# L1Calo Upgrade Mapping Issues

## thoughts provoked by the recent meeting

Murrough Landon

16 December 2008

- Background to mapping issues
- Reminder of present layout
- Assumptions for upgrade
- ROD Crates and "Network RODs"
- Ideal future layout?
- Summary

# Implications of On-Detector Digitisation

- Digitisation on detector seems to be favoured
  - This has many advantages
  - But will bring trigger and readout closer together
- Present L1Calo relies on towers built by calo front end electronics
  - Result of much detailed and careful work by calo groups
  - Thereafter any additional reorganisation, fanout, grouping is handled by the trigger
  - The readout branch is currently able to keep neatly aligned with detector geometry from front ends to RODs
- Upgrade needs to include the trigger requirements
  - Little simulation: we dont yet know what we will need
  - So start by assuming the worst – maybe it wont be so bad!

# Perspectives

- Different perspectives in LAr/Tile and L1Calo
- L1Calo sees EM and hadronic layers, not LAr and Tile
- Current calorimeter layout typically has N:1 mappings
  - from channels to FE boards, FE boards to crates/drawers, FE boards to RODs, RODs to crates and to TTC partitions
- In contrast, L1Calo and Receivers see multiple TTC partitions in the same crate and the same module
- Combining the trigger into the main readout will require importing some of this mixing and merging into the organisation of the calorimeter readout
  - Will be unfamiliar and uncomfortable (or impossible)
  - Try to find ideas to make it easier or even attractive!

# The Basic Mapping Problem

- Barrel, endcap & FCAL have many different geometries between (and within them)
- Transition regions span boundaries in both eta & phi
- Diagram shows LAr front end crate layout in the EM layer
  - Groupings by layer in FEBs not seen by present L1Calo towers
  - Will all need to be done again...

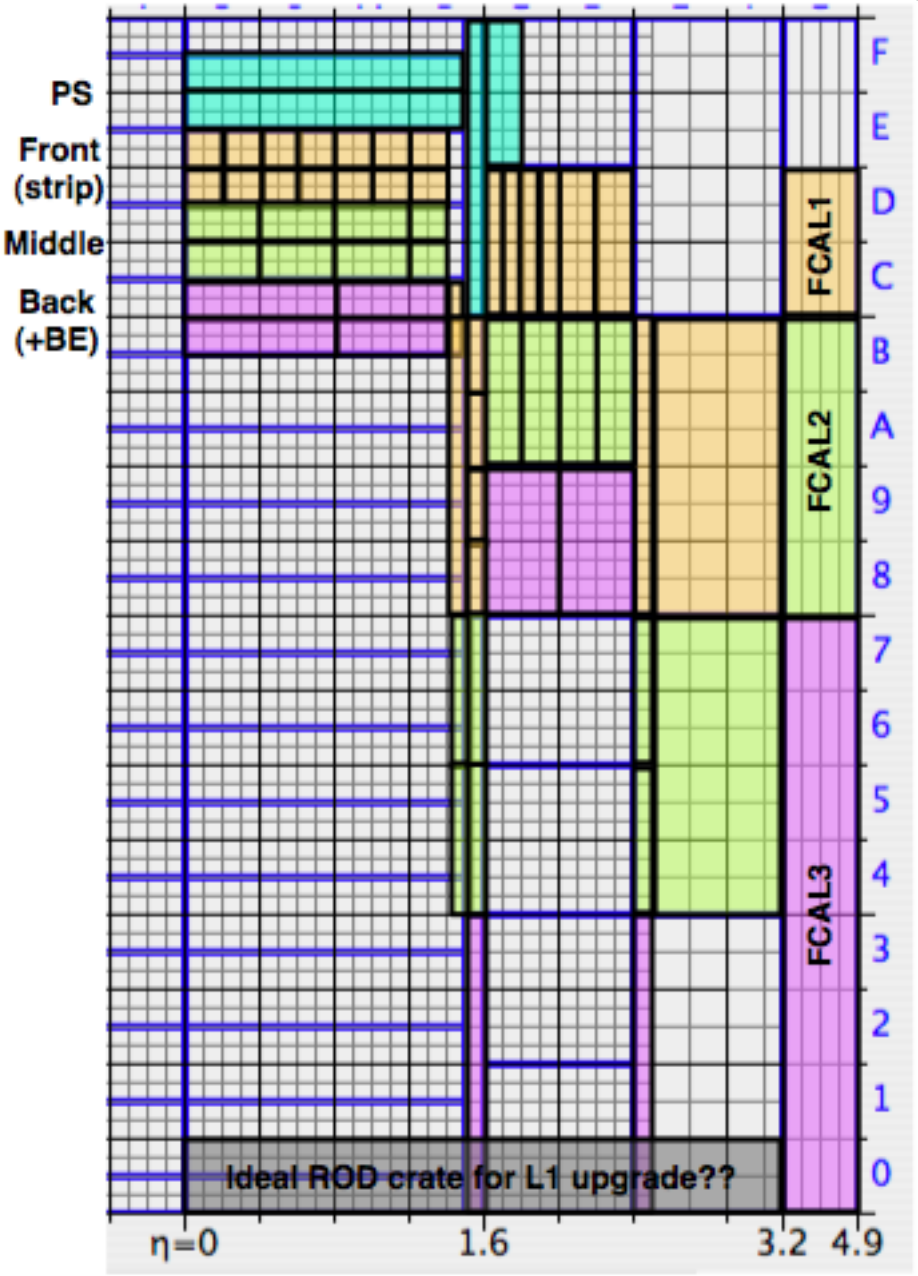

# EM Endcap Geometries

- 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)

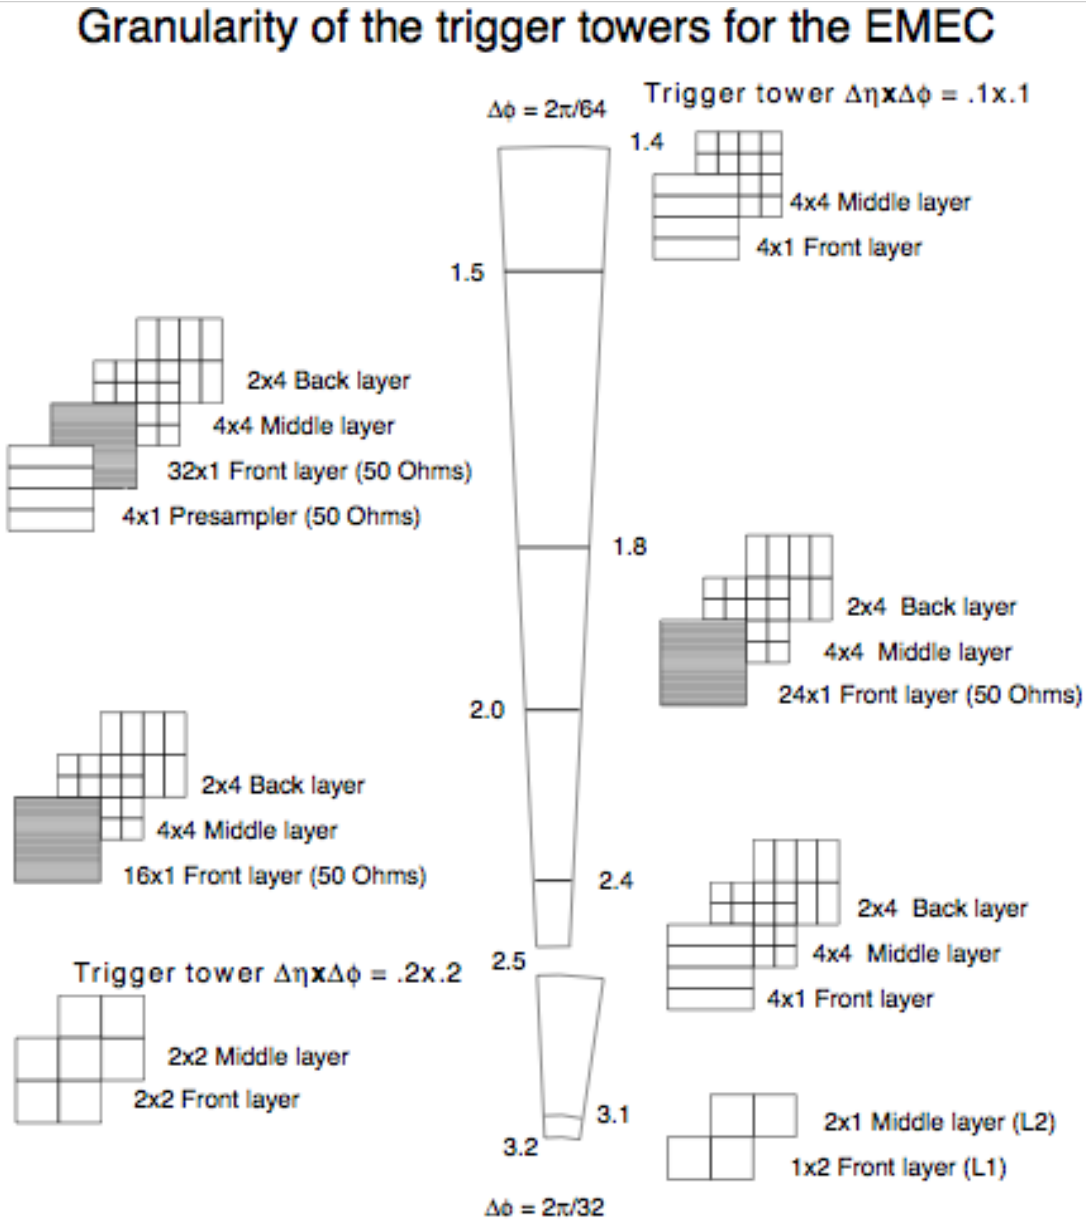

# Present Mapping Stages

- 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

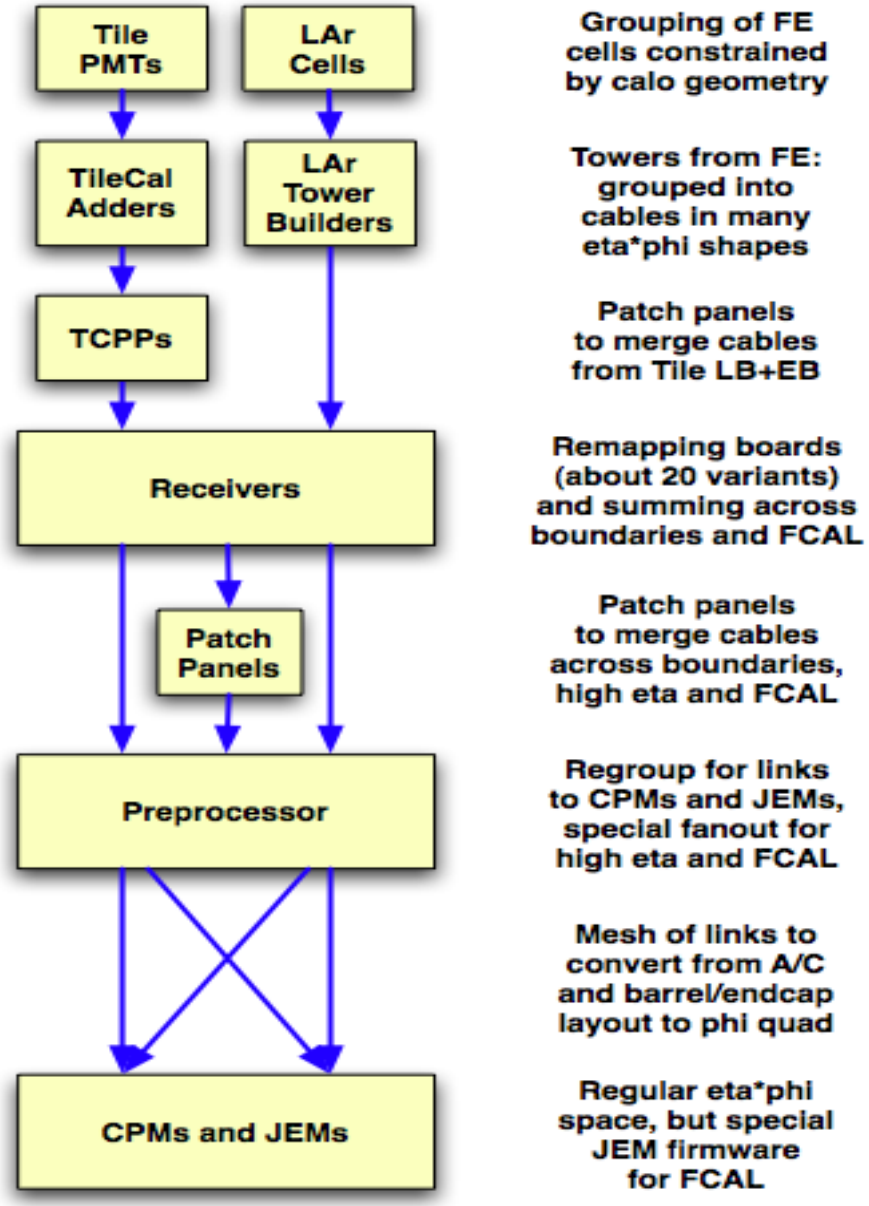

# Aside: The FCAL Story

- 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

# Lessons from the Past

- Lessons from existing L1Calo
  - Worry about the difficult areas early in the design process
    - It only gets worse later
  - Reorganising is easier with fewer (mini)towers per link
    - But obviously requires more links…
    - Trade off with bulk, ease of installation, space, efficiency
  - Do as much as possible at the first stage in the chain
    - Irreducible constraints from calo geometry will hit later
- Data organisation: order of preference
  - In same chip, on same module, in same crate (parallel data), in same crate (serial data), between crates (assume serial)
    - L1Calo avoids the last option by data fanout at source
  - NB if the data is from different TTC partitions, the calo preferences will probably be exactly the opposite!

# Assumptions: Feature Extraction

- Extract L1 features on the RODs
  - Otherwise massive number of links out of EM RODs
  - May need to extract L1 "features" using full granularity
    - Not subsequent processing of lower granularity sums
  - Especially across boundaries
  - Maybe even need full EM+hadronic depth?

# Assumptions: Latency

- Latency may still be tight
  - If L1Muon doesnt change it will still have ~2μs latency
  - Should upgraded L1Calo aim for the same to allow maximum time for L1Track trigger and fancy topological stuff?
  - High speed Serialisation/Deserialisation steps add a lot
    - EMB->EMEC and LB->EB not so bad as barrel signals come earlier
  - Minimise number of stages, maximise reorganisation/stage
    - Ideally without extra serialisation/deserialisation

# Assumptions: 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 L1 feature extraction?
  - Add in crack scintillators? (Currently read out via Tile EB)
    - Any possibility of local detector upgrade 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

# Link Speeds and Granularity

- 1Gbit/s is 25 bits at 40MHz
  - Roughly one tower with Et and feature bits?
  - 6Gbit/s would easily cover 4 towers (or mini-towers)
    - Likely possible to cover calo eta phi space with 0.2*0.2 groups
    - Regularisation easier with fewer smaller granularity "towers"/link
- Viewpoints...
  - View expressed at upgrade meetings: dont design until the link speeds are known
    - Dont know the granularity until then
  - Alternative view: granularity, mapping and handling calo boundaries may be a real problem
    - Find granularity and layout that works and calo groups will accept
    - If link speeds increase, send more feature bits, write better trigger algorithms?! Or just ignore new zillion Gbit/s links!!

# Possible Upgrade Mapping Stages?

- Fewer steps available?
  - Unless we add latency with an additional reorganisation
- Start with FE boards
  - Add remapping board like on present receivers?
- Remap FE to ROD links
  - Signals in depth and across boundaries to same place
- Low latency backplane transfers in ROD crates?
- Remap ROD to L1Calo links
  - It might all work!

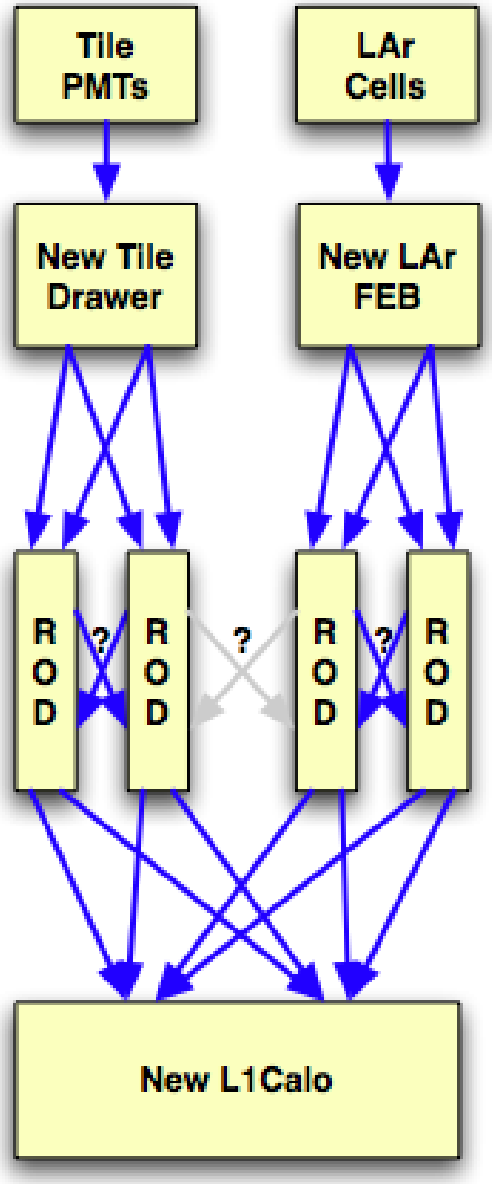


Grouping of FE cells constrained by calo geometry

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

Regroup inputs to RODs?

Data transfer between RODs? Regroup towers on links to L1Calo

Mesh of links to L1Calo

Any remaining mapping issues resolved by firmware

# Issues for Front End Electronics

- Fibre ribbons
  - LAr suggest 12 fibre ribbons: reduces ability to reorganise links from parts of one FEB to different RODs?
    - Consider using ribbons with few fibres? At least in tricky regions?
    - Or duplication of links? Or "patch panels" for reorganising ribbons?
      - May be useful if it is really required to split signals between crates
- Grouping cells onto links?
  - Can this be changed in different regions?
    - Add receiver-style (passive) remapping board?
    - Single FEB design but multiple routing personalities?
  - For reorganisability its best to have 2**n cells per link
    - Especially in phi(?)
  - Are there any constraints against sending signals from the same FEB/FEC to different ROD crates (in tricky regions)?

# Issues for RODs

- Collect all depth layers on one ROD
  - Suggested layout for EM barrel ROD (from one half FE crate) shows layers split between different RODs
    - All strips to one ROD, all other layers to a separate ROD
  - Is this just for neatness or is it a requirement?
  - L1 feature extraction would need all layers from some eta*phi space in one ROD
  - This may involve splitting links from one FEB (eg PS) between these two RODs
    - Any constraints on how FEBs get their TTC/configuration?
- Collect cells across boundaries in one crate or ROD?
- Ratio of bandwidth in from FE to bandwidth out to L1
  - EM: up to 30-60 cells per tower, Had: 3-5 cells per tower
    - ROD for HEC more similar to Tile than LAr EM?

# 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

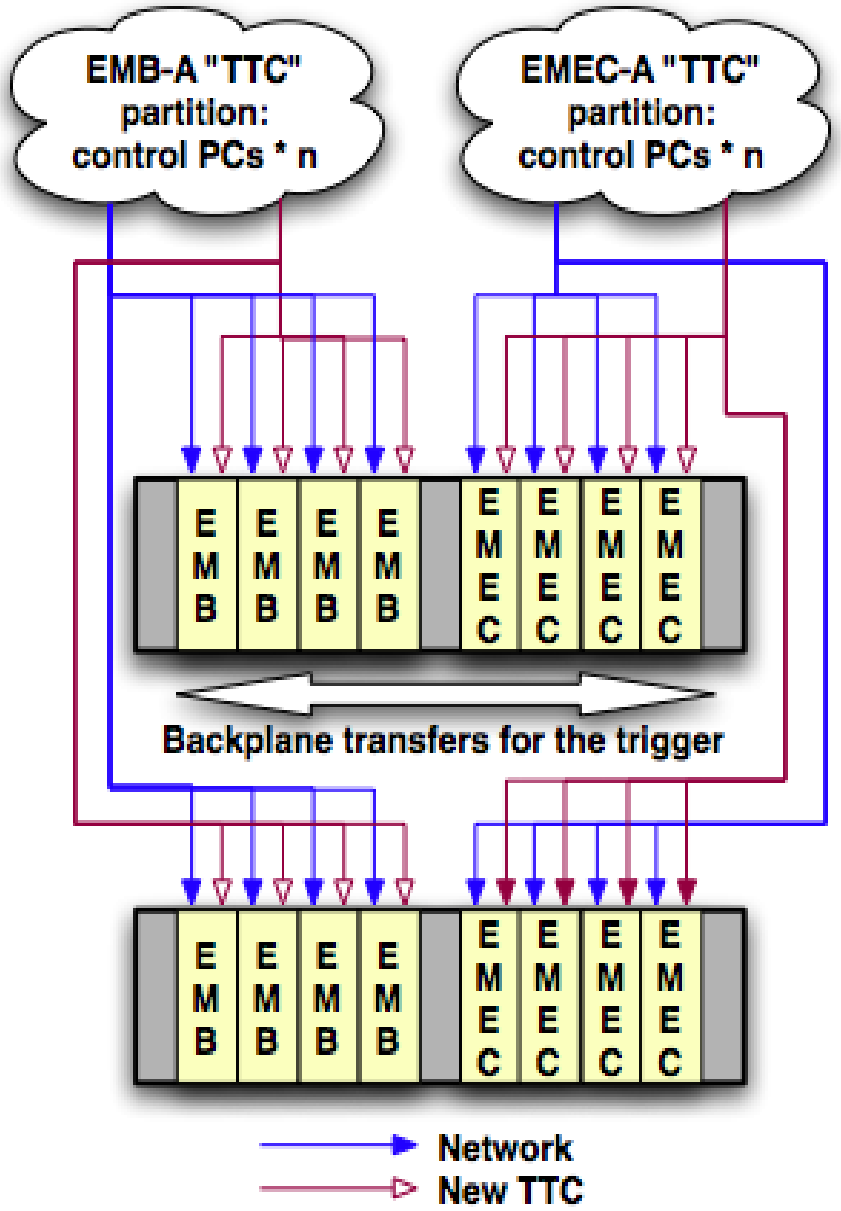

# Network ROD

- Assume feature extraction for L1 is on the new ROD
  - Otherwise huge number of links from new EM RODs
- May share parallel (ideally) or serial data (otherwise) with RODs in same crate
- Independent network and TTC++ per ROD

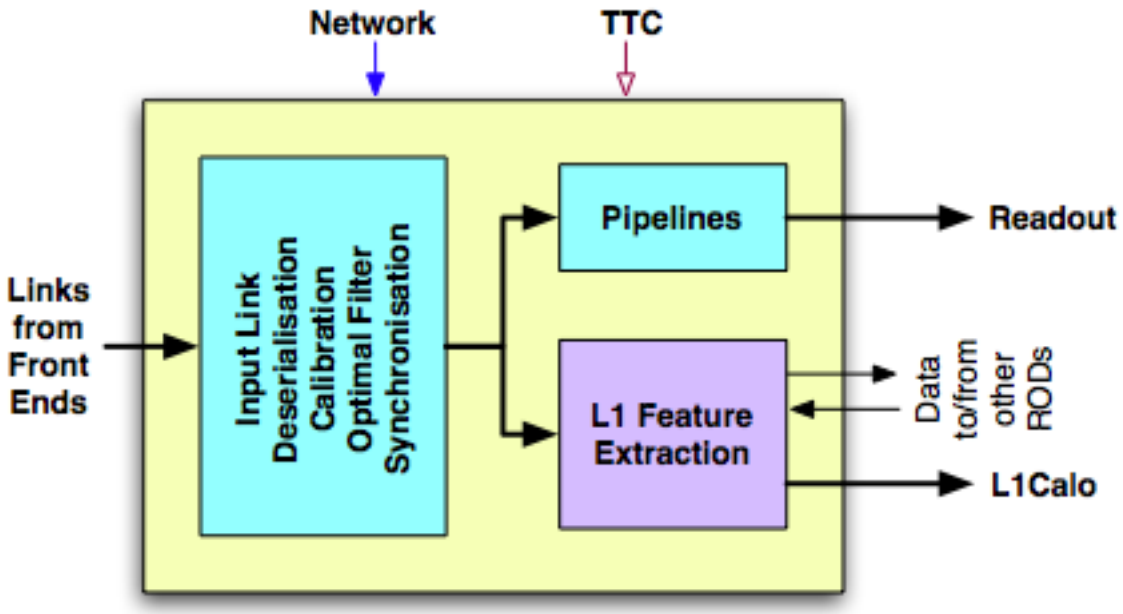

# Network Half&Half ROD?

- If RODs have their own network and TTC++ connections, why not two?!
- Two partitions on the same ROD?
- Run independently (standalone calib)
- Sum across LB/EB (or EMB/EMEC) before L1 feature extraction

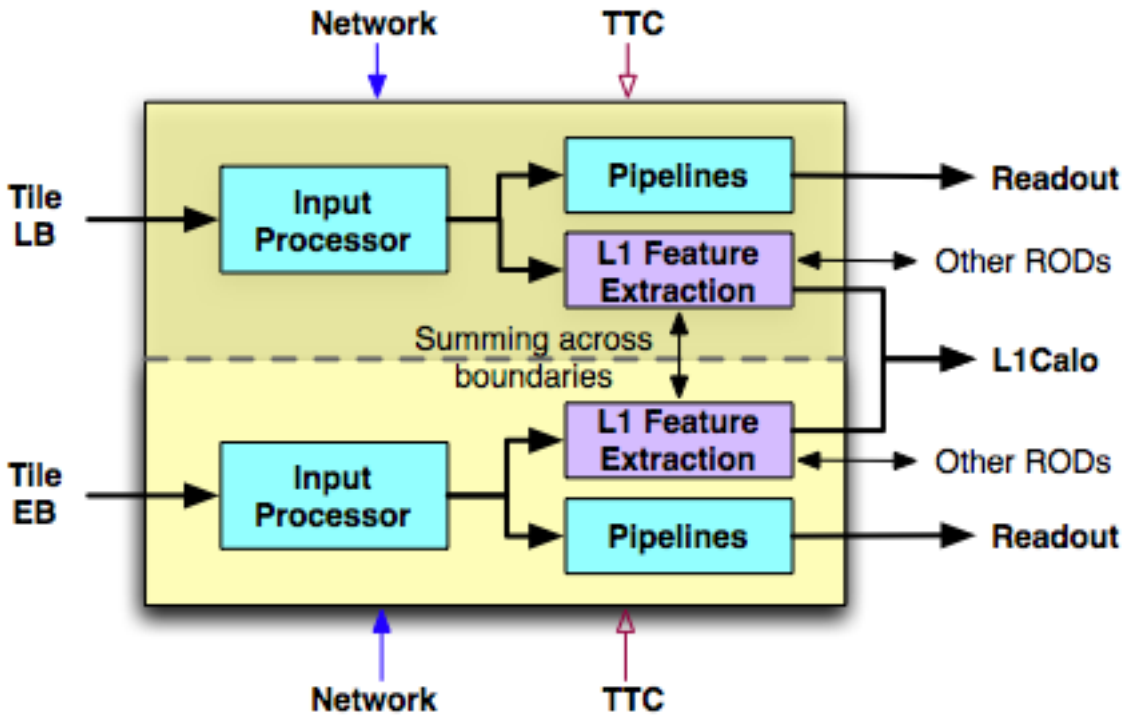

# L1 Dream Layout (Calo Nightmare?)

- Mixed ROD crate: 0.4 in phi, all eta (A or C), EM+had
  - LAr ROD talk: 2 RODs could cover 0.2*1.6 in eta*phi
  - Suggest 2 purely EMB RODs, 2 EMB+EMEC half&half RODs, 3 or 4 purely EMEC RODs, 1 Tile LB+EB half&half ROD, 1 HEC(+FCAL?) ROD: total ~10 RODs/crate, 5 TTC partitions
    - Most likely need special treatment for FCAL (unless 16 links/FEB)
      - Only fair if L1Calo shares some of the pain and the FCAL is the traditional instrument of torture!
  - Need 16 crates (8 racks) per A or C side
    - TCPPs+PPMs+RPPPs+Receivers currently use 7 racks per side
    - Leaves four racks for L1Calo processor and readout (CP+JEP+ROD)
  - Could do full depth feature extraction (backplane transfer)
    - Send duplicated features (0.4*0.4?) to Uli-style phi octant L1Calo?
  - Doubtless dozens of devilish details overlooked
  - Doubtless dozens of likely Calo objections ignored

# L1 Dream Layout(s)

- Ideal layout(?) shown at the top
  - few FCAL RODs in subset of all crates
- Alternative scheme below if it helps to have the endcap standard region handled differently
- Both have RODs with LB+EB and EMB+EMEC partitions

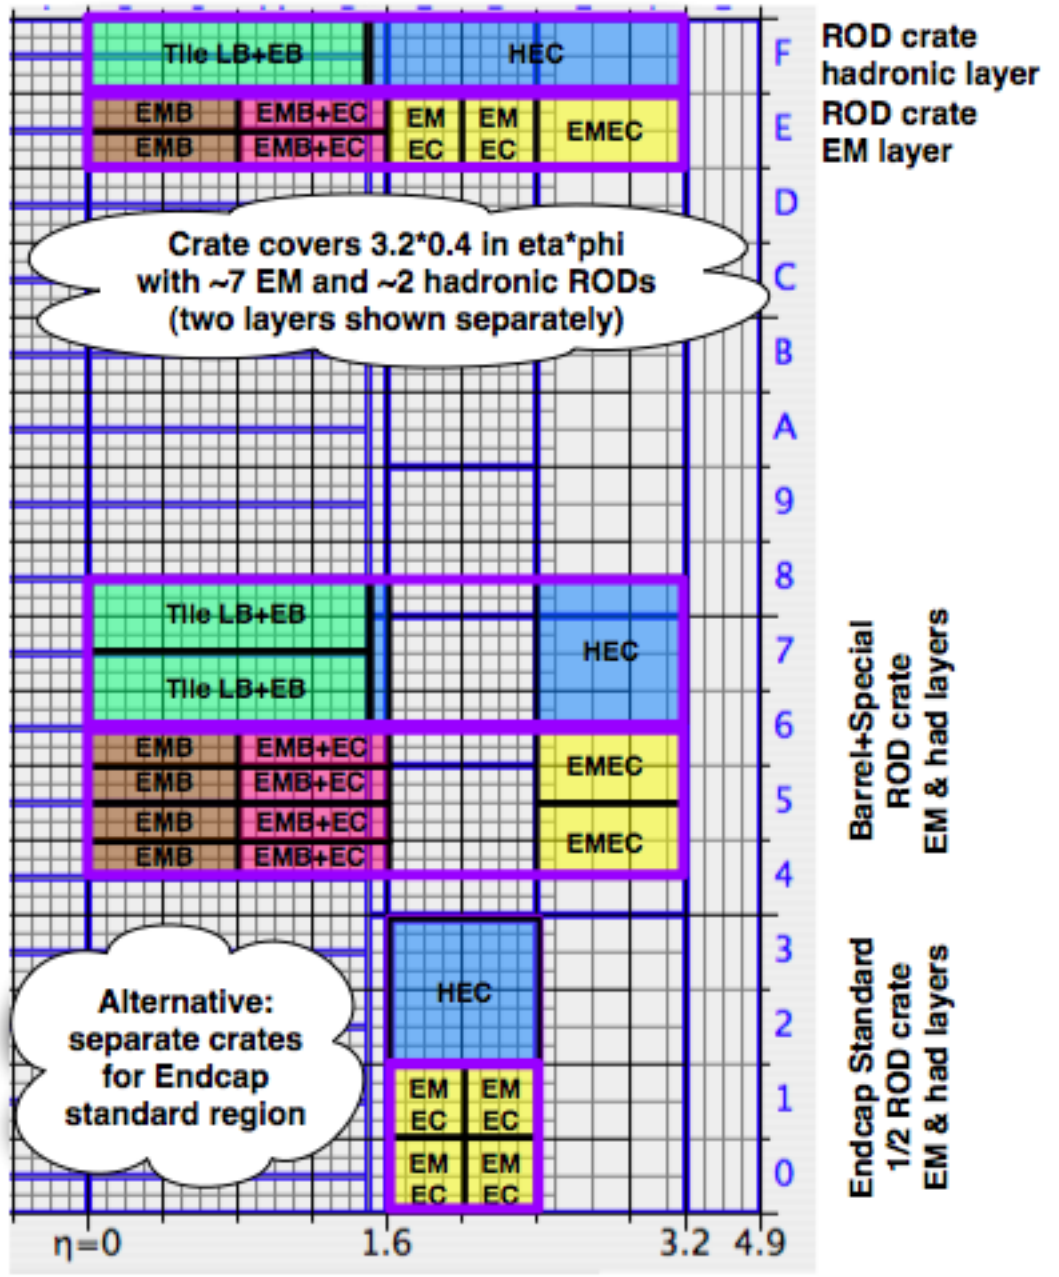

# Requirements For "Dream" Layout

- Need all LAr depth samplings together in same ROD
  - Diagram at workshop showed strip layer in separate ROD
- Crate covering 0.4 in phi means splitting fibres from some FEBs (eg EMEC/HEC 1.4<eta<1.6) across several RODs in different crates
  - How will FEBs get TTC? What about control/configuration?
- Such splitting requires 4*N links per FEB with same number of channels per link
  - May not be optimal use of huge required bandwidth
- Can expect disquiet about mixing TTC partitions and LAr/Tile RODs in the same crate and/or same module
  - Clearly must be able to run separately and standalone
    - Any fatal reasons why "network ROD" model would not allow this?

# Summary

- Exploring various ideas
  - In the absence of guidance from simulation
- Calo groups may find them:
  - Desirable?
  - Acceptable?
  - Undesirable?
  - Unacceptable?
- Get their feedback
  - And any alternative suggestions...