Lattice QCD on Perlmutter

Steven Gottlieb
Indiana University

GPUs for Science
NERSC
October 25, 2022

Introduction

- ✦ Lattice QCD
- ✦ Some of our accomplishments on Perlmutter
- ✦ Benchmarks
- ✦ Software development for GPUs
- ✦ Conclusions

Lattice QCD

- Quantum Chromodynamics (QCD) is a 50 year old quantum theory of the strong interaction
  - quarks (matter) and gluons (force carrier)
  - quarks carry quantum number called color (red, green, blue)
  - responsible for bound states: mesons and baryons
    - the bound states are color neutral objects
    - nuclear force is residual force color neutral protons and neutrons
    - Our mass comes from these interactions, not the Higgs field
- Ken Wilson developed lattice QCD to go beyond perturbation theory
  - continuum is replaced by a 4-dim grid of space-time points
  - quarks described by complex fields with 3 or 3X4 components
  - gluons 3X3 complex unitary matrices

Lattice QCD II

- ✦ The basic calculation is related to the Feynman path integral, but as we change time to imaginary, theory is more like a statistical mechanical partition function
- ✦ Numerical methods:
  - Monte Carlo
  - Sparse matrix solvers
  - molecular dynamics in a new simulation time evolution
- ✦ First step of the calculation is to create ensembles of gauge fields
  - these are like properly weighted paths in the path integral
  - we take averages over the gauge fields in the ensemble
  - larger ensembles gives better average of quantum fluctuations

Control of Systematic Errors

- To generate an ensemble, we must select several physical parameters:
  - lattice spacing (a) or gauge coupling β
    - actually set β in code and determine lattice spacing later
  - grid size $N_s^3 \times N_t$
  - sea quark masses (m_{ud}, m_s, m_c , sometimes $m_{ud} \to m_u, m_d$
- To control systematic errors, we must:
  - take continuum limit, i.e., $a \to 0$
  - take infinite volume limit
  - tune sea quark masses to reproduce known masses

Why We Use So Much Time

- ✦Controlling each systematic error requires investing more time
  - halving the lattice spacing increases time by about 2^6
  - doubling the linear size increases time by 2^4 (or 2^3 if not increasing time extent)
  - we can now tune the up and down quarks to be at their physical value
    - for many years that was too expensive
- ✦In addition to creating the costly gauge ensembles as a stochastic evolution of the system, we do costly calculations on the ensembles
  - The iterative solver takes much of the time
  - Once the ensemble is generated, these jobs can be run simultaneously
- ✦Generating gauge configurations favors high speed, while doing physics analysis on stored configurations can be done in parallel and is more of a capacity problem.

Some of our accomplishments

- I was late in asking the lattice QCD crowd for accomplishments, so this presentation is heavily weighted to work done by Carleton DeTar and myself for the Fermilab Lattice and MILC Collaborations.
- Chris Kelly did refer me to his recent talk at Lattice 2022, so I will also summarize that.
- DeTar and I are interested in decay of mesons containing a bottom quark. These mesons decay weakly and can help to determine elements of the CKM quark mixing matrix. I was creating new gauge configurations and he was running meson decay measurements.

Power of Perlmutter

- Started generating configuration in 2014, by mid-2018 it was half done.
- Goal was 6000 time units.
- Late summer 2021 we were able to resume running during Perlmutter early science period.
- Note the remarkable change in slope due to power of Perlmutter.

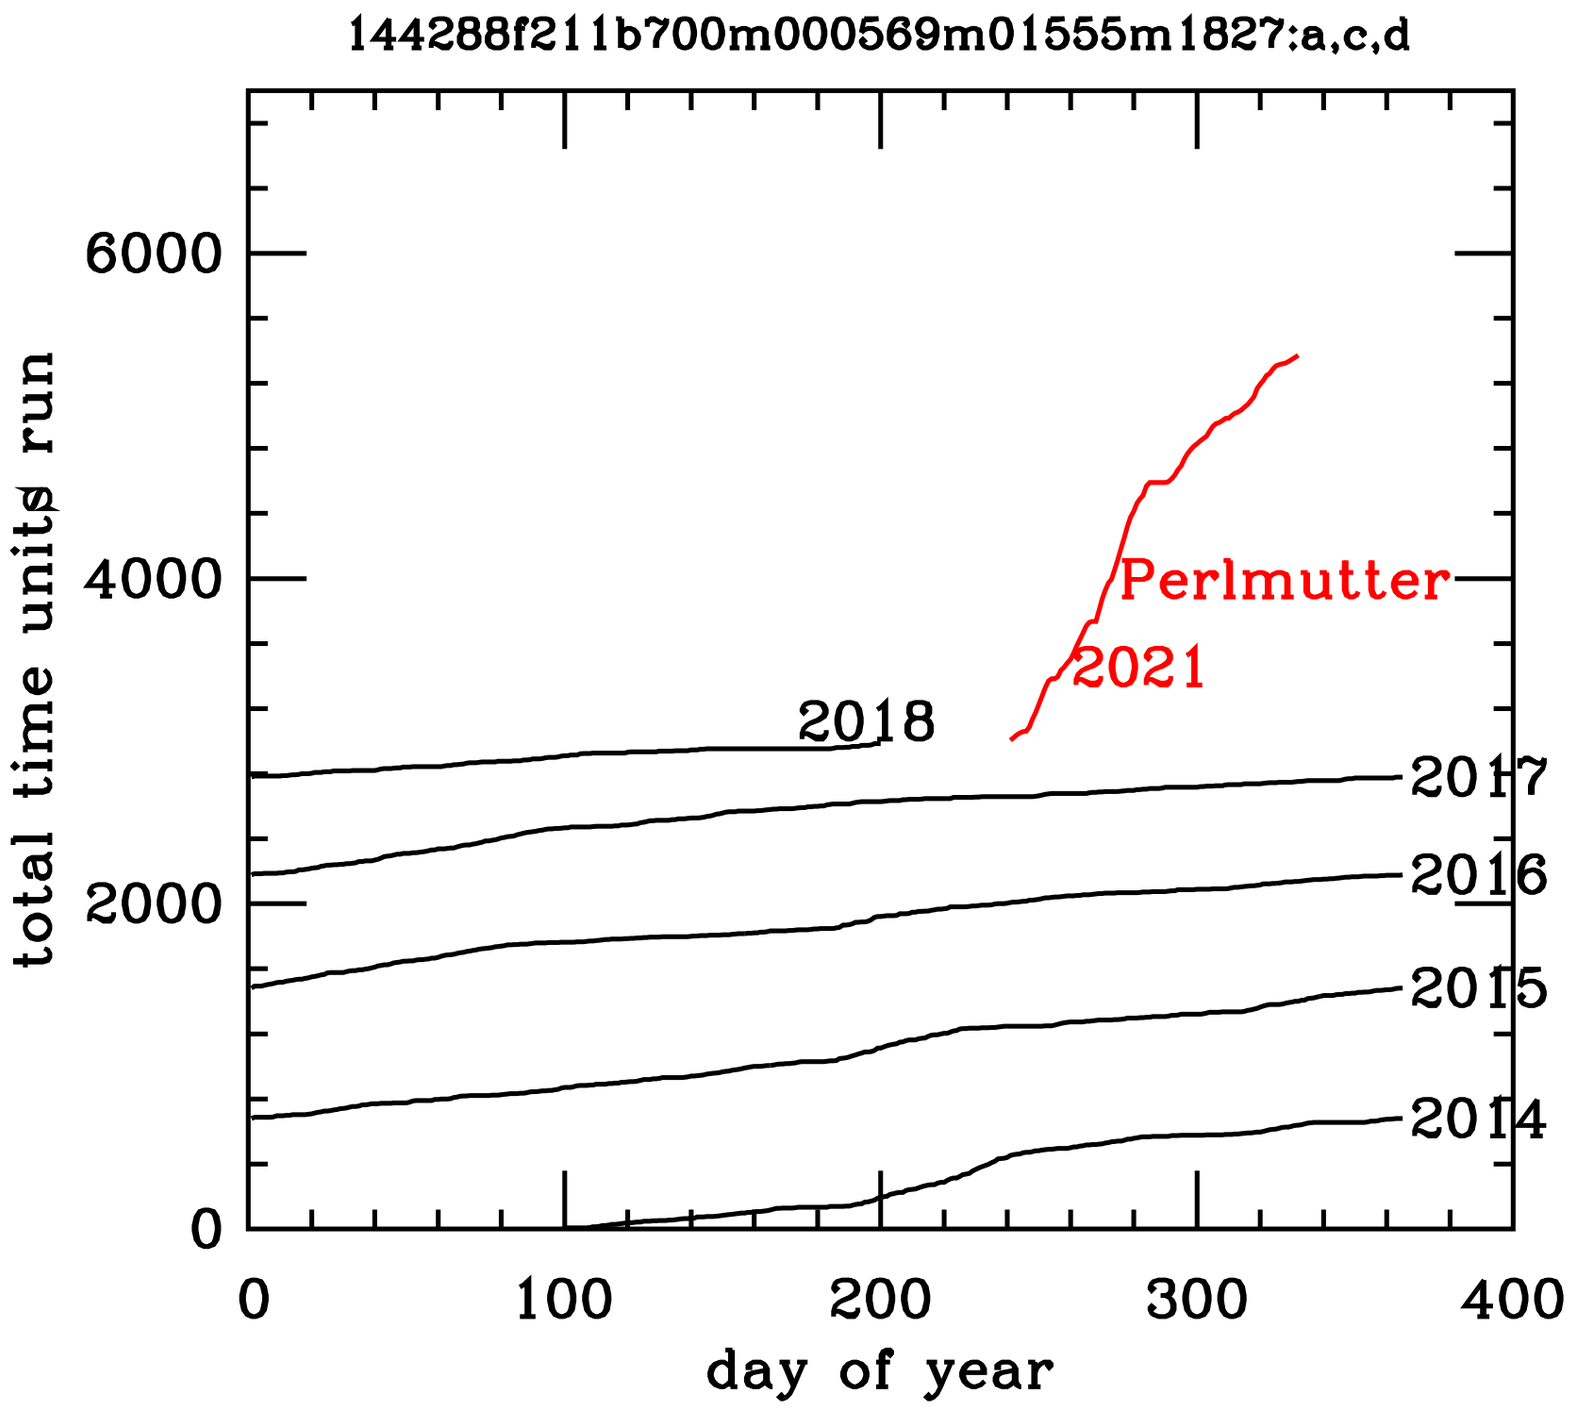

Power of Perlmutter II

- Previous graph prepared early Dec., 2021.
- Slope decreased as Perlmutter became much busier.
- For lattice QCD, need both a fast computer and allocation to use it frequently.
- We created 500 new configurations and analyzed about 50.

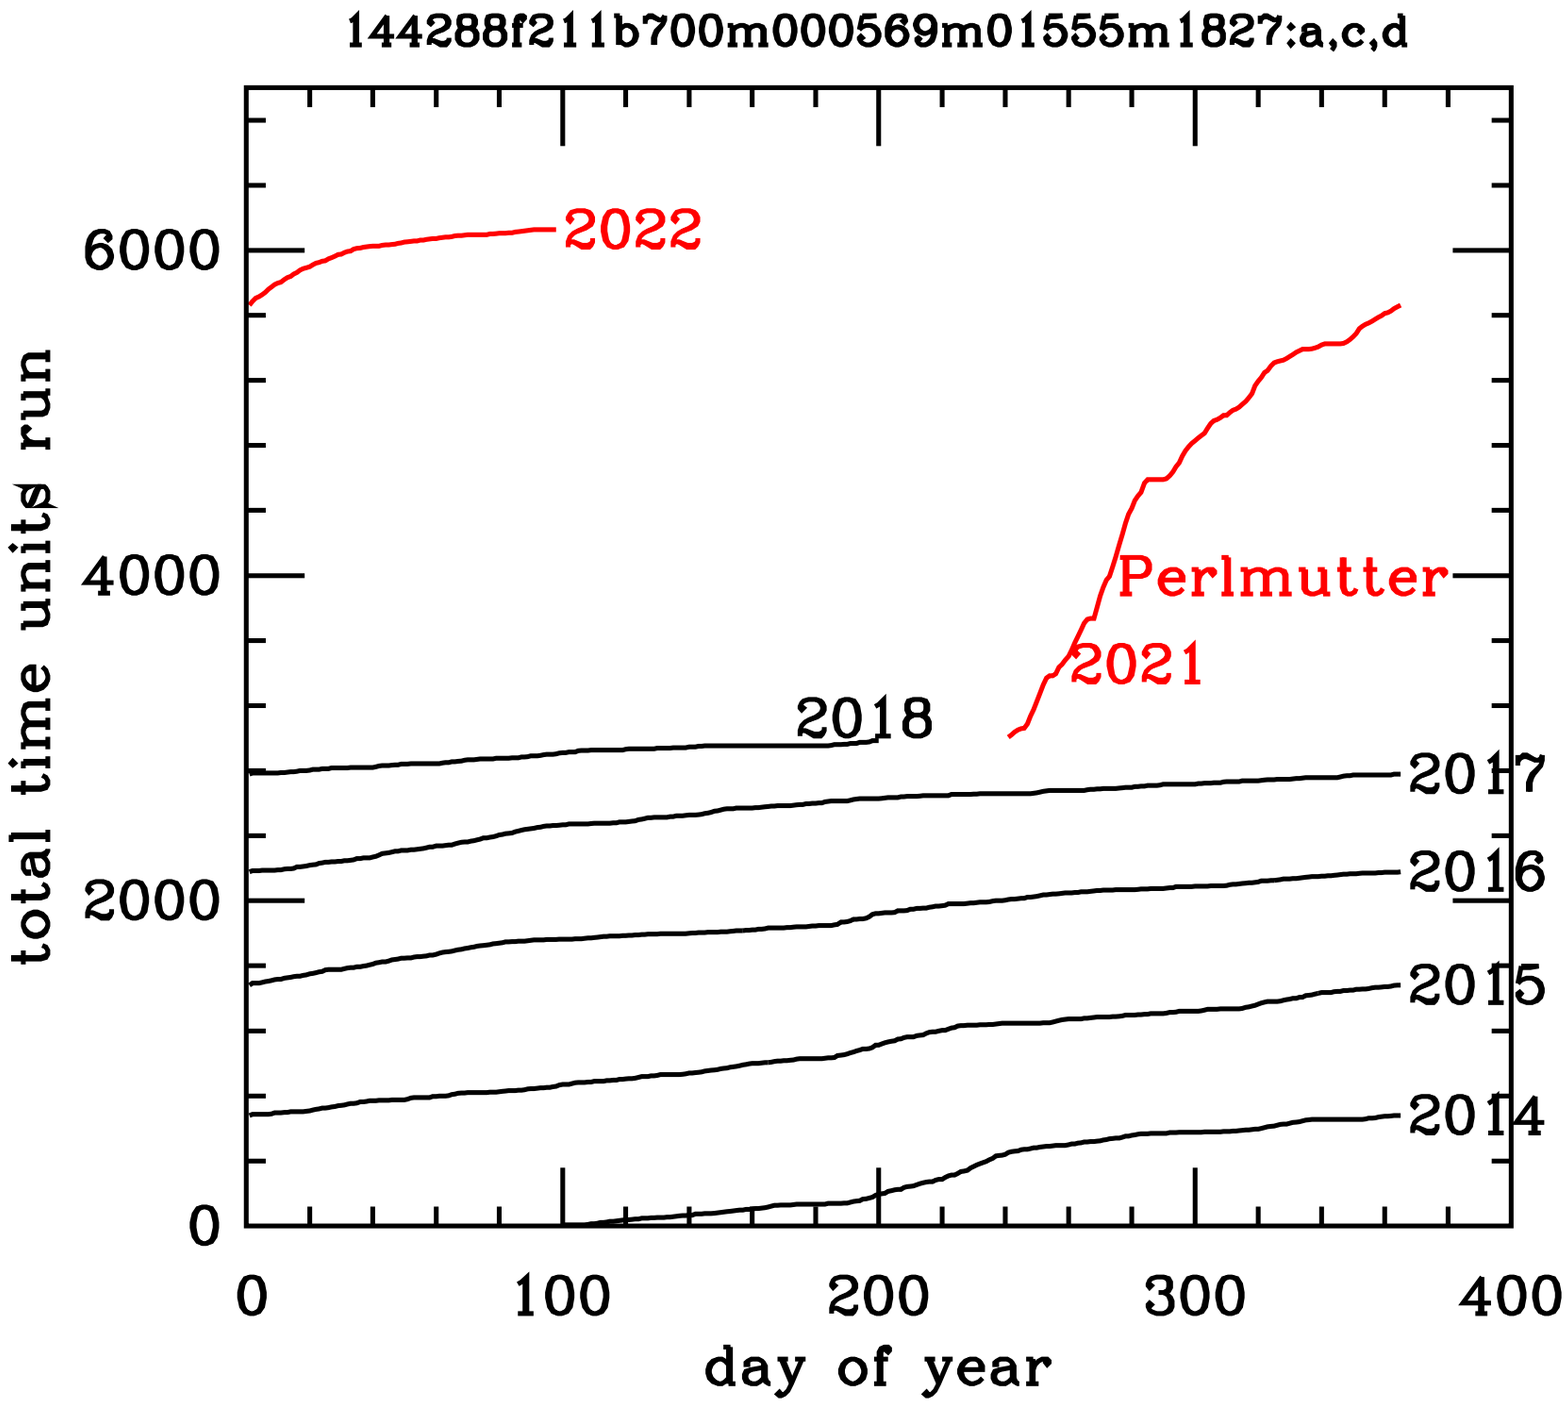

Power of Perlmutter III

- ✦ Jefferson Lab group uses QCD-JIT, QUDA, and Chroma for their work.
- ✦ Introduced new multigrid solver and GPU optimization to speed up configuration generation with Wilson-clover quarks.
- ✦ Job that took
  - 1,192 seconds on 256 Edison nodes
    - now takes only
  - 80 seconds on 32 Perlmutter nodes
- ✦ Combination of algorithmic and hardware improvements!

Power of Perlmutter IV

- RBC/UKQCD group has been studying muon anomalous magnetic moment and CP violation in $K \rightarrow \pi\pi$ decays.
- For first project, using 256 nodes to analyze configuration of size $96^3 \times 192$.
- For latter, running 32 node jobs to generate domain wall fermion ensembles with size $40^3 \times 64 \times 12$ and $48^3 \times 64 \times 12$.
- Using the smaller volume, a 5.9 hour job using the Slingshot 10 network was reduced to 4.7 hours after upgrade to Slingshot 11.

Cross Platform Comparison

- Table compares times to run a trajectory of length 2 and save the configuration on four different computers

Computer	nodes or cores	MPI ranks	generate (hr)	save (hr)	total (hr)
Edison	18432 cores	36864	7.10	0.24	7.34
Cori	1024	65536	4.34	0.98	5.32
Blue Waters	1536(?)	49152	8.45	0.39	8.84
Perlmutter	128	512	1.46	0.07	1.53

MILC Performance

- All running on 128 nodes, i.e., 512 GPUs on 144^3 X 288 configuration
- ON EACH NODE (RANK) 144 x 36 x 18 x 18
- Multimass CG: 285 GF/gpu single precision (mixed)
- Link smearing: 150 GF/gpu
- Gauge Force: 1.5-1.7 TF/gpu
- Fermion force: still don’t have flop count in code. Need to fix.

CG Volume Study

- local volume L^4
- early study without any compression
- generic single precision (no mixed)
- no binding script
- really need to redo these studies
- lattice cut in all four dimensions for ≥16 cpus

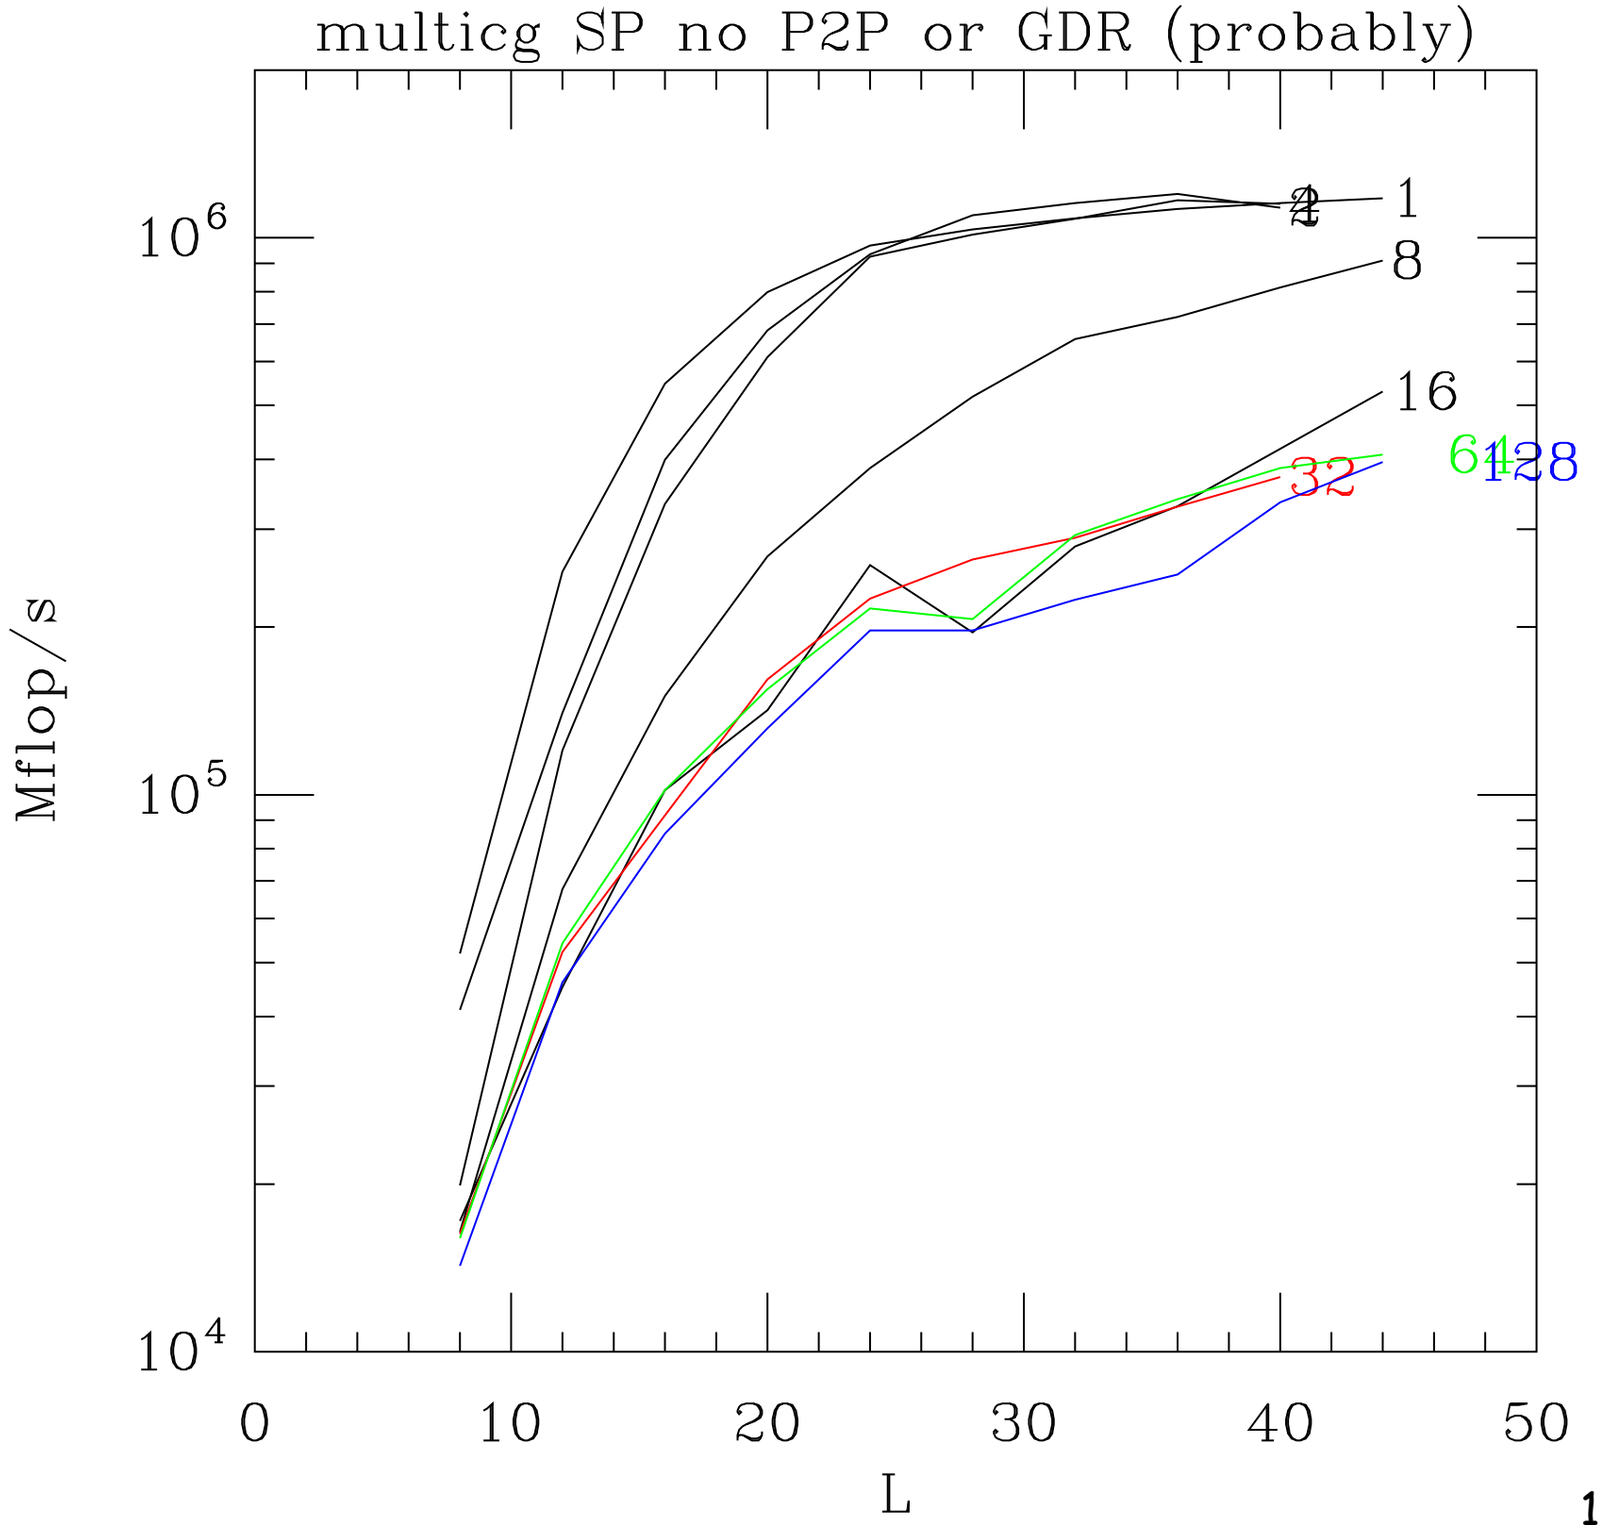

Gauge Force Volume Study

- similar limitations compared with previous plot

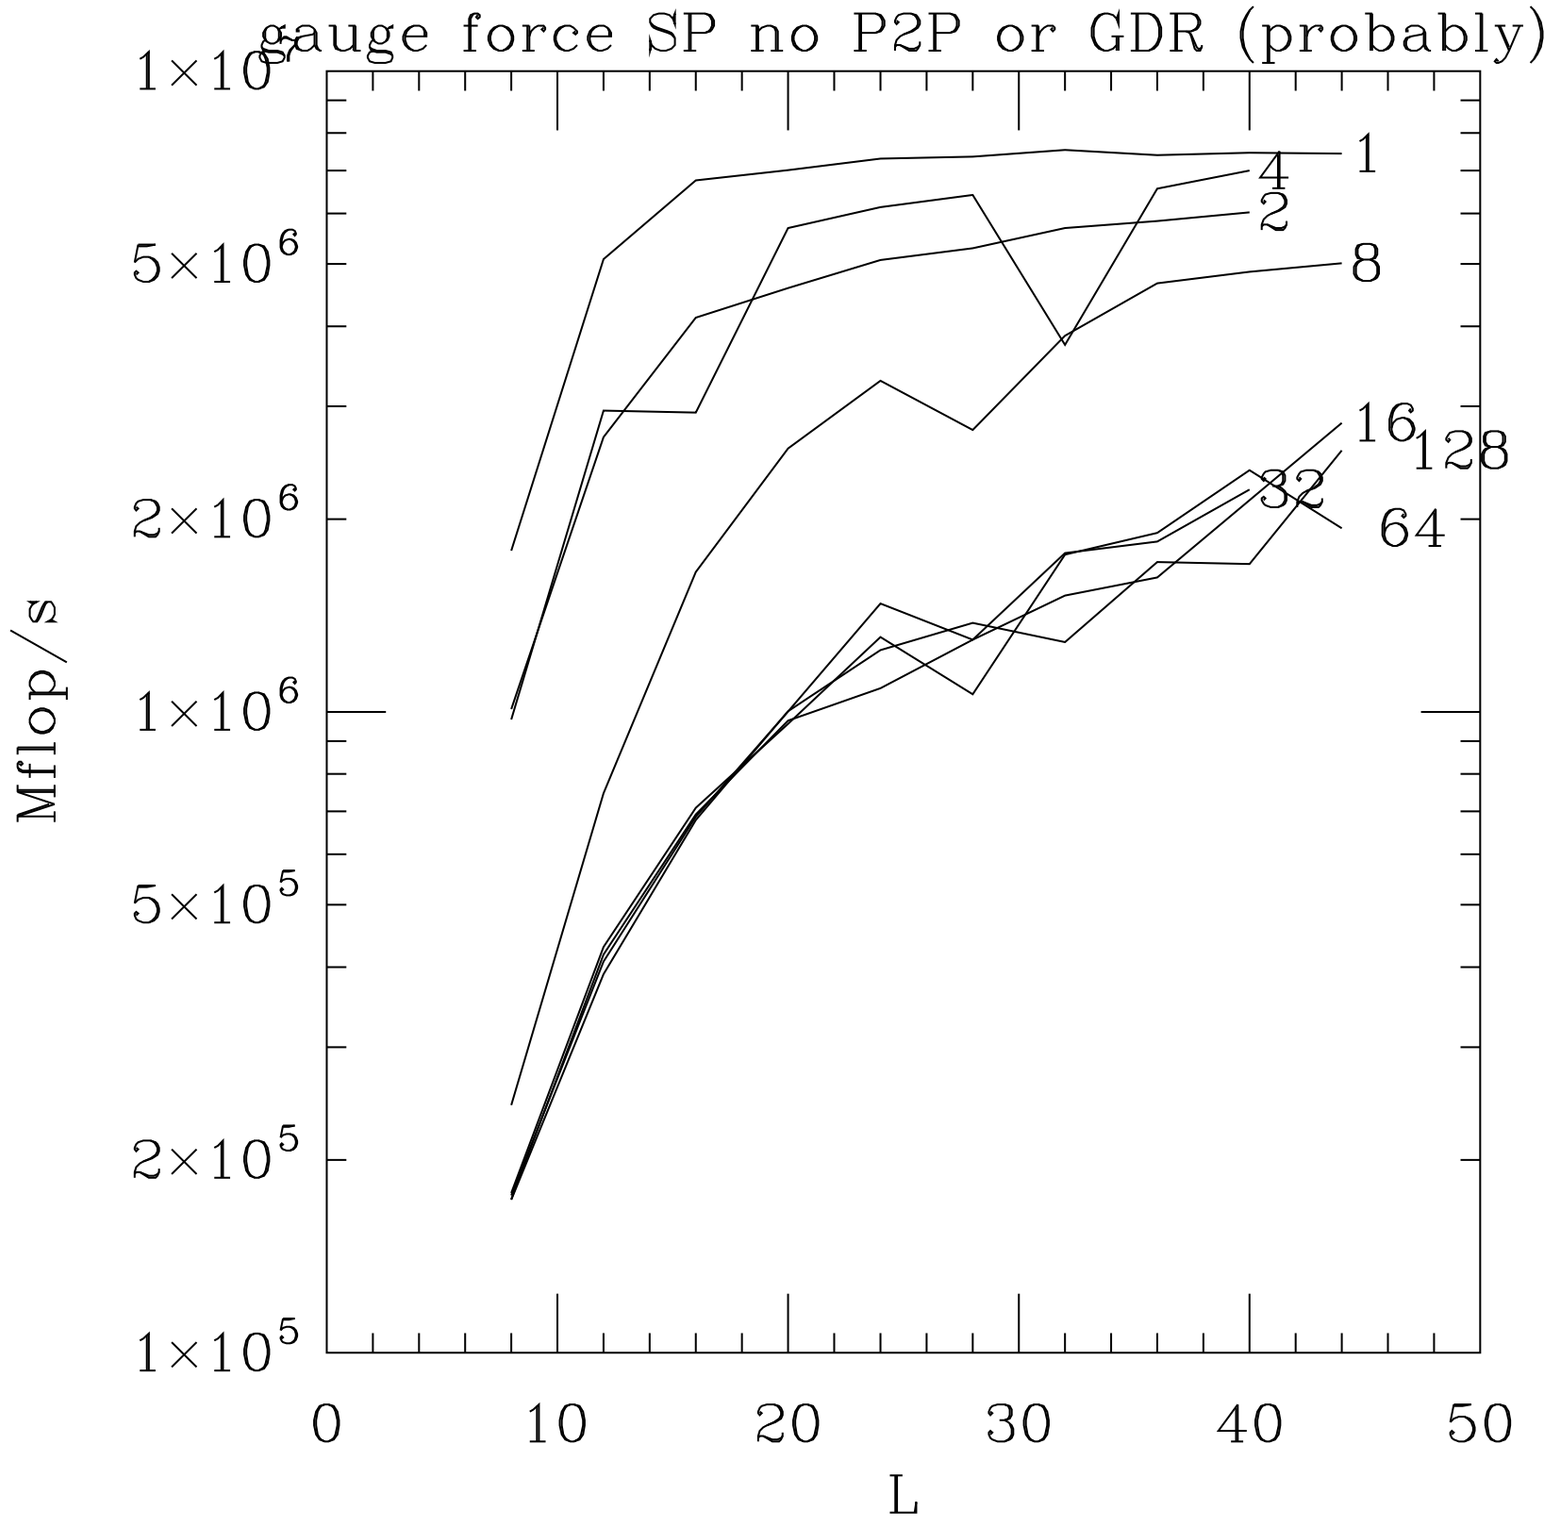

Software Development

- ✦ Lattice QCD has a long history of developing and sharing community software.
- ✦ QUDA project began in 2008 at Boston University.
  - Two of the main early developers Kate Clark and Ron Babich now work for NVIDIA
  - One of my former postdocs Mathias Wagner and a former BU postdoc Evan Weinberg also work for NVIDIA.
  - My work to support staggered quarks done with Guochun Shi while on sabbatical at NCSA for Blue Waters project
- ✦ QUDA is being generalized to support HIP, SYCYL, and Open MP.
- ✦ Our community clearly benefits greatly from this, and probably not easy to replicate in other areas of science.

From Jiqun Tu, Kate Clark, Evan Weinberg, and Mathias Wagner

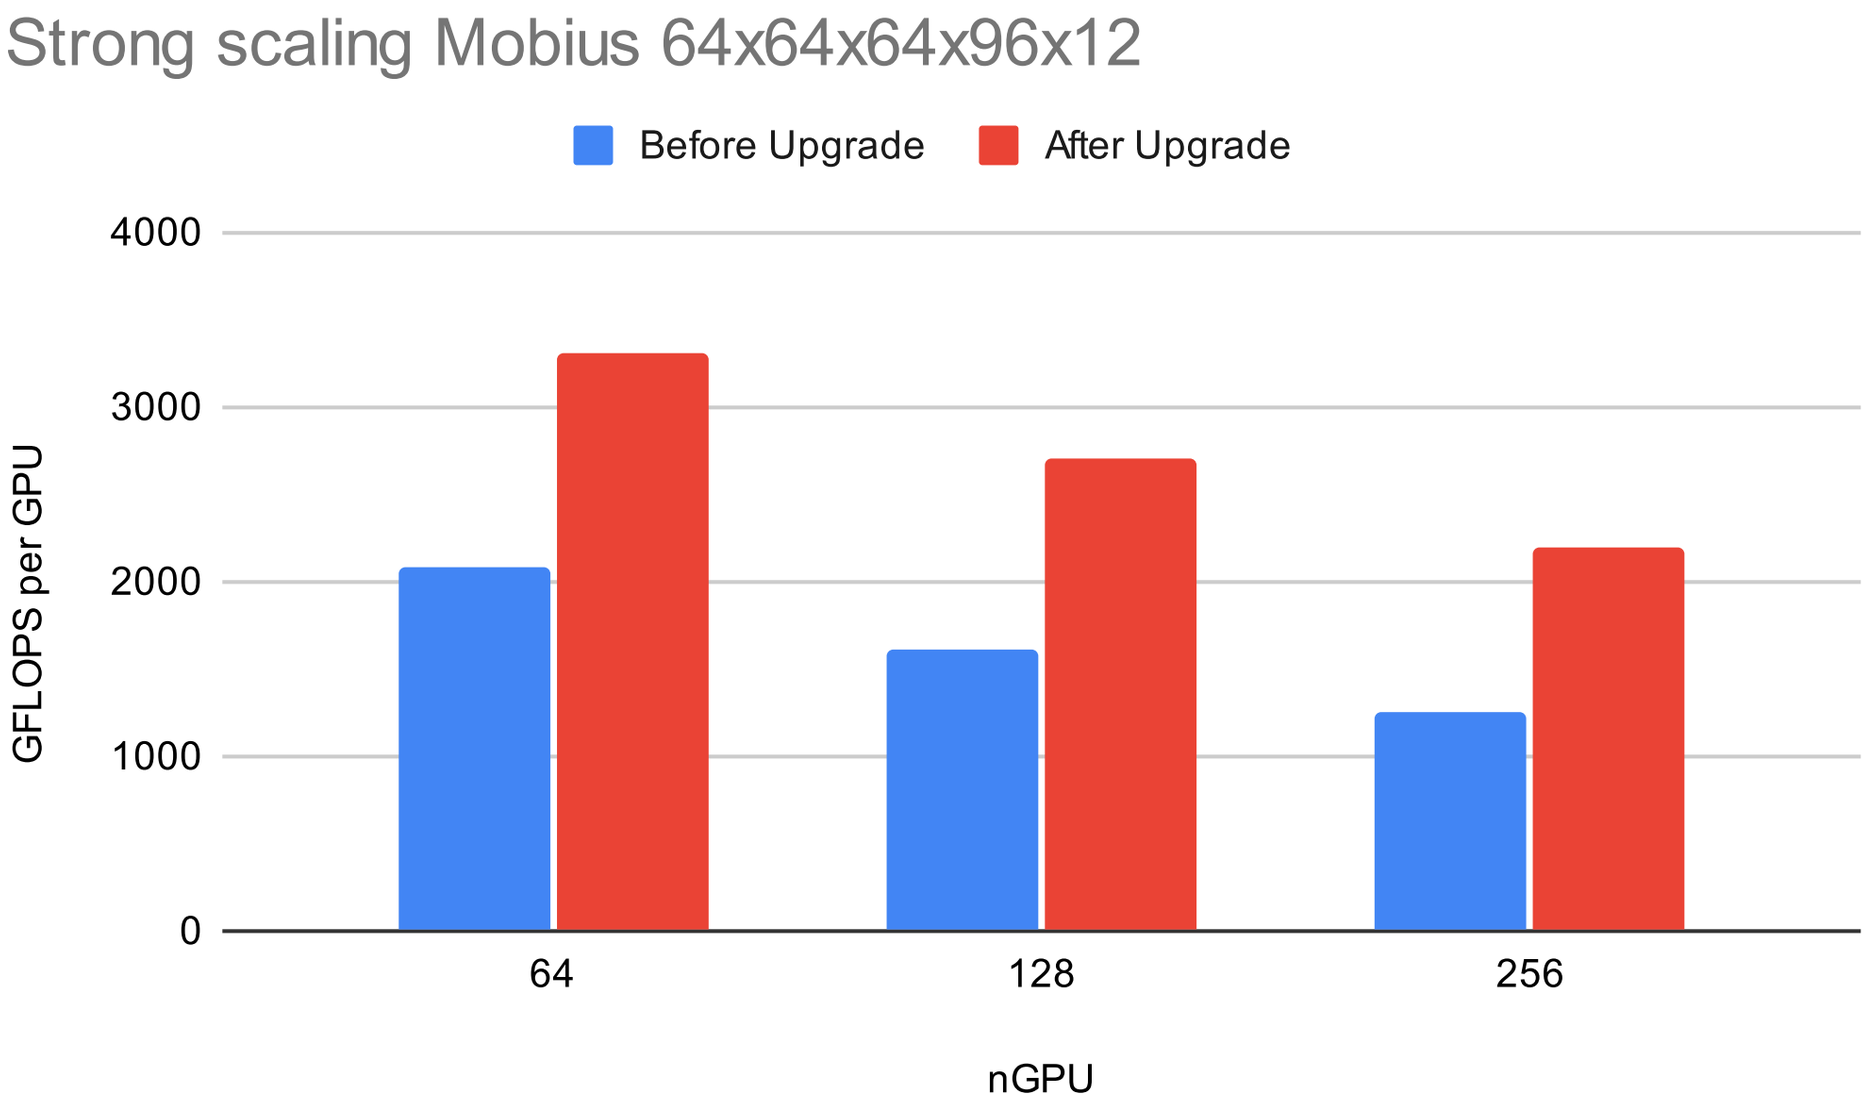

- Strong scaling CG (in double-half precision) performance on a 64x64x64x96x12 lattice with Mobius domain wall fermion using QUDA. Results are preliminary: by using a special CPU-GPU-NIC binding, the "After Upgrade" performance can be increase by about 30% over the numbers shown here, but currently do not have access to those data due to ongoing Perlmutter maintenance.
- For a full MatPCDagMatPC operator with Ls = 12, QUDA FLOP = 12976, Grid FLOP = 8064.

Performance Portability

- ✦Not at all clear what is the right path to develop performance portable software in near future
- ✦Kokkos, RAJA, OpenMP, DPC++, HIP, Alpaka ??
- ✦There has been a series of DOE meeting on performance portability.
- ✦I co-edited an issue of Computing in Science & Engineering with Doug Doerfler (NERSC), Bill Gropp (NCSA), Barry Schneider (NIST), and Alan Sussman (NSF); Sept-Oct, 2021)

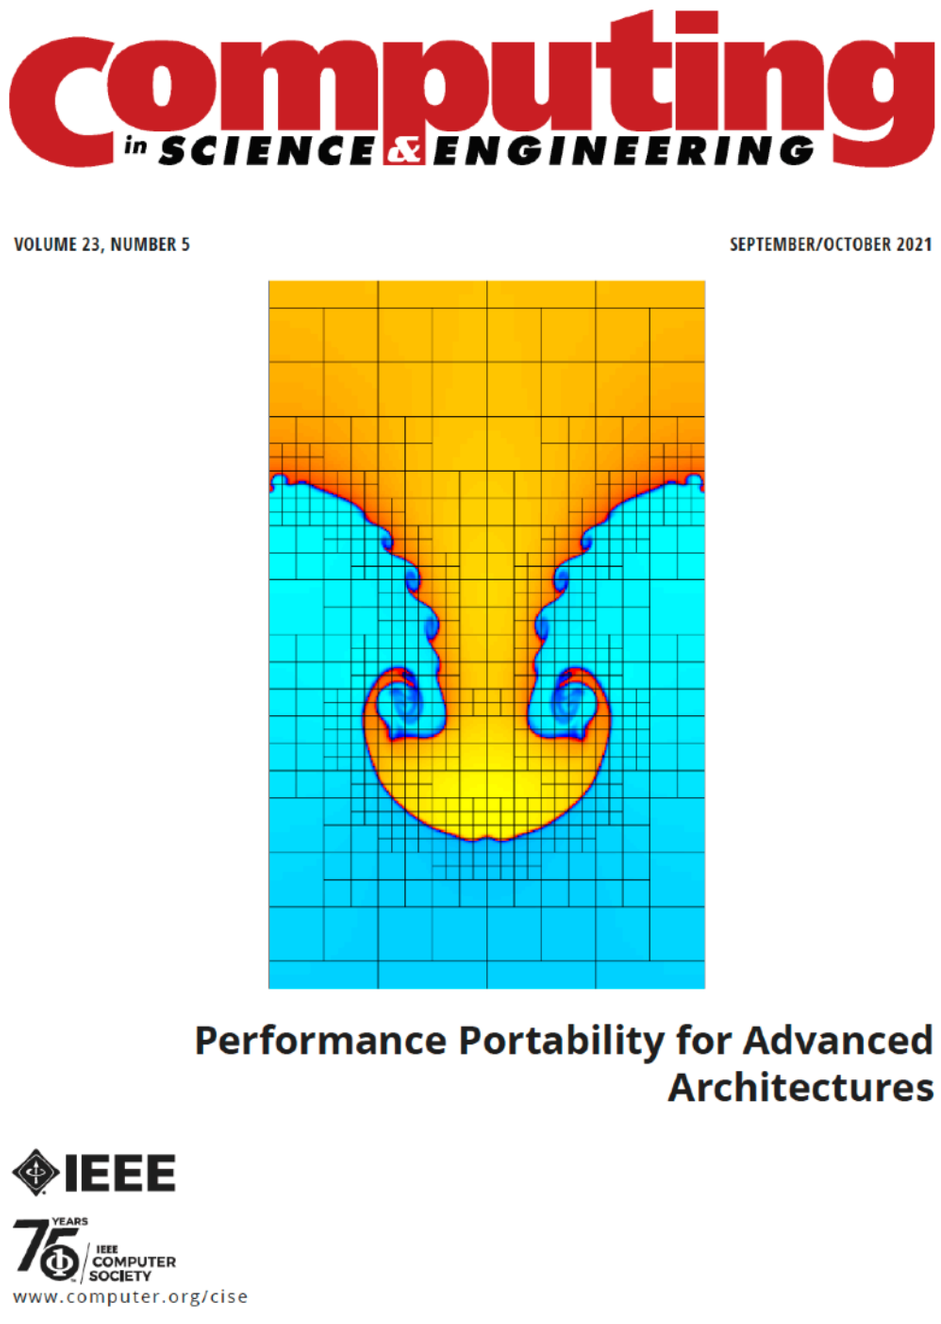

Conclusions

- ✦ Perlmutter is a very powerful platform for scientific computing
- ✦ If you have been using Cori GPU, the transition should be easy.
- ✦ If you are just starting out with GPUs, study the various approaches before committing to a porting strategy

- ✦ Have fun, but please…
- ✦ Leave some of the time for us! (Lattice QCD physicists)