ECP Application Development

Andrew Siegel, ANL

HPC User Forum

Argonne National Laboratory

Sept. 10, 2019

Exascale Computing Project: Application Development

Goal: Ensure that exascale hardware impacts DOE science/engineering mission

Approach: Significant investment in scientific applications well in advance of exascale machines

- ~~Code Porting~~
- Algorithmic Restructuring
- New Numerical Approaches
- Alternate choice of Physical Models

Hardware has significant impact on all aspects of simulation strategy

Portfolio of ECP Applications

Application Categories	Number of Projects
Chemistry and Materials	**6**
Energy (generation)	**5**
Earth and Space Sciences	**5**
Data Analytics and Optimization	**4**
National Security	**4**

24 Domain Science/Engineering Simulation Projects

50+ separate codes

2/3 C ; 1/3 Fortran

Most pure MPI, or MPI+OpenMP at outset

Well defined, evolving dependencies on ECP software technology projects

ECP EXASCALE COMPUTING PROJECT

What defines an application project?

1. Scientific or Engineering exascale *challenge problem*.
2. Detailed completion criteria for (1) on exascale platform
3. A Figure of Merit (FOM) > 50 for project success

New physics capabilities – Not just faster/bigger version of existing codes

Quick Flyover of all 21 non-NNSA AD Application Projects

Energy Applications

ExaWind: Turbine Wind Plant Efficiency

Harden wind plant design and layout against energy loss susceptibility; higher penetration of wind energy

Lead: NREL
DOE EERE

ExaSMR: Design and Commercialization of Small Modular Reactors

Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior

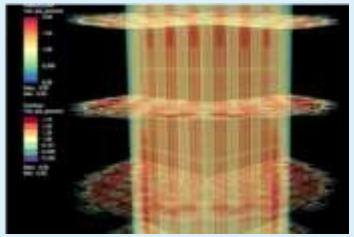

Lead: ORNL
DOE NE

WarpX: Plasma Wakefield Accelerator Design

Virtual design of 100-stage 1 TeV collider; dramatically cut accelerator size and design cost

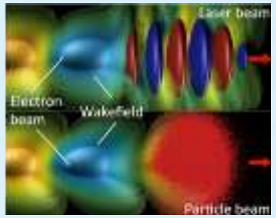

Lead: LBNL
DOE HEP

MFIX-Exa: Scale-up of Clean Fossil Fuel Combustion

Commercial-scale demo of transformational energy technologies - curbing CO_2 emissions at fossil fuel power plants by 2030

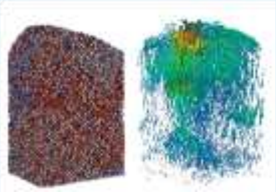

Lead: NETL
DOE EERE

Combustion-PELE: High-Efficiency, Low-Emission Combustion Engine Design

Reduction or elimination of current cut-and-try approaches for combustion system design

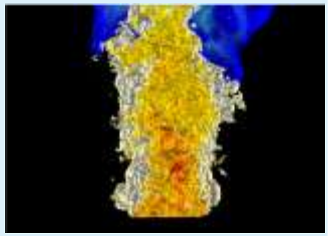

Lead: SNL
DOE BES, EERE

WDMApp: High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas

Prepare for ITER experiments and increase ROI of validation data and understanding; prepare for beyond-ITER devices

Lead: PPPL
DOE FES

Chemistry and Materials Applications

QMCPACK: Find, Predict, Control Materials & Properties at Quantum Level

Design and optimize next-generation materials from first principles with predictive accuracy

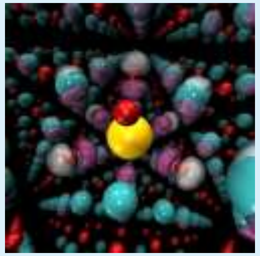

Lead: ORNL
DOE BES

ExaAM: Additive Manufacturing of Qualifiable Metal Parts

Accelerate the widespread adoption of AM by enabling routine fabrication of qualifiable metal parts

Lead: ORNL
DOE NNSA / EERE

LatticeQCD: Validate Fundamental Laws of Nature

Correct light quark masses; properties of light nuclei from first principles; <1% uncertainty in simple quantities

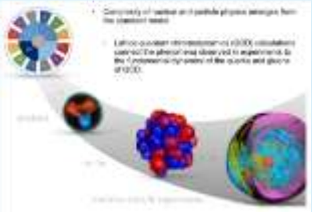

Lead: FNAL
DOE NP, HEP

NWChemEx: Catalytic Conversion of Biomass-Derived Alcohols

Develop new optimal catalysts while changing the current design processes that remain costly, time consuming, and dominated by trial-and-error

Lead: PNNL
DOE BER, BES

GAMESS: Biofuel Catalyst Design

Design more robust and selective catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower

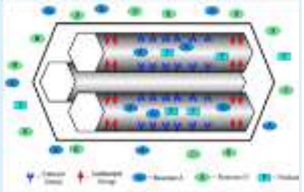

Lead: Ames
DOE BES

EXAALT: Materials for Extreme Environments

Simultaneously address time, length, and accuracy requirements for predictive microstructural evolution of materials

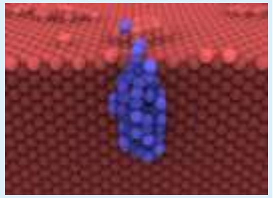

Lead: LANL
DOE BES, FES, NE

Earth and Space Science Applications

E3SM-MMF: Accurate Regional Impact Assessment in Earth Systems

Forecast water resources and severe weather with increased confidence; address food supply changes

Lead: SNL
DOE BER

Subsurface: Carbon Capture, Fossil Fuel Extraction, Waste Disposal

Reliably guide safe long-term consequential decisions about storage, sequestration, and exploration

Lead: LBNL
DOE BES, EERE, FE, NE

ExaStar: Demystify Origin of Chemical Elements

What is the origin of the elements? Behavior of matter at extreme densities? Sources of gravity waves?

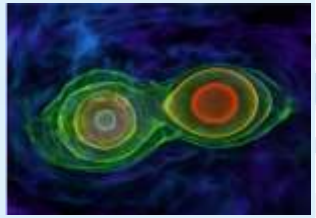

Lead: LBNL
DOE NP

ExaSky: Cosmological Probe of the Standard Model of Particle Physics

Unravel key unknowns in the dynamics of the Universe: dark energy, dark matter, and inflation

Lead: ANL
DOE HEP

EQSIM: Earthquake Hazard Risk Assessment

Replace conservative and costly earthquake retrofits with safe purpose-fit retrofits and designs

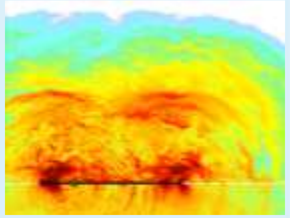

Lead: LBNL
DOE NNSA / NE, EERE

Data Analytics and Optimization Applications

ExaSGD: Reliable and Efficient Planning of the Power Grid

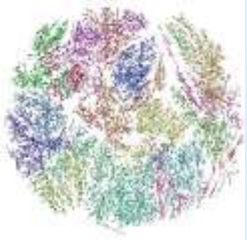

Optimize power grid planning, operation, control and improve reliability and efficiency

Lead: PNNL
DOE EDER, CESER, EERE

ExaFEL: Light Source-Enabled Analysis of Protein and Molecular Structure and Design

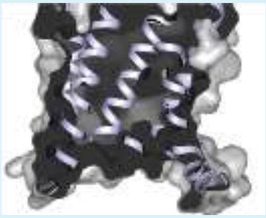

Process data without beam time loss; determine nanoparticle size & shape changes; engineer functional properties in biology and material science

Lead: SLAC
DOE BES

ExaBiome: Metagenomics for Analysis of Biogeochemical Cycles

Discover knowledge useful for environmental remediation and the manufacture of novel chemicals and medicines

Lead: LBNL
DOE BER

CANDLE: Accelerate and Translate Cancer Research

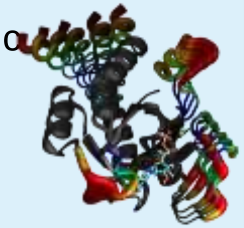

Develop predictive pre-clinical models & accelerate diagnostic and targeted therapy thru predicting mechanisms of RAS/RAF driven cancers

Lead: ANL
NIH

Application Development Milestones

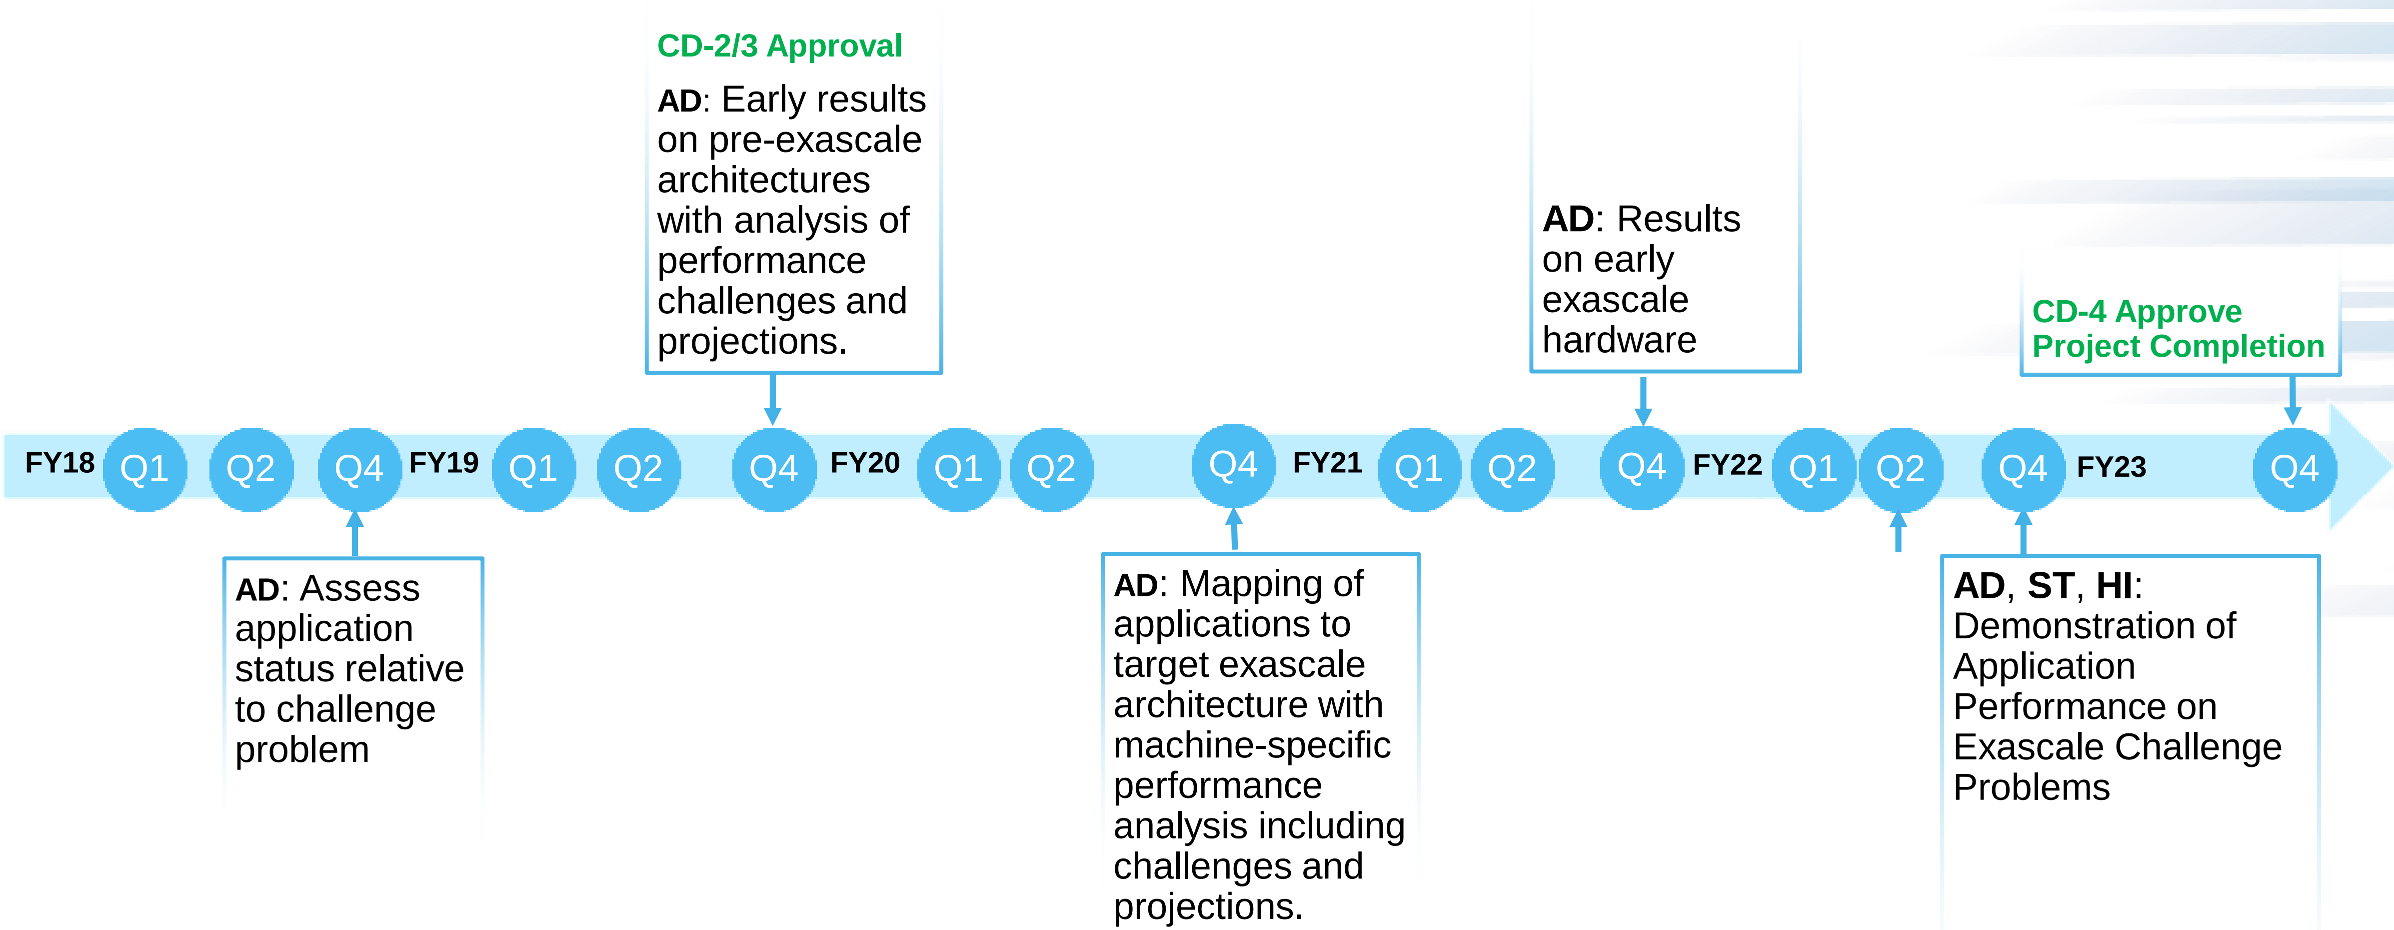

AD: Assess application status relative to challenge problem

CD-2/3 Approval

AD: Early results on pre-exascale architectures with analysis of performance challenges and projections.

AD: Mapping of applications to target exascale architecture with machine-specific performance analysis including challenges and projections.

AD: Results on early exascale hardware

AD, **ST**, **HI**: Demonstration of Application Performance on Exascale Challenge Problems

CD-4 Approve Project Completion

Baseline Platforms

O (10PF)

Current ECP Focus

CROSSROADS

O (100PF)

ECP Target Exascale Platforms

FRONTIER

O (1EF)

Current Figure of Merit Improvements on Summit/Sierra

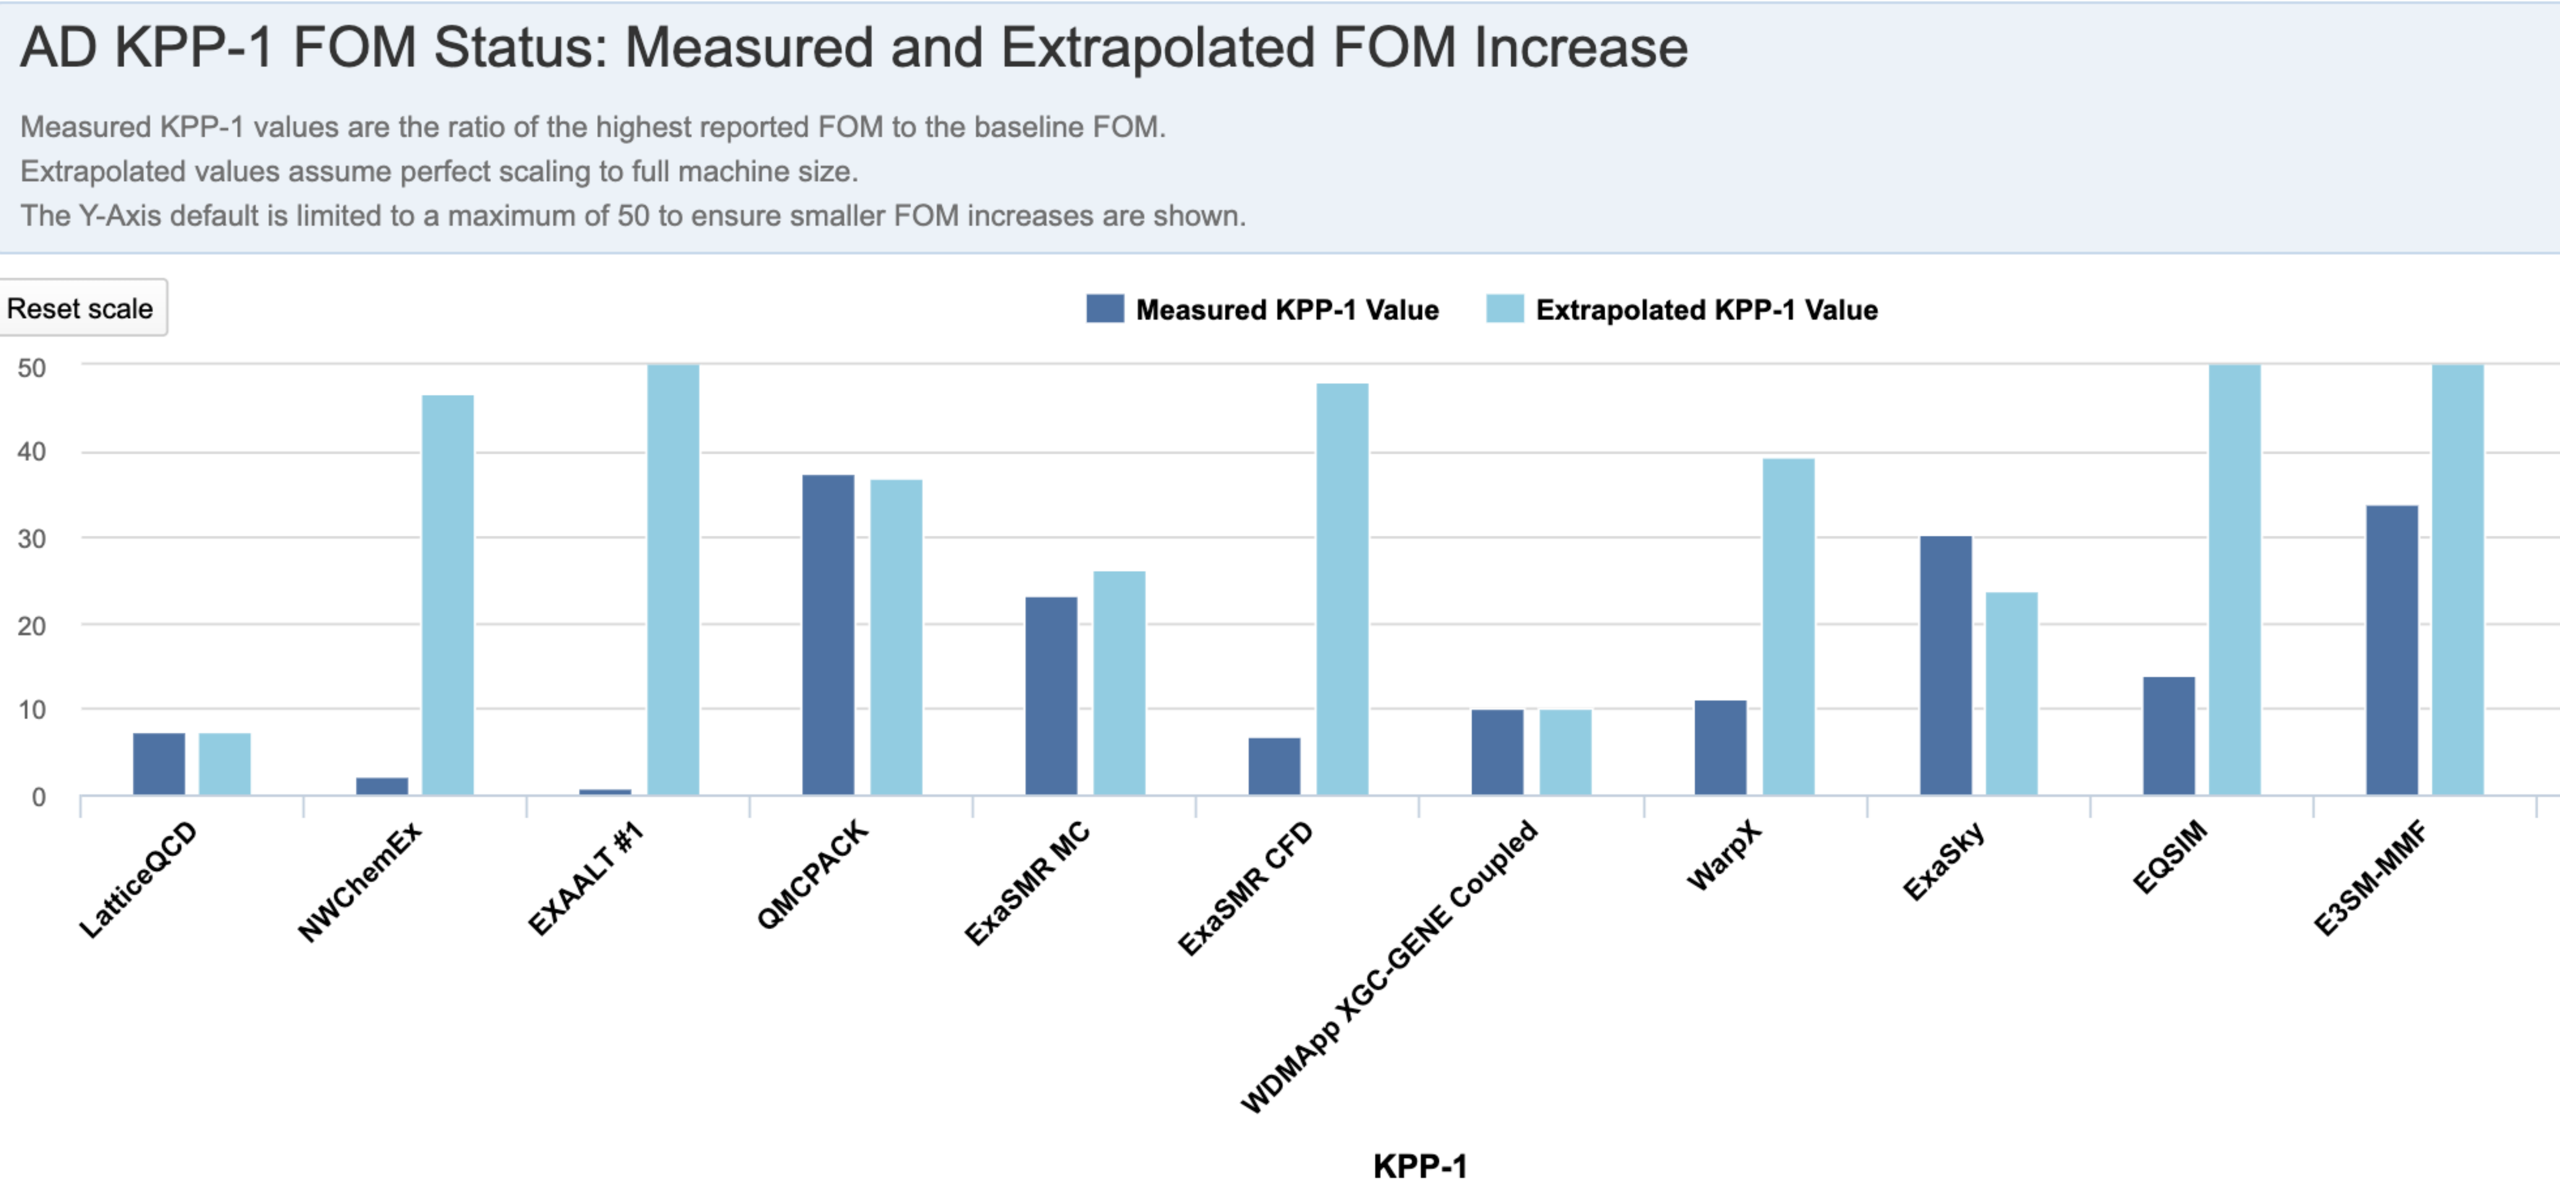

AD KPP-1 FOM Status: Measured and Extrapolated FOM Increase

Measured KPP-1 values are the ratio of the highest reported FOM to the baseline FOM.
Extrapolated values assume perfect scaling to full machine size.
The Y-Axis default is limited to a maximum of 50 to ensure smaller FOM increases are shown.

Applications Face Common Challenges

1) Flat performance profiles
2) Strong Scaling
3) Understanding/analyzing accelerator performance
4) Choice of programming model
5) Selecting mathematical models that fit architecture
6) Software dependencies

Strong scaling on modern high throughput cores

- High throughput processors do not perform near their peak in starvation limit
  - High value of $n_{1/2}$
  - Require abundance of fine-grained tasks that are efficiently scheduled on available resources
  - Otten *et al. e.g.* demonstrate that, on Titan, certain problems can run faster by exploiting the additional granularity afforded through the all-CPU model rather than using the highly-tuned GPU code (albeit with a 2.5 increase in power)

- Important because work is linear in time, so 50x has major performance impact

Fast reductions is another key component of strong scaling

- Conjugate Gradient
  - If vector reductions performed in software
    - η=.5 n/P ≥ 8500–12000 for P = 10^6–10^9
  - If vector reductions performed in hardware
    - η=.5 n/P ≥ 1200 for P = 10^6–10^9 !

- Multigrid
  - η=.5 n/P ~ 10000-20000 on machine like BG/Q
  - 2-4 times faster if hardware support for addition prefix ops

  - Bottom line: enables same simulation to run faster

Exemplar: Nuclear Engineering (ExaSMR)
Steve Hamilton (ORNL)
Approach: Monte Carlo Method

- Instead of solving equation, simulate individual neutrons directly
- Use known probability distributions for events (distance to collision, reaction, etc.)
- Count (or “tally”) the number of events that occur
- Simulating many (think millions+) particles gives average behavior

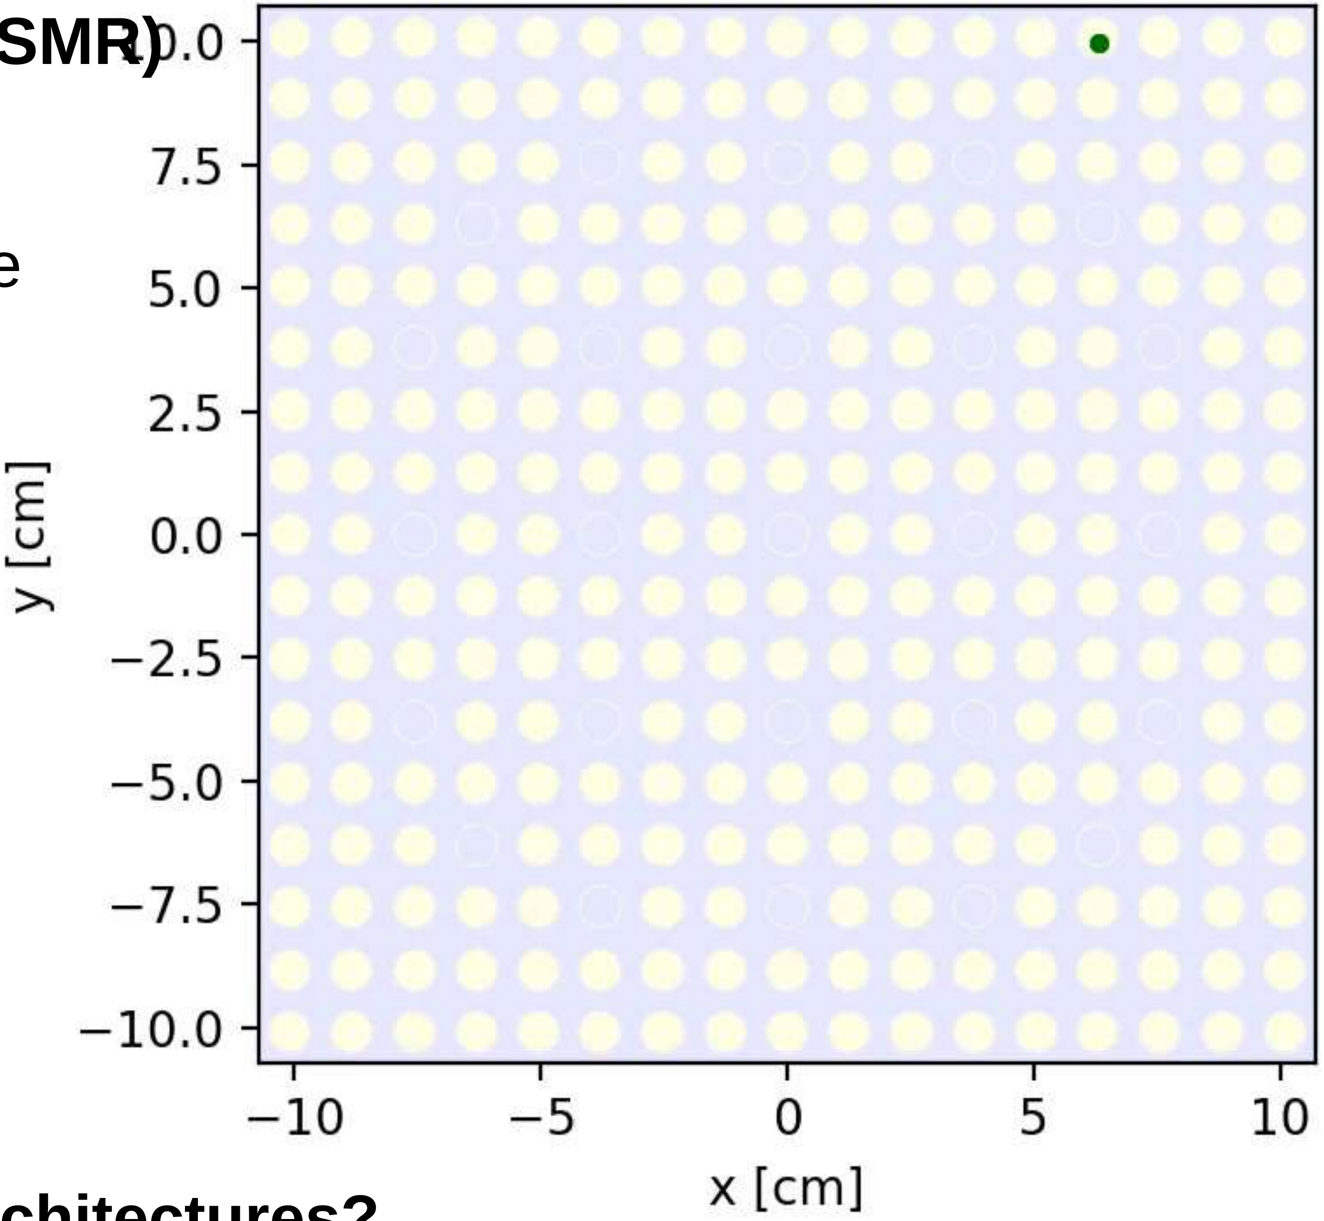

Why is this hard on accelerator architectures?

History-based Algorithm

One particle at a time

```
for each particle do
  while particle is alive
      calculate next interaction
  endwhile
endfor
```

Entire life of a particle *history*

Thread divergence: not a natural fit for GPUs

Event-based Algorithm

- Do one step at a time
- Sort by event type
- Process as SIMD

```
Get vector of particles
while any particle alive do
    for each event type do
        for particle ∈ event queue do
            Process event
        end for
    end for
end while
```

Data-level parallelism?

Algorithmic mapping to hardware – neutron particle transport

- Reduce thread divergence – change from history- to event-based algorithm
- Flatten algorithms to reduce kernel size; smaller kernels = higher occupancy
- Partition events based on fuel and non-fuel regions
- Take advantage of other architectural improvements

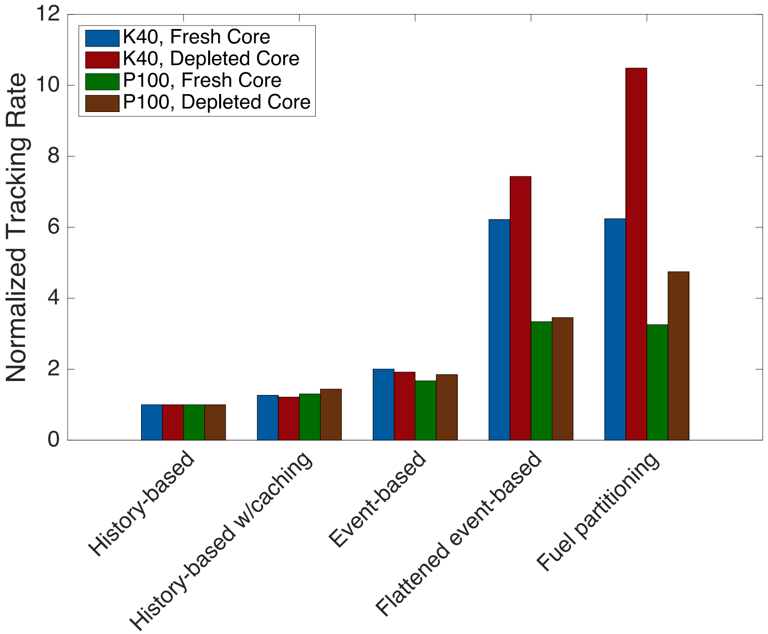

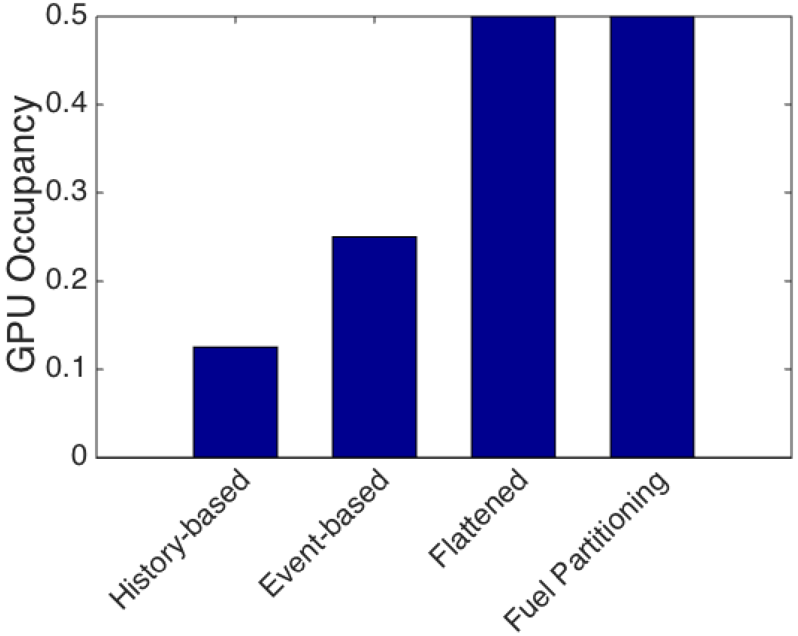

Exemplar: SNAP Potential (Danny Perez, LANL)

- Machine-learned MD potential that seeks for quantum-chemistry accuracy
- Neighbors of each atom are mapped onto unit sphere in 4D

$$(q_0, q, f) = \left(q_0^{max}\, r/r_{cut}, \cos^{-1}(z/r), \tan^{-1}(y/x)\right)$$

- Density around each atom is expanded in a basis of **4D hyperspherical harmonics**
- Bispectrum components of the 4D hyperspherical harmonic expansion are used as the geometric descriptors of the local environment
  - Preserves universal physical symmetries
  - Invariant to rotation, translation, permutation
  - Size-consistent
- SNAP uses linear regression to fit coefficients to DFT data

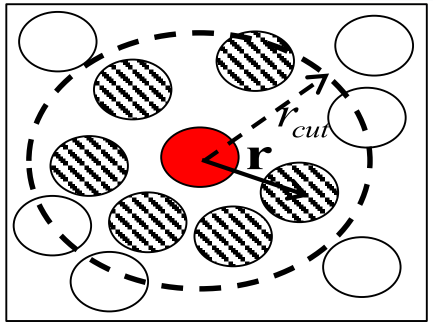

$$u^j_{m,m'} = U^j_{m,m'}(0,0,0) + \sum_{r_{ii'} < R_{cut}} f_c(r_{ii'}) w_i U^j_{m,m'}(\theta_0, \theta, \phi)$$

$$B_{j_1,j_2,j} = \sum_{m_1,m'_1=-j_1}^{j_1} \sum_{m_2,m'_2=-j_2}^{j_2} \sum_{m,m'=-j}^{j} (u^j_{m,m'})^* H^{jmm'}_{\substack{j_1 m_1 m'_1 \\ j_2 m_2 m'_2}} u^{j_1}_{m_1,m'_1} u^{j_2}_{m_2,m'_2}$$

SNAP GPU Performance Over Time

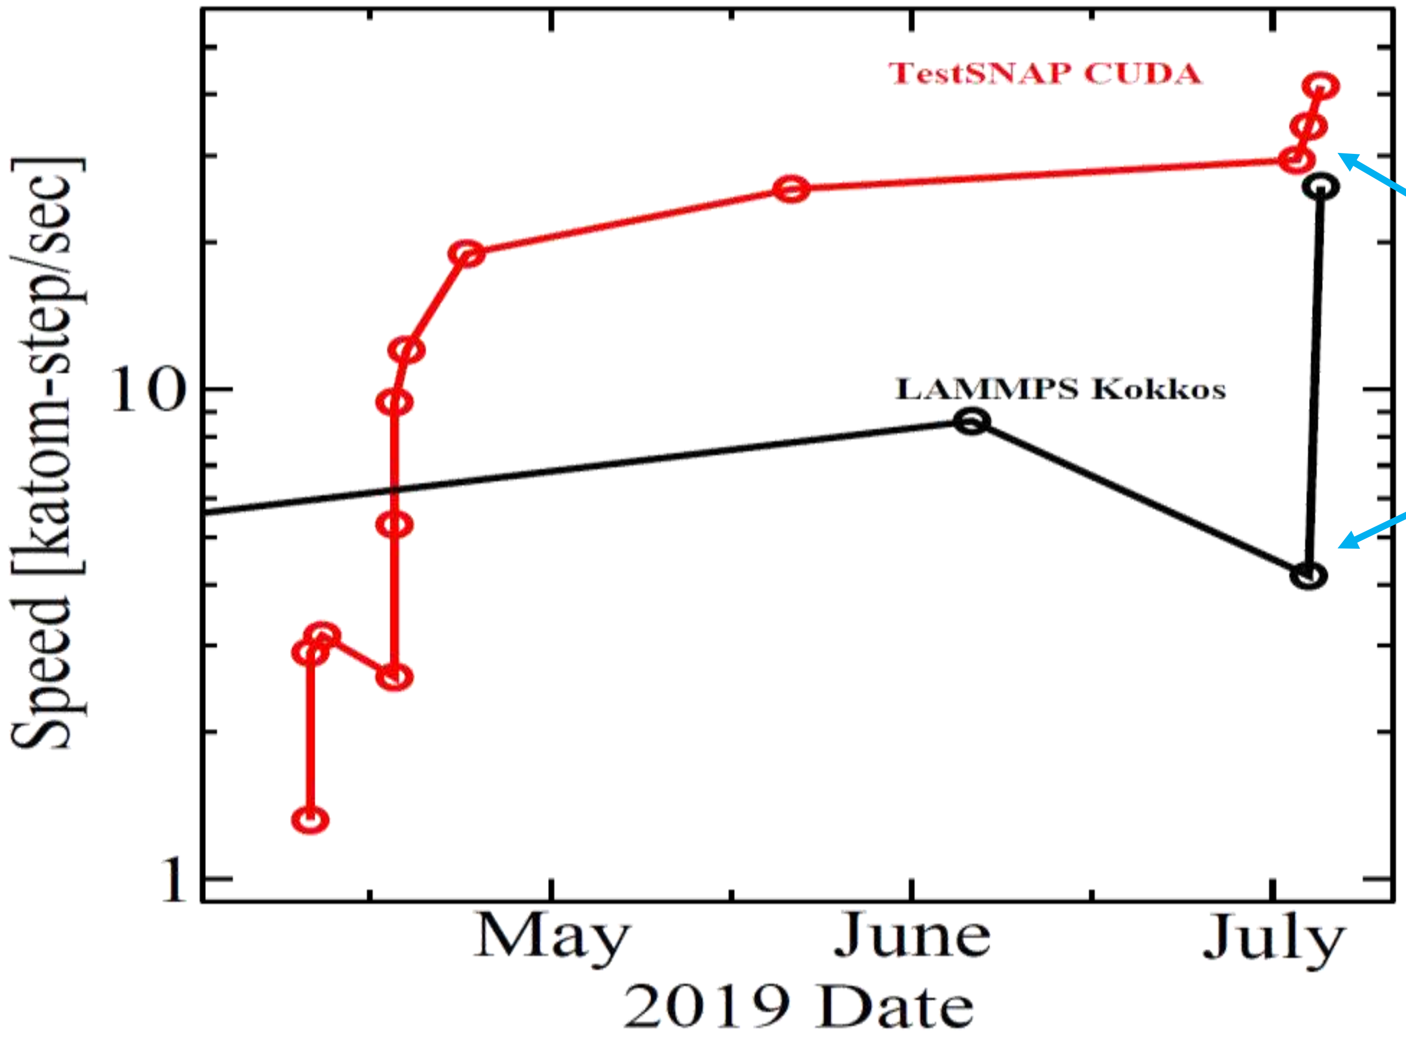

OLCF GPU Hackathon

SNAP Performance Improvements

- Aidan Thompson (Sandia) took the SNAP CPU code out of LAMMPS → **TestSNAP** stand-alone (realistic) force kernel, includes correctness check
- Idea from Nick Lubbers (LANL) →Aidan made algorithmic improvements that reduced FLOP count and eliminated some intermediate storage → ~2x speedup on CPUs
- Aidan reduced memory use by collapsing multidimensional arrays into compact lists
- Rahul Gayatri (NERSC):

1. broke up the one monster kernel into many smaller kernels, reduces register pressure and allows tailoring launch parameters for each kernel, but blows up the memory
2. inverted loops and changed data layouts to improve memory access

- Also had help from Sarah Anderson (Cray) and Evan Weinberg (NVIDIA)
- These improvements were ported to Kokkos SNAP in LAMMPS by Stan Moore

EXAALT FOM/KPP Projection for Summit

- Mira (IBM BG/Q) FOM baseline: 0.182 Katoms-steps/s/node * 49152 Mira nodes
- 2018 LAMMPS performance on Summit: 33.7 Katom-steps/s/node * 4608 Summit nodes: projected **17.4x faster than Mira baseline**
- New LAMMPS performance on Summit: 175.1 Katom-steps/s/node * 4608 Summit nodes: projected **90.2x faster than Mira baseline**
- Recently ported energy minimization in LAMMPS to Kokkos, which is needed by ParSplice
- Danny Perez (LANL) planning to validate these projections with large-scale Summit run soon

Overall …

- ECP is a very difficult project with many moving parts: specialized node architectures, system software, programming models, application level libraries, etc. enabling ambitious science and performance goals.

- Early adoption of intermediate (100PF) systems, test hardware, and hardware simulators critical to lowering risk by enabling progress tracking and early identification of issues.

- Surprisingly good progress to date, need to continue to push early adoption of exascale-type hardware, ensure proper balance of domain expertise and performance engineering. Facilities engagement programs are critical to achieving this.