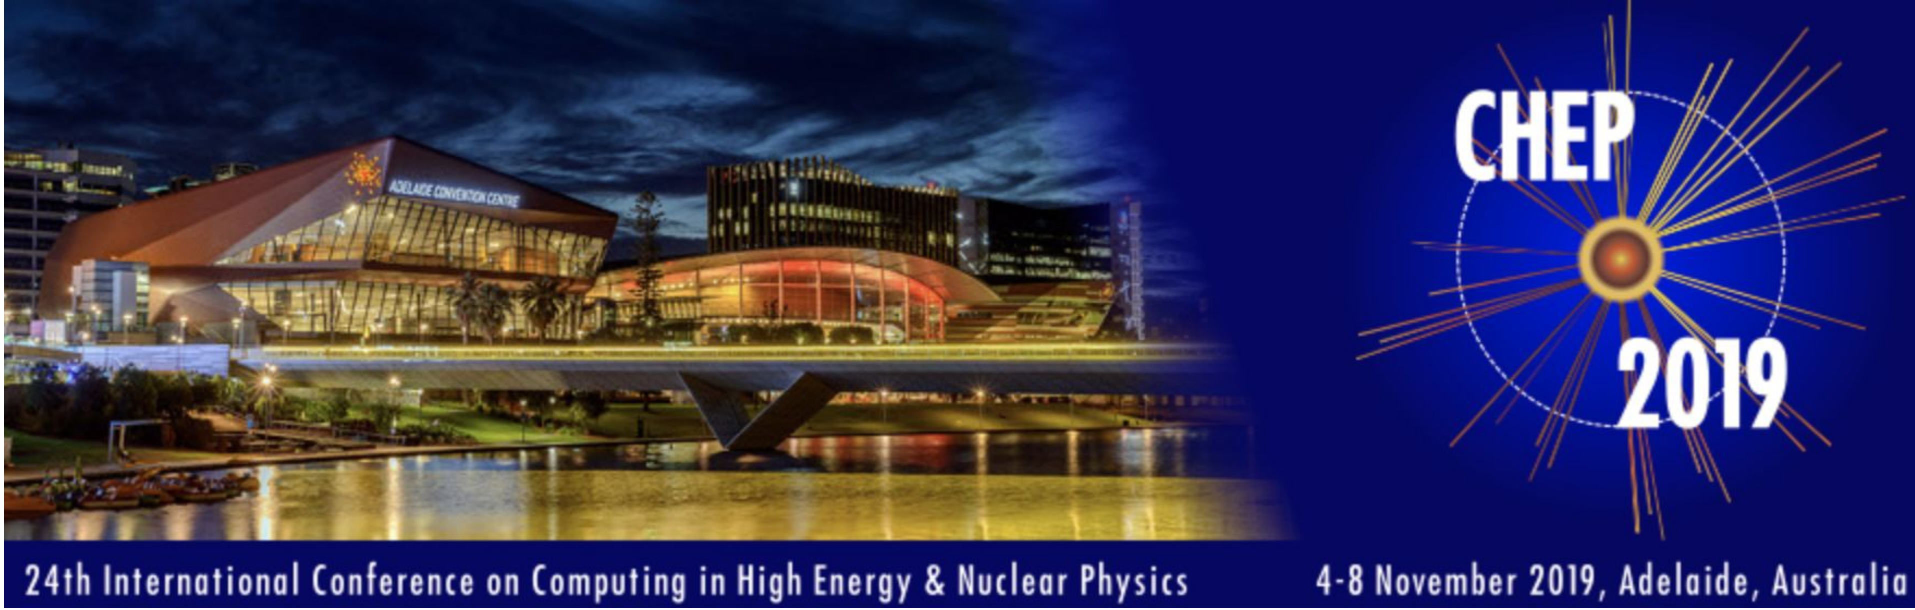

CHEP 2019 Highlights

David Lawrence

Nov. 12, 2019




HDR

Monitoring and control systems

- Community converging around two message passing options
  - FairMQ for heavy ion experiments, ZeroMQ for everyone else
- Live monitoring is a CPU-expensive task
  - Huge amount of monitoring data (often histograms), updated frequently
  - Effort is being invested into reducing computing costs
- Increasing effort to reduce shifter dependencies
  - LHCb already requires minimal shifters due to being predominantly automated
  - CMS is developing DAQExpert to resolve DAQ-related problems faster than shifters
  - Belle II is working to automate PXD and SVD DAQ, ideally reducing shifters

Trigger farms and networks

UNIVERSITÉ DE GENÈVE

- Trigger farms seem to increasingly be integrating GPUs
  - ALICE: farm is based upon GPUs, note that no actual "trigger" rather just data compression
  - CMS: HLT farm will contain 1x GPU per server
  - LHCb: full software trigger (no hardware trigger), first level potentially fully GPU-based
- Servers and their configurations are very complex to thoroughly optimize
  - Reading the server specifications is not even close to the full story
  - ALICE presented a very nice study showing dependencies on severs and their configurations
- Trigger networks primarily studying usage of 100 Gbps, mixture of RoCE and Infiniband
  - Very few places where 200 Gbps Infiniband is used

Hardware acceleration and machine learning

- Growing usage of GPUs and FPGAs in trigger applications
- Key use-case #1: track reconstruction on GPUs or FPGAs
  - Tracks are key to success in high-pileup environments
  - This is an expensive and slow process on CPUs $\Longrightarrow$ accelerate it!
  - GPU-based tracking planned by multiple groups in the coming years (more on this later)
  - CMS has shown a FPGA-based tracking study for the hardware trigger (running at 40 MHz) for the HL-LHC; FPGA tracking is working up to 300 simultaneous collisions
- Key use-case #2: machine learning in FPGAs
  - Regressions, classification, etc all appear very promising to improve trigger performance
  - Integration into currently-simplistic hardware trigger algorithms allows for large gains
  - FPGA implementations of NNs studied with several different approaches: HLS4ML or direct

Common Problems, Common Solutions

- LHC experiments do not dominate in data volume
- FTS is de-facto THE data transfer service
- Rucio is Breaking ATLAS Boundaries
  - missing functionality contributed by comminuty

3rd Party copy (the problem)

- GridFTP (Globus Toolkit) is not supported anymore (Dec 2018)
- Working alternatives are required at all sites …
  - DPM, StoRM, dCache, XrootD, EOS, ECHO, Dynafed
  - FTS, Rucio, gfal ...

Performance

Large effort to utilize all dimensions of performance

- multithreaded frameworks established
- thread performance by e.g. ATLAS, ILC, CMS, LArSoft
- GPU effort by e.g. LHCb, CMS, ATLAS

⇒ Shift to optimize in multiple dimensions

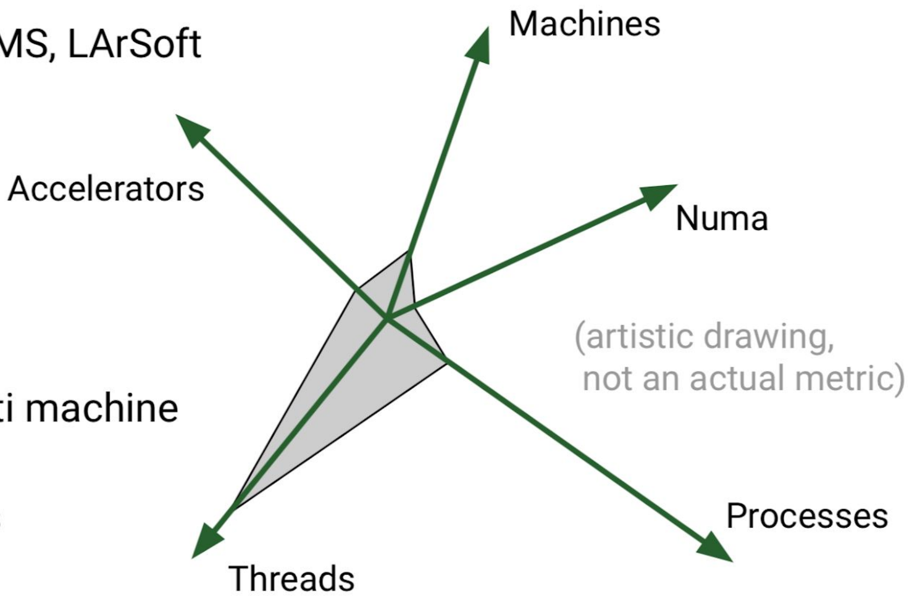

Two complementary approaches

- monolithic framework augmented by multi machine workflow management (e.g. Raythena)
- loosely coupled microservice frameworks (e.g. ALFA, CLARA)

Workflow management

Evolution in workflow management automation

- shift to more and advanced automated workflow decisions
- from early in the process (e.g. LHCb trigger configuration)
- up to analysis level (e.g. law)
- including the big frameworks (e.g. ATLAS, CMS, ALFA, CLARA, ...)

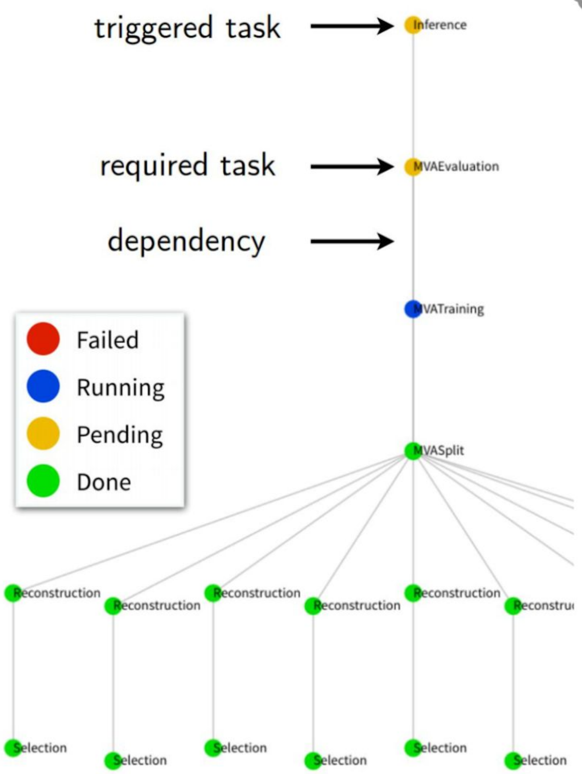

Spack Package Manager

Picking up momentum in HEP

- used by FCC, Key4HEP, SupeNEMO
- FAIR moving to spack
- CMS has proof of concept
- ATLAS considering but things still to be understood
- LHCb and Belle II willing to follow

spack.io

SpackDev: Multi-Package Development with Spack

“Coordinated build & test for integration, or initialize an environment for rapid build & test cycles of a particular package.”

- extension to Spack to help with development of interdependent packages

Packaging: Other Options and Use Cases

Spack is not the only possible solution

- Gentoo Prefix and many more

Conda and Conda Forge

"Reliably install ROOT in under 5 minutes on any machine"

Chris Burr

- user centric software management
- increasing ease of use of many of our software packages

```
conda install python=3.8 root uproot boost-histogram
```

High Level Libraries (and more)

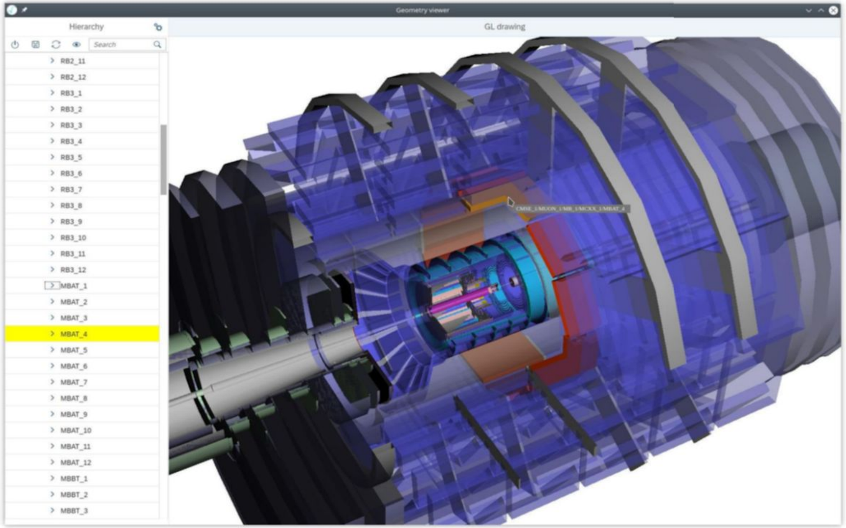

ROOT moving to Web technologies for Visualization

- reduce dependency on system libraries
- experimental but feedback welcome!

Boost Histogram

Standalone, high-performance, header-only histogramming library added to boost

- good python bindings

```
histogram(regular(20, -3, 3, options=underflow | overflow))
                  +----------------------------------------------------------+
[-inf,   -3) 9    |                                                          |
[  -3, -2.7) 19   |=                                                         |
[-2.7, -2.4) 36   |==                                                        |
[-2.4, -2.1) 110  |=====                                                     |
[-2.1, -1.8) 191  |=========                                                 |
[-1.8, -1.5) 275  |=============                                             |
[-1.5, -1.2) 518  |=========================                                 |
[-1.2, -0.9) 644  |===============================                           |
[-0.9, -0.6) 914  |============================================              |
[-0.6, -0.3) 1107 |=====================================================     |
[-0.3,    0) 1183 |========================================================= |
[   0,  0.3) 1185 |========================================================= |
[ 0.3,  0.6) 1120 |======================================================    |
[ 0.6,  0.9) 874  |==========================================                |
[ 0.9,  1.2) 663  |================================                          |
[ 1.2,  1.5) 491  |========================                                  |
[ 1.5,  1.8) 322  |===============                                           |
```

MACHINE LEARNING FOR DATA ANALYSIS

- **Domain aware / physics informed / physics inspired ML algorithms**
- BSM physics searches at the LHC
  - Unsupervised ML for anomaly detection
  - Neural networks, reinforcement learning for jet tagging
- Nuclear physics track fitting ➡ importing tools from HEP to NP
- ML-based Monte Carlo event generators
- High-dimensional ML inference – beyond summary statistics

LATTICE QCD

- **Hadron and nuclear physics from the Standard Model**
- **Extreme-scale computations**
  - ~10^7 GPU hours
  - 100-1000 GPUs in parallel
  - >10% of open-science supercomputing in e.g., USA
- **Era of precision calculations, fully-controlled uncertainties**
  - Hadron structure relevant to LHC, EIC
  - New physics searches e.g., muon g-2
  - Moving towards nuclei for intensity frontier experiments

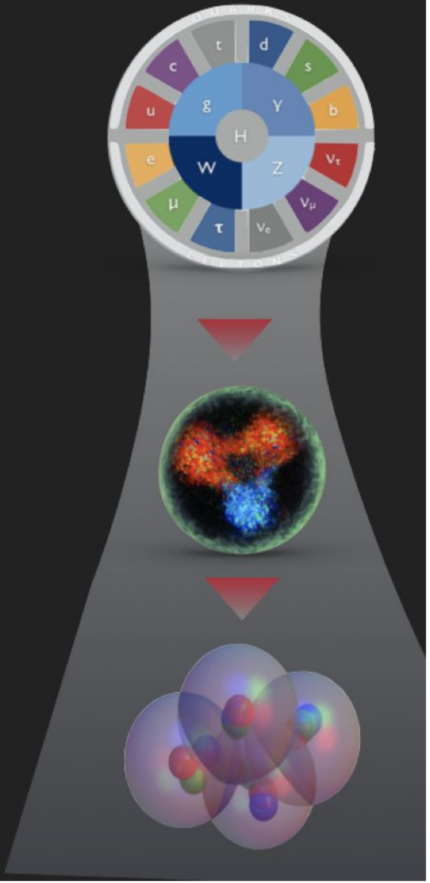

TRACK 6 POSTER HIGHLIGHT – JOSH CHARVETTO (UNI OF ADELAIDE)

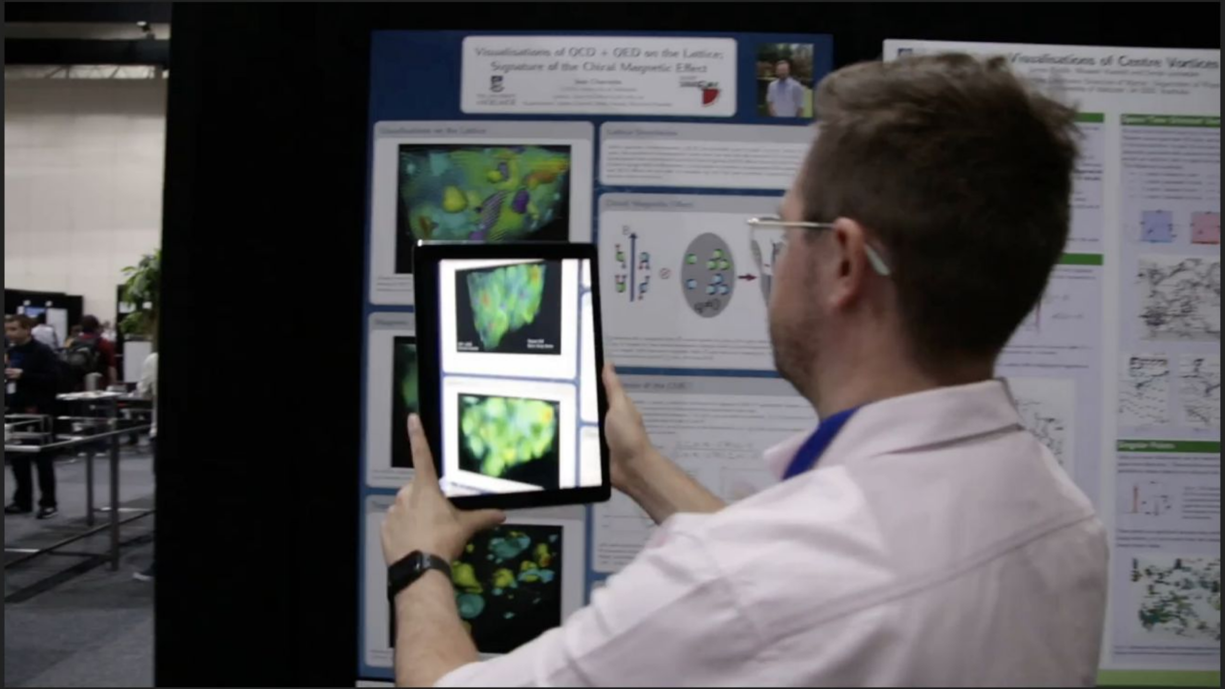

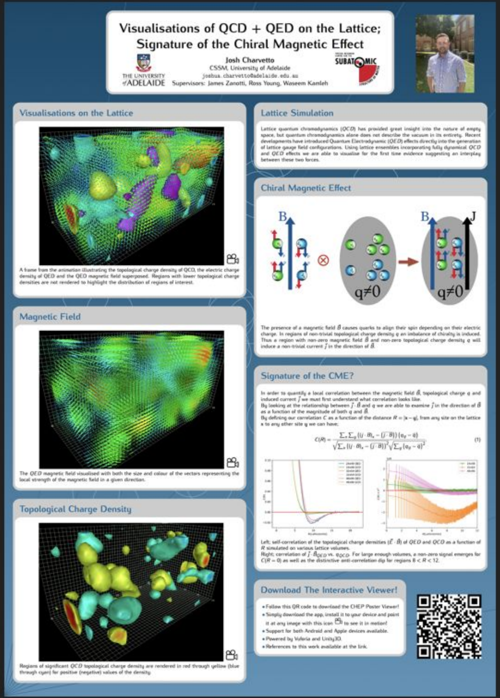

Trends

- Jupyter
  - Increasing interest on the user side
    - Interactive, fast learning curve
    - Easy development and sharing
  - Jupyterhub enables access to diverse HTP and HPC resources
- Overview by Hepix TechWatch working group
  - Hyperscales (Google, Amazon) drive the market
  - x86 market: AMD is back
  - Magnetic disk: Market is shrinking
  - Tape: Risks – essentially one company left for R&D
  - Ethernet evolving very fast
    - Pace of change exceeding IEEE standards process

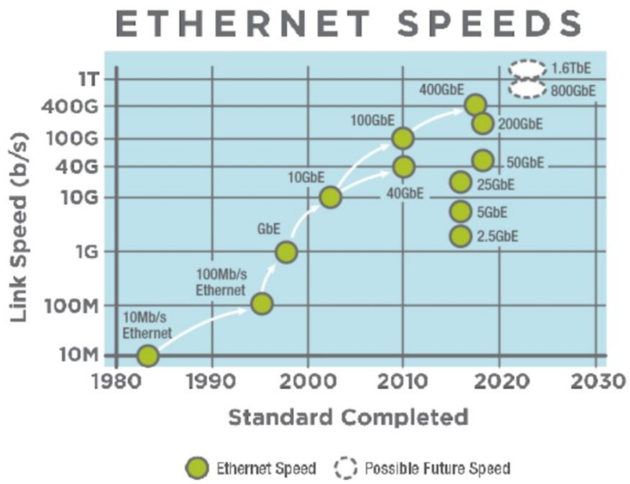

Containers

- Kubernetes, kubernetes, kubernetes, …
  - Intensive deployment of containers based on Kubernetes engine
  - Job scheduling without batch system, simply relying on Kubernetes
    -> effective and promising, aspect to simplifying site operations, no CEs...
  - Registry solutions to deploy container images is one of concerns
- ScienceBox
  - Complete solution for scientific set of services
    from highly-scalable storage solutions (EOS)
    to user-friendly application, Jupyter notebook
  - All nicely packaged in containers
- Container technology facilitates use of various resources
  - HPC, HTC, Grid resources, etc.
- Moving CERN batch from Openstack VMs to Kubernetes
  - First benchmarks indicate 5% performance gain
- All major experiments use containers in production

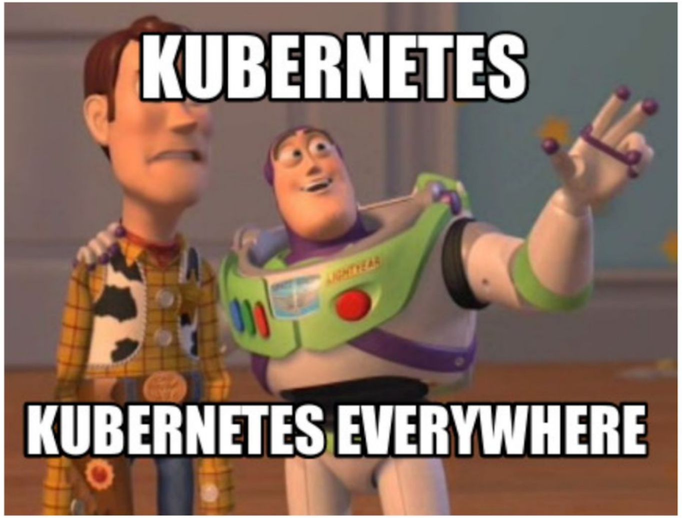

Open data is growing

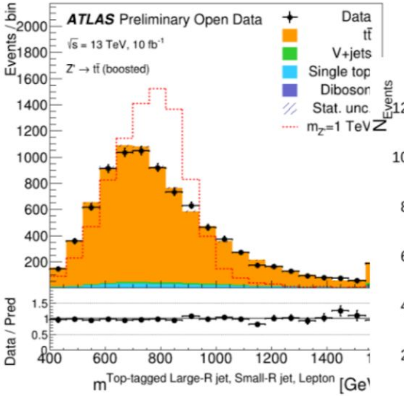

ATLAS

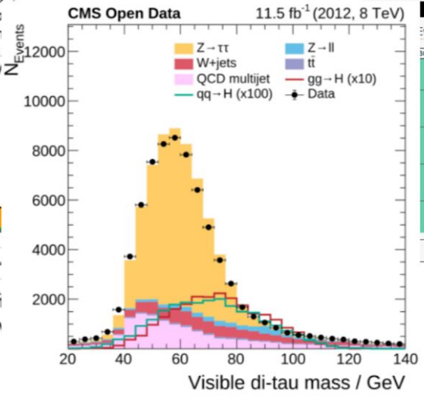

CMS

OPERA

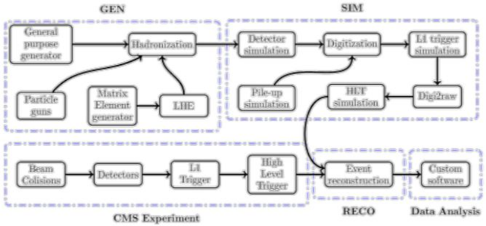

Experiments releasing more open data batches

- ATLAS: New 13 TeV 2016 data samples
- CMS: Machine Learning, RAW/AOD/AODSIM
- OPERA: tau-neutrino events

Growing use of open data for research

- Independent research papers on INSPIRE
- CMS released full provenance information as well as raw data samples

Saito M.

Quantum Tracking

Quantum Computer

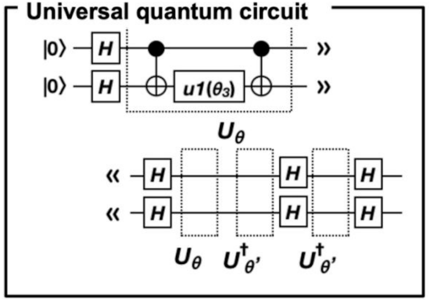

- Arrange gates for each problem
- General-purpose computer

D-Wave

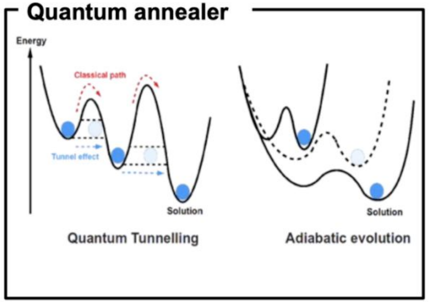

- Find the minimum energy state of a given Hamiltonian
- Suitable for an optimization problem

Focus on quantum annealer

4

Results

1600 particles (20% of HL-LHC)
- 11000 hits

Input

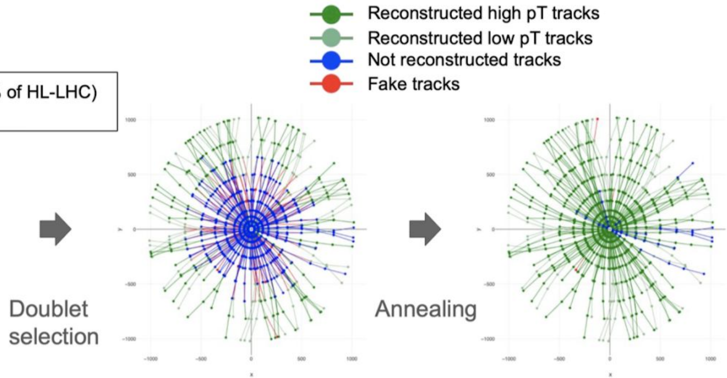

390000 Doublets	
Purity	0.22 %
Efficiency	99.5 %

2445 Doublets

1424 Doublets	
Purity	98.5 %
Efficiency	96.4 %

11

Not just a pipe dream!?

J-R Vlimant

Tracking ML challenge

Leaderboard

	RESULTS										
	#	User	Entries	Date of Last Entry	score ▲	accuracy_mean ▲	accuracy_std ▲	computation time (sec) ▲	computation speed (sec/event) ▲	Duration ▲
HEP people	1	sgorbuno	9	03/12/19	1.1727 (1)	0.944 (2)	0.00 (14)	28.06 (1)	0.56 (1)	64.00 (1)
HEP people	2	fastrack	53	03/12/19	1.1145 (2)	0.944 (1)	0.00 (15)	55.51 (16)	1.11 (16)	91.00 (6)
PH+CS	3	cloudkitchen	73	03/12/19	0.9007 (3)	0.928 (3)	0.00 (13)	364.00 (18)	7.28 (18)	407.00 (8)


Norfolk, Virginia, USA
May 10-14, 2021
25TH INTERNATIONAL CONFERENCE
CHEP
2021
Computing in High Energy & Nuclear Physics
LOCAL ORGANIZING COMMITTEE
Amber Boehnlein - Chair
David Abbott
Mark Ito
Brent Morris
Graham Heyes
Rachel Harris
VISIT Norfolk
Jefferson Lab
JSA
chep2021.org

Elizabeth River Waterfront, Norfolk

- Scenic, walkable downtown with restaurants, lodging, shopping, museum and entertainment conveniently located near a major airport
- Light Rail Line to access to additional restaurants and museums

Social Activities

- Monday Night Reception will be held at the conference hotel
  - Feature Virginia Craft Beer and Virginia Wine: 4 excellent craft brewers and several vineyards in local area
- Banquet to be held at the Blue Moon Taphouse, Waterside
  - Music provided by Jae Sinnett, local Jazz legend
- Excursion/free afternoon options
  - Trips to the beach or seashore nature preserves
  - Norfolk Botanical Gardens
  - Geek tours: Port of Virginia and NATO Allied Command
  - Kayaking on the James

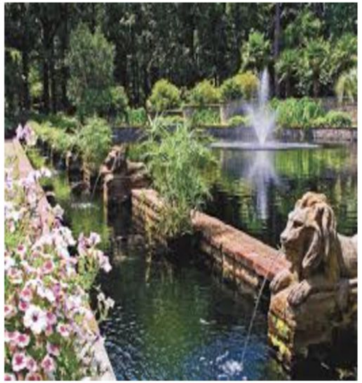

- Dining
  - Over 45 Restaurants near the Conference Hotel