Lattice QCD Application Development within the US DOE Exascale Computation Project

Carleton DeTar

University of Utah

EXASCALE COMPUTING PROJECT

Exascale Computing Project Goals

- US Department of Energy
  - Office of Science (DOE-SC)
  - National Nuclear Security Administration (NNSA)
- Develop computers, software, algorithms providing 50 X the computing power of today's fastest machines
- Initial "advanced architecture": 2021
  - At least one – perhaps two
  - Procurement process is separate from ECP
- "Capable exascale" computing based on ECP R&D: 2023
- Note: other initiatives in Europe, Japan, China

ECP EXASCALE COMPUTING PROJECT

Key Technical Hardware Challenges

- Parallelism: 1000 X greater than in today's systems
- Memory, network, and storage efficiencies: consistent with increased computational rates and data movement requirements
- Reliability: the ability to adapt and recover from faults in systems of vastly greater size and complexity
- Energy consumption: Must be much reduced from today's standards
- Industry innovation required.

"Capable Exascale System"

- Not interested in peak flop rate!
- Delivers 50X the performance of today's 20 PF systems for a wide range of applications
- Consumes power in the range 20–30 MW
- Is sufficiently resilient (perceived fault rate: ≤1/week)
- Includes a software stack that supports a broad spectrum of applications and workloads

US DOE ECP Laboratory Partners

- Six major US DOE laboratories
- All with strong HPC experience

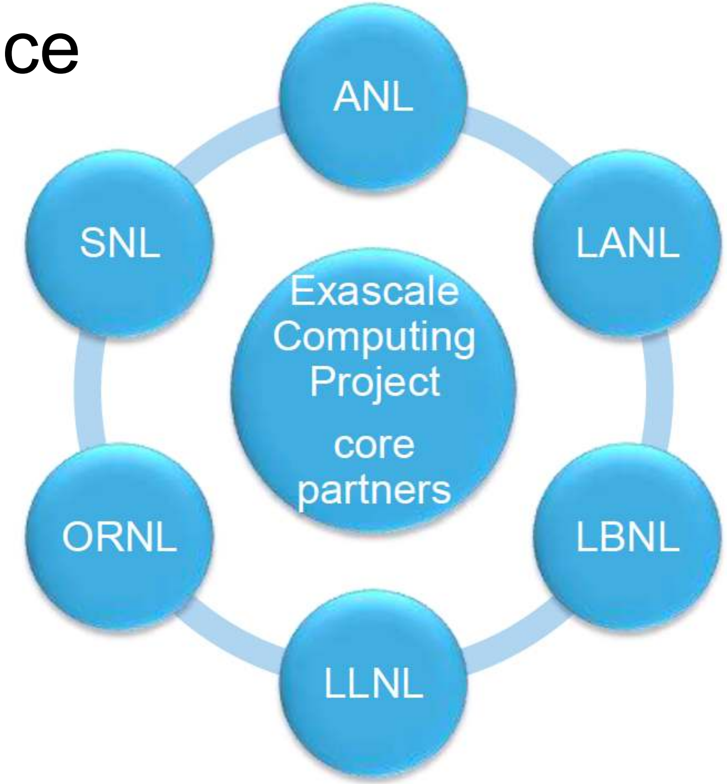

“Holistic co-design” and integration

- Application development (science and “mission” applications)
  - Some 30 teams with subjects ranging from cosmology to electrical power-grid design
- Software technology
  - Some 80 teams with subjects ranging from MPI to solvers to resilience
- Hardware technology
  - Vendor participation “PathForward” hardware contracts: AMD, Cray, Hewlett Packard Enterprise (HPE), IBM, Intel, and NVIDIA
  - > $430 M with $258 M from DOE in stages
- Exascale systems
  - Testbeds are being set up at national labs.

Leading Solution Providers

Capable Exascale System Applications Will Deliver Broad Coverage of 6 Strategic Pillars

National security

- Stockpile stewardship

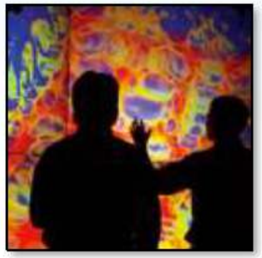

Energy security

- Turbine wind plant efficiency
- Design and commercialization of SMRs
- Nuclear fission and fusion reactor materials design
- Subsurface use for carbon capture, petro extraction, waste disposal
- High-efficiency, low-emission combustion engine and gas turbine design
- Carbon capture and sequestration scaleup
- Biofuel catalyst design

Economic security

- Additive manufacturing of qualifiable metal parts
- Urban planning
- Reliable and efficient planning of the power grid
- Seismic hazard risk assessment

Scientific discovery

- Cosmological probe of the standard model of particle physics
- Validate fundamental laws of nature
- Plasma wakefield accelerator design
- Light source-enabled analysis of protein and molecular structure and design
- Find, predict, and control materials and properties
- Predict and control stable ITER operational performance
- Demystify origin of chemical elements

Earth system

- Accurate regional impact assessments in Earth system models
- Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols
- Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

Health care

- Accelerate and translate cancer research

Some Exascale Science Goals for Lattice QCD

- Aid the search for physics beyond the Standard Model by comparing precise calculations with precise experimental results
  - B physics at lattice spacing smaller than 1/m_b
  - CP violation in the K system
  - “Non-standard” weak K decays (Xu Feng)
  - 10X reduction in uncertainties in weak matrix elements
- Precision study of the structure and interactions of light nuclei
- Calculations at physical quark masses with QED.

Lattice QCD

- Some 50 Lattice physicists. (Several postdocs and lab staff receive some salary support)
- Not just US lab members and academics: includes Edinburgh, NVIDIA, applied mathematicians
- Steering Committee
  - Paul Mackenzie (Chair and PI)
  - Rich Brower: Solvers (Multigrid, etc)
  - Norman Christ: Critical slowing down (HMC)
  - CD: Software (New data parallel API and implementation)
  - Robert Edwards: Wick contractions (nuclear matrix elements)

Solver research

- Communication reduction/avoidance is a key consideration at exascale
- Staggered fermion multigrid
  - Monday talk by Evan Weinberg
  - Monday talk by Kate Clark on GPU considerations
- Multigrid with deflation (Stathopoulos et al)
- Domain decomposition
- All mode averaging with eigenvector compression
  - Talk by Christoph Lehner

Mitigating Critical-slowing-down

- Topology freezing becomes an increasing concern at lattice spacings smaller than 0.05 fm
  – Poster by Rainer Sommer
- Not just global – perhaps slow local diffusion as well [McGlynn and Mawhinney, Phys Rev D 90, 074502 (2014)]
- We are exploring a variety of approaches
- Also studying the effect of topology freezing on observables

Contractions

- Wick contractions in light nuclear matrix elements present a huge combinatoric problem
- Do “edge” reductions

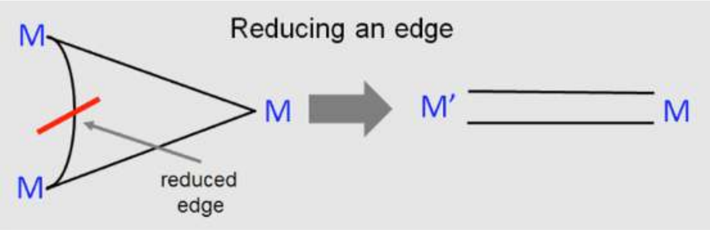

- Find unique set of reductions to maximize reuse
- Rearrange graph order to lower number of temps

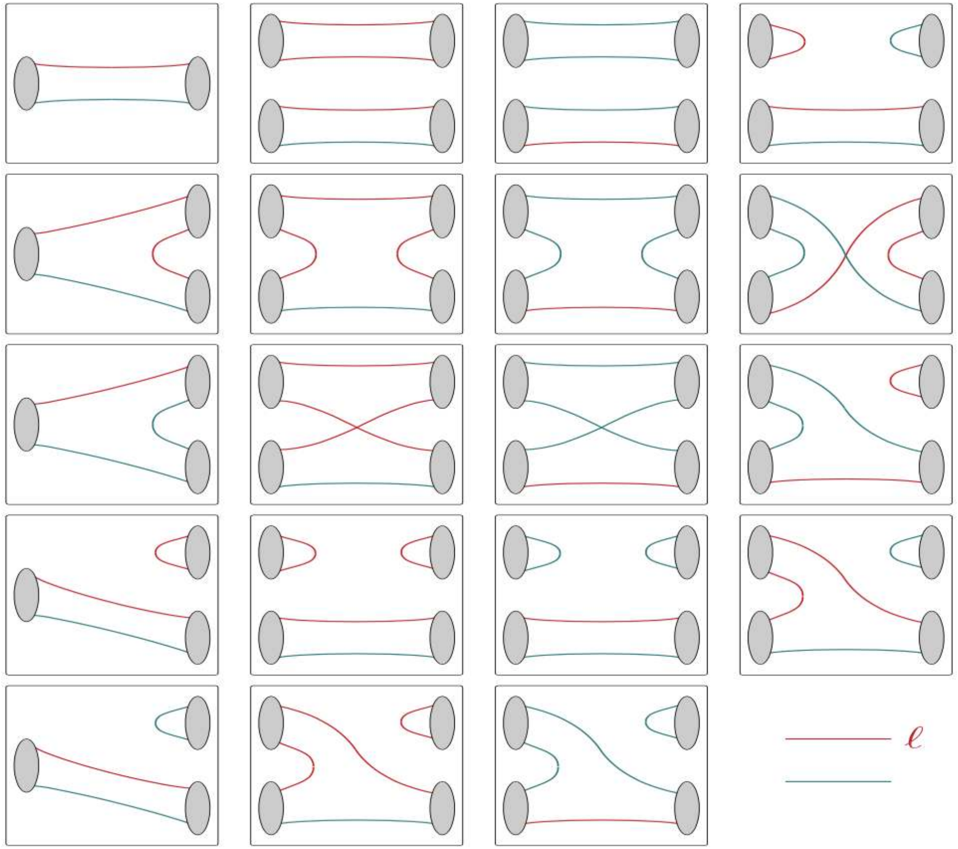

I=1/2 K*π arXiv:1406.4158

Software development goals

- Data parallel API to support ECP algorithms under development, e.g.
  - Multigrid
  - Domain decomposition
- Portable and performant software
  - Many-core to GPU portability is the main challenge
  - See talk by Meifeng Lin
- Extensible
- Enable rapid prototyping of new algorithms

Past USQCD-developed APIs

- SciDAC data parallel API: QDP/C, QDP++, QLUA
- Plus a variety of software packages with internal APIs
  - QUDA
  - CPS
  - MILC

Software activities

- List data parallel features needed to support new algorithms
- Conduct experiments with expression-template portability
- Start an assessment of semantic alternatives for data-parallel API
  - C++-11 (QDP++, Grid, etc)
  - QEX(Nim) (James Osborn, Xiao-Yong Jin)
    - https://nim-lang.org

EXASCALE COMPUTING PROJECT

Examples of semantics

Covariant shifts

C++/Grid

```
tmp = CovShiftForward(U[mu], mu, src);
```

Nim/QEX

```
let T = newTransporters(U, src, 1)
tmp = T[mu] ^* src
```

Examples of semantics

Fine to Coarse-grid projection

Subdivide a fine grid into blocks and use n basis vectors to project a fine-grid field onto a coarse-grid field.

C++/Grid

```
blockProject( coarseData, fineData, Basis );
```

Nim/QEX

```
let R = newRestrictor(coarseLattice, Basis)
coarseData := R ^* fineData
```

Examples of semantics

- Loop fusion (Still under discussion)

C++/Grid

```
   PARALLEL_FOR_LOOP
     for(s = 0; s < nsites; s++){
         p[s] = r[s] + beta * p[s];
          q[s] = p[s] + alpha * r[s];
 }
```

Macro for OpenMP pragma

Site loop goes below data parallel level

Expression template for single line

Nim/QEX

```
  fuse:
      p = r + beta * p
      q = p + alpha * r
```

Sociology of management

- DOE "projects" require a robust management structure
  - Atlassian management software: Jira, Confluence
  - Milestones with deliverables, deadlines, "epics", "stories", "sprints"
  - Monthly progress reports
- Academics are unaccustomed to corporate-style management
- Too much "management" risks hampering creativity
- Too little risks failing to reach goals

EXASCALE COMPUTING PROJECT

ECP Integration

- ECP Integration meeting Jan 2017 in Knoxville, Tennessee
- Brought together applications, software technology, hardware, and systems groups – some 450 participants
- “Match-making”
- Lattice QCD integration interests so far
  - Kokkos, Trilinos project
  - OpenMP and OpenACC design
  - MPI design
  - HDF5
  - Checkpointing

Measure of success

- The success of the entire DOE ECP program will be measured by the ability of most of the included applications to reach the goal of a 50X increase in computing capability.
- All improvements included: hardware, system design, software, algorithms
- Measured by application-specific Key Performance Parameters (KPP) and Figures of Merit (FOM)
- QCD FOM: Rate of generating lattices and making physics-related measurements for a defined problem using a fixed fraction of a “Leadership Class Facility” such as Argonne ALCF or Oak Ridge OLCF.
  - HISQ action
  - DWF action
- Look for a 50X increase in that rate.

EXASCALE COMPUTING PROJECT

Outlook

- The US Dept of Energy Exascale Computing Project is helping to accelerate the arrival of wide-spread exascale computing
- Coordination between hardware vendors, systems development teams, and application developers is crucial
- We hope our efforts contribute not just to the success of the US lattice QCD program, but to that of the entire world.