WP9: DAQ

THE ROOKIES

Tommaso Colombo on behalf of the WP9ers

EP R&D day
CERN, 22 May 2024

What? Who?

DAQ for HEP experiments is

- The interface between detector and rest of the world
- A varied domain (from readout FPGAs to shifter interfaces)
- Largely based on commercial components (a relatively new development, introduced mainly with LHC)
- Large-scale: significant economic and environmental impact

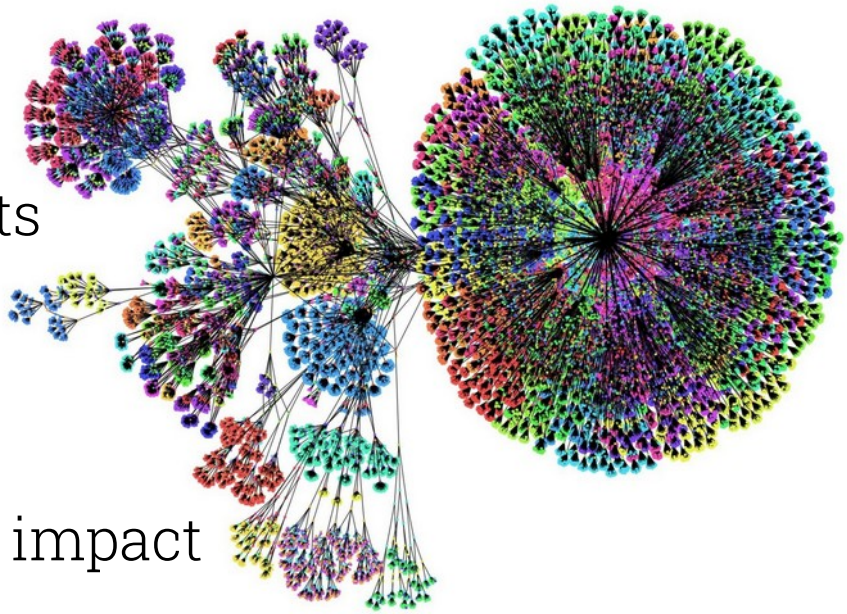

What? Who?

Goal: “shape HEP DAQ systems of the future”

- Transform commercial products into DAQ systems building blocks
  - high-performance, low-cost, energy-efficient
- Establish blueprints that can be used by future experiments instead of bespoke systems
- Pool the knowledge of DAQ experts from DT and all LHC experiments
  - 6 activities
  - 38 people in the “group”

WP9.1: Artificial intelligence in experiments

Activity leader:
G. Avolio (ADT)

- Despite current hype, not a new thing in HEP:
  - In analysis: neural networks, deep learning algorithms
  - In DAQ: rule-based expert systems and complex event processing
    - configured by experts
    - powerful and reliable, but not adaptive
- Goal: apply latest AI tech to experiment control systems
  - Increase automation, improve operational efficiency
  - Enhance accuracy: learn from historical data
  - Detect anomalies and changing conditions in real time
  - Predict trends and potential issues in advance

WP9.1: Artificial intelligence in experiments

- 9.1.1: Exploitation of Generative AI
  - Build a LLM-based chatbot to assist operators sifting through documentation (often scattered and poorly indexed)
  - Create summaries of interesting events at the end of a run, of a day, …
- 9.1.2: Monitoring, anomaly detection and prediction
  - Improve on rule-based anomaly detection (deals badly with unexpected conditions or complex contexts)
  - Assist data quality monitoring
  - Anticipate failures

WP9.1: Artificial intelligence in experiments

- 9.1.3: Experiment Autopilot
  - Combine Generative AI and anomaly detection and prediction to build a real “self-driving” DAQ

- Note: we are still in the phase of thinking how to tackle this whole activity, so more detailed plans will follow later

WP9.2: DAQ with COTS technologies

- COTS: commercial-off-the-shelf
  - HEP experiments generate huge amounts of data to read and process in real-time
  - Data-centre industry faces similar challenges and delivered cost-effective solutions
- Goals:
  - Demonstrate the usability of industry technologies for HEP DAQ
  - Exploit the new features that they offer to our advantage

Activity leaders: F. Pisani (LBC)
E. Pozo (ADT), R. Sipos (DT)

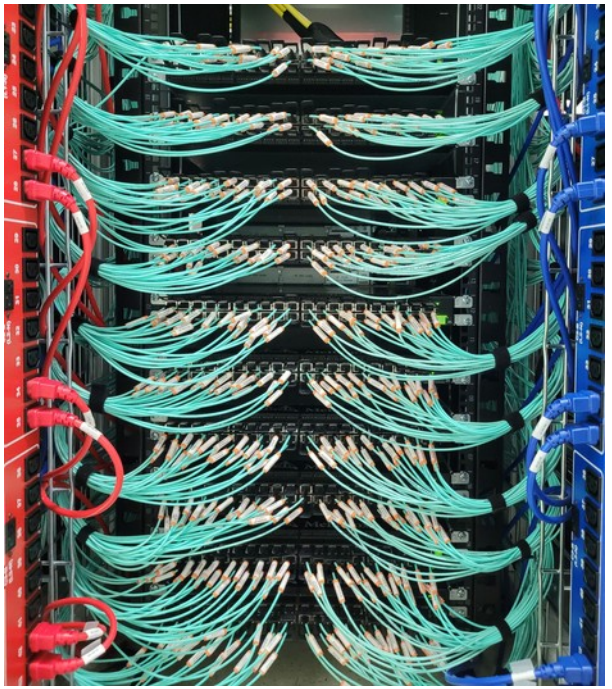

WP9.2: DAQ with COTS technologies

Total readout throughput of HEP experiments:

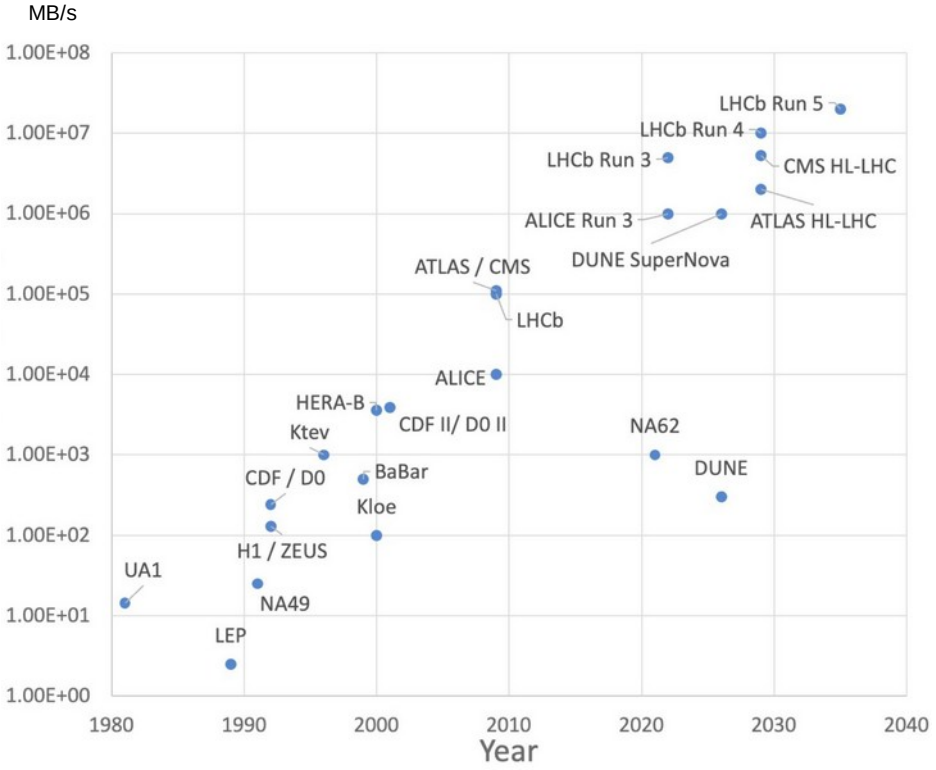

Throughput of a **single-chip** network switch:

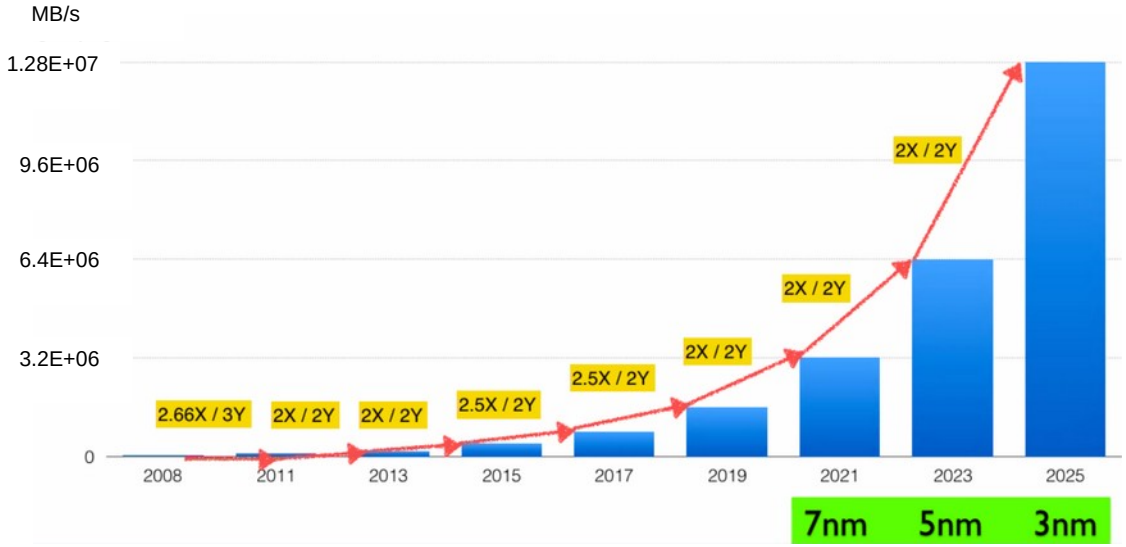

A. Bechtolsheim (Arista Networks)

WP9.2: DAQ with COTS technologies

- 9.2.1: Study and optimisation of DAQ architectures
  - Identify strength and weaknesses of current and past DAQ systems
  - Study the architectural changes in DAQ enabled by:
    - PCIe Gen6
    - 800+ Gb/s networks
    - SmartNICs (i.e.: network cards with an accelerator CPU or FPGA on them)
    - Fine-grained time synchronization for event building (i.e.: exploit the precise timing of DAQ links to use EB networks more efficiently)
- 9.2.2: software readout (Ethernet readout based on DPDK)
  - Create a COTS back-end for the Ethernet-based front-end being developed by DRD7.5

WP9.2: DAQ with COTS technologies

- 9.2.3: Ethernet remote DMA (RoCEv2) for event building
  - Remote DMA (RDMA) accelerates networks by minimising the amount of data copying
  - Ethernet is the only real multi-vendor network standard
  - Using some form of RDMA over Ethernet is almost unavoidable for high-throughput event building at low cost
  - Goal: evaluate and compare performance, cost, and scalability of event builder architectures based on RDMA over Ethernet networks

WP9.3: Generic FPGA-based readout board

Activity leaders: F. Costa (AID)
P. Durante (LBC), E. Gamberini (DT)

- Readout back-ends are almost always experiment-specific
  - different requirements, different timescales
- Increased requirements from new and upgraded experiments correspond to a growing complexity of back-end FPGA boards
- Goal: identify a way to share back-end building blocks
  - Identify problems and evaluate the effort for porting readout firmware used in different experiments between Altera and Xlilinx FPGAs
  - Identify the challenges in using a common readout board across experiments

WP9.3: Generic FPGA-based readout board

- 9.3.1: Firmware portability framework
  - Improve the portability of the firmware by removing as much dependency as possible on specific hardware
  - Evaluate these efforts by developing a readout firmware that can work on both the FELIX board (Xilinx) and the PCIe40 (Altera)
- 9.3.2: Conversion from custom to Ethernet protocol
  - Implement an lpGBT-to-Ethernet converter on a suitable FPGA
  - Evaluate the cost and feasibility of a readout based on such a converter (as opposed to a more “classic” approach with a PCIe card) for a CERN experiment

WP9.4: Alternative Programming Languages

Activity leaders:
F. Le Goff (ADT)
F. Pisani (LBC)

- C++ everywhere: high performance, hard to use
  - Easy to interface with raw memory, drivers, wide platform support, know-how in HEP community
  - Easy to introduce memory access bugs, complex language rules
- Goal: explore the use of alternative languages for DAQ
  - Improve software robustness
  - Easier development and maintainability
  - Attract younger software experts

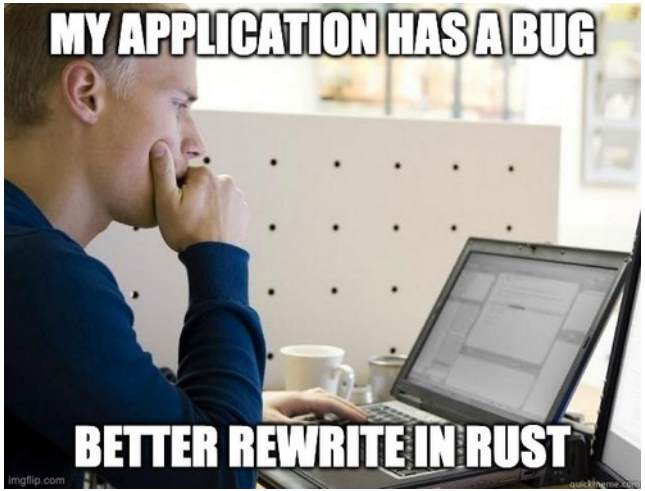

WP9.4: Alternative Programming Languages

- 9.4.1: Rust Implementation of LHCb Run 3 Event Builder
- 9.4.2: Rust bindings for DIM
- 9.4.3: Rust implementation of ATLAS asynchronous network messaging library
- Common objectives:
  - Attempt to port important pieces of DAQ infrastructure to Rust
  - Document the difficulties encountered and the advantages gained
  - Demonstrate that performance isn't impacted
  - Prove interoperability with existing software

WP9.5: Power efficient DAQ solutions

Activity leader:
W. Vandelli (ADT)

- Motivation: obvious
  - Largest energy user in DAQ systems is not DAQ itself, data filtering stages dominate
  - DAQ systems can still do their part
- Goals:
  - Reduce energy usage of data movement and storage components
  - Better use of the power-management features built into typical data filtering hardware

WP9.5: Power efficient DAQ solutions

- 9.5.1: Energy efficient CPUs for data-flow
  - Port subset of existing DAQ software to energy-efficient ARM processors
  - Evaluate co-processors for offloading data-flow tasks
- 9.5.2: Energy efficient architectures for storage
  - Measure power consumption of existing deployments at ATLAS / LHCb
  - Evaluate low-power alternatives: ARM, storage accelerators, SmartNICs
- 9.5.3: Power management solutions for distributed event processing
  - Establish benchmarks representative of data filter workload
  - On various hardware platforms, evaluate the impact of power management features on both energy savings and performance

WP9.6: Direct liquid cooling

Activity leader:
F. Sborzacchi (LBC)

- As compute power grows, so do physical density and thermal power density
  - Cooling dense data-centres efficiently:
    Either: high ΔT air (25° C in, 45° C out)
    Or: direct liquid cooling
  - Some DAQ components (timing, optics) need either more controlled temperature

- 9.6.1 goals:
  - Evaluate the power efficiency gains of liquid cooled data centres
  - Prototype with commercial solutions
  - Develop solutions to adapt custom-made readout boards to liquid cooling

Where do we go from here?

- WP9 planning documents are in the final stage of review
- Hiring students and graduates soon
- See you at the next EP R&D day, with results instead of plans!