Jefferson Lab LQCD Computing

April 2020 All Hands Meeting

Robert Edwards & Graham Heyes

Nuclear and Particle Physics Computing Initiative

Reminder

- Single lab, NP funded, serves all of USQCD,and is complementary to the modified (2 lab, IC based) HEP LQCD project
- $1M per year, about half hardware, half labor
  (equals average NP investment per year at JLab for last 10 years, so no real change in funding for the lab)
- FY2018: upgraded Jlab's KNL resources (added to system from FY2016)
- FY2019: upgraded Jlab's GPU resources
- (both of these steps were specified in the initiative proposal)
- FY2020: looking to increase GPU footprint

Nuclear and Particle Physics Computing Initiative

- Chip Watson was head of JLab LQCD project & operations for 20 years
  - Retired in Feb. 2020
- LQCD Project duties split
  - Robert Edwards - PI of LQCD project
  - Graham Heyes - new Head of Scientific Computing - including operations of LQCD
  - Bryan Hess - site manager

- Existence proof:
  - JLab Farm computing system supports Experimental Physics Div.
    - Scientific oversight is Assoc. Dir. Expt. Physics Division
    - IT division manages & operates computing system
- Similar arrangement for LQCD
  - Theory gathers computational requirements & systems purchased with theory funds
  - IT division manages/operates

Evolution of JLab/LQCD program

- When LQCD@JLab started in ~FY01, all funds directed through JLab/Theory
  - Theory and LQCD/IT group have worked closely over 20 years
  - IT program very responsive to needs of LQCD Theory
  - Systems deployed based on computational requirements for LQCD Theory program
- Theory group leader is actually Assoc. Dir. for Computations and Nuclear Theory
- NPPLCI initially originally 4-years with it's own reporting
- Going forward, DOE asked JLab to merge LQCD into JLab Theory Group Annual reporting structure
  - First report (in Feb. 2020) included highlights of scientific program and future plans
  - Expect to carry this structure going forward

Some highlights of scientific progress

- Proton & hadron structure - core mission of DOE/NP research program
  - Axial charge - neutron lifetime (Bhattacharya, PRD94) & other charges constrain BSM searches (Gupta, PRD98)
  - Tensor charge constraints on quark transversity - first global QCD analysis (Lin, PRL120)
  - Near light-cone parton distributions & structure funcs (Karpie, JHEP 12; Sufian, PRD99)
  - Gluon pressure in proton (Shanahan, PRL122)
- Spectroscopy
  - Resolution of scalar, vector and tensor meson octets (Briceno, PRD97; Wilson, PRL123)
  - Electromagnetic structure of resonances (Briceno, PRL115)
- Nuclear structure
  - Gamow-Teller contributions to weak decay of triton (Savage, PRL119; Chang, PRL120)
- Extreme matter
  - Chiral crossover in T - mu_B plane relevant to BES program @ RHIC (HotQCD, PLB795)

Future plans - JLab at intersections

Community white-papers

- Hadrons and Nuclei: http://inspirehep.net/record/1730506
- Thermodynamics: http://inspirehep.net/record/1730501
- Lepton & flavor physics: http://inspirehep.net/record/1730491
- Physics beyond the Standard Model: http://inspirehep.net/record/1730601
- Neutrinos: http://inspirehep.net/record/1730504
- Fundamental symmetries: http://inspirehep.net/record/1730510
- Computing: http://inspirehep.net/record/1730494
  - Trend is ever-more complicated analysis programs
  - Computing WP stresses role of capacity computing
  - A desire going forward - increase the versatility of the capacity computing

2016 KNL cluster

- LQCD science is computing capacity constrained - especially NP
  - Local facility leverage - provide that capacity computing
- In 2016, JLab purchased 260 KNL nodes
  - Single partition - larger mid-range
- Good price performance for a range of LQCD apps
  - Supports range of NP apps not well suited for GPUs (2016 snapshot)
  - Good solution for BNL

2018 KNL cluster upgrade

- LQCD science is computing capacity constrained - especially NP
  - Local facility leverage - provide that capacity computing
- In 2016, JLab purchased 260 KNL nodes
  - Single partition - larger mid-range
- Good price performance for a range of LQCD apps

- In 2018, JLab purchased an addition 180 KNL nodes at ~half cost of the 2016 cost
  - Configured as 4 mini-clusters of 40 - 44 nodes per 100g Intel Omnipath switch
  - Well suited for smaller jobs
  - About 4x performance of SkyLake clusters of similar costs
    - KNL high on-package memory bandwidth

- Now total of 440 nodes in one queue

2019 GPU cluster

- Old quad k20 cluster (2012, Kepler)
  - Reached 6 years in 2018 decision time.
  - Modulo servers slowly dying - deliver 32 K20 node-years

- FY 2019 possibilities considered
  - Quad or Octal v100 (Volta) - roughly 1.5x Pascal or 6x our 2012 Kepler system
  - Octal RTX-2080 (newest gamer at the time)

- Goal
  - Bring new GPU resource online
  - Continue supporting mixed architecture workflows

Factors in GPU selection

- RTX strengths
  - Highest memory bandwidth / dollar
  - Highest multi-grid inverter performance / dollar

- RTX weaknesses
  - No ECC memory (tends to limit to inverters-only)
  - No GPU-GPU communication (must use CPU)
  - No GPU-fabric communications (must use CPU)
  - Small memory per GPU (8 GB; other models had rumored overheating issues)

FY 2019 GPU selection optimization

- With available budget, JLab could have added 0.4M V100 based K80 hours (27% of USQCD GPU resources)
- NP community typically larger user of GPUs
- Multiple NP users favored the higher aggregate performance of RTX-2080 despite weaknesses
- Final configuration - “19g”
  - 32 nodes of octal RTX 2080
  - 19 GB memory / node
  - 100 Gps Omnipath (single rail, single switch)
  - 256 GPUs -> 280 - 380 Gflops/GPU on multi-grid: > 64 TFlops total

2018-2019 Projects @ JLab

P.I.	KNL usage	GPU usage	Project Title
Robert Edwards	60	0.4	Meson Resonances from Anisotropic Clover Lattices
Patrick Steinbrecher	46		Non-Gaussian Cumulants of Conserved Charges Fluctuations
David Richards	36		Pion Properties from Lattice QCD
Tom Blum	30		pi-pi Scattering and K to pi-pi Decay Calculations at the Physical Point
William Detmold	8	0.3	Nuclear Physics from the Standard Model
Rajan Gupta	13		Nucleon Matrix Elements with 2+1 Flavor Clover Fermions
Aaron Meyer	11		Exclusive Study of (g - 2)μ HVP and Nucleon Form Factors with Distillation
Tanmoy Bhattacharya	6		Contribution of Theta, Chromo EDM and Weinberg Operators to nEDM
Swagato Mukherjee		0.1	Non-Gaussian cumulants of conserved charges fluctuations
Oliver Witzel	3		Simulations with 4 6 Flavors Using Smeared Möbius DW Fermions

Mix of Cold, Hot NP, IF & BSM - Thermo projects relied on ECC

2019-2020 Projects

P.I.	KNL award	GPU award	Project Title
Robert Edwards	15	1.1	Meson Resonances and their Couplings from Anisotropic Clover Lattices
David Richards	17	0.4	Parton Distribution Functions and Amplitudes of Pseudoscalar Mesons and Nucleon from Lattice QCD
William Detmold	15	0.3	Nuclear Physics from the Standard Model
Luchang Jin	19	0.1	Neutrinoless Double Beta Decay from di-pion to di-baryon System
Patrick Steinbrecher	24		The Total Chiral Susceptibility at Finite Chemical Potential
Jian Liang	10	0.2	Lattice Calculation of Nucleon Form Factors and EDM using Overlap Fermions
Tom Blum	22		K to Pi Pi Decay Calculations at the Physical Point
Rajan Gupta	21		Nucleon Matrix Elements with 2+1 Flavor Clover Fermions
Tanmoy Bhattacharya	17		Contribution of Theta, Chromo EDM and Weinberg Operators to nEDM
Sergey Syritsyn	12		Calculation of Nucleon Axial Form Factors, Proton Decay Amplitudes, and Nucleon EDMs Induced by QCD theta term and Quark Chromo-
Aaron Meyer	10		Contribution of Theta, Chromo EDM and Weinberg Operators to nEDM
Oliver Witzel	10		Composite Higgs Model with Four Light and Six Heavy Flavors
Ethan Neil	9		Investigation of Near-conformal Anomalous Dimensions using Gradient-flow RG

New GPU awards in Cold NP + IF-NP intersection

Hardware options - FY20

- Conventional cluster w Intel Xeon processors (SkyLake) + Infiniband
- Conventional cluster w AMD EPYC 2 processors (Rome) + Infiniband
- New NVIDIA GPU cluster w Intel or AMD host, NVIDIA Tesla (ECC) GPU + Infiniband
- Upgrade portion of 2019 GPU cluster: RTX-2080 (non ECC) to Tesla GPU (ECC)
- New AMD CPU+ AMD GPU accelerated cluster + Infiniband

Hardware options - FY20 (CPU)

- Conventional cluster w Intel Xeon processors (SkyLake) + Infiniband
- Conventional cluster w AMD EPYC 2 processors (Rome) + Infiniband
- New NVIDIA GPU cluster w Intel or AMD host, NVIDIA Tesla (ECC) GPU + Infiniband
- Upgrade portion of 2019 GPU cluster: RTX-2080 (non ECC) to Tesla GPU (ECC)
- New AMD CPU+ AMD GPU accelerated cluster + Infiniband

- Key NP performance metrics
  - mixed half-double or mixed single-double precision AMG clover
  - batched zgemm complex matrix multiplies
- CPU/KNL - JLab (16p & 18p KNL) + USQCD/CPU+KNK - large capacities
  - low cost performance
  - will continue to evaluate

Hardware options - FY20 (GPU)

- Conventional cluster w Intel Xeon processors (SkyLake) + Infiniband
- Conventional cluster w AMD EPYC 2 processors (Rome) + Infiniband
- New NVIDIA GPU cluster w Intel or AMD host, NVIDIA Tesla (ECC) GPU + Infiniband
- Upgrade portion of 2019 GPU cluster: RTX-2080 (non ECC) to Tesla GPU (ECC)
- New AMD CPU+ AMD GPU accelerated cluster + Infiniband

  - Strong demand for GPU (SPC + NP community), but with ECC
  - Pros/cons of alternatives - goal to increase TF-yr delivered on science
  - Testing has been held up by current MedCon
    - Just notified yesterday - new evaluation systems now available
    - AMD+AMD evaluation system (finally) enroute!

  - Goal is (still) to deploy a resource by originally scheduled Nov. 1

Next up is JLab LQCD operations

- Organizational structure:
  - Graham Heyes - new Head of Scientific Computing - overseas operations of LQCD
  - Bryan Hess - site manager (assumed position last year)

- LQCD & expt. computing already well integrated, so has been a smooth transition

- Presentation by Graham…

Many thanks to Chip Watson!

- Trained as an Expt. Nuclear Physicist - worked at BNL
- ~30 year career at JLab
- Headed development of accelerator control systems
- RGE & DGR arrived in 1999 just as Chip started LQCD group
- Chip principal lead in proposal for original SciDAC project (~2000)
  - Helped shape original USQCD software stack
- Keenly aware of computing trends & harnessing for computations
  - Early proponent/adopter - new systems, including GPUs (ARRA 2009) + KNLs
- Lead growth of JLab computing center and future directions
- Rare talent at combining technical skills + science advocacy in US communities

- Wish him well on his next adventures (not really retiring of course…)