CMS Update Dec 09

g.j.kunde@lanl.gov

Thanks to Andi, Pat and Matt

OUTLINE

- LHC Status
- Bolek's Email
- Pixel talk from CERN CMS HI Readiness Workshop
  - Simulation Results
- Recent Testing
  - Data
  - Interpretation
- New Simulations
- Proposal for Testing Hardware
  - FPGA driven testing of the FED
- Outlook

CMS First Collisions 11/23/09

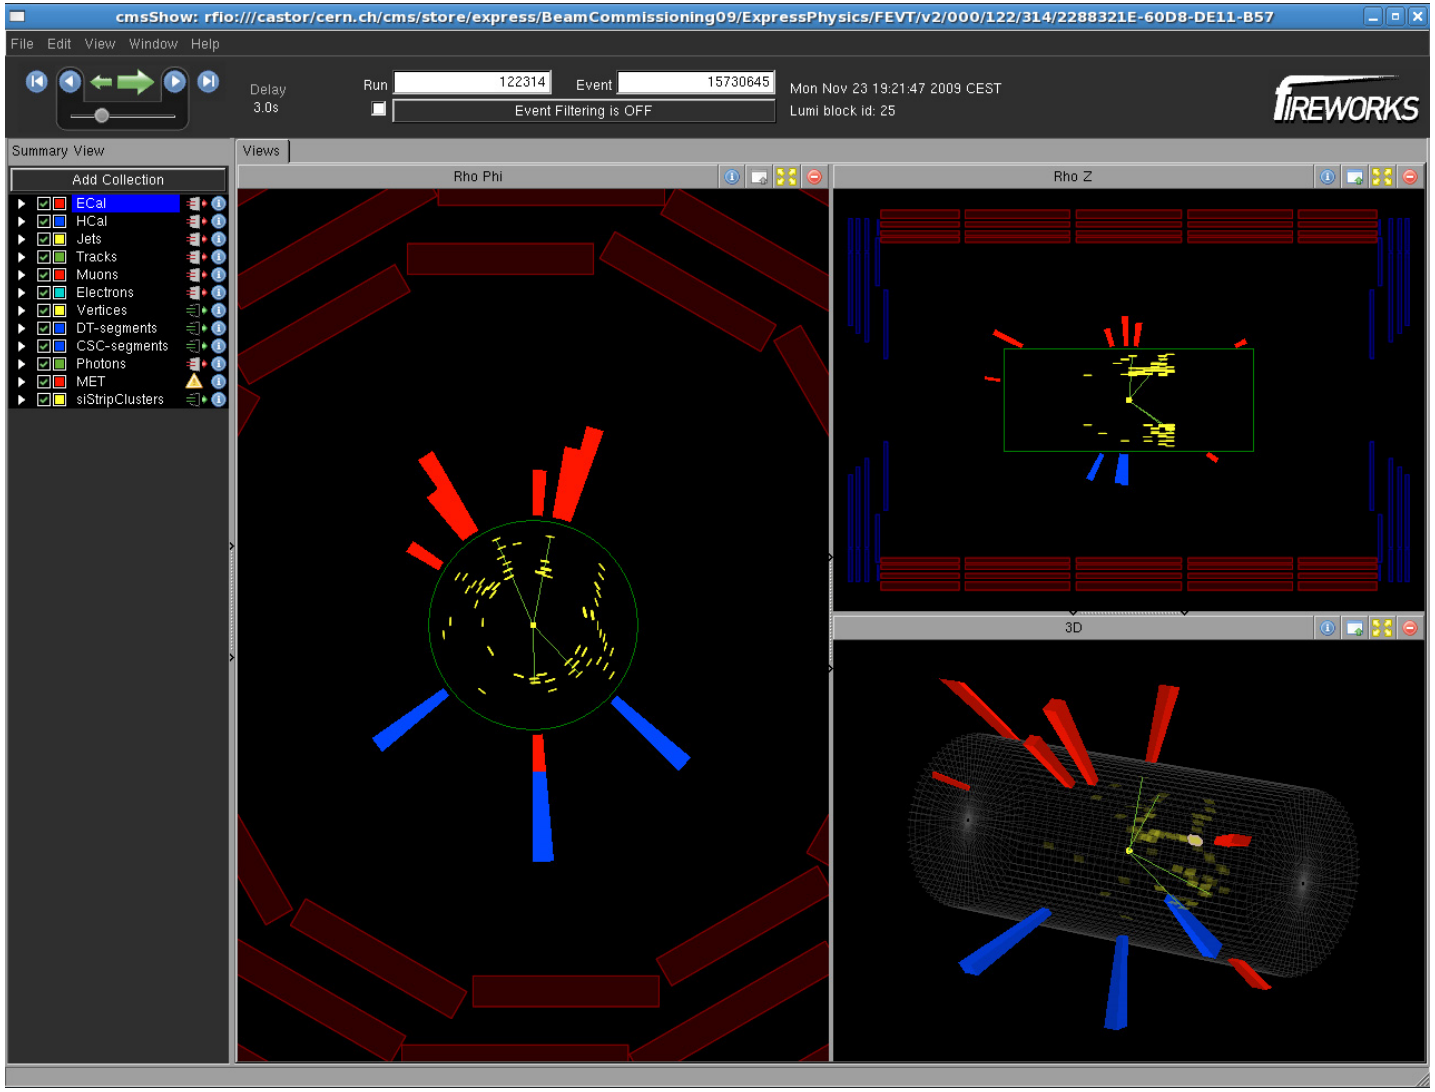

CERN LHC Progress

- It's been an eventful day at CERN. It started with two circulating beams and ended with first collisions in four detectors.
  - 11:39 AM Nov 23rd from web
- There was a press conference earlier in the day, taking stock of what happened over the weekend
  - 12:04 PM Nov 23rd from web
- CERN's press release has pictures of some of the first collisions
  - 12:38 PM Nov 23rd from web
- The LHC accelerated a beam to 540 GeV overnight.
  - 12:46 AM Nov 24th from web

CERN LHC Progress II

- RT @CMSexperiment: World Record!! Tonight at about 22:00 the LHC accelerated a beam of protons to 1180 GeV - a new record energy!
  - 3:59 PM Nov 29th from web
- A new record. Both beams in LHC reach 1.18 TeV at 00:42 on 30 November.
  - 4:04 PM Nov 29th from web
- LHC set world record for beam energy last night. Twin beams circulated at 1.18 TeV
  - 12:59 AM Nov 30th from web
- LHC beam-commissioning progressing, with first "multi-bunch" beams last night (2 bunches per beam).
  - 2:16 AM Dec 5th from web
- LHC has started running with stable beams at 450 GeV (4 bunches per beam). Collision data is being recorded by the experiments.
  - 1:18 AM Dec 6th from web
- **The LHC has run with solenoid magnets in the ALICE, ATLAS and CMS detectors turned up to full field.**
  - **5:42 AM Dec 1st from web**

CERN LHC Progress III

- Last night the LHC accelerated both beams to 1.18 TeV with 2 bunches per beam for the first time.
  - 2:48 AM Dec 9th from web
- LHC beam commissioning continues, aiming towards higher intensities at 450 GeV.
  - 2:56 AM Dec 9th from web
- LHC has started higher intensity running at 450 GeV. Target is 1 million events for the experiments.
  - 5:19 AM Dec 11th from web

- **A very good weekend for the LHC: well over 1 million collisions at 900 GeV and about 50 000 collisions at 2.36 TeV**
  - **Dec 11th from web**

CERN Press Release

The LHC circulated its first beams of 2009 on 20 November, ushering in a remarkably rapid beam-commissioning phase. The first collisions were recorded on 23 November, and a world-record beam energy was established on 30 November. Following those milestones, a systematic phase of LHC commissioning led to an extended data-taking period to provide data for the experiments. Over the last two weeks, the six LHC experiments have recorded over a million particle collisions, which have been distributed smoothly for analysis around the world on the LHC computing grid.

"Council is extremely pleased and impressed by the way the LHC, the experiments and the computing Grid have operated this year," said President of Council Torsten Åkesson. "The laboratory set itself an ambitious but realistic programme at its February planning meeting. The fact that all the objectives set back then have been achieved is a ringing endorsement of the step-by-step approach adopted by the CERN management."

A technical stop is needed to prepare the LHC for higher energy running in 2010. Before the 2009 running period began, all the necessary preparations to run up to a collision energy of 2.36 TeV had been carried out. To run at higher energy requires higher electrical currents in the LHC magnet circuits. This places more exacting demands on the new machine protection systems, which need to be readied for the task. Commissioning work for higher energies will be carried out in January, along with necessary adaptations to the hardware and software of the protections systems that have come to light during the 2009 run. Taking advantage of the stop, the CMS experiment will upgrade part of its water cooling system.

"So far, it is all systems go for the LHC," said CERN Director General Rolf Heuer. "This first running period has served its purpose fully: testing all the LHC's systems, providing calibration data for the experiments and showing what needs to be done to prepare the machine for a sustained period of running at higher energy. We could not have asked for a better way to bring 2009 to a close."

Heavy Ion Future is Brighter ….

- What hasn't been discussed before is the heavy-ion schedule, which in this plan has 2+4 weeks instead of just 4 weeks, with collisions scheduled for the whole month of November after 2 weeks of setup in October. Clearly, if this is the "official" plan and eventually becomes reality, we might get significantly more data (more than factor 2) compared to estimates based on the 2009 schedule, where the 2-week setup was part of the 4 week HI beam time.

Bolek's Email ….

On 12/01/2009 06:54 AM, Bolek Wyslouch wrote:

> Here they come (FED stuck)
>
> -------- Original Message --------
> Subject: pixel monster events
> Date: Tue, 1 Dec 2009 14:51:40 +0100
> From: Danek Kotlinski <danek.kotlinski@psi.ch>
> To: <hn-cms-pixelOnlineSW@cern.ch>
>
> *** Discussion title: Pixel Online Software Development
>
> Hello,
> From last evening we have a few monster events.
> From 260 events preselected by Vincenzo about 10 have number of hit pixels comparable to the high-lumi (10^34) running, 2-3 exceed it.
> The occupancy is also not uniformly spread like in high-lumi simulations. It is concentrated in 50%-25% of the modules, which means that some modules have a very large pixel occupancy.
> This is not noise, these are tracks going along detector planes.
> I expected some hallo tracks, but that would be just a few tracks at a time.
> Instead we get a massive shower (from forward region?) with many tracks
> in 25ns. These tracks also go through the fpix but there they just make small clusters, in the bpix they form very long clusters with many pixels. Are we hitting something?
>
> We do not see in DQM because these are just a few events, but these are
> the events which get the FEDs stuck.
>
> I have hard time to visualize these events in a nice way.
> If somebody has a good way to visualize it these are the event numbers: 26, 117, 127, 299, 291, 262 from Vincenzos express file.
>
> Danek

Expected pp Performance

PAUL SCHERRER INSTITUT

Simulated data loss

- high luminosity LHC: [10^{34}]
- 11 cm / 7 cm / 4 cm layer
- total data loss @ 100kHz L1A:
  - 0.8%
  - 1.2%
  - 3.8%

Pixel waiting:
0.04% / 0.08% / 0.21%
pixel insensitiv until hit transferred to data buffer (column drain mechanism)

Column busy:
0.004% / 0.02% / 0.25%
Column drain transfers hits from pixel to data buffer. Maximum 3 pending column drains requests accepted

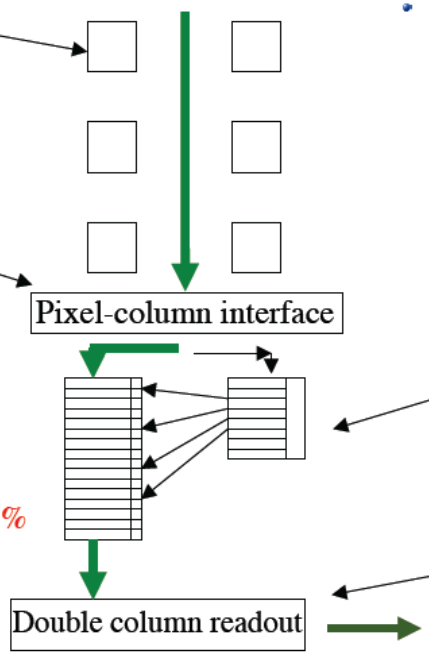

Timestamp Buffer full:
0 / 0.001% / 0.17%

Data Buffer full:
0.07% / 0.08% / 0.17%

Readout and double column reset:
0.7% / 1% / 3.0%
for 100kHz L1 trigger rate

Hans-Christian Kaestli Pixel 2005, Bonn 5.9.2005

Detector Readiness

- Detector Readiness Workshop in October
  - Two LANL talks !
    - Catherine on RPC (Muon Detectors) Readiness
    - Gerd on Pixel Readiness (+Andi at CERN)
      - FPGA limits simulated, predicted and discussed
- Detector Readiness Review in April
  - Plan is to present measurements and solutions for the Pixel FED problem

Current Readout

- Transfer from FIFO I only after complete event is received for group of 4/5 FIFO I
  - That means waiting for the longest row of hits to arrive in any of the 4 or 5 before continuing in a **SEQUENTIAL** pattern
- Transfer from FIFO II only after complete FIFO I is received
- FIFO II is the slink buffer

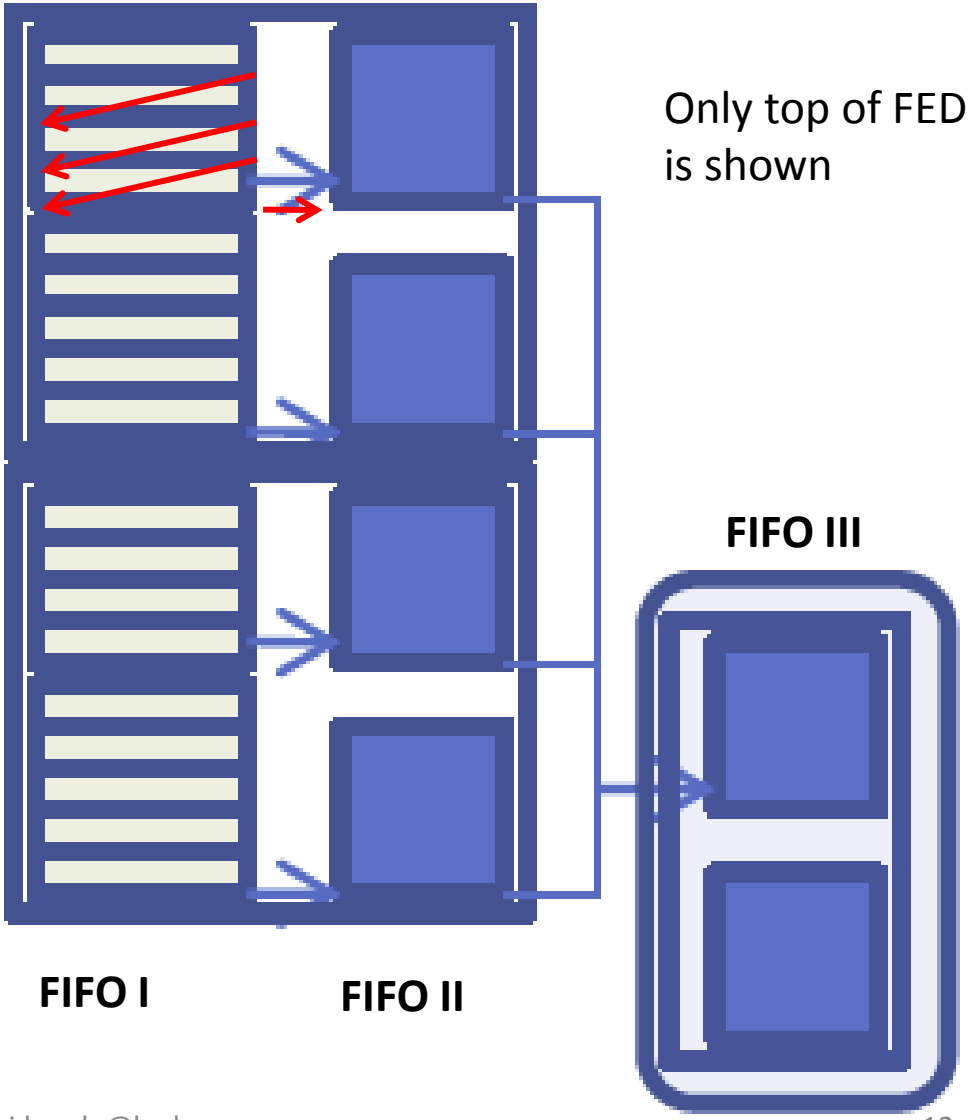

Only top of FED is shown

Obvious bottleneck is FIFO I length …

Definitions

- There are two ways to measure the fill level of FIFO I
  - Fill level at time of next event
    - Used in Ivan's studies
  - Maximum fill level between events
    - Obviously different in 10^25 running
    - Fifo I is always empty at the next event but obviously filled up in between
  - Proposed early, after initial simulation results

Silicon Pixel Readiness

Ivan A. Cali MIT, Gerd J.Kunde,
A. Klein, C. Mironov LANL

Outline

- Simulation Details
- Occupancies 4 TeV and 5.5 TeV
- ROC Performance
- Data Volume
- FED Performance Year One
- FED Performance Year Two
- Conclusions

Pixel Simulations

4 TeV and 5.5 TeV

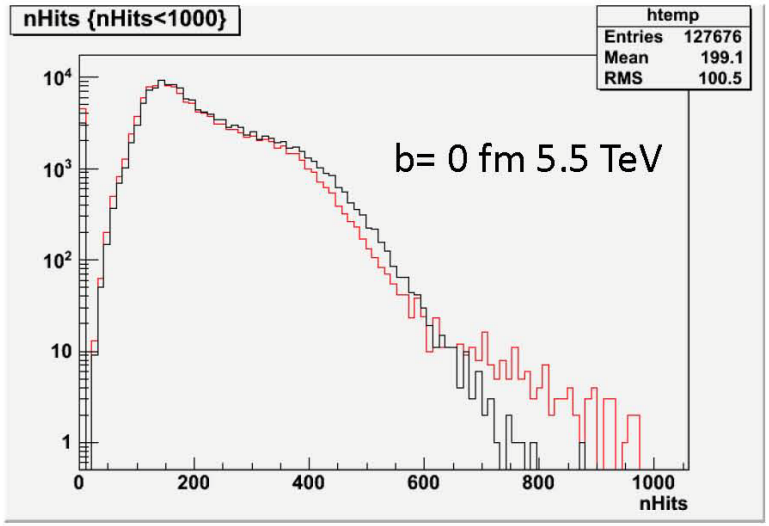

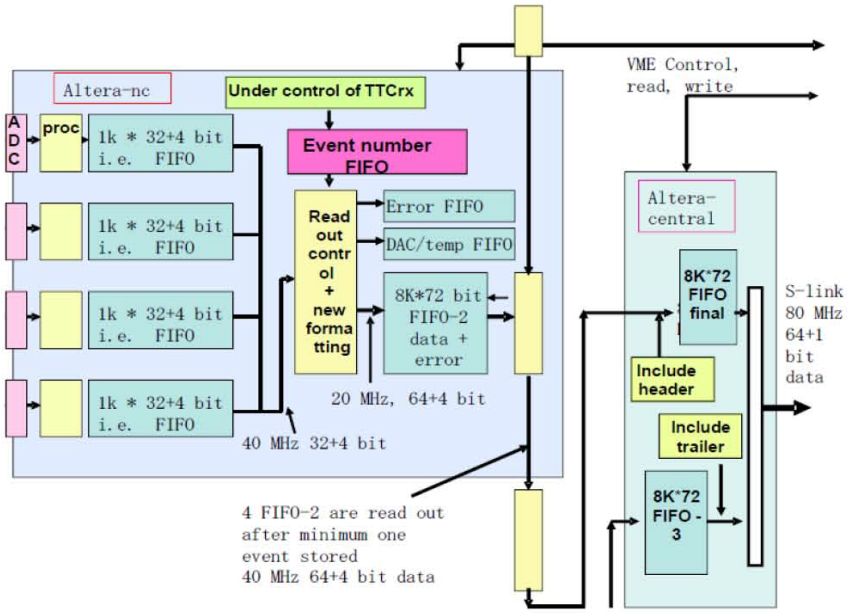

2_2_11 hydjet at 5.5 TeV from Oct 2008
98 events for b0 and about 8600 events for minbias
3_1_1 hydjet at 5.5 TeV simulated by LANL
100 b0 events and 1000 minbias

FED Simulation of FIFOI, II and III performance in ROOT by I.A.Cali and G.J.Kunde

Desirable to perform VHDL simulations and actual test with HI simulation input files

g.j.kunde@lanl.gov 10/14/2009

3

Pixel Occupancies 4 TeV and 5.5 TeV

central

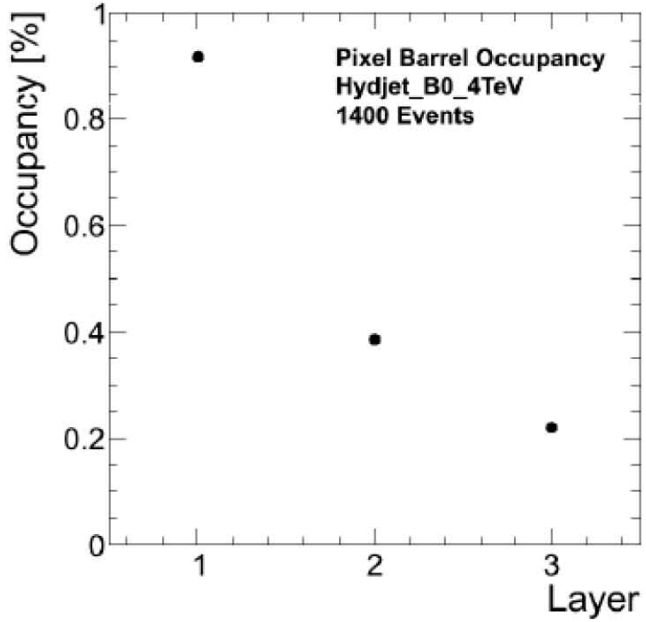

minbias

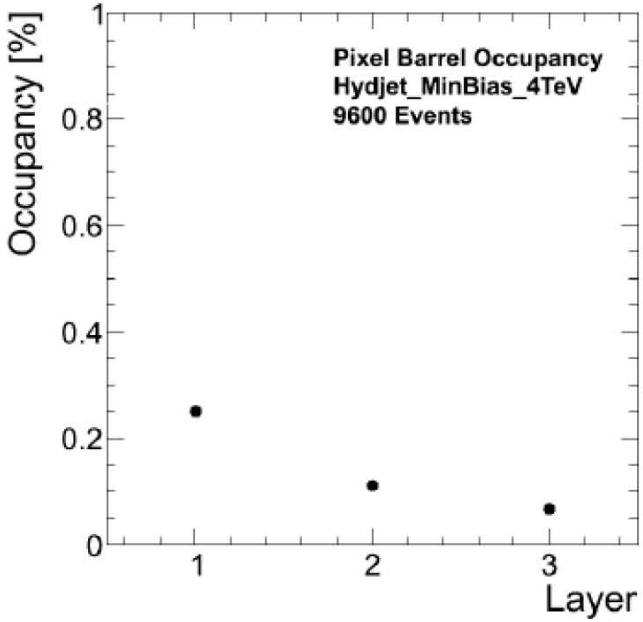

The occupancy in 5.5 TeV is ~ 1.1 times the 4 TeV results

Occupancy low, even for central collisions

Roc Multiplicities b=0 for 5.5 GeV

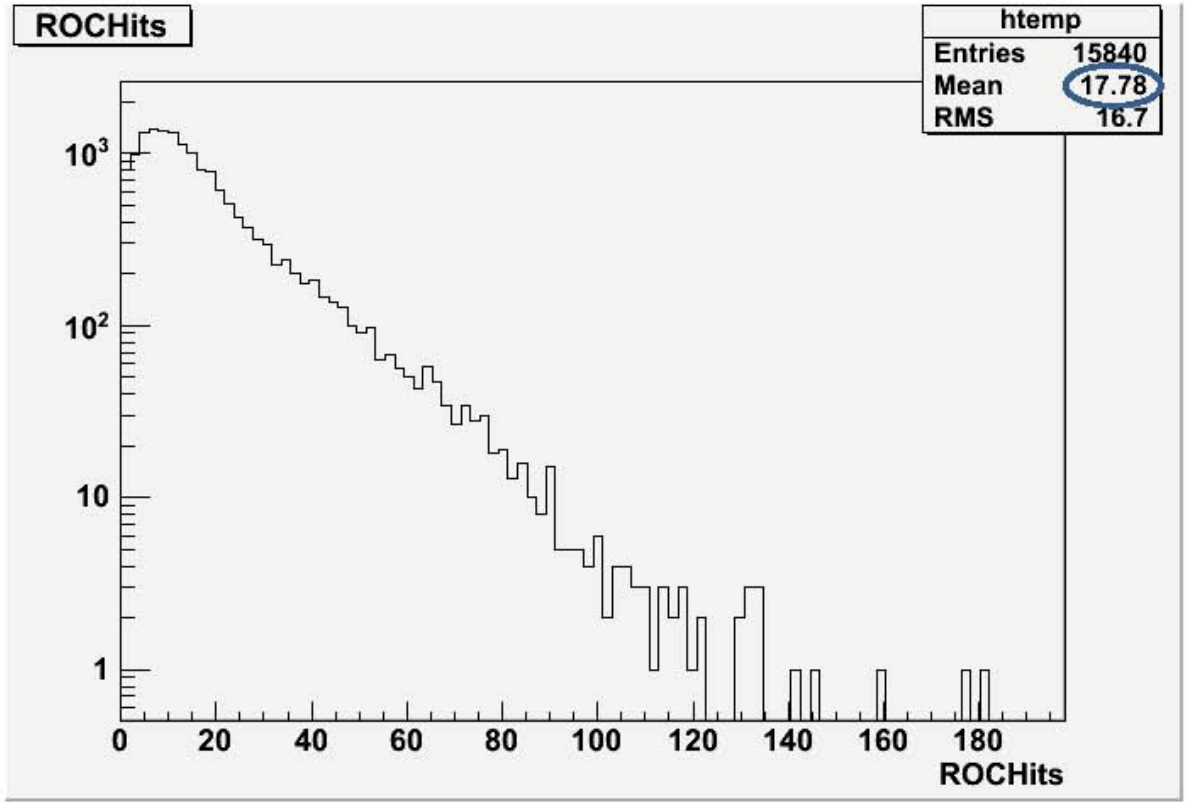

No issue per se , except for possible amplitude loss

Speculative ROC I_d Effects

- The current distributed inside the ROC may not be able to feed all the activated pixel cells in case of high multiplicity. This effect could introduce:
  - Reduction of ADC counts – order 10% in first tests
  - Hit loss ?
  - Uniformity of these effects over the ROC surface is not clear ?
- Instantaneous I_d drop due to high multiplicity following "long" (>micro second) period of ROC inactivity:
  - Threshold reduction for next event ?
  - ROC Reset ?
- Need manpower for tests

The D-Column

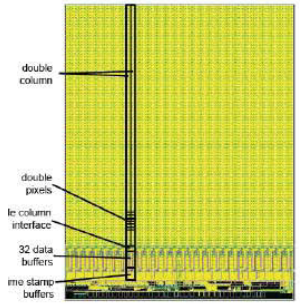

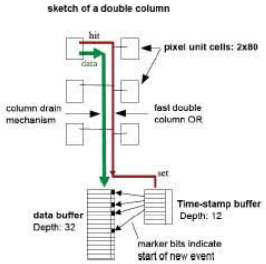

Pixel over threshold: creates column OR :-> leads to column drain through token (2 cycles), pixel insensitive until column drain.

At this time bunch crossing stamp latched into Time stamp

Column drain is 50ns*#pixels, (can have max of 1 drain and 2 pending)

If L1 accept and trigger time corresponds to timestamp: readout of ROC starts. Governed by TBM : TBM sends token to ROC1 … ROC 8 sends it back to TBM. Data transferred: 6 cycles/pixel: dcol (2), pix add (3), pix analog(1). Each ROC transfers at least a header. Only way to determine the ROC id is counting in a sequence

g.j.kunde@lanl.gov 09/04/2009

D Column Hits b=0

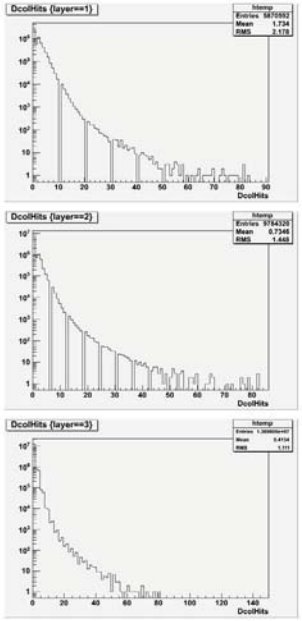

~4*10^-5 loss of dcolumns in layer 1

~2*10^-5 loss of dcolumns in layer 2

~1*10^-5 loss of dcolumns in layer 3

Considered acceptable

g.j.kunde@lanl.gov 10/14/2009

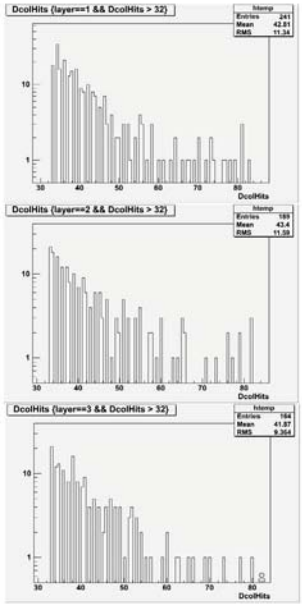

Pixel Data Volume for PbPb

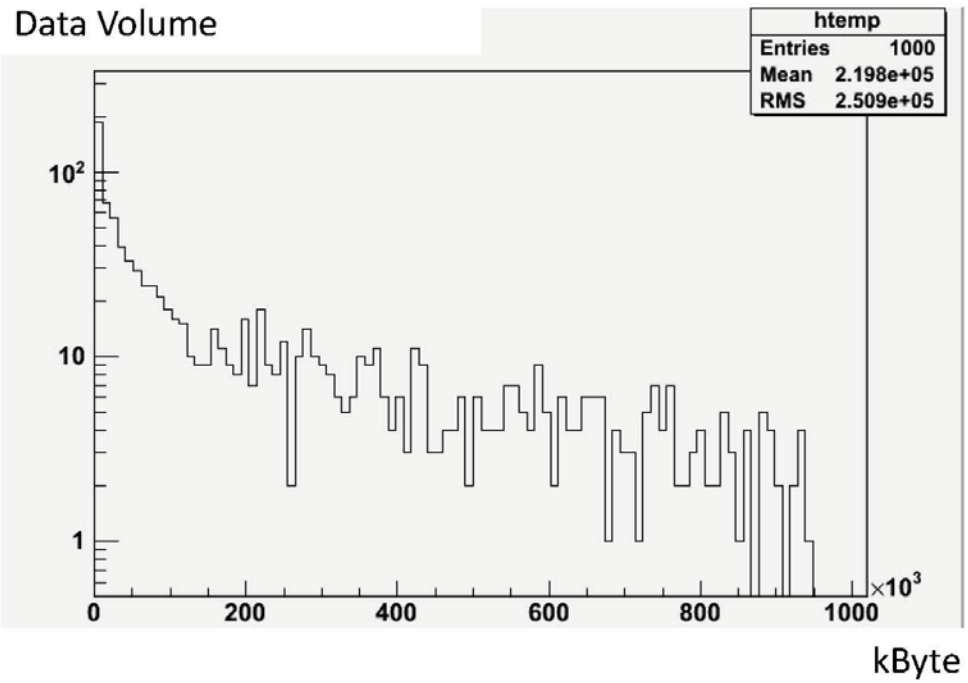

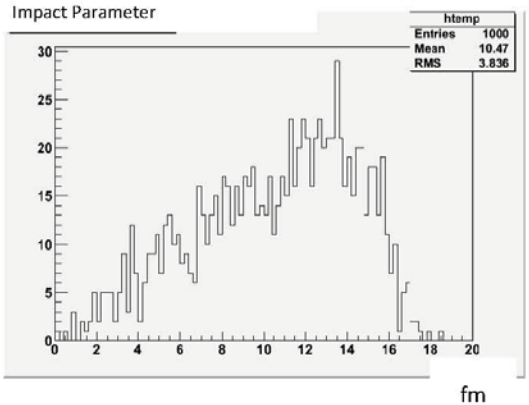

The minbias average is 220 kByte, the max around 1 Mbyte

Multiplicities vs. LINK # for b=0

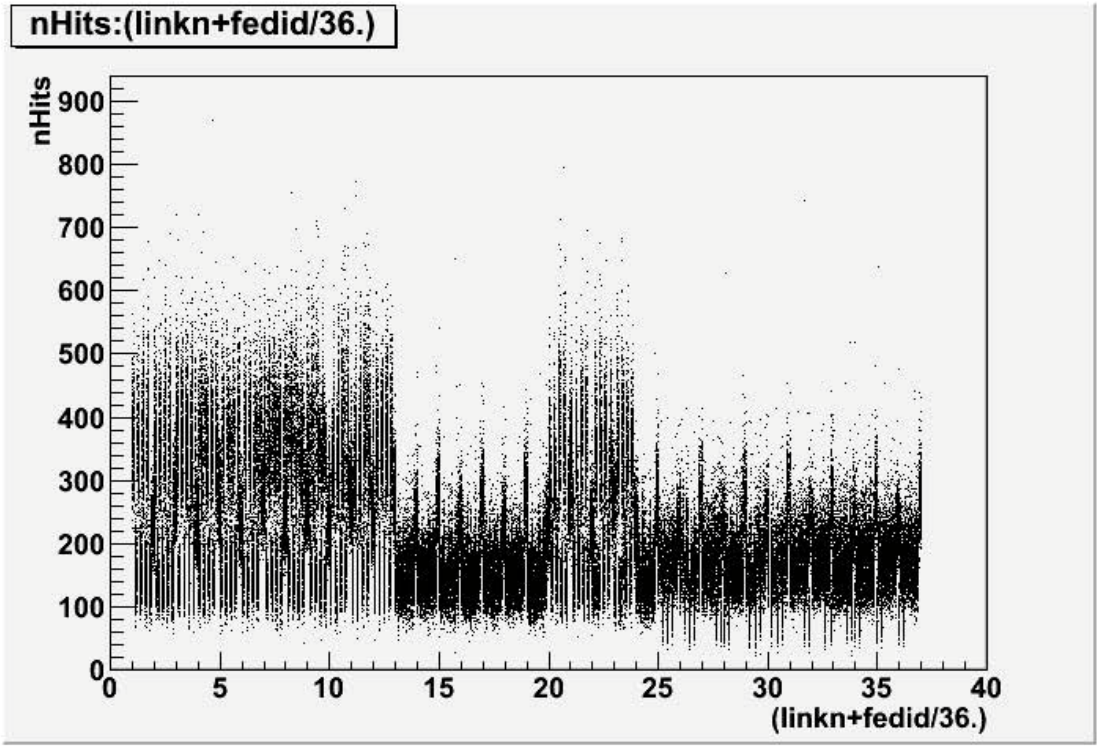

g.j.kunde@lanl.gov 10/14/2009 10

Total Hits per FED for b = 0 fm

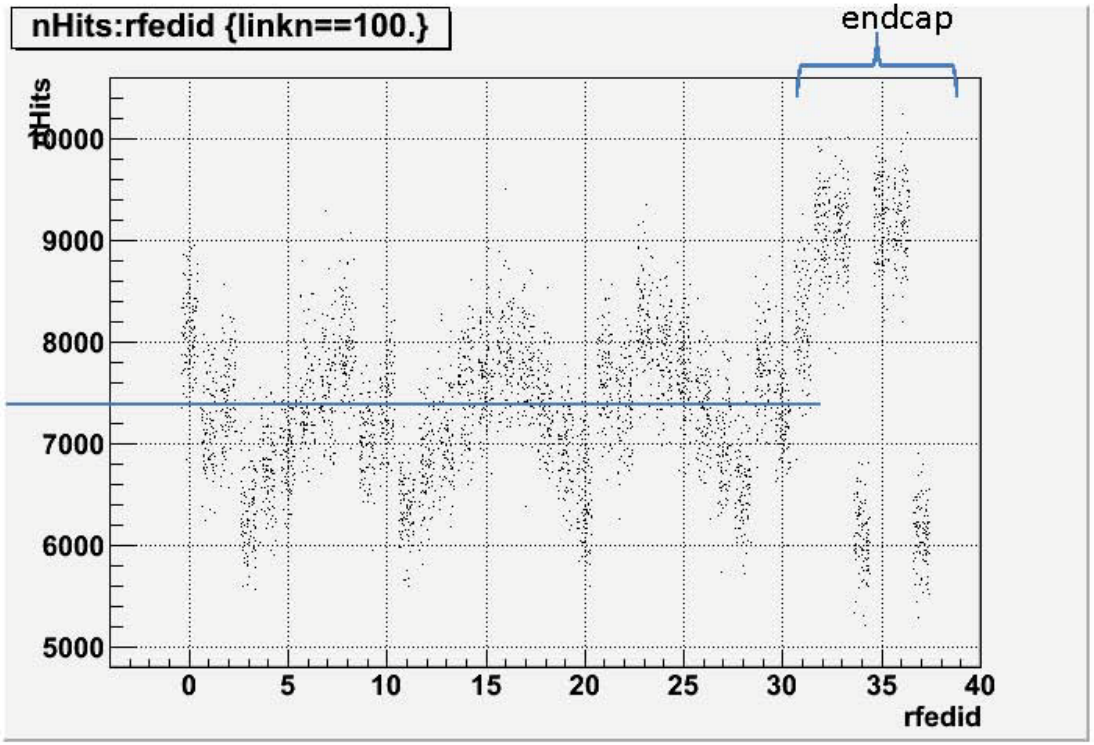

The system was designed for pp with 500 hits/FED and 100 Khz L1, safety factor in HI ? 11

g.j.kunde@lanl.gov 10/14/2009

Year one

- Luminosity of 5 * 10^25
- Translates to and average of 385 Hz hadronic minimum bias interactions
- Instantaneous rate is higher !
- Only 62 filled bunches for P5

Reminder : 7.7 barn Pb+Pb at *L=5*10^25*

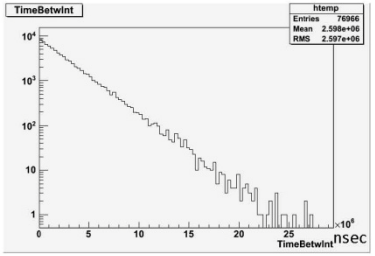

This is the time between filled bunches for *L=5*10^25* used in the FED simulations

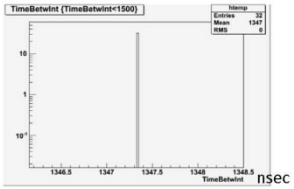

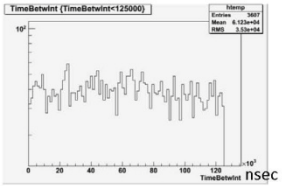

4.12 * 10^-4 of the events are two subsequent bunches

4.68 % of the events are "faster" than 8kHz

Fifol Max Level for 5*10^25

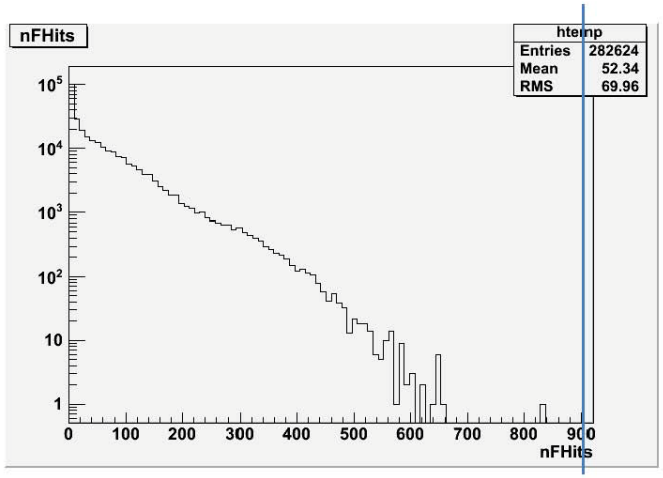

- Looks (as expected) like the link occupancy distribution
- Safety factor ? Foresee dead time and busy ?

"Year 2"

- Luminosity of 10^27
- 592 bunches

Comparing to the fixed 8 kHz Assumption

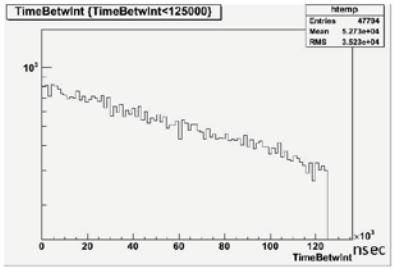

61 % of the collisions have an instantaneous rate faster than 8 kHz

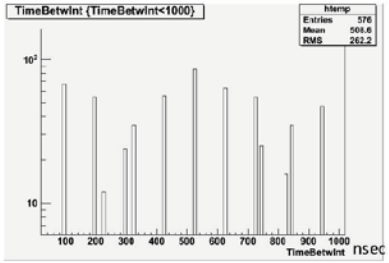

0.75 % happen within a micro second (different from year one)

~0.087 % happen in the next filled bunch at 99.803 nsec

Reminder:
Time Between Hadronic Interactions (10^27)

10*Flhc*Nbb events to recover 10 * 7700 mbarn

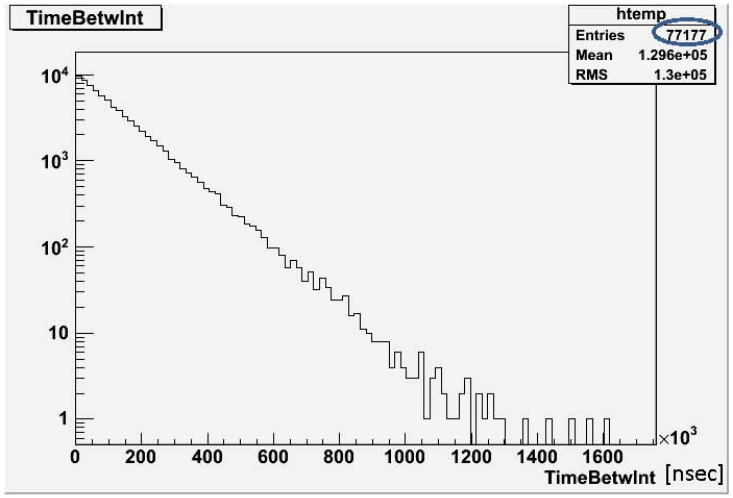

~8 Khz

Fifo Fill Level at Time of Next Event

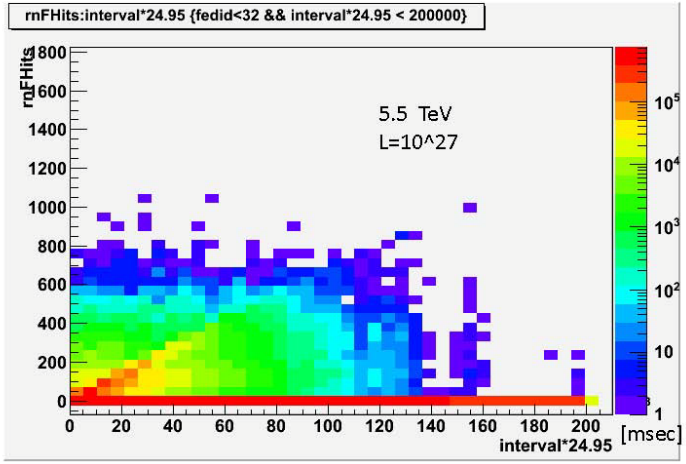

Diagonal structure due to 150 nsec fill speed of FIFO I

Fifo I Overflow
Confirmed by two independent MCs

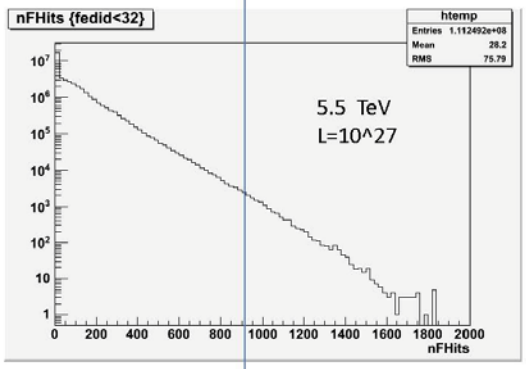

- ~ 150 k minbias events at 5.5 TeV Pb Pb at 10^27
- Fifo I overflows occur with an average rate of ~600/sec
- Better algorithm or better hardware or reduced Level I rate ?

Preliminary Fifo II, Fifo III and SLink …

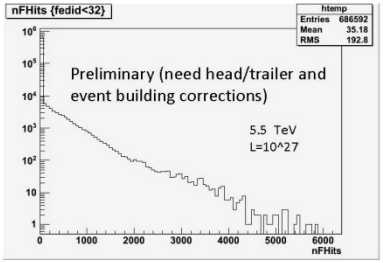

FIFO II

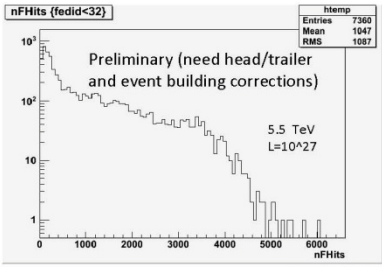

FIFO III -> SLink

Within limits, safety factor ?

Heavy Ion Running Conclusions

- Occupancies below one percent
- ROC should perform reasonably well
  - Caveat are possible I_d effects
- Average data size ~200 Kbyte for central PbPb
- "Year one" running should be ok (dead time ?)
- "Year two" running needs FED modification to allow triggering at full level I speed
  - Are firmware changes enough ?
  - New FPGA mezzanine boards ?
- Need FED diagnostics and realistic testing soon

CERN Discussions

- Met several times with the Vienna Group that build the FED
- Met with Danek from PSI
- Met Will from Vanderbilt

- Conclusions are many fold:
  - Need a Heavy-Ion test system
  - Need probably a new set of mezzanine boards
  - HEP buy-in essential ….

Alternate Readout with Templates

- Transfer from FIFO I only after complete event is received for group of 4/5 FIFO I
  - That means waiting for the longest row of hits to arrive in any of the 4 or 5 before continuing in a **PARALLEL** pattern
- Transfer from FIFO II only after complete FIFO I is received
- FIFO II is the slink buffer

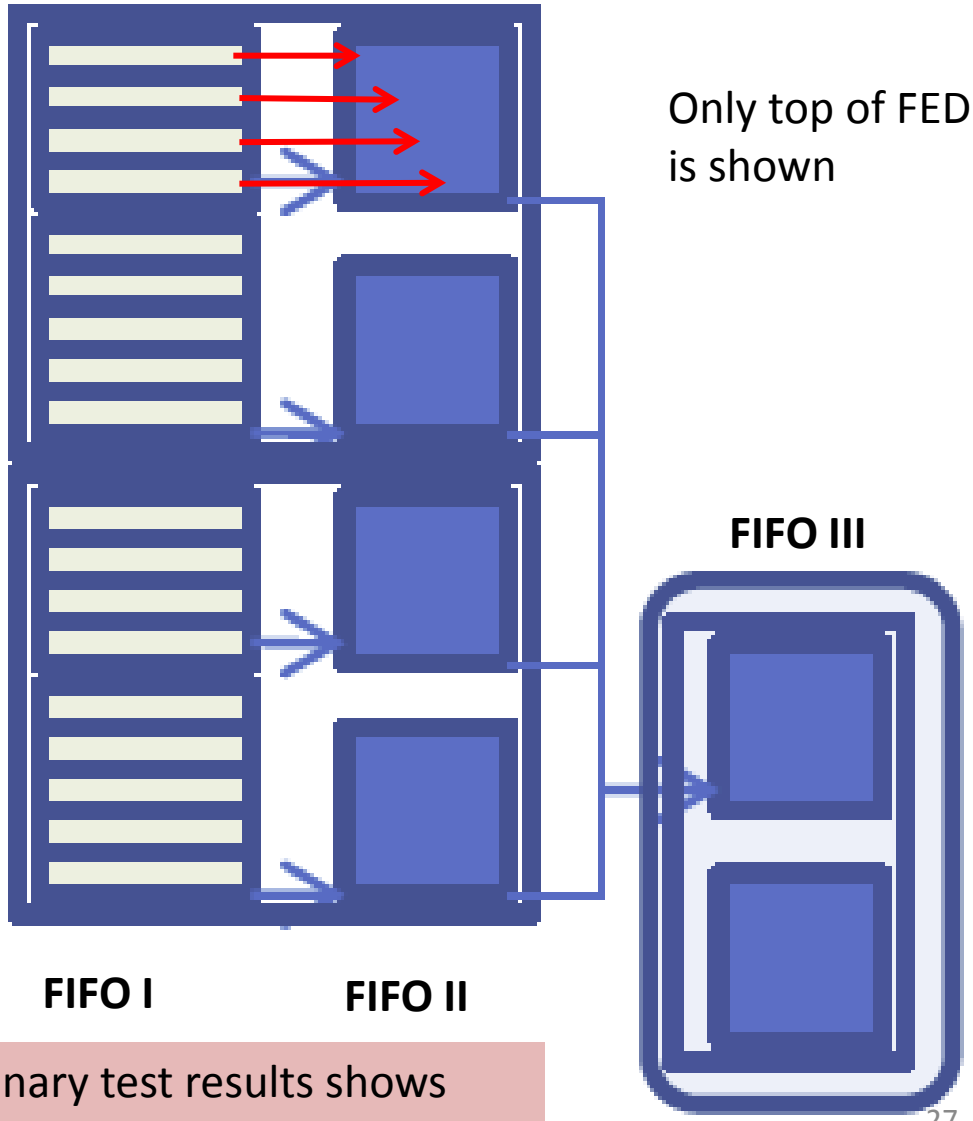

Only top of FED is shown

Bottleneck will be the FIFO II length but that is bigger, system more complex

Preliminary test results shows nearly no gain !

Alternate Readout with Vienna Idea

- Transfer from FIFO I **as soon as a hit** is received **in any** FIFO I
  - That means **NO** waiting for the longest row of hits to arrive in any of the 4 or 5 before continuing in a **PARALLEL** pattern
- Transfer from FIFO II only after complete FIFO I is received
- FIFO II is the slink buffer

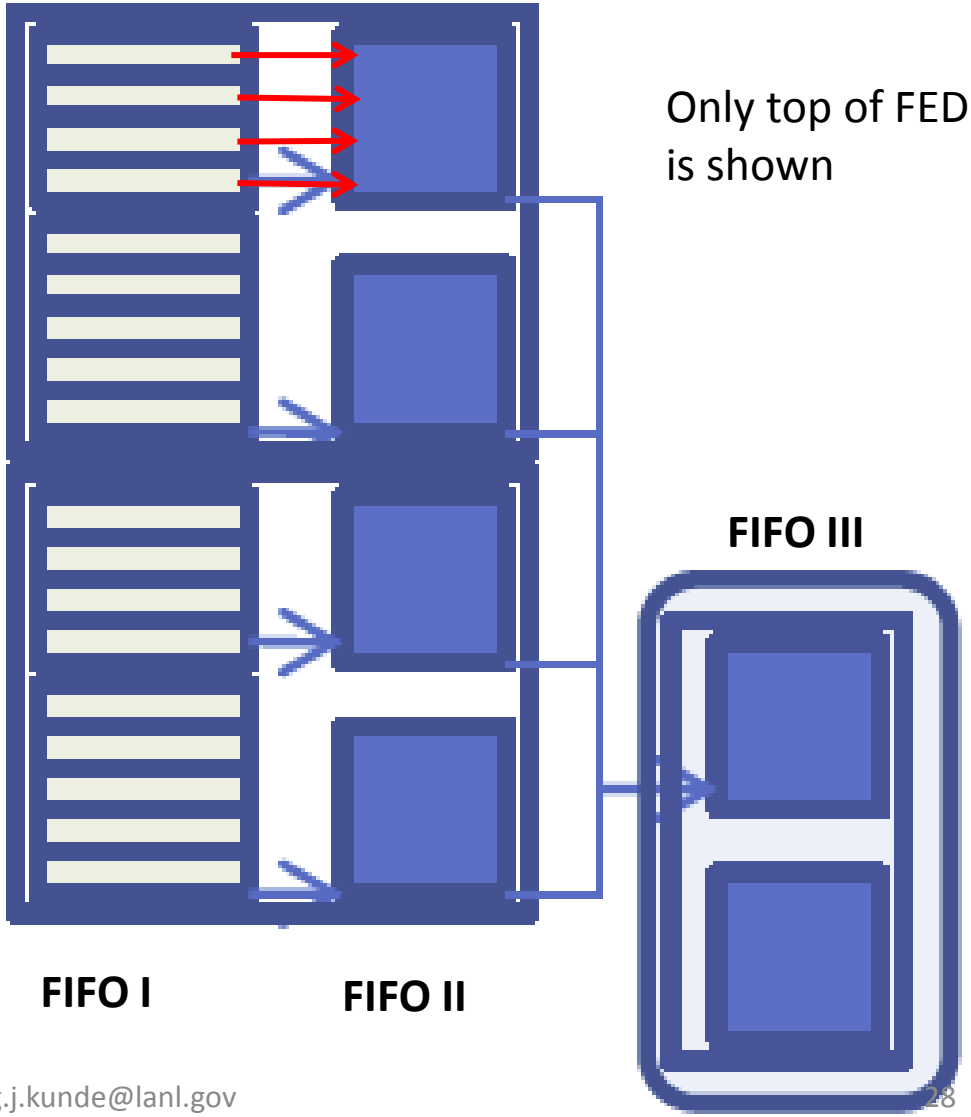

Only top of FED is shown

Bottleneck will be the FIFO II length but that is bigger, system more complex …

g.j.kunde@lanl.gov

Testing of the Front End Driver

9 Fifo-1's on each FPGA

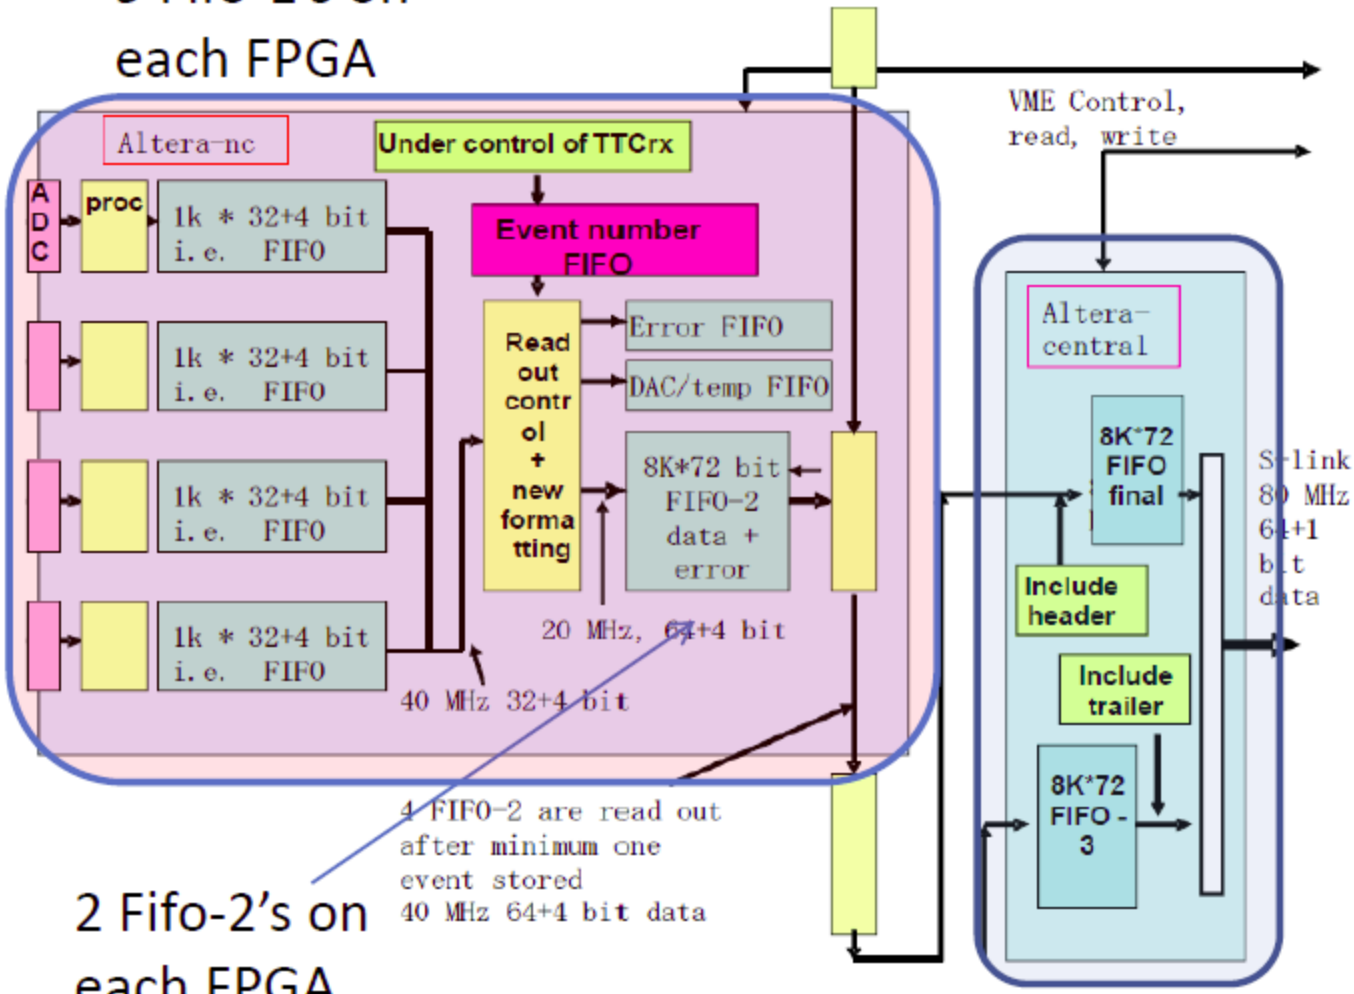

2 Fifo-2's on each FPGA

Testing of the Front End Driver

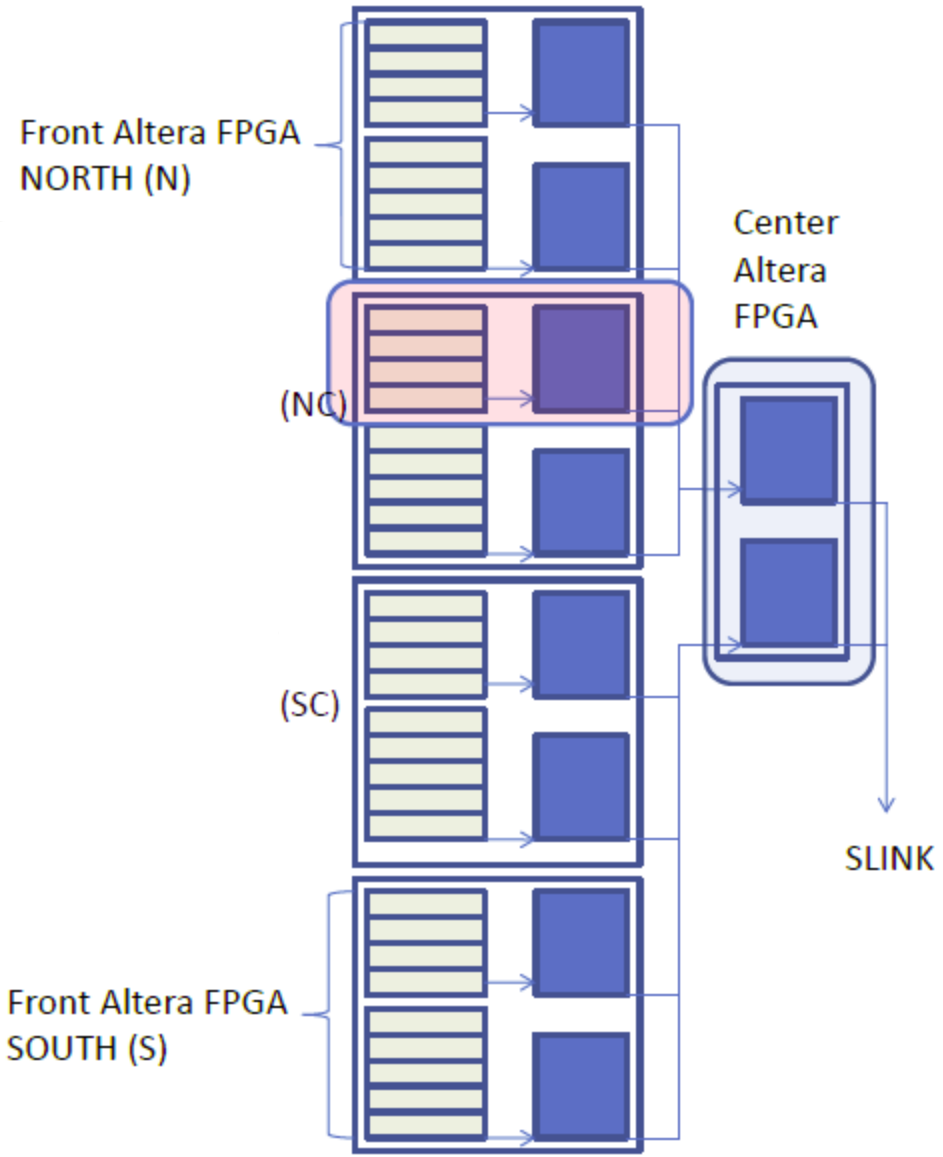

36 Channels Test by Will

- Early Test by Will
- Used missing trailer to overload all 36 channels with 800 hits
- Checked for error messages from FED
- Minimum error free period is 10,450 cycles
- Was translated in 300 micro second dead time requirement
- But this will never happen in even a central event …

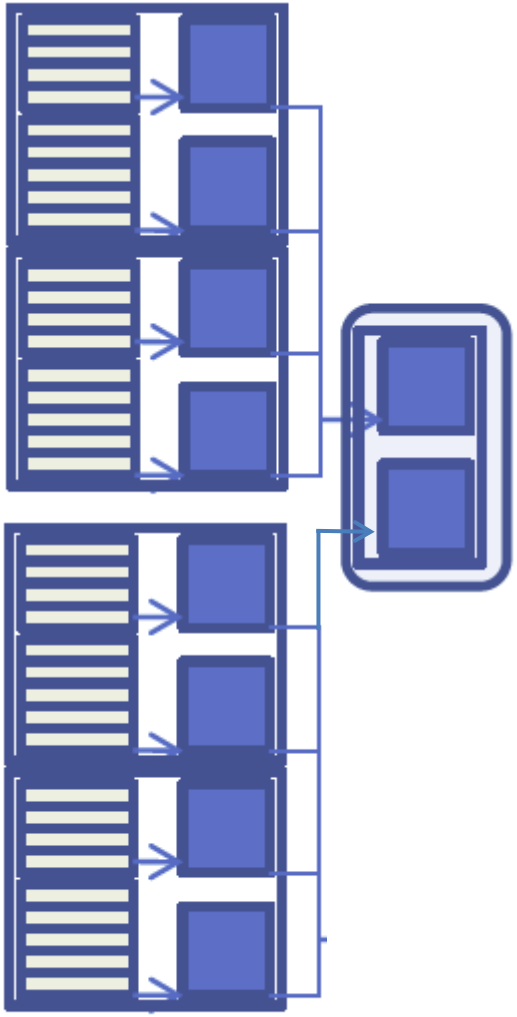

New Testing

- The 300 micro seconds was thought to be due to the fifo I overload
- Designed new suit of testing to test hypothesis
  - Max number of hits in fifo I
    - 920 and reduced 828
  - 18 channel test to verify fifo I problem
  - 9 channel test to check fifo II
  - 4 channel test to check fifo I -> fifo II
  - 3 channel test
  - 2 channel test

Testing 18 Channels

- Minimum save interval is 7555
- NOT 10,450 as expected

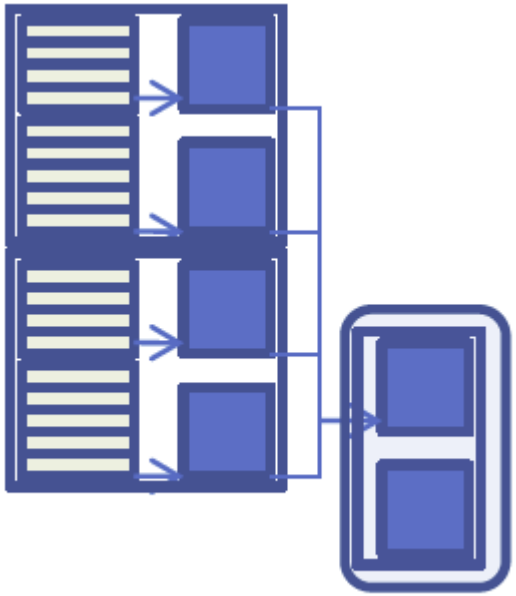

Testing 9 Channels

- Minimum save interval is AGAIN 7775
- Checked FIFO I length dependence
  - Reduced 920 to 828
  - Minimum goes up to 8078

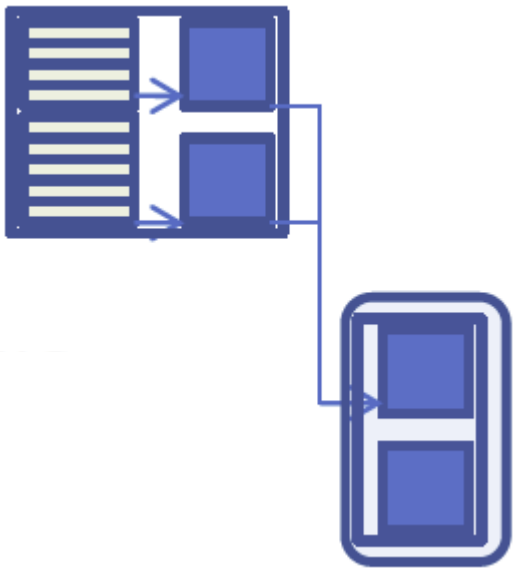

Testing 4+ Channels

- Minimum save interval is 6736
- Implemented small additional load of 30 hits per channel in software
- SAME minimum interval of 6736
- Changed FIFO I length
  - 920 to 828
  - Same (as in 9 channels) incremental change in interval +522

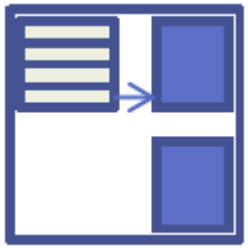

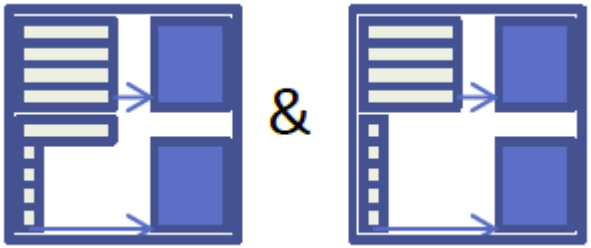

Testing 3 Channels

- Minimum save interval is 5918
- No difference if there are additional small counts in extra channels

Testing 2 Channels

- Minimum safe interval is 5100

- Minimum safe interval is 5148

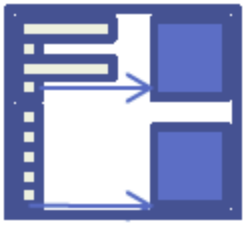

Results in Table Form

Fifo I Channels	Minimum Safe Interval	Delta to previous
18	7555	
9	7555	
4	6736	819
3	5918	818
2	5100	818
1	can't measure	

- Same FIFO II for 18 and 9 channels
- Delta is #hits + #Last DAC(16) + 2 Error words
- State Machine needs more # Intervals to process an 800 hit event
  - 800 hits*6 + Header (8) + Trailer(8) + 16 ROC*3 = 4864

Scary Result ….

- Initial testing with ALL channels at 800 gave 10,450 as an error free save minimum interval
- This translated into ~ 300 micro second dead time ..
- More complete testing shows that even if only one channel per FED has a central event load then we measure 5100 as a minimum safe interval …
- This still translates into ~ 150 micro seconds or loosing about 50 percent of the min bias triggers

Testing with Minbias Average of about 250 hits …

- More precise 36 ch minimum safe interval: 2274
- Group of 18 (top) ch minimum safe interval : 1494
- Group of 9 (top) ch minimum safe interval : 1494 (same)

- Did not get x 2 speed increase, time to process 250 hits in state machine
  - 250*6 + 66 (from prev) = 1566

- Conclusion is complicated because of imperfect testing method !
  - Test buffer is likely getting cut off for 250+
  - Rate can increase since next event can mimic a trailer IFF it happens to come in the right spot

- 1494 translates into a 164 hit/event on average ..
- One Fifo-II should be able to deliver 1 hit every interval to fifo-III
- Expect ultimate limit of (250+2(errors))*9 = 2268 + (gap/filler/header)
  - 2274 limit above agrees with expectations

LANL Simulation

- Installed Altera FPGA software for FED
- Obtained the FED firmware from Vienna
  - Surprise …. Mostly schematic only few blocks in VHDL
- In the process of understanding it (tough)
- Updating the simulations with more details and better understanding of Fifo II and Fifo III
- In the process of reproducing the test results

Remove Fifo I Bottleneck

- Manfred's new method suggests to empty groups of fifo-1's faster than they are filled
- Problems occur now if there is back pressure from the Slink
  - Need new Busy/Warning definitions
- Backlog of hits/triggers from a TBM or two in a fifo-1 grouping ch'swith small #hits will back up in fifo-1
  - Also indicates New Busy/Warning definitions

- Currently debugging LANL simulation to test this concept

Proposal for New Hardware

- Problem with current testing is that it is too limited …
  - Either missing ultra black that is interpreted as missing trailer and translated as limit
  - Or small number
- Idea is to build an FPGA driven system that feeds the FED with test data
  - Same way the simulator works, with HI events and the LHC time structure
- Best would be to plug in directly after the analoge fiber receiver …
- Idea was discussed and confirmed with Vienna and pixel group

Bypassing the Analog Optical Links

Readout

- chip header: 3 clock cycles:
  - ultra black: separates chips in output data stream
  - black: output driver recovers from large negative signal
  - last DAC: analog level of last addressed DAC
- hit information for each hit pixel: 6 clock cycles
  - 5 address cycles. It is encoded in 6 discrete levels
    - double column address: 2 cycles
    - row address: 3 cycles
  - analog pulse height information
- ROC internal risetime: t=2ns

superposition of readouts for all pixels

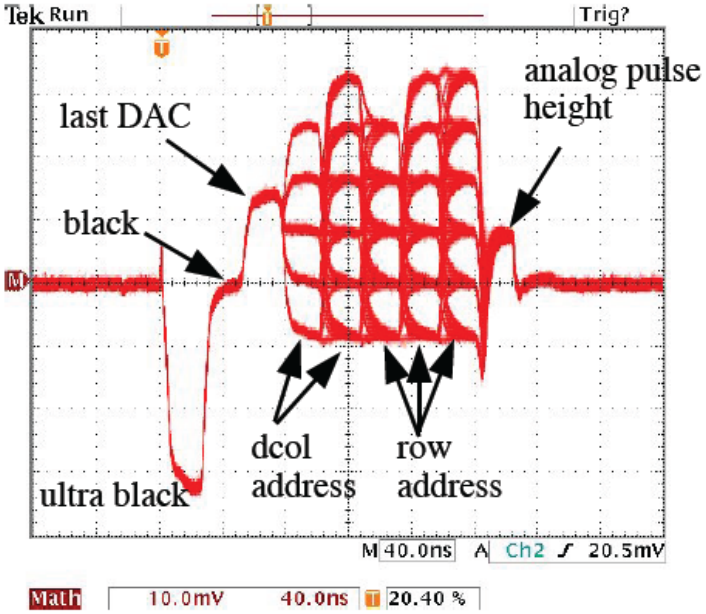

Proposal for new LANL Test Hardware

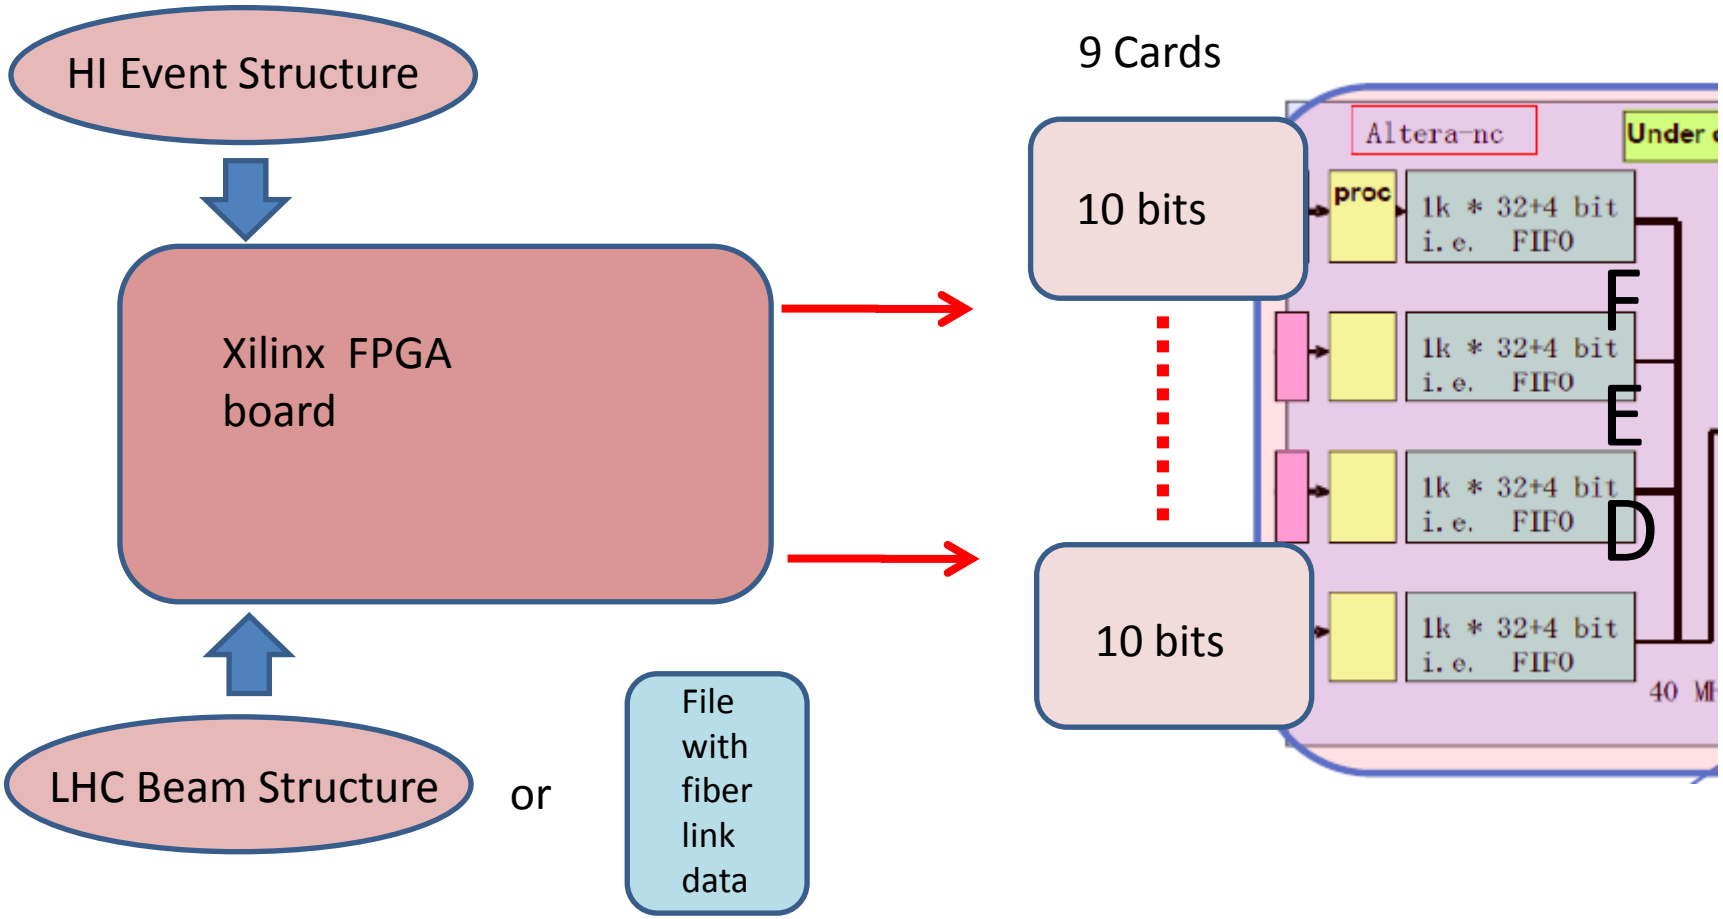

- Replace the fiber translators with serial receivers for FGPA driven HI events with realistic time structure

The FED and the Daughter Board

9U VME Card

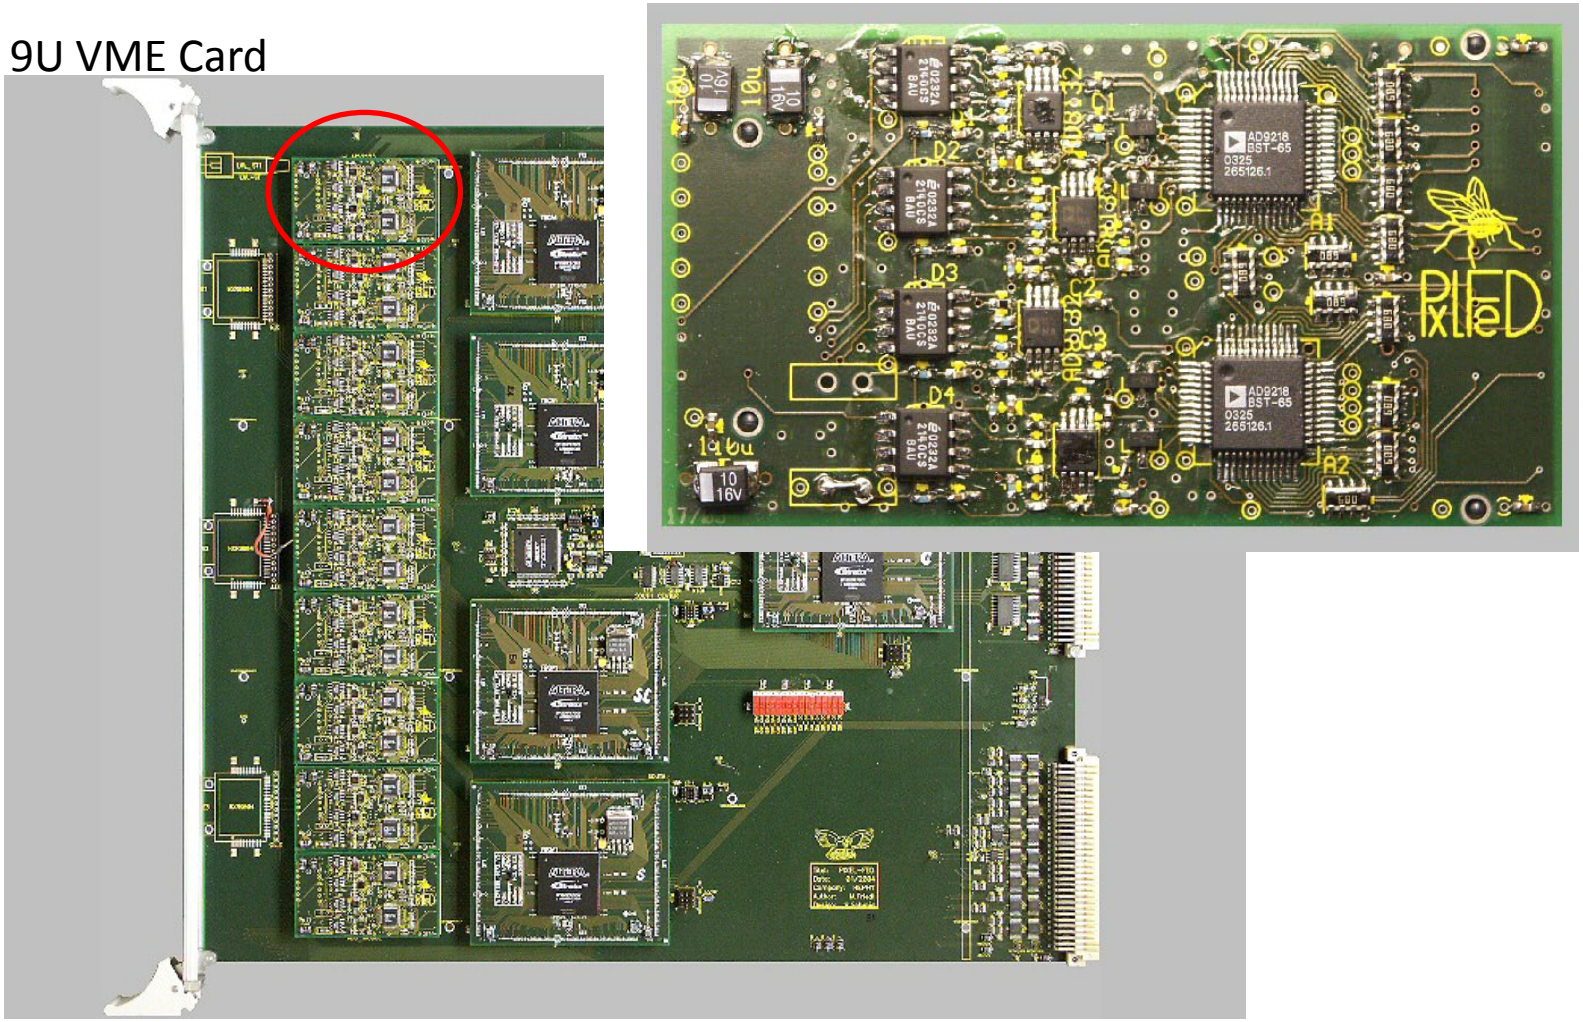

This is the ADC daugther borad

M. Pernicka Hephy Vienna

Where to plug in I

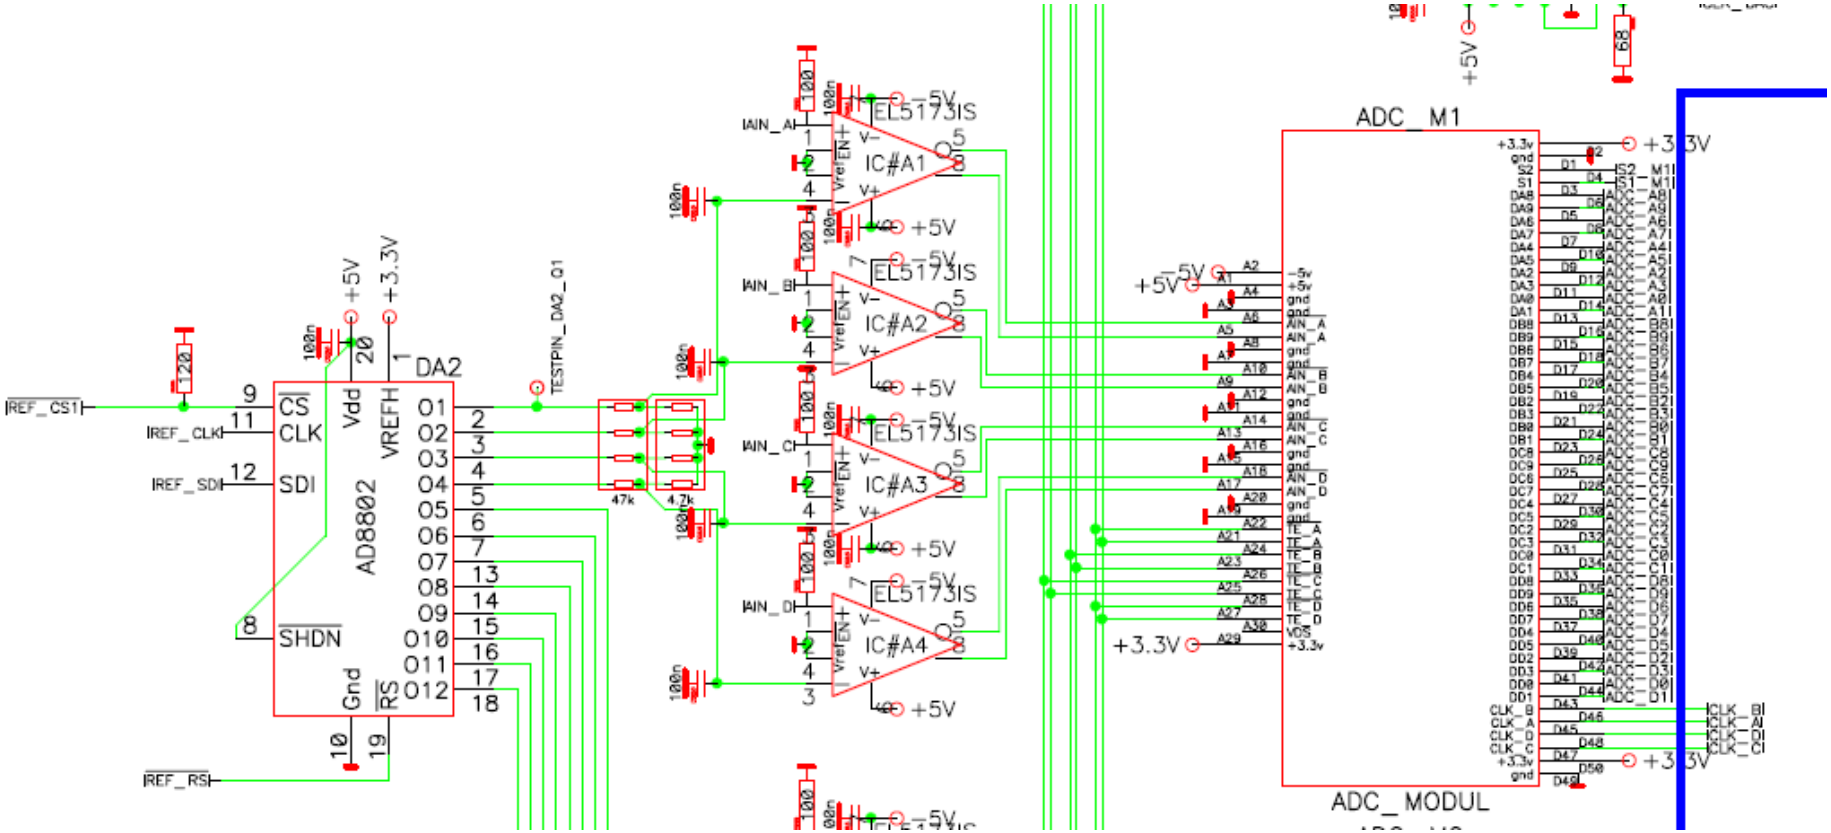

This is card to replace

The FED Board

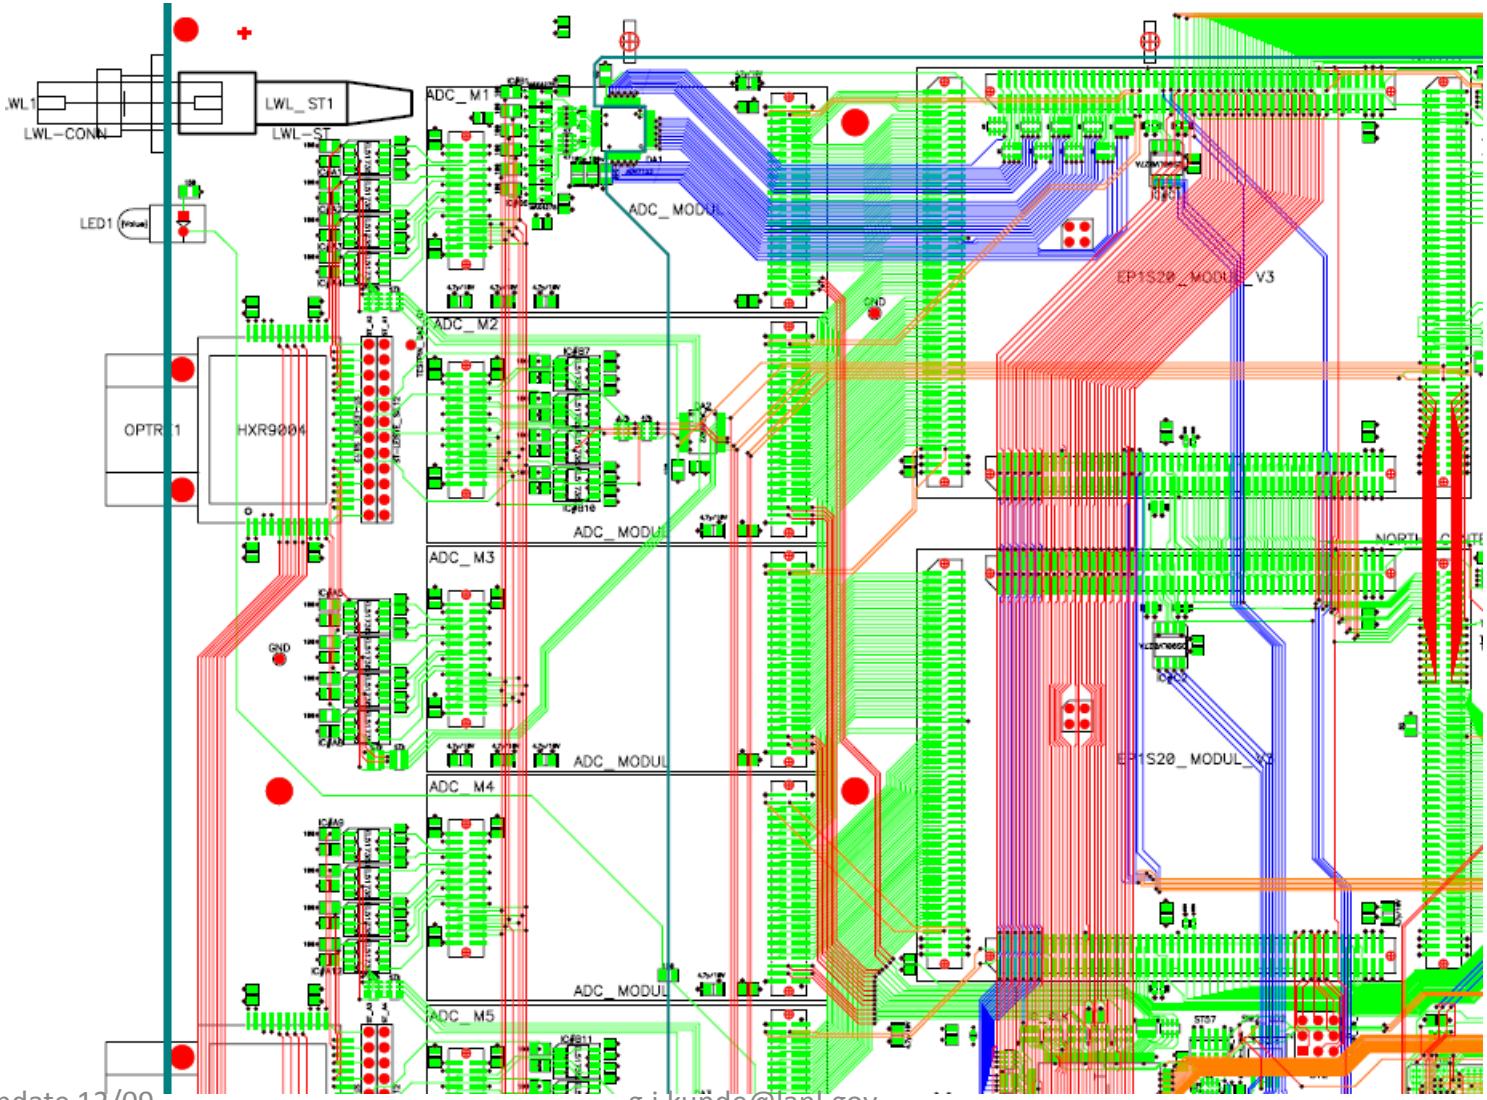

Where to plug in III

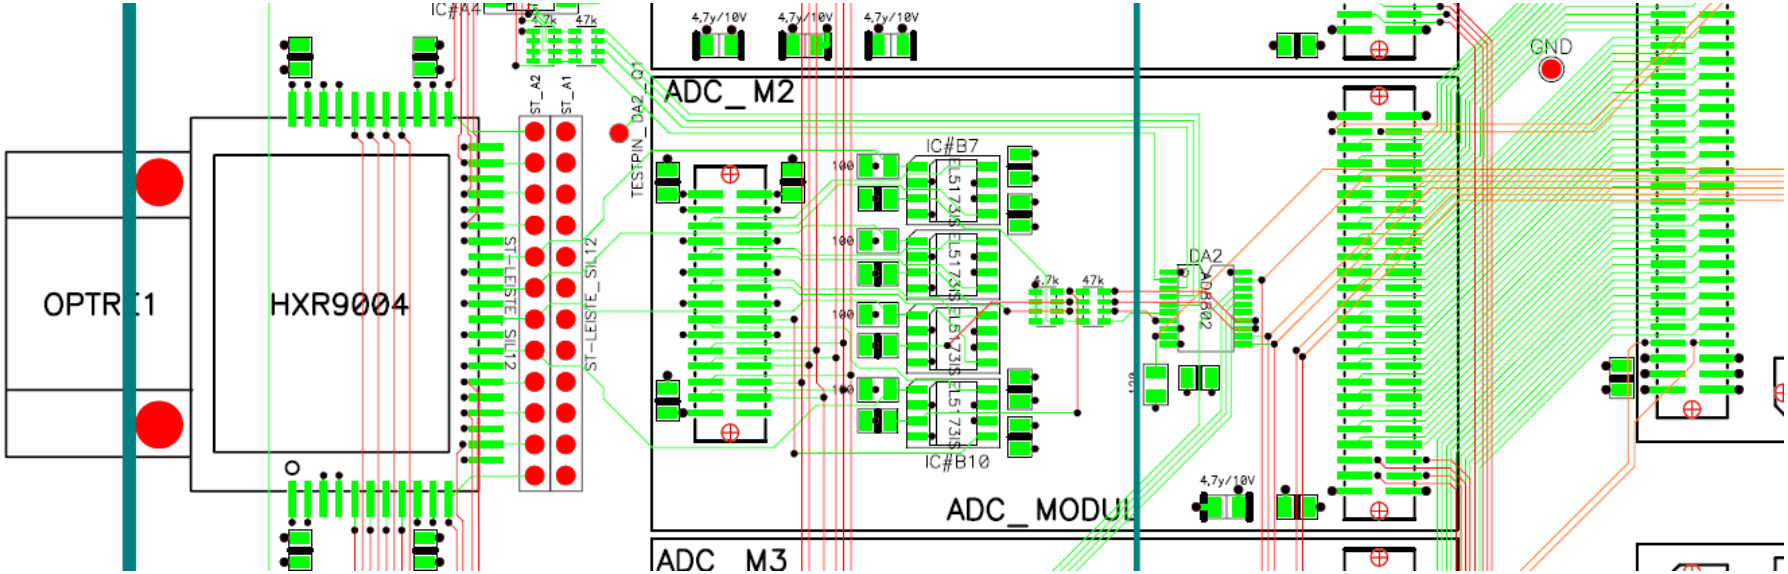

Acquired all the relevant schematics to build the card, including header type and pinouts.

What the ADC Daughter Board does

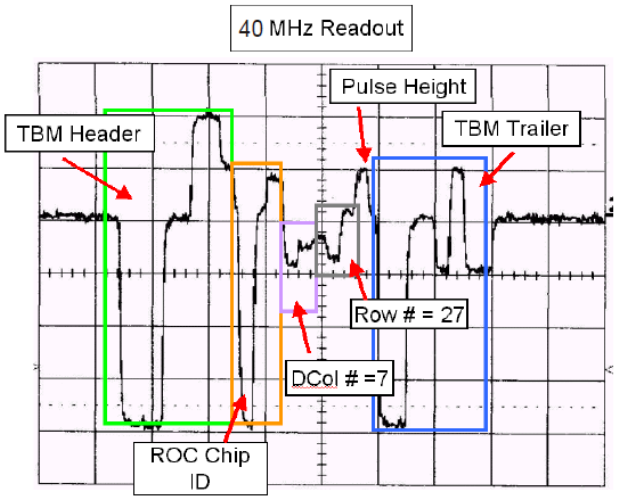

Figure 2: Single ROC hit with TBM header and trailer

IV. Pixel FED Data Flow

This chapter describes the functionality of the Pixel FED, following the main data flow from the optical inputs until the output to the DAQ system (Fig. 3). Finally, controls and monitoring are discussed.

A. Analog Part

Each Pixel FED has 36 optical input channels. A pin diode with a customized amplifier [7] converts the light signal into a single line current output. The working point and the frequency behavior of this receiver can be adjusted under VME control. In addition, its output signal offset can be shifted using a slow DAC in order to achieve an operating range suitable for the subsequent ADC. From this point on, the analog signals are fully differential to minimize noise and crosstalk.

2) ADCs

Each ADC daughter board (Fig. 6) has four channels and contains independent amplifiers and two commercial dual-channel ADCs (Analog Devices AD9218). All analog signals are fully differential, and each channel has its own test input which is added to the main signal.

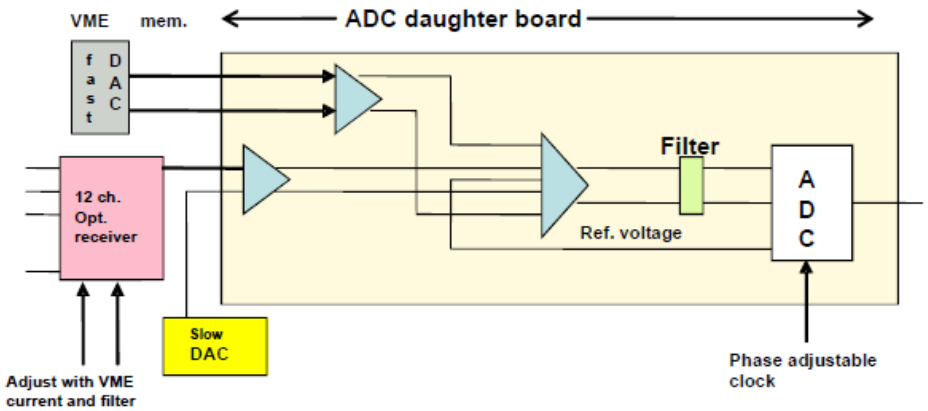

Figure 6: Block diagram of the ADC daughter board

As the incoming signal timing will be skewed between channels, each channel has its own, adjustable clock. In total, 16 clock phases with a spacing of approximately 1.6 ns are generated by the four Alteras in the front and collected by a fast Altera (Max family) located in the center. From there, 36 individual, serially terminated, clock lines deliver a selectable phase to each ADC channel.

Output of the ADC Board

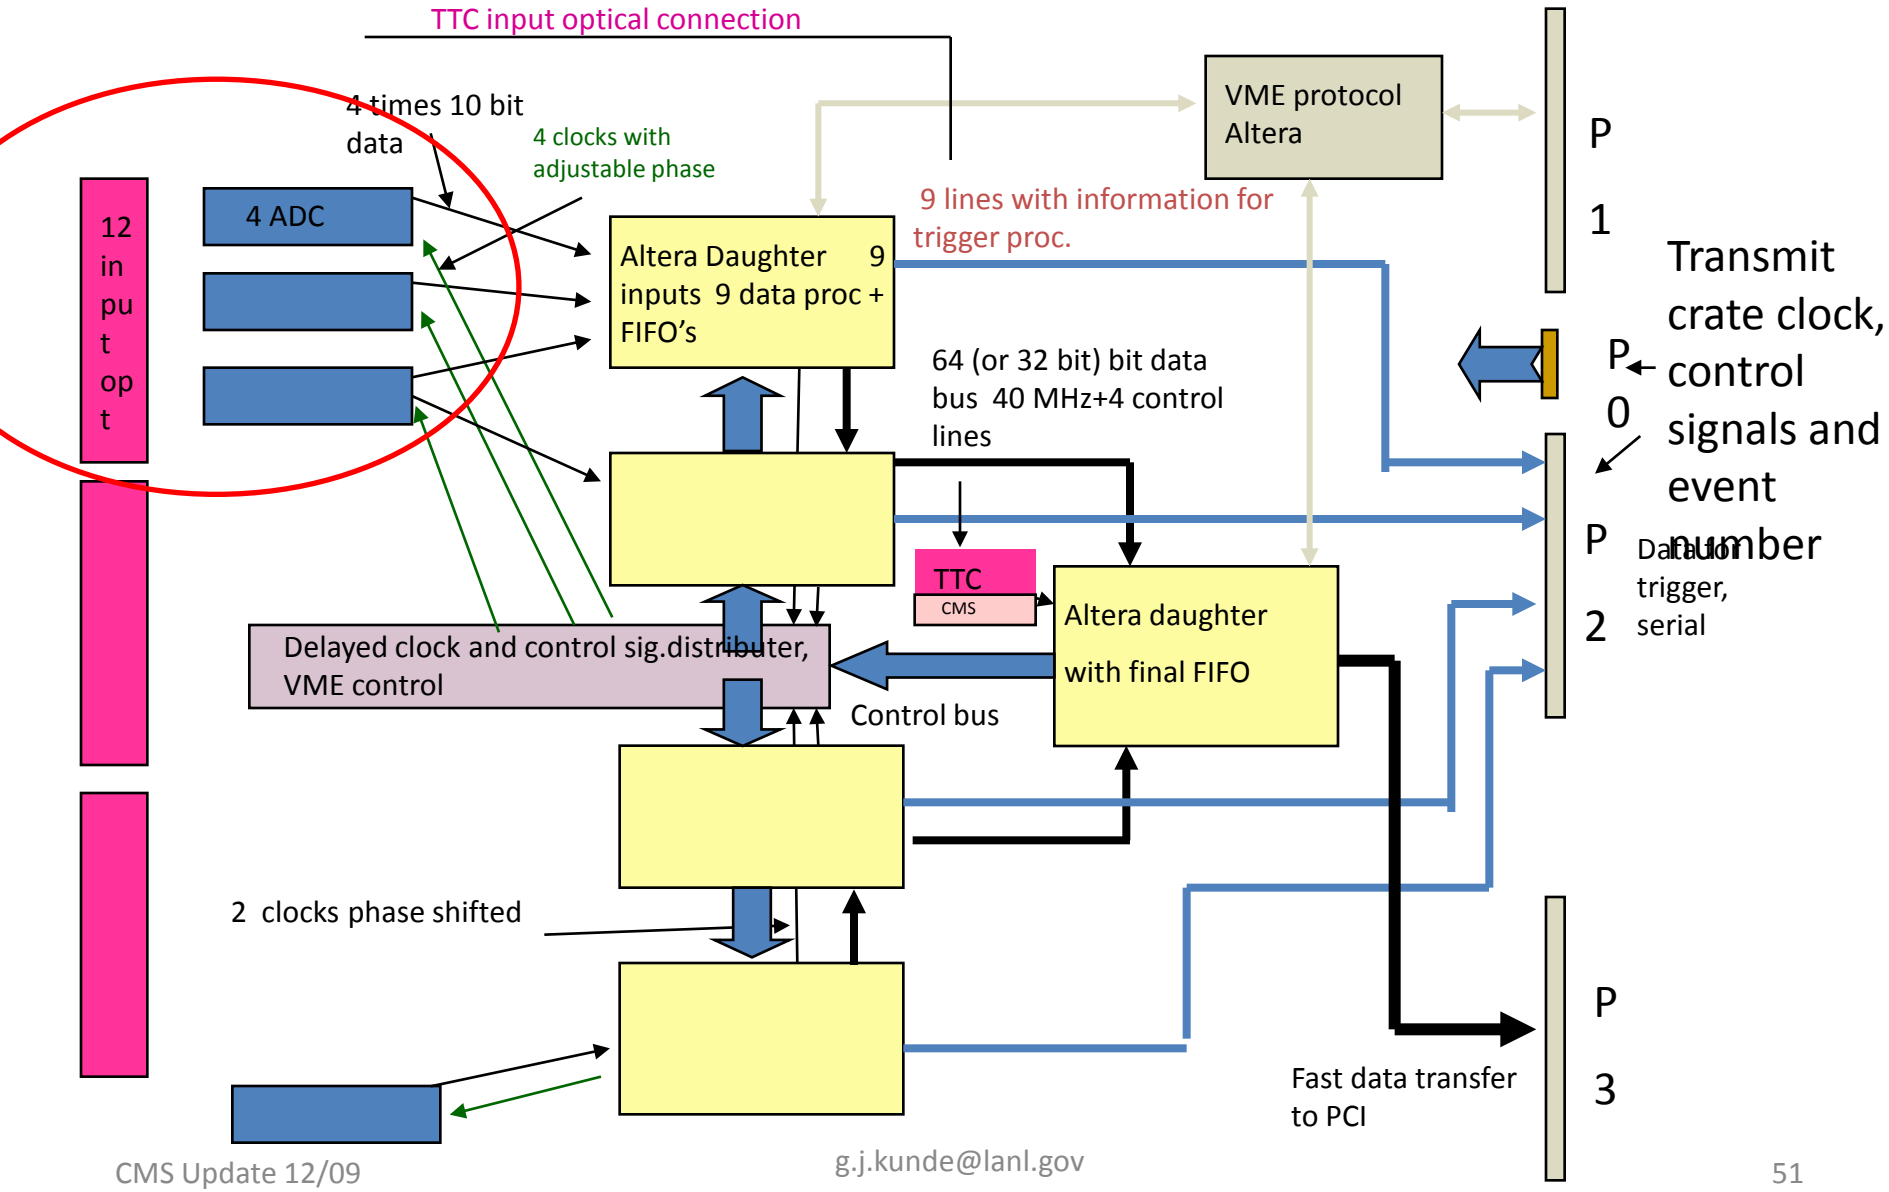

g.j.kunde@lanl.gov

The Daughter Board

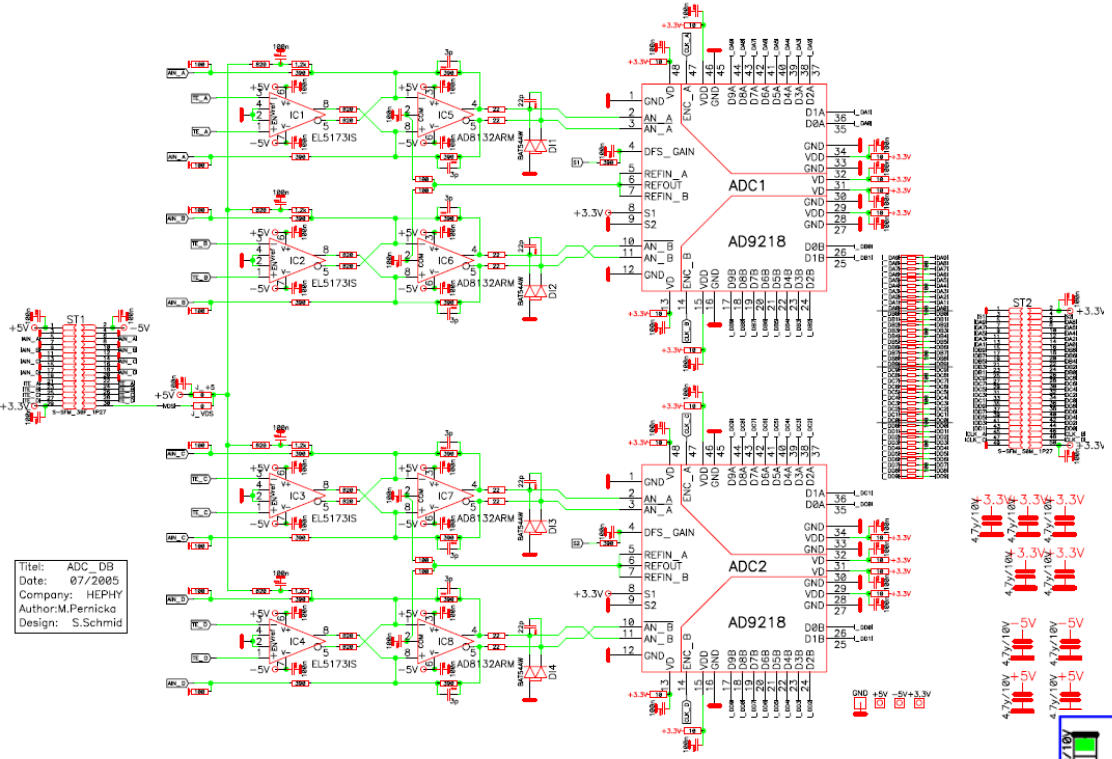

10 bits of ADC information with a clock, we need to simulate the optical levels to fake the fiber input.

There are nine cards with 4 adcs each

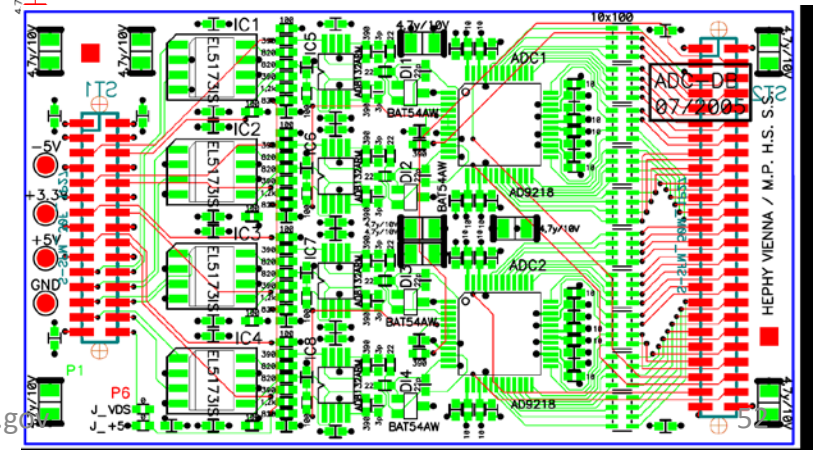

g.j.kunde@lanl.gov

The FPGA Firmware

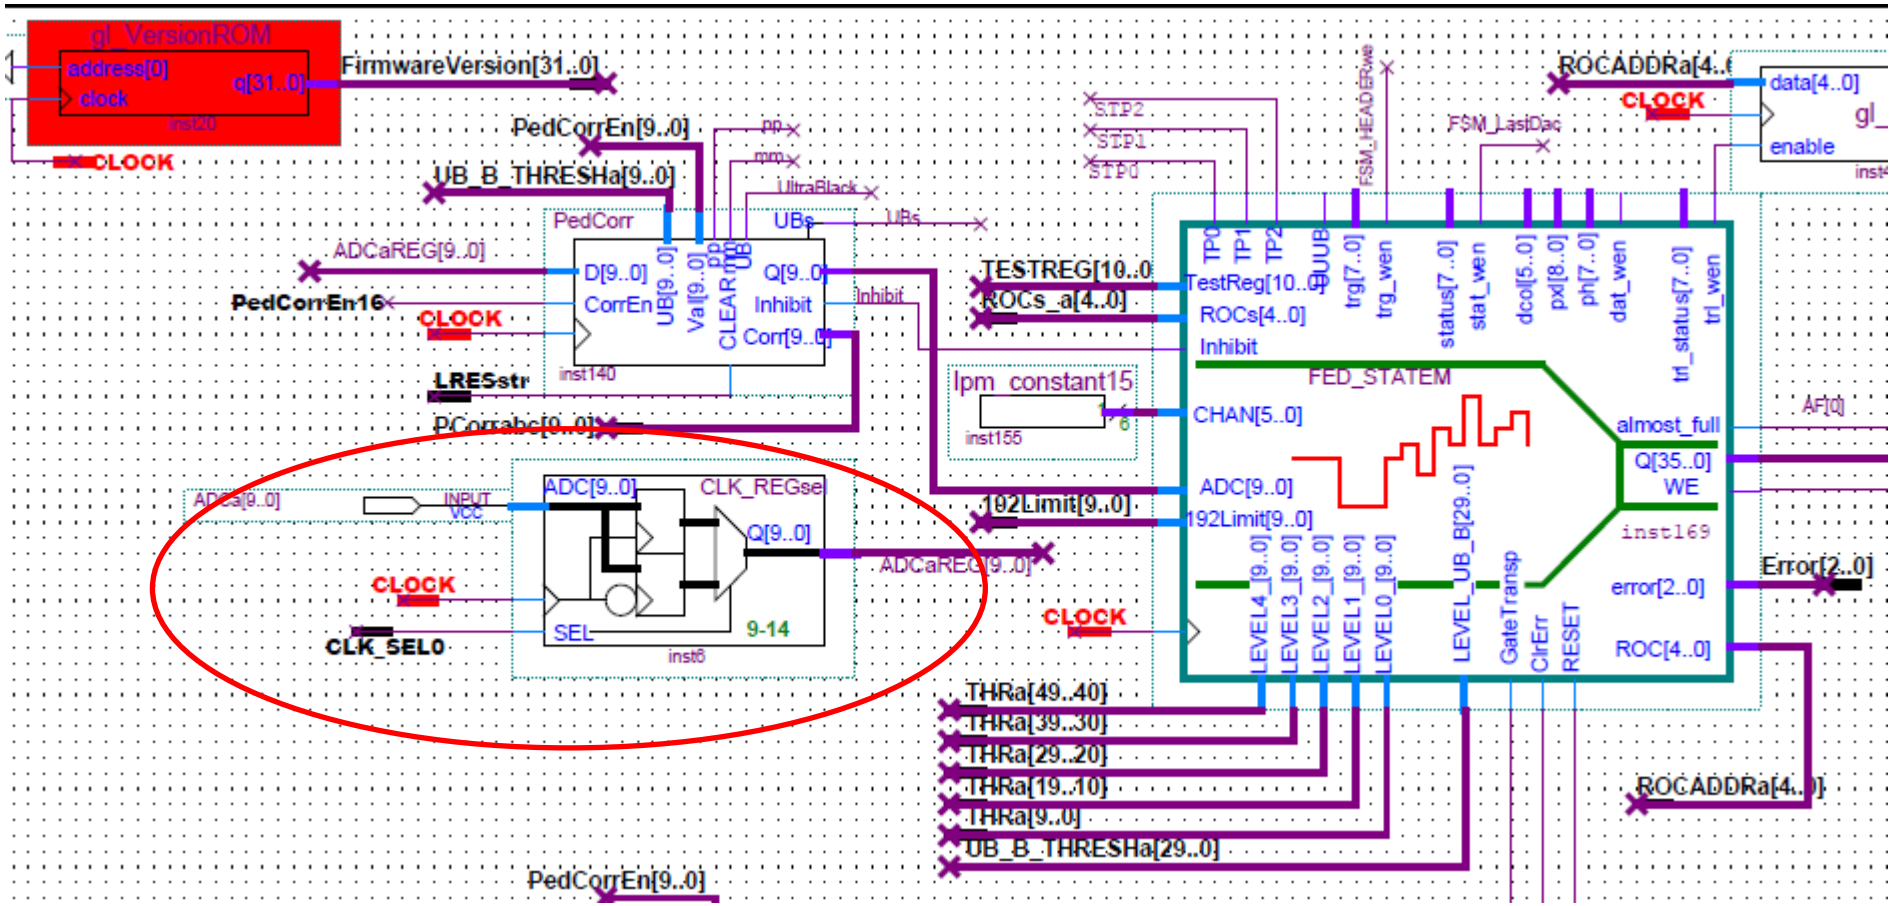

This is where we want to tie in, 10 bit ADC signals with a clock !

The Level Decoding

Setting thresholds for the 6 discrete levels :

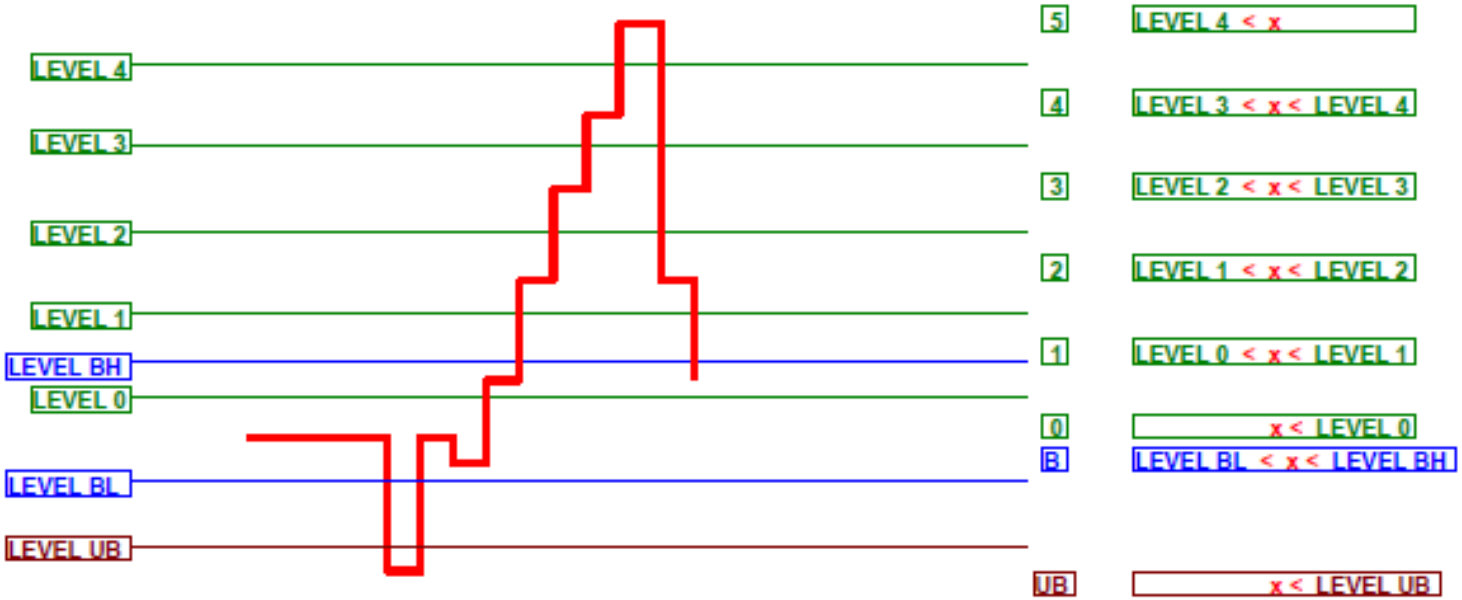

Discussion with the Vienna engineers on whether to drop the 2 LSB

The FED State Machine

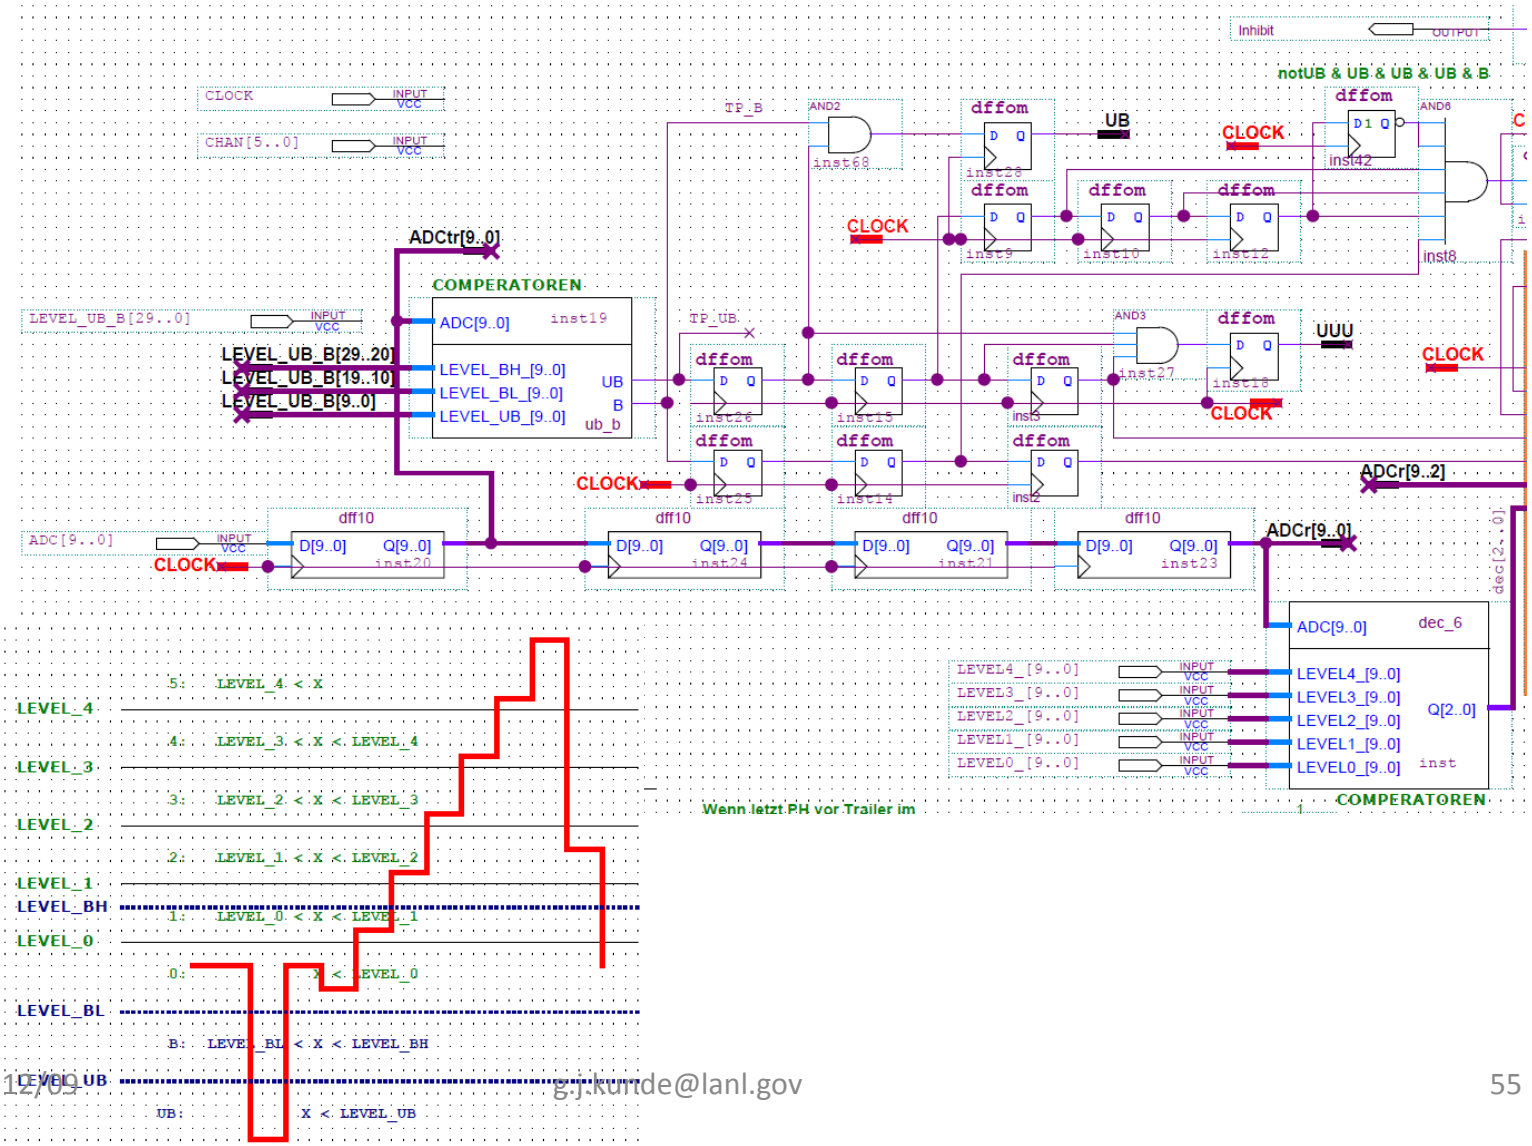

Daughter Board Output Format I

DECODING FSMs

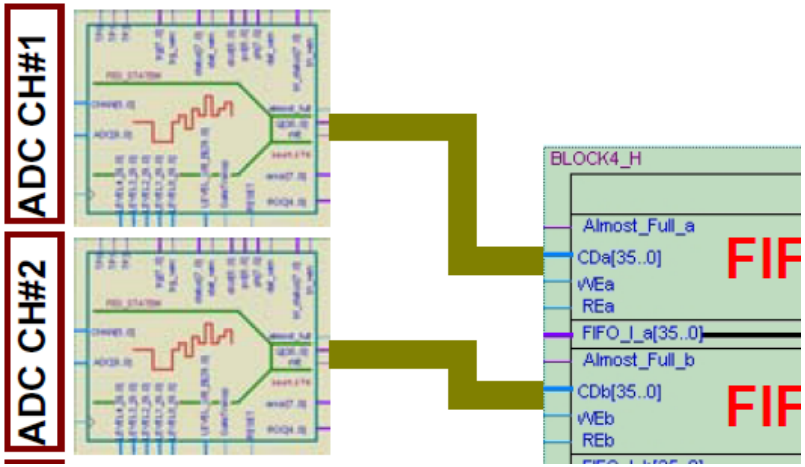

“ FRONT “ FPGAs Decoding FSM :

The encoding of the ADC Data is done by a State Machine written in VHDL. Necessa multipliers are implemented as ALTERA library blocks to ensure proper synthesize results.

STATE DIAGRAMM:

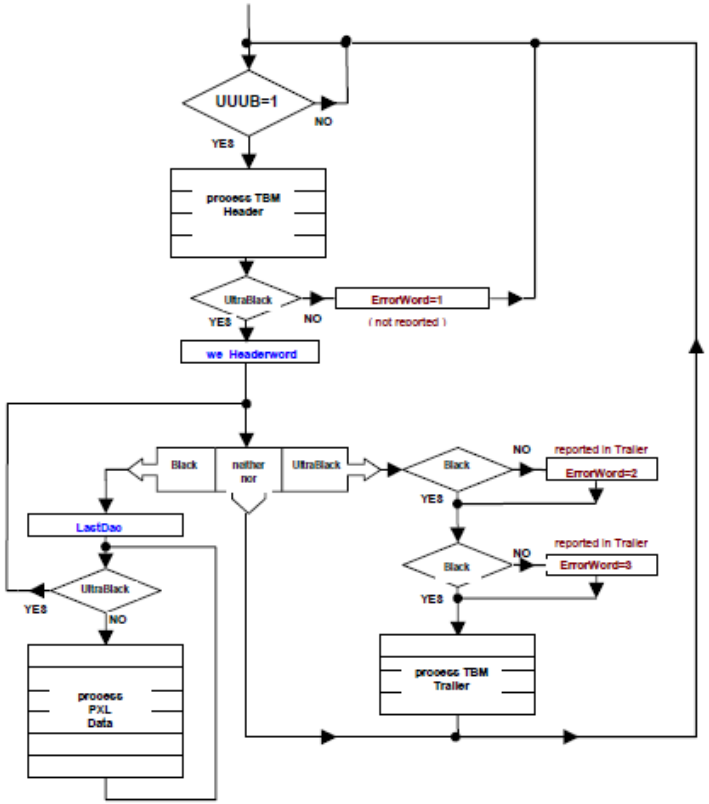

Output of the FSM

Decoding FSM output data format "normal mode":

```
BitPos: 35 34 33 32|31 30 29 28 27 26|25 24 23 22 21|20 19 18 17 16 15 14 13 12 11 10 09 08|07 06 05 04 03 02 01 00|

Header:  1  0  0  0| < -CHANNEL #-- >| 1  1  1  1  1| 0--------------------------------- 0|< --TBM HEADER TRG#-- >|
            0x8                            0x31
         HeaderID                      Header Marker


BitPos: 35 34 33 32|31 30 29 28 27 26|25 24 23 22 21|20 19 18 17 16 15 14 13 12 11 10 09 08|07 06 05 04 03 02 01 00|

LastDac: 1  1  0  0| < -CHANNEL #-- >| < --ROC #-- >| 0--------------------------------- 0|< ----- LAST DAC ---- >|
            0xc
         LastDacID


BitPos: 35 34 33 32|31 30 29 28 27 26|25 24 23 22 21|20 19 18 17 16|15 14 13 12 11 10 09 08|07 06 05 04 03 02 01 00|

Data:    0  0  0  1| < -CHANNEL #-- >| < --ROC #-- >| < --DCOL -- >|< ------- Pxl ------- >|< ----PulseHeight---- >|
            0x1
          DataID


BitPos: 35 34 33 32|31 30 29 28 27 26|25 24 23 22 21|20 19 18 17 16 15 14 13 12 11 10 09 08|07 06 05 04 03 02 01 00|

Trailer: 0  1  0  0| < -CHANNEL #-- >| 1  1  1  1  0| 0 --------------------- 0 e2 e1 e0  0|< --TBM Trailerword-- >|
            0x4                            0x30                                 e2          => invalid number of ROCs
         TrailerID                     Trailer Marker                              e1 e0  => FSM ErrorByte
```

What drives the Boards

- Xilinx ML605 evaluation package $ 2k
- Plugs into PC or stand alone
- Handles all the test driving
- Either via high speed serial or lvds
- FMC connectors

What plugs into the Xilinx

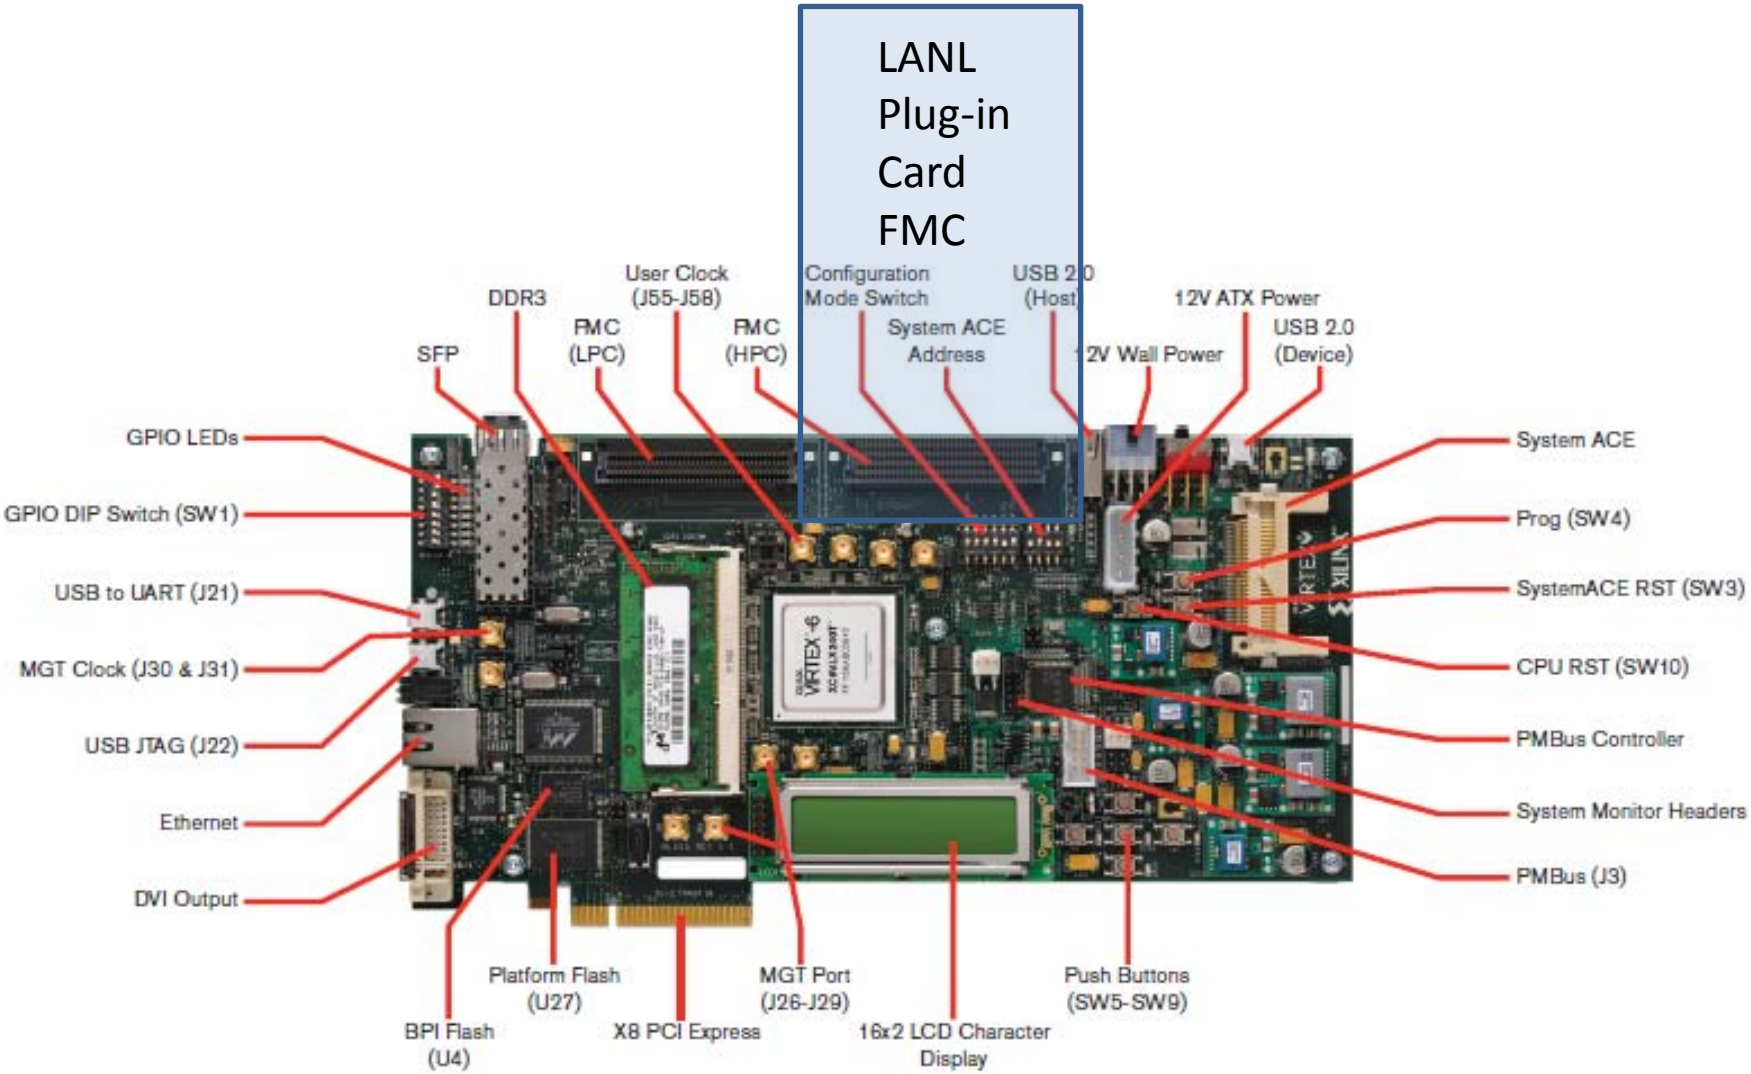

Three Solutions I (Matt)

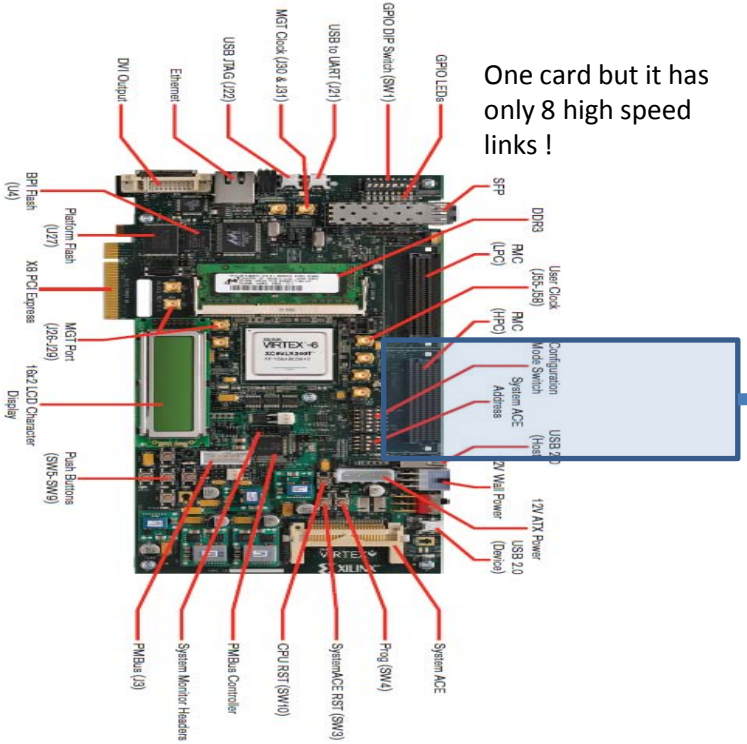

One card but it has only 8 high speed links !

copper

X 9

Spartan 6 T

Fiber would add a serdes each end

- Xilinx FMC via high speed serial to Spartan 6
  - Pros: Most flexible and elegant
  - Con: Need board with ball grid array ……

Three Solutions II (intermediate)

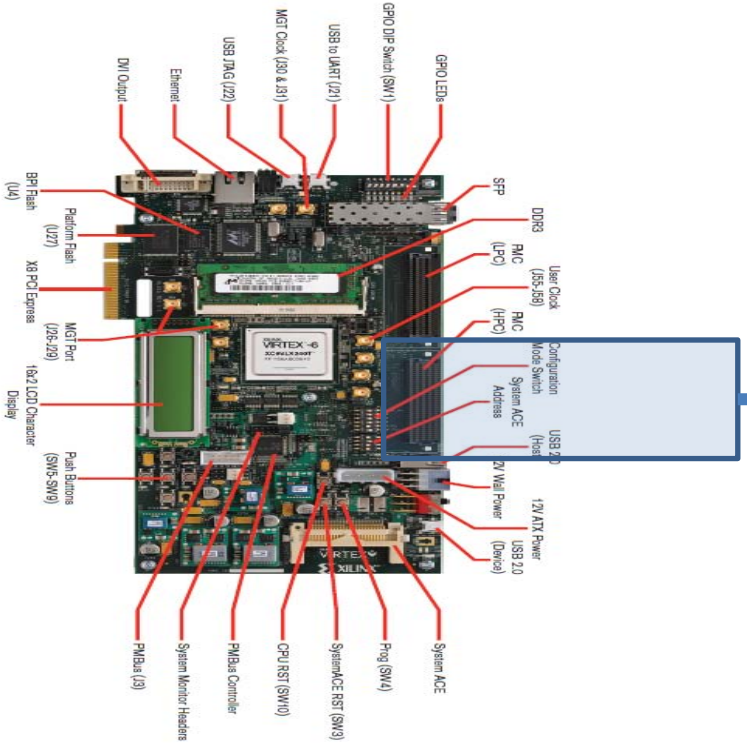

X 9

TLK1501 Serdes

Spartan 3AN

- Xilinx FMC via high speed serial to Serdes and Spartan 3 AN
  - Pros: QFP packages, can do the layout with express PC, eprom on FPGA
  - Con: less flexible ….

Why we need the Spartan 3

- Need to recover the signal for the right clock which comes from the FED board
- Plus allows for clears in synch loss etc ……

Three Solutions III (proposed)

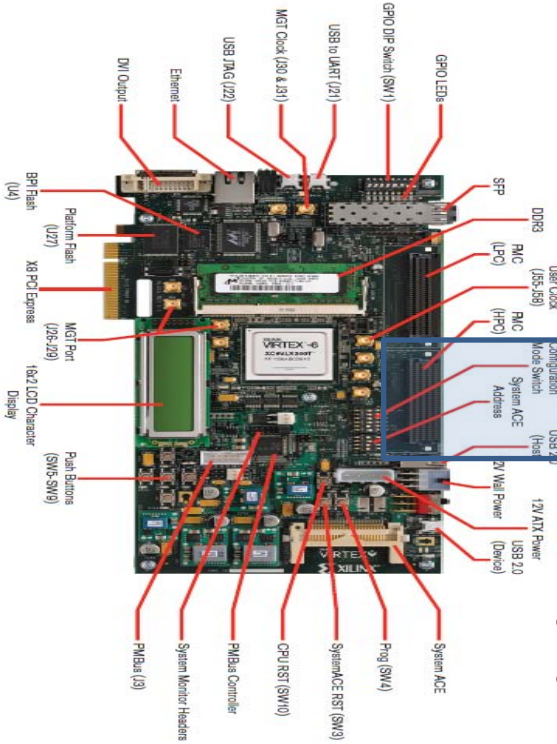

8 bit parallel LVDS
with clock on
Samtec cable

X 8

Spartan 3AN

FED

8 bit at 200 MHz gives us 40 MHz 40 bits = 4 * 10 bits
8 bit at 160 MHz gives us 40 MHz 32 bits = 4 * 8 bits

- Xilinx FMC via 8 bit parallel+clk and Spartan 3 AN
  - Pros: QPF packages, can do the layout with express PC
  - Straight forward ….
  - Con: ‘need custom protocol’ but that should be simple ……

The Spartan 3-AN

Package	TQ144 TQG144	
Body Size (mm)[4]	20 x 20	
Device	User	Diff
XC3S50AN	108 *(7)*	50 *(24)*

Device	System Gates	Equivalent Logic Cells	CLBs	Slices	Distributed RAM Bits[1]	Block RAM Bits[1]	Dedicated Multipliers	DCMs	Maximum User I/O	Maximum Differential I/O Pairs	Bitstream Size [1]	In-System Flash Bits
XC3S50AN	50K	1,584	176	704	11K	54K	3	2	108	50	427K	1M

The Board

33.7 mm

20.9 mm

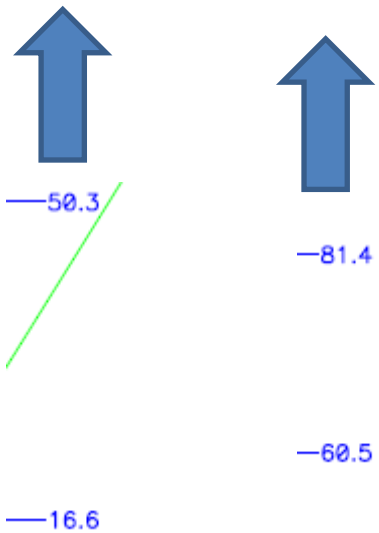

Single ended to FPGA on FED

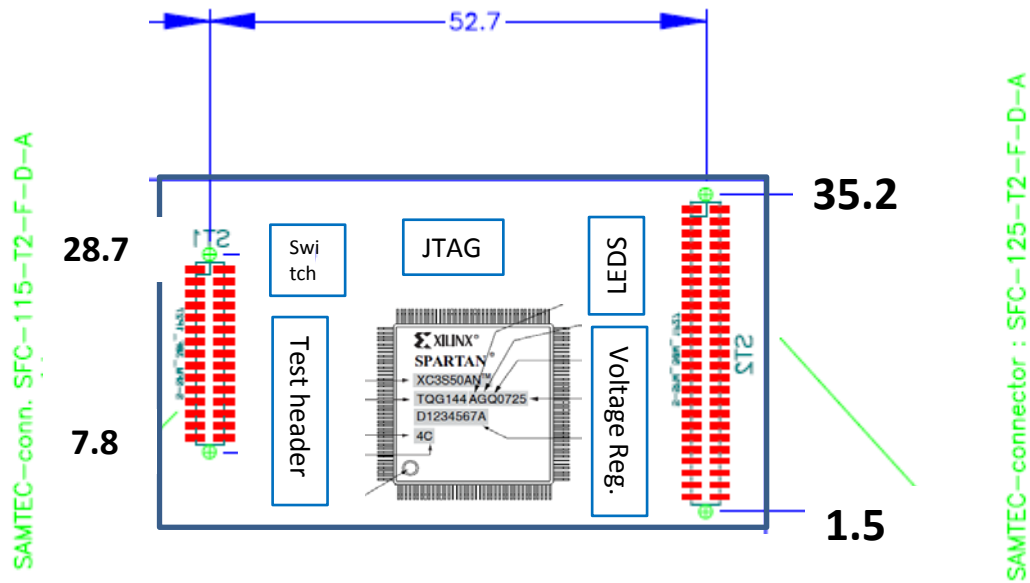

We have -5V, + 5 V and 3.3 V on the connector, need power for the FPGA ….

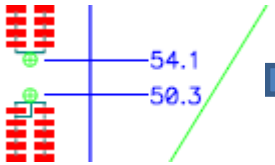

Connector to connector is 3.8 mm, i.e. <1.9 to edge , 1.5 chosen which give 0.8 mm gap between boards

Needs

- Xilinx ML605
- FPAG software
  - LANL expertise
- Event and time structure files
  - Exist
- Receiver board that is pin compatible with FED front end receiver
  - LANL development
  - Need 9 boards total
  - See details on the previous slides

Present Working Plan for Holdoff

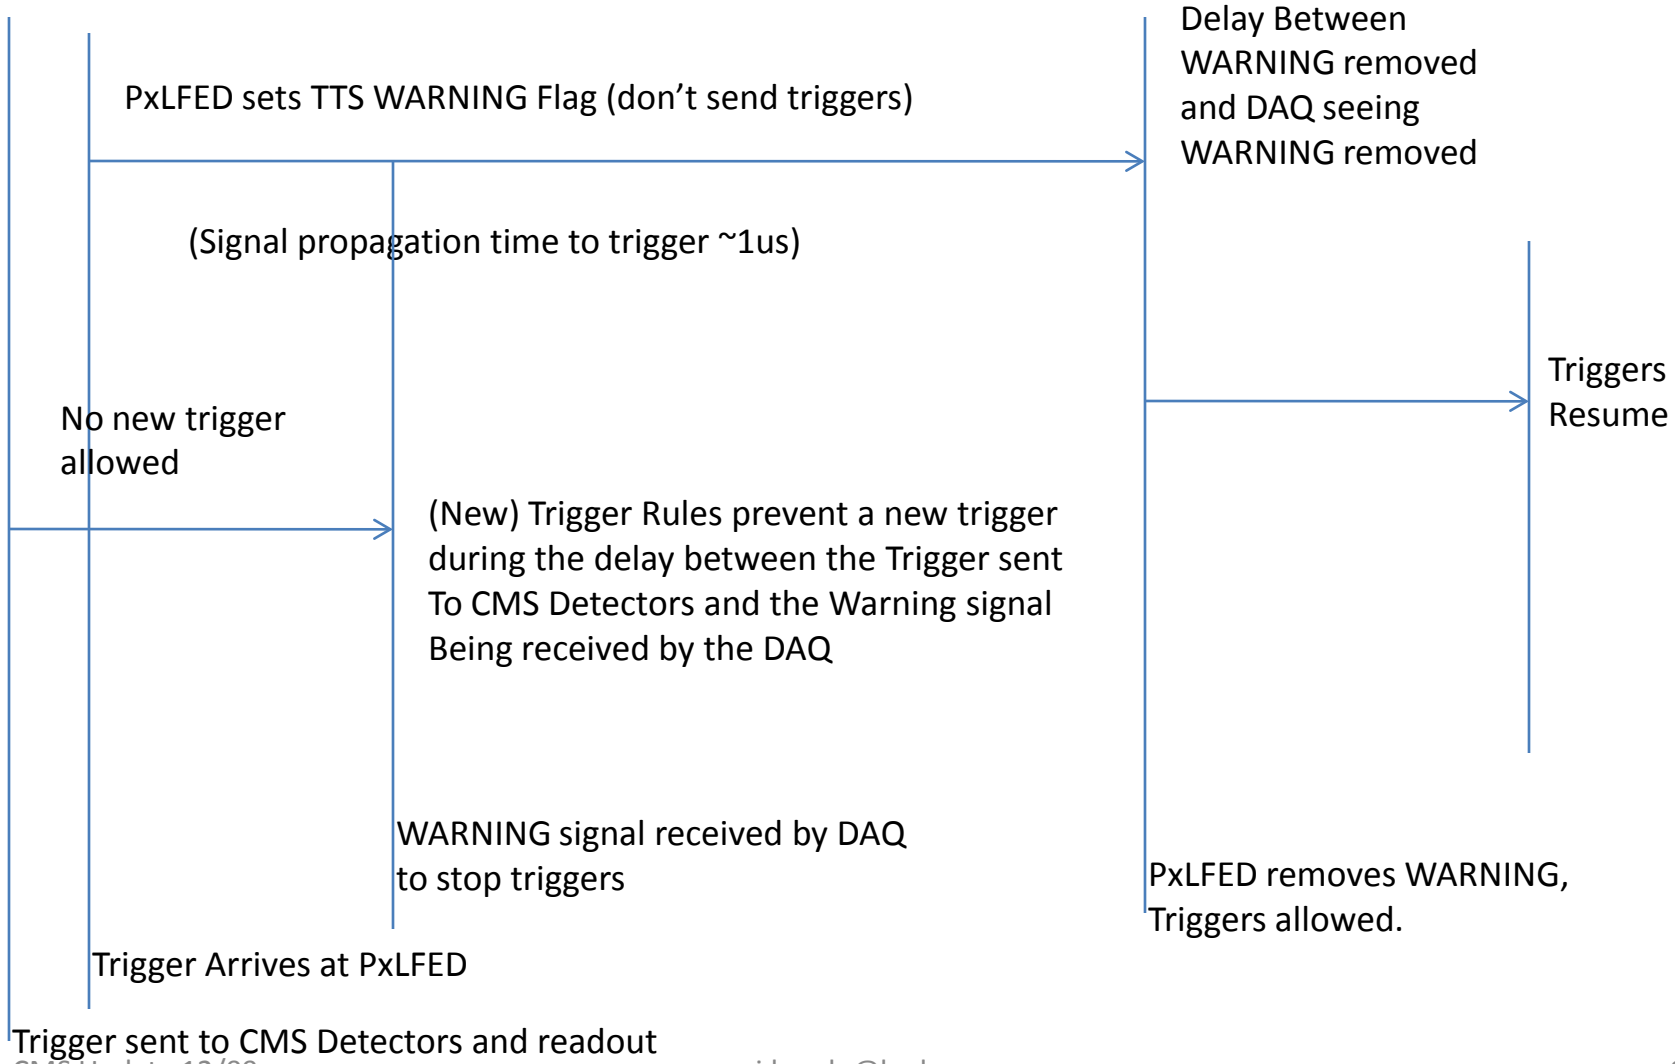

Conclusions

- Pixel FED has problems in HI
- And EVEN in pp commissioning already !
- Detector readiness workshop in April
- Realistic simulations of
  - Current firmware and limits
  - Proposed changed firmware
- Develop working idea of a dead time implementation
- Need testing hardware to verify any solution before there is HI beam ……
- Propose LANL test system
- Next step is to propose a new FPGA mezzanine board …