Streaming Electronics
December 3, 2018

Benjamin Raydo
Electronics Group (Physics Division)

JLAB DAQ

Triggered DAQ Systems

- 1kHz to 100kHz triggered experiments
- 100k DAQ channels (Hall B & D)
- Trigger latency
  - ~400 ns for Hall A,C
  - ~8us for Hall B
  - ~3us for Hall D
- Data rates
  - <100MB/s for Hall A,C
  - ~500MB/s for Hall B
  - ~1GB/s for Hall D

FADC250 DAQ Crate

- Sample rate: 250MHz
- Voltage resolution: 122μV
- 256 FADC Channels, 12bits

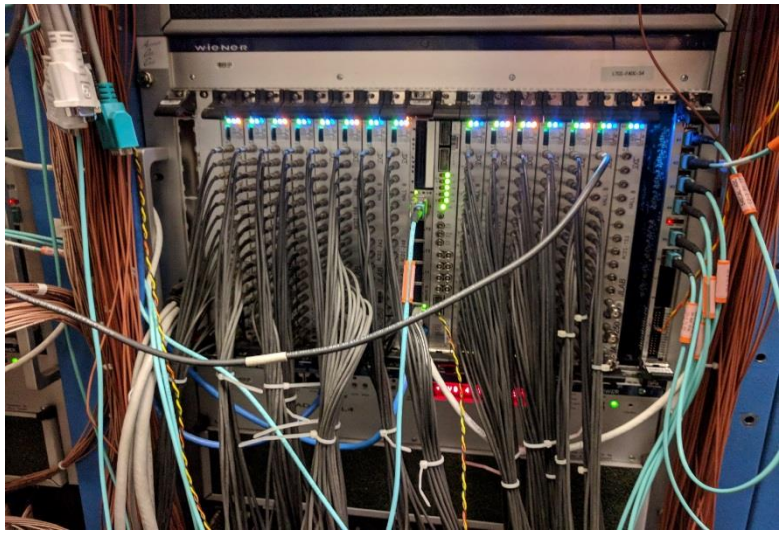

Example waveform readout from FADC250 (from PMT lead-glass based calorimeter)

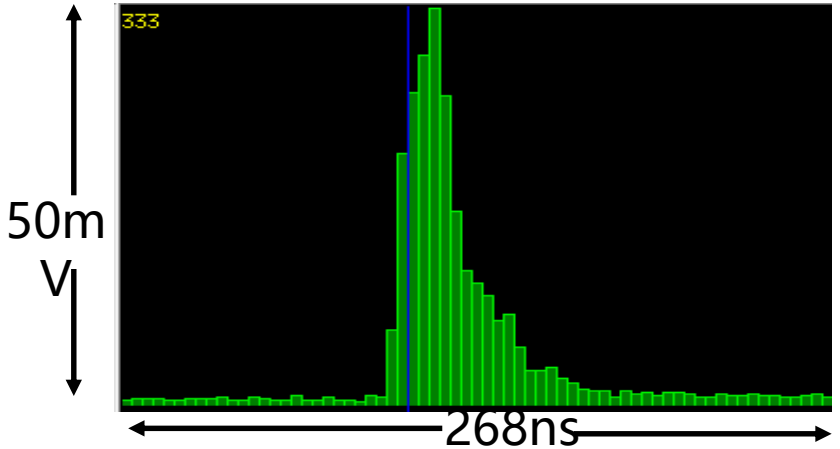

JLAB DAQ Example Hardware

Front-end Digitizer Modules

FADC250
16 channel, 12bit Flash ADC
250MHz Sample rate
200MB/s VME Readout
20Gb/s Trigger Serial Link

RICHFPGA
192 Channel MAROC3A
1ns time-to-digital Converter
2.5Gb/s optical Readout

L1 Trigger Processing Modules

VTP
16 Payload Ports 34Gb/s = 544Gb/s
4 Fiber Ports 34Gb/s each, 1 40Gb/s
4GB DDR3 SDRAM 200Gbps
1GHz Dual Core ARM Processor

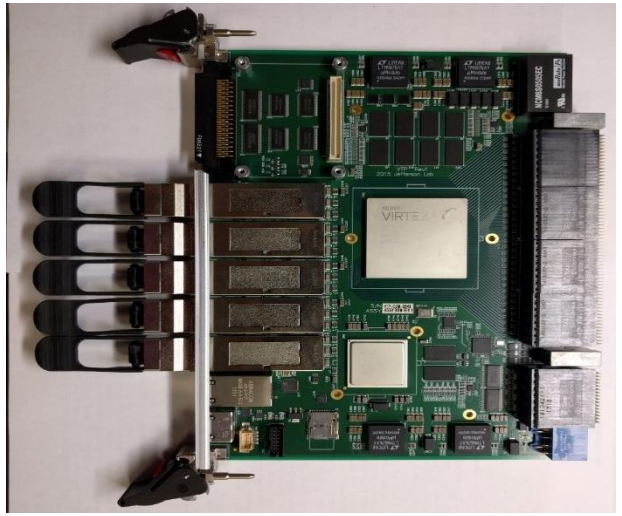

Clock & Trigger Distribution & Synchronization Modules

TS
Global Oscillator (250MHz)
64 Trigger Inputs

TI/TD
8 Fiber Port Optical Fanout
Clock/Trigger/Sync
Low Jitter Clock (<3ps)

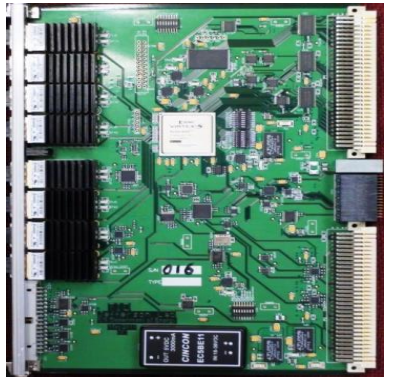

JLAB Example Trigger System (CLAS12)

Stage 1: Front-end Processing (Clusters, Roads, Hits)

Stage 2: Sector/Central Coincidence/Geometry Matching

Stage 3: Global Trigger bits

DCRB

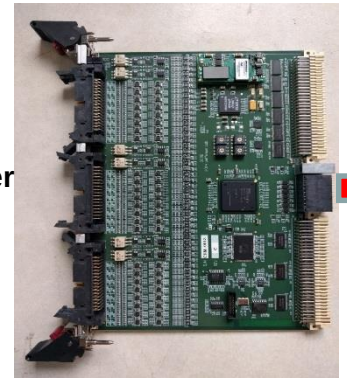

Drift Chamber

VTP

FADC250

FTCAL
FTHODO
ECAL
PCAL
FTOF
CTOF
CND
HTCC

VXS Backplane 8Gbps

DC Roads

Fiber (up to 300m) 20-34Gbps

ECAL Clusters
PCAL Clusters
PCAL Pixels
PCAL U Mask
HTCC Mask
CTOF Mask
CND Mask
FTOF Mask
FTCAL Clusters

SSP

Sector Trigger Bits
Central Trigger Bits

VXS Backplane 8-16Gbps

VTP

32 Trigger Decisions to TS

Front-end ASICs – Challenges with streaming model

Many front-end ASICs require a trigger signal

- **APV25, DREAM, PETIROC, …**
- **Not possible to use in a streaming model system – they have a large dead-time after triggering to perform digitization**
- **These are low-power digitizers (typ. few mW per channel)**

Commercial fast waveform sampling ADCs work streaming model

- **Large data, cost, power, and real-estate**
- **Only practical for off-detector electronics where channel counts are <~10k (imagine 100k channels of on-detector 250Msps ADCs: 100kW, 0.5Pbps)**

Front-end ASICs – Better Solutions

FSSR

- Used in CLAS12 – zero suppress and stream hit channel, charge, and time directly from low-power ASIC mounted on detector

SAMPA

- ASIC for ALICE TPC – shaping, waveform sampling, DSP in low power front-end chip.
- Ed Jastrzembski talk will discuss its planned use at Jlab

Nalu Scientific

- Low power, extremely fast capacitor array sampling ASICs
- Earlier releases not compatible with streaming model, but new developments underway…
- Likely some self triggering and buffering tricks could make this a practical low-power solution for on-detector streaming solutions

Pacific Microchip

- 32ch 12b 500Msps ADC ASIC at <4mW/ch
- This is a massive drop in power consumption compared to existing commercial options, but this low power is without the interface (which would likely by > ~50mW per channel based on using Serdes)
- Adding pulse feature extraction logic into this ASIC could make a very generic readout solution

Front-end Pulse Processing

Ideally front-end reports a time and charge per hit, but this can be challenging in the cases of:

- **Large crosstalk/coherent noise => may require information from many channels to correct**
- **Large background or shaping times => pileup**
- **Pedestal fluctuations**

Some of these issues might require raw samples to be processed in a non-local way which can require large bandwidth from front-end

On the other hand, detectors without these concerns can benefit from front-end processing to massively reduce the front-end bandwidth

Planned Tests in the near term

256 Channel 250MHz 12bit FADC streaming crate

Use only FADC250 trigger path: zero suppress and stream to VTP then to server over 10GbE and/or 40GbE using UDP or TCP

i. **This allows software testing to begin with real data source and where performance tuning/optimizations can start**
ii. **Implement real-time pulse time fitting for improved timing resolution (sub ns)**
iii. **Implement special case pulse reporting**
    i. **Report raw samples for pulse pipe-up for computer based time/energy extraction**

Will plan to use a pseudo-random pulse waveform generator to emulate various channel occupancies.

Should serve as a good test bed for software development and tests…
Hardware is nearly all in place in the INDRA lab area!

Test Setup

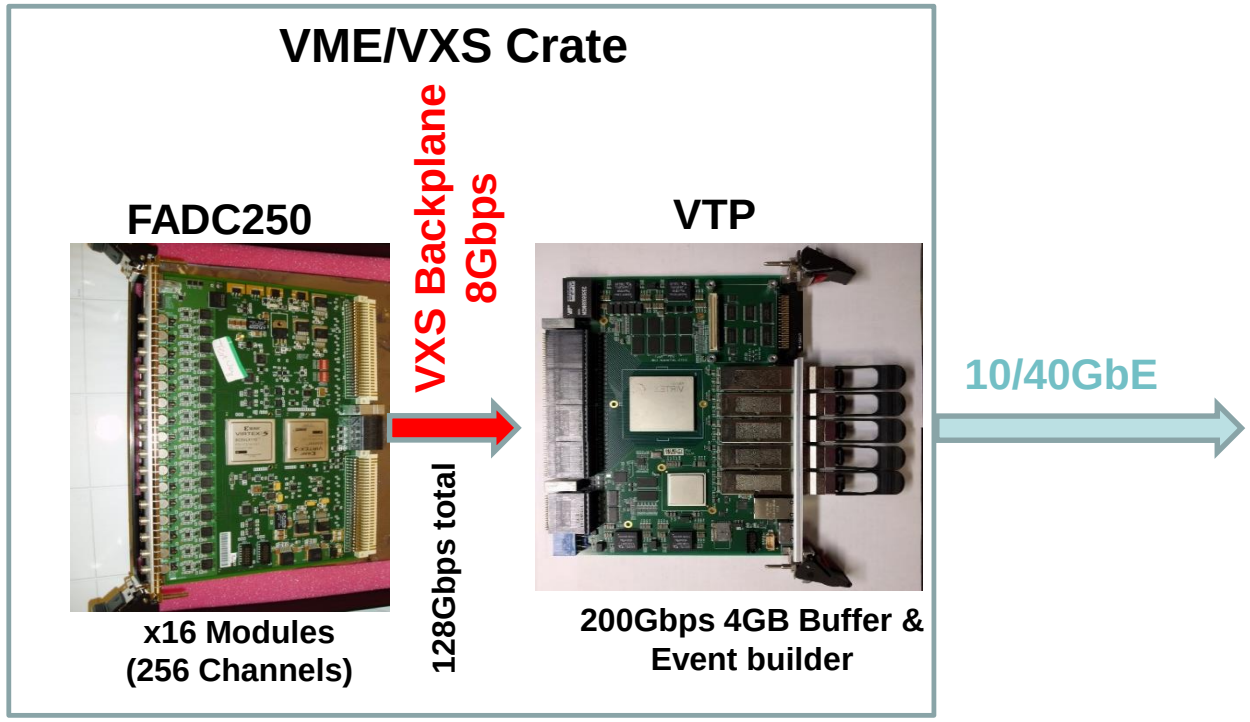

Would prefer this for long-term which would allow simplified & customizable front-ends. Single server serves specific detector channels and is responsible for interleaving events to next layer.

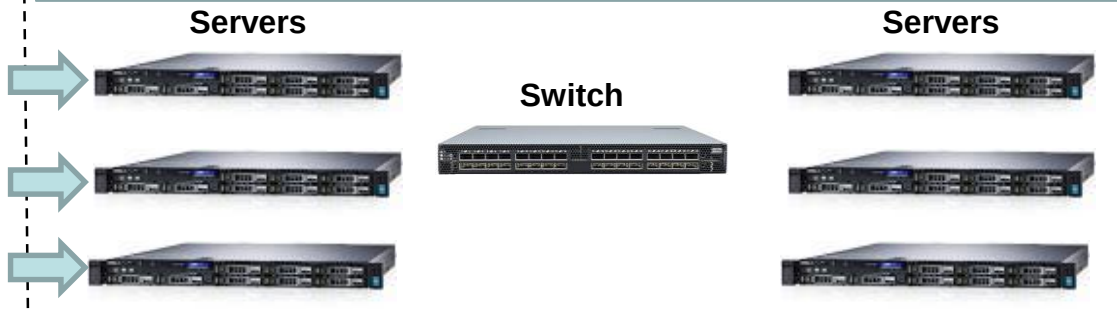

or

Could do this given VTP buffering and TCP capability – VTP interleaves events across servers, but unlikely/unwanted front-end hardware complexity for future developments.

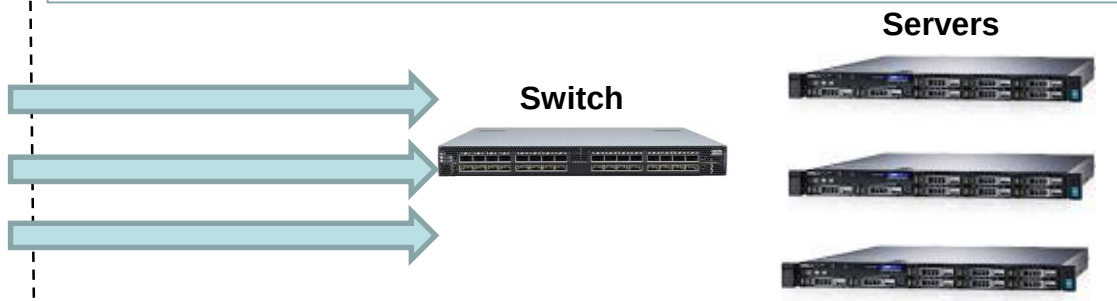

Behavior

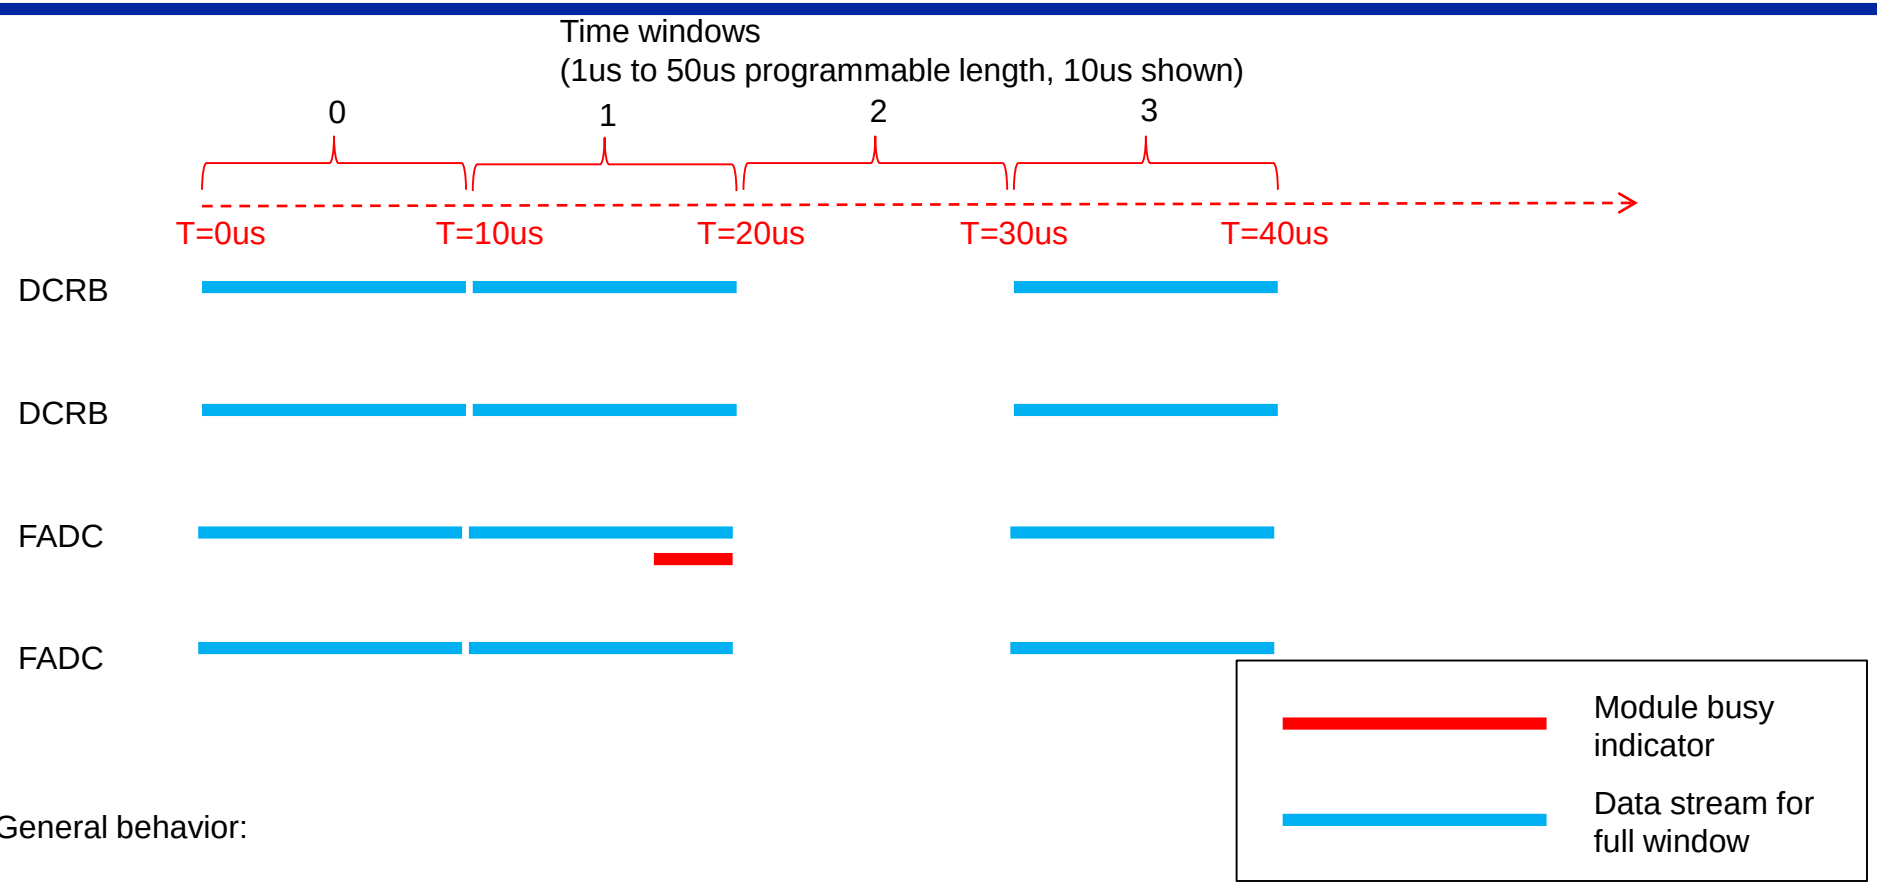

General behavior:

1) Modules stream full time windows
2) Module asserts busy if buffers too full that it can another full time window – due to data streaming back pressure
3) If any module becomes busy then all modules don't send data for the next time window
4) So this is an "all or nothing" style of streaming that supports deadtime.
5) If the network links and downstream processing can handle the data, then no deadtime
6) For zero-suppressed data this means occupancy will dictact the needed CPU and network capacities, not the front-end sample rates

Possible Use Case?

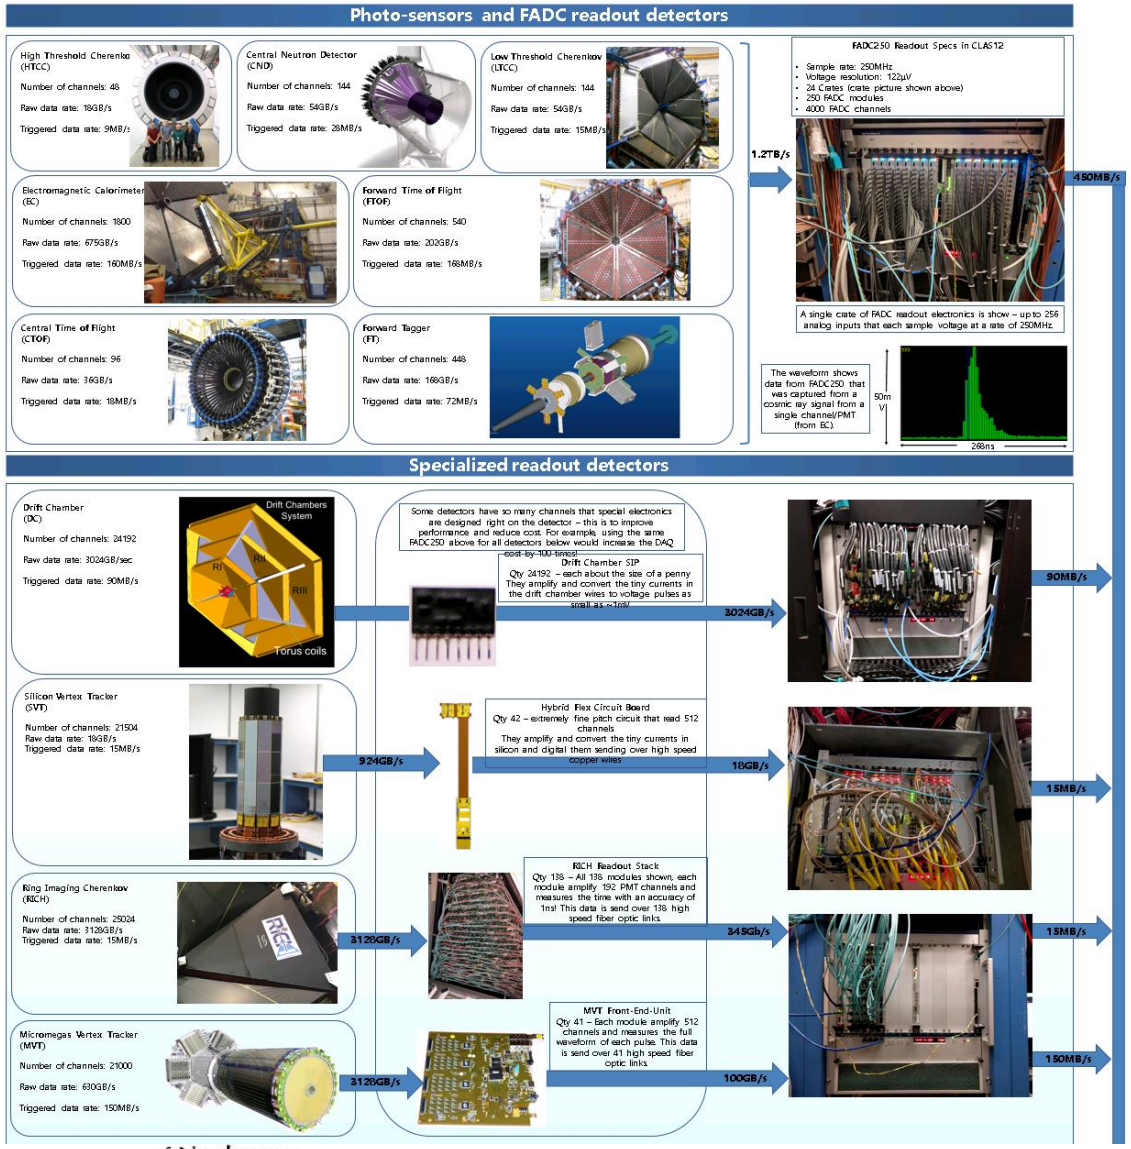

This is the full CLAS12 DAQ System

- **Many crates already equipped with VTP modules and could support the streaming readout concept presented earlier**
- **Rough estimates would put the streaming data rate in the ~50GByte/s ballpark**
- **Only 1 detector would not be compatible – the Micromegas tracker (which uses the DREAM ASIC)**

Nobody is pushing for this, but maybe someone would be interested

Front-end Protocol/Interfaces

GBT & SCA chips allow FPGA-less front-ends. Very nice for rad-hard and low power solutions

- **Hope to see continued developments from this group (28Gbps???)**

Certainly at this time it is too early to settle on any front-end interface protocol and in general my feeling is that streaming directly directly into servers (and not on the network) will keep the front-end interface open allowing a variety of solutions which seems the most flexible plan

Timing & Synchronization

Prefer (at the moment) to keep this separate from the front-end streaming data interface:

- **Would place restrictions on streaming interface protocol/hardware options early on**
- **Using recovered clocks from serialized data sources isn't the best quality if we want to minimize jitter and baseline wander – depends on required stability by a particular detector**

Conclusion

- Using existing hardware we have at JLAB we can build a functional generic streaming DAQ system based on the JLAB FADC250 readout module (but compatible with other VXS based JLAB hardware)

- If made working, it could serve as an upgrade for existing DAQ systems already at JLAB and will gain valuable experience when the time comes to build future hardware.

- From the hardware point-of-view, low-power ASIC solutions strike me as one of the bigger hurdles for streaming on high channel count detectors

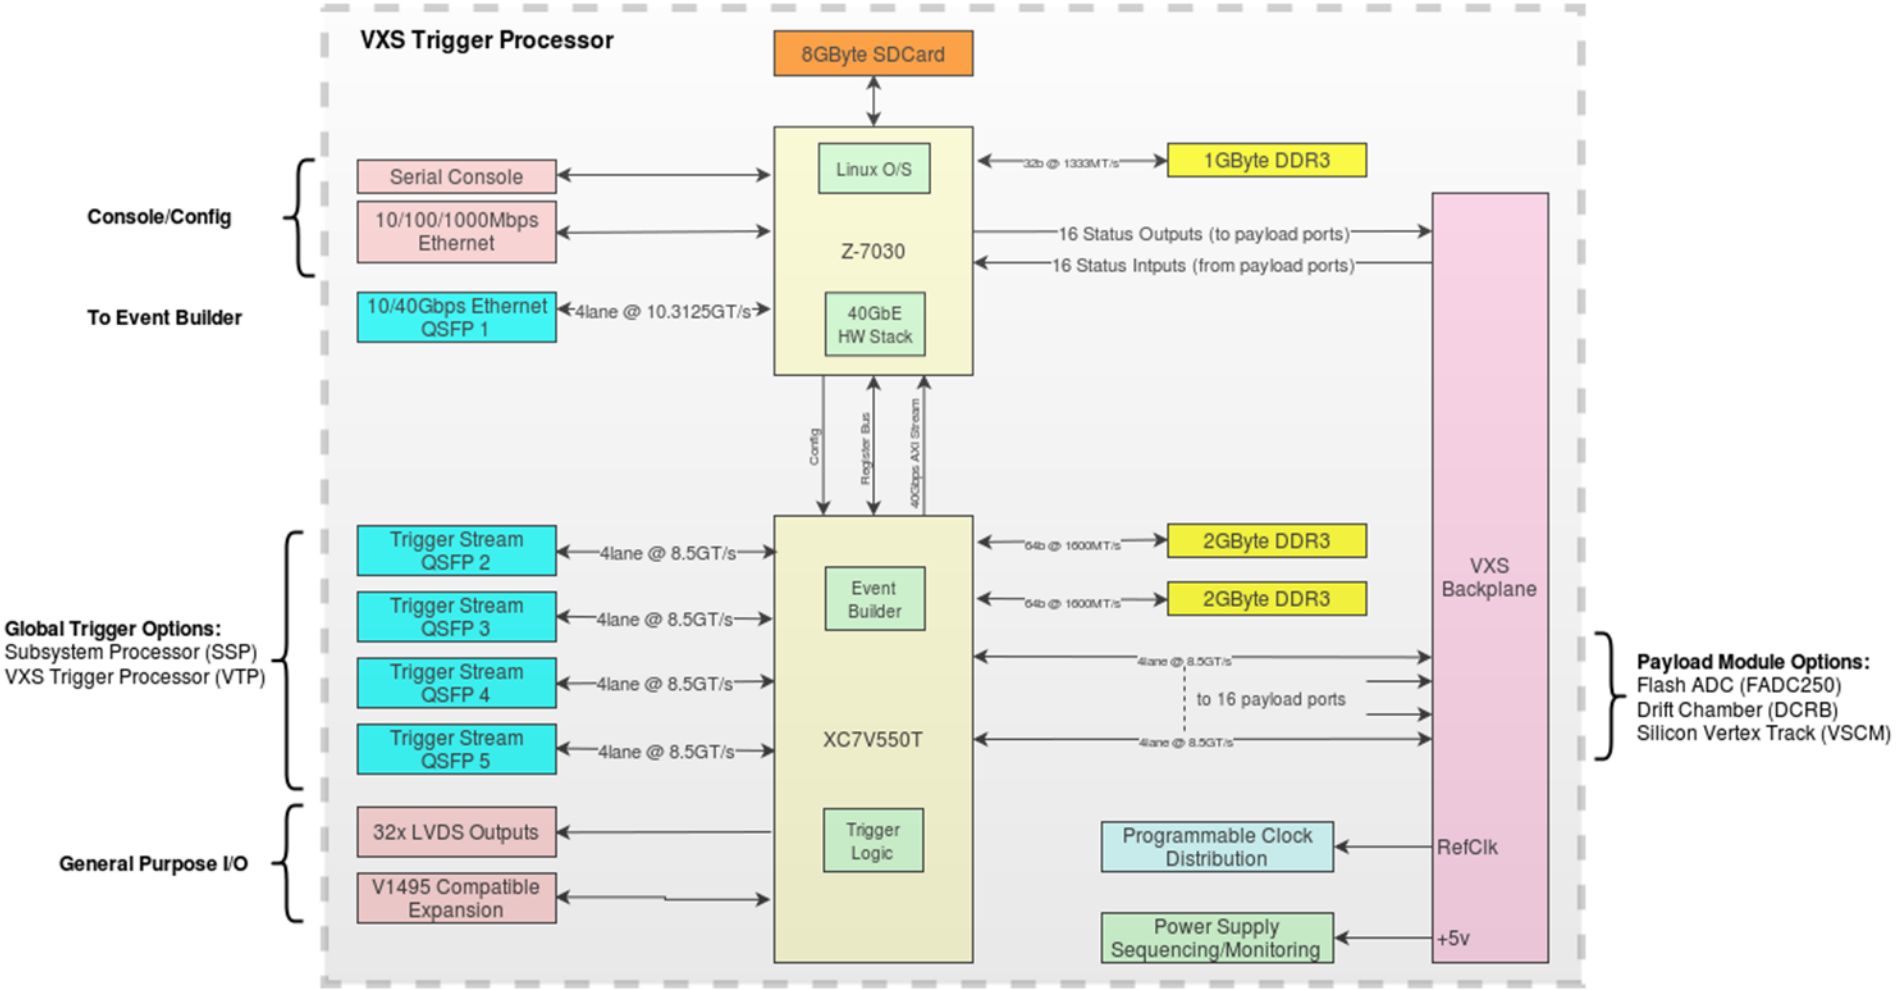
VXS Trigger Processor
8GByte SDCard
Linux O/S
32b @ 1333MT/s
1GByte DDR3
Serial Console
10/100/1000Mbps Ethernet
Console/Config
Z-7030
16 Status Outputs (to payload ports)
16 Status Intputs (from payload ports)
10/40Gbps Ethernet QSFP 1
4lane @ 10.3125GT/s
To Event Builder
40GbE HW Stack
Config
Register Bus
40Gbps AXI Stream
VXS Backplane
Trigger Stream QSFP 2
4lane @ 8.5GT/s
Trigger Stream QSFP 3
4lane @ 8.5GT/s
Trigger Stream QSFP 4
4lane @ 8.5GT/s
Trigger Stream QSFP 5
4lane @ 8.5GT/s
Event Builder
64b @ 1600MT/s
2GByte DDR3
64b @ 1600MT/s
2GByte DDR3
4lane @ 8.5GT/s
to 16 payload ports
4lane @ 8.5GT/s
Global Trigger Options:
Subsystem Processor (SSP)
VXS Trigger Processor (VTP)
XC7V550T
Payload Module Options:
Flash ADC (FADC250)
Drift Chamber (DCRB)
Silicon Vertex Track (VSCM)
32x LVDS Outputs
Trigger Logic
General Purpose I/O
V1495 Compatible Expansion
Programmable Clock Distribution
RefClk
Power Supply Sequencing/Monitoring
+5v

TCP/IP Accelerated Stacks for FPGA

	TCP IPv4 only: 0 UDP rx, 0 UDP tx, 1 TCP client, ARP, Ping, routing table, IPv4 only, 32KB TCP buffers
Flip Flops	2224
LUTs	2381
36Kb block RAM	9.5
DSP48	0

IPv4 TCP single client, uni-directional st
MTU = 8252 Bytes, equal length maxim
frames, buffer size = 32K Bytes:
9.88 Gbits/s

		Z-7030
		XC7Z030
Programmable Logic (PL)	7 Series PL Equivalent	Kintex®-7
	Logic Cells	125K
	Look-Up Tables (LUTs)	78,600
	Flip-Flops	157,200
	Total Block RAM	9.3Mb
	(# 36Kb Blocks)	(265)
	DSP Slices	400
	PCI Express®	Gen2 x4

Block Diagram

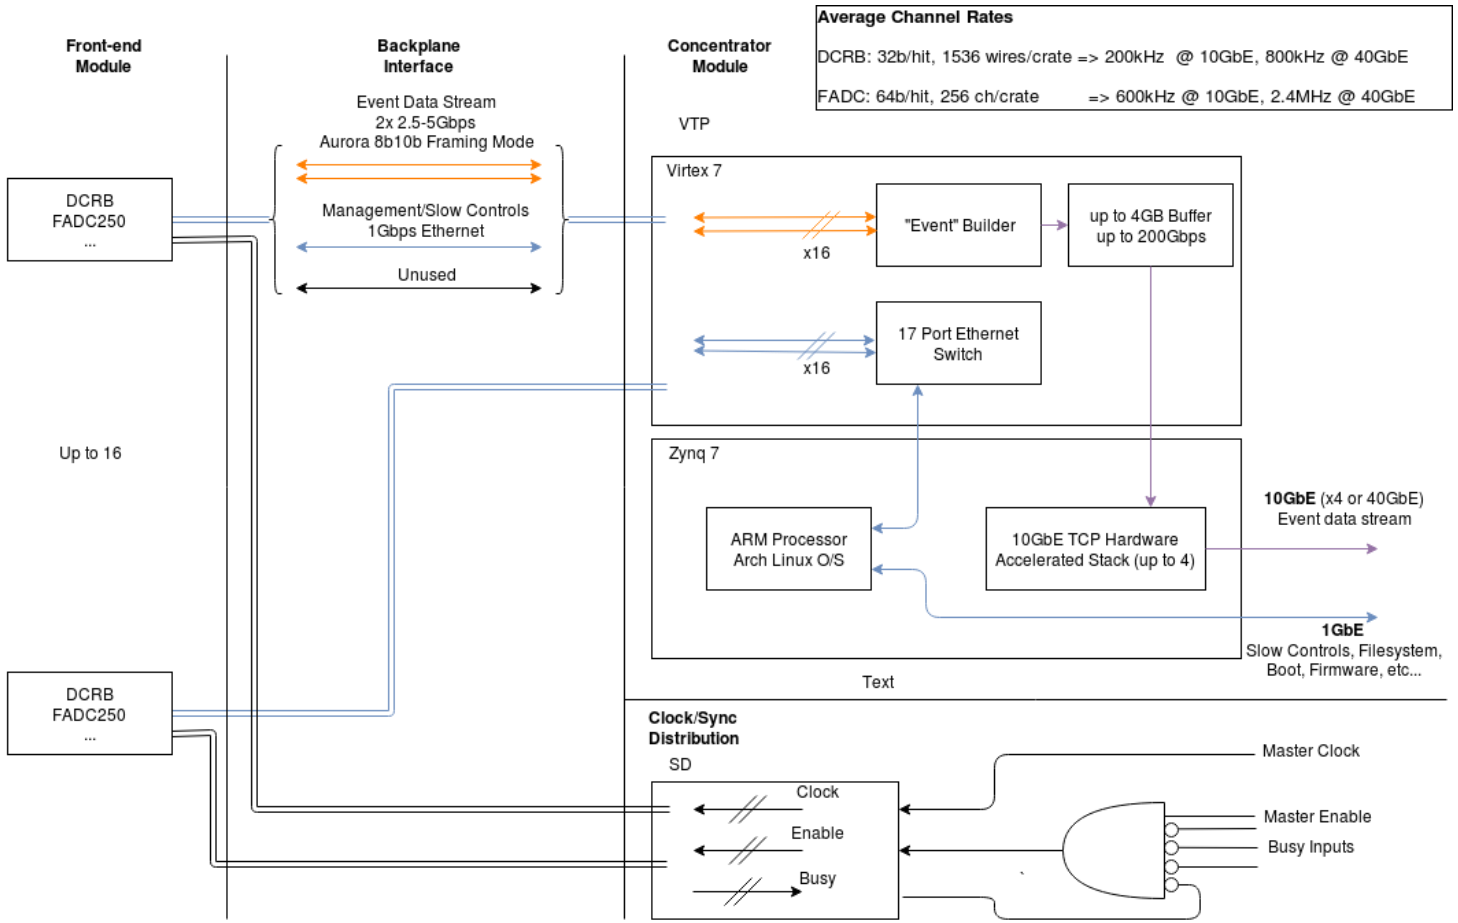

Notes:

1) Front-end modules stream complete time windows when "Enable=1" at the start of a new time window.
2) Time windows programmable size somewhere between 1us and 50us
3) Front-end modules buffer at least 2 full windows.
4) Busy assertion must happen at least 1us before next window start time if any module does not have room to buffer 1 more full window.
5) Front-end module can stream all raw data, zero suppressed, and processed data.
6) At some point the system will have deadtime if the average data rate exceeds the available bandwidth. Deadtime will inhibit streaming windows synchronously across all front-end modules.
7) Use of serial links allow VXS backplanes to be replaced with uTCA backplanes, stand alone boxes with optical links, or other serial bus structures without the need to redesign firmware.
8) Ethernet slow controls and event data streams allow scaling using commercial off-the-shelf hardware.
9) Dedicated master clock/synchronization signals is very stable (drift between modules will be minimized to ~10ps stability for stable temperature).
10) Embedded master clock/synchronization signals in ethernet links is possible, but there will be low frequency drifts on the order of 100ps or more between modules.