CHEP 2018 Summary

Stewart Martin-Haugh

RAL Particle Physics Seminar
19 September 2018

Science & Technology Facilities Council

Introduction

- Random, non-comprehensive set of things I found interesting
- Some local colour

Traditional Bulgarian folk dance

Temporary hotel room upgrade

Conference venue

Conference dog

Nearby hills

Next Generation Generative Neural Networks for HEP

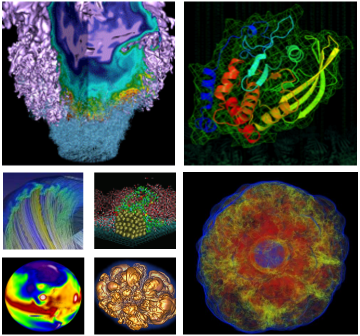

Steve Farrell, Wahid Bhimji, Thorsten Kurth, Mustafa Mustafa, Debbie Bard, Zarija Lukic, Ben Nachman, Harley Patton

CHEP 2018, Sofia Bulgaria

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

HEP simulation

- **Simulation is an essential application for HEP**
- **We have very powerful tools for simulation**
  - And active R&D programs
- **But gall darn is it *expensive*!**
  - Large cost in CPU resources
  - Large manpower cost in developing fast-simulation methods

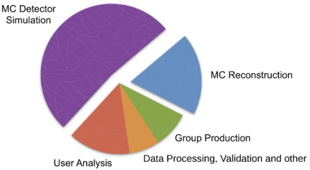

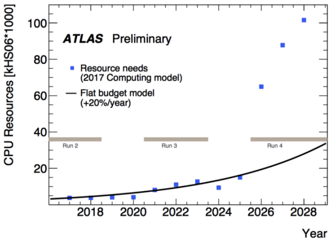

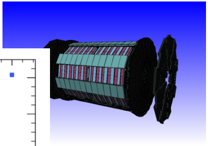

Deep learning generative models

- **Deep neural networks that learn to sample from a data distribution**
  - Transform a simple "noise" distribution to the target distribution
- **Popular examples:**
  - Variational Autoencoder (VAE)
  - Generative Adversarial Network (GAN)
- **The GAN framework poses the problem as a trainable two-player game**
  - A *generator* tries to produce realistic samples
  - A *discriminator* tries to distinguish real from fake samples
- **GANs are notoriously unstable to train**
  - Difficult to define good metrics for learning

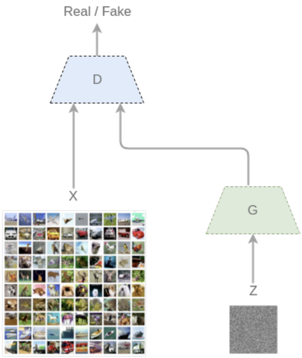

Pileup GAN - μ=20

Real

Fake

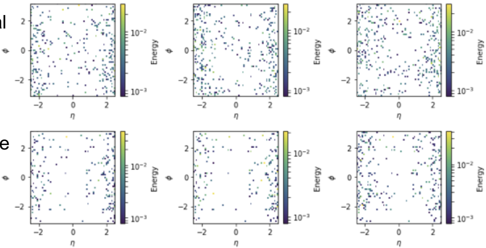

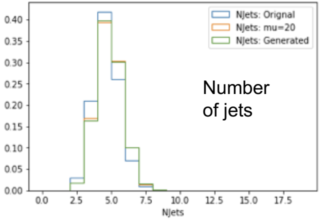

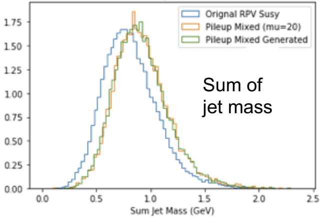

- **The GAN gives realistic looking pileup images**
- **When overlayed onto RPV events, we see realistic shifts in the distributions**

U.S. DEPARTMENT OF ENERGY | Office of Science hls4ml

Deep learning on FPGAs for L1 trigger and Data Acquisition

CHEP 2018, 9-13 July 2018, Sofia, Bulgaria

Javier Duarte, Sergo Jindariani, Ben Kreis, Ryan Rivera, Nhan Tran (Fermilab)
Jennifer Ngadiuba, Maurizio Pierini (CERN)
Edward Kreinar (Hawkeye 360)
Phil Harris, Song Han, Dylan Rankin (MIT)
Zhenbin Wu (University of Illinois at Chicago)
Sioni Summers (Imperial College London)

CERN

CHEP 2018
Sofia, Bulgaria

Data processing @ LHC

Latency

1 s	Offline
1 ms	HLT
1 us	
	L1 Trigger
1 ns	

ML methods typically employed in offline analysis or longer latency trigger tasks

Many successes in HEP: identification of b-quark jets, Higgs candidates, particle energy regression, analysis selections, ….

Exploration of ML algorithms in low-latency, real-time processing has just begun!

What can we do in < us on one FPGA?

Case study: jet tagging

- Study a 5-output multi-classification task: discrimination ***q, g, W, Z, t*** initiated jets
- Fully connected neural network with **16 inputs**:
  - **mass** (Dasgupta et al., arXiv:1307.0007), **multiplicity, energy correlation functions** (Larkoski et al., arXiv:1305.0007)
  - expert-level features not necessarily realistic for L1trigger, but the lessons here are generic

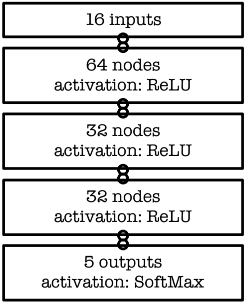

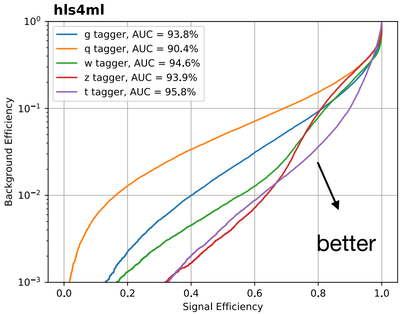

AUC = area under ROC curve
(100% is perfect, 20% is random)

Efficient NN design for FPGAs

- We focus on tuning the neural network inference such that it uses the FPGA resources efficiently and meets latency constraints

- We have three handles:

  - **compression:** reduce number of synapses or neurons
  - **quantization:** reduces the precision of the calculations (inputs, weights, biases)
  - **parallelization**: tune how much to parallelize to make the inference faster/slower versus FPGA resources

Summary

We introduced a new software/firmware package **hls4ml**

Automated translation of everyday machine learning inference into firmware in ~ minutes

Tunable configuration for optimization of your use case

First application is single FPGA, <1 us latency for L1 trigger or DAQ

Explore also applications for acceleration with CPU-FPGA co-processors for long latency trigger tasks

For more info

- https://hls-fpga-machine-learning.github.io/hls4ml/
- https://arxiv.org/abs/1804.06913

Fermilab U.S. DEPARTMENT OF ENERGY Office of Science

Blockchain for Large Scale Scientific Computing

Lindsey Gray, Lothar Bauerdick, Bo Jayatilaka, Gabe Perdue (FNAL),
Josh Bendavid (CERN)

12 July, 2018

CHEP 2018
23RD INTERNATIONAL CONFERENCE ON
COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS
9-13 July 2018
National Palace of Culture
Sofia, Bulgaria

Blockchain in a nutshell

- Pseudonymous paper and code (Satoshi Nakamoto) led to concept of Bitcoin in 2009
- Relies on functions that are difficult to compute, but easy to verify
  - e.g. factorisation: $5775 = 3 \times 5 \times 5 \times 7 \times 11$
- Blockchain: distributed graph of participants
- Does not require a central trusted "ledger" of transactions
- Most applications financial: can it be used in HEP?

Putting it all together for HEP Computing (simple example!)

Nodes define needed resources and available jobs

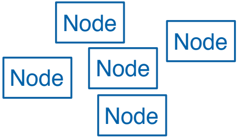

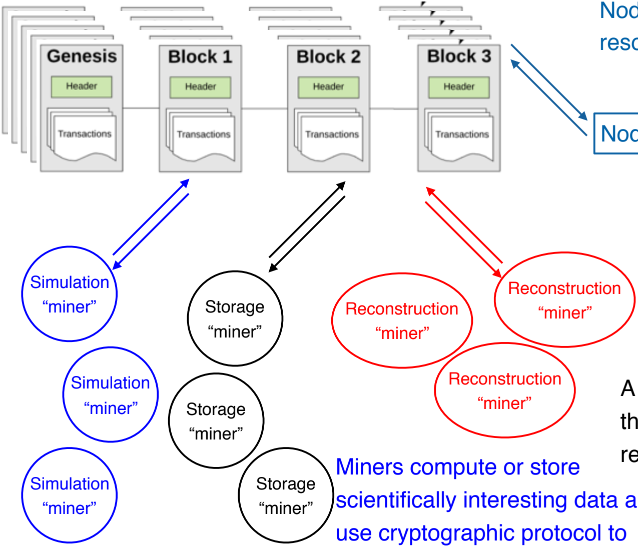

The blockchains maintain consensus over job states and quality of results with smart contracts and transaction logs

A further blockchain maintains the consensus between all of respective sub-chains.

Miners compute or store scientifically interesting data and use cryptographic protocol to prove fidelity of results for reward!

All powered by useful computations rather than hashing!

Quantum Computing

Elizabeth Sexton-Kennedy for James Amundson, *Fermilab, Batavia, Illinois USA*

CHEP 2018

2018-07-12

Quantum Computing Excitement

Nov. 13, 2017

The New York Times

Yale Professors Race Google and IBM to the First Quantum Computer

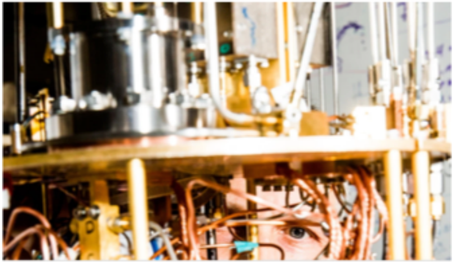

More Quantum Computing Excitement

October 16, 2017 THE WALL STREET JOURNAL.

Quantum Computing Excitement Has Reached the U.S. Congress

June 8, 2018

GIZMODO

Two Quantum Computing Bills Are Coming to Congress

Ryan F. Mandelbaum
6/08/18 5:26pm · Filed to: QUANTUM COMPUTING

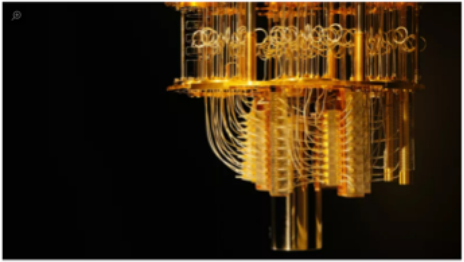

A dilution refrigerator from an IBM quantum computer.
Photo: IBM Research (Flickr)

…and Congress Appropriates money

Where the Excitement Started

- Peter Shor: A general-purpose quantum computer could be used to efficiently factor large numbers
  - Shor's Algorithm (1994)
  - Resource estimates from LA-UR-97-4986 "Cryptography, Quantum Computation and Trapped Ions," Richard J. Hughes (1997)

num size	1024 bits	2048 bits	4096 bits
qubits	5124	10244	20484
gates	3×10^{10}	2×10^{11}	2×10^{12}

Analog of clock cycles in classical computing

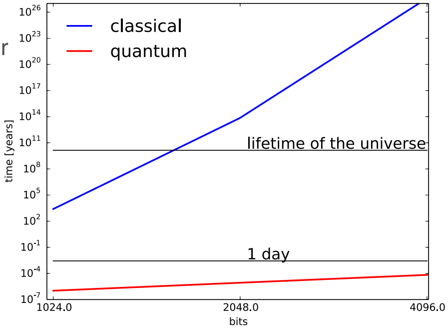

n.b. This is an old estimate; improvements have been made in the meantime.

Fermilab

Theoretical Computer Science

- Classical Computing
  - "Easy" problems can be solved in "polynomial time" (**P**)
  - "Hard" problems require "nondeterministic polynomial time" (**NP**)
    - Proving P ≠ NP is a great unsolved problem in computer science
- Quantum Computing
  - Some problems are easy in quantum computing, but hard in classical computing -> quantum complexity classification
  - Some problems appear to be hard either way

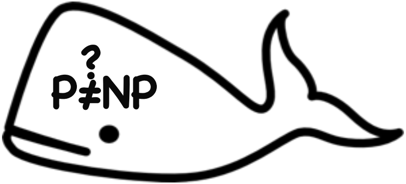

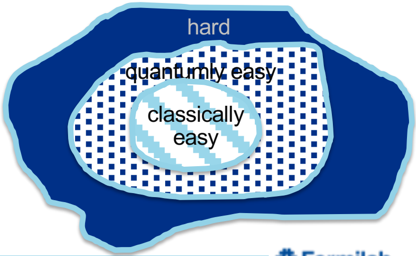

Quantum Algorithms

- Shor's Algorithm: factorization -- Speedup: Superpolynomial
- Grover's Algorithm: search -- Speedup: Polynomial
- If there exists a positive constant α such that the runtime C(n) of the best known classical algorithm and the runtime Q(n) of the quantum algorithm satisfy $C=2^{\Omega(Q\alpha))}$ then the speedup is superpolynomial, otherwise it's polynomial.
- Many more available at the Quantum Algorithm Zoo https://math.nist.gov/quantum/zoo/
  - **A catalog of 60 quantum Algorithms in 3 categories:**
    - **Algebraic and Number Theoretic Algorithms -> cryptography**
    - **Oracular Algorithms→ optimization and machine learning**
    - **Approximation and Simulation Algorithms -> quantum physics and chemistry**

Current and Near-term Quantum Hardware

Ion trap

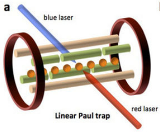

Scientific Reports 4, 3589 (2014)

NMR

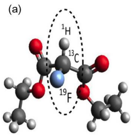

Sci. China Phys. Mech. Astron. 59:630302 (2016)

NV center

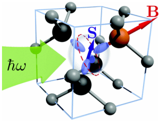

Phys. Rev. B 86, 125204 (2012)

Quantum dot

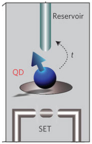

4 Nature Nanotechnology 9, 981–985 (2014)

Linear optical

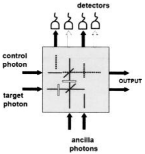

J. Opt. Soc. Am. B, 24, 2, 209-213 (2007)

Superconducting

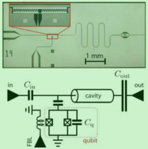

Ann. Phys. (Berlin) 525, 6, 395–412 (2013)

- Thanks to Andy Li
  - Fermilab Scientific Computing Division's first quantum computing postdoc!

… many more

- Superconducting is the most prominent commercial HW and was presented at CHEP2016

Counting Qubits is Only the Beginning

- The number of gates that can be applied before losing quantum coherence is the limiting factor for most applications
  - Current estimates run few – thousand
  - Not all gates are the same
    - The real world is complicated
- IBM has a paper proposing a definition of “Quantum Volume”
  - Everyone else seems to dislike the particular definition
  - The machines with the largest number of qubits are unlikely to have the largest quantum volume

From the earlier factoring estimate

num bits	1024 bits	2048 bits	4096 bits
qubits	5124	10244	20484
gates	$3x10^{10}$	$2x10^{11}$	$2x10^{12}$

- “Logical qubits” incorporating error correction are the goal
  - Probably require ~1000 qubits per logical qubit
    - Minimum fidelity for constituent qubits is the current goalpost

Successful Quantum Simulation

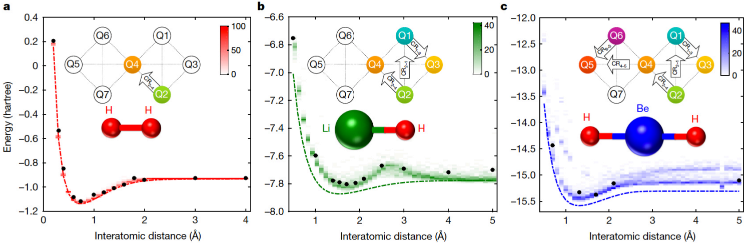

- Quantum Chemistry has the first big successes in quantum simulation.
- GitHub has a project for general simulations of interacting fermions.
- However, interesting HEP systems, e.g., QCD, also require boson-fermion interactions.

Open Fermion

https://github.com/quantumlib/OpenFermion

Conclusions

- Quantum computing holds the promise of remarkable new computational capabilities
  - The future is not here yet
    - … but we are getting there
- Fermilab has quantum computing efforts on many fronts
  - Quantum Applications
  - HEP technology for QC
  - QC technology for HEP experiments
  - Quantum Networking

https://www.smbc-comics.com/comic/the-talk-3

Square Kilometre Array

SKA and its data processing

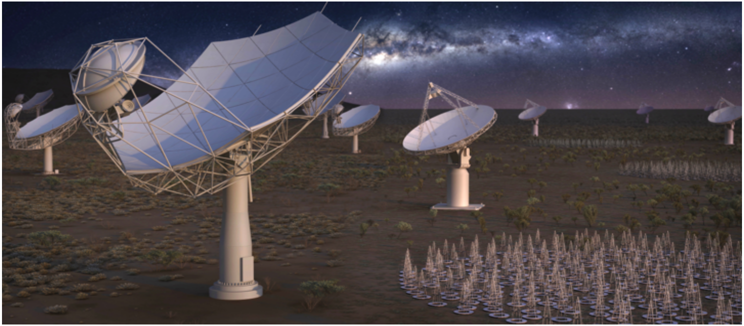

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Rosie Bolton

SKA's Global footprint

11 member countries: AUS, CAN, CHN, ESP, GBR, IND, ITA, NED, NZL, RSA, SWE

Currently in discussion with others: FRA, GER, JPN, KOR, POR, SUI

Why not?

Exascale era – data centres approaching Exaflops and ExaByte capacities

Exa-sky astronomy era?

- Better spatial resolution (1 arcsec)
- Better frequency coverage (~0.05-15 GHz)
- Better frequency resolution (~1-10kHz)
- 20 Exabyte sky!

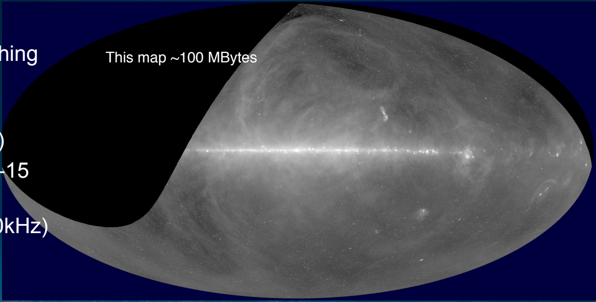

Calabretta, Stavely-Smith and Barnes, 2014 (Re-processing of archival HI Parkes All-Sky Survey data)

(Light reading: https://what-if.xkcd.com/63/)

SKA Science Drivers

Very broad range of science

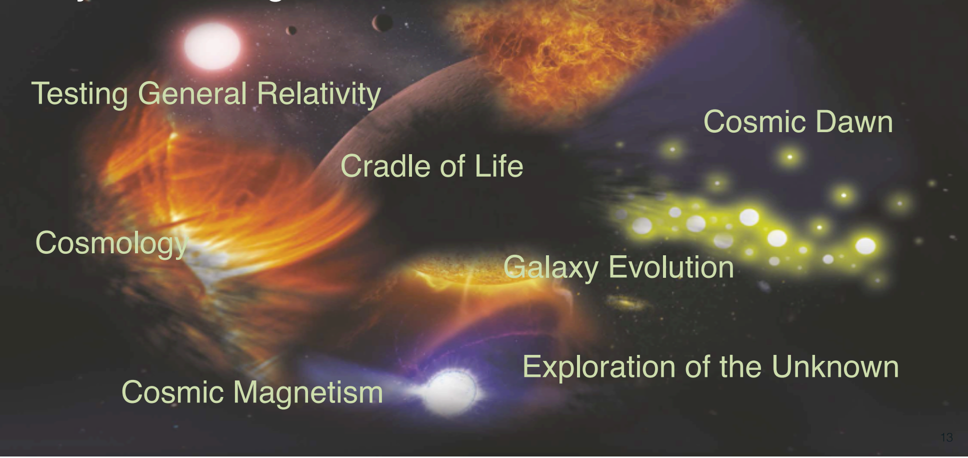

SKA science performance

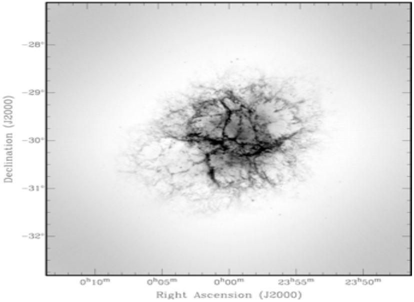

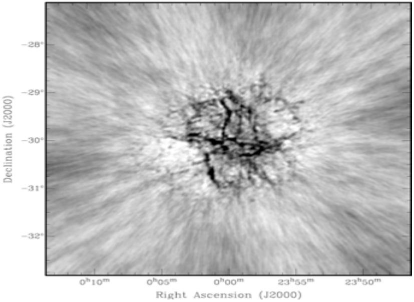

Crab Nebula: Simulated SKA1-LOW snapshot image compared to LOFAR snapshot (noiseless images; target PSF ~10" at 140 MHz)

Data flow challenges

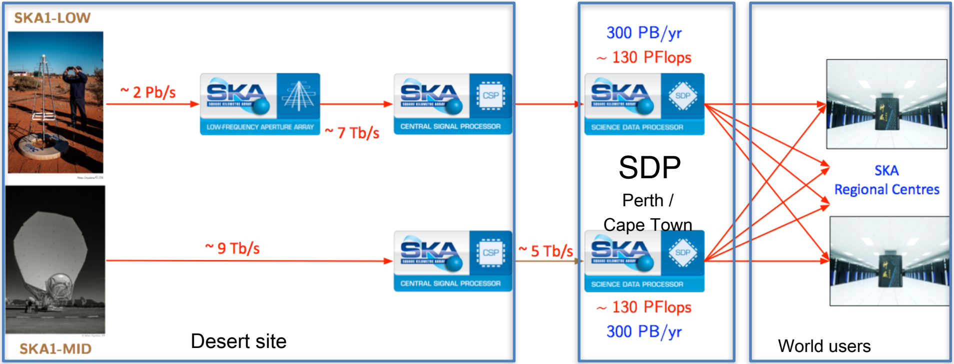

10 – 50 x data rate reduction by SDP

Data flow challenges

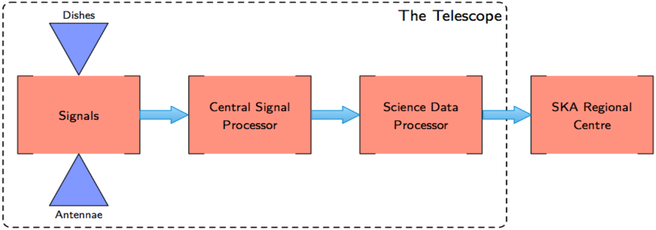

For SKA data rates and volumes emerging from central signal processor are so high that we will not be in a position to store the raw data from the Central Signal Processor

- it will be cheaper to re-observe than store the raw data indefinitely

The science data processor becomes a schedulable resource of the telescope for observation planning

Summary

SKA will be the biggest radio observatory in the world

- 50 lifetime – commissioning data ~2024, Key Science start 2028
- Use cases - science led, very open

Data processing within the observatory ~250 Pflop/s, challenging power, memory bandwidth, computational efficiency.

Delivery to Regional Centers up to 6-700 Pbytes/year (100 gbit/s out of each site)

Users do not "get" their data

- SRC science gateway to visualise data products, develop SRC pipelines and extract final products (e.g. plots)
- SRCs will be community led, hopefully based on a collaborative model
- SRCs are unlikely to be dedicated HW – must support a mix of different models and architectures.

Summary

- More machine learning: fast, doing more things, using dedicated hardware
- Quantum computing growing
- Blockchain?
- SKA has a similar but distinct set of problems to LHC