# L1Calo Phase 2 Upgrade

## and implications for channel organisation

Murrough Landon

25 February 2009

- Baseline Concept
- L1Calo Architecture Idea(s)
- Constraints
- Implications for ROD Crates and Front Ends
- Summary

# Introduction to L1Calo Phase 2 Work

- Work done so far (a little)
  - Survey of ATCA, links and other technology
  - Survey of HLT algorithms we might steal
  - Thoughts on general L1Calo phase 2 architecture
  - Issues relating to mappings from FE and RODs
- Work not done yet (lots more!)
  - Simulation, simulation, simulation
  - More simulation – especially of pile up
  - Detailed thinking about
    - Granularities
    - Algorithms
    - Architecture
    - Bandwidths
    - etc

# Baseline Phase 2 Concept

- Strong preference to digitise and transmit all cells every BC from the front end to off-detector pipelines
- ROD/preprocessor:
  - generates Level 1 primitives (towers++) from calibrated Et for the correct BC
  - sends them to a separate L1Calo trigger processor
  - different LAr/Tile views of the ROD?

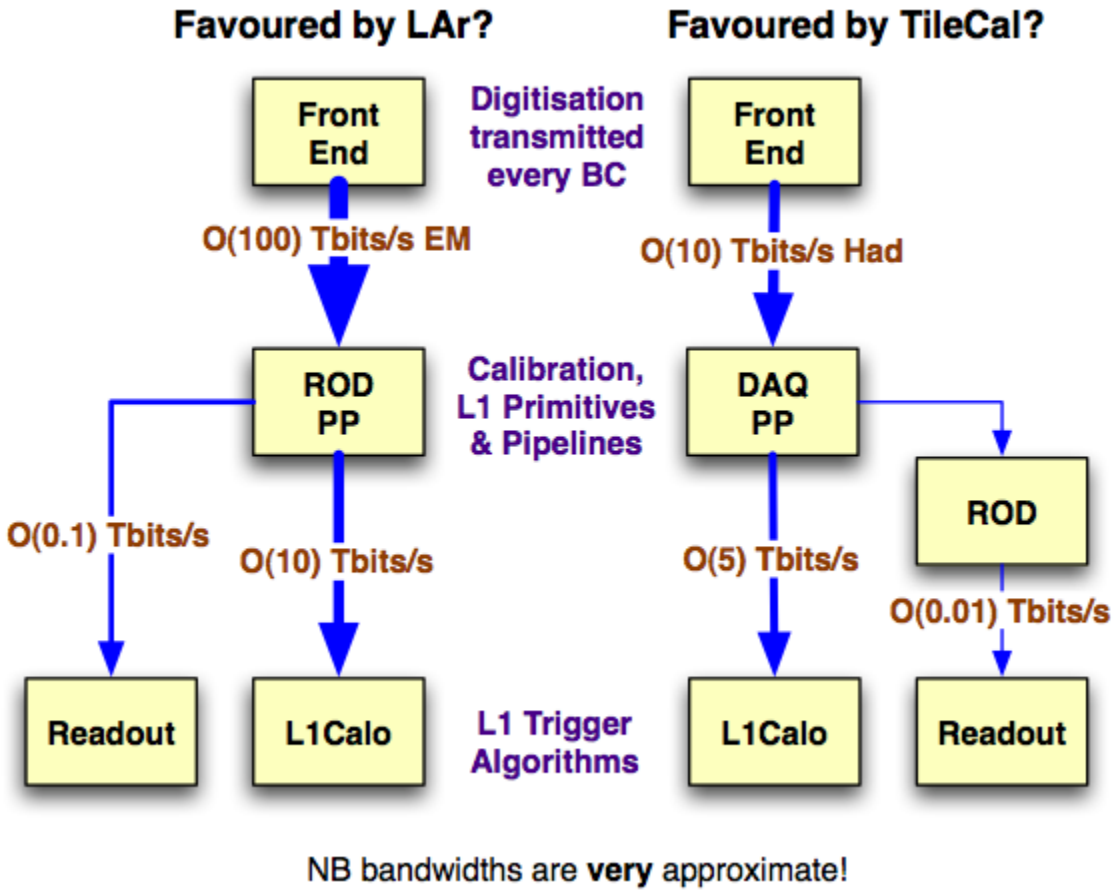

# Implications of On-Detector Digitisation

- Allows much more sophistication in forming "towers"
  - New digital "L1 Primitive" could be a bit field with
    - Et (to greater precision than before if required)
    - Depth and lateral shower profile information, quality flags?
  - Better handling across boundaries?
  - Finer granularity (in EM layer)
    - Different granularities possible in EM and Hadronic layers
- Single calibration for trigger and main readout
- But brings trigger and readout closer together
  - Present architecture allows each branch complete freedom to optimise the organisation of their own system
  - Phase 2 upgrade will impose trigger constraints on layout of RODs and mapping of FE to ROD links
  - Also no (completely) independent readout path

# Present L1Calo Architecture

- Separate EM/Tau and Jet/Energy processors
- Sliding window algorithms
  - Requirement for environment
- Phi quadrant layout
  - O(30%) fanout at source (PPM)
  - O(75%) fanout at CPMs/JEMs
  - Strong requirement on eta,phi shape covered by all modules
    - NB orthogonal to detector layout
- Many remapping stages
  - Receivers (20 mapping variants), patch panels, PPMs

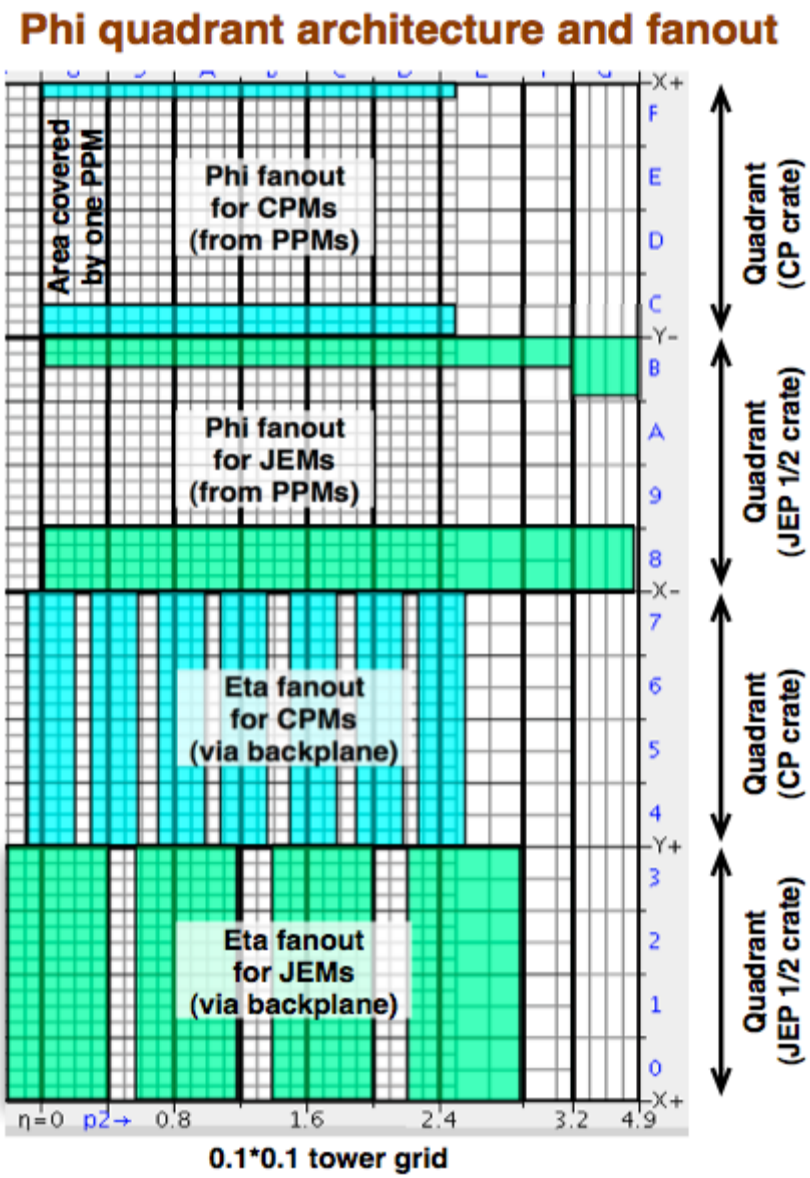

# Possible Phase 2 Architecture (1)

- Single processor module?
  - For all objects: EM,tau,jet
- Still use sliding windows
  - Unless there is a better idea?
- Fewer remapping stages?
  - May want fibre ribbon PPs?
- Consider phi octant layout?
  - Similar fanout in L1 modules
    - Unless modules wider in eta?
  - O(75%) fanout from RODs
  - Fewer restrictions on eta,phi shape covered by RODs
    - But still need regularity

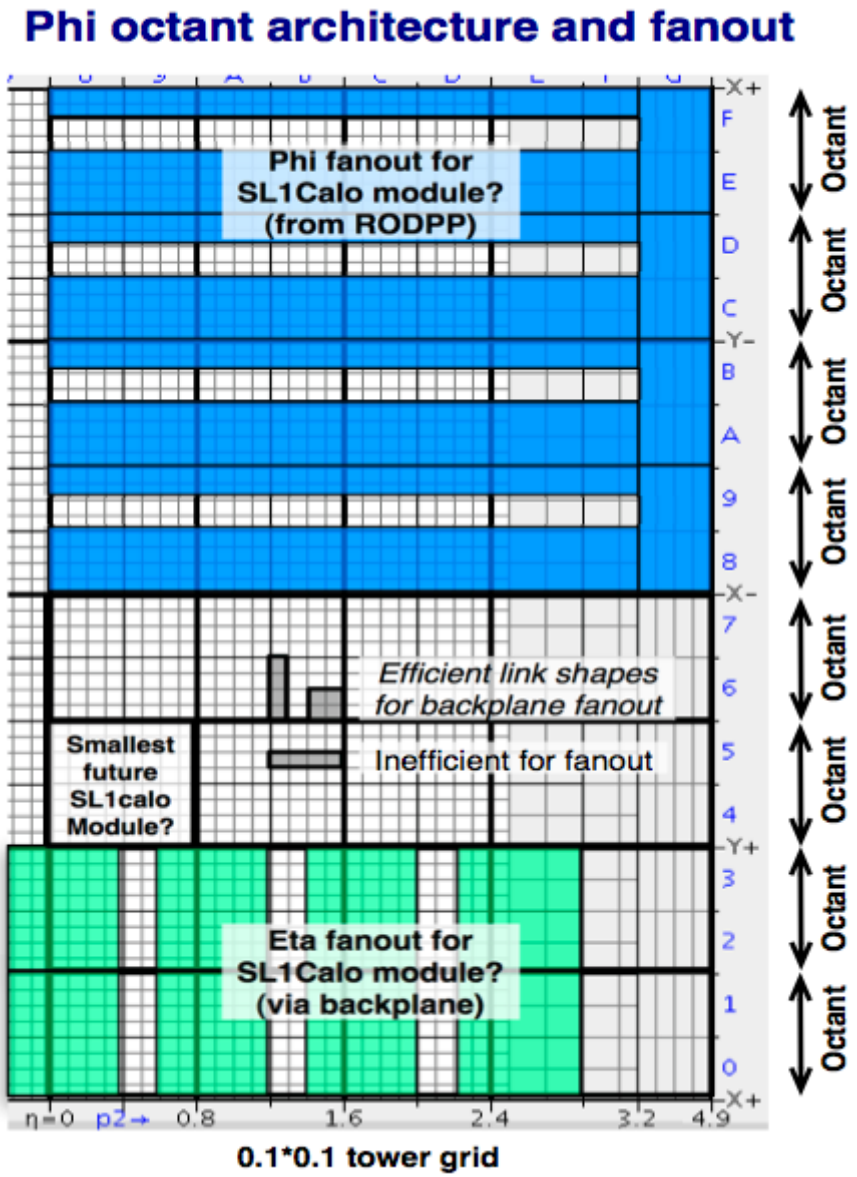

# Possible Phase 2 Architecture (2)

- Links from RODs duplicated to neighbouring octants
- Links to neighbouring L1 modules duplicated without reserialisation
- Example optical bandwidths
  - Including factor 2 phi fanout
  - Excluding 1.75 electrical fanout
  - Assume 0.8*0.8 trigger module

| Option | Tbit/s (total) | Gbit/s (module) |
|---|---|---|
| EM,had towers+bits | 14 | 200 |
| Minitowers*samplings | 60 | 1000 |
| cf LAr ROD | | 600 |

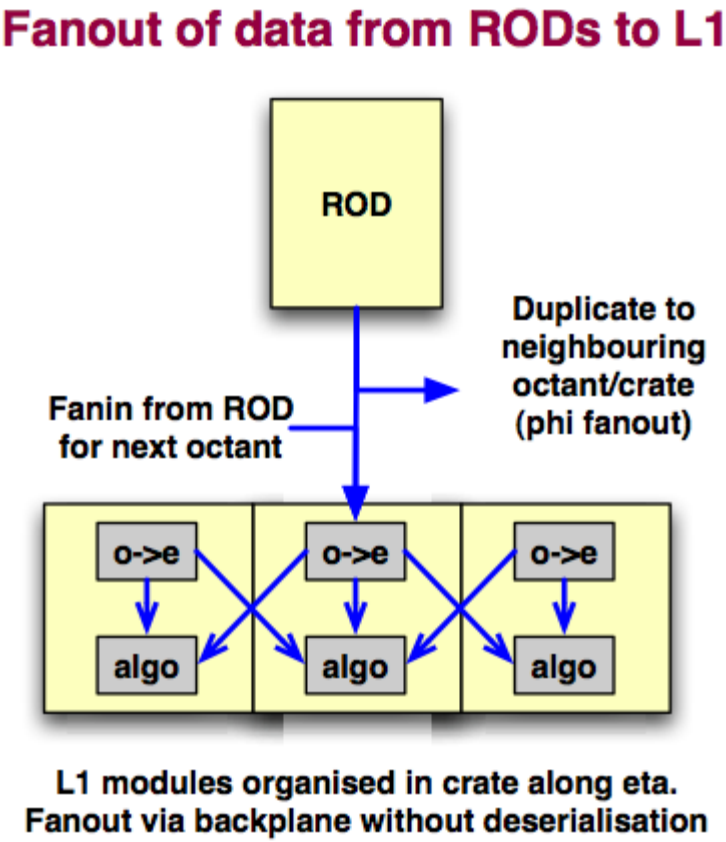

# Latency

- Two main upgrade scenarios:
  - No L1 track trigger: small latency increase possible (constraint from muons?)
  - With L1 track trigger: need fast L0 seed from calo+muon (latency same as now?)
- Latency is still critical
- No unnecessary deserialisation/reserialisation
  - Significant latency penalty (6 BCs?) at each such step
- Unavoidable: FE->ROD, ROD->L1, L1->Merger/CTP(?)
- Avoidable: everything else!
  - Eg between RODs, between L1 modules

# Granularity: Towers

- Need something better than existing analogue sums
  - "L1 primitive": Et plus additional information bits?
- Eta,phi options:
  - Keep present 0.1*0.1 towers (required for hadronic layer)
  - Fine granularity 0.05*0.05 EM "mini-towers"
    - Strip layer covers 0.1 in phi: how to assign energy?
  - Intermediate 0.05*0.1 EM "semi-towers"
    - Easier to handle the strip layer
  - Lateral shower profile, position and quality bits for EM layer
- Depth (sampling) options:
  - Send single object with depth profile bits
    - Requires all data in the same place in the ROD
  - Or send separate depth samplings?
    - Fewer constraints on RODs, but may be too much data for L1?

# Granularity: Links

- ROD->L1 links
  - Sliding window algorithms require lots of fanout
  - For phi octant layout, this is most efficient if L1 link contains contains more "towers" in phi (2**n) but is narrower in eta
    - Not how the detector is organised (especially TileCal)
  - Small number of towers per link easier to handle
    - But greater number of serial streams to fanout
  - 1 Gbit/s is 25 bits at 40 MHz
    - Roughly one EM tower with Et and profile bits?
    - 6 Gbit/s would easily cover 4 towers (or mini-towers)
      - Likely possible to cover eta,phi space with 0.2*0.2 links
- FE->ROD links
  - Group together cells in towers
    - Follow existing tower builder or Tile adder layouts?
      - Projective geometry in TileCal, not division by z

# Channel and Link Organisation (1)

- Lessons from existing L1Calo
  - Worry about the difficult areas early in the design process
    - It only gets worse later (and dont forget about the FCAL!)
  - Do as much as possible at the first stage in the chain
    - Irreducible constraints from calo geometry will hit later
- Link organisation
  - Data processed together needs to be brought the same chip!
  - Best to bring links directly to the right chip
  - If not, at least to the same module
  - Or from a module in the same crate (fast parallel transfer)
  - Avoid need for high latency serial transfers
    - Either between modules in the same crate or different crates
- L1 constraints affect the layout of RODs and FE links

# Calorimeter Boundaries

- Worst case (by far) is EM barrel/endcap transition
  - Anything we can possibly do will probably be needed
  - Sum cells across EMB/EMEC before making L1 primitives?
  - Add in crack scintillators? (Currently read out via Tile EB)
    - Upgrade beig considered in that region?
- Cant do anything about crack at eta=0
- Next worst is Tile LB/EB transition
  - Currently cells are deliberately misorganised to adjacent eta bins to avoid analogue summing across the boundary
  - Probably need to do it properly for the upgrade
  - Add in the gap/crack scintillators?
- Least worst is Tile EB/HEC transition
  - Again, currently misorganised – could do better digitally

# Channel and Link Organisation (2)

- Little guidance yet from simulation
  - Assume the worst cases (from FE and ROD viewpoint) and look at the implications
- Assume L1 primitives formed from all depth samplings
  - For EM and hadronic layers separately (at this point)
    - Sending separate depth samplings to L1 is easier for organising links
- Assume L1 primitives must cross calo boundaries
  - Process Barrel/Endcap cells together in same chip
  - Assume crack scintillators for EMB/EMEC boundary
    - Implies LAr and Tile RODs sharing crates
- Worst possibility? Full EM+hadronic depth summing
  - Inevitable latency penalty, high degree of convergence between EM and hadronic RODs (and shared crates)

# Possible Upgrade Mapping Stages?

- Fewer steps available?
  - Unless we add latency with an additional reorganisation
- Start with FE boards
  - May need several different channel to link mappings?
- Remap FE to ROD links
  - Signals in depth and across boundaries to same place
- Minimal (low latency) transfers in ROD crates?
- Regroup (and duplicate) ROD to L1Calo links

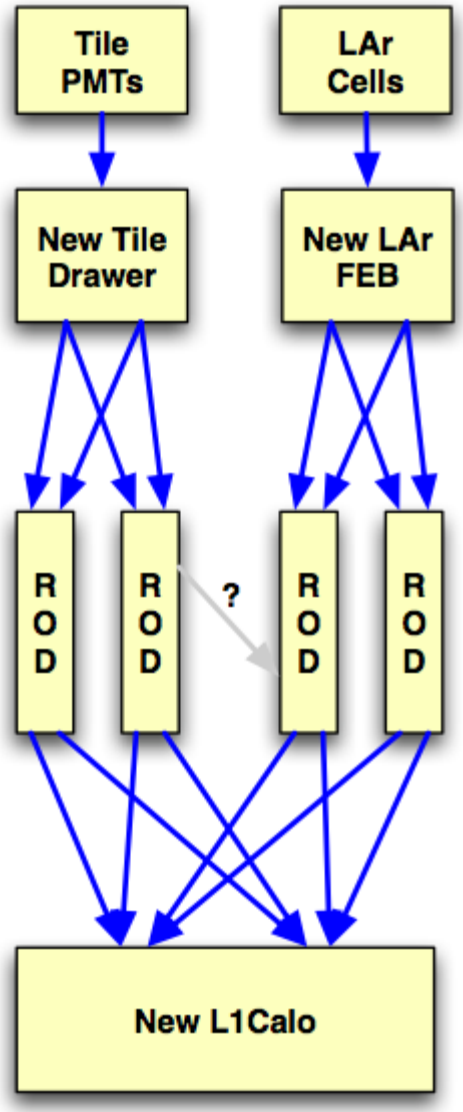


Grouping of FE cells constrained by calo geometry

Add remapping board on LAr FEB to regroup cells on links?

Regroup inputs to RODs? Split/merge fibre ribbons?

Little data transfer between RODs? Regroup towers on links to L1Calo

Mesh of links to L1Calo (duplicate data for phi fanout)

Any remaining mapping issues resolved by firmware

# LAr Front End Board Layout

- Diagram shows eta,phi sizes of FEBs in different regions
  - Sketched on 0.1*0.1 tower grid
- Barrel, endcap & FCAL have many different geometries between (and within them)
- Transition regions span boundaries in both eta & phi
- Bring all layers to one ROD requires splitting some FEBs between two or four RODs
  - Is this a problem (in principle)?

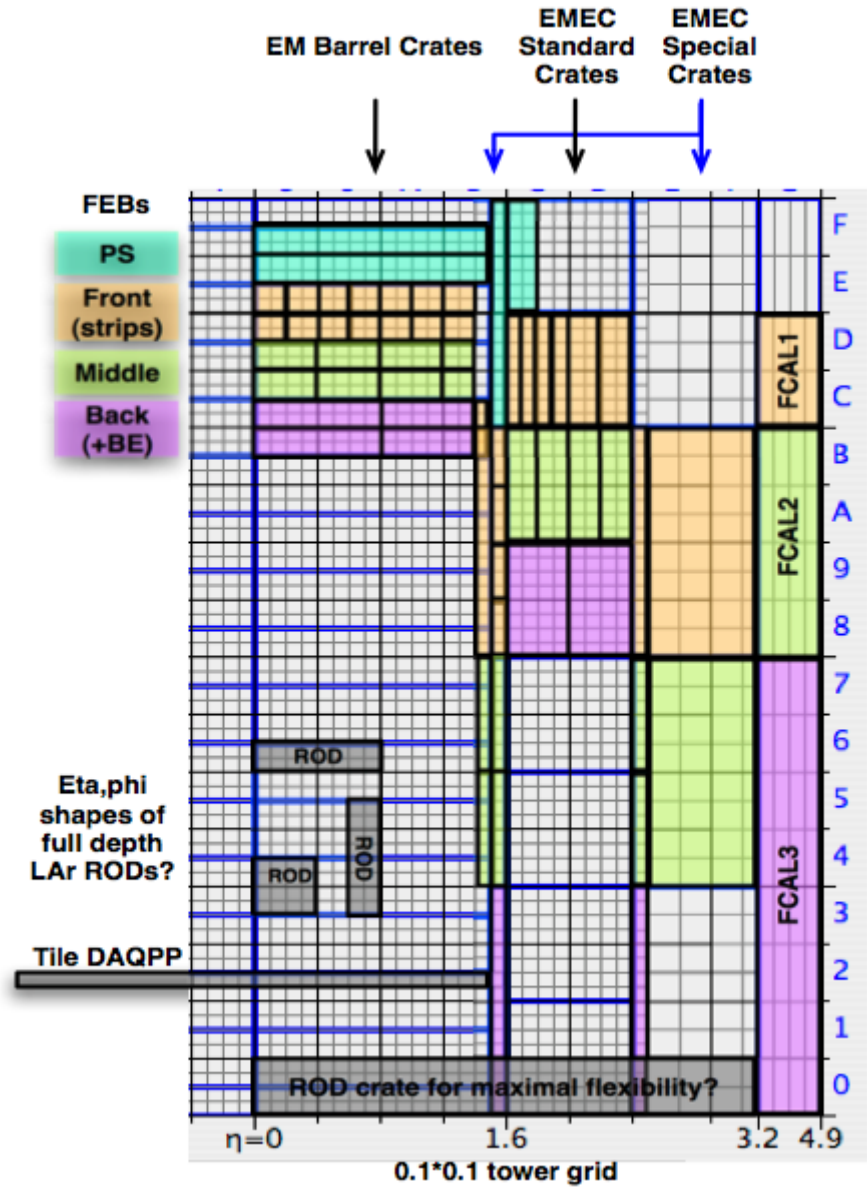

# ATCA-based ROD Crate?

- New crate: new architecture for control/configuration?
- No crate CPU or control bus
- Separate network and TTC++ connection to each ROD
- Flexible and scalable set of PCs to configure N RODs/PC
- Different TTC partitions can (but need not) share crates
- Can run separate standalone partitions for calibration
- Many configurations possible

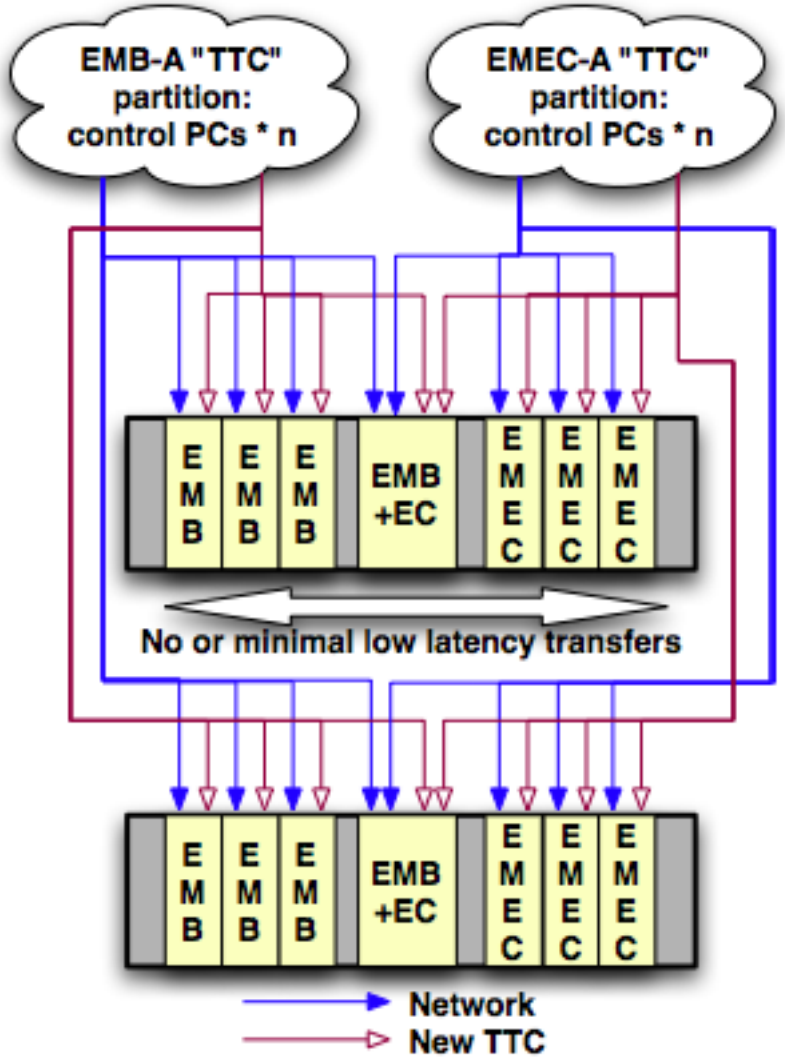

# ROD Issues (1)

- Two TTC partitions in one ROD at boundaries
- Tile baseline (DAQPP) has all four partitions:
  - Q1 (curiosity): how to run partitions independently?
  - Q2 (request!): can LAr do the same?

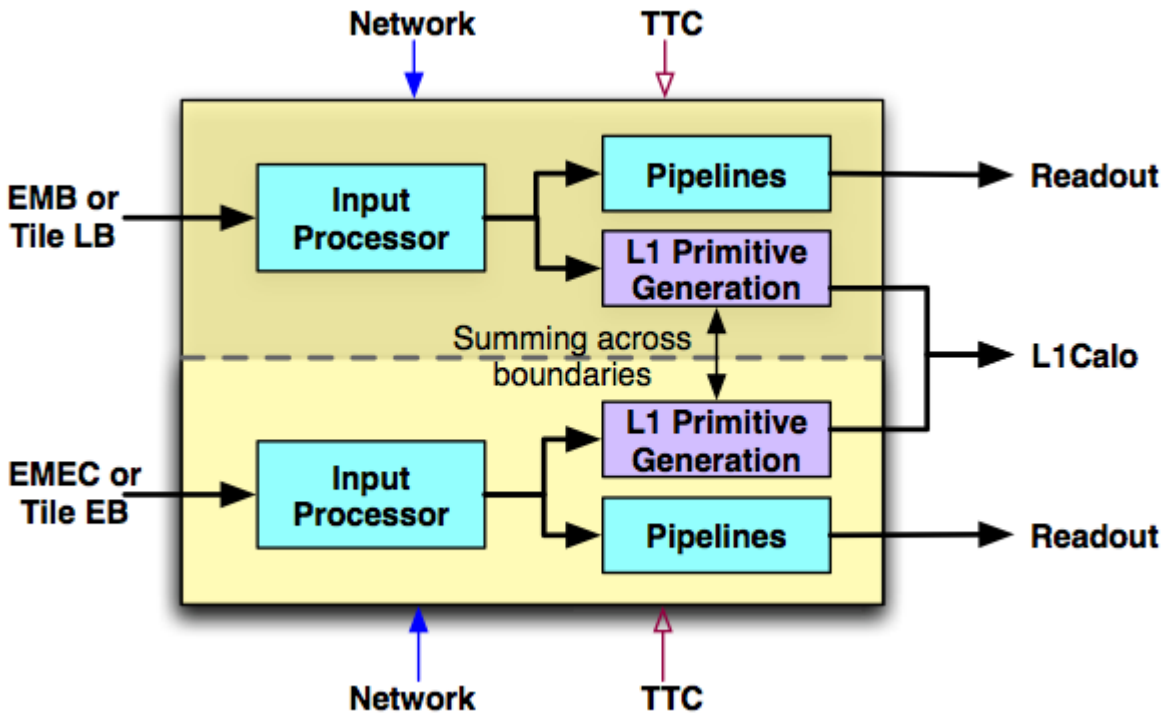

# ROD Issues (2)

- Granularity of FPGAs:
  - Probably want little or no transfer between chips?
  - If so, N links input to one chip define the eta,phi space of n links output to L1
  - Can all depth layers be really brought together?
    - Depends on cells/link, links/chip, chips/ROD
- Eta,phi space covered by whole RODs:
  - Any need to transfer data between RODs requires congruent eta,phi spaces covered by those RODs
    - Eg for Tile RODs to send crack scintillators to EMB/EMEC RODs
- Sparseness at higher eta
  - Changing ratio of input cells per "tower" with eta
  - Underutilised RODs or reconfiguration of input:output links

# Summary

- (Much) more simulation and thought required to:
  - Identify optimal and workable algorithms
  - Derive viable L1Calo architecture in more detail
    - Whats in a tower?
  - Discard any unnecessary "worst case" scenarios
- Need to discuss implications of FE and ROD layouts
  - What is desirable/acceptable/undesirable/unacceptable?
  - Keep each other in touch with proposed changes
  - Responsibility for L1 primitive generation??

# Backup Slides

# EM Endcap Geometries

- Seven different layouts between eta=1.4 and eta=3.2
- Many different ways cells are grouped into FEBs
  - Always(?) by layer
- NB two granularities in the EM barrel
- One in the FCAL
- (Plus similar in the hadronic layer)

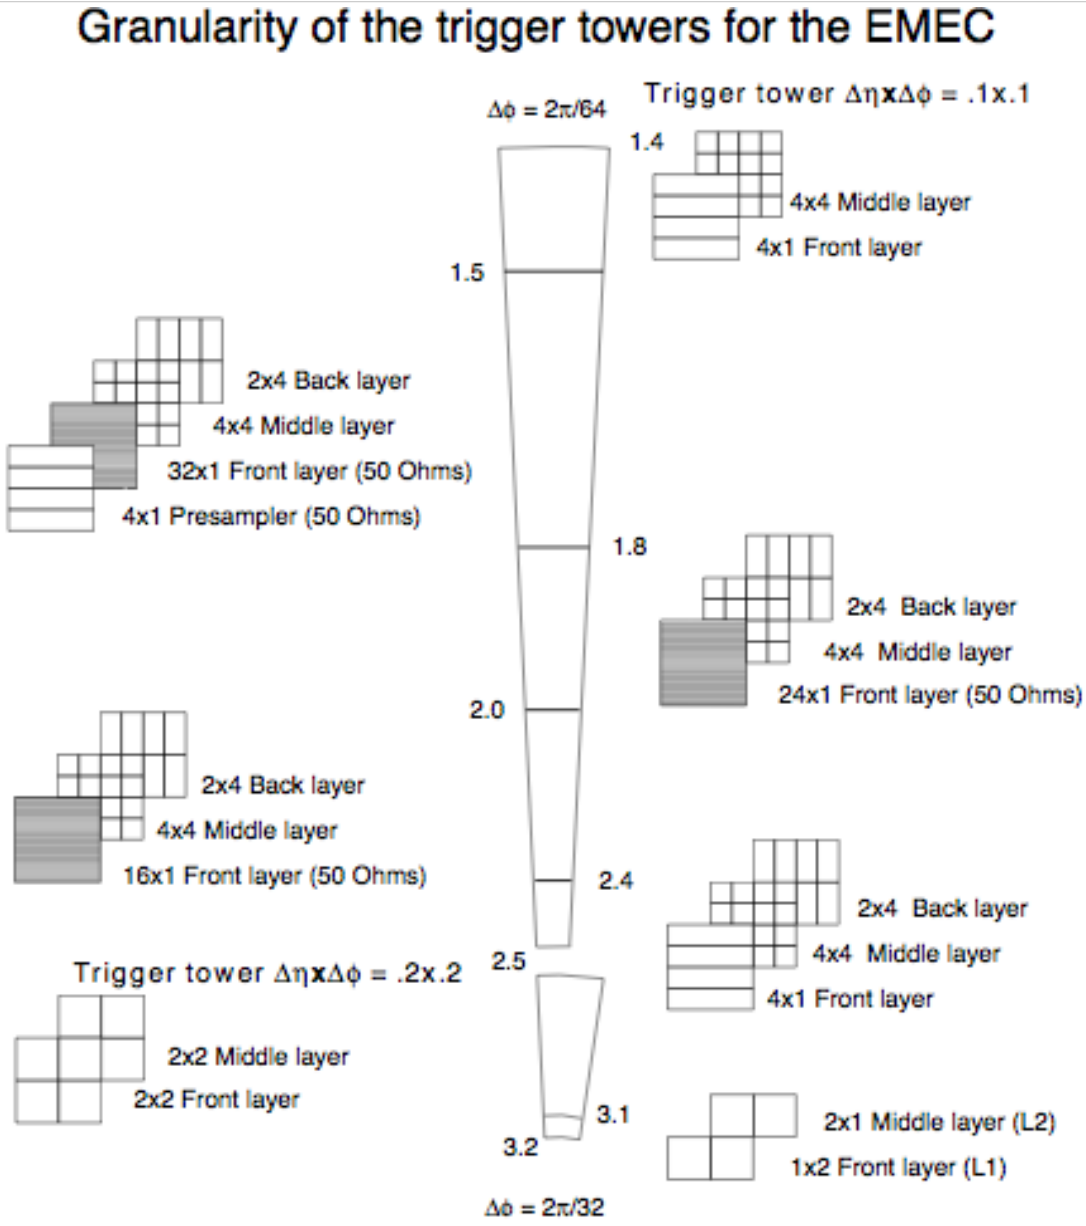

# Present Mapping Stages

- Many stages
- Lots of patch panels
  - Humble TCPP is ~2Gbit/s remapping device with ~0 latency and power!
- Easy areas regularised in one step at receivers
- Tricky areas needed many successive steps
- Never really managed it with the FCAL

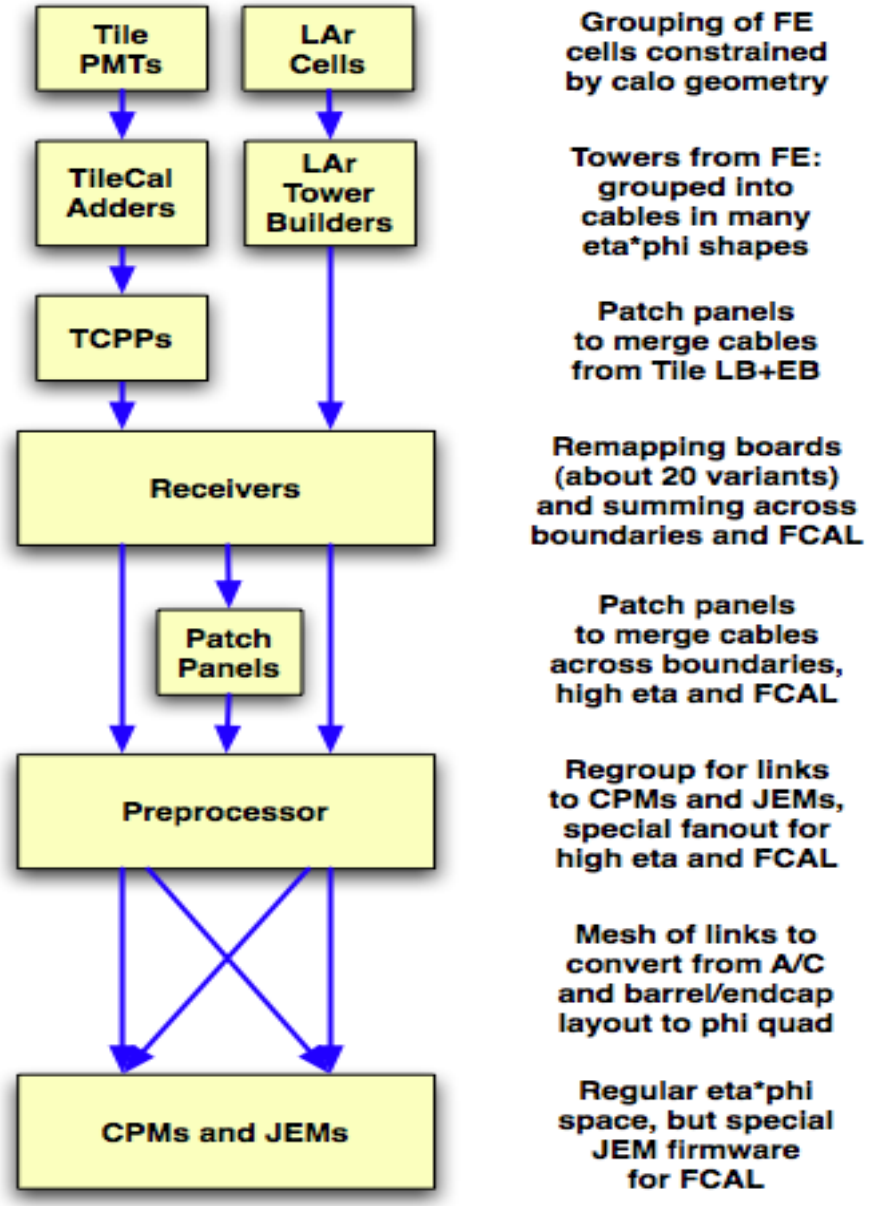

# Aside: The FCAL Story

- FCAL was late entering the trigger design
- (Partly) As a result it needs
  - special summing on the Receivers
  - special patch panels before the PPM
  - special summing on the PPM
  - special PPM outputs for phi fanout to the JEM
  - special firmware in the JEM
  - special software in mappings and especially in graphics
- Less than 0.4% of L1Calo towers
- Between 10 and 100 times that in design effort!
- Hope for less than that next time
  - Though it will always be a difficult area