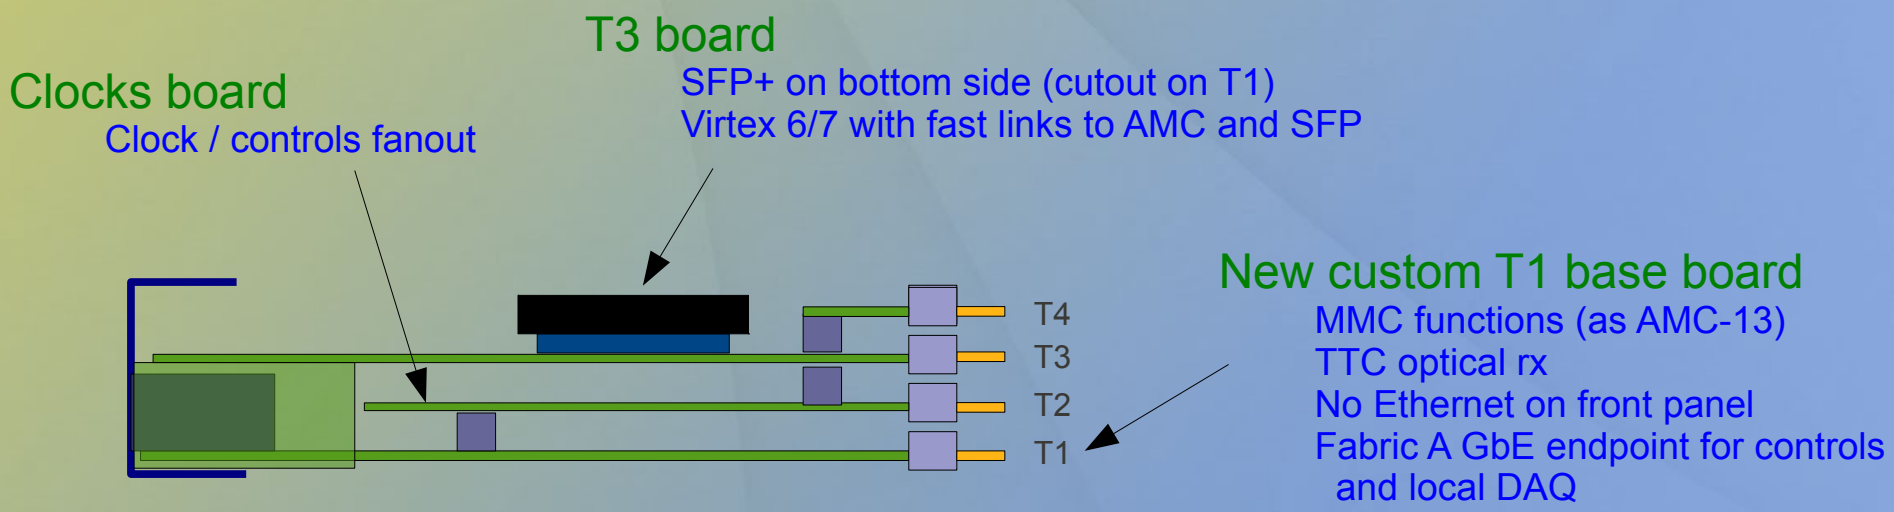

Plans

- Finish commissioning HCAL prototype system
  - DAQ path; fast links from AMCs, fiber to DAQ output
- Develop an S-Link / TTC adapter
- Develop a T1 replacement for NAT-MCH
  - Simpler design; not tied to specific vendor
  - Can accommodate some fiber I/O for i.e. timing
- Further down the road – add support for CMS upgraded DAQ (i.e. “S-Link Express”)

Reserve Slides

Legacy DAQ Adapter

One i.e. 9U VME module handles two uTCA crates, each with 400MBytes/sec out

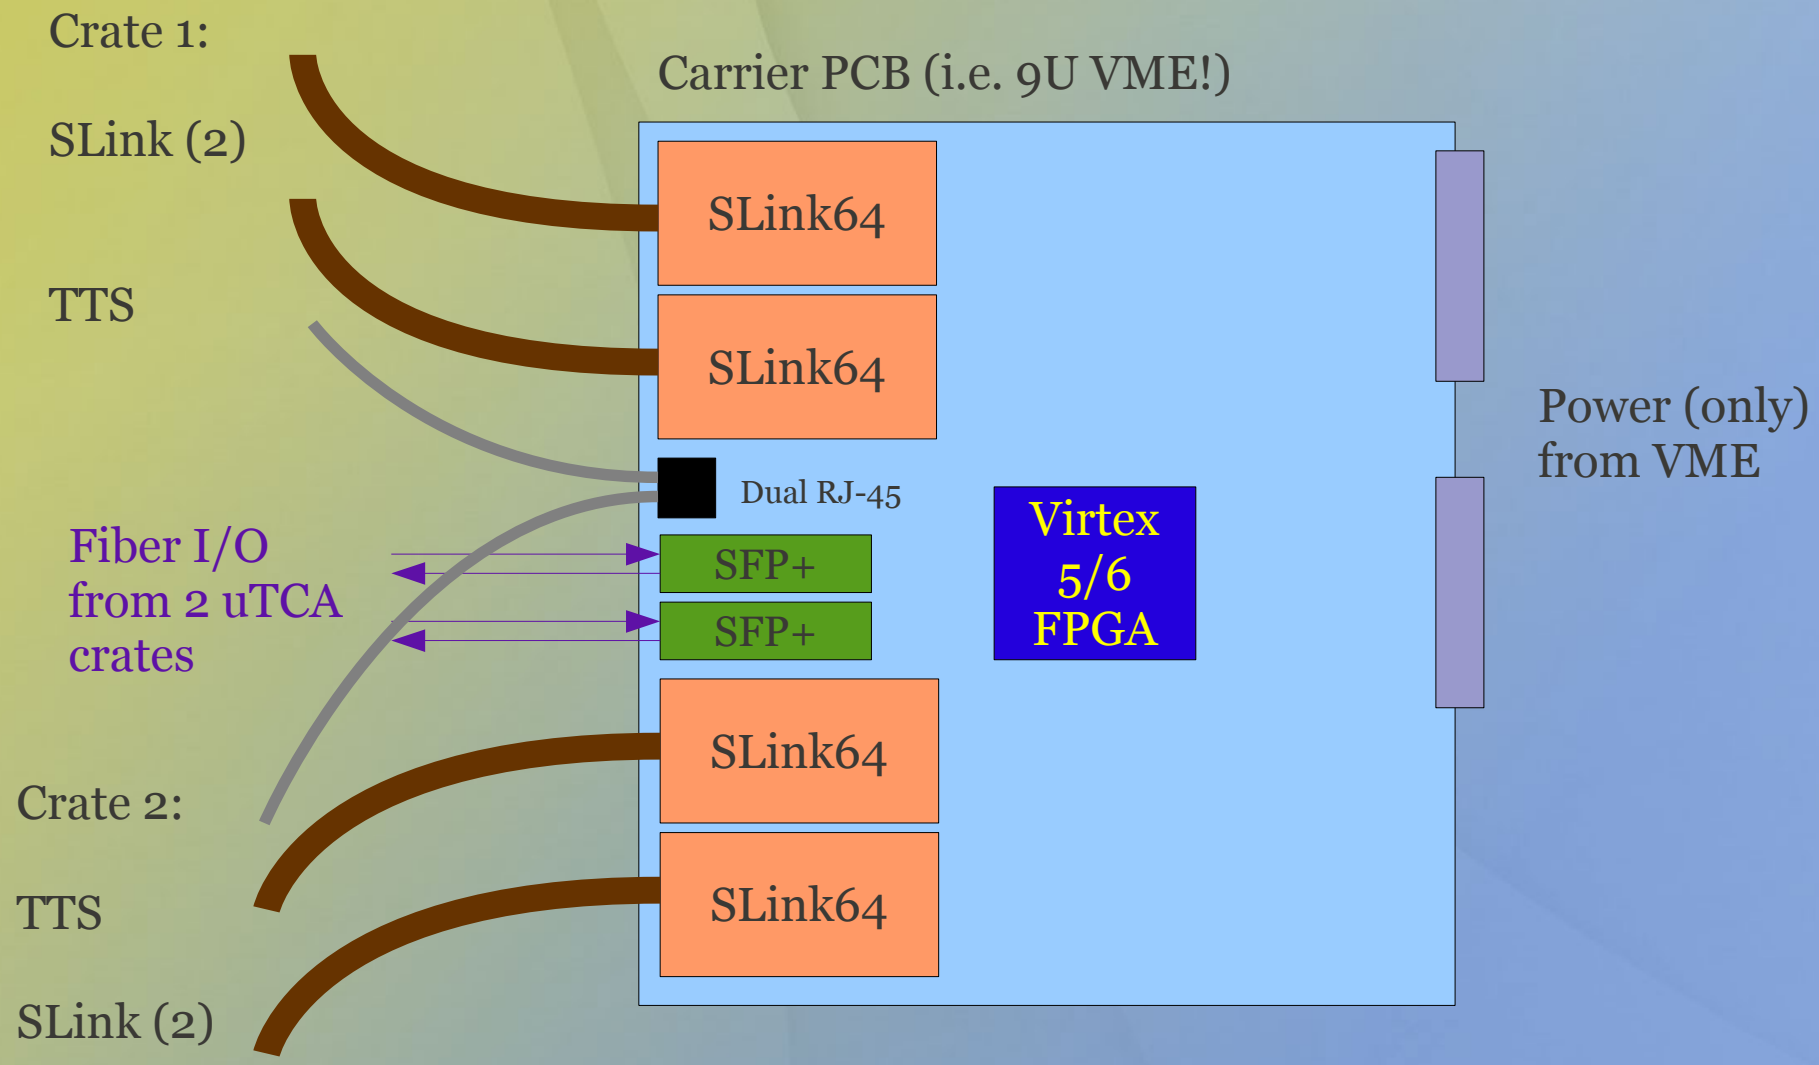

SFP/TTC board

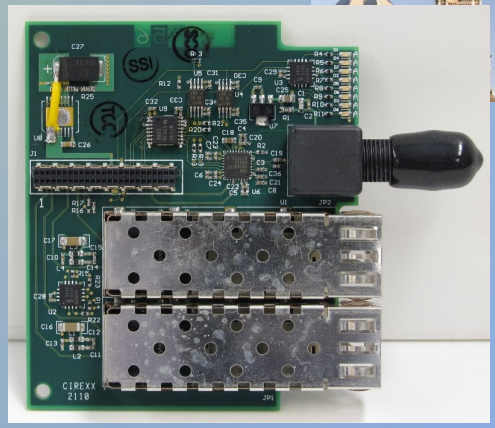

T2 (clocks) board

T3 (fabric) board

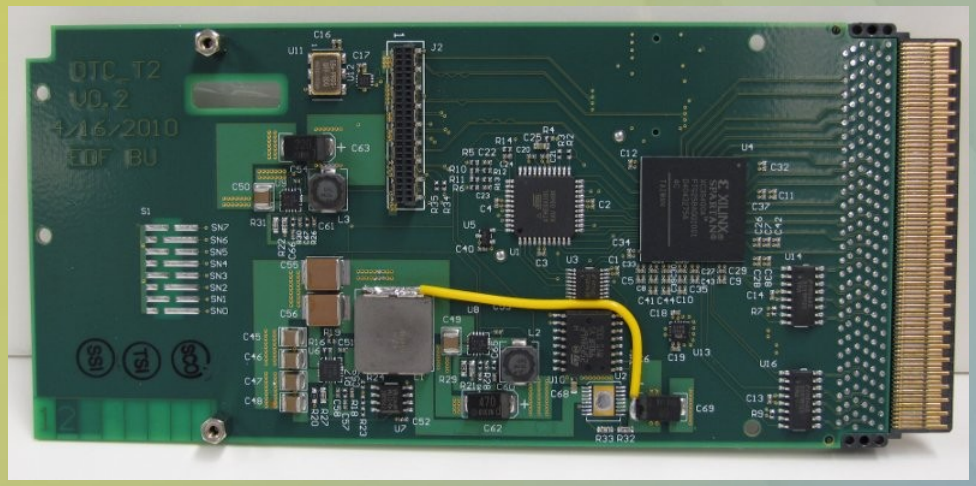

Prototype DTC (CMS-MCH)

SFP+ (2)

TTC Fiber

GbE (2)

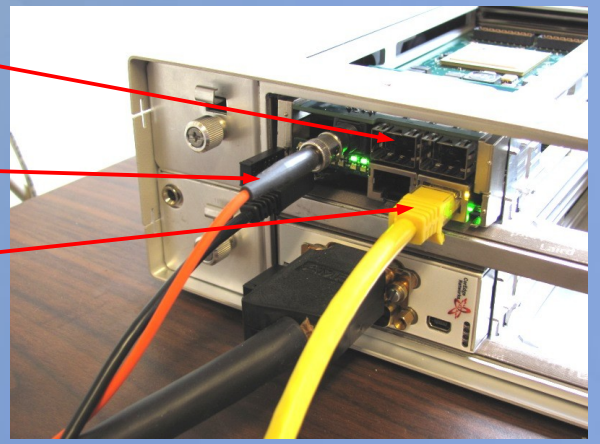

NAT-MCH Mezzanine Connector

GbE (switched)

IPMB (I2C)

Payload power

SPI (flash programming)

Note that there are many more GND on uTCA connectors