Trigger Upgrade Planning

Greg Iles, Imperial College

We've experimented with new technologies...

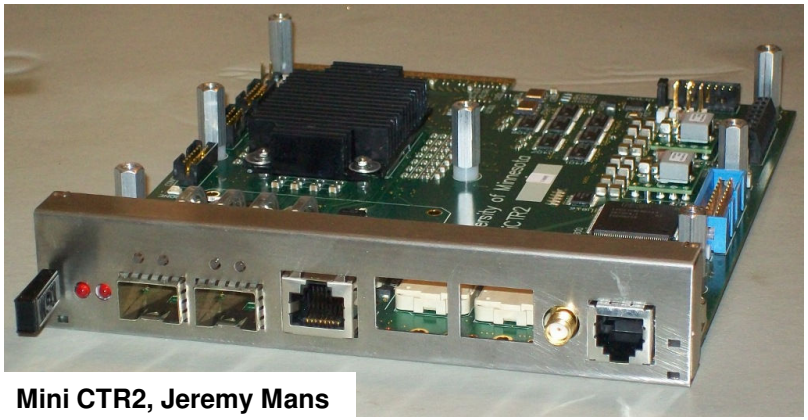

Mini CTR2, Jeremy Mans

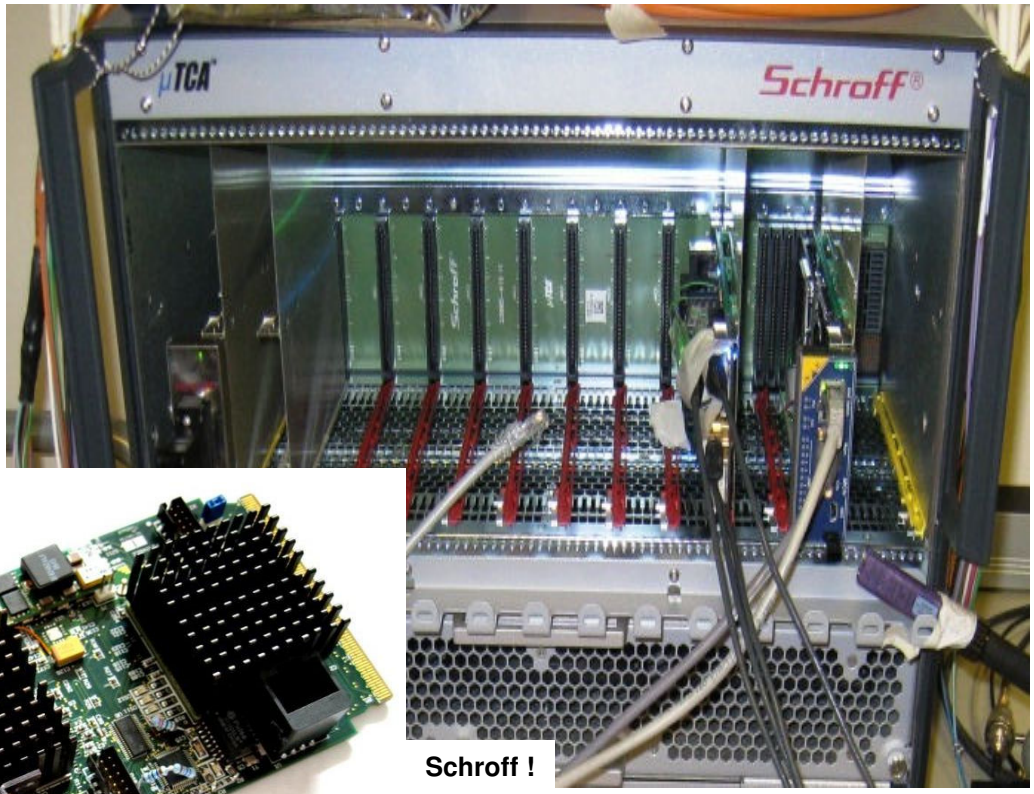

Schroff !

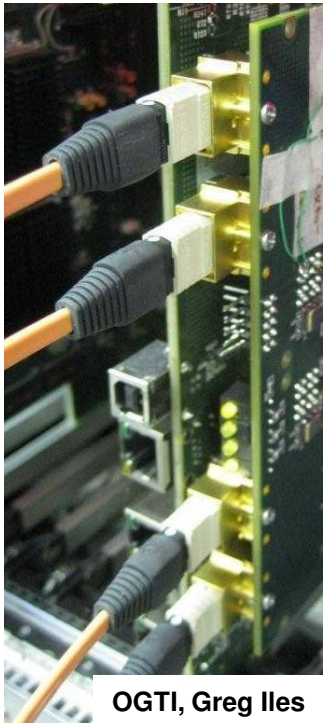

OGTI, Greg Iles

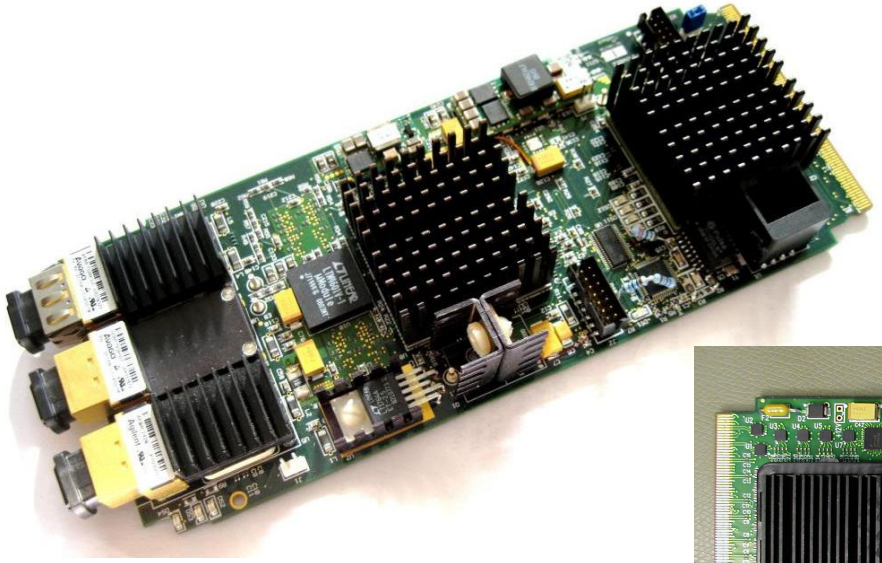

Matrix, Matt Stettler, John Jones,
Magnus Hansen, Greg Iles

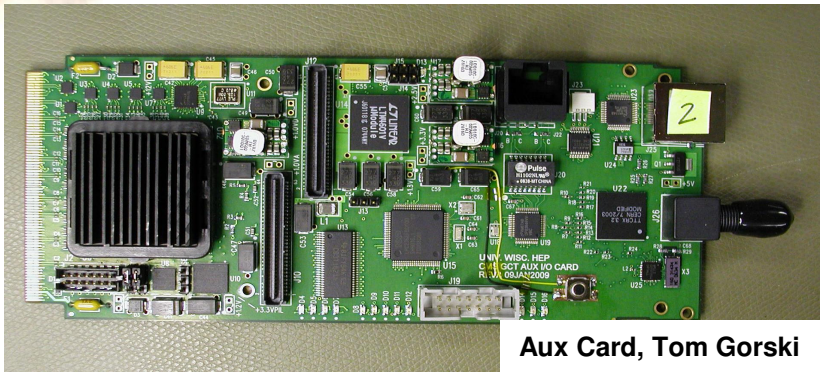

Aux Card, Tom Gorski

and discussed ideas...

Time for a plan...

Why is it urgent?

- Subsystems already designing new trigger electronics.
  - In particular HCAL, but others also have upgrade plans.
- If comprehensive plan not adopted now we will squander the opportunity to build a next generation trigger system
  - Interface between front-end, regional and global trigger is critical
  - Requirements driven by Regional Trigger (RT)

=> Would be best to try and define new Trigger now before HCAL plans fixed.

Regional Trigger Part I

How best to process ~5Tb/s?

The initial idea was to have a single processing node for all physics objects

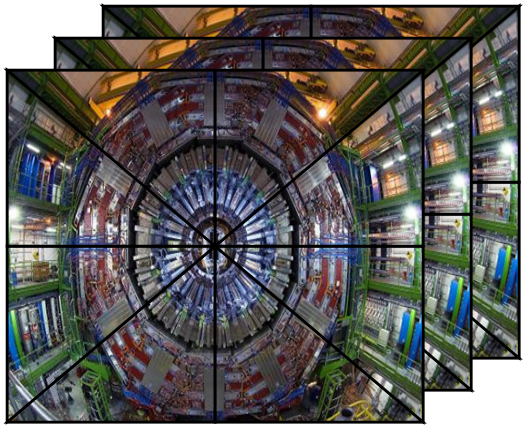

Red boundary defines single FPGA

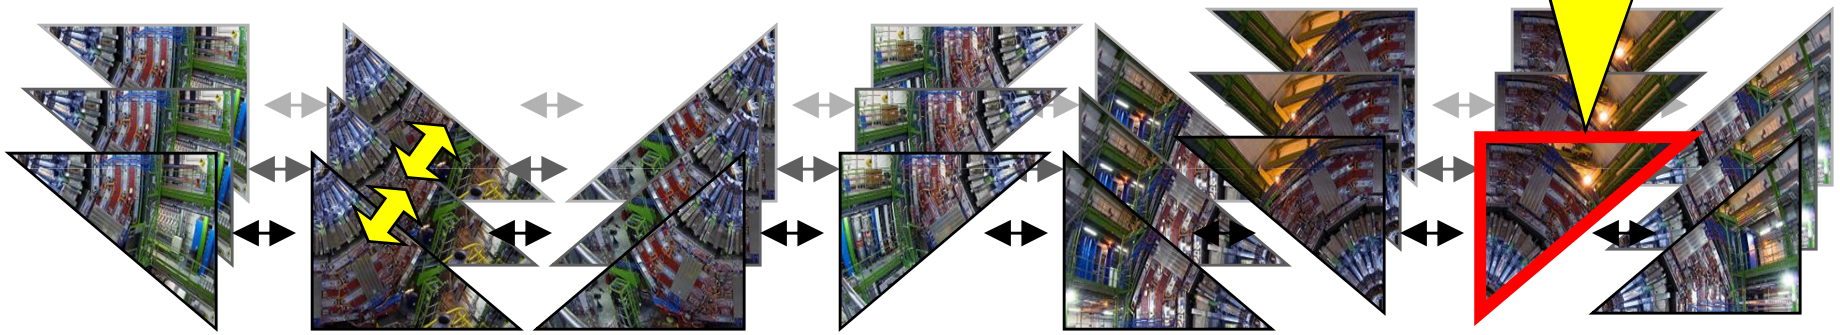

Single processing node:

- Share data in both dimensions to provide complete data for objects that cross boundaries

Very flexible, but Jets can be large ~15% of image in both η and φ

Very large data sharing required => Very inefficient => Big system

Regional Trigger Part II

Alternative solutions:

- Split processing into 2 stages
  - Fine procesing for electrons / tau jets
  - Coarse processing for jets
- Pre cluster physics objects
  - Build possible physics objects from data available then send to neigbours for completion
  - More complex algorithms. Less flexible.
- Both currently applied in CMS
  - But believe we have a better solution...

Time Multiplexed Trigger - TMT

Time Multiplexed Trigger

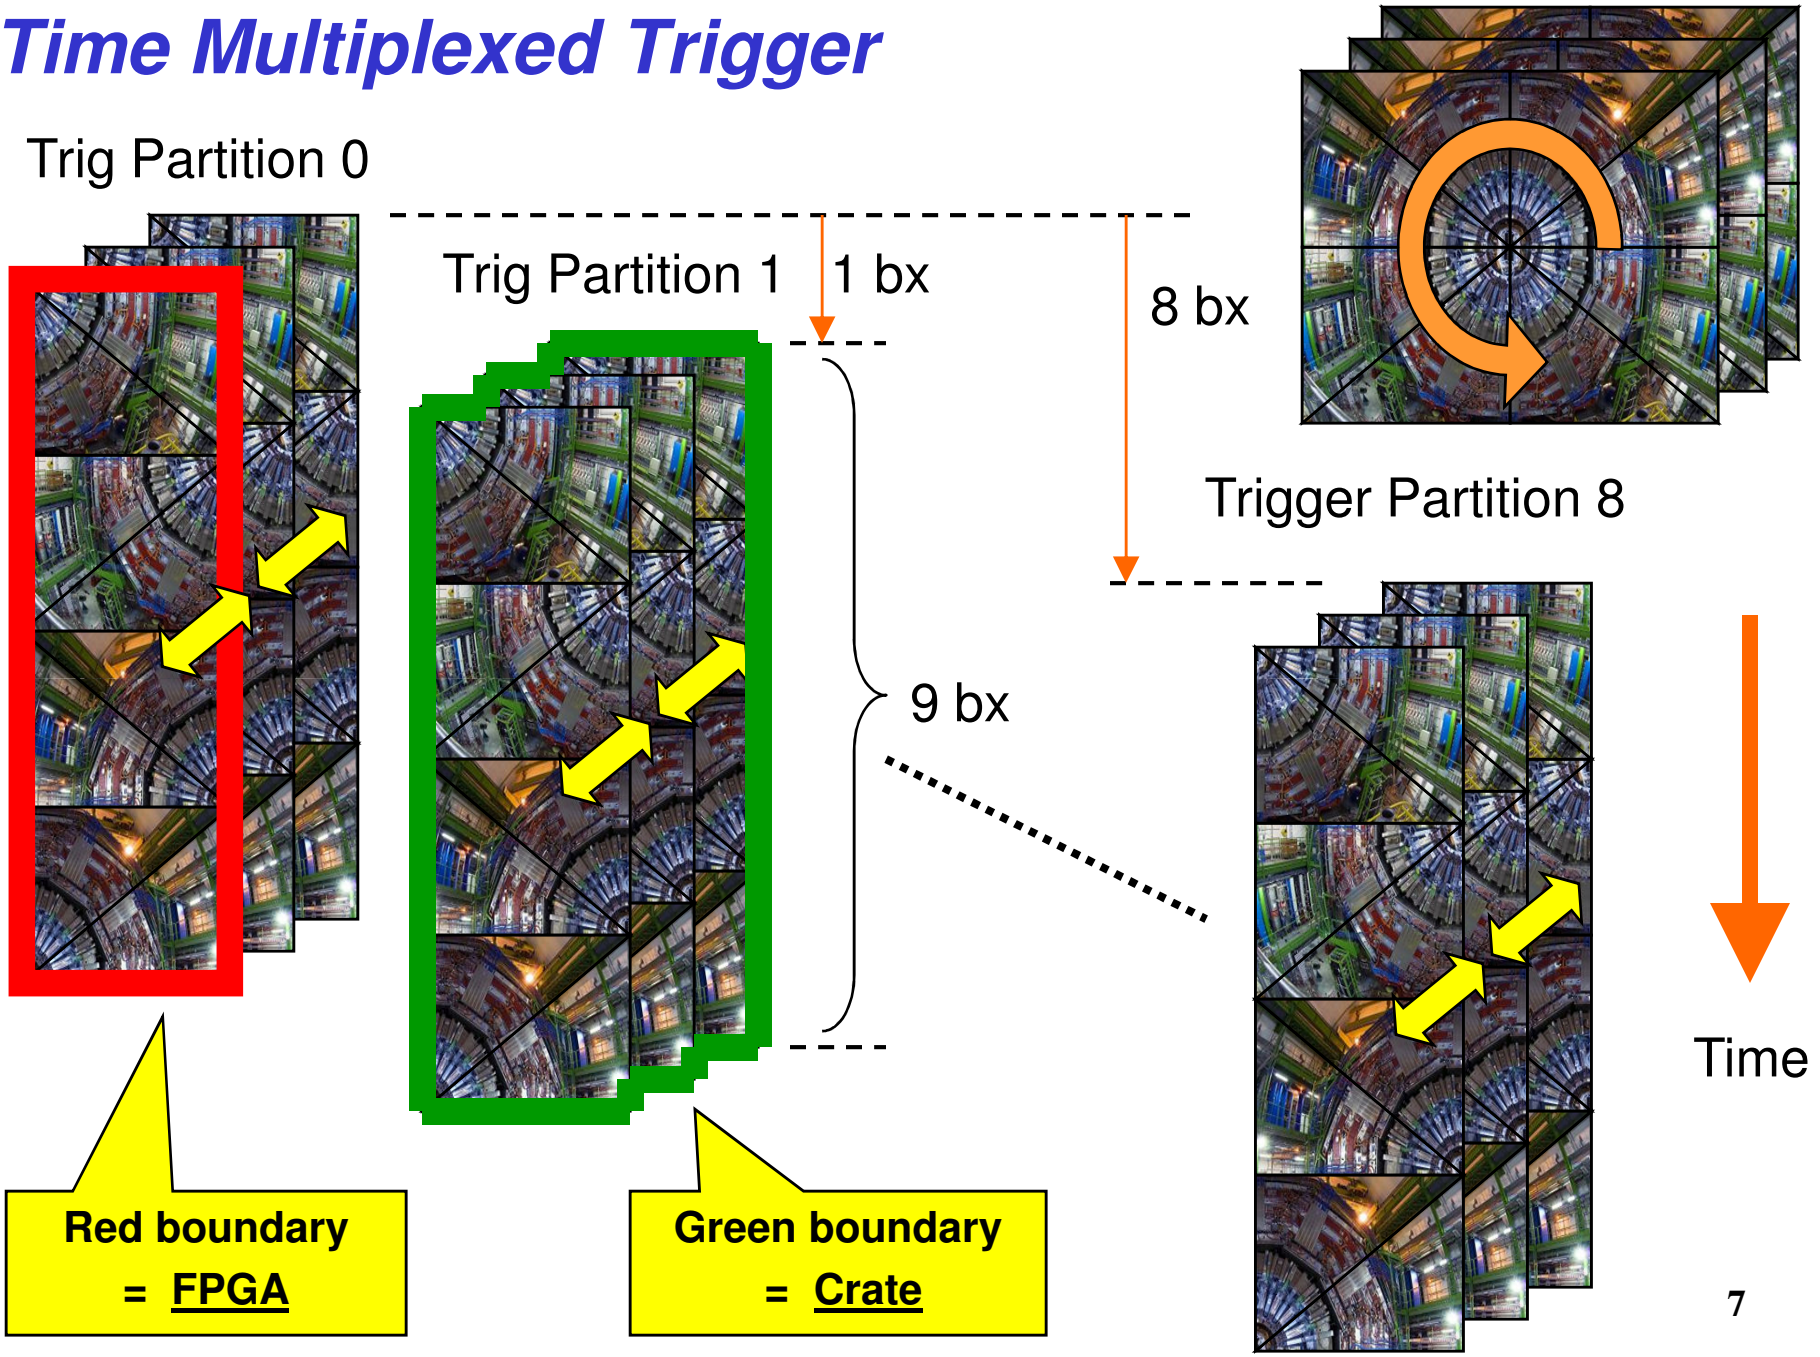

The Trigger Partition Crates

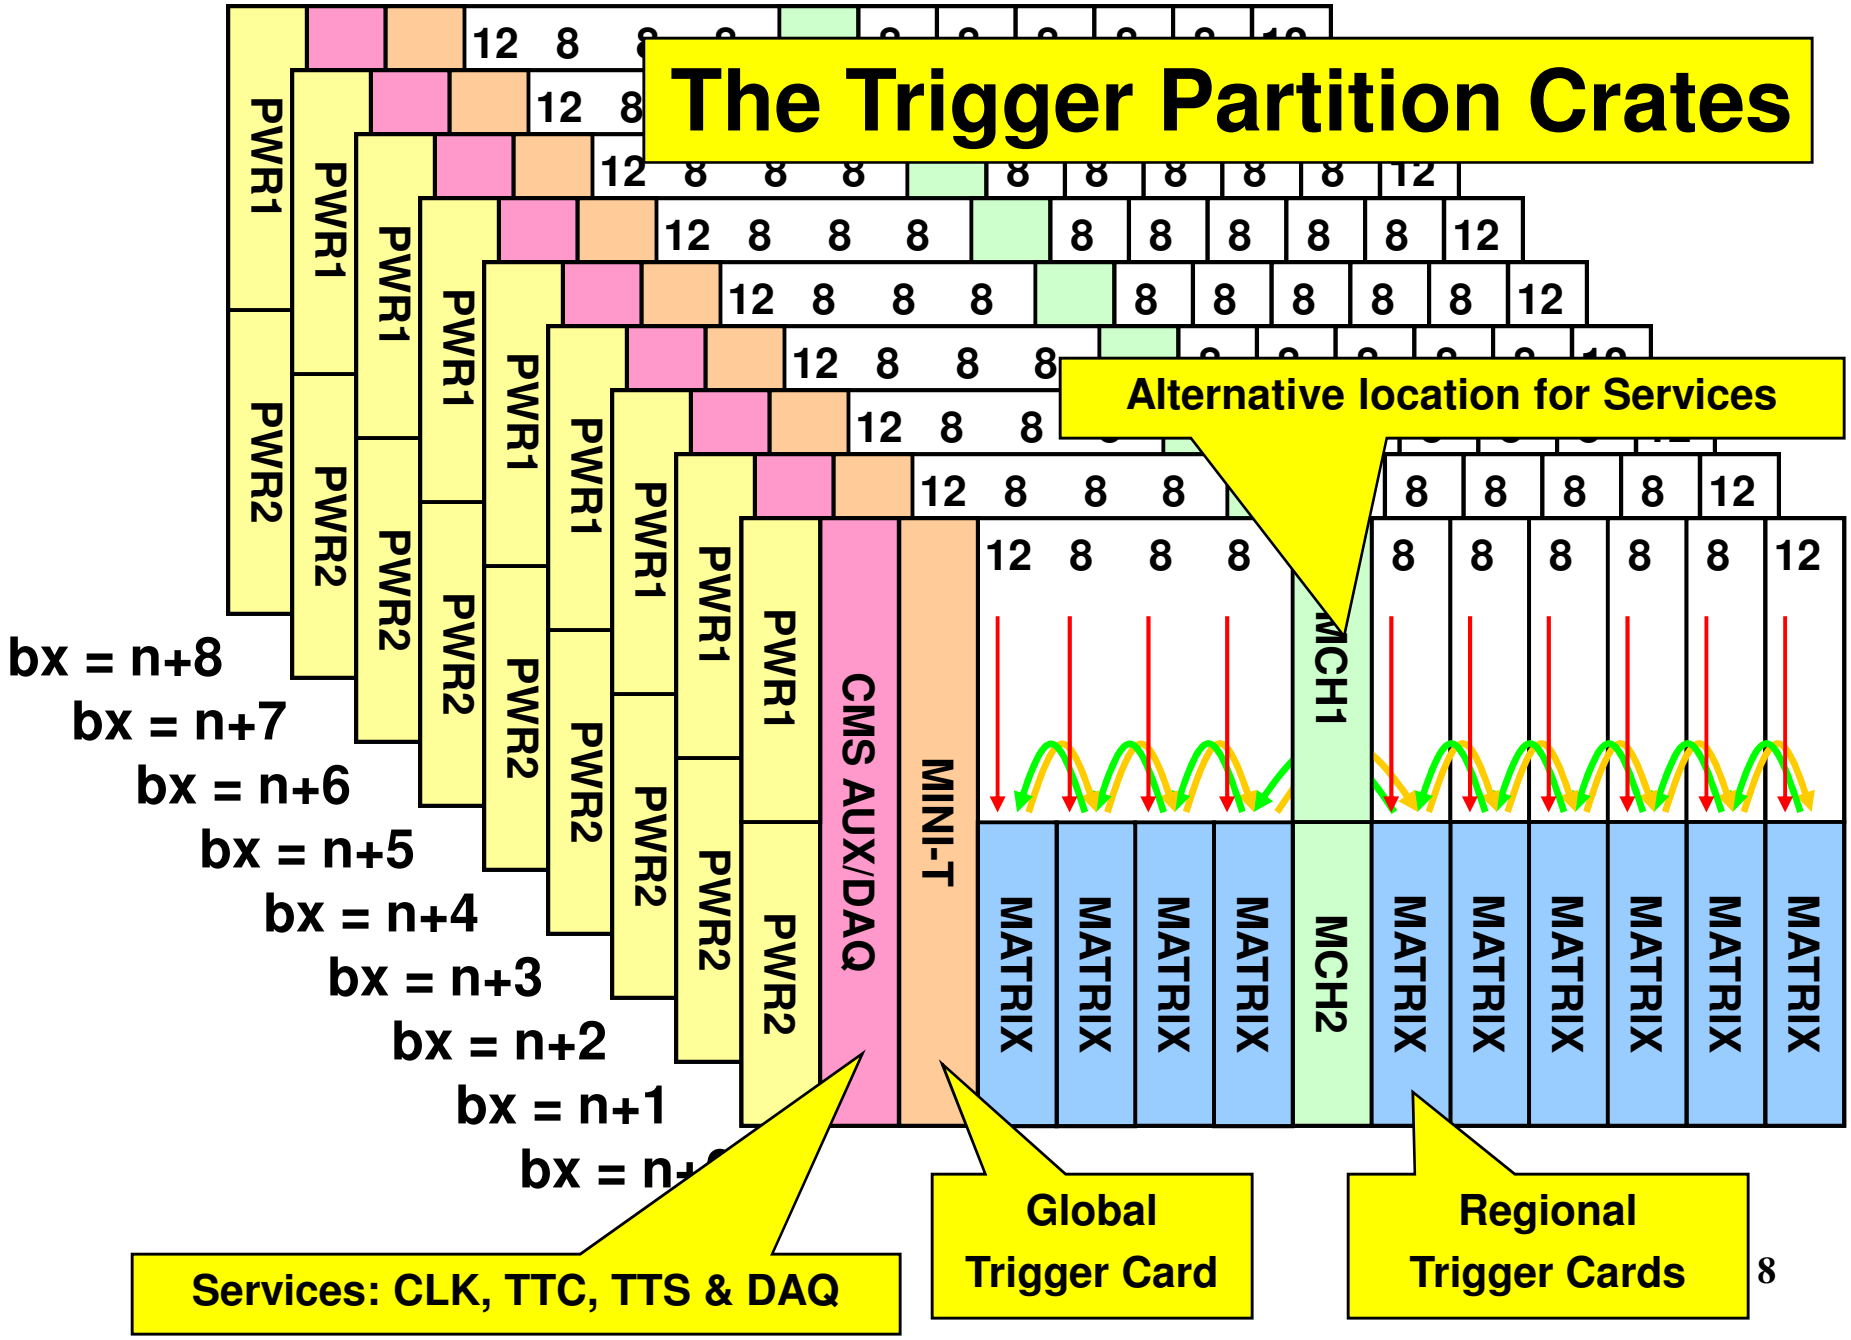

Regional Trigger Requirements:

- Use Time Multiplexed Trigger (TMT) as baseline design
  - Flexible
    - All HCAL, ECAL, TK & MU data in single FPGA
  - Minimal boundary constraints
  - Can be spilt into seprate partitions for testing
    - e.g. ECAL, HCAL MU, TK can each have a partition at the same time
  - Simple to understand
- Redundant
  - Loss of a partition results in loss of max 10% of the trigger rate
  - Can easily switch to back up partition (i.e. software switch)
- But what constraints does this impose on TPGs...

Services

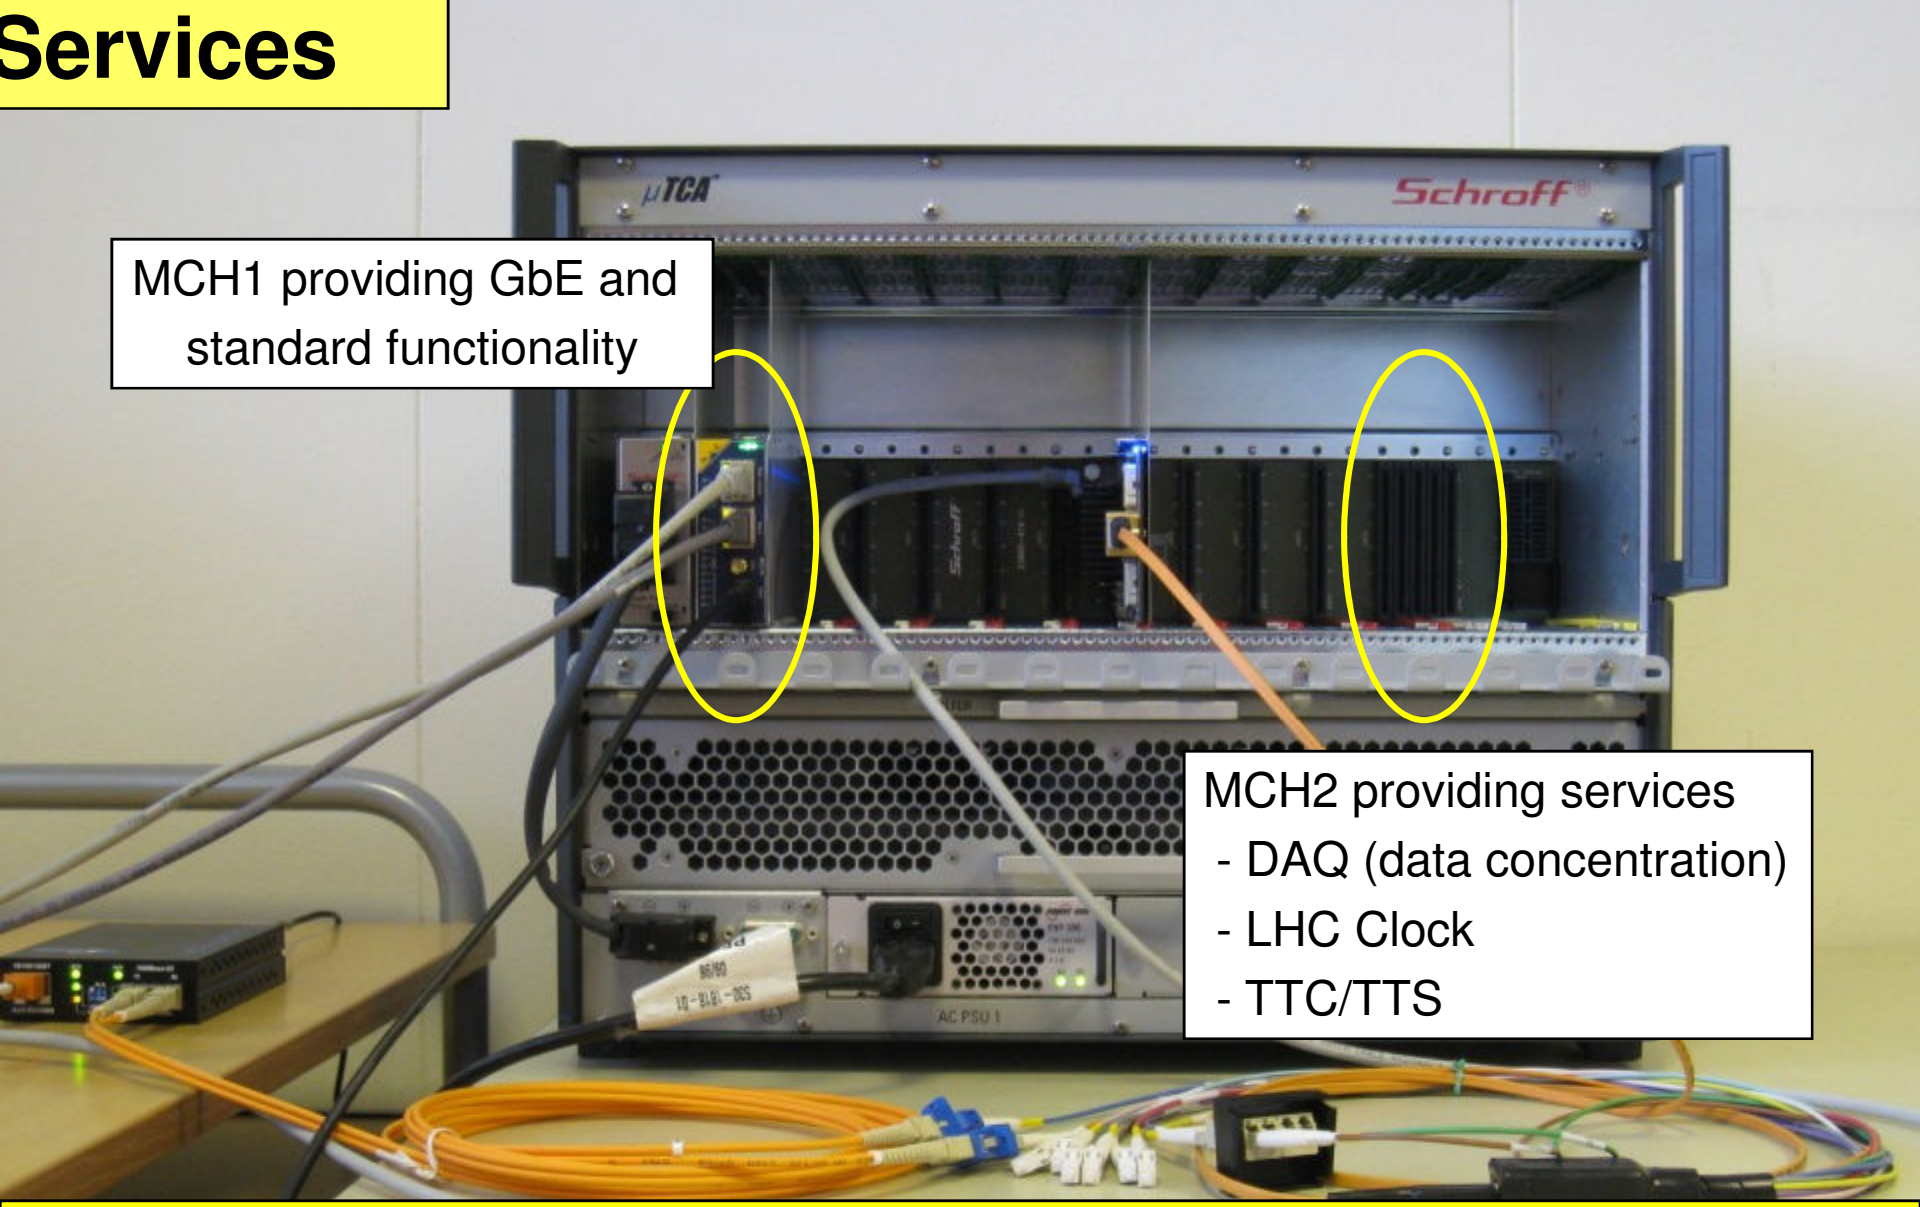

See also:

DAQ/Timing/Control Card (DTC), Eric Hazen,

uTCA Aux Card Status - TTC+S-LINK64, Tom Gorski

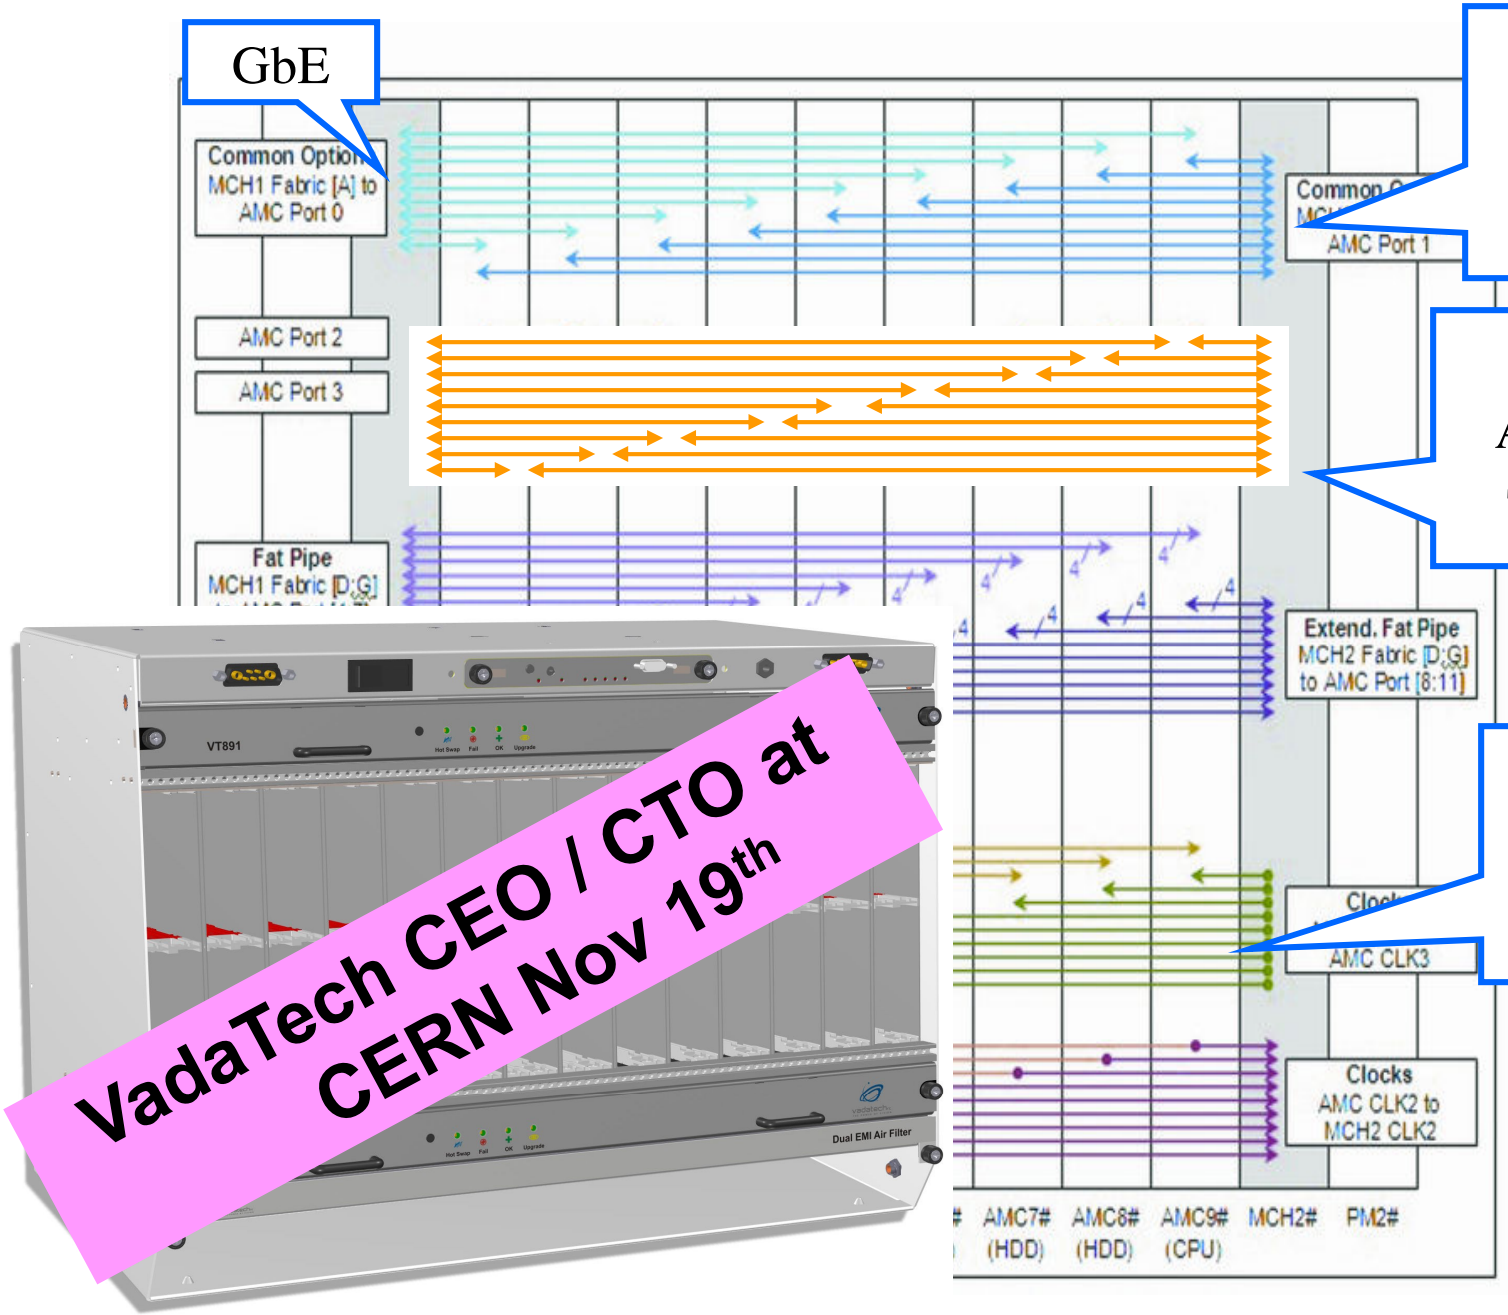
GbE
Common Option
MCH1 Fabric [A] to
AMC Port 0
AMC Port 2
AMC Port 3
Fat Pipe
MCH1 Fabric [D:G]
Common
AMC Port 1
Extend. Fat Pipe
MCH2 Fabric [D:G]
to AMC Port [8:11]
Clock
AMC CLK3
Clocks
AMC CLK2 to
MCH2 CLK2
AMC7#
(HDD)
AMC8#
(HDD)
AMC9#
(CPU)
MCH2#
PM2#
VT891
Dual EMI Air Filter
VadaTech CEO / CTO at
CERN Nov 19th
Tongue 1:
AMC port 1
DAQ
Tongue 2:
AMC port 3
TTC / TTS
Tongue 2:
AMC Clk3
LHC Clock

CMS MCH: Service

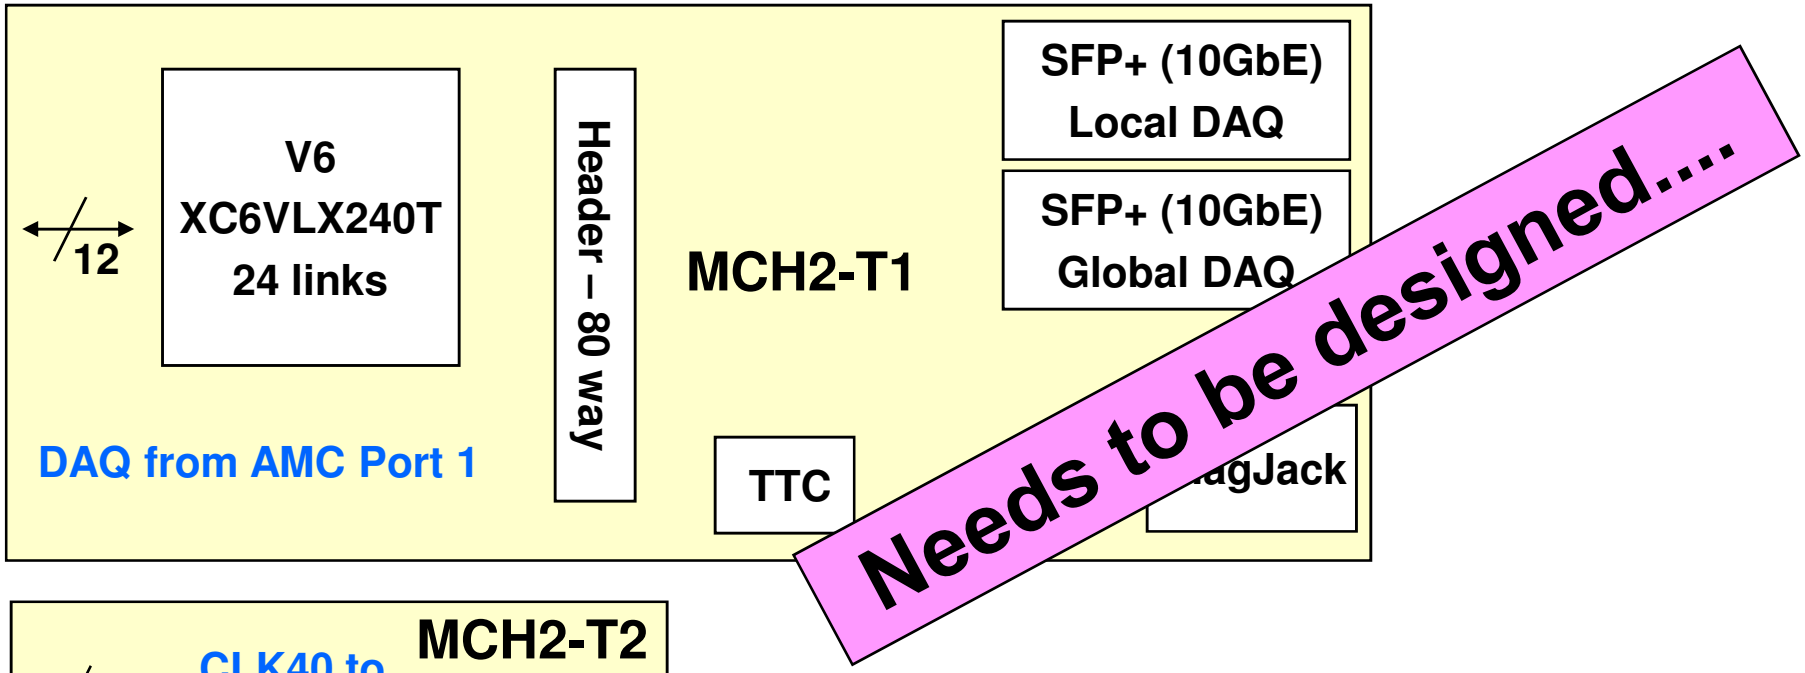

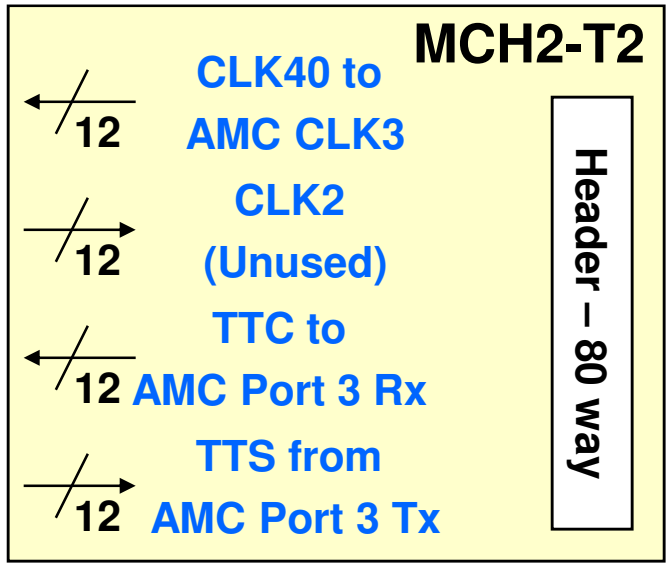

Advantages:

- All services on just one card
- High Speed Serial on just Toungue 1
- Dual DAQ outputs
- Tounges 3/4 (fat, ext-fat pipes) spare for other apps

Installation scenario

- Ideally wish to install, commission and test new trigger system in parallel with existing trigger

- Following is an example of how this might happen

  - HCAL Upgrade + Test Time Multiplexed Trigger
  - ECAL Upgrade
  - Full Time Multiplexed Trigger System
  - Incorporate TK + MU Trigger

Current System

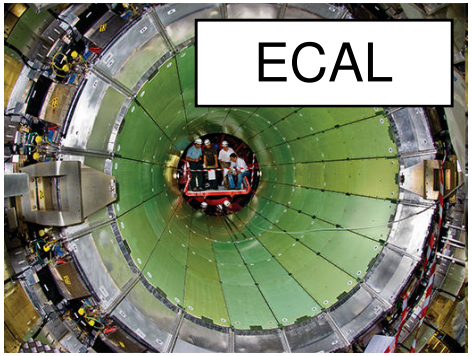

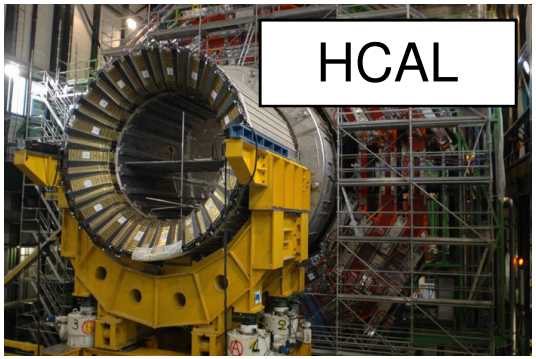

ECAL TPG

4x 1.2Gb/s

Copper

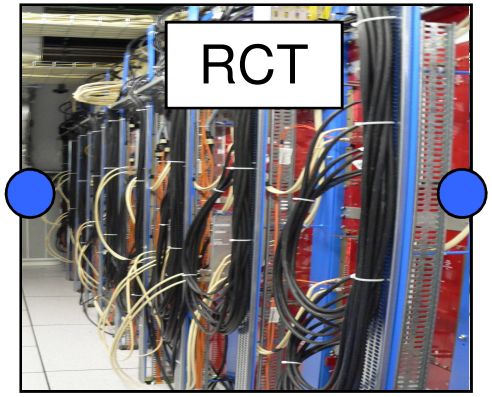

4x 1.2Gb/s

HCAL TPG

To GCT and GT

Stage 1: HCAL Upgrade

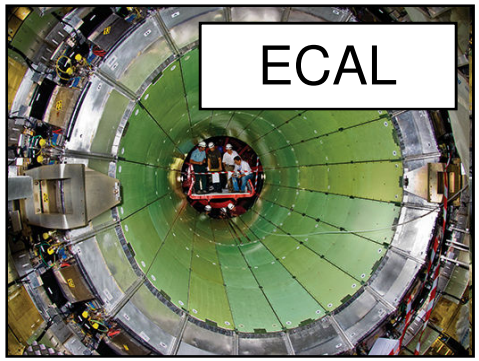

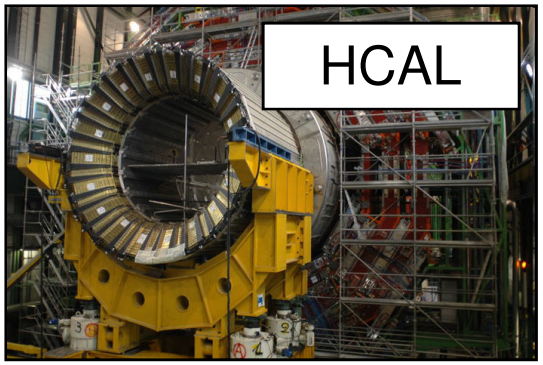

ECAL TPG

4x 1.2Gb/s

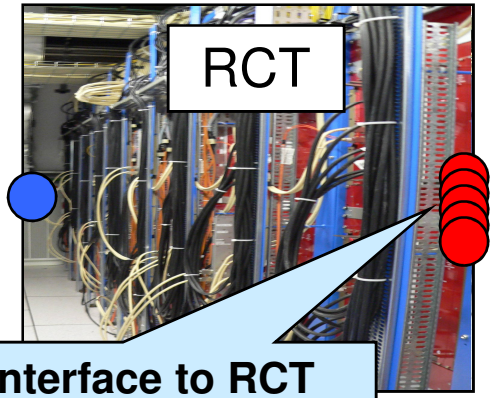

Optical interface to RCT

Many 4.8Gb/s

Fibre

HCAL TPG

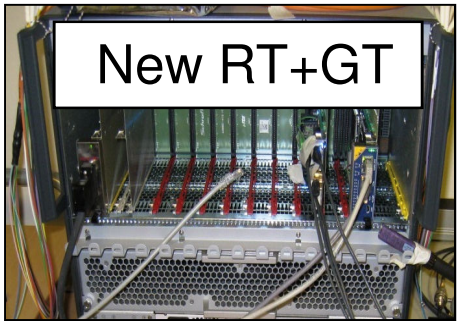

Time multiplexed data
1/10 to 1/20 of events

Tech Trigger to existing GT

Stage 2: ECAL Upgrade

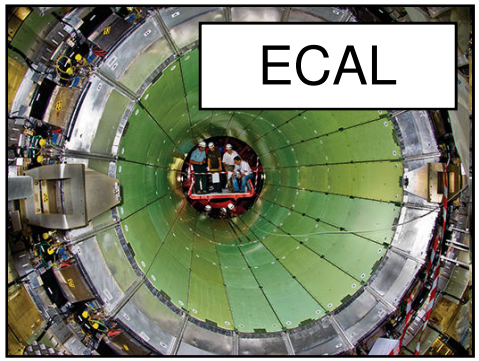

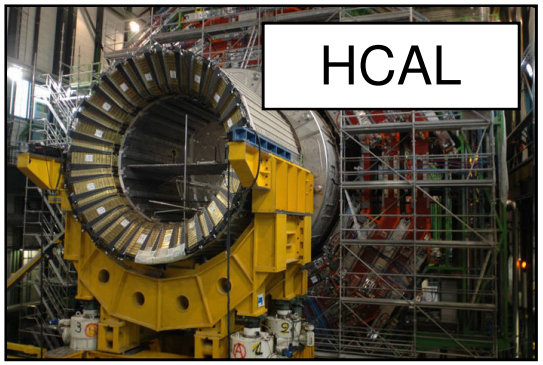

ECAL TPG

4.8 Gb/s

Duplicate

Other ECAL TPGs

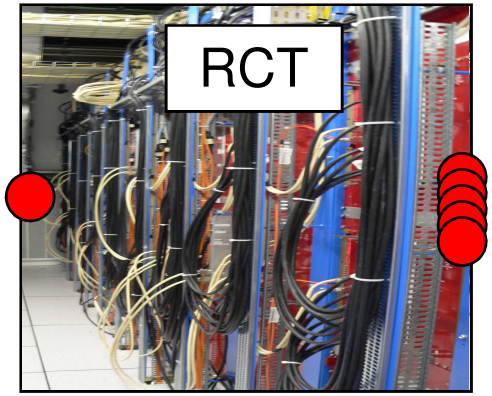

Many 4.8Gb/s

HCAL TPG

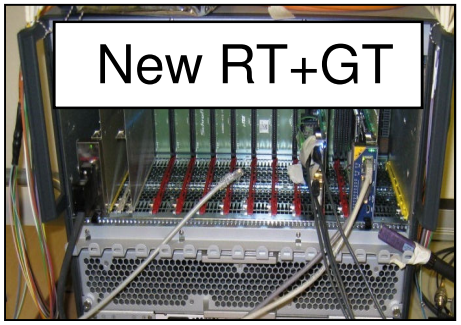

Time multiplexed data
1/10 to 1/20 of events

ECAL TM

Time Multiplexer Rack

Tech Trigger to existing GT

Stage 3:
Upgrade to New Trig

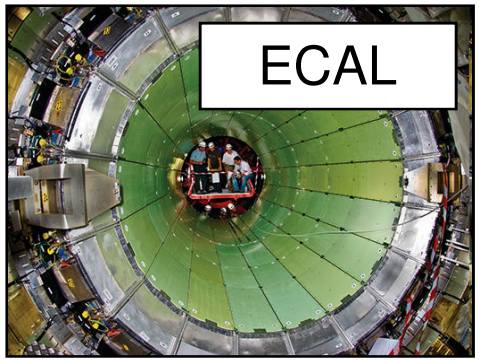

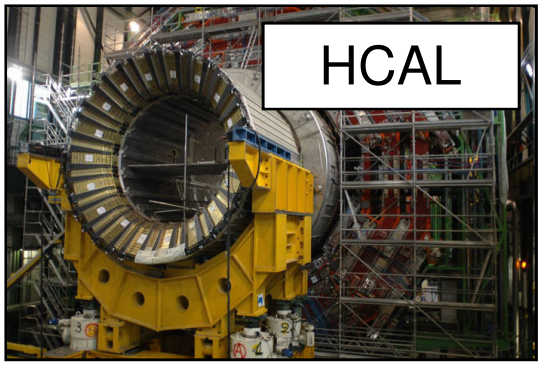

ECAL TPG

Other ECAL TPGs

Duplicate

ECAL TM

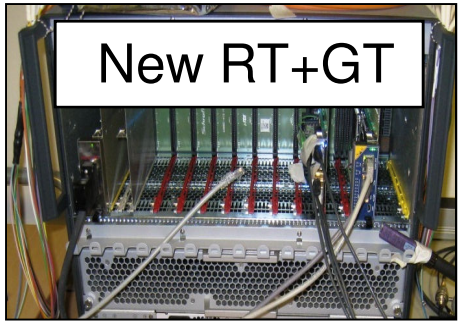

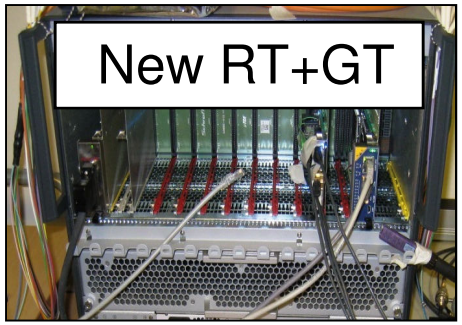

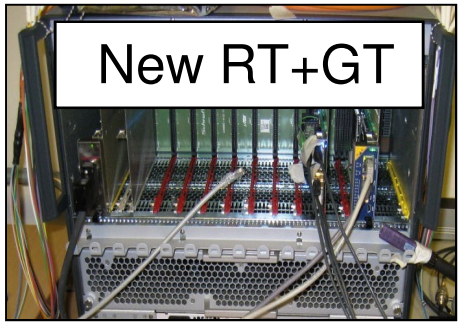

HCAL TPG

Stage 4:
Add Tk + Mu Trig

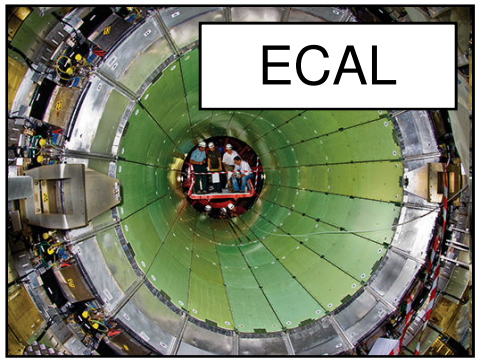

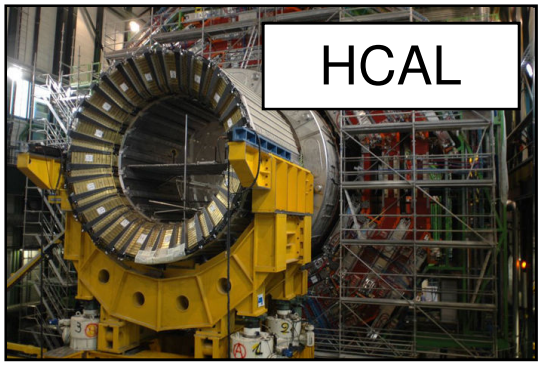

ECAL TPG

Other ECAL TPGs

Duplicate

ECAL TM

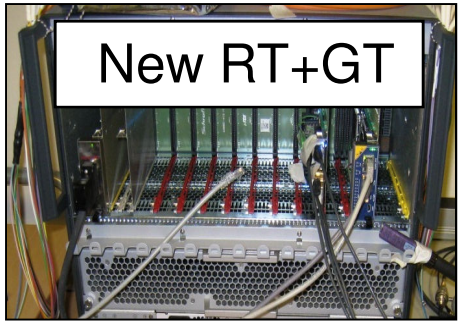

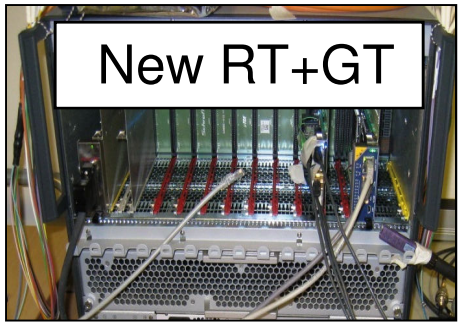

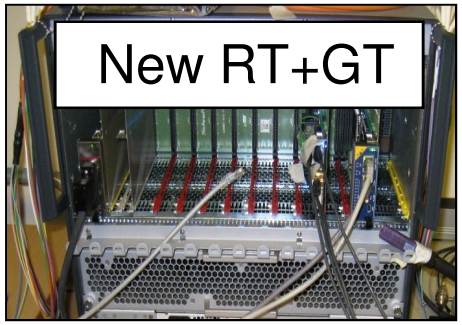

HCAL TPG

MU TPG

TK TPG

Technology choice.. Use V6 not V5

	XC5VTX150T	XC5VTX240T	XC6VLX550T
Links	40 @ 5.0Gb/s	48 @ 5.0Gb/s	36 @ 6.5Gb/s
Slices (k)	23	37	86
Logic Cells (k)	148	240	550
CLB FlipFlops (k)	92	150	687
Distributed RAM (kbits)	1,500	2,400	6,200
BRAM (36kbits)	228	324	632
Cost	N/A	4.7k	4.5k

A single Virtex-6 FPGA CLB comprises two slices, with each containing four 6-input LUTs and eight Flip-Flops (twice the number found in a Virtex-4 FPGA slice), for a total of eight 6-LUTs and 16 Flip-Flops per CLB.

A single Virtex-5 CLB comprises two slices, with each containing four 6-input LUTs and four Flip-Flops (twice the number found in a Virtex-4 slice), for a total of eight 6-LUTs and eight Flip-Flops per CLB

Plan for the nexy 5 years:

- Technical Proposal
  - Review of Proposal
- Resources Required
  - Financial / Manpower / Time

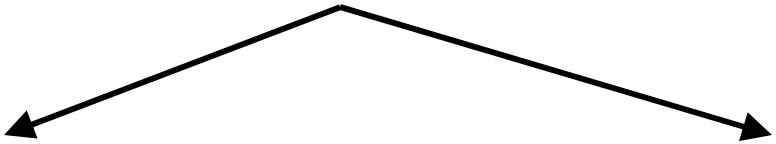

- Hardware development
  - Prototypes
    - Procesing Card
    - Service Card
    - Optical SLBs
    - Custom Crate
  - Cooling
    - uTCA crates have front to back airflow

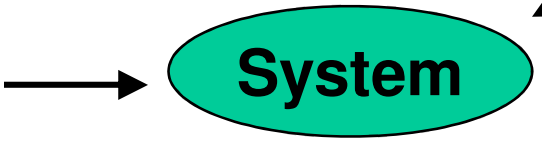

- Firmware
  - Serial Protocol
  - Slow Control
  - Algorithms
- Software
  - Low level drivers
  - Communication
  - Interfaces to legacy systems

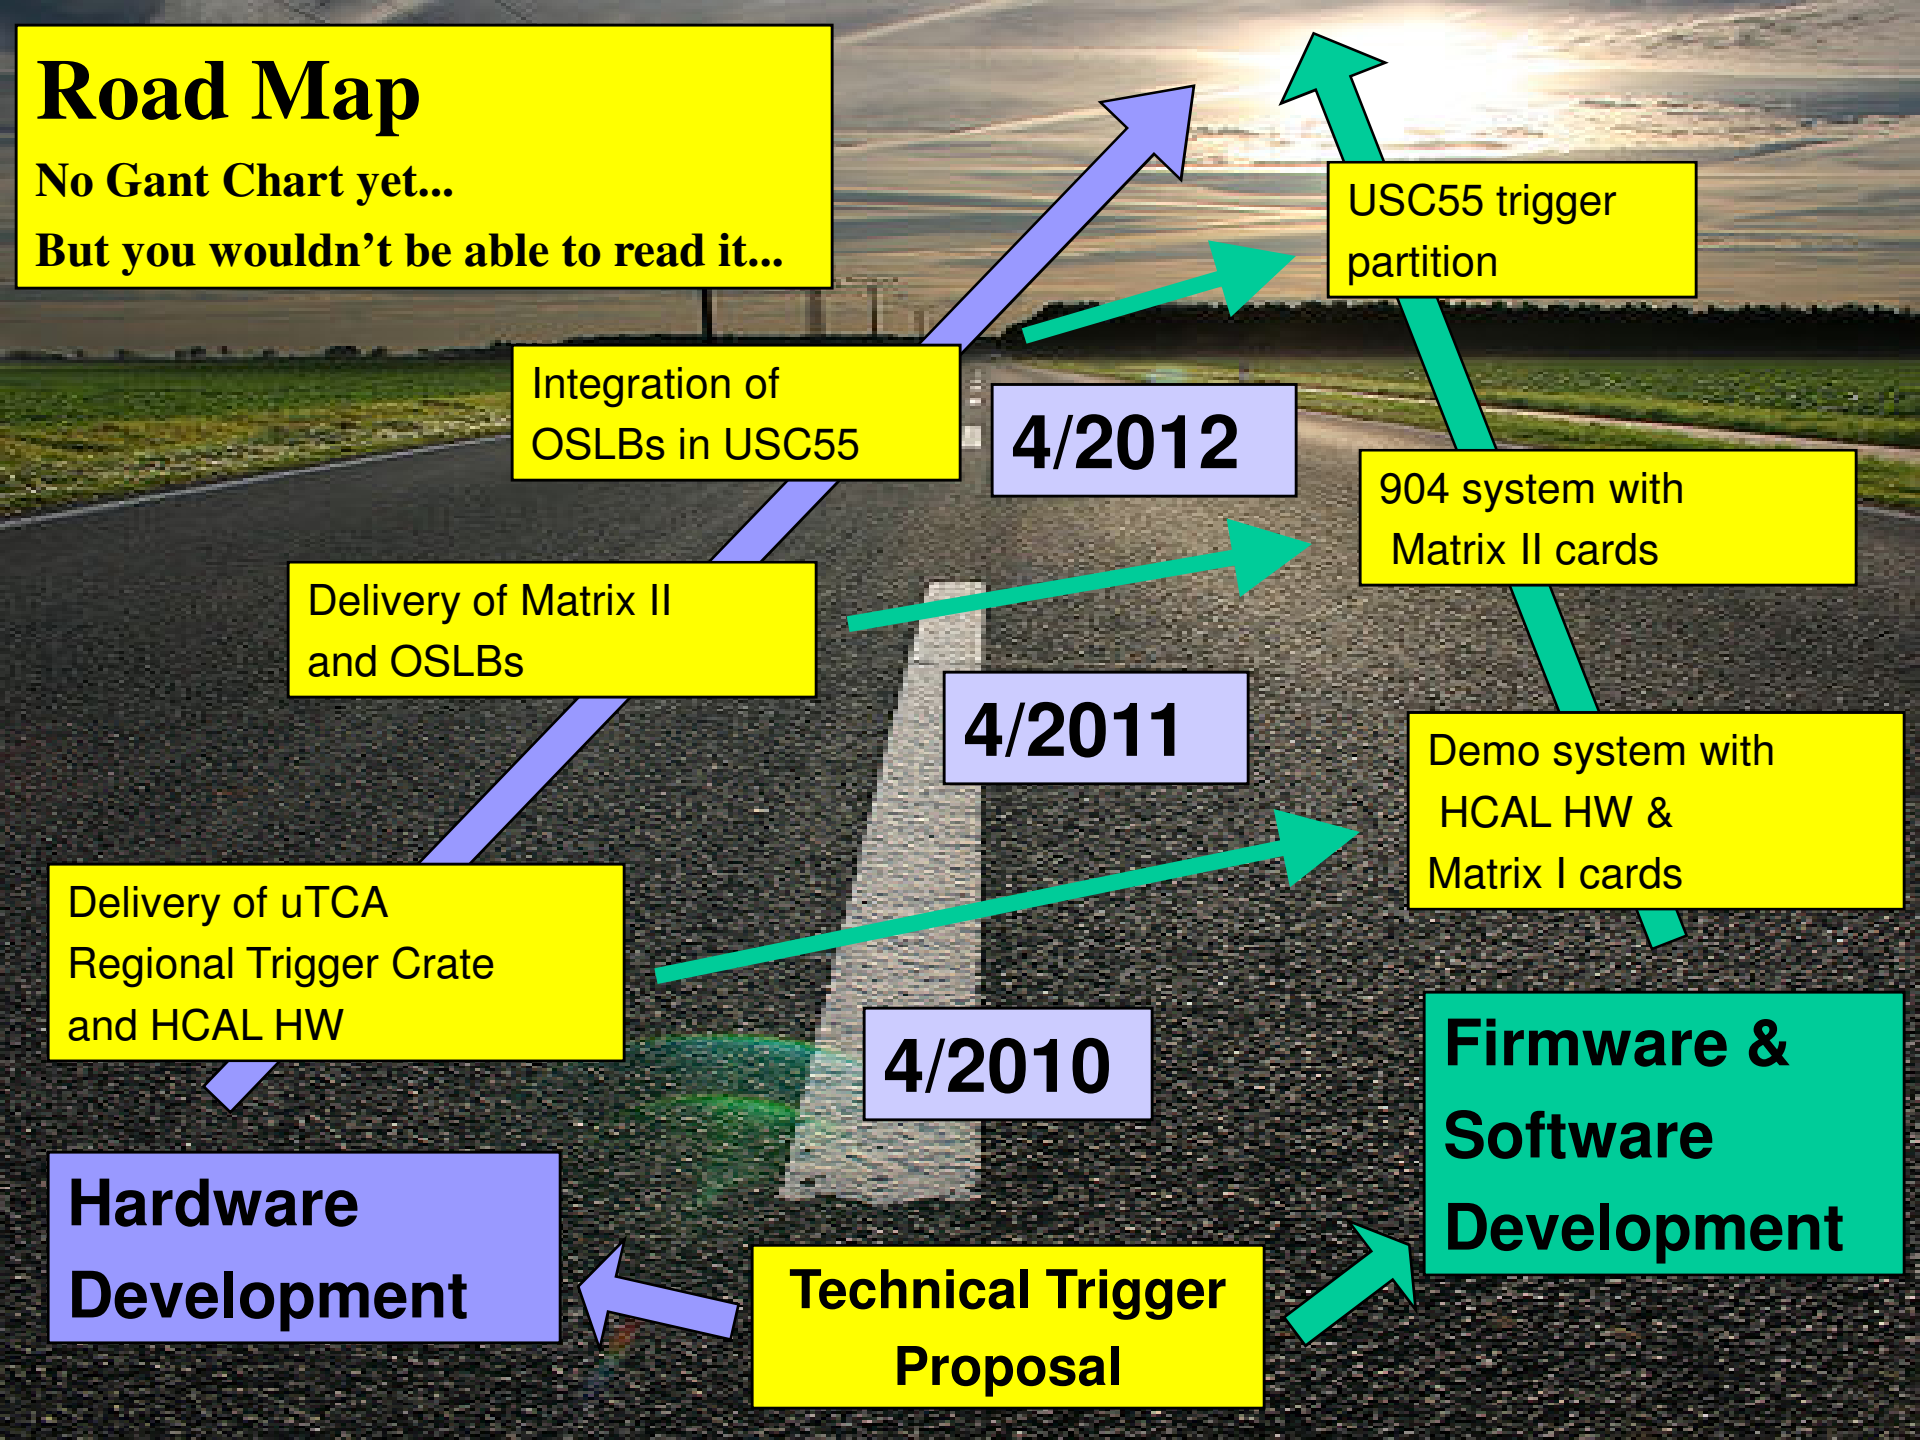
Road Map
No Gant Chart yet...
But you wouldn't be able to read it...
USC55 trigger partition
Integration of OSLBs in USC55
4/2012
904 system with Matrix II cards
Delivery of Matrix II and OSLBs
4/2011
Demo system with HCAL HW & Matrix I cards
Delivery of uTCA Regional Trigger Crate and HCAL HW
4/2010
Firmware & Software Development
Hardware Development
Technical Trigger Proposal

=> Trigger Upgrade Proposal by April 2010

End of Part I

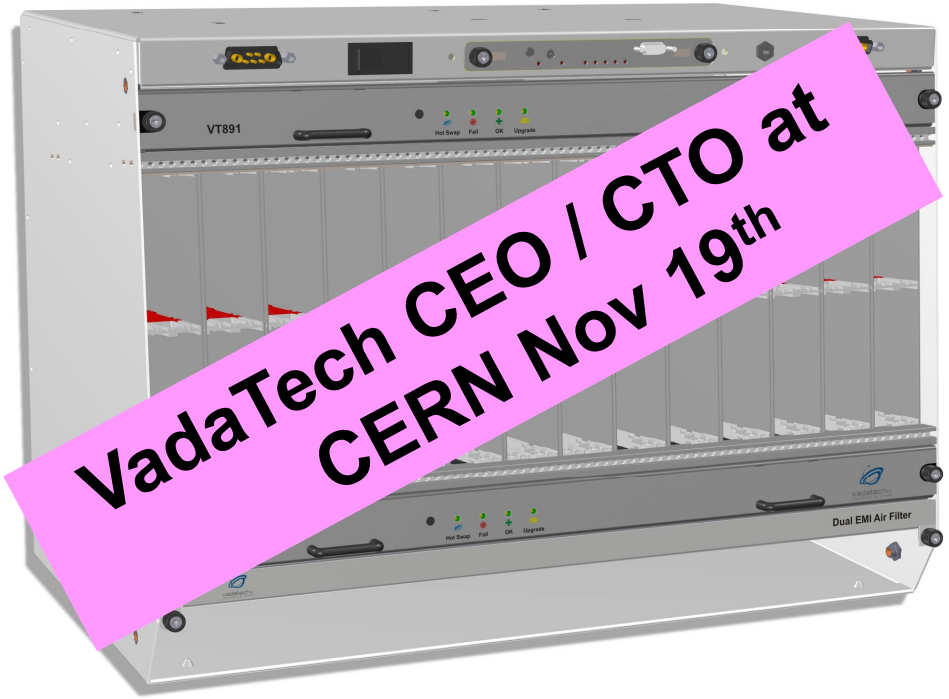

The next bit is confusing so before I lose you...

Implications for on TPGs

Can we put time multiplexing inside TPG FPGA?

- Maximum number of trigger partitions limited by latency.
  - Assume between 9 and 18 partitions
- Hence each TPG requires a minimum of 9 outputs
  - Each fibre carrying 8 towers/bx x 9 bx = 72 towers
  - Transmist 4 towers in η over ¼ of φ
    - Assumed time multiplex in φ, but η also possible
- But HCAL TPG designed to generate 36-40 towers
  - This is very impresive, but still need factor of 2
  - Revist TPG idea...

TPG design

Assumed TPG is single FPGA
Why: simplicity, board space, power,

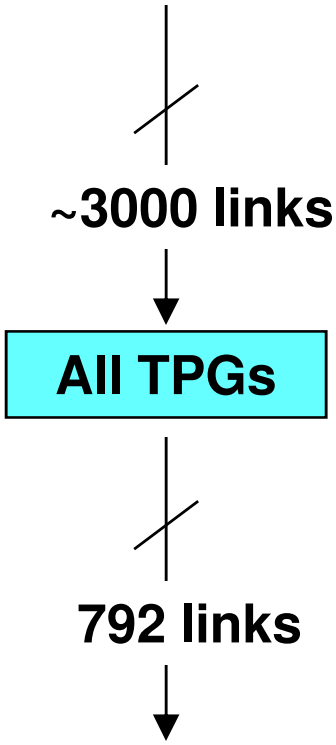

Approx
4:1 input to output ratio

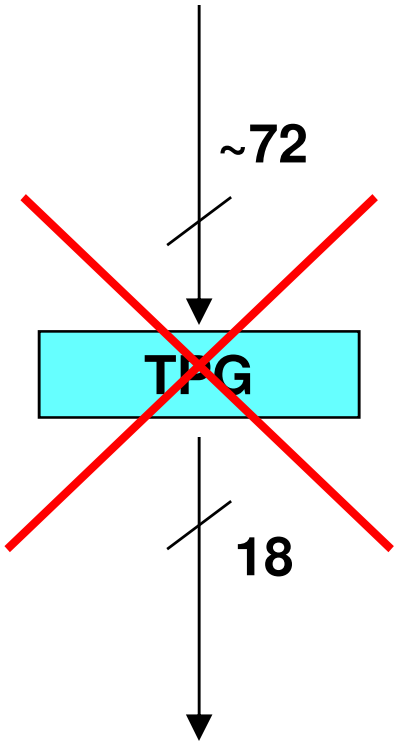

Too much input data?
No, too many serial links.
But only just....
Close to FPGA data BW limit
Assumes input is 3.36 Gb/s effective GBT

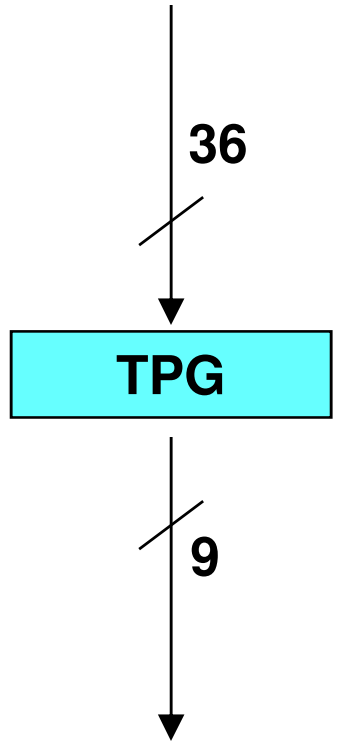

Trigger partitions reduced to 9.
Requires 5.0Gb/s links.
Not possible with GBT protocol
(insufficient FPGA resources)

TPGs with GBT protocol...

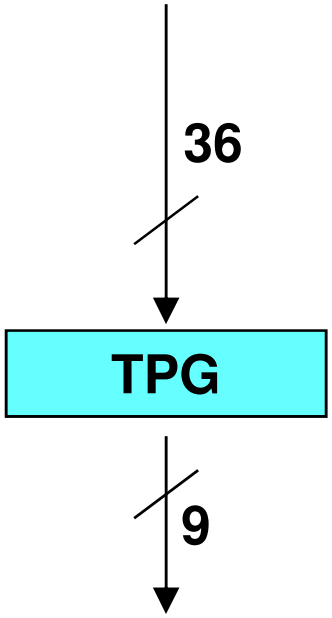

Possible, but not with GBT protocol

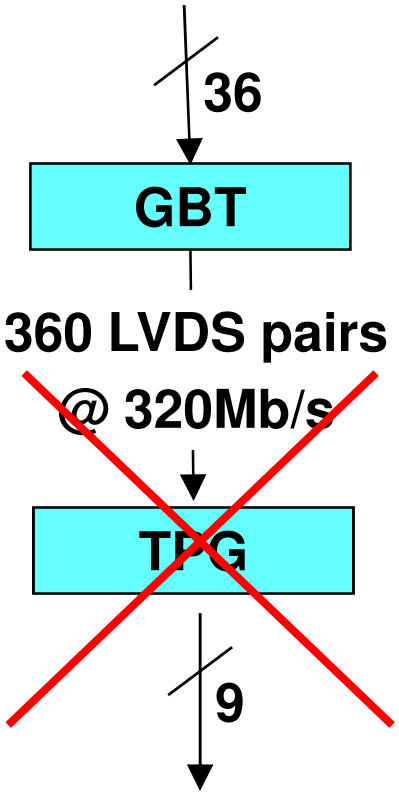

36x GBTs
power ~36W to 72W
real estate = 92cm2
cost = 4320 SFr

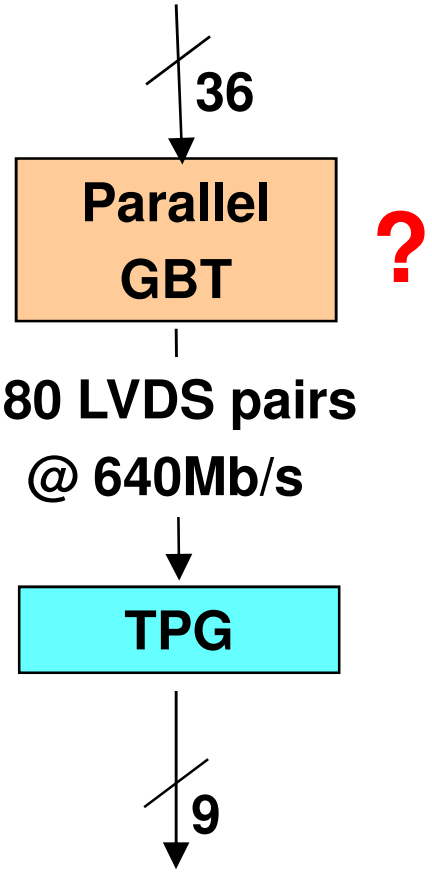

Can select output clk
i.e. doesn't have to be LHC
Will future FPGAs have std I/O?
XC6VLX240T IO = 720 (24,0)

The alternative: Time Multiplexing Rack

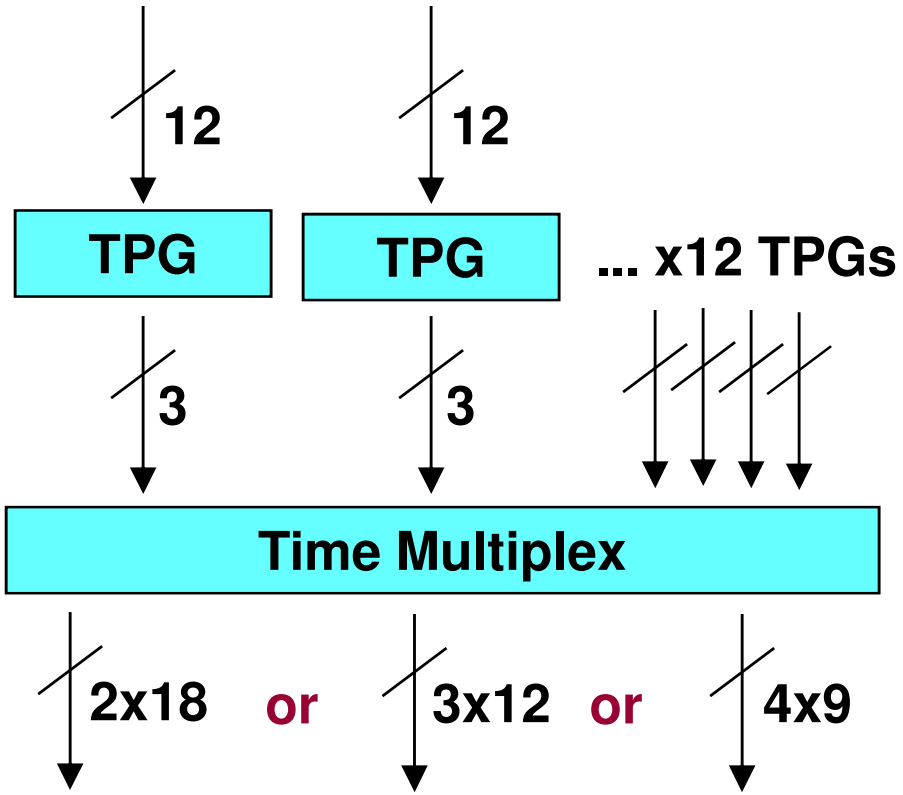

- Use FPGA to decode GBT data
  - Only use ½ of available links to decode GBT otherwise no logic left for other processing
- Use different GTX blocks for GBT-Rx and RT-Tx inside TPG
  - Allows trigger to decouple from LHC clock if required.
- CMS arranged in blocks of constant φ strips
  - Would allow mapping to constant η strips
- Consequence
  - At least x3 to x5 more HCAL hardware
  - Time muliplex rack
  - Larger latency

Time multiplex rack

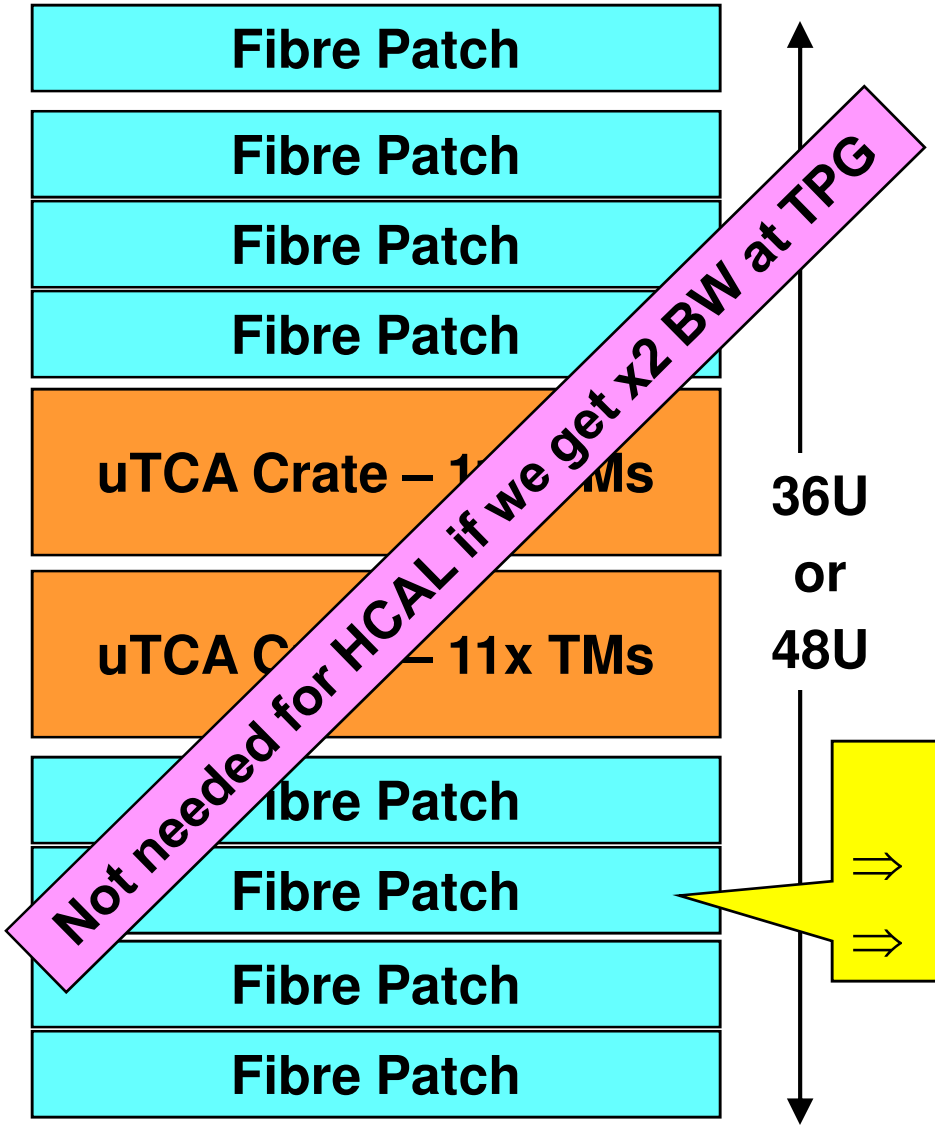

- Assumes tower resolution from HF
  - 18x22 regions
  - 16 towers per region
  - 792 fibres
  - 12bit data, 8B/10B, 5Gb/s
- Time multiplex card may have
  - 36 in/out (2 crates x11 cards)
  - 18 in/out (4 crates x 11 cards)

12 way ribbon = ¼ η, 1 region in φ

⇒ 9 ribbons can cover ½ φ

⇒ Hence patch-panel = ¼ η, ½ φ

Are there are alternatives to Time Multiplex Racks if HCAL TPG stays the same...

- Yes:
  - e.g. return to concept of fine/coarse processing for electrons/jets processing
  - i.e. jets at 2x2 tower resolution

- Requires less sharing (e.g. just 2, rather than 4 tower overlap for fine processing)
  - ⇒ Alllows a more parallel architecture
  - ⇒TPG would have to provide 2 towers in η and ¼ of φ

But... Is it wise to separate fine & coarse procesing ?

Still requires some thought

& discussion with TPG experts

End...

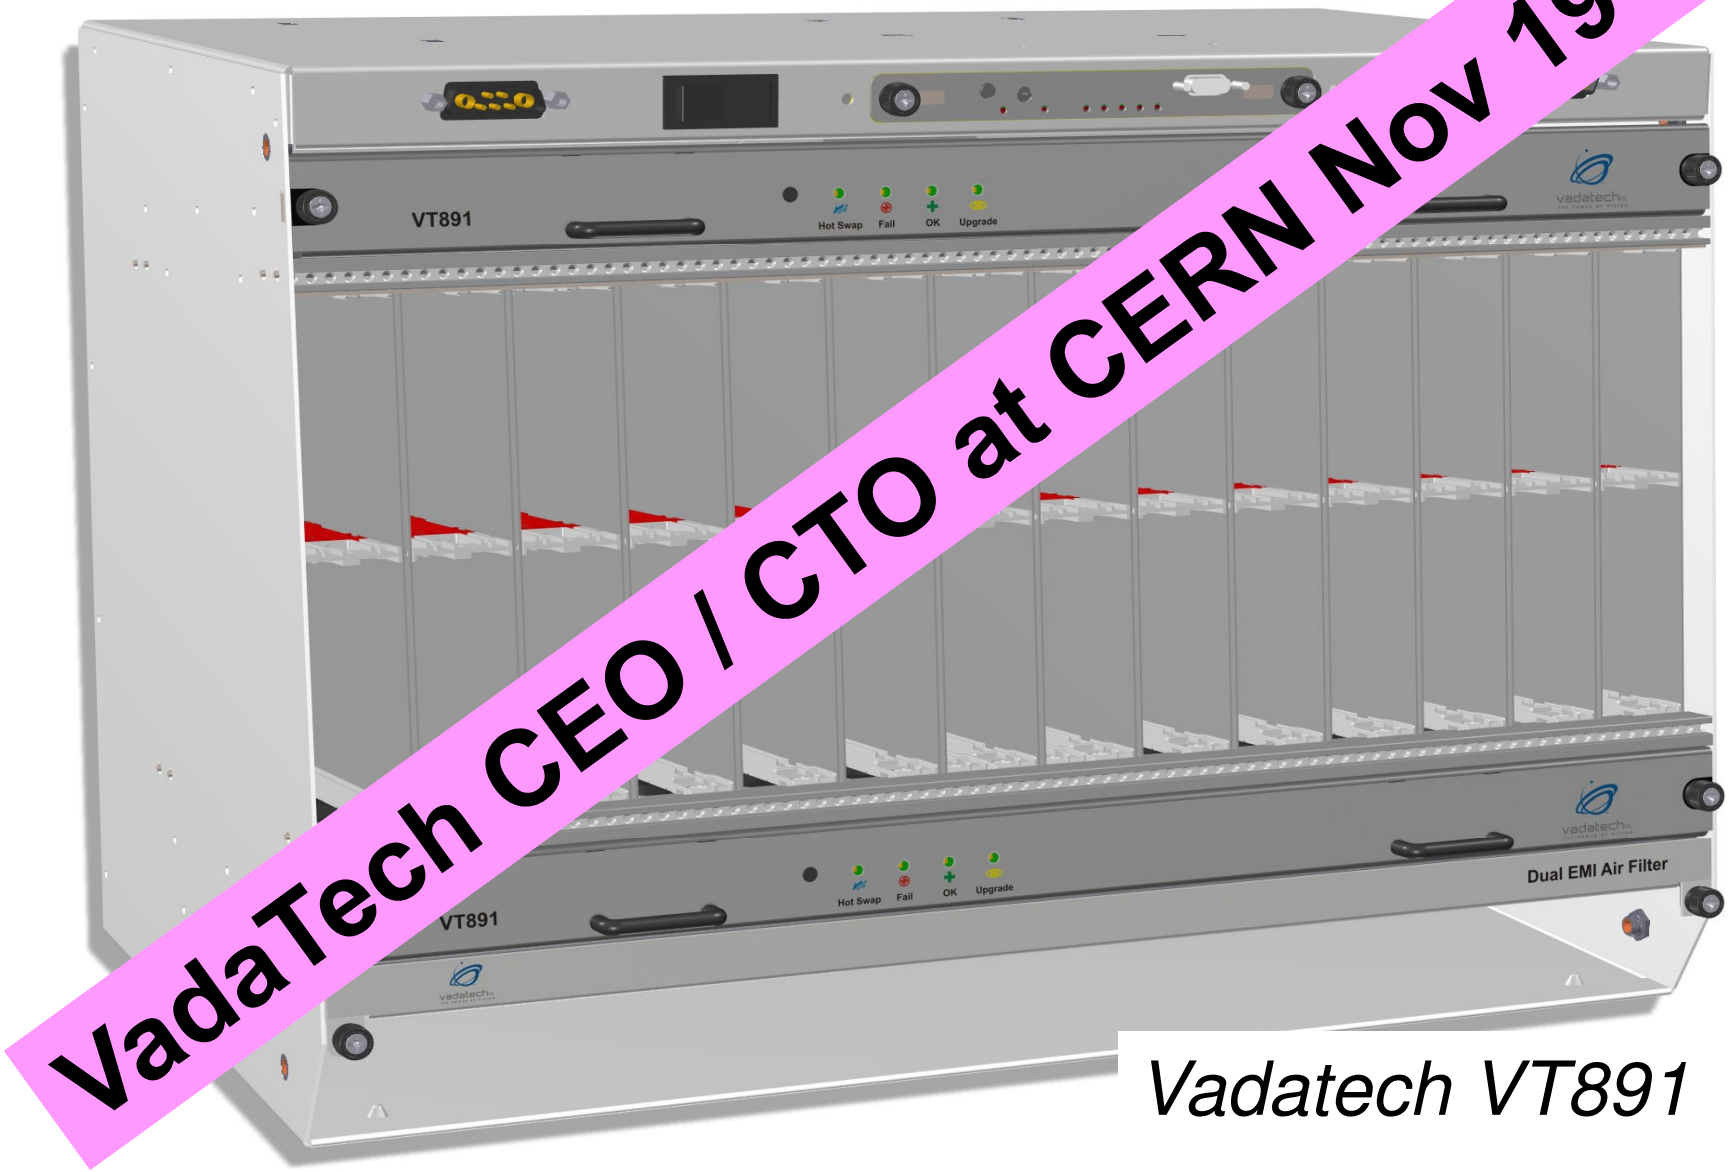

VadaTech CEO / CTO at CERN Nov 19th

Vadatech VT891

Extra

Cavern Data into TPG

½ region in η, all φ regions OR 1 region in η, ½ of all φ regions

Mux: Put all data from 1bx into single FIFO

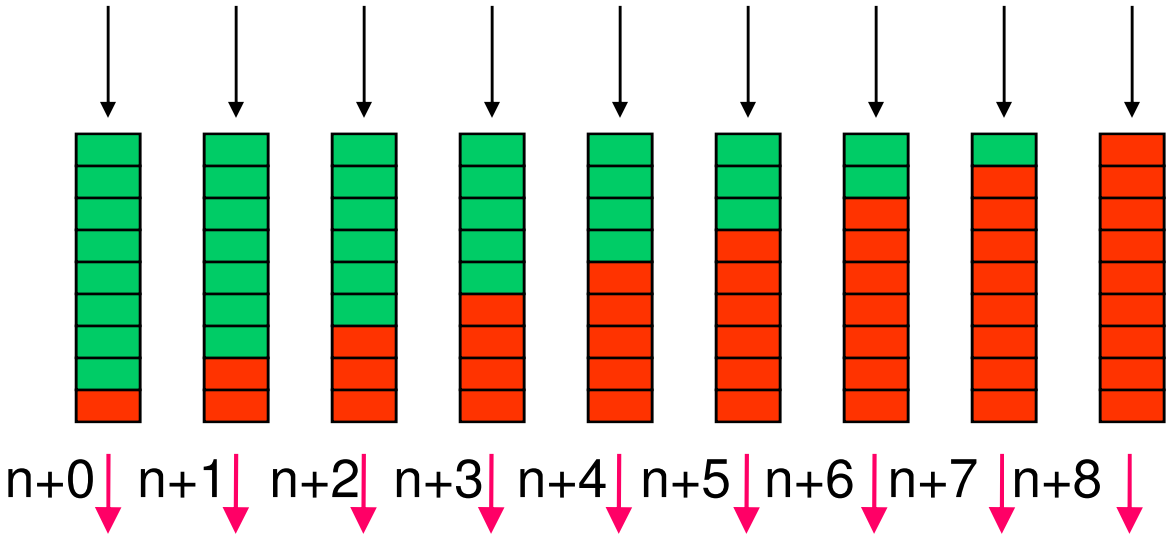

Single fibre to each processing blade

Decode GBT

- Only instrument 20/24 links of XC5VFX200T, 96% logic cell usage
  - F. Marin, TWEPP-09. Altera results slightly better

- Cannot use GBT links to drive data into processing FPGA
- Options
  - (a) Use powerful FPGAs to simply to decode GBT protocol
    - Seems wasteful.
    - SerDes Tx/Rx may have to use same clk
  - (b) Use mulltiple GBTs
    - Diff pairs = 200 @ 320Mb/s lanes. OK if diff pairs remain on FPGAs!
    - Power = 20W (assume less power because less I/O)
    - Components = 20 (16 mm x 16 mm, 0.8 mm pitch)
    - Cost = 120x20 = 2400 SFr !
  - (c) Use Xilinx SIRF both on & off detector
    - If they let us...