MUON & CORRELATOR TRIGGER ARCHITECTURE OPTIONS

Achieving Muon Universality?

D.Acosta

Science Drivers and Requirements

- Maintain sensitivity to electroweak scale physics at higher luminosity and pileup of the HL LHC
  - Report all standalone muon coordinates and momenta in convention to facilitate global correlation with tracks from the Track Trigger
    - The tracker will have far better P_T resolution for rate reduction
  - Incorporate additional HL LHC forward muon detectors to improve
    - Efficiency, redundancy, and improved standalone PT measurement
  - Maintain standalone muon trigger (without track combination) for sufficiently high P_T threshold
    - HL LHC is "only" 3-4X higher lumi, and we increased the max L1 rate
- Add sensitivity to new physics scenarios, i.e. acceptance to displaced muons and HSCPs from long-lived particle decays
  - Additional patterns/logic (displaced tracks, timing)
  - Expanded momentum assignment (vertex constrained and not)

Comment about displaced muons in endcap

- A large source of displaced muons for the endcap regions is beam halo, which is <u>not included</u> in our HL LHC simulations

(Thanks to Osvaldo Miguel Colin)

- Looking at rates in 2015, the halo muon rate per bunch was measured using the CSCTF (predecessor to EMTF).
- Scaling to 2500 bunches gives ~30 kHz !
- Algorithms for displaced muons will need to consider that other background source
  - non-pointing, but <u>not</u> parallel to beamline (e.g. displaced in x-y)

Muon-Track Matching Architectures

Original baseline

Another variant

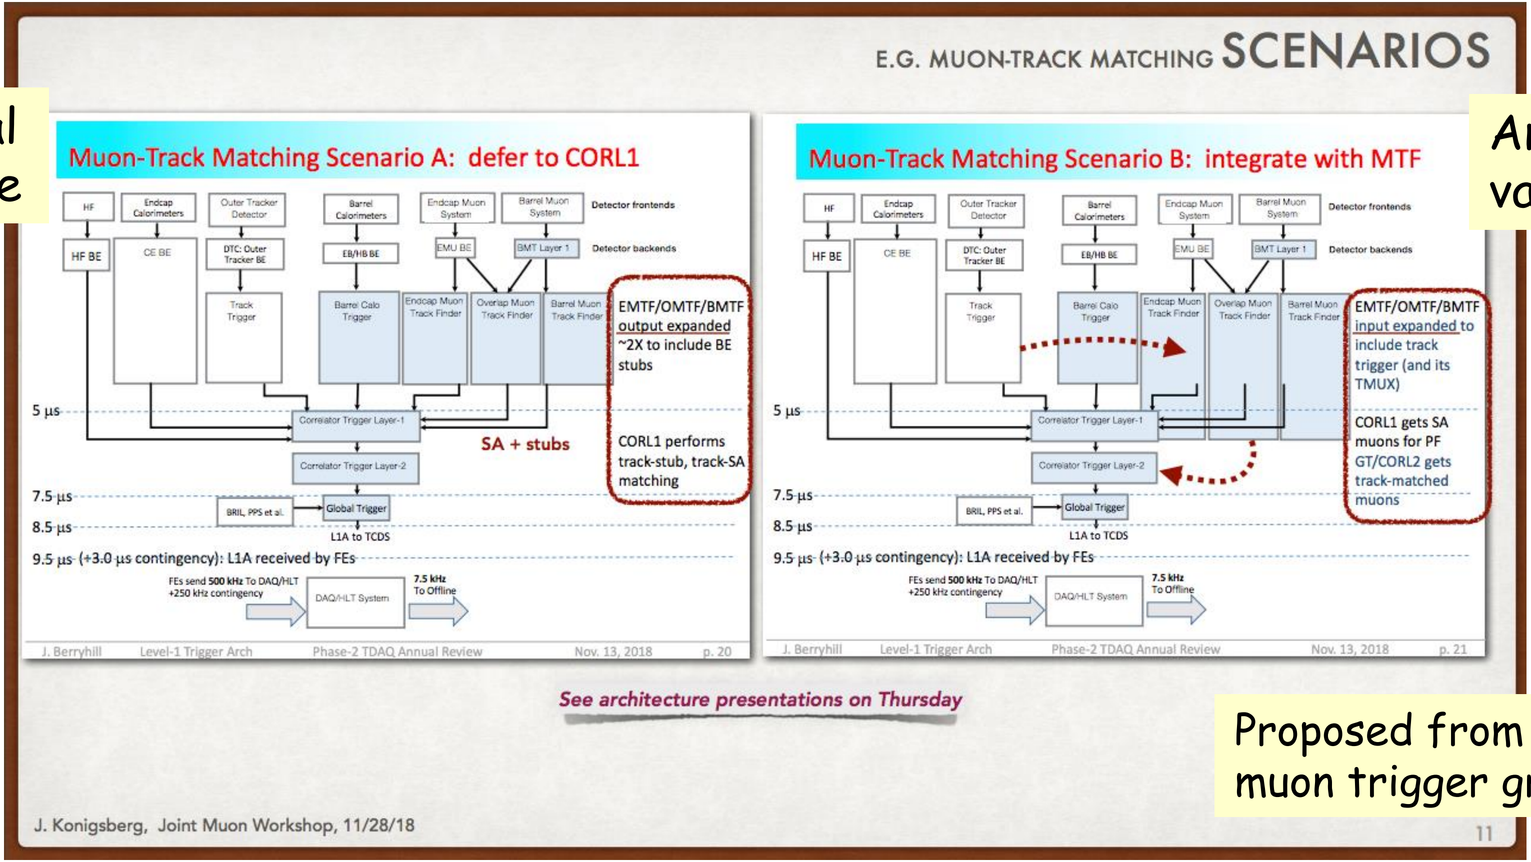

Proposed from barrel muon trigger groups

Endcap Muon Standalone Track-Finding

- Proposed baseline EMTF++ for Phase-2 does regional standalone (SA) track-finding in the endcaps ("Layer-1")
  - 12 sectors, 6 per endcap
  - Takes trigger primitives from existing or new muon detector electronics
- Recent studies show that "tracker muons" (track + muon stub matches) are interesting to include for performance reasons beyond just track + SA muon
- Could have EMTF++ send out muon stubs in addition to standalone muons to the Correlator for matching with tracks
  - If time multiplexed, send all stubs within a time window around current BX (HSCPs, monitoring,…)

Variation of Muon Track-Finding

- The proposed variation of this is to postpone SA muon track-finding and only send stubs from Layer-1 (concentrated to higher bandwidth links).
  - Then the Correlator layer finds both standalone muons and track+muon matches
  - Matches current barrel region architecture proposal
  - Allows smaller/cheaper FPGAs for data concentration (or re-use old hardware)
  - But adds to latency, and project cost (who does this concentration layer?)
    - One option is to use current Phase-1 EMTF as a concentrator for Phase-2

Input Link Count for full Endcap Coverage

- Prior assumptions, per 60° sector, with neighbor sharing (+N):
  - CSC:
    - 3.2 Gbps links from legacy MPCs (40+9)
  - RPC endcap via CPPF/New RPC electronics:
    - 12.5 Gbps links from LBs + RE1/3+RE2/3 (8+2)
  - iRPC forward:
    - RE3/1: 10 Gbps links (3+1)
    - RE4/1: 10 Gbps links (3+1)
  - GEM:
    - ME0: 10 Gbps links (3+1)
    - GE1/1: 10 Gbps links (8+2)
    - GE2/1: 10 Gbps links (9+2)
  - DT:
    - 10 Gbps links, 1/30° for MB2/1+MB2/2 (2+1)
- Summary for Regional SA tracking:
  - 95 links per 60° sector [Includes sharing]
  - 12 cards total, ~1100 total input links for ~7 Tbps bandwidth

Assumes no CSC data concentration

Assumes link concentration for RPC as current CPPF

iRPC electronics will be distinct from RPC Link board replacement

Provision to cover overlap with barrel

Endcap Data Concentrator Only

- Remove neighbor inputs (and barrel) for concentrator mode (no need)
  - CSC:
    - 3.2 Gbps links from legacy MPCs (40)
  - RPC endcap via CPPF/New RPC electronics:
    - 12.5 Gbps links from LBs + RE1/3+RE2/3 (8)
  - iRPC forward:
    - RE3/1: 10 Gbps links (3)
    - RE4/1: 10 Gbps links (3)
  - GEM:
    - ME0: 10 Gbps links (3)
    - GE1/1: 10 Gbps links (8)
    - GE2/1: 10 Gbps links (9)
- Summary:
  - 74 links per 60° sector
  - Could fit into existing EMTF system, for example
    - Optical link input limit is 84 x 10 Gbps for an MTF7 uTCA processor, and output is 24 x 10 Gbps

An Endcap Data Concentrator Layer-1

- Despite the large number of input links, the stub occupancy in the muon system is rather low even at HL LHC
  - The input links are mostly sending "zeroes"
- The first layer can zero suppress, concentrate, and time multiplex muon stub data for transmission to Layer-2
  - If this is an "intelligent" Layer-1, it also can perform standalone muon track-finding in the endcap in parallel to concentration (as currently)
- What constitutes a stub?
  - Could be a multilayer CSC LCT, or ME0 GEM stub
  - RPC and GE1/GE2 hits might need a coincidence with a CSC because of background noise rates (for now don't include in calculations)
    - Good case for forming "super-primitives" in Concentration Layer
  - Allows for finding "tracker muons" in Correlator (TT tracks + single stub)
- Concentrator Layer-1 is essentially equivalent to a fast DAQ path, but sending data for every BX rather than on L1Accept (as current EMTF)
- <u>How much data</u>?

CSC LCT Occupancy from Data

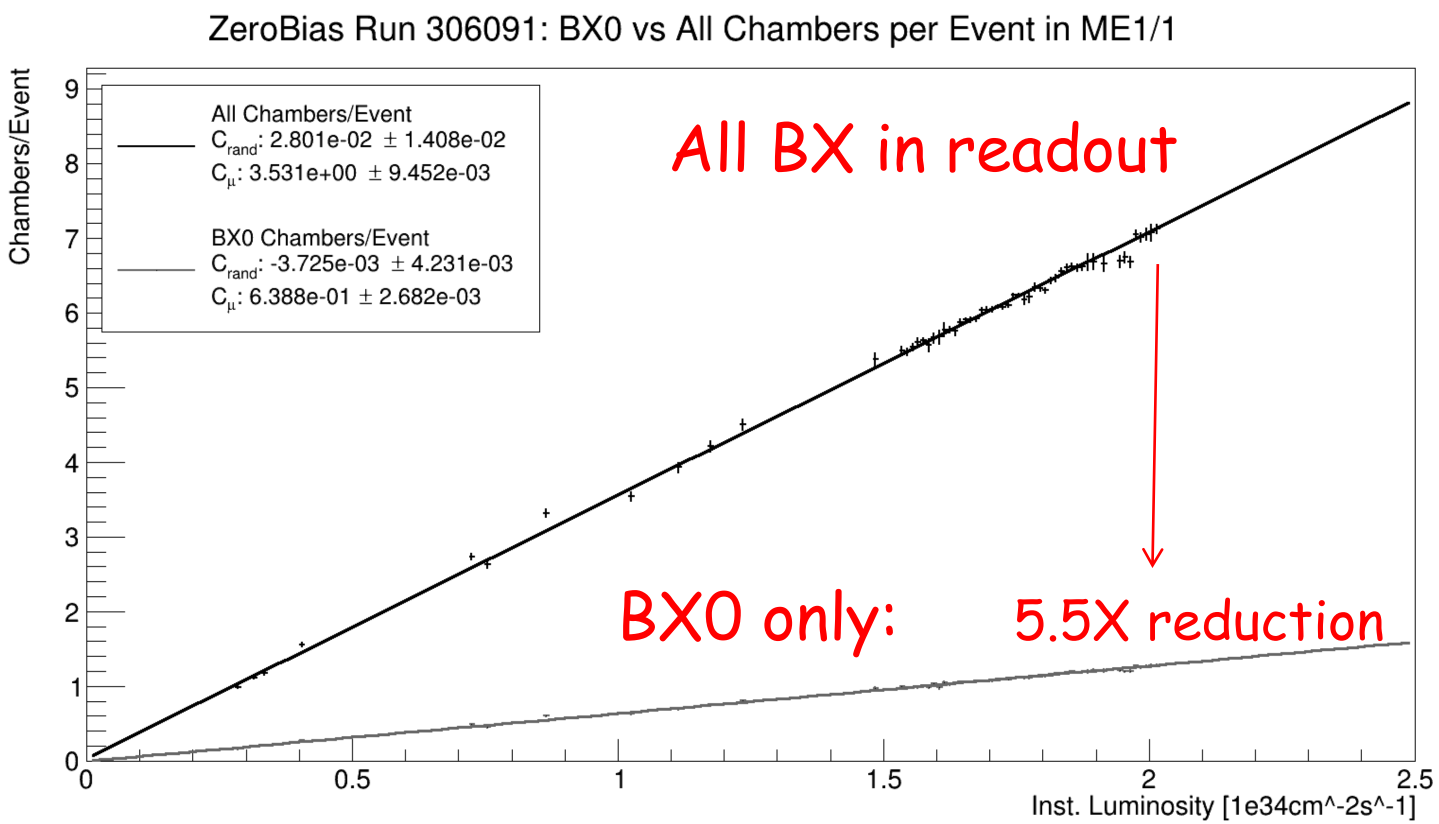

A.Aubuchon, Northeastern U

CSC chambers with ≥1 LCT

→ 6.5 stubs/BX at HL LHC

MC Study of CSC+ME0 LCT Occupancy

(ii) Without neighbors

Jia Fu Low,
PU=200 sample

chmb type	avg occupancy
ME1/1	3.810
ME2/1	0.682
ME3/1	0.387
ME4/1	0.303
ME1/2	0.270
ME2/2	0.076
ME3/2	0.092
ME4/2	0.118
ME1/3	0.013
ME0	3.653
TOTAL	9.404

MC yields 5.8 CSC stubs/BX, in good agreement with data (6.5)

ME0 adds 3.7 stubs/BX for full eta region (1.7 for $|\eta|<2.4$)

Sum is ~10 stubs/BX, or ~0.8 stubs/BX/sector

Layer-1 Output Data for Standalone Muons

"Intelligent" Layer-1

- Assume up to 3 muons / sector / BX
  - Any combination of prompt or displaced muons
- Start from Phase-1 Muon Trigger Detector Note:
  - 192 bits for 3 muons → 64 bits/μ (Phase-1)
- To this add another P_T word for displaced μ hypothesis, (+9 bits), plus more spare bits for future use
- Propose: ~100 bits/μ (a là Track-Trigger)
  - Up to 3μ /sector/BX at 64/66 encoding → 300 bits → 1x 16 Gbps link/sector
  - One output link ≥ 16 Gbps per sector per target destination for standalone tracks, even without longer time multiplexing latency
- How many stubs need be sent on (perhaps) additional links?

CSC+ME0 stubs per 60° sector

- For now assume one stub takes as much data as a CSC LCT sent from the MPC: ~40 bits
  - 32 bits per correlated LCT
  - Plus additional bits [7] for labeling which CSC LCT of the possible 5*18 = 90 CSC LCTs/BX/sector

Table 1. MPC to SP Data Format

Frame		15	14	13	12	11 … 08	07	06 … 00
	1	Quality			L/R	CLCT Pattern #	CLCT Pattern ID	
	2	BC1	BC0	SE	VP	CSC ID	AM	Wire Group ID

- Unassociated stubs from pileup are ~0.8/BX/sector
  - ≤ 2 @ 95% CL → 80 bits / BX / sector
  - If time-multiplexed, need to send stubs from a time window around current BX
    - Useful for HSCP trigger plus commissioning and operations monitoring to see stubs before an after in-time crossing
    - Including +/- 2 BX → ~4 unassociated stubs/sector (≤ 8 @ 95%CL)
    - → 8 stubs * 40 bits = 320 bits / 5BX / sector
- Allow stubs from up to 3 muon tracks per sector for a signal in event
  - Assume muons (and their segments) are in-time (1 BX, even if TMUX)
  - Assume each track can have 4 stubs
  - 3 tracks * 4 stubs * 40 bits ~ 500 bits

Endcap Data Bandwidth to Layer-2

~1 kb of data/sector

- Endcap muon data to send to Layer-2, per sector
  - Tracks: 300 bits (if done in Layer-1) → 12.5 Gbps
  - Stubs: 600, 800 bits (TMUX=1, >1) → 25, 33 Gbps ÷TMUX
  - 46 Gbps @ 64/66 encoding [33 Gbps stubs only]
- For a latency of 1 BX+serdes for transmission (e.g. TMUX=1):
  - Tracks + stubs: 3 x 16 Gbps links per sector, or 2 x 25 Gbps
  - Stubs only: 1 x 25 Gbps links, TMUX=1
    or 2 x 16 Gbps links, TMUX>1
- For a latency of 3 BX+serdes (e.g. for TMUX ≥ 3) ← Or even more data for longer latencies
  - 1 x 16 Gbps link per sector (or higher BW)
    - Or 1 x 10 Gbps (old EMTF concentrator, 8b10 encoding), 4 BX latency
- For entire endcap muon tracks + stubs to a target Layer-2 node
  - 36 x 16 Gbps @ 1 BX [24 x 16 Gbps or 12 x 25 Gbps, stubs only]
  - 12 x 16 Gbps @ 3 BX

Muon+Correlator Architecture, TMUX Model

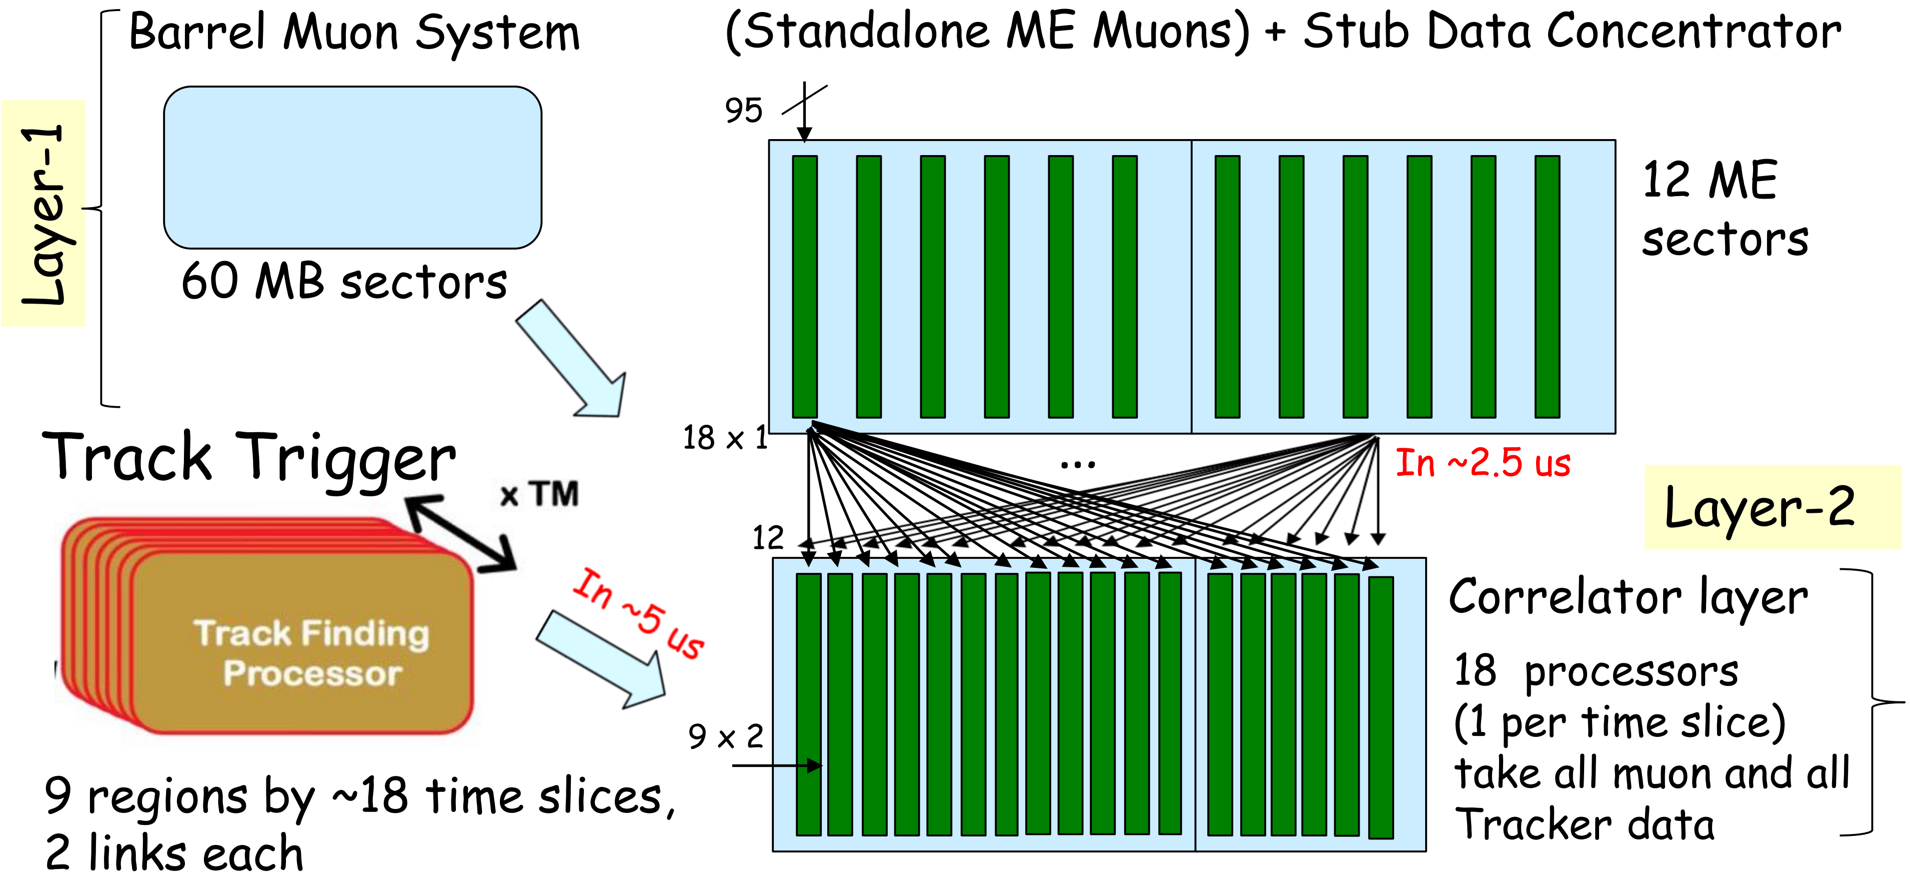

Muon Correlator Board Inputs, TMUX Model

- Track Trigger
  - 18 x 25 Gbps (9 regions by 2 links per time sample)
- Endcap Muon (New Layer-1)
  - 12 x ≥16 Gbps for 3BX + SerDes latency (stubs, or stubs+muons)
  - or 36 x ≥16 Gbps (24 x 25 Gbps) for 1BX + SerDes latency
- Endcap Muon (Re-use old EMTF as concentrator)
  - 12 x 10 Gbps, stubs only

options

- Barrel Muon
  - 60 x ≥16 Gbps (60 layer-1 processors)

Barrel muon dominates, could use its own concentrator layer?

- Total links per Concentrator (+standlone muon) module
  - 90 links for each of 18 cards (fits in APx card)
    - Does not allow +12 more links from endcap if needed, in APx
    - Does not allow links from other subsystems for correlation

Muon Correlator Board Logic and Latency

- Note that latency to receive muon data should be ~2.5μs, whereas that for the Track Trigger is ~5μs
  - Can receive (and buffer) all muon data for a given BX before tracks are available
  - Standalone muon reconstruction can start in advance of track matching and not add to overall latency
    - But now done globally, with TMUX, and not just regionally
      - Increased logic, but also increased latency available
  - Can immediately test tracks with all SA muons + stubs as tracks are received

With Dedicated Barrel Track-Finder/Concentrator

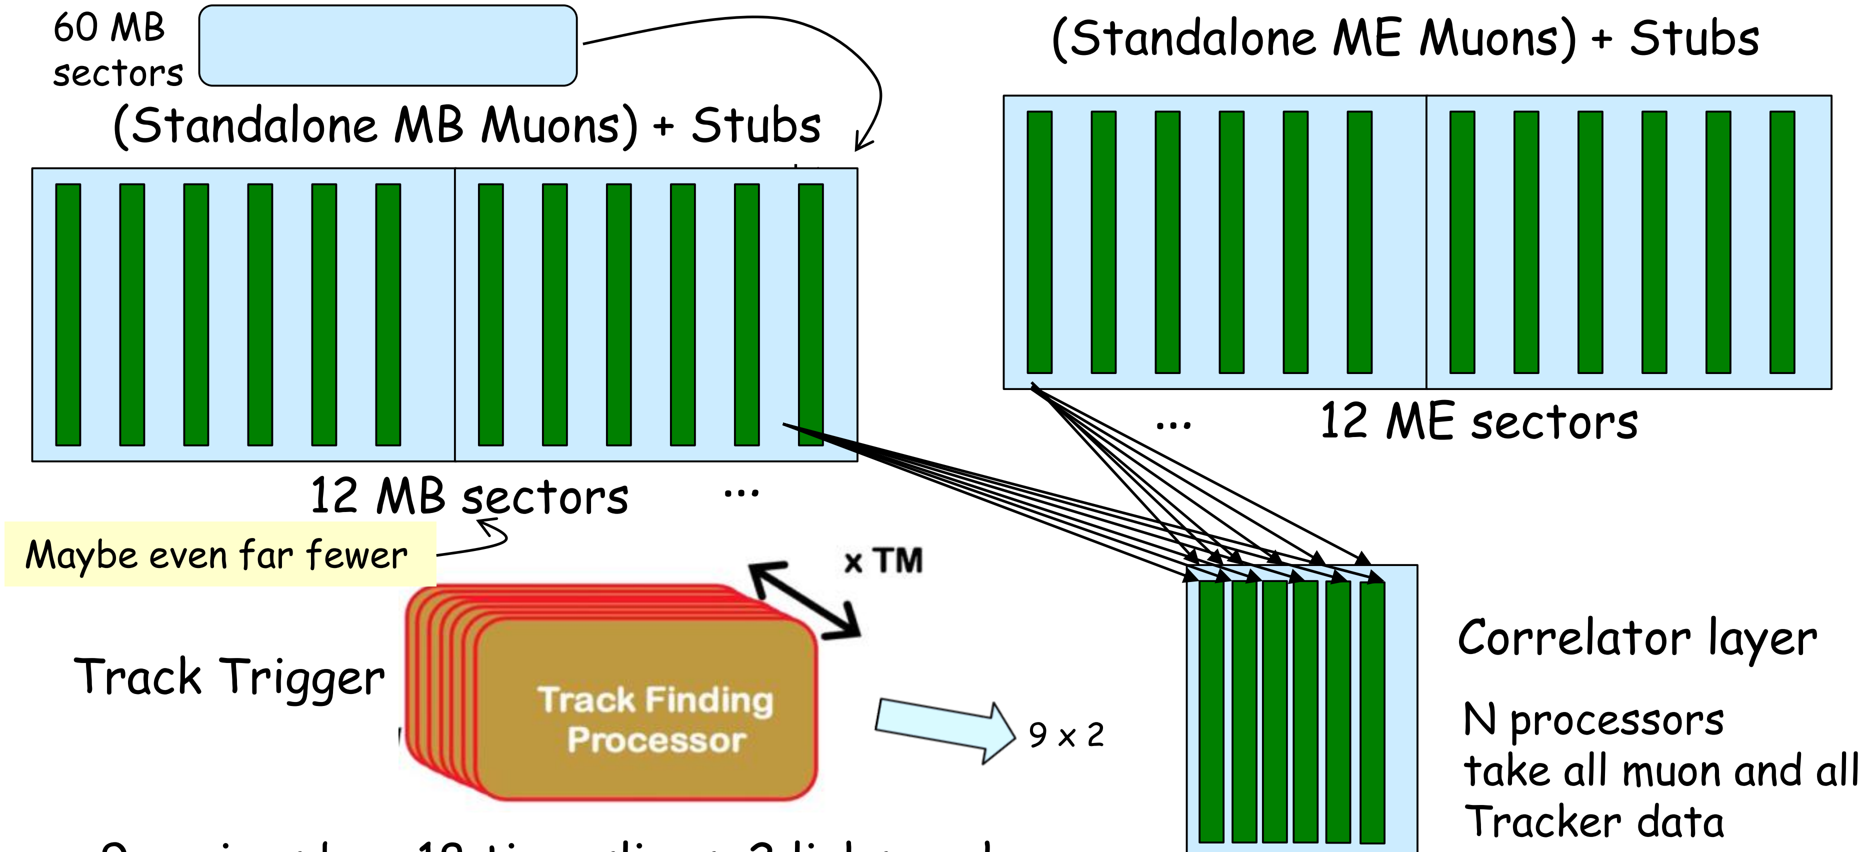

9 regions by ~18 time slices, 2 links each

Muon Correlator Board Inputs with MB Concentrator, TMUX Model

- Track Trigger
  - 18 x 25 Gbps
- Endcap Muon
  - 12 x ≥16 Gbps (or 10 Gbps)
- Barrel Muon
  - 12 x ≥16 Gbps (assume same as for endcap)
- Total links per Muon Correlator time slice
  - 42 links
- Could fit more than one time slice per APx card:
  - 9 cards (2 time slices) → 60 links for longer muon serdes (+3BX)
  - 6 cards (3 time slices) → 78 links for longer muon serdes (+6BX)
- Or could fit additional subsystem inputs for same 18 cards

Regional Track-Finding and Matching

Another alternative

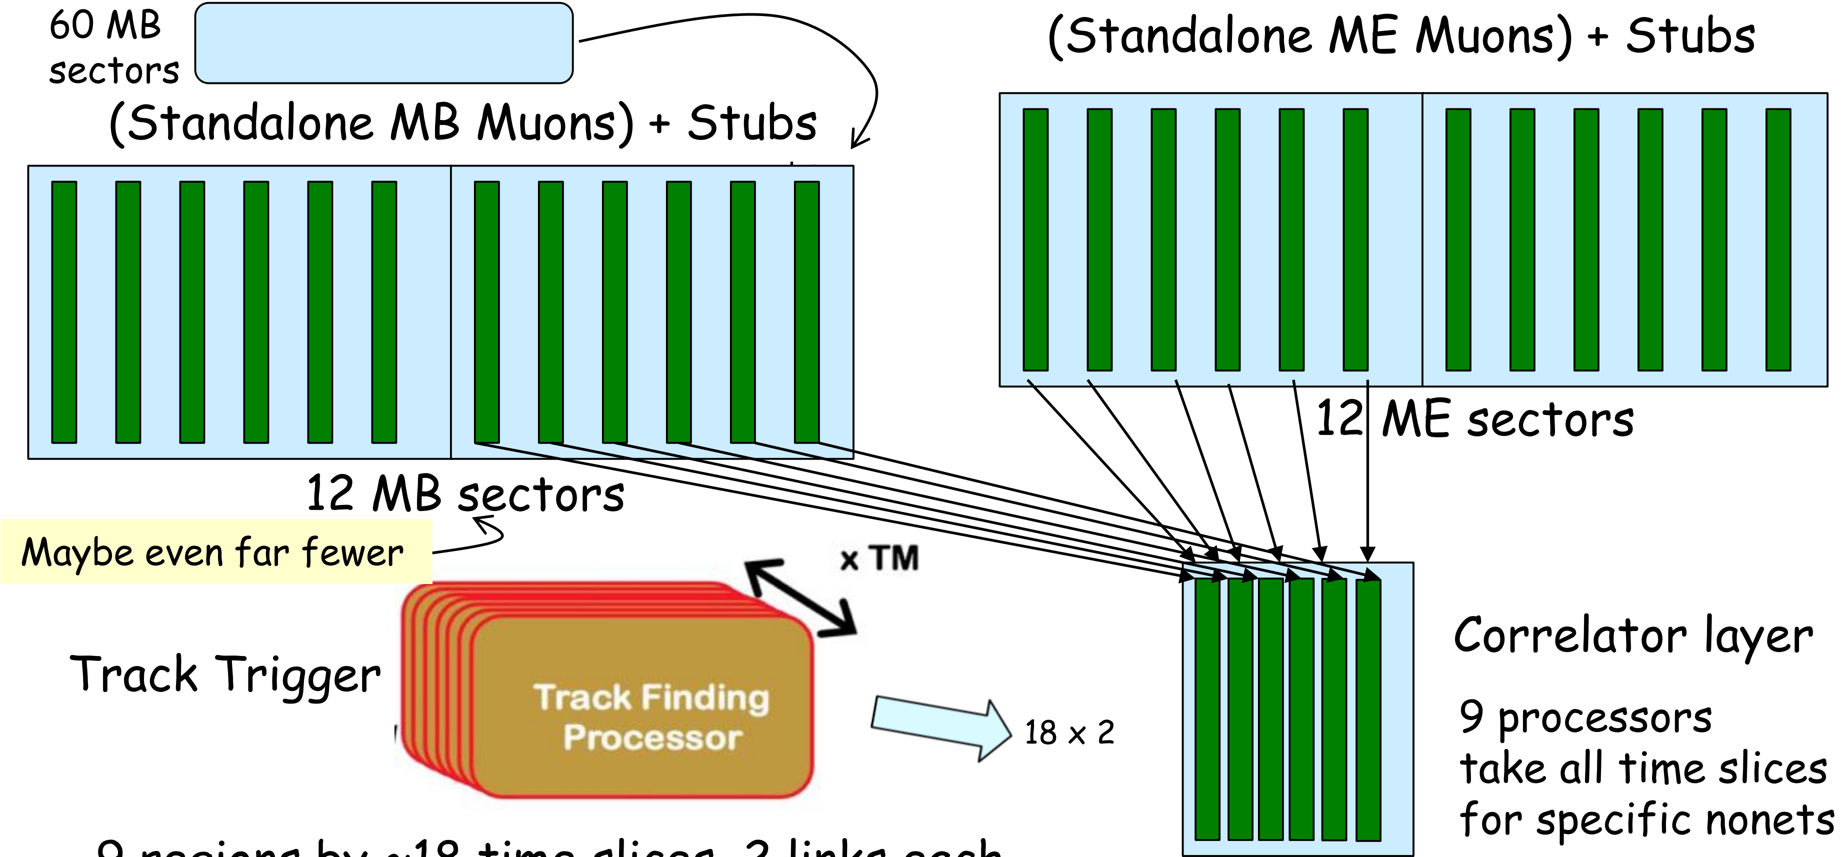

9 regions by ~18 time slices, 2 links each

Muon Correlator Board Inputs with MB Concentrator, Regional Model

Take all time slices into one board

- Track Trigger
  - 36 x 25 Gbps (18 time slices by 2 links per 40° nonet region)
- Endcap Muon
  - (<) 12 x 25 Gbps (up to 12 sectors, higher bandwidth to get stubs in 1 BX without TMUX)
- Barrel Muon
  - (<) 12 x 25 Gbps (assume same as for endcap)
- Total links per Muon Correlator time slice
  - (<) 60 links (+12 if sending SA tracks from endcap Layer-1)
- Target 9 APx cards (one per 40°)
- SA Muon Logic in Layer-2
  - Could limit SA muon reconstruction to relevant sector around Tracker nonet → move SA muon tracking also to 40° sector ?
- Correlator logic
  - Will need stubs from wider region of muon system (~180°), but covered by links above
  - Leaving tracker data time multiplexed implies correlator target would receive links from 9 regions per time slice
- Could even double region to 2 nonets (80°) for 5 APx cards
  - (<) 96 links (one board only gets a 40° region)

Architecture Conclusions, 1

- Could perform standalone muon reconstruction <u>and</u> track matching in one layer for barrel-overlap-endcap in a set of 18 cards → unifies muon trigger
  - I/O takes 90 links per card
  - But need to verify logic usage in FPGA to check if logic for entire detector can fit in the FPGA resources and in the latency
  - Requires rewriting EMTF algorithm from regional → global+TMUX
  - Requires group(s) to contribute an endcap concentrator layer, or re-use the Phase-1 EMTF
    - Could be used to form super-primitives
- Adding a BMTF Concentrator/TF Layer can reduce number of Layer-2 cards required for track matching, or accommodate additional subsystem inputs
  - Requires group(s) to contribute a barrel concentrator/TF layer

Architecture Conclusions, 2

- Having an "intelligent Layer-1 that finds SA muons in addition to stub concentration does not significantly increase the data BW to Layer-2
  - Still one output link per Correlator target
  - Reduces logic resources required in Layer-2, but adds to Layer-1
  - Would allow use of baseline regional EMTF++ algorithm
  - But breaks barrel-endcap symmetry unless barrel region does same
    - And would need to share data across boundary
- A regional approach (vs. TMUX) is also feasible for track+muon correlation
- Standalone muon reconstruction can take place before Tracker Trigger tracks are available for correlation
  - ~2.5 μs vs 5.0 μs, either in Layer-1 or Layer-2 (assuming SA latency ~ Phase-1 latency)

BACKUP