Nuclear & Particle Physics Lattice QCD Computing Initiative

Jefferson Lab

Amitoj Singh

Thursday, April 20, 2023

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

Nuclear & Particle Physics Lattice Computing Initiative

- NPPLC Initiative started FY18
  - Mission - "*deploy & operate a significant dedicated computing resource for LQCD calculations.*"

- NPPLCI organization structure
  - Robert Edwards, PI
  - Graham Heyes, Deputy
  - Amitoj Singh, JLab LQCD site manager
  - Edwards & Singh, JLab LQCD site architects

- NPPLCI reporting structure
  - Report annually to Research Division of DOE Office of Nuclear Physics
  - First report (in Feb. 2020) included highlights of scientific program and future plans
  - DOE responds to the reports; however, not under an annual panel review process
  - For the first time NP will participate along with HEP in the annual DOE review this May 2023
  - NPPLCI renewal review planned in 2024; this is significant

- Operations
  - Take part in annual surveys & bi-weekly facility coordination

Nuclear & Particle Physics Lattice QCD Computing Initiative

- Model: dedicated facility - initiative purchases systems and operates them
  - Allows detailed tailoring of hardware & software development
  - Benefits compared to two-lab IC-based model (less constraints)
  - Amortize overhead costs with large system acquisitions

- Annual budget of ~$1M - about 50:50 labor:hardware
  - Hardware acquisition over FY boundaries for a bigger purchase
  - This FY23 we carried over hardware money from FY22 and will make one big purchase in FY23

- History of past procurements:
  - FY2018: upgraded Jlab's KNL resources (added to system from FY2016)
  - FY2019: upgraded Jlab's GPU resources (gamer-card system)
  - FY2021: new AMD CPU + MI-100 GPU system (generation before Frontier)

NP-LQCD & JLab-LQCD programs

- Computing - Computational Sciences & Technology Division (CST) [Amber Boehnlein]

- Theory - Computations and Nuclear Theory
  - LQCD@JLab started in ~FY01 - all funds directed through JLab/Theory
  - Theory and CST group have worked closely over two decades
  - CST division very responsive to the needs of Lattice QCD Theory
  - E.g., CST covers the cost of tape hardware, we only pay for media

- Systems - based on computational requirements for Lattice QCD Theory and USQCD
  - Funding directly from DOE Office of Nuclear Physics
  - Program is coordinated with USQCD Exec. Comm and Scientific Program Comm.
  - E.g., allocations taken from SPC

JLab & software development

- Heavy involvement since 2001 & SciDAC-1, then SciDAC-2 thru 4
- Currently, lead institution for ASCR/NP SciDAC-4 project
- Also, direct NP portion under Exascale Computing Project

- Leverage software development & local/commodity resources to efficiently utilize national resources
  - Our community well positioned for leadership systems
  - E.g., early adopters of NVIDIA GPUs and KNLs -> Titan and Cori-KNL

- Exascale systems - now AMD and Intel GPUs
  - Opportunity/necessity to diversify our codes

- LQCD commodity systems ↔ on-ramp to LCF systems + cycles for USQCD

Exascale

- Leadership + Exascale
  - Perlmutter - AMD CPU + NVIDIA GPU + Cray network
  - Frontier - AMD CPU + AMD GPU + Cray network
  - Aurora - Intel CPU + Intel GPU + Cray network
- Exascale Computing Project (LQCD is one of ~20 applications)
  - Performance metric - improve application performance by 50x
- LQCD code bases
  - QUDA refactorized - now “backend” support for NVIDIA, AMD, Intel GPUs
  - Grid support for all, but focused more on Intel
  - CPS & MILC - rely on QUDA
- NP & Chroma - refactored to support more GPU systems, but prioritized AMD
  - Production ready on AMD, NVIDIA, not quite ready for Intel
  - Software available on GitHub

New Hardware Acquisition

- Is a successor to KNLs.
- Considered two main metrics to quantify price-performance
  - Wilson dslash - proxy for inverters
  - ZGEMM based contractions considering possible swapping to main memory
- Request for Information (RFI) was released on March 21 requesting information on the following hardware configuration:
  - ~75-100 compute nodes
    - an x86 (such as AMD, Intel) CPU based
    - ARM (NVIDIA Grace Superchip) based
  - A few must haves in either server design are as follows:
    - Top tier HBM (High Bandwidth Memory).
    - Low latency, high bandwidth network fabric between the compute nodes.
- 9 responses to the RFI. Supermicro included a price quote which is helpful to set a baseline.

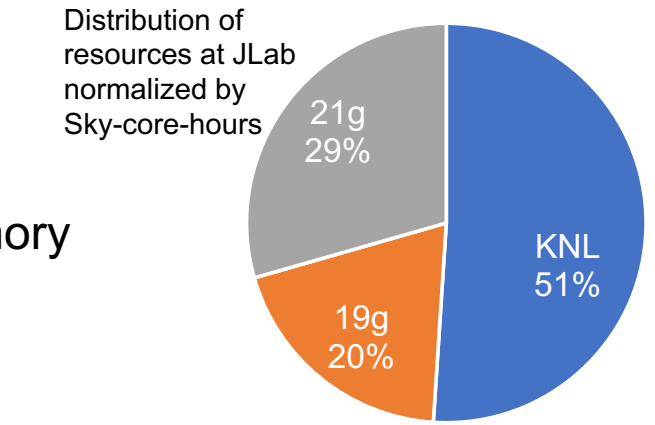

Acquisition Steps	Task Days	Due Date
Benchmarking on shortlisted hardware configurations and configuration(s) finalized	-	6/16/2023
RFP approvals completed and released for bids.	31	7/17/2023
RFP Responses due	30	8/16/2023
PO Awarded	44	9/29/2023
Delivery of equipment	120	1/27/2024