



Science & Technology
Facilities Council



ATLAS experience with FPGA (TDAQ's perspective)

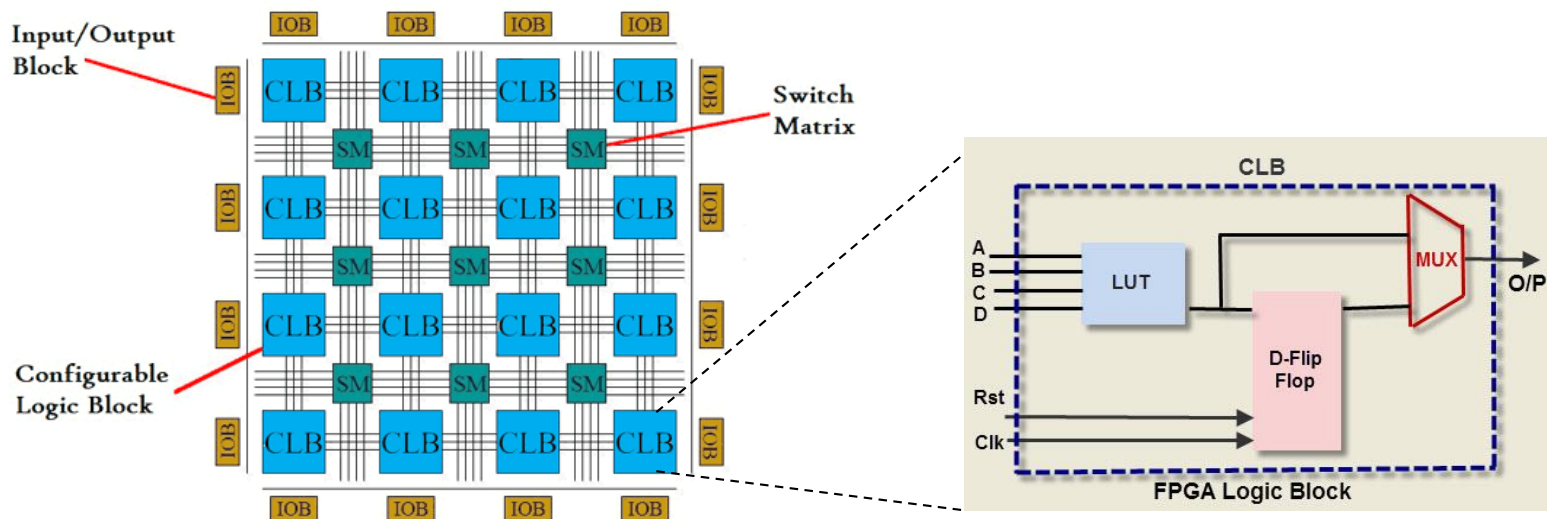
Weiming Qian
Rutherford Appleton Laboratory

Outline

- FPGA introduction
- ATLAS TDAQ history with FPGA
- Modern large FPGA design consideration
 - Power
 - Cooling
 - Clocking
 - High-speed PCB design
 - Test methodology
 - Control/configuration
 - Monitoring
 - F/W management and collaboration
 - Hardware design platform for fast prototyping
- FPGA future trend

What is FPGA

- Field Programmable Gate Array (FPGA)
 - a matrix of configurable logic blocks (CLB) connected via programmable interconnects
 - Can be reprogrammed to desired functionality requirements after manufacturing
- Application Specific Integrated Circuit (ASIC)
 - custom manufactured for specific design tasks
 - Can not change function after manufacturing



Architecture of early FPGA

Modern Large FPGA

- Modern large FPGAs integrate many other features in addition to unprecedented logic density
 - Embedded processors
 - DSP clocks
 - High-speed transceivers
 - PCIe blocks
 - EMAC blocks
 - Advance clock manager
 - High bandwidth Memory
- Modern large FPGAs have become System-On-Chip design platforms




LHC timeline




- LHC first started serious data taking in 2011
 - Run 1, running at 75% nominal luminosity
- LHC Run 2 finished at the end of 2018
 - 2 times nominal luminosity, shut down at the end of last year for upgrade for Run 3
- Long shutdown foreseen in 2024 for upgrade to HL-LHC
 - 10 times nominal luminosity

ATLAS



- ATLAS one of two general purpose detectors at the LHC
 - different subsystems wrapped concentrically in layers around collision point to record trajectory, momentum, and energy of particles

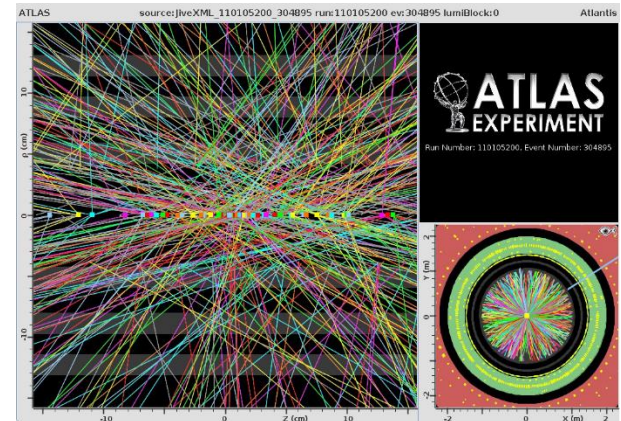

ATLAS Calorimeters



- Inner layers of fine-grained LAr electromagnetic calorimeters
 - ~200 thousand cells to measure energy of electrons and photons
- Outer layer of hadronic calorimeter
 - ~10 thousand cells to sample energy of hadrons as they interact with atomic nuclei
- Different particles leave different signatures in each layer

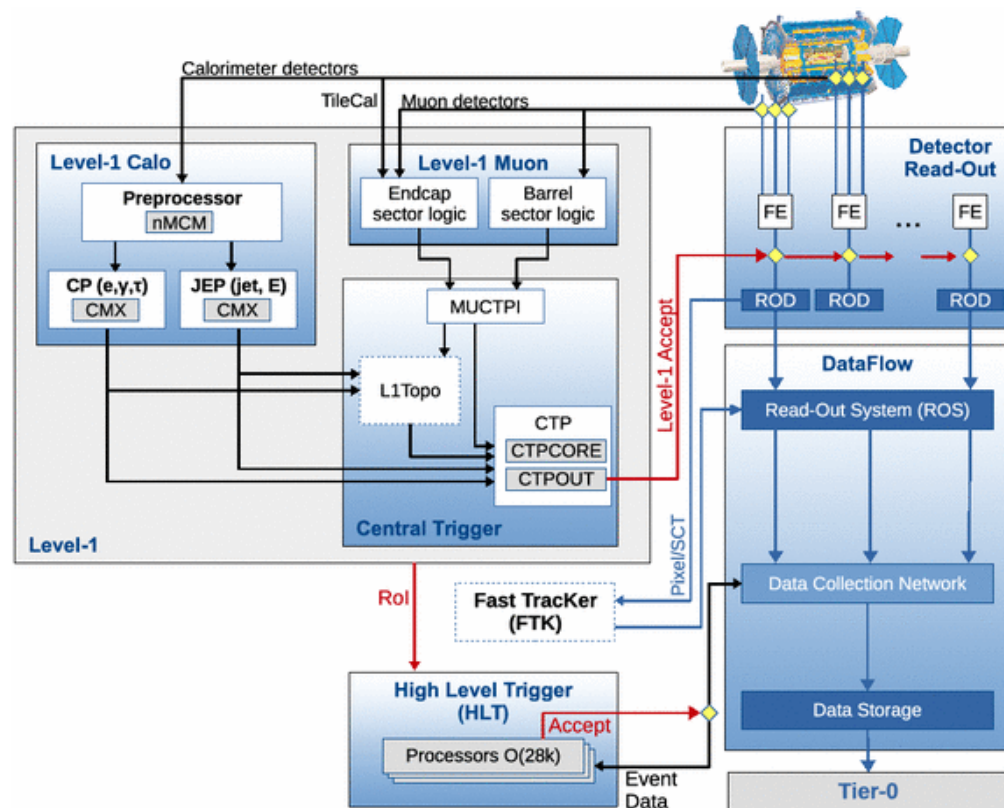
Triggering at ATLAS

- LHC collides bunched beams of protons @40MHz
 - Every time the bunches cross some protons collide
 - Real example from 2016 with 10 collisions
 - *At Run2, LHC delivered around 1 billion proton-proton collisions per second*
- From 2026 we expect an average of 200 collisions each time bunches cross
 - This a simulated example with 200 collision
 - *up to 10 billion collisions per second*
 - Interesting collisions at much much lower rate
 - 1 Higgs boson per 10 billion collisions
- We want just the interesting collisions





Triggering at ATLAS

- ATLAS trigger system
 - Level-1
 - custom hardware
 - 40 million decisions per second
 - up to $2 \mu\text{s}$ for each decision
 - $2.5 \mu\text{s}$ buffer on detector
 - High Level Trigger
 - CPU farm
 - 100 thousand decisions per second
 - 1 thousand events to storage per second



Initial prototypes for Calorimeter trigger

- Mid-1990's
- Calorimeter trigger processes 6000 coarse grain 'trigger towers' (summed cells)
 - Initial design based on multichip modules and ASICs
- Initial prototyping in FPGA
 - Xilinx XC4008 devices
 - 0.35 μm CMOS
 - 8000 gates
 - Each FPGA here 4 channels of Bunch Crossing ID
 - "Whilst FPGAs provide an ideal solution for the demonstrator system, they are too large and costly to be used in the final system"





Actual calorimeter trigger for Runs 1 and 2

- Mid-2000's
 - FPGAs displaced almost all ASICs
 - 9U VME production modules using Xilinx FPGA
 - Virtex-E, Virtex-II, Virtex-II Pro
 - 0.13 μm CMOS
 - up to 30000 logic cells
 - Up to 8 high-speed transceiver pre FPGA
 - each FPGA covers half entire upper prototype on previous slide


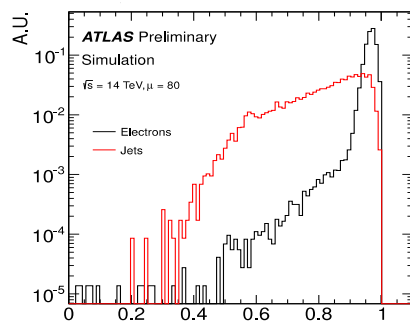


Upgrades for Run 3

- 2021
- *Need to improve sensitivity to electroweak physics in face of increased pileup*
 - Maintain thresholds close to original LHC by increasing data into trigger by factor 10
- LAr SuperCells with finer azimuthal and depth segmentation






- Allows lateral R_η and depth f_3 shower shape discriminants between electrons and jets



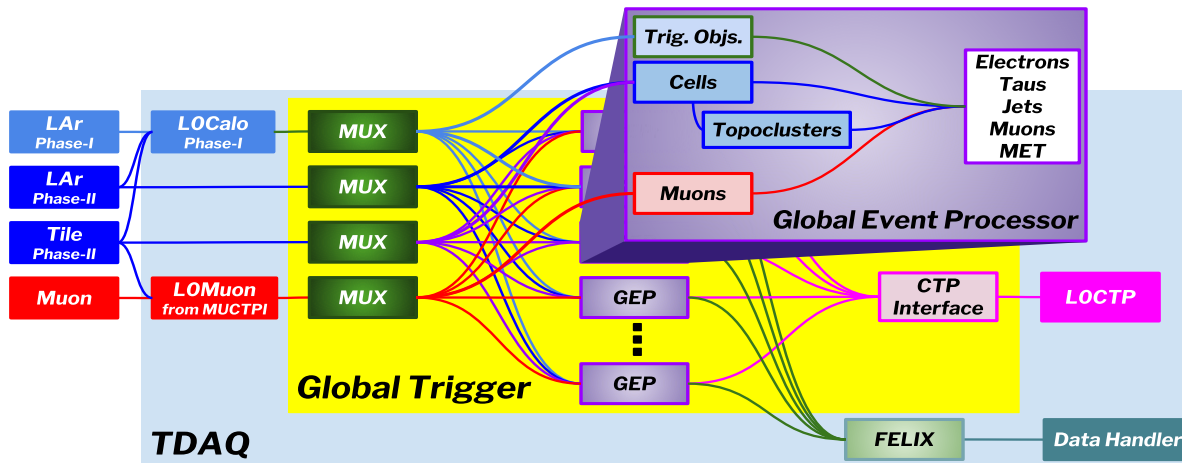
L1Calo trigger modules for Run 3

- Factor 10 increase in data rate to 30 Tb/s requires many very high-speed links
- ATCA-based boards with up to 240 inputs each running at 11.2 Gb/s
 - Altera Arria 10, Xilinx Virtex 7 and UltraScale+
 - 22nm-14nm FinFET
 - Up to 2.5 million system logic cells
 - up to 120 high-speed serial transceivers per FPGA



Upgrades for Run 4


- 2026
- Ultimate LHC luminosity requires even finer sensitivity
 - Addressed by enabling use of individual calorimeter cells above threshold in trigger
- Longer latency for hardware trigger allows time-multiplexed event building
 - Complete data for a full event on a single FPGA



Global Event Processor for RUN4

- Concept
 - ATCA board with ~200 link @25Gbps
 - First prototype with current volume production FPGAs

Next generation FPGAs




Note: What defines a process "generation" and the start of "volume" production varies from company to company, and may be influenced by marketing embellishments, so these points of transition should be used only as very general guidelines.

Sources: Companies, conference reports, IC Insights

Modern large FPGA design - Power

- Up to 350W per board for Run 3
 - Very large current for FPGA core voltage rail and high-speed transceiver voltage rail
 - Up to 100A per power rail for a board with 4 large FPGAs
 - Very stringent power noise requirement
 - $\pm 3\%$ pk-pk for low voltage rails




Power Integrity simulation very important



eFEX preproduction VMGTA VCC Power Plane (1 oz.)


Modern large FPGA design - Cooling

- Cooling capacity of ATLAS counting room infrastructure
 - 350W (front board) + 50W (RTM)
 - 80% took away by water cooling
 - 20% leakage into environment
 - Took away by aircon
 - A lot of thermal simulation and measurement and optimization






Modern large FPGA design - Clocking

- High-speed transceivers need very low jitter reference clock
- Jitter is more complicated than most thought
 - Evaluation in frequency domain




Modern large FPGA design - PCB

- Low Dk and low Df Material
- PCB stackup
 - Microvia/blindvia vs backdrill
 - Complexity vs yield
- High-speed signal simulation
- PCB measurement



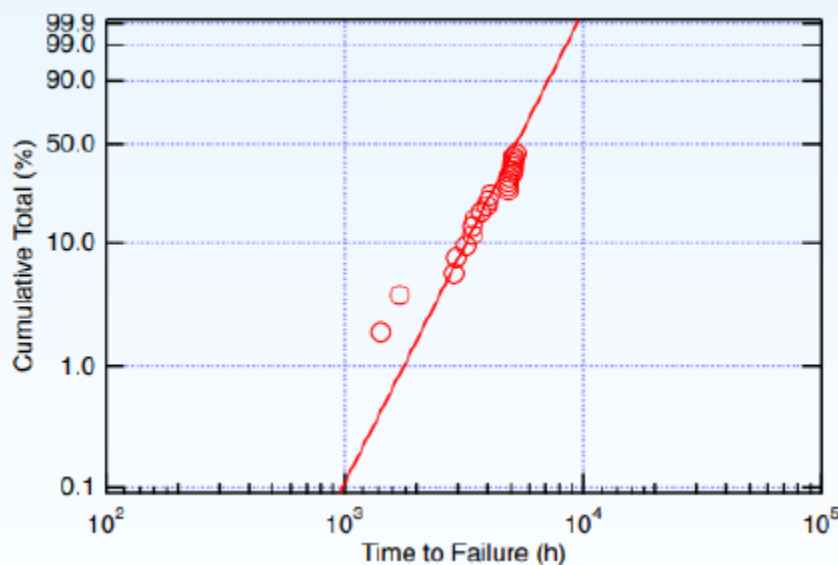
Modern large FPGA design - Test

- Test all the parameter corners
- Accelerated aging test
 - Capture design/manufacture error before production



Wrong voltage connected to XADC!

*Die backside laser imaging of
EOS-induced diffusion damage*



CMS FC7 AMC

Modern large FPGA design - control

- Control via Ethernet in ATCA shelf
 - IPBus
 - Simple
 - Hardware isolated from software
 - Zynq SoC + Linux
 - Flexible
 - Long term support issue
 - IT security
- FPGA configuration
 - ATLAS counting room will be a micro radiation zone
 - Access will be restricted
 - F/w upgrade remotely
 - Needs fail-safe mechanism
 - Reliable golden image keep the board always accessible

Modern large FPGA design - monitoring


- Trigger board for RUN3 cost \$40,000 ~ £100,000 each
- Accident can happen and did happen
- Monitoring and auto-protection is very important
 - Safety infrastructure should be installed on day one
 - However, it is often retro-fitted after accident happens

```
[atlun01] ~/ipmc % ipmitool -I lan -H shelf1.pp.rl.ac.uk -A NONE -t 0x8a sensor
```

Hot Swap	0x0	discrete	0x1080	na	na	na	na	na	na
IPMB Physical	0x88	discrete	0x0880	na	na	na	na	na	na
Version change	0x0	discrete	0x0080	na	na	na	na	na	na
Internal temp.	29.000	degrees C	ok	na	na	na	40.000	60.000	80.000
LM82 internal	28.000	degrees C	ok	na	na	na	40.000	60.000	80.000
LM82 FPGA temp	49.000	degrees C	ok	na	na	na	75.000	80.000	80.000
QBDW033 Vinut	48.250	Volts	ok	na	40.000	na	na	55.000	na
QBDW033 Voutut	11.640	Volts	ok	na	1.112	na	na	16.152	na
QBDW033 Ioutut	7.568	Amps	ok	na	na	na	na	9.998	na
QBDW033 temp	40.485	degrees C	ok	na	na	na	53.145	50.191	53.145
MDT040 Vinut	11.983	Volts	ok	na	5.039	na	na	14.029	na
MDT040 Ioutut	1.892	Amps	ok	na	na	na	na	2.500	na
MDT041 Vinut	12.045	Volts	ok	na	4.543	na	na	14.029	na
MDT041 Ioutut	1.902	Amps	ok	na	na	na	na	2.500	na

Modern large FPGA design - firmware


- F/w repository
 - CERN gitlab
- F/w build system
 - [IPBB](#)
 - Used by CMS
 - [HDL make](#)
 - Used by ATLAS LAr
 - [HOG](#)



- Used by ATLAS L1Calo
- A set of TCL (Tool Command Language) scripts manage firmware repository
- A Gitlab Continuous Integration script automatically synthesizes and implements HDL projects when a Git Merge Request is opened

Modern large FPGA design – fast prototyping

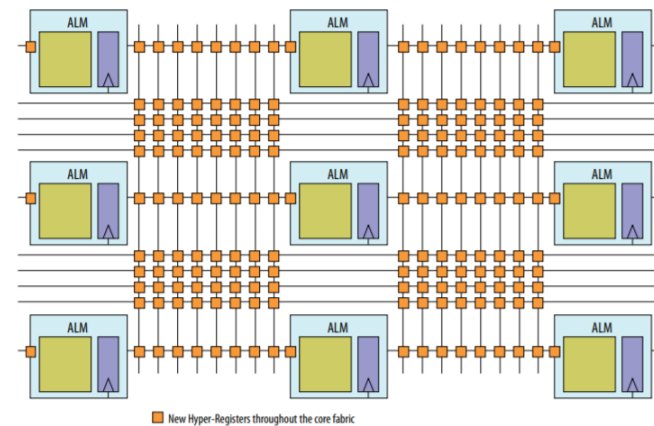
- Fast prototyping platforms with ATCA
 - [Serenity](#) consortium lead by Imperial college
 - Carrier card
 - Daughter cards
 - [Apollo](#) consortium lead by Boston University
 - Command module
 - Service module




Future FPGA

- Intel Agilex
 - Architecture closer to traditional FPGA

Intel Hyperflex Core Architecture



- Xilinx Versal
 - Prime Series
 - Upgrade to Zynq SoC
 - AI Core Series
 - Automatic search for new physics?
 - Needs R&D



Summary

- FPGAs was originally used as the prototyping for ASIC in ATLAS TDAQ.
- At Run1, FPGAs displaced most of ASICs in the original technical design.
- At Run3, FPGAs with high density high speed links dominated trigger upgrade design.
- At Run4, increased L1 trigger latency (due to detector frontend upgrade) allows a step change in trigger architecture.
 - Next generation FPGAs will run iterative algorithms in real-time system.
 - Even higher speed links (25G or 50G) allows data aggregation for full event process.
- ATLAS time scale much longer than industrial norm
 - Typically running electronics for some 20 years without large-scale upgrade
 - Original L1Calo was base on 9U VME designed in early 2000, and part of it will continue until the end of RUN3 in 2024.
- FPGA technology has been advancing very fast, and designing with modern large FPGAs is challenging.