Lattice QCD at NERSC by the MILC & Fermilab Lattice Collaborations

Steven Gottlieb
Indiana University
MILC and Fermilab Lattice Collaborations

NERSC@50 Seminar Series
January 27, 2025

Outline

- ✦ Introduction
- ✦ Lattice QCD
- ✦ Some ancient history
- ✦ Physics goals and some notable results
- ✦ Benchmarks
- ✦ Software development for GPUs
- ✦ Thoughts about the future

Introduction

- We are celebrating 50th anniversary of NERSC
  - I am neither a historian nor an archivist
  - Tried to recover some of the early notes and results
  - But, we don’t just want to review the past
  - At least a brief discussion of lattice QCD
  - Plan to present material about the present and future
- Very pleased that the MILC code has been used for several generations of benchmarking for NERSC’s computer acquisitions.

Lattice QCD

- Quantum Chromodynamics (QCD) is the 50+ year old quantum field 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 $m_{ud} \rightarrow m_u, m_d$

✦ To control systematic errors, we must:

- take continuum limit, i.e., $a \rightarrow 0$
- 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 Ensembles

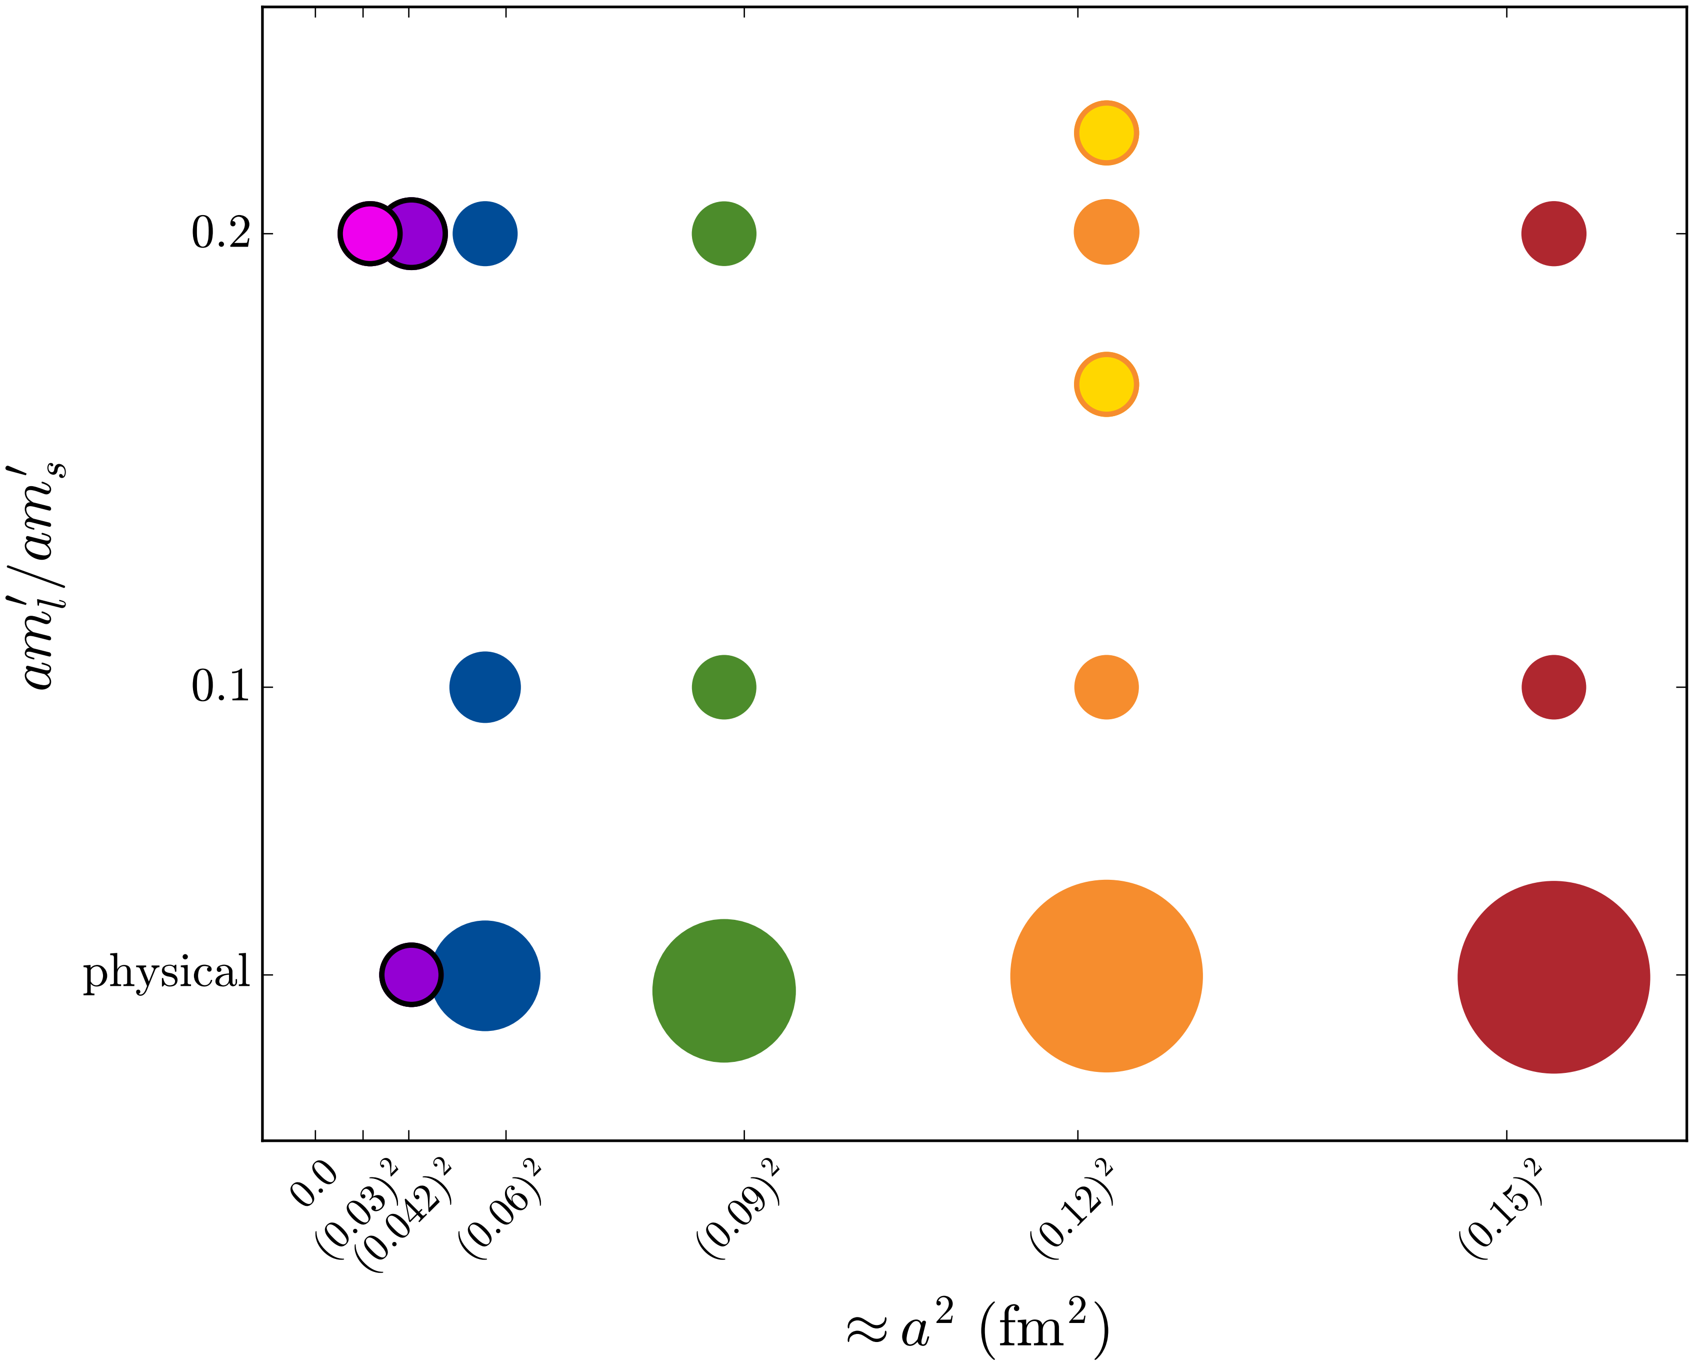

This is incomplete, but shows how we cover a range of lattice spacing and light quark mass.

In the beginning…

- I checked my disk to try to find earliest references to our computing at NERSC
  - April 28, 1995, email from Bob Sugar
  - I just spoke to Greg Tomaschke, who is now running the NERSC T3D. One or more of the paying users has complained that we are interfering with his (their) work. The upshot is that we are essentially going to be running on background until our proposal is approved, which should be soon.
  - May 11, 1995, email to Bob Sugar

UC has about 16% of the machine under the current operating mode. We are trying not to give any single user more than 3% of the 16%. Your request is for 3840*128 hours (that's right, isn't it?) putting your request orders of magnitude above what we would normally allocate.

I would like to see you get your work done, so we will try to accommodate you if at all possible. On the other hand, I don't want the other users getting upset with me. My sense is that a few projects are starting to show signs of life. So I hope you can get a lot of time now, while the getting is still good.

Please write a nice paper for us for our year end report! You are one of the few true production users, and I need someone to say this was worth the effort.

Some things just have not changed over the years. We always want more time!

Systems I've Used at NERSC

- ✦ Aliases on my desktop: H4P was accessed by telnet! All others by ssh
- ✦ 1994: H4P, Cray T3D
- ✦ 1996: Franklin, Cray YMP
- ✦ 1997: MCurie, Cray T3E-900
- ✦ 1997 or 2001: Seaborg, IBM SP
- ✦ 2009: Hopper, Cray XT5 (Opteron quad-core)
  - XE6 (Opteron hex-core)
- ✦ 2011: Carver, IBM iDataPlex (Nehalem quad-core)
- ✦ 2013: Edison, Cray XC30 (Ivy Bridge 12-core, dual socket)
- ✦ 2014: Babbage, (a testbed)
- ✦ 2015: Cori, XC40 (KNL & Haswell)
- ✦ 2021: Perlmutter, HPE Cray EX325n (AMD EPYC & NVIDIA A100)

Fun Fact

- Richard Gerber sent me a list of top users over the years.
- My friend and retired collaborator Doug Toussaint is #1!
- Seven of top 11 do lattice QCD.

Top Users

Most Hours			NERSC Hours	NERSC Days	Top User	
toussain	Doug	Toussaint	4,503	187.6	2004, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013	LQCD
dks	Donald	Sinclair	1,986	82.7		LQCD
compo	Gilbert	Compo	1,672	69.7	2014	Climate
zlin	Zhihong	Lin	1,573	65.5	2003	Fusion
detar	Carleton	DeTar	1,490	62.1		LQCD
adrianne	Adrianne	Middleton	1,430	59.6		Climate
u6338	David	Trebotich	1,364	56.8		Subsurface
wdetmold	William	Detmold	1,361	56.7		LQCD
sungwoo	Sungwoo	Park	1,331	55.5	2020, 2021	LQCD
psteinbr	Patrick	Steinbrecher	1,271	53.0	2017	LQCD
hwlin	Huey-Wen	Lin	1,080	45.0	2015, 2016	LQCD

Big Picture I

- ✦Mid-1980s to mid-1990s, using two flavors (up and down) naive staggered (or Kogut-Susskind) quarks
- ✦Late 1990s, started using improved actions and added a strange quark
  - asqtad action program summarized in Reviews of Modern Physics **82** (2010) 1349-1417; e-Print: 0903.3598 [hep-lat]
- ✦Third generation of calculations uses Highly Improved Staggered Quark (HISQ) action and adds charm quark
  - Most of this work is done in collaboration with the Fermilab Lattice Collaboration and some with High Precision QCD (HPQCD)
  - We can now run with the light quark mass at it physical value
  - Soon look at some aspects of generating an ensemble with the finest lattice spacing (0.042 fm) we have tried with physically light quarks
- ✦Looking forward to NERSC 10 for reducing lattice spacing

Big Picture II

- ✦Multiple physics goals:
  - determine quark masses
  - study weak decays of heavy quarks to determine Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix
  - calculate hadronic vacuum polarization contribution to muon anomalous magnetic moment
  - study QCD at non-zero temperature and the quark-gluon plasma
- ✦Bullets 2 and 3 represent search for beyond the Standard Model physics
- ✦Next few slides feature some notable results

Quark Masses

- Fermilab Lattice and MILC, PRD **98**, 074512 (2018)
- Ratio of strange to light quark mass
- The dashed horizontal lines separate calculation with different numbers of dynamical sea quarks

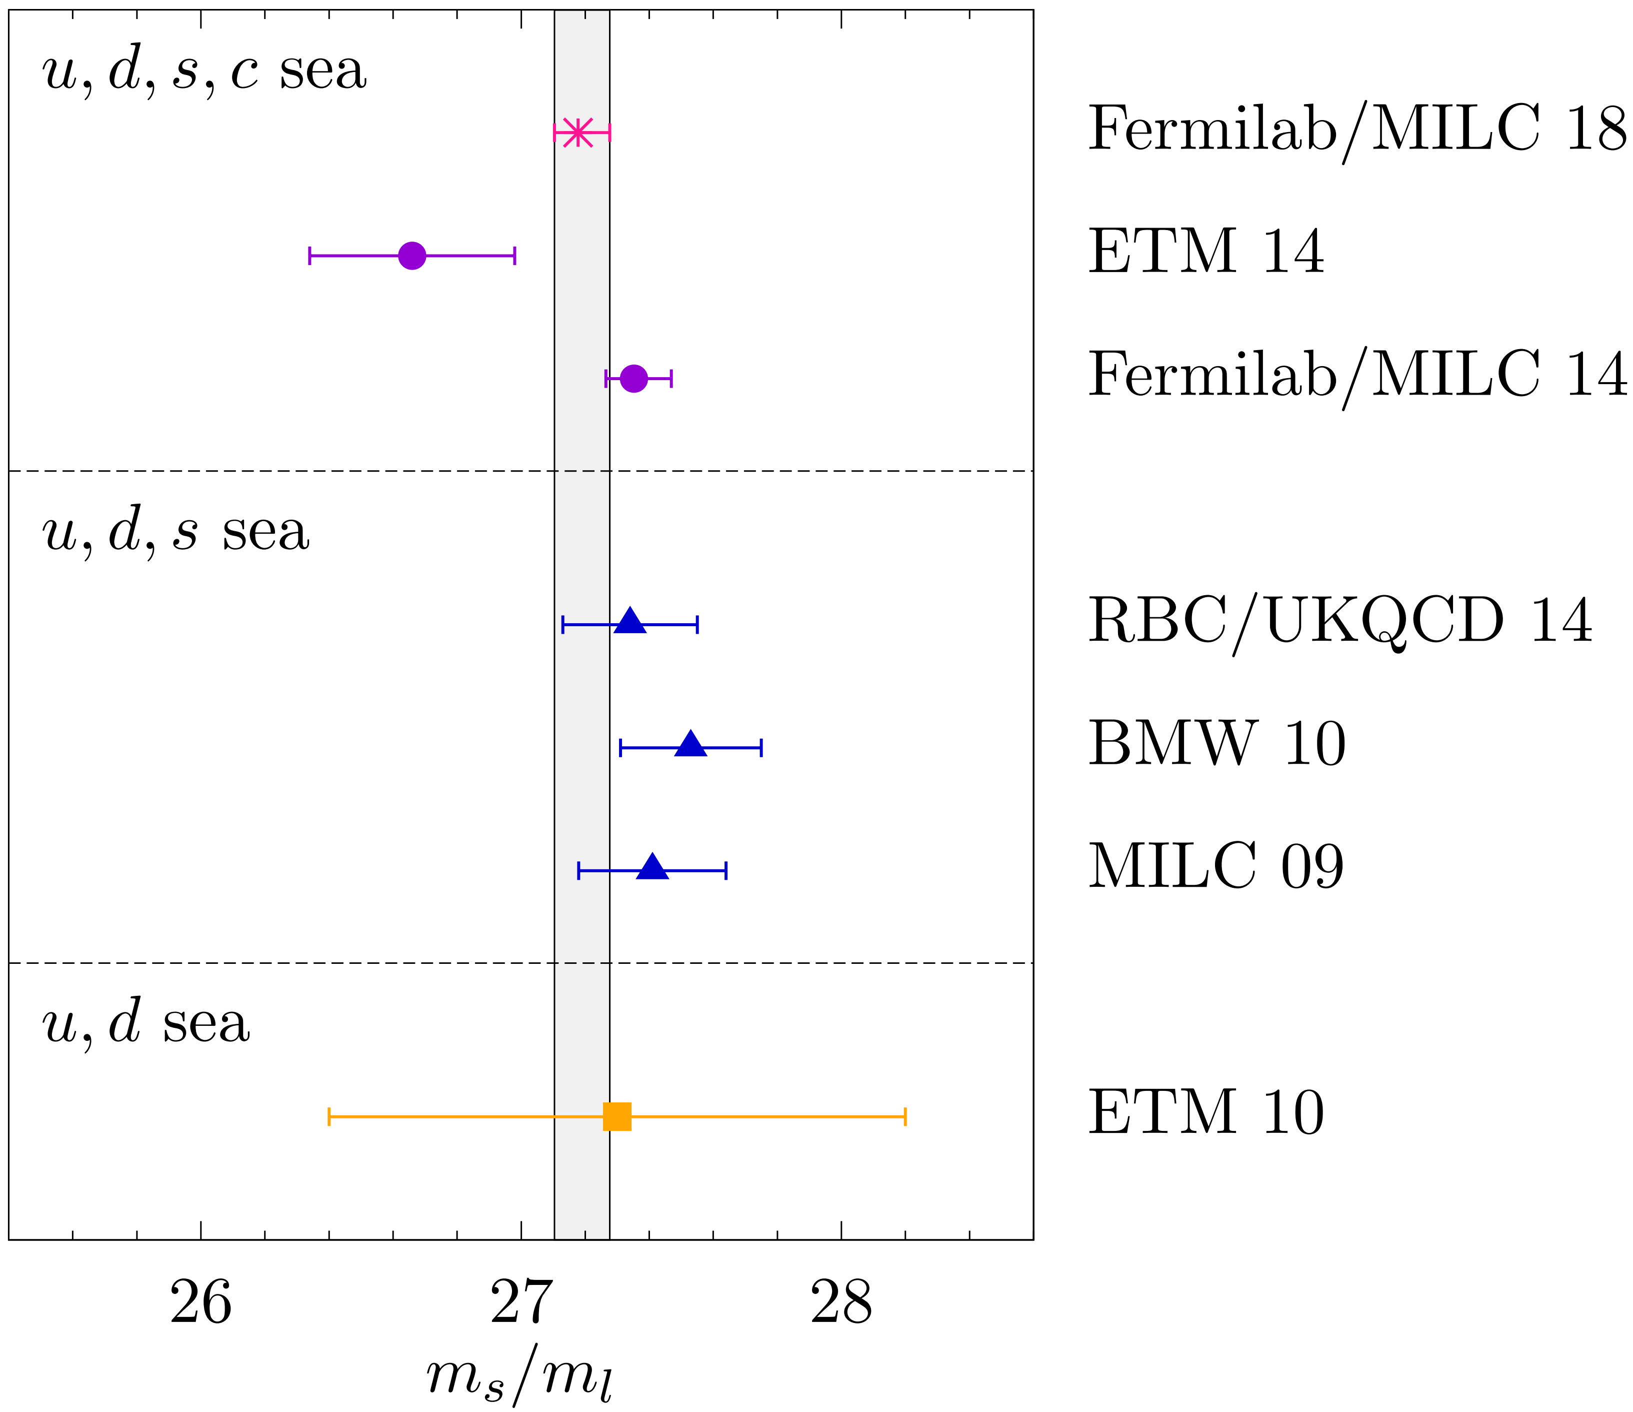

Decay Constants I

- Quantity related to decay constant for heavy-light meson with light (lower) or strange (upper) antiquark
- Note that as the lattice spacing gets smaller, we can increase the heavy quark mass
- Open symbols are not part of global fit to all the filled points

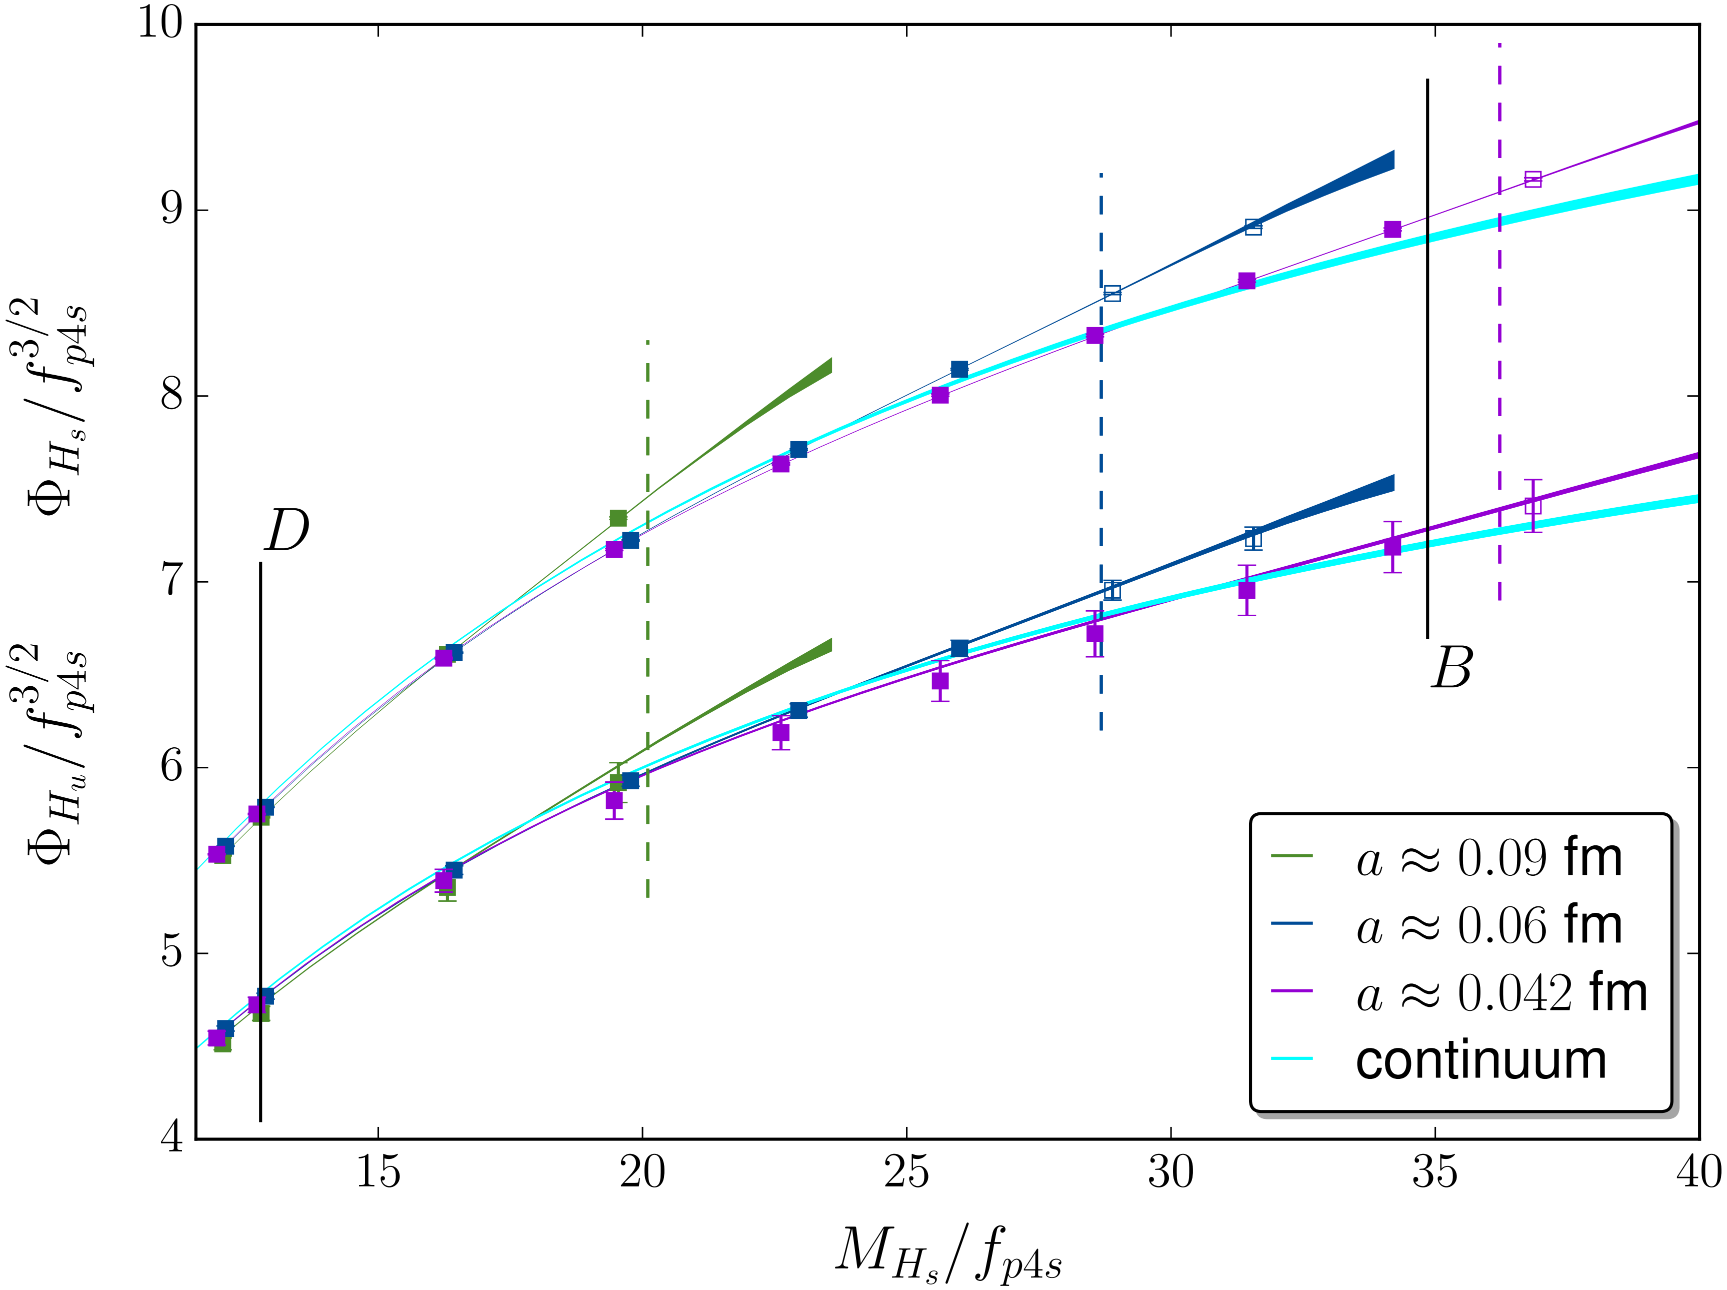

Decay Constants II

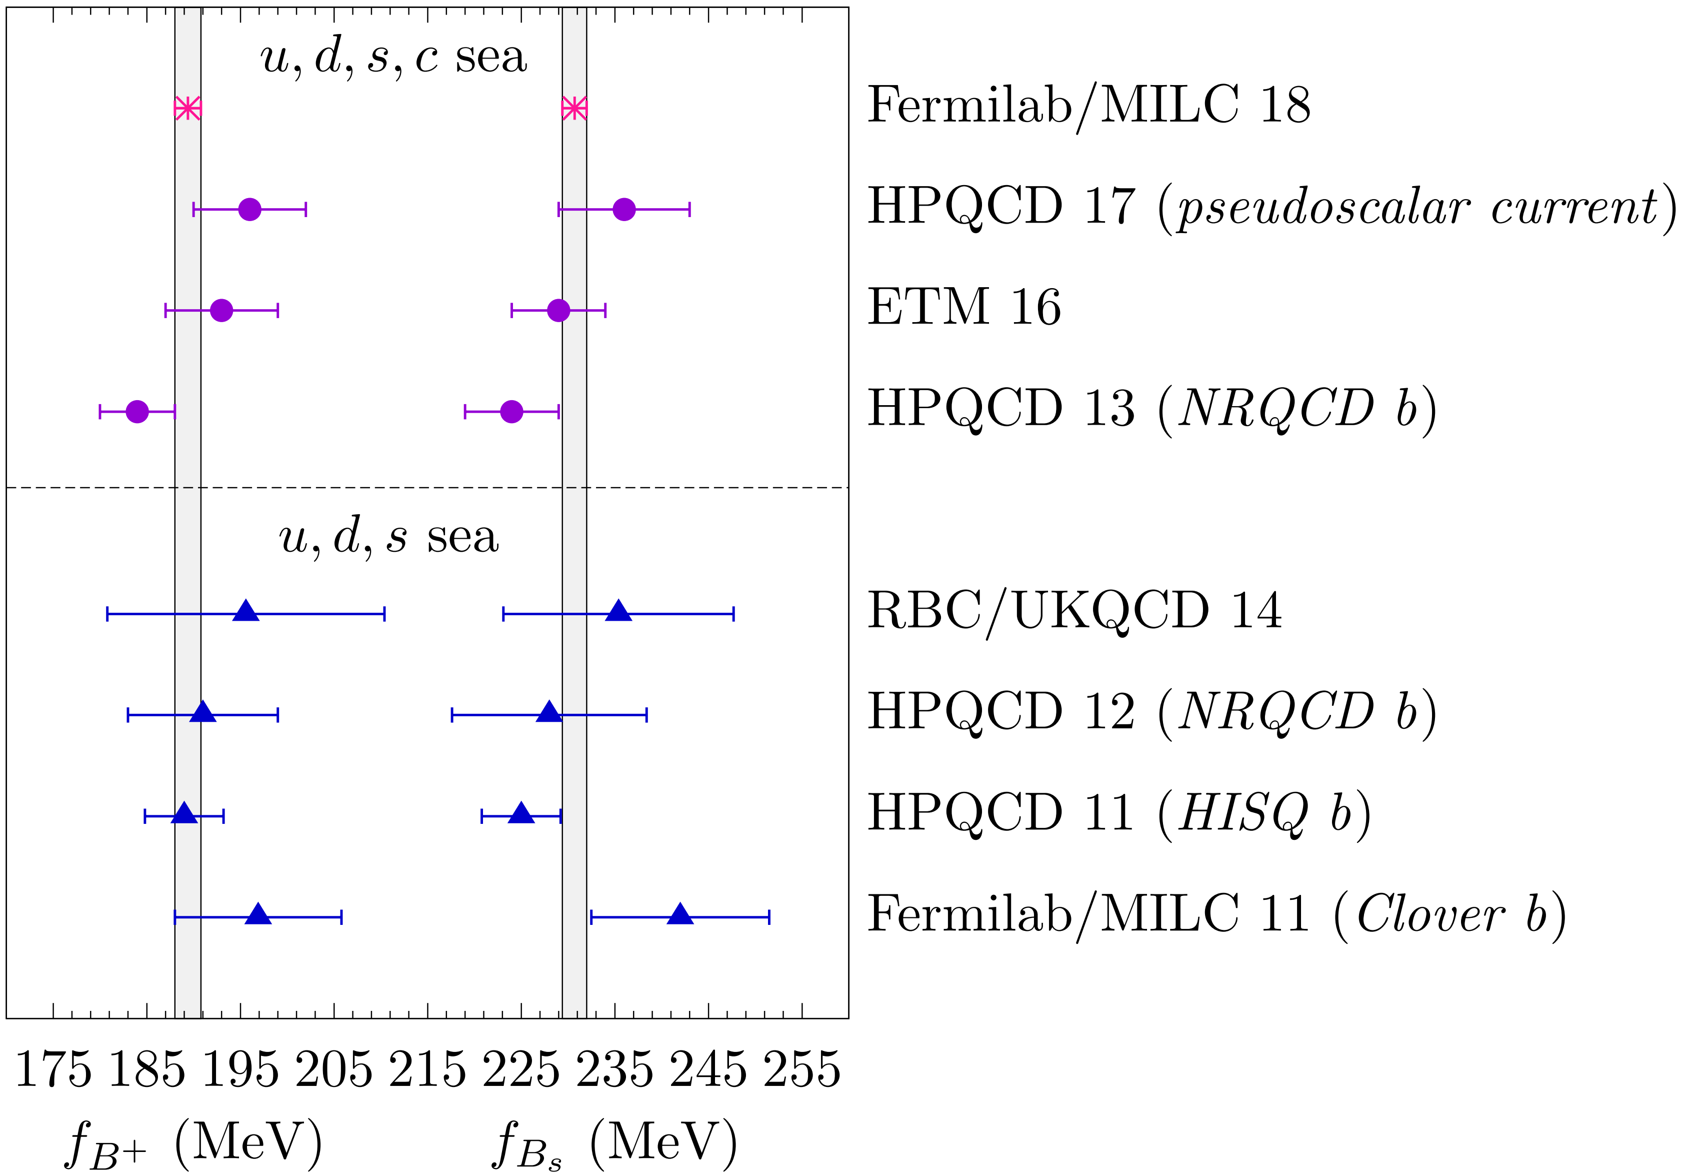

- Decay constants for B^+ and B_s mesons
- Note increase in precision compared to previous results
- The Belle II experiment at KEK lab in Japan is measuring these decays. Combining experiment with theory determines CKM matrix element.

Muon Anomalous Magnetic Moment

- ✦ The E821 experiment was done at Brookhaven National Laboratory ending around 2000
  - One of the most precisely measured quantities in physics
  - However, the result was about 3 standard deviations from theory
- ✦ The superconducting magnet from BNL was shipped around Florida and up the Mississippi River to Fermilab where a new experiment E989 was done.
  - Spectacular agreement with the BNL result
  - Final results expected to be announced in a couple of months
- ✦ In the meantime, two theoretical approaches, but they are not in great agreement.
  - Complicated story
  - Lattice QCD may now be the leading theoretical approach

Muon HVP Results

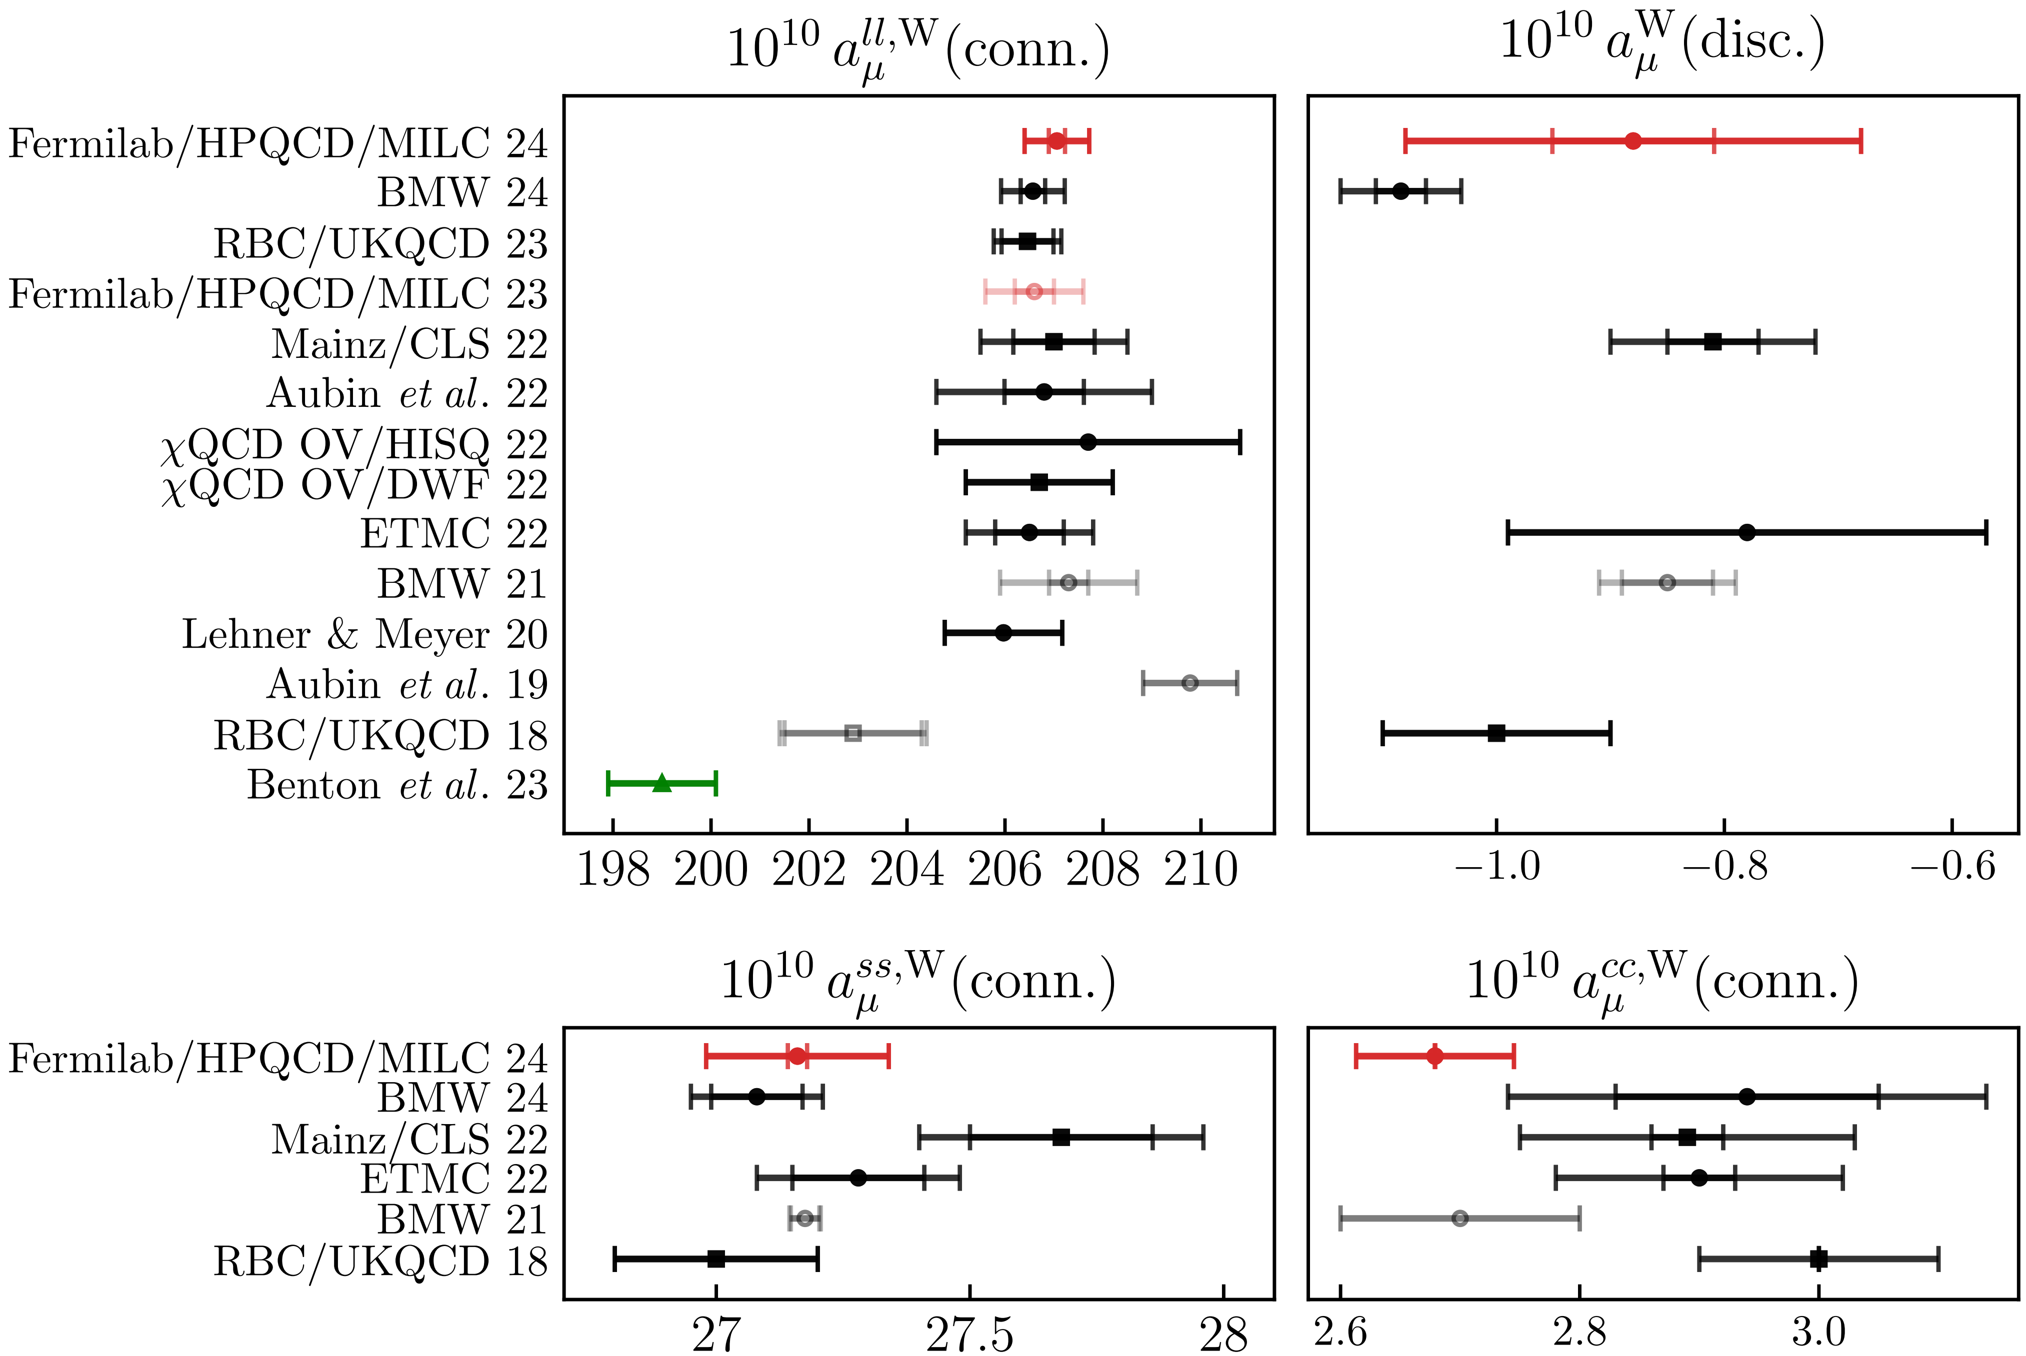

arXiv:2411.09656: Our results in red. Improving point on upper right is one of our 2025 goals on Perlmutter.

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.

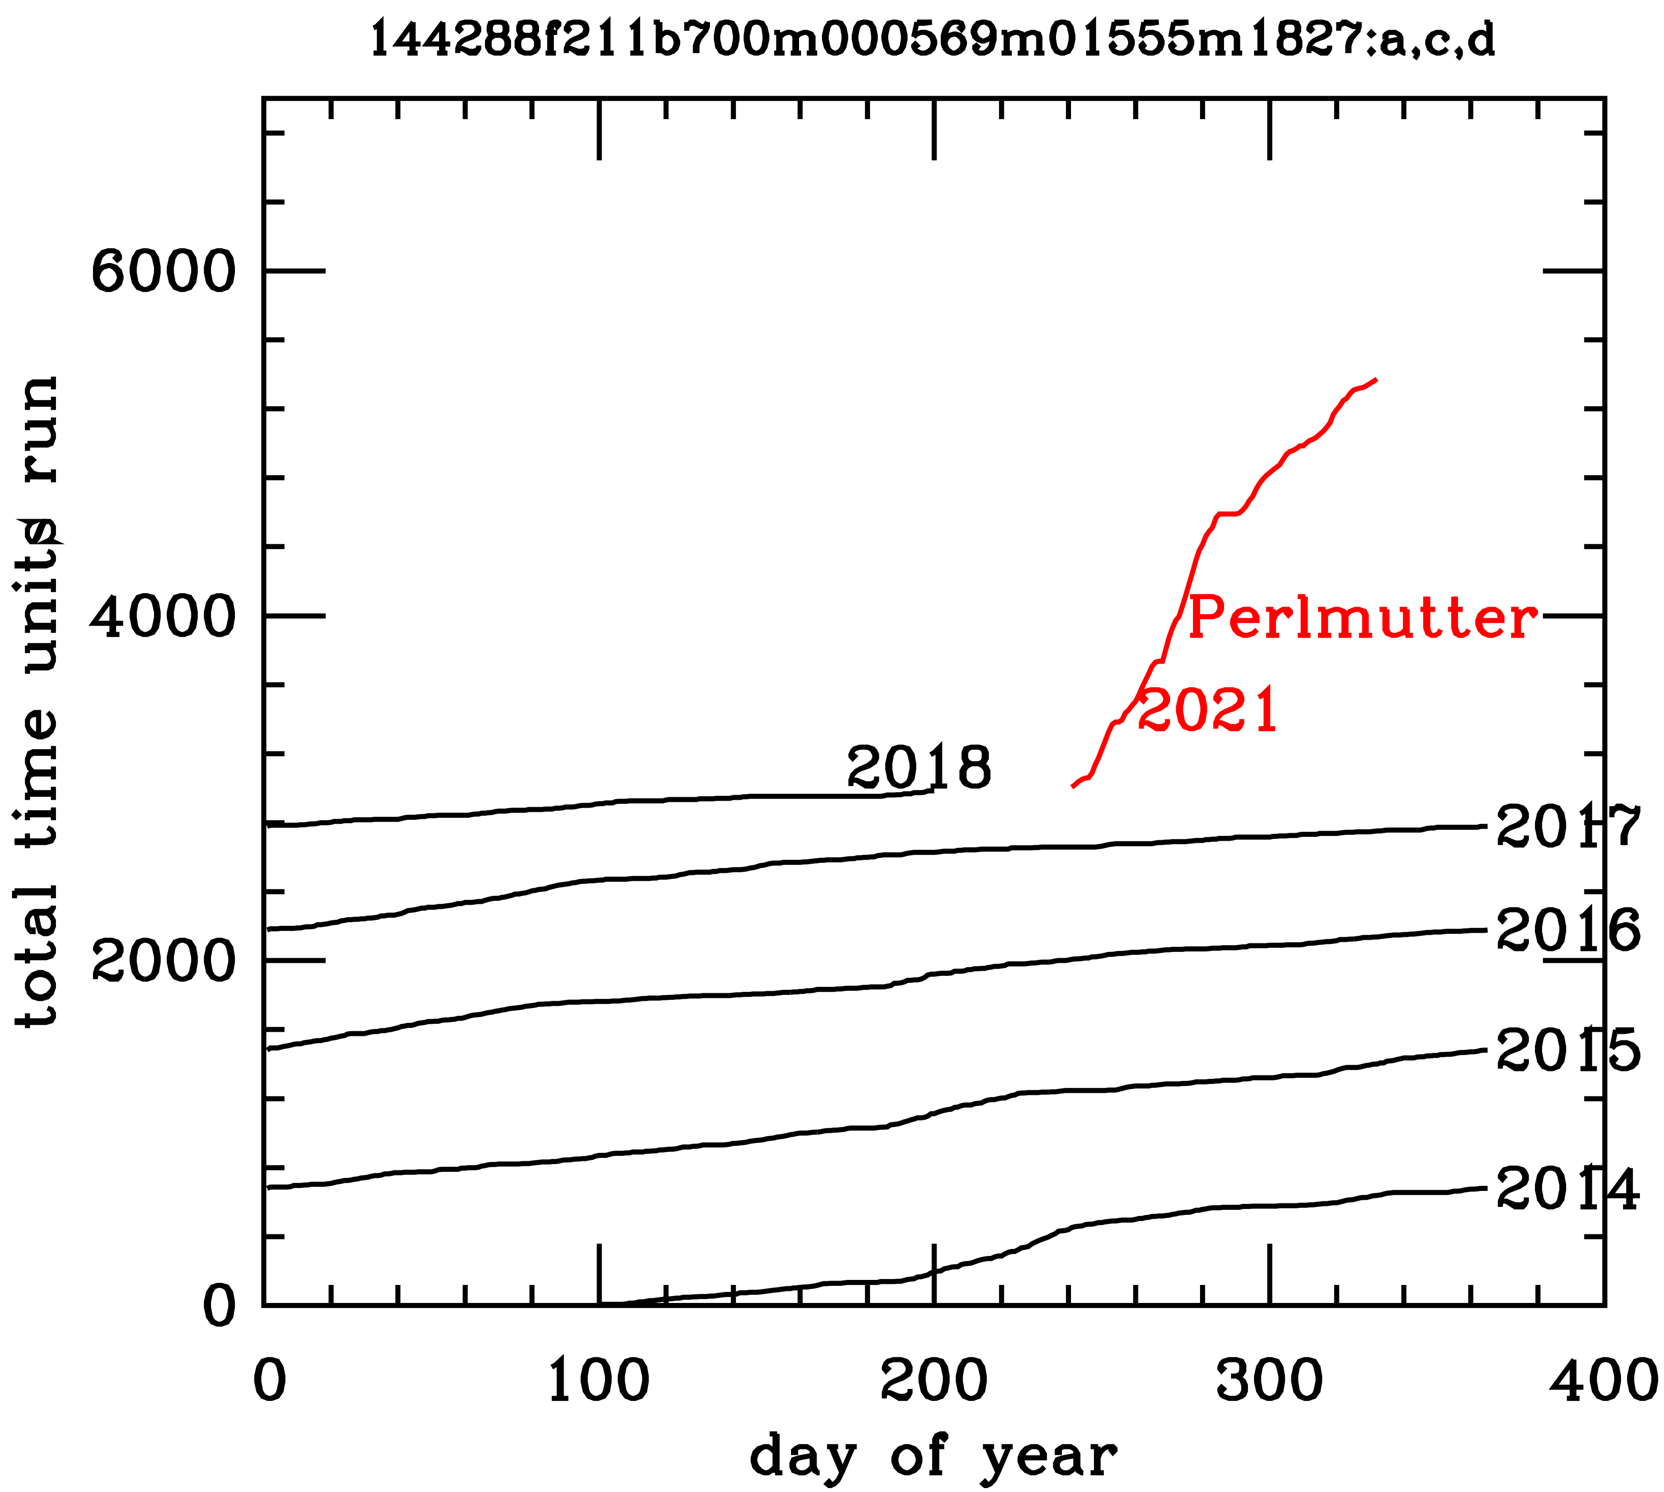

Power of Perlmutter II

- 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.

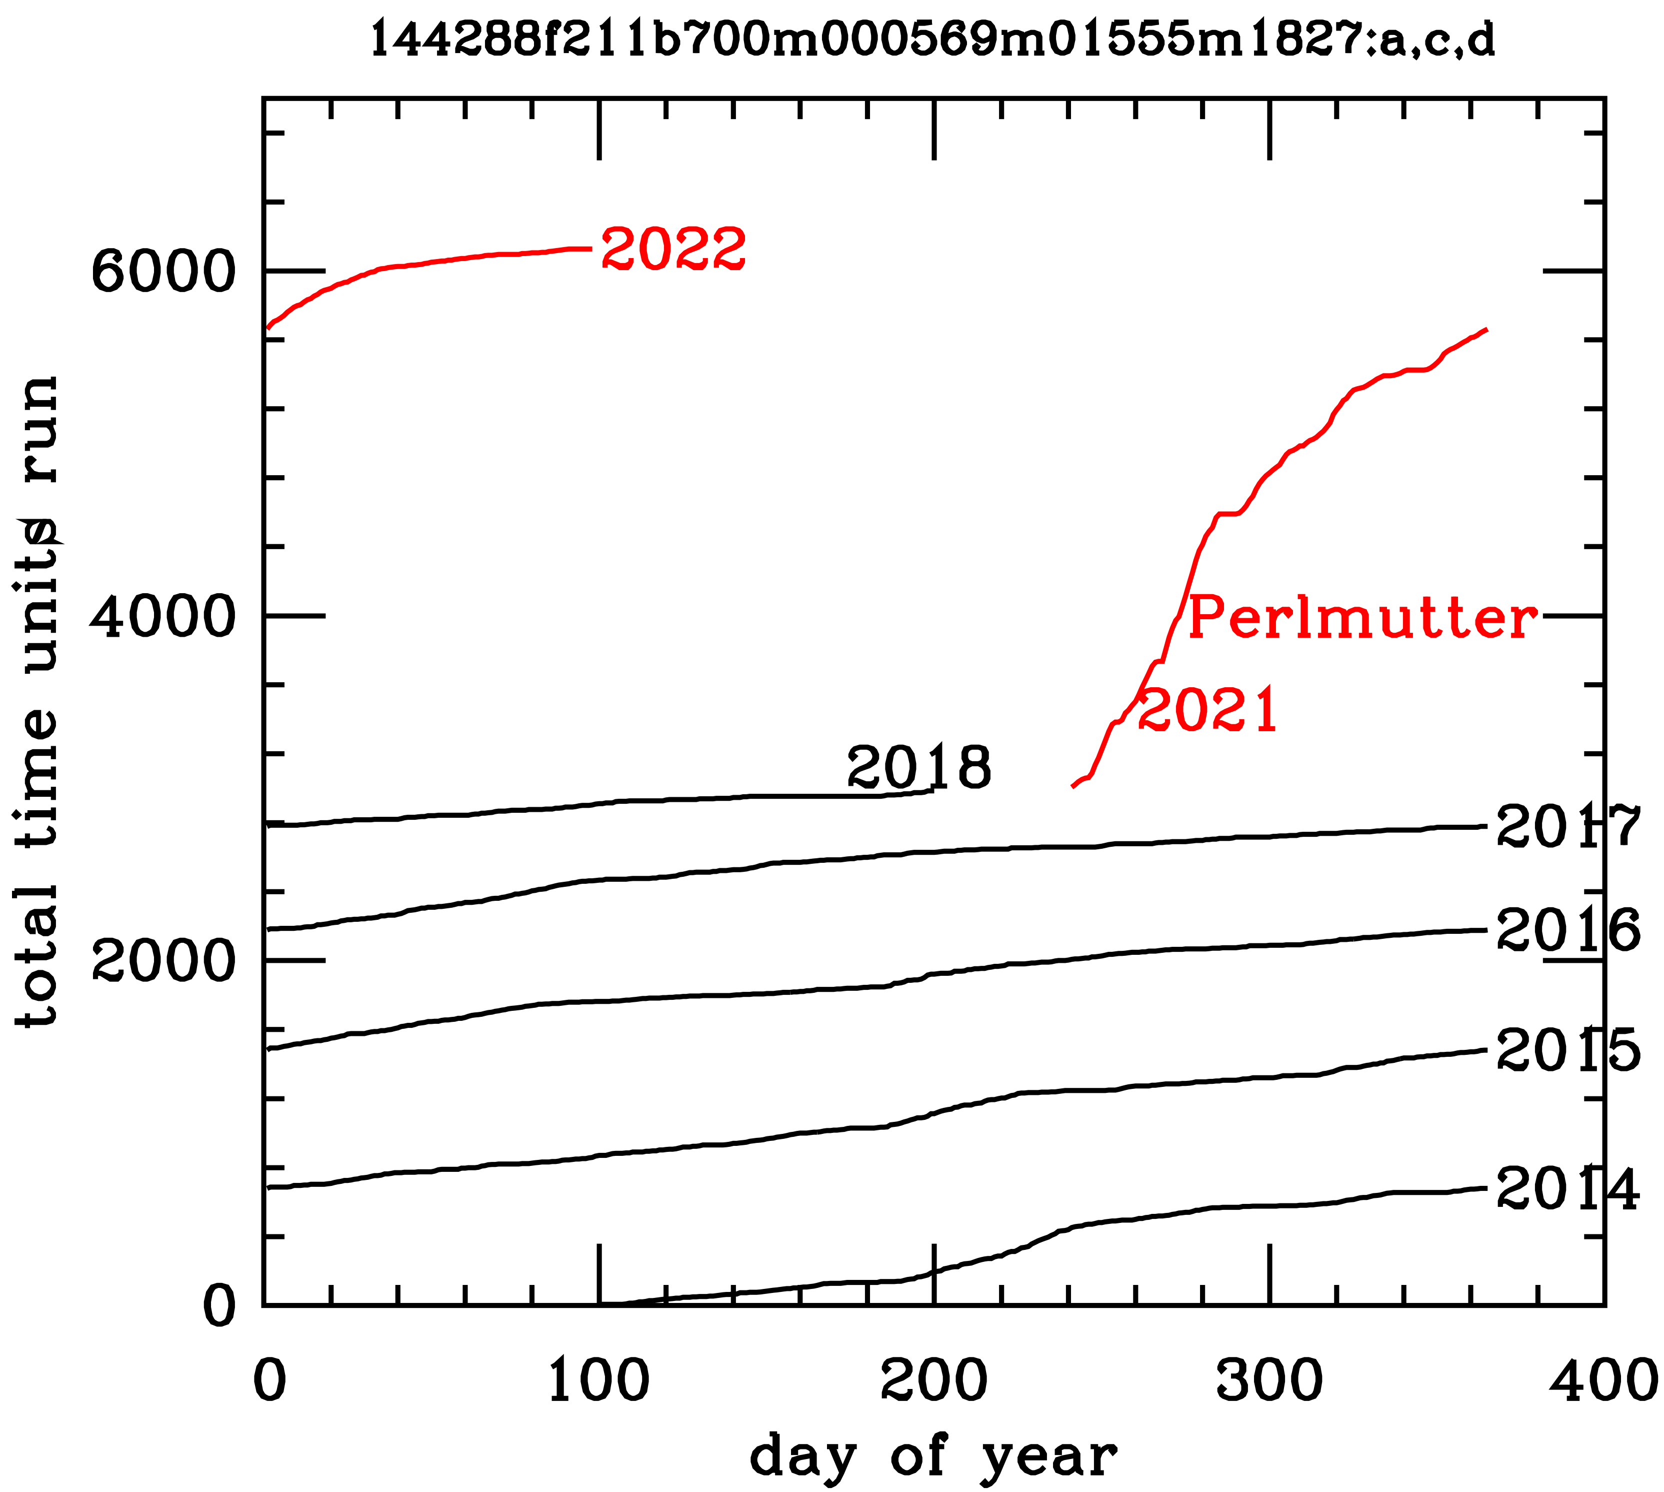

Cross Platform Comparison

- Table compares times to run a trajectory of length 2 and save the configuration on five different computers
  - second Perlmutter line reflects network upgrade from Slingshot 10 to Slingshot 11

Computer	nodes or cores	MPI ranks	generate (hr)	save (hr)	total (hr)
Edison	18432	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 (SS10)	128	512	1.46	0.07	1.53
Perlmutter (SS11)	128	512	0.99	0.04	1.03
Aurora	32	384	2.21	0.06	2.27

GPU Performance Comparison

- After Perlmutter (NVIDIA A100, 4 GPUs/node), we gained access to
- Frontier (AMD MI250X, 2 GCDs per GPU, 4 GPUs/node)
- Aurora (Intel Data Center GPU Max Series, 2 tiles per GPU, 6 GPUs/node)
- Vista, DeltaAI (NVIDIA GH200, 1 and 4 GPUs/node, respectively)
- In next slide, compare performance on a single task for staggered quark conjugate gradient performance in double precision.
- Volume is L^4 grid points
- Use a disordered start and report maximum performance during run (explain why)

GPU Comparison II

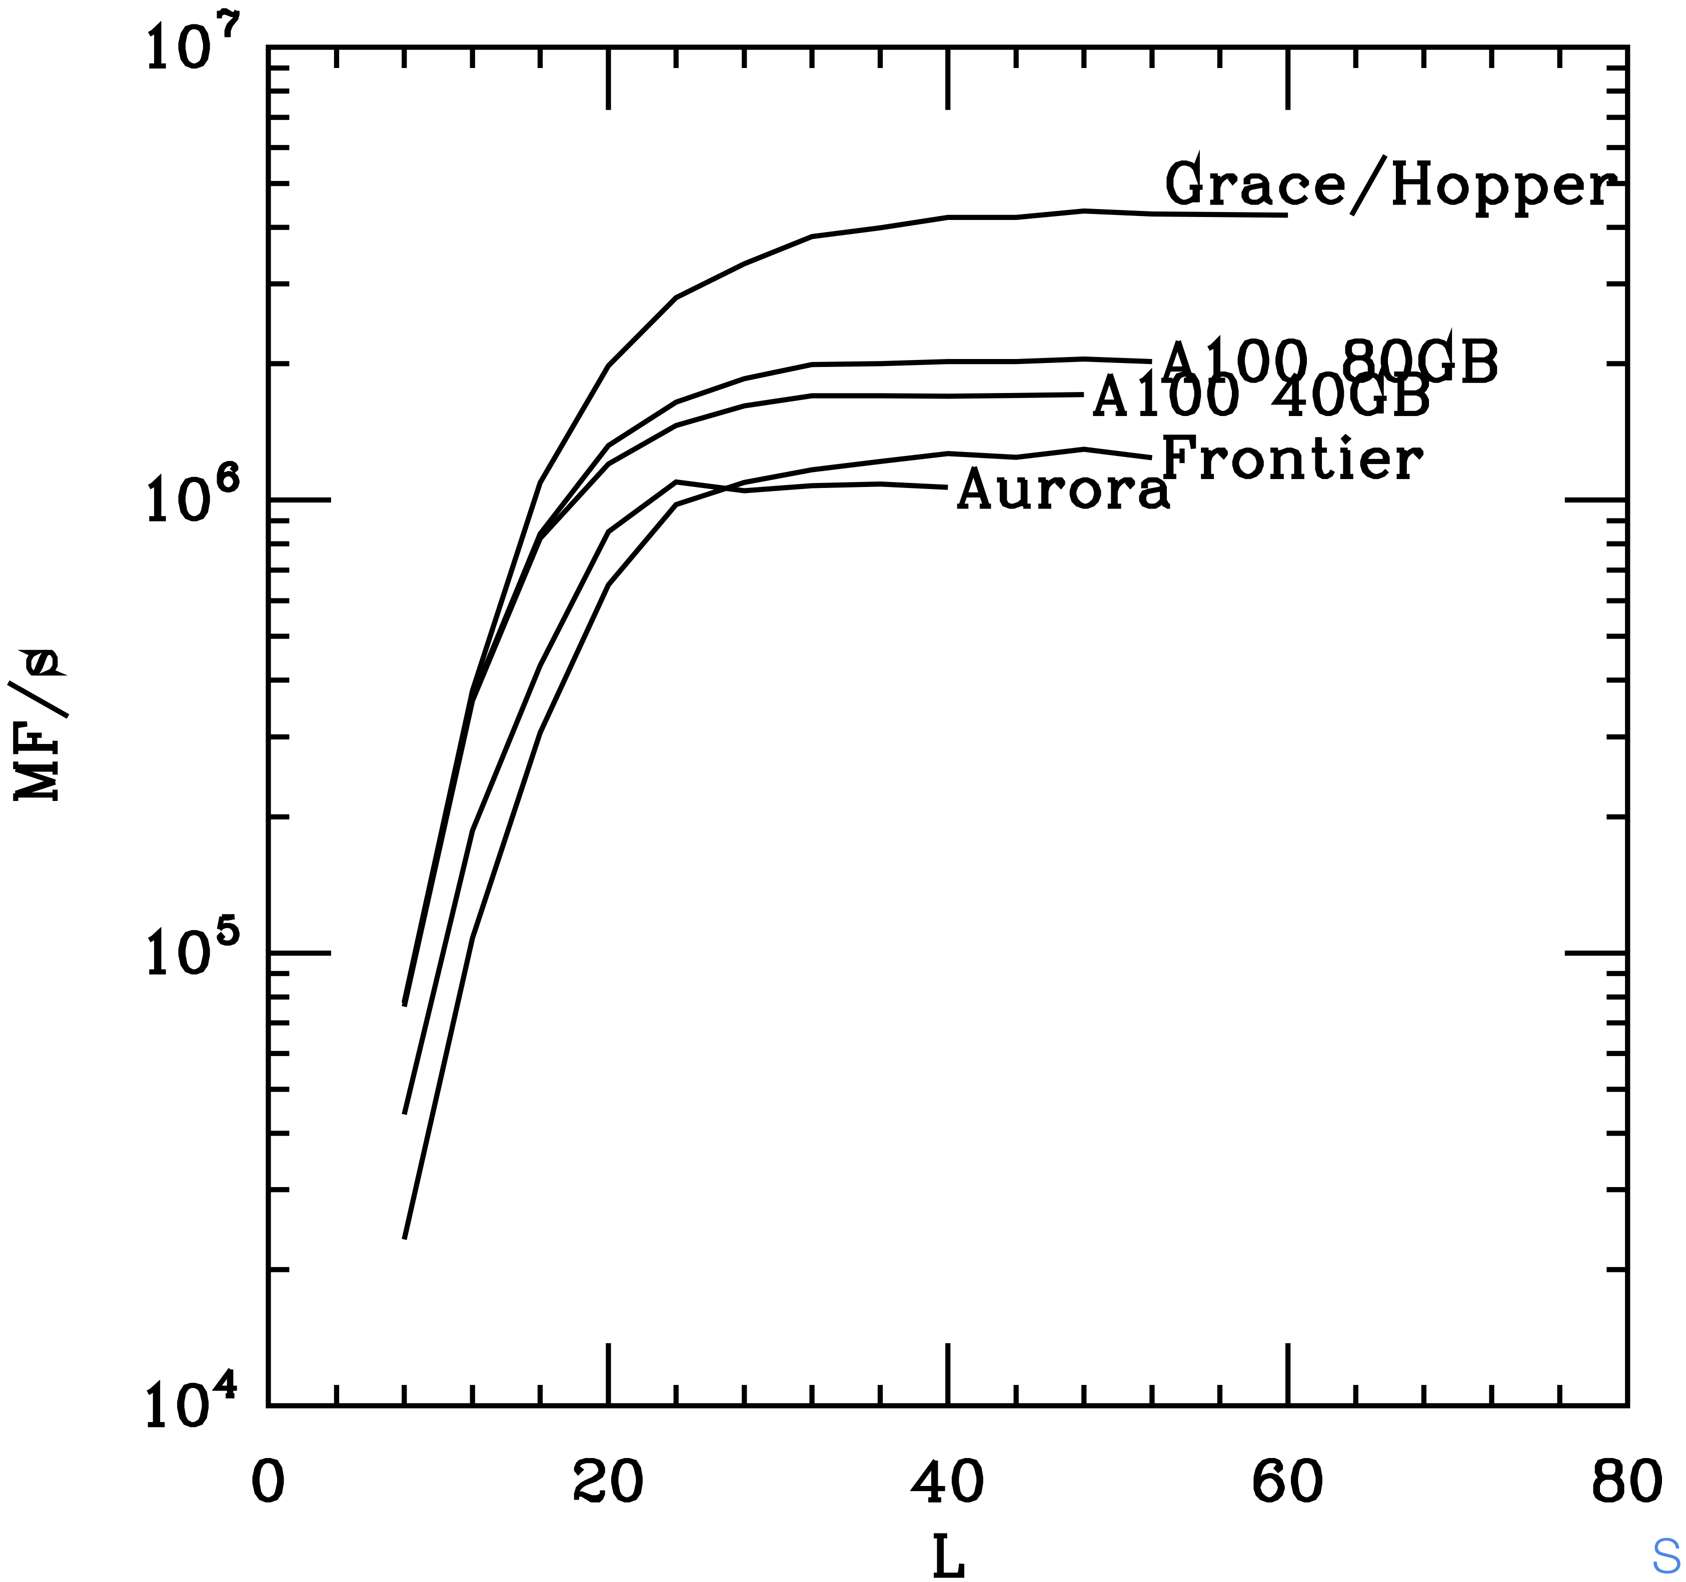

- CG performance in single (GPU, A100), (CGD, MI250X), (tile, Intel Ponte Vecchio)
- Each reaches about 1 TF or more
- Need to have large enough local volume to have enough parallelism
- This is **not** a node to node comparison
- see details on previous slide

MILC Performance

- All running on 128 nodes, i.e., 512 GPUs on 144^3 X 288 configuration (using 11.6 GB of GPU memory)
- ON EACH NODE (RANK) 144 x 36 x 18 x 18
- Multimass CG: 525 GF/gpu single precision (mixed)
- Link smearing: 350 GF/gpu
- Gauge Force: up to 7.1 TF/gpu (quite variable)
- Fermion force: still don’t have flop count in code. Need to fix.

Software Development I

- ✦ 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 has been generalized to support HIP, SYCL, and Open MP.
- ✦ Our community clearly benefits greatly from this, and probably not easy to replicate in other areas of science.

Software Development II

- ✦ We also use Grid for a number of projects
  - Peter Boyle is the development chief
  - Hadrons is combined with Grid for some projects
- ✦ We have been working on eigenvalue deflation and multigrid
- ✦ Hwancheol Jeong has been working on eigenvalue compression as I/O and storage become significant with deflation
- ✦ Leon Hostetler has been working with Evan Weinberg and Kate Clark on deflation and multigrid to try to find the best solution for current and future ensembles.
  - need to optimize set up costs, i/o, and solve time in work flow

Thoughts About Future

- ✦ We still love high memory bandwidth and GPUs.
- ✦ In the past, I would have said we care little about I/O bandwidth, but both deflation with eigenvectors and multigrid place more demands on I/O.
- ✦ NERSC has played an essential role in supporting lattice QCD calculations for multiple collaborations.
- ✦ I look forward to continuing to compute at NERSC, but probably not for the next 50 years.
- ✦ Thanks to Richard Gerber for inviting me to speak
- ✦ Thanks to **all** my Fermilab Lattice and MILC collaborators for so many years of working together.
- ✦ We still have jobs to run…