L1Topo

Run-2 → Phase-1 (→ Phase-2 L0Topo)

Uli / Mainz

L1Topo in its habitat

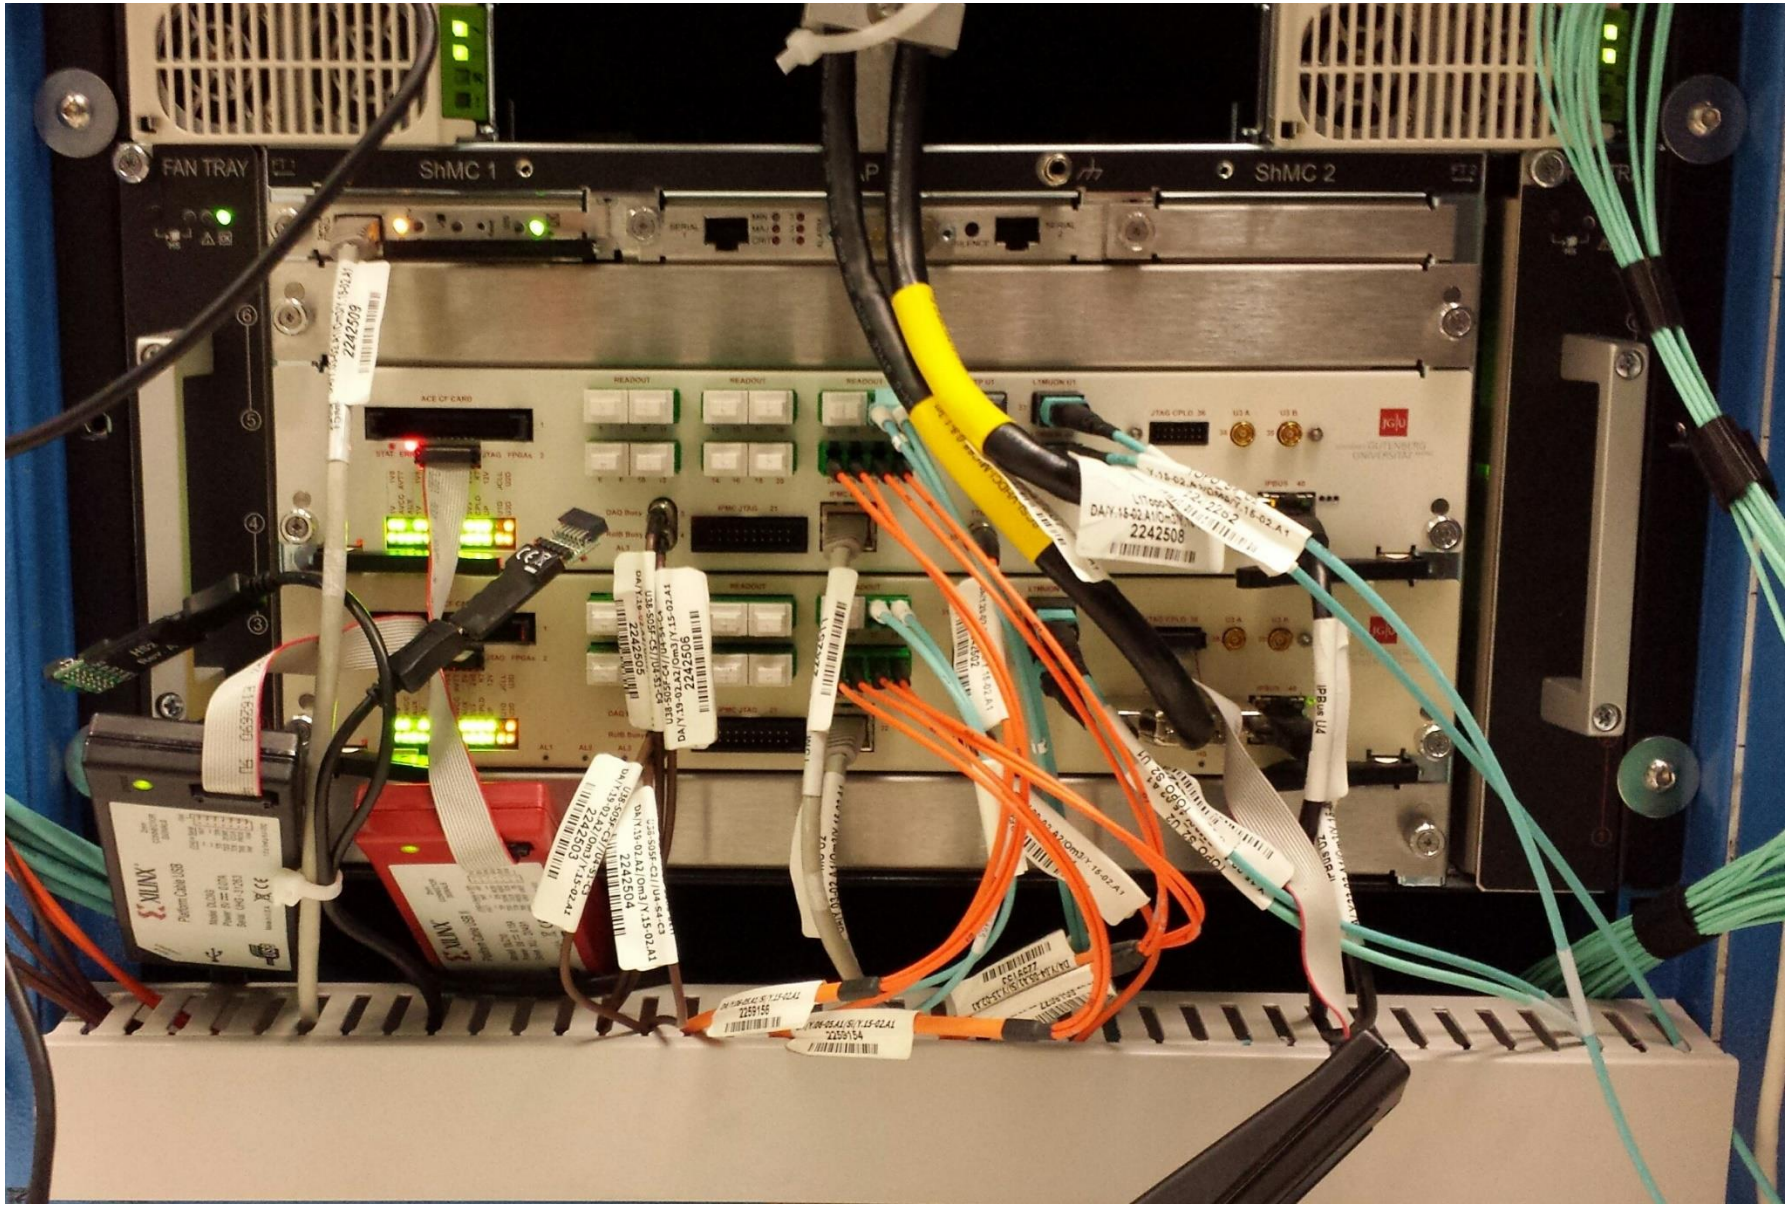

L1Topo / run 2

L1Topo / Run-2

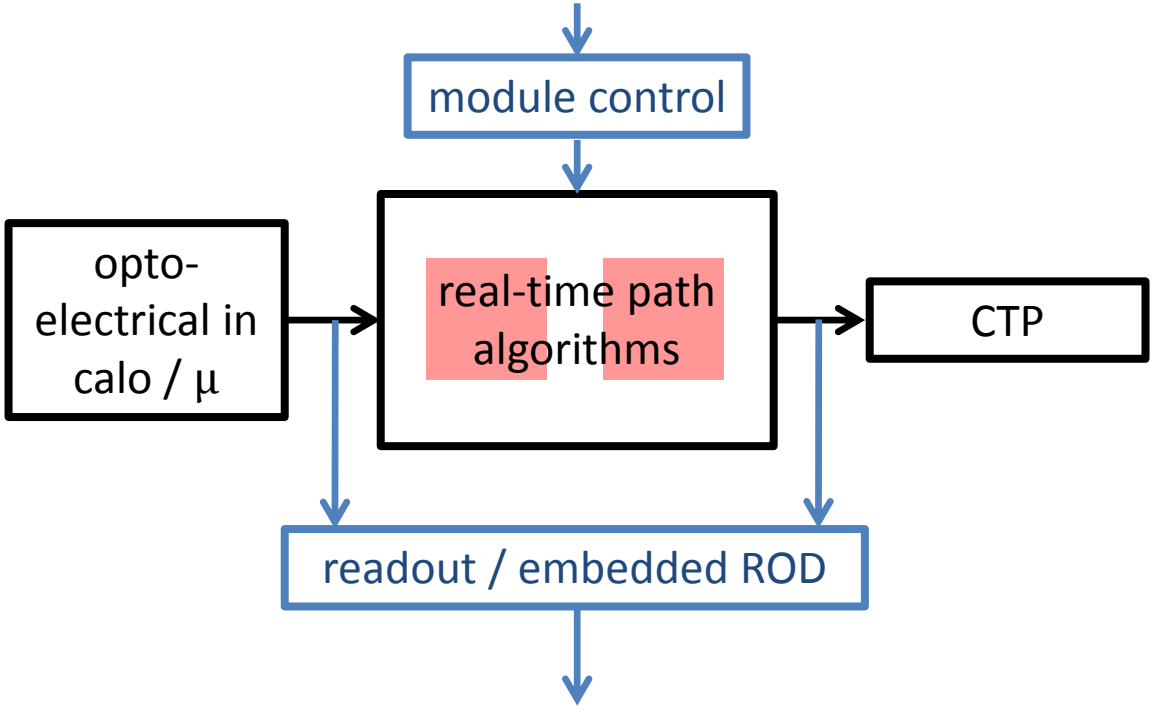

- five modules built for run 2
  - two algorithmic processors
  - one DAQ/control FPGA
  - designed for phase-1/phase-2 compatibility
- two modules deployed at point 1
- one module in test rig part of link speed tests (12.8Gb/s)

L1Topo unplugged

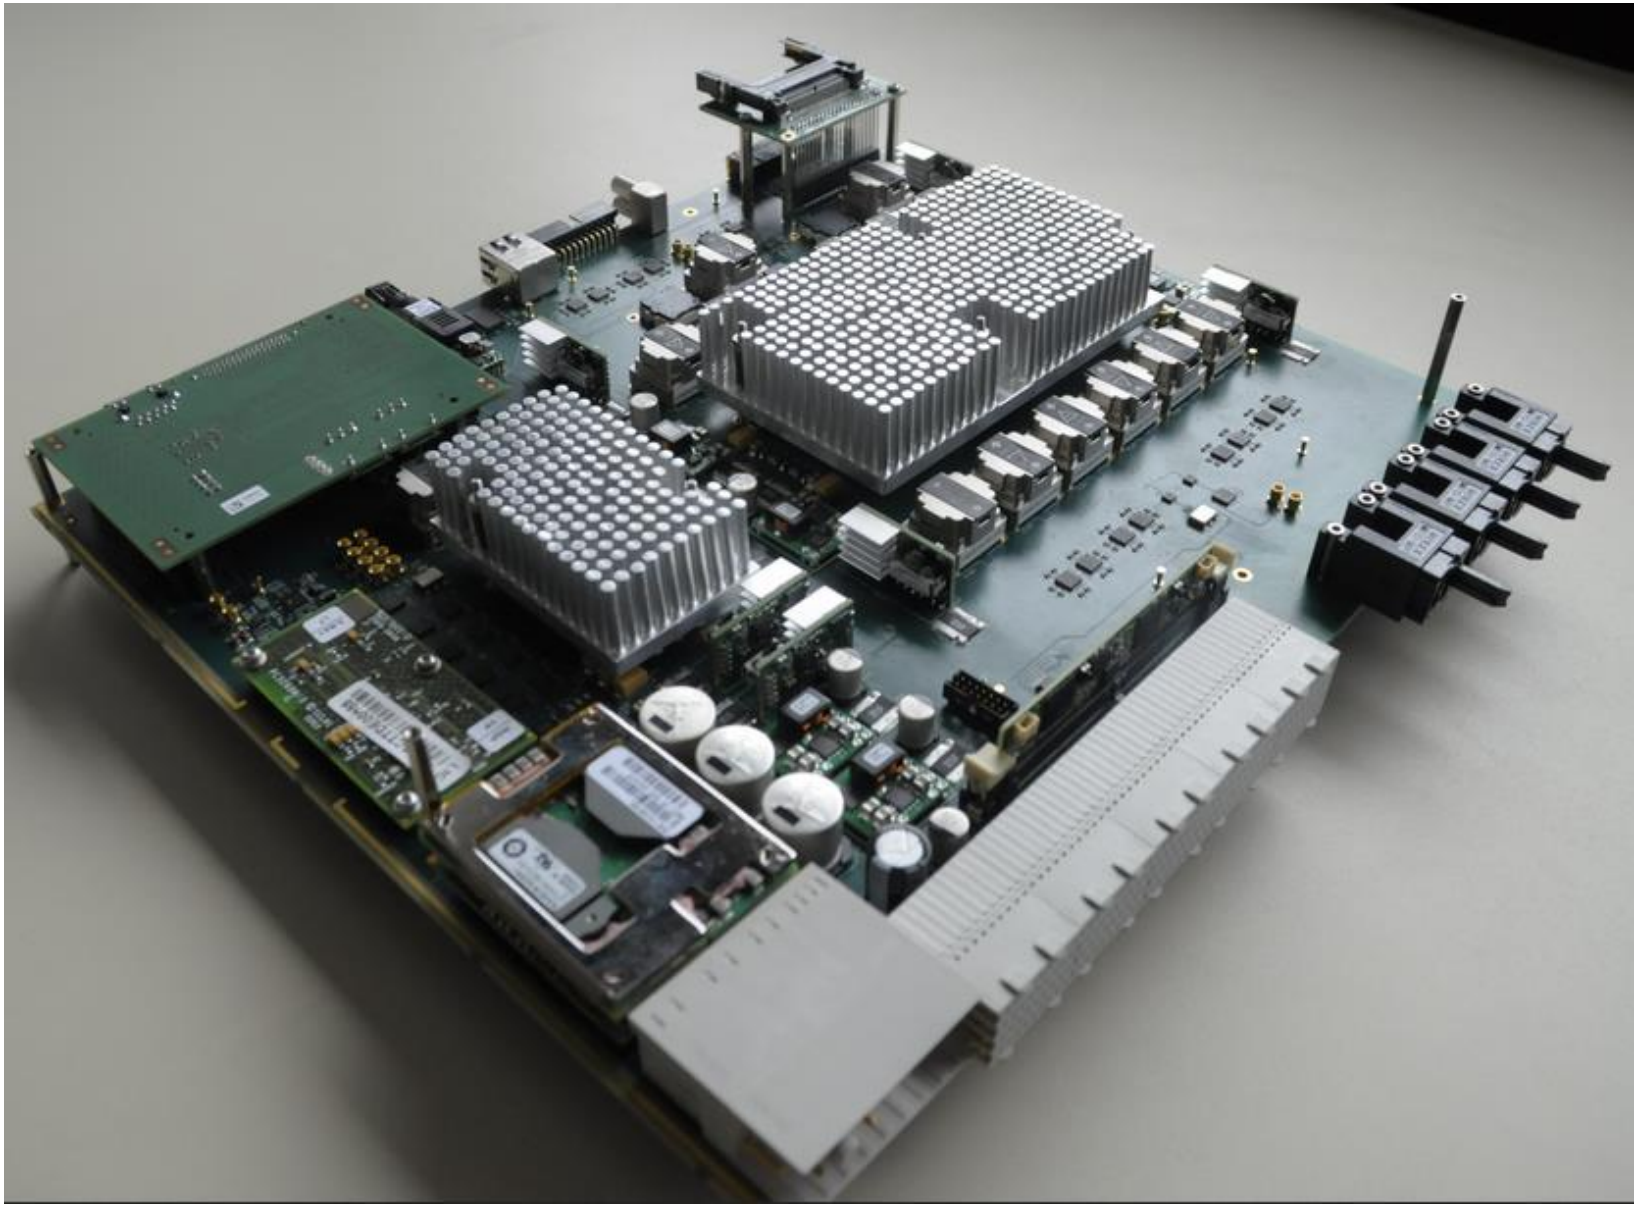

L1Topo floor plan / run-2

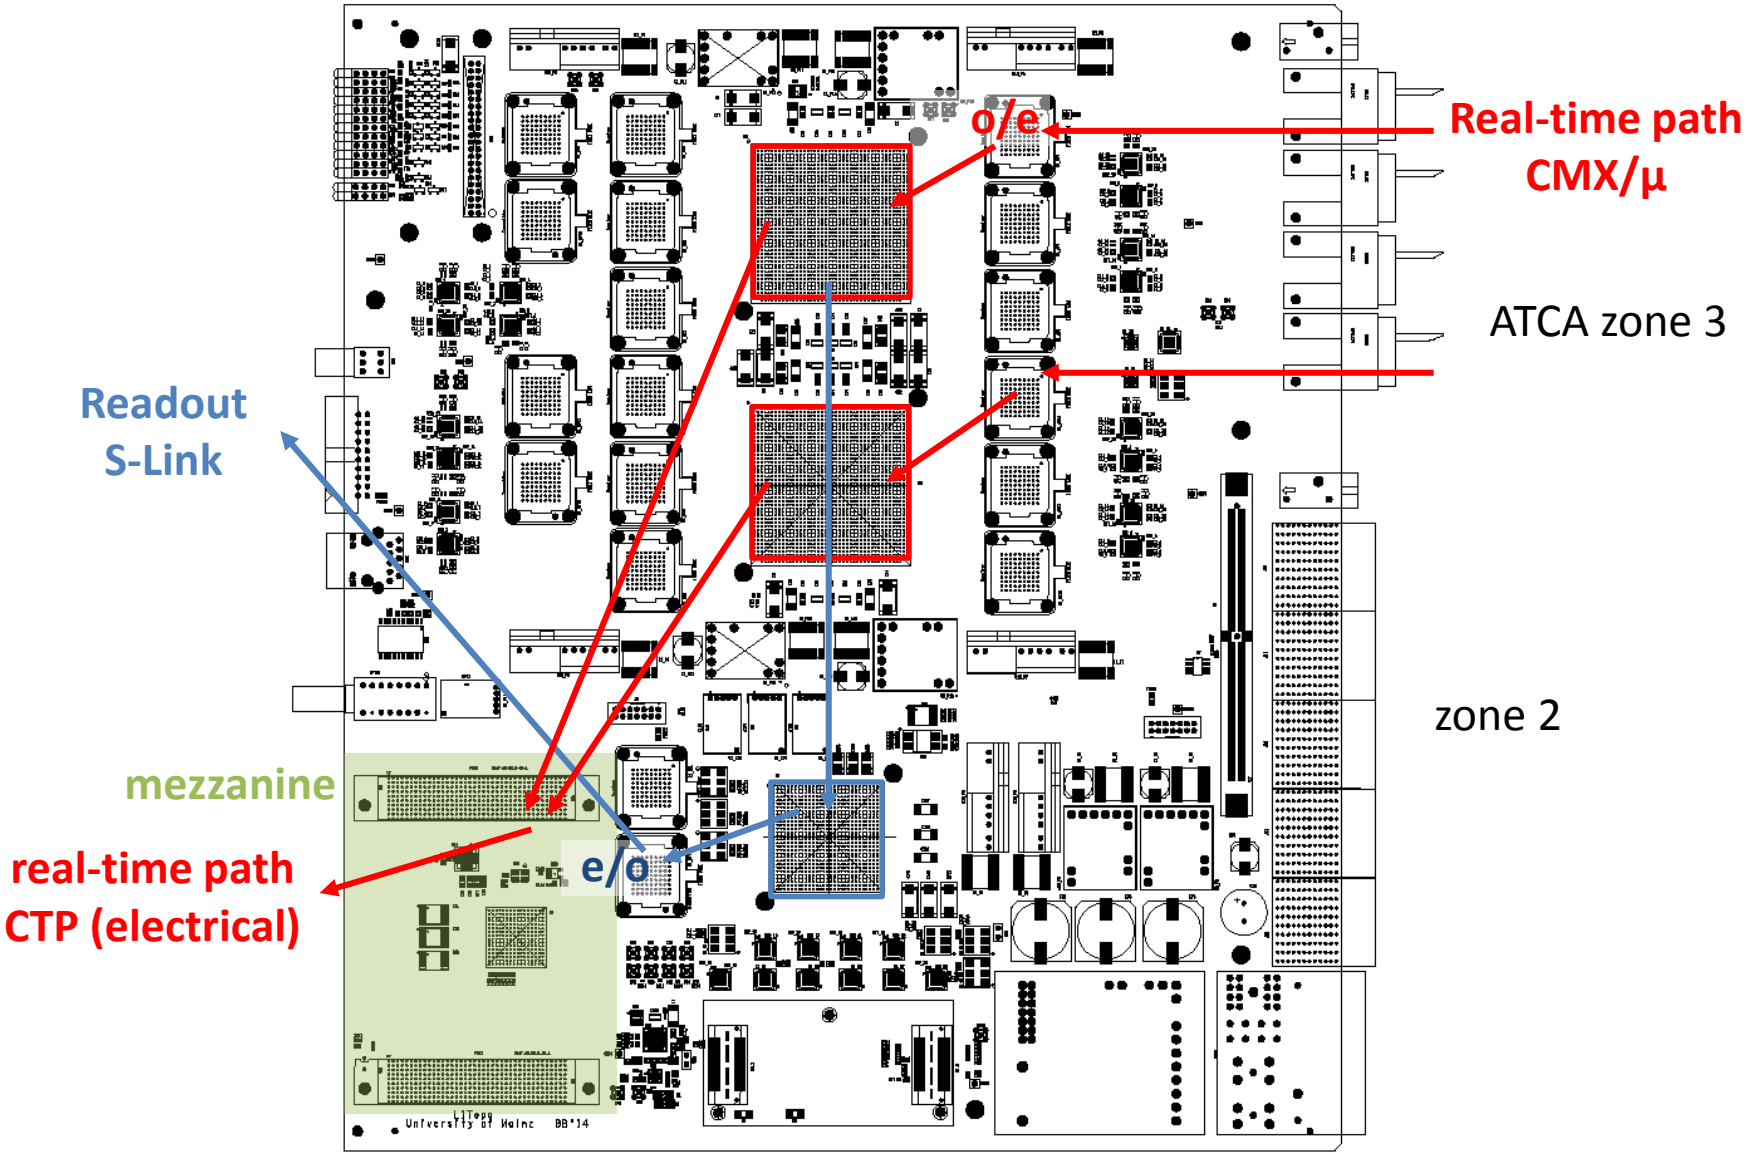

L1Topo → phase 1

- Additional L1Topo production for phase 1
- 3 new modules were initially planned for phase 1
  - 2 × topology
  - 1 × hit merger for L1Calo multiplicity triggers
- Current L1Topo can act as a prototype:
  - Existent modules expected to match needs for Phase-1
  - No changes to concept
  - Will need to run new PCB production anyway
  - Will allow for improvements / adaptation to phase-2 requirements

- Expected to live in an ATCA water cooled "standard" shelf
  - Space
  - Power / cooling

L1Topo floor plan / phase-1

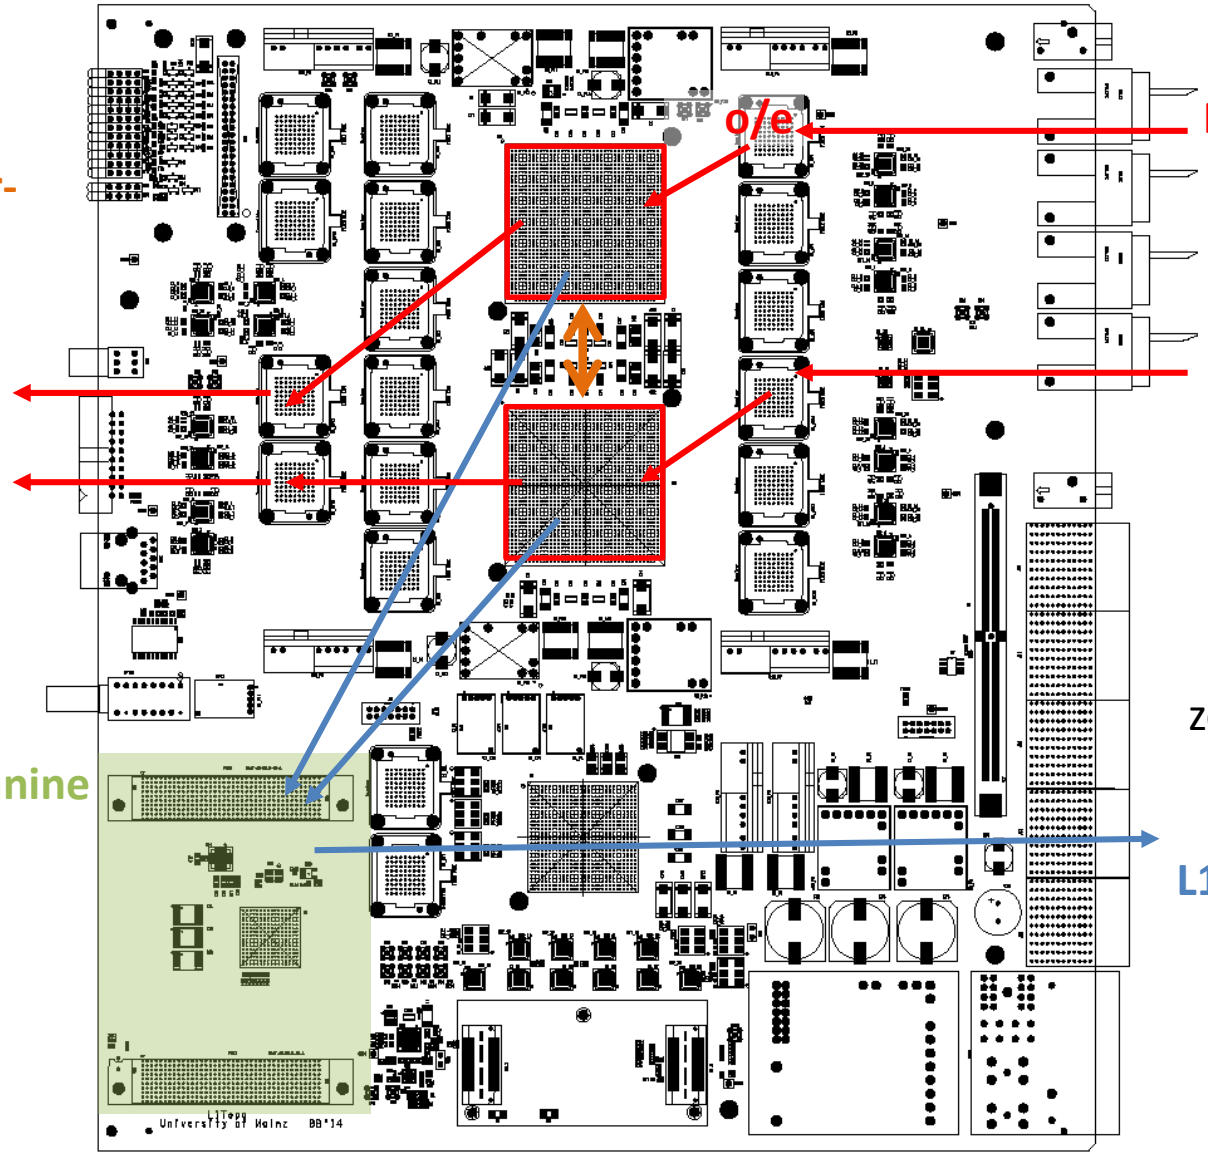

Latency at Phase-1

Table 13: L1Topo

	ns	BCs	Sub Total	Total	
Optical Input available from CMX, eFex, jFex & Muctpi				64,9	
L1Topo Input Deserialisers	50	2,0			
Synchronize to local clock - 320-> 40 MHz	25	1,0			
Algorithmic Processing	125	5,0			
			8,0		
Electrical Output to CTP (multiplexed) (if us	25	1,0			
Electrical Cable to CTP (if used) (2m)	10	0,4			
			1,4		
L1Topo electrical input available at CTP				74,3	L1Topo_Electrical
Output Multiplexers 40-320 MHz (if used)	25	1,0			
Output Serialisers for optics (if used)	50	2,0			
Fibres to CTP (if used) (2m)	10	0,4			
			3,4		
L1Topo Optical inputs to CTP available			11,4	76,3	L1Topo_Optical

L1Topo:

TDR figures

- Latency is tight
- Inter-FPGA fan-out is part of algo latency ! ~2BC ?
- Incoming signals on latency critical path (Muons / NSW) might bypass inter-FPGA fan-out (upstream duplication)
- Check input MGT/deserialization latency assumptions against current reality (Topo I)
- Electrical fast output path available (limited bandwidth)

Summary: Modifications wrt run-2

- Connect to L1Calo hub/ROD modules via backplane
  - Requires different extension mezzanine
  - Basically route-through of 16 high-speed signals
    - DAQ/ROI
    - Clock / TTC data
- Switch-on inter-processor communication
  - Signals duplicated between FPGAs rather than forward duplication upstream
  - Adds bandwidth for incoming signals
  - Latency: be prepared to continue duplicating a fraction of signals upstream. The ones on latency critical path… NSW…
- Add output bandwidth for signals into CTP
  - Optical link
  - No plan to remove low-latency electrical link
- One possible modification under consideration already now:
  - We might prefer the CERN recommended jitter cleaner (used on jFEX prototype) over the one used on L1Topo

Needs / changes for phase-2

… as known by now

- Challenging processing: more/complex algorithms
- Fortunately less stringent latency requirements
  - Far less duplication of internal logic resources required
  - Time multiplexing might be possible

- Plus the need to supply ROIs to "Regional Readout Request" (R3) logic
  - Near real-time
  - Latency approximately on the same scale as L0 trigger
  - Will not be possible through the backplane / L1Calo hub/ROD modules
  - ROIs basically TOBs that fired topo algorithms
  - *… how many ?*

L0Topo floor plan / phase-2

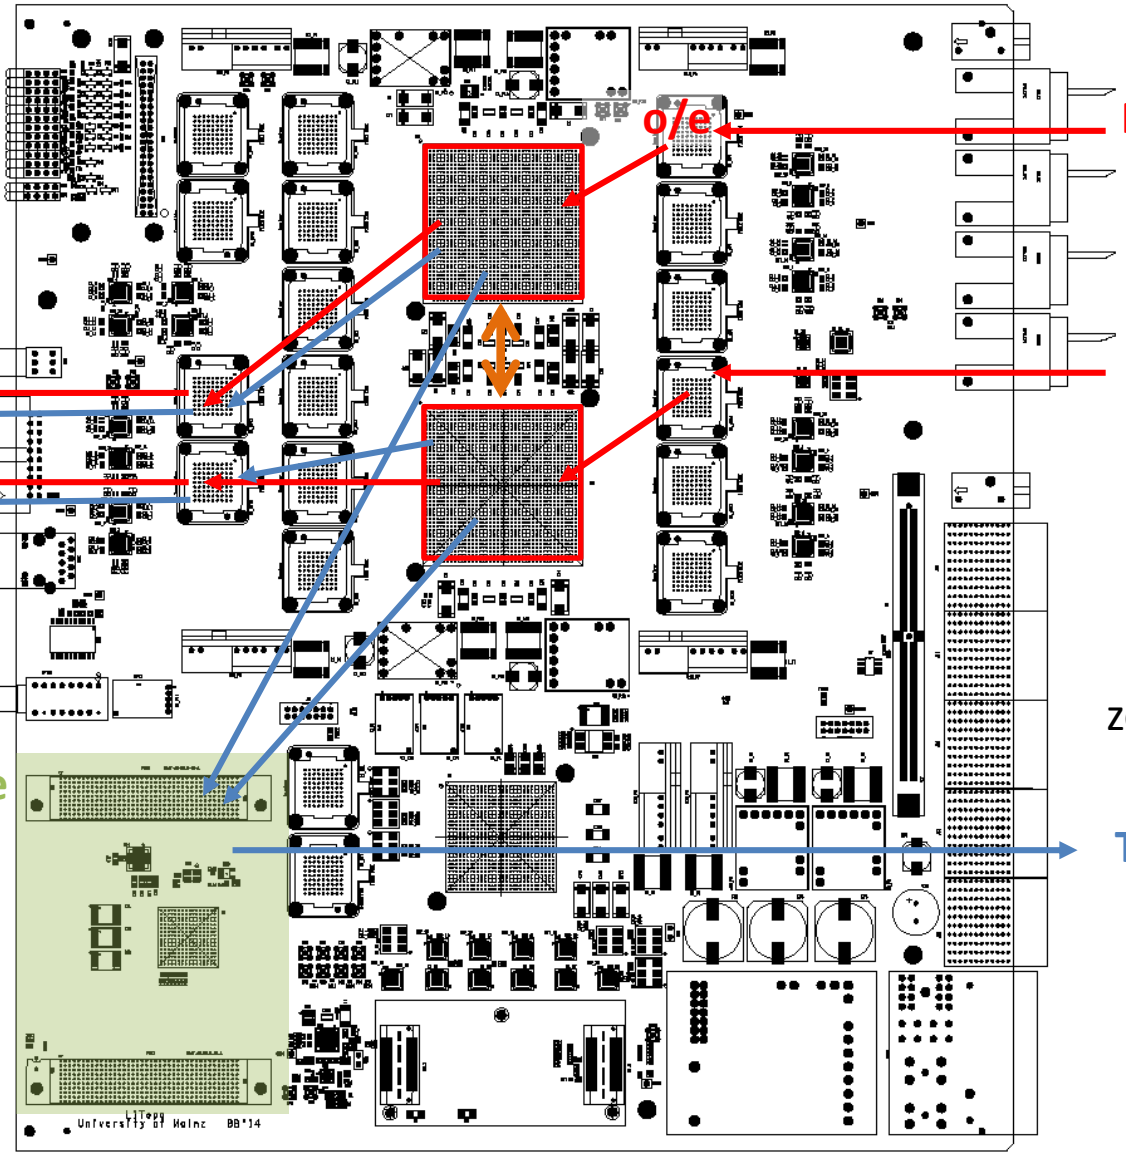

Summary of changes for phase-2

- Optical output bandwidth needs to be shared between L0 and R3 data
  - Total current bandwidth per module : 24 * ~10Gb/s
- Algorithmic firmware will need to be modified to address
  - Need for complex algorithms
  - Need to identify TOBs that fired the triggers
  - Involves tagging / pipelining of TOBs such that at time of decision the information is still available
  - Route the TOBs / ROIs out to high-speed links

In the light of two well-filled L1Topo modules in run-2:

- The system is scalable
  - Standard ATCA crate has 12 slots available for processor blades
  - Will allow for additional modules to be installed

Stop press – NSW latency

Recent e-mail thread with Alexander Oh, new small wheel latency task force. Suggestion as of yesterday evening (inspired by David S., Stefano V.):

- There are 4 topo FPGA.
- We use 1 FPGA for muon algs only.
- No sorting required.
- Process muon algs within 1 BC.

The algs would be mainly DR cuts between muons in jets for tau->mu gamma and tau related triggers, and invariant mass cuts for B physics. For these we would use only objects above a certain threshold (jets/egamma) but all muon objects, so no sorting is required.

The muon information would come in later, but if we need only one BC for processing we could solve the latency problem.

Uli: … decision latency is currently 1BC for *all* algos, but that cannot be guaranteed if object count and TOB width (granularity) are going up. "All with many" type algos require resources/latency, though simple dR cuts are relatively harmless (at least if box cuts are acceptable !!)… Eduard:… currently the muon algos are well distributed over 3 FPGAs…

Summary / outlook / discussion

- L1Topo module concept will allow smooth transition up to phase-2
- Current L1Topo part of link speed test setup (successful so far !?!)
- Some hardware modifications possible before new production run
- No current plans for major rework
- System is scalable, larger module count anticipated
  - Upstream modules should allow for sufficient output bandwidth to supply larger numbers of L1Topo modules
- Latency at Phase-1 !!! Not just NSW !
- Phase-2 : ROI/R3 bandwidth of ~200Gb/s per module adequate ?
- If not: need to understand now and upgrade design
  - Spare output bandwidth on processor FPGAs available
  - Need to be routed out to additional opto/electrical converters
- # of L1Topo modules: can/should we try to shrink module to single width ? Not unless required !!!
- Planning to start design work for pre-production module soon after jFEX prototype done !!!!!!! → require feedback *now*
- Firmware will require major changes to benefit from increased latency and to support R3 concept
  - Data formats and protocols into R3 ?
  - Unidirectional links! – otherwise eat into real-time bandwidth !