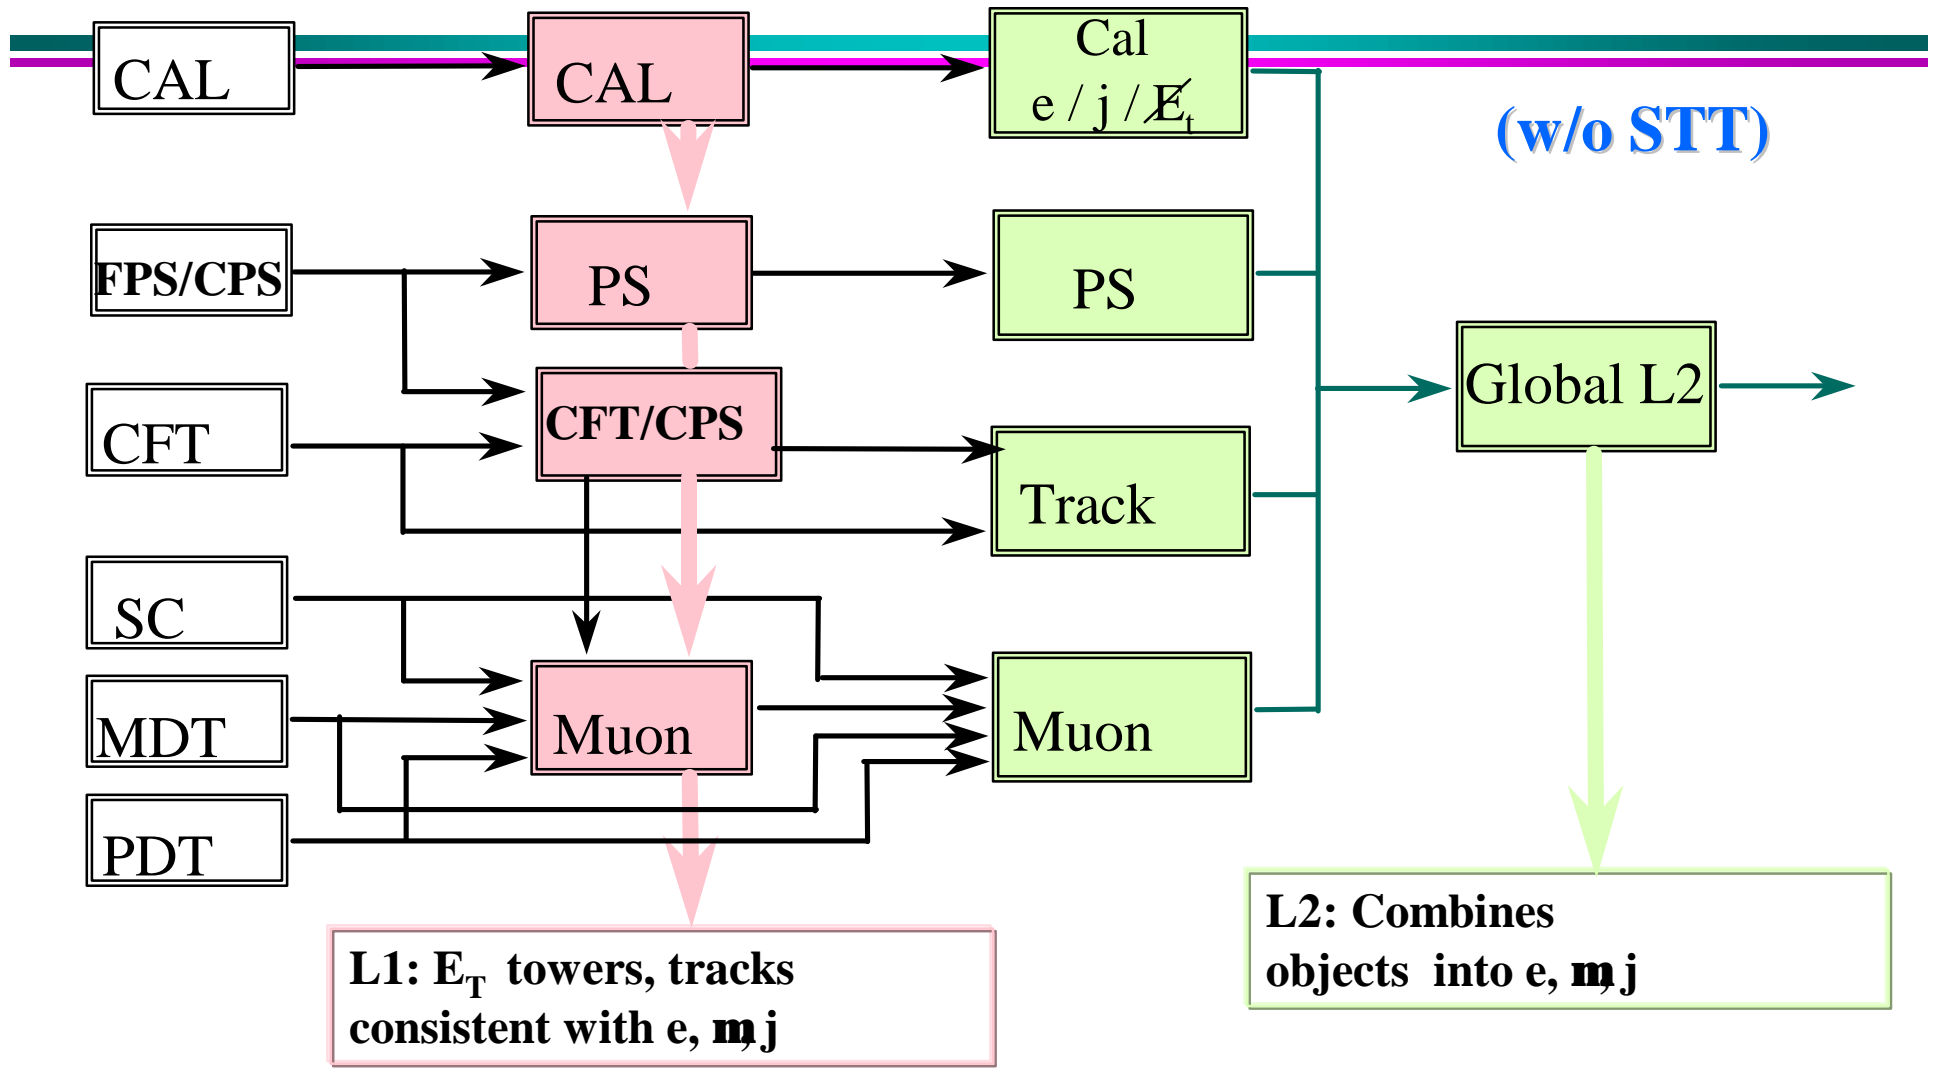

Overview of Level 2

James T. Linnemann
Michigan State University
Level 2 Review
Feb 6, 1999

Requirements

- 10 KHz input rate
 - 100 μ sec decision nominal time budget
- Reject 90% at acceptable efficiency
 - read out at 1KHz
- Deadtime < 5%
 - 16 buffers for events awaiting decision
- Flexibility in trigger configuration

L2 Trigger



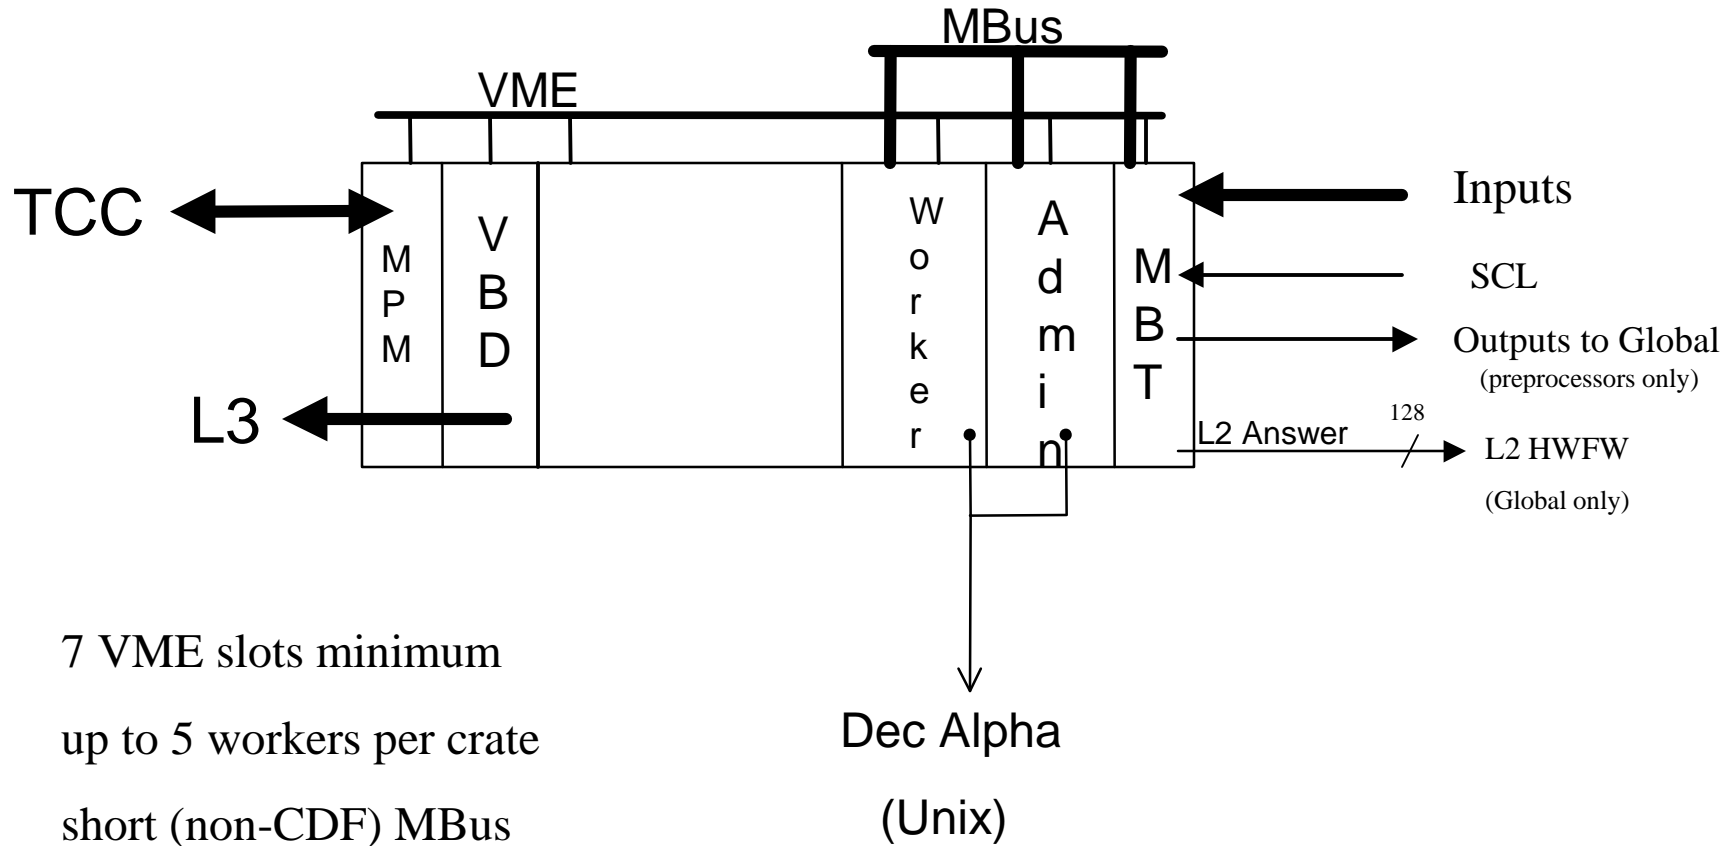
Architecture

- 2 to 3 stage stochastic pipeline 100 $\mu\text{sec}/\text{stage}$
 - Preprocessors for individual detectors
 - Global processor to combine detectors
 - 128 trigger conditions (1 to 1 with L1)
 - each programmable
 - series of conditions (e, j, μ) and cuts ($E_T > 20$)
- 16 buffers in Front end (real events)
 - 16 buffers in front of preprocessors, global
 - Busy raised by Front Ends
- Hardware frame drives readout

Queuing Simulations: Effective Time Budget

- 50-75 $\mu\text{sec}/\text{stage}$: tails of processing time
- RESQ + standalone to check simple cases
- Preprocessors: not event synchronous
 - avoid worst of n distribution
- Need the buffers in front of each element
- Avoid long tails in processing time
- Farms feasible only if surrender event order
 - our front ends required order preservation

Standard Crate



JTL, MSU 12/18/97

Standard Crate

VME Slot Assignments

- 1: Bit3 (Crate Controller) no J3 (1 slot)
- 2: VBD (2 signals from J3 to Admin)
 - through hole in blank MBus
- 3-6 J3 connector for VTM
 - up to 4 FIC's, or any non-MBus cards (SLIC/SFO)
- 7-21 J3 Magic Bus:
 - 20-21 Administrator (all Alphas 2 slots)
 - 19 Pilot MBT (preproc. : 1MBT for 2 Workers)
 - 18 down [Assistant MBT as needed]
 - [need 1 MBT per 2 Workers for output]
 - 7-8 up **up to 5 Workers** (or non-MBus cards)

Alphas

- Up to 1 GIP Alpha 21164 on VME card
 - small local disk for bootup
 - Enet to Dec Unix Alpha for user .EXE, debugging
- Most MBus I/O via MBT card
 - MBus DMA input 80-100 MB/s (Input “for free”)
 - MBus bi-directional programmed I/O 20 MB/s?
 - preprocessor output to Global
 - but interprocessor communication w/o MBT
- 2 per crate
 - Worker formatting, Output to Global
 - Administrator housekeeping, L3 R/O

Alphas, continued

- VME for L3 readout, monitoring, downloading
- 32 bits ECL output
 - scaler gates for monitoring in L1 Scalers
 - available even if alpha crashed to tell states
- J2 Inputs
 - miscellaneous communication
 - e.g. “you have a message from MBT”

MBT

Magic Bus Transceiver

- VME slave; MBus master and slave
 - Administrator controls card(s)
- 7 Cypress Hotlink inputs
 - 16 MB/s each (Gigabit Ethernet UTP)
 - broadcast to Alphas (Workers & Admin) on MBus
 - normal data Input path
- 3 Cypress Outputs
 - 2 Preprocessor outputs to L2 Global
 - 1 Echo of L1 SCL info

MBT, continued

- Serial Command Link (SCL) Receiver
 - broadcast L1 to Alphas on MBus
 - synchronization check
 - L1 Qualifiers (basic info on handling events)
 - echo'd on Cypress output for SLIC
 - Queue L2 accept/rej for Administrator MBus reads
- Parallel Output (16-128 b)
 - Global uses to send L2 decision to L2 HWWF
 - handy for monitoring/debugging

Other Cards

(Not unique to L2)

- Bit3 is commercial VME interface
 - multiport for indirect communication with TCC
 - parameter download, monitoring, error logging
- VBD is standard DØ VME Readout to L3
 - tolerable constraints on how Alphas read out
 - Forces interprocessor communication to MBus

Bit3 MPM

- Commercial; fiber optic connection
- To PCI of a PC; VME master, crate controller
- Add Multiport Memory Module
- Perform general VME I/O, generate interrupts
- Download parameters for run
- Run begin/end commands
- Collect Monitoring information
 - preferably, already placed in MPM by Administrator Alpha
 - If necessary, can collect from other modules

VBD

- Standard DØ card
- VME Master to read out to L3 (standard card)
- Not interruptible during Readout
- Probably 10-20 MB/s effective (more?)
- Must read from SAME set of VME addresses every event
 - intent is readout from Worker Alpha
 - move data, or map to actual location
 - some wordcounts may be zero
 - faster if fewer addresses

Standard Crate Uses

- Global JUST Standard Crate described so far
- Cal: more workers
- Standard Crate can also be used with non-Alpha, non-MBus pre-preprocessor
 - Cypress inputs to Worker via MBT
 - format, message data for Global
 - handle L2, L3 buffering & I/O, most of monitoring
 - *Completely standard data movement software*
 - *User code testable once data structure fixed*
 - Penalty: extra latency (lose a buffer)
 - 3-stage pipeline as in L2Mu, L2STT

L2 Inputs

- Cypress Hot Links 160 Mbit/s UTP
 - well-defined protocol
 - begin, end event special characters
 - compatible with muon (except cable: CIC)
- Standard L2 header/trailer defined
 - some header info repeated in trailer
 - allows more error detection/correction
 - Hardware Longitudinal Parity Check in trailer

L2 Header

B0	# objects (NOT IN HEADER)	[note 255 max!]
B1	Header Length in 4B words (1B)	[=3 for default]
B2	Object Length in 4B words (1B)	[ALL same size!]
B3	Header/ Trailer Format # (hi 3 bits) Object Format # (lo 5 bits)	[ONLY changes if new format] [ONLY changes if new format]
B4	Data Type # (1B)	[unique in all L2 MBT inputs]
B5	Bunch # (1B)	
B6-7	Rotation# (2B)	[B6 is LSB of rotation]
B8	Algorithm Major Version (1 B)	[e.g. 7 from Version 7.1]
B9	Algorithm Minor Version (1B) or Processor Specific Bits (1B)	[e.g. 1 from 7.1] [esp. if hardware data source]
B10	Processor Specific Bits (1B)	
B11	Status Bits	[b7 on means some error] [some standard for L2 Proc]

Standard Status Bits

b7, b0 for all; others if L2proc

7 error on event (any kind): use at own risk

6 no processing attempted (none required)

5 object list truncated (any reason)

4 Receiver error on some input physical trailer

3

2

1 more data-type info (processor-specific)

other test modes; unbiased-sample data...

0 0 for real data, 1 for MC data

L2 Trailer

B0 Bunch # (1B) = B5 of Header

B1 **Data Type** # (1B) = B4 of Header

(Swapped even/odd from Header)

B2 **Longitudinal Parity** of even Bytes

B3 **Longitudinal Parity** of odd Bytes

or--if parity too slow to calculate, Turn # (B6-7 of Header)

MBT Out, SLIC, FIC will append physical trailer with 8-bit hardware-generated longitudinal parity

Zero padding to 16 B group FOLLOWS trailer, before End of Event

L2 Physical Trailer

- FIC, SLIC, MBT Out: add a physical 2B trailer
 - after logical trailer, before End Event
 - This BREAKS 16B boundary, but handled by MBT
 - B0 8 bit longitudinal parity of received data
 - B1 Status Bits [b7 on if any receive error]
 - not included longitudinal parity!
 - b0, b1 are type ID: 0 = FIC, 1 = SLIC, 2 = MBT
- MBT inputs place this in B0, B1 of 16B physical trailer
 - adds B14, its own longitudinal parity of everything received
 - B15 its own Error Bits [b7 on if any receive error]
 - reserves 4B for incoming, may give error locations in B4-13
 - MBT Outs produce 2B physical trailer like FIC

SLIC:

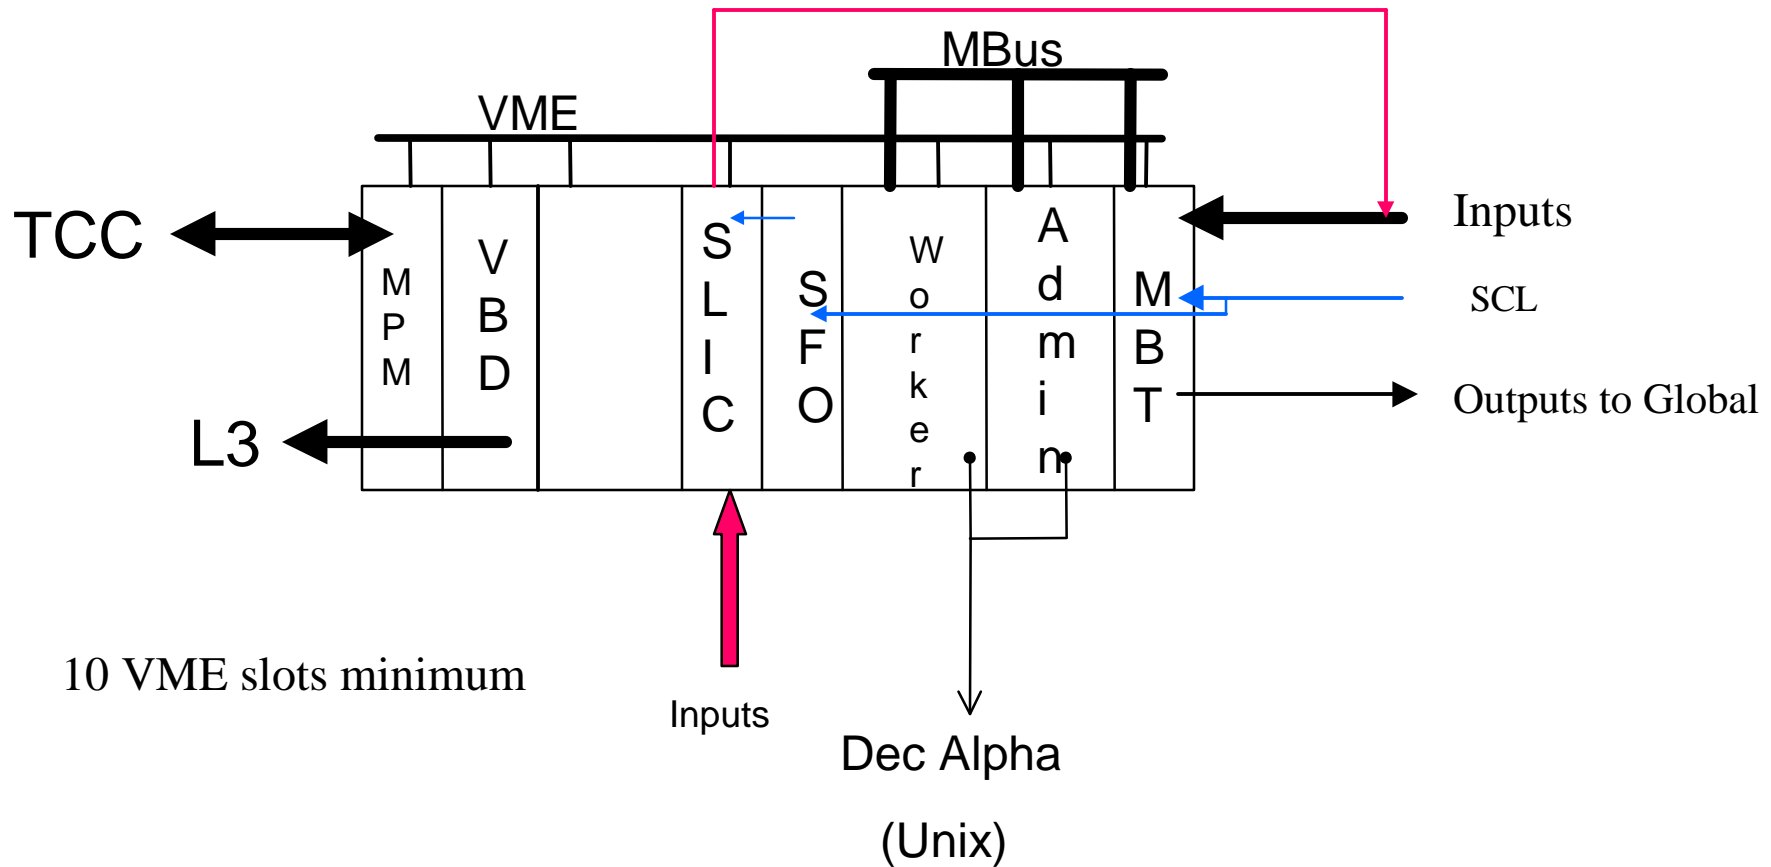
Serial Link Input Card

- 16 Cypress serial inputs
 - 1-slot VME slave card
 - 4 TI DSP's, up to 2 GIPs each
- more inputs, CPU / slot than Alpha
- output via Hotlink to MBT (avoids VBD R/O)
- Readout via Worker Alpha via MBT
 - Acts as pre-preprocessor
- test registers on all inputs (eg. SCL)
- NO MBus! (big simplification)

SFO: SCL Fanout (Really: Cypress Fanout)

- Receives L1 SCL information
 - from MBT as Cypress Hotlink
- Fans out as Cypress output to 12 SLIC cards
 - event synchronization
 - L1 Qualifiers
- purely analog fanout
- can be used to fan out *any* Cypress signal
 - L1HWWF messages to L2
 - potential use in L2STT?

Standard Crate with SLIC



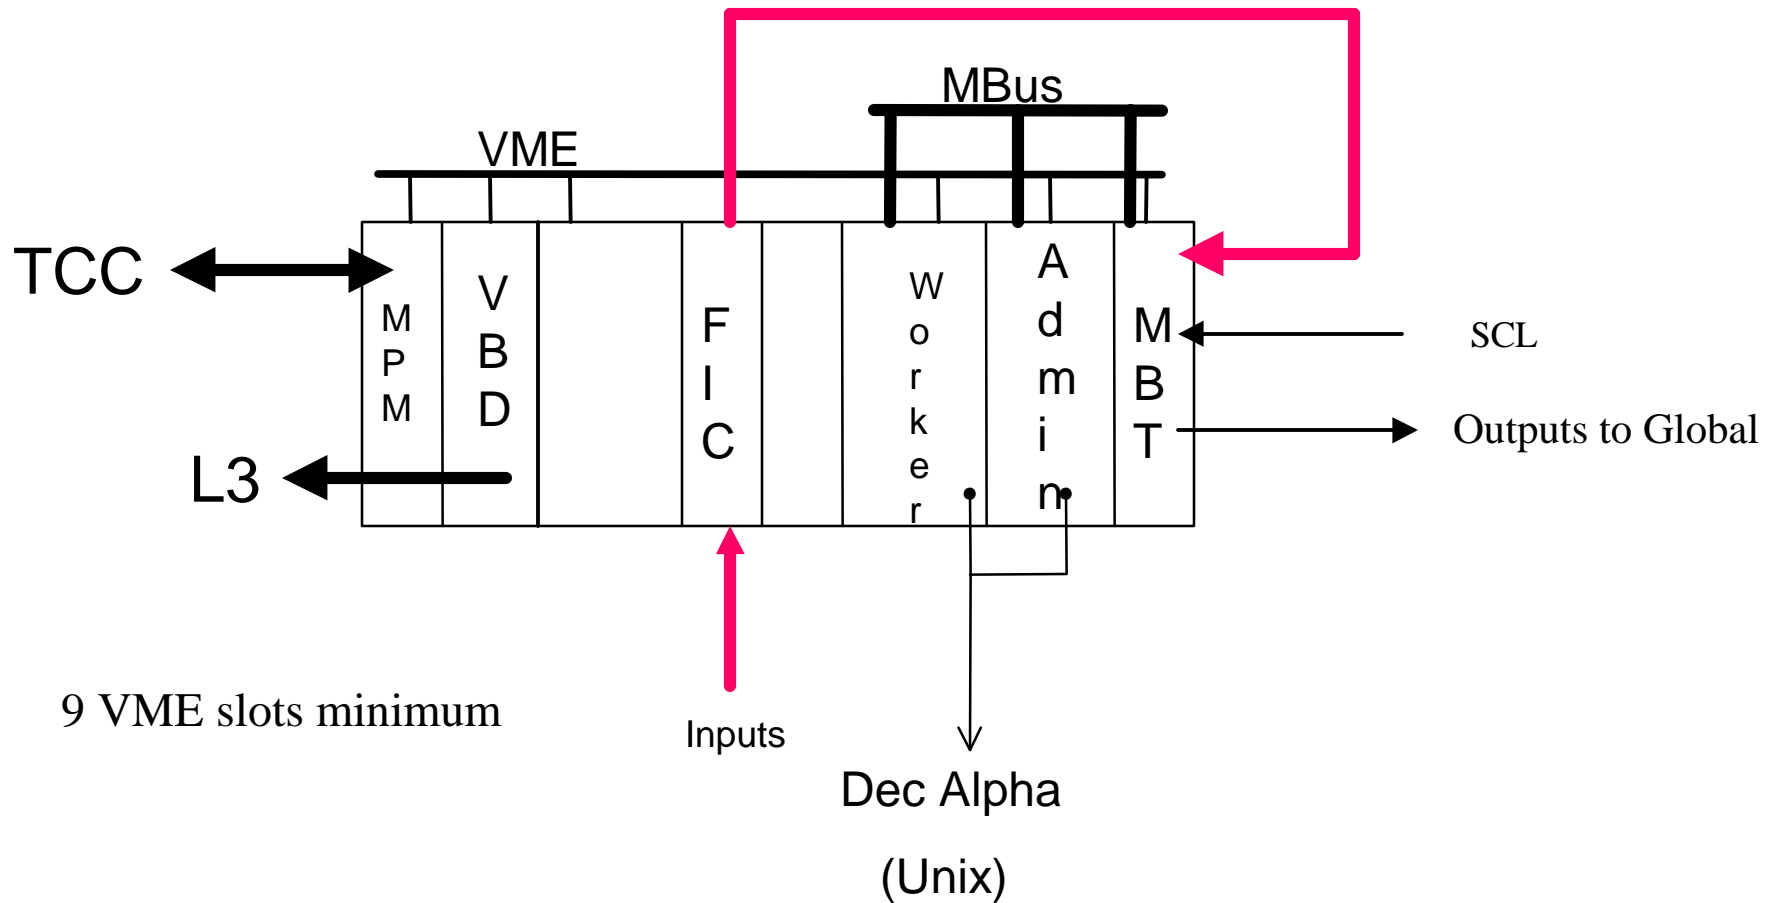
Fiber Input Converter (FIC)

- Convert Fiber Input to Cu Cypress 160 Mb/s
 - G-link input 16b data in 20b data frame (24b total)
 - input thru J3 by **standard VTM** (hard G-link engineering done)
 - implement g-link input via VRB card
 - allows passive split for fanout to L3 or STT
 - adds physical trailer with longitudinal parity
- Front end to *either* SLIC or MBT
 - avoids variants of complex card
 - used in L2PS, L2CaI, L2CTT
- 4 independent channels per card
- VME control, monitoring

FIC: L2CFT from L1 CFT trigger (& L1 Cal)

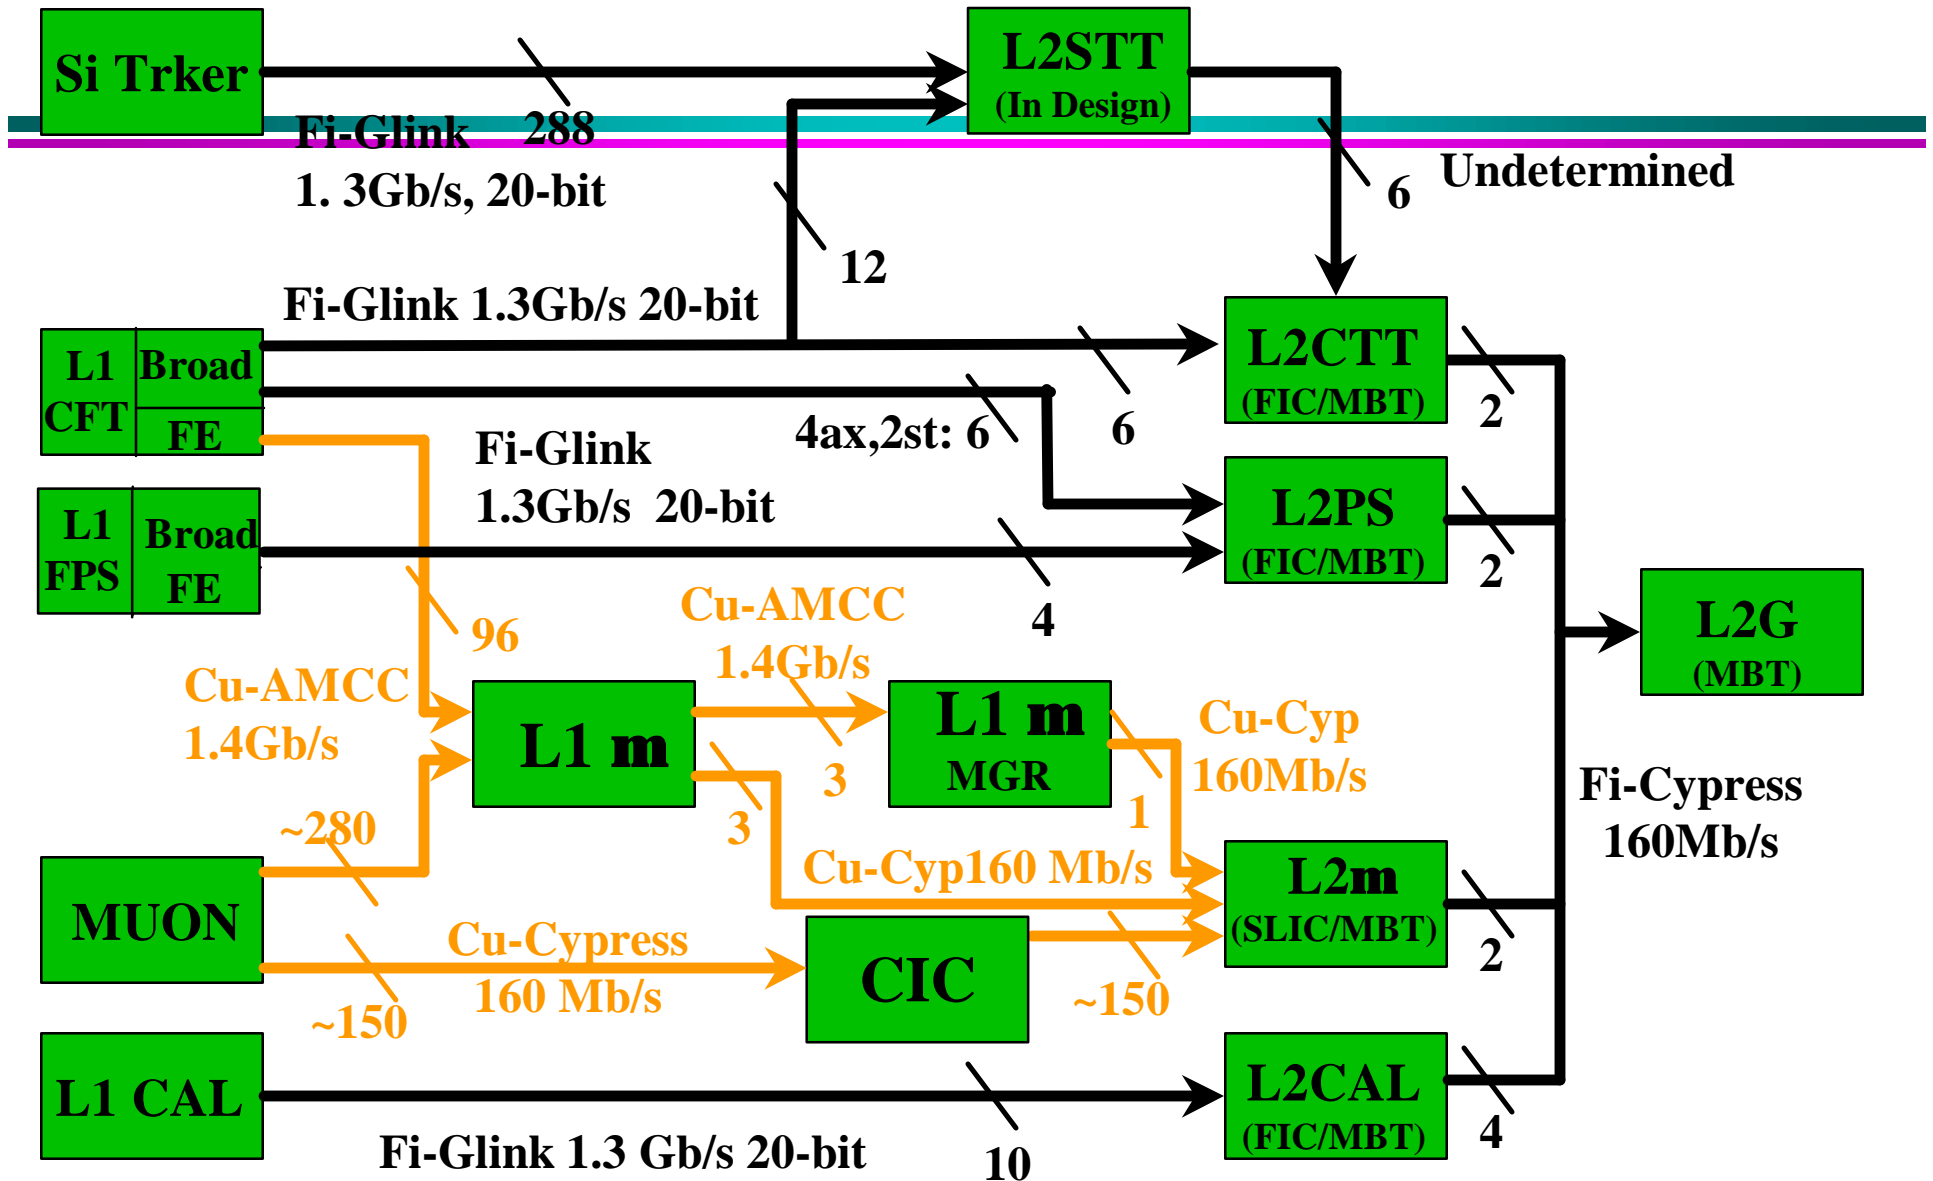
- g-link 1.3Gb/s = 106MB/s
 - 16b=2B data in 24b frame, frames at 53MHz
 - L1CFT: 100B (50 tracks)/fiber to STT in 1 μ s
 - standard L2 header
 - trailer includes 2B longitudinal parity
 - pad w/ trailing zeros
 - L1Cal:
 - similar format, fixed-length data
 - optical split from data for L3 readout

Standard Crate with FIC to MBT



JTL, MSU 12/18/97

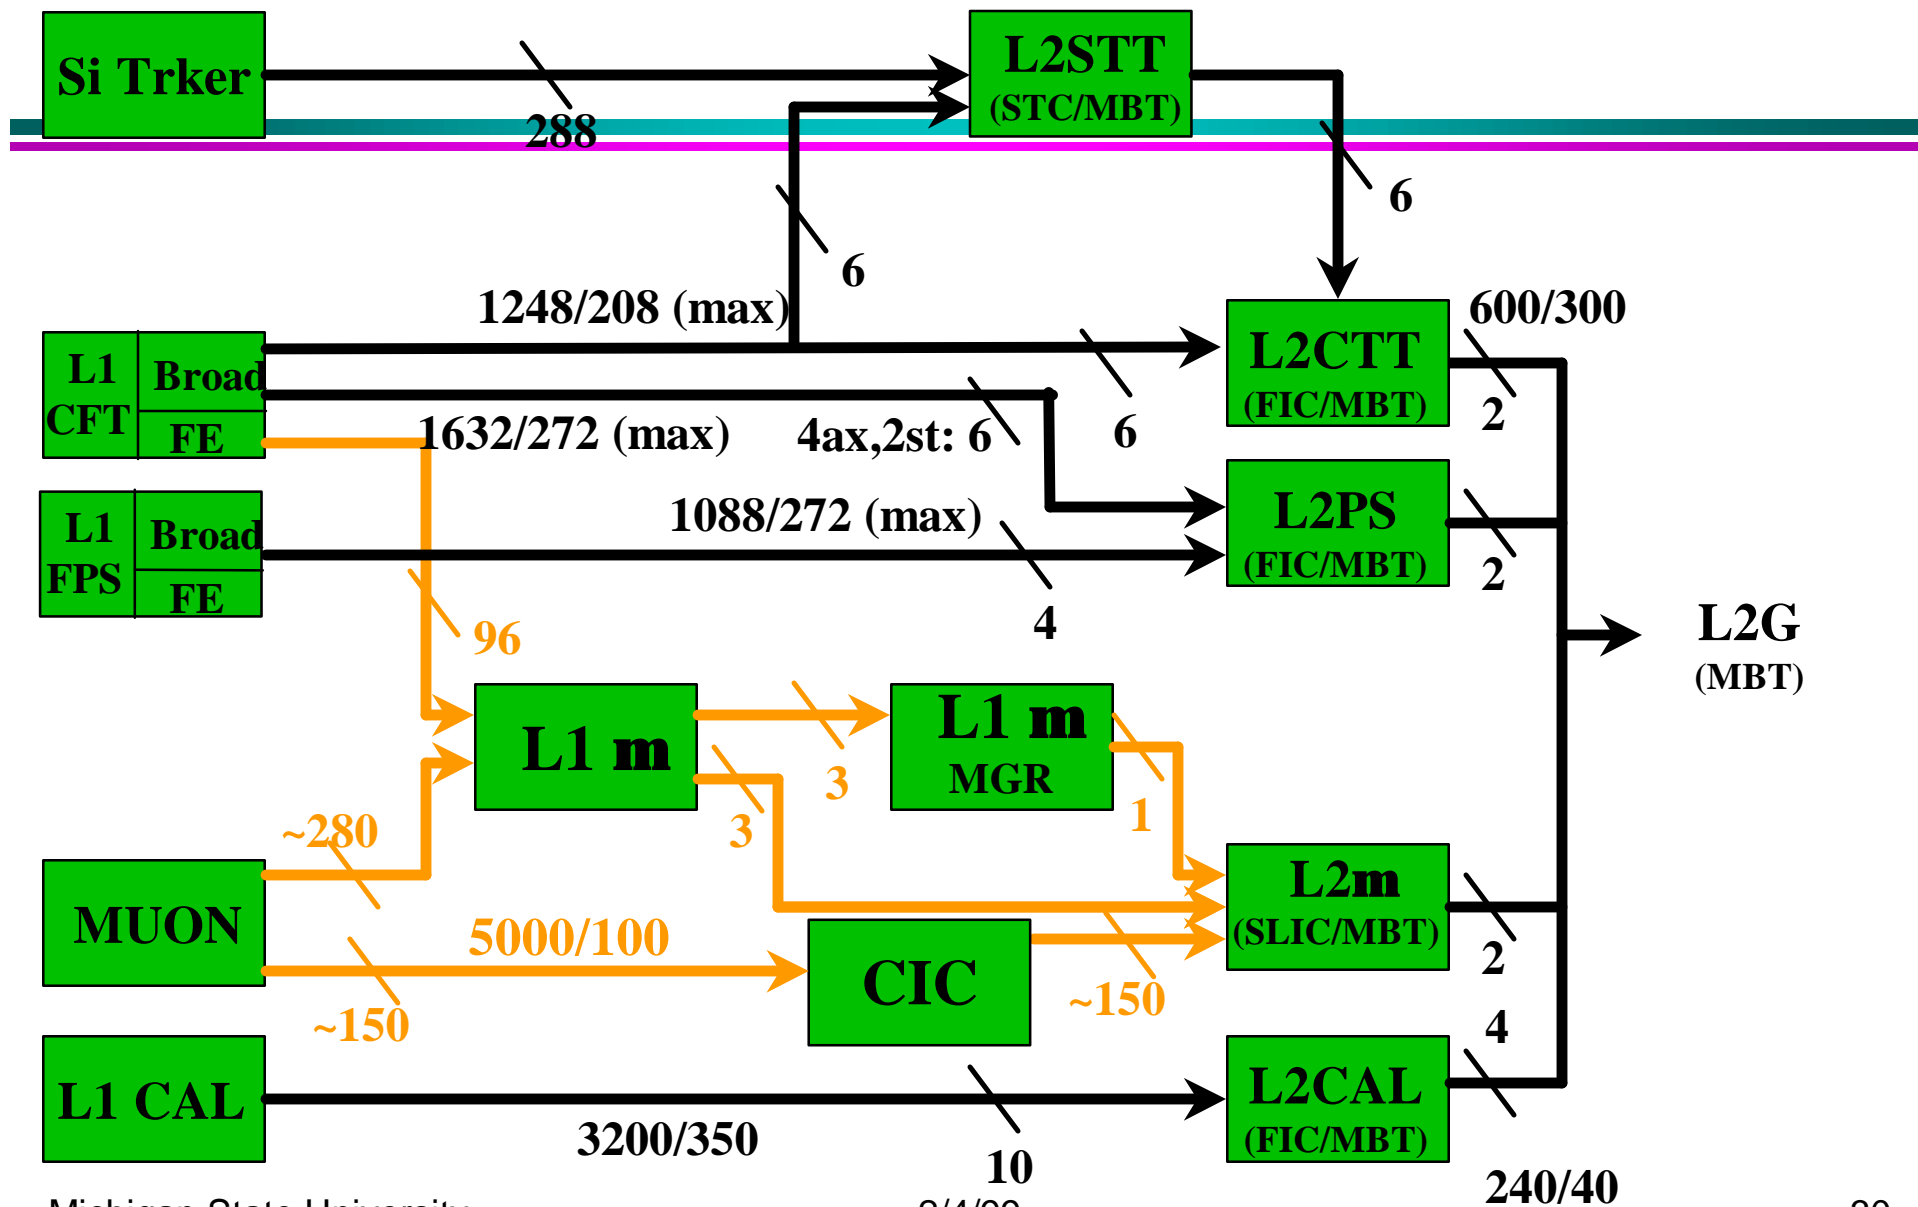
Trigger Connections



L2 Bandwidth and budget/event @ 10KHz

<u>Link</u>	<u>MB/s</u>	<u>KB/event</u>	
G-Link	106	10.6	
Cypress Hotlink	16	1.6	
Mbus DMA	~80	8	320MB/s nominal
Mbus Prog	~20	2	.5-1 μ sec/16B
VME Block	>10	1	
VME Prog	~1	.1	.5-1 μ sec/4B

Loading of Paths



% Capacity used

Crate	Worst In /1.6KB	Mbus In/ 8KB	Mbus Out/2KB	Worst Out/1.6 KB	L3 /1KB (10%)	
Cal	22 % fixed	40% fixed	12%	5%	2%	
CTT	13% max	16% max	30%	19%	6%	18 x 8B /list
PS	12% max	26% max	6%	2%	1%	Avg < .3*max
MU to SLIC	6%	X	X	1%	X	
Slics to Mu Alpha	1%	10%	2%	3%	1%	
Global	19%	13%	X	X	3%	