EXALAT Benchmarking

Ed Bennett

ExaTEPP Workshop, 2023-06-22

Group theory

- Gauge (gluon) fields are represented on the lattice by matrices, e.g.
  - $SU(N)$: complex, $U^\dagger U = 1, \det U = 1$
  - $SO(N)$: real, $O^{\mathrm{T}} O = 1, \det O = 1$
  - $Sp(N)$: complex, $U^\dagger U = 1, \det U = 1, U^{\mathrm{T}} \Omega U = \Omega$ $\left(\Omega = \begin{pmatrix} 0 & \mathbb{I} \\ \mathbb{I} & 0 \end{pmatrix}\right)$
- Fermion (quark) fields represented by vectors acted on by some representation of these
  - Fundamental representation: N elements
  - Higher representations (antisymmetric, adjoint, …): M elements, $M > N$

Codes for Lattice QFT

HiRep Pica et al.

- Started ~2007
- C, with C++/Perl code generator
- Heavy use of preprocessor macros
- Support general SU(N), SO(N)
  - Sp(2N) theories in fork, since 2018
- Supports general 2-index rep
- Historically CPU-only

Grid Boyle et al.

- Started ~2013
- C++
- Heavy use of templating, expression templates
- Supports SU(N) (esp. SU(3))
  - Sp(2N) theories in fork, since 2023
- Supports general 2-index fermion representation
  - Bias towards fundamental rep
- CPU (AVX2/512/QPX), GPU (CUDA, HIP, SYCL)

SOMBRERO

Jarno Rantaharju,
Supercomputing Wales

- Benchmarks the CG inversion of the Dirac operator
  - Primary consumer of FLOPs in production code
  - Fixed iteration count
- Tests six different theories (varying gauge group, fermion rep)
  - Changes the relative demands on on-node compute, inter-node comms
- Strips down HiRep
  - Removes code generator
  - Removes unused function
  - <10k SLOC
- First version released 2018

Shoplifter

Michele Mesiti, EXALAT

- Extracting by hand is time-consuming
- Instead, use pycparser to extract and prune call graph
- Automatically generate headers and remove any dead code
- Shoplifter then generates a packaged SOMBRERO

Grid benchmarks

- Extensive set of unit benchmarks built into Grid
  - Dirac operator application (at various volumes, fermion formulations)
  - Disk I/O
  - Memory bandwidth

DiRAC Technical Commissioning

Michele Mesiti + EB, EXALAT

- DiRAC procured new systems in 2021
  - Tursa: 112 nodes, 4 × A100 40GB, 4 × Infiniband HDR
  - COSMA8: 360 nodes, 2 × AMD 7H12 + 1TB RAM
  - DIaL3: 200 nodes, 2 × AMD 7742
  - CSD3: 267 nodes, 2 × Intel Ice Lake + 10 nodes, 4 × A100 80GB
- Initial acceptance testing done by vendors + DiRAC RSE team
  - 2 × committed performance achieved for Grid on Tursa
- Next: community tests software
- Aim: Test both SOMBRERO and Grid on Tursa, COSMA8, DIaL3

DiRAC Technical Commissioning: Approach

- Test all available compilers, MPI implementations
  - On 1 node; on many nodes
- Test scaling of best-performing setup
- Collate data from output files to create comparison plots
- Automate as much as possible
  - Shell scripts, Jupyter notebooks
- Push all data and scripts to DiRAC Technical Commissioning repo

Comparing compilers for SOMBRERO

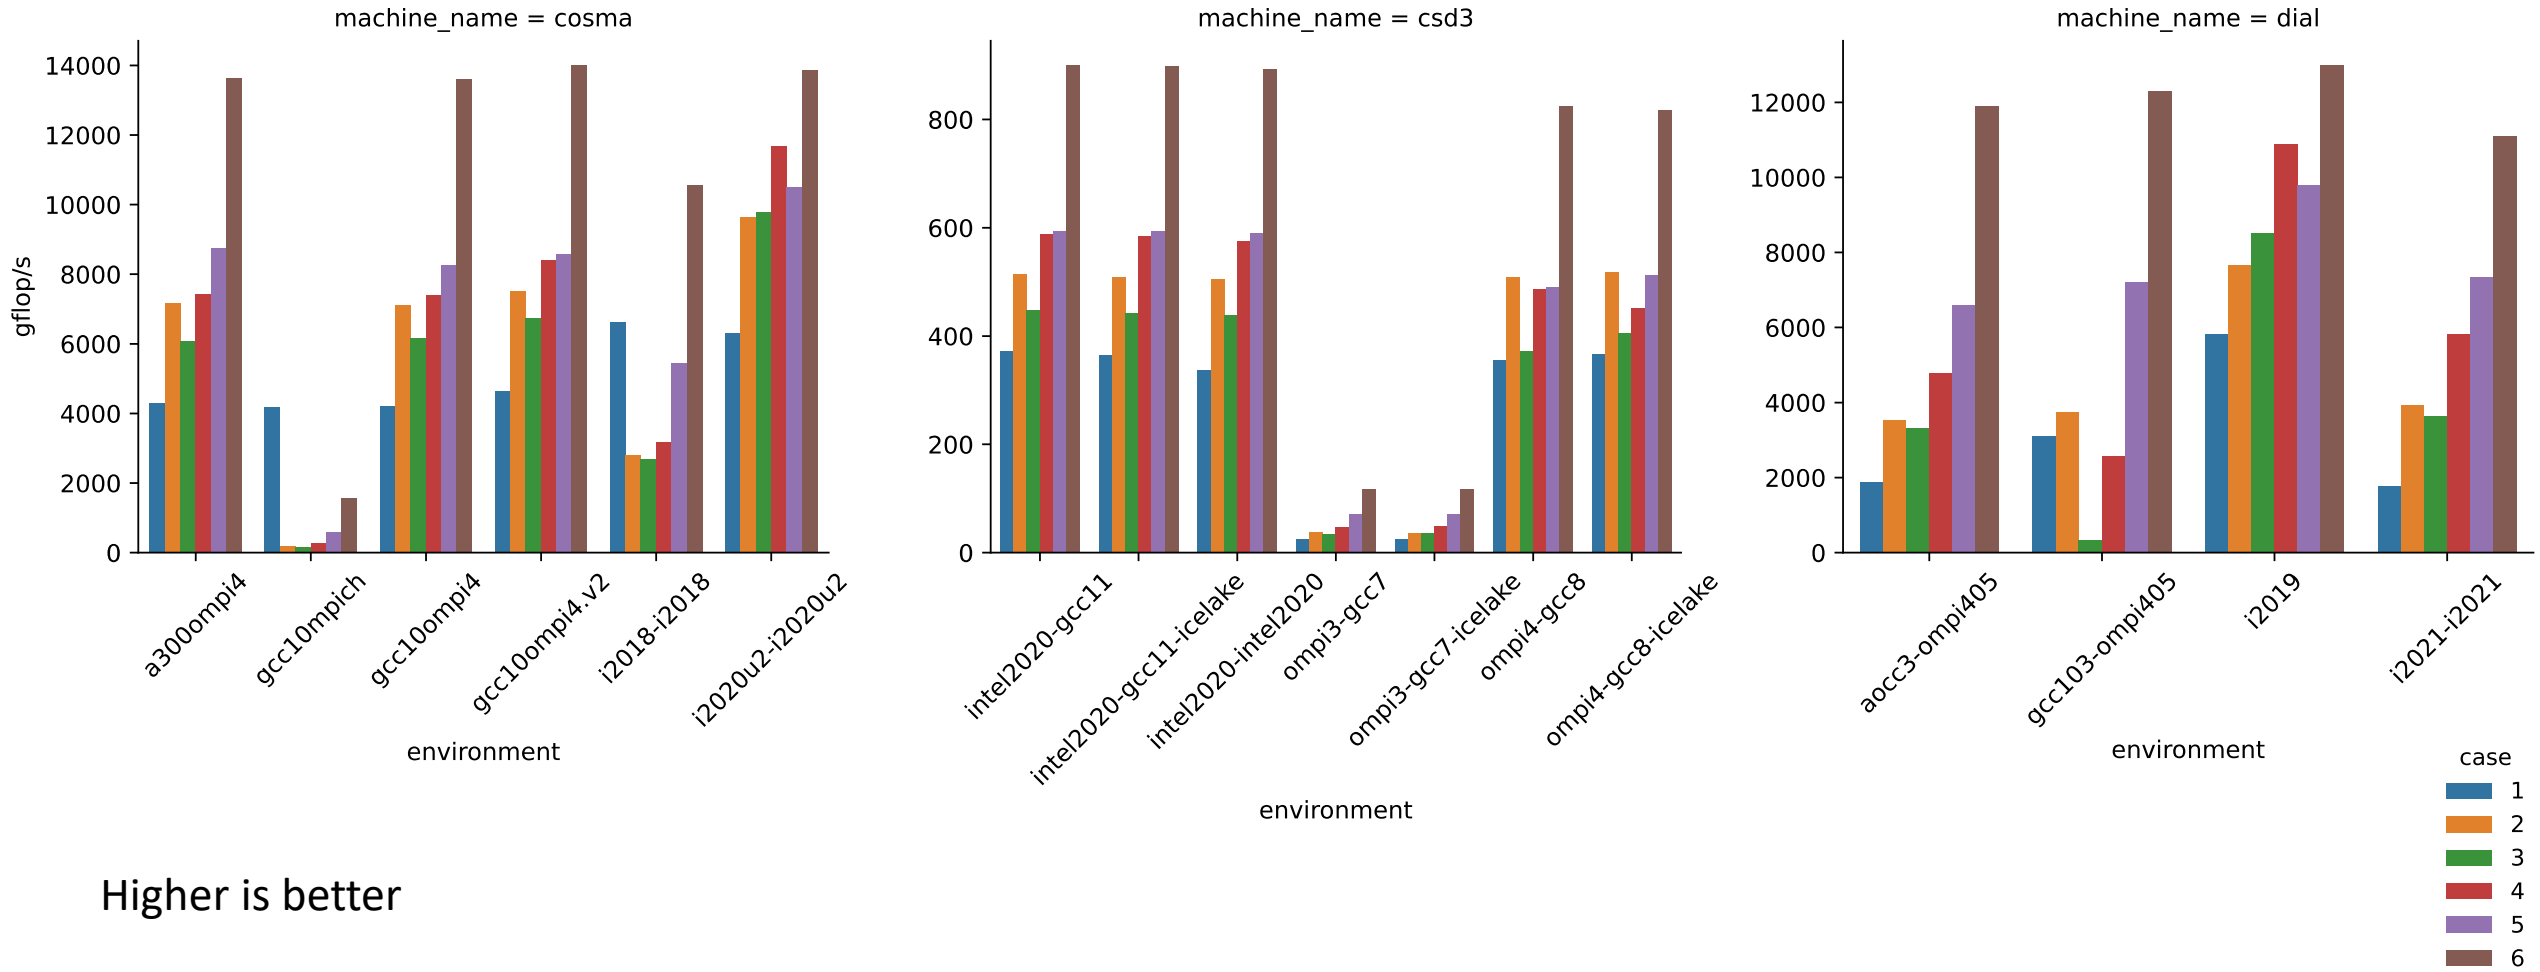

Higher is better

Comparing compilers for Grid

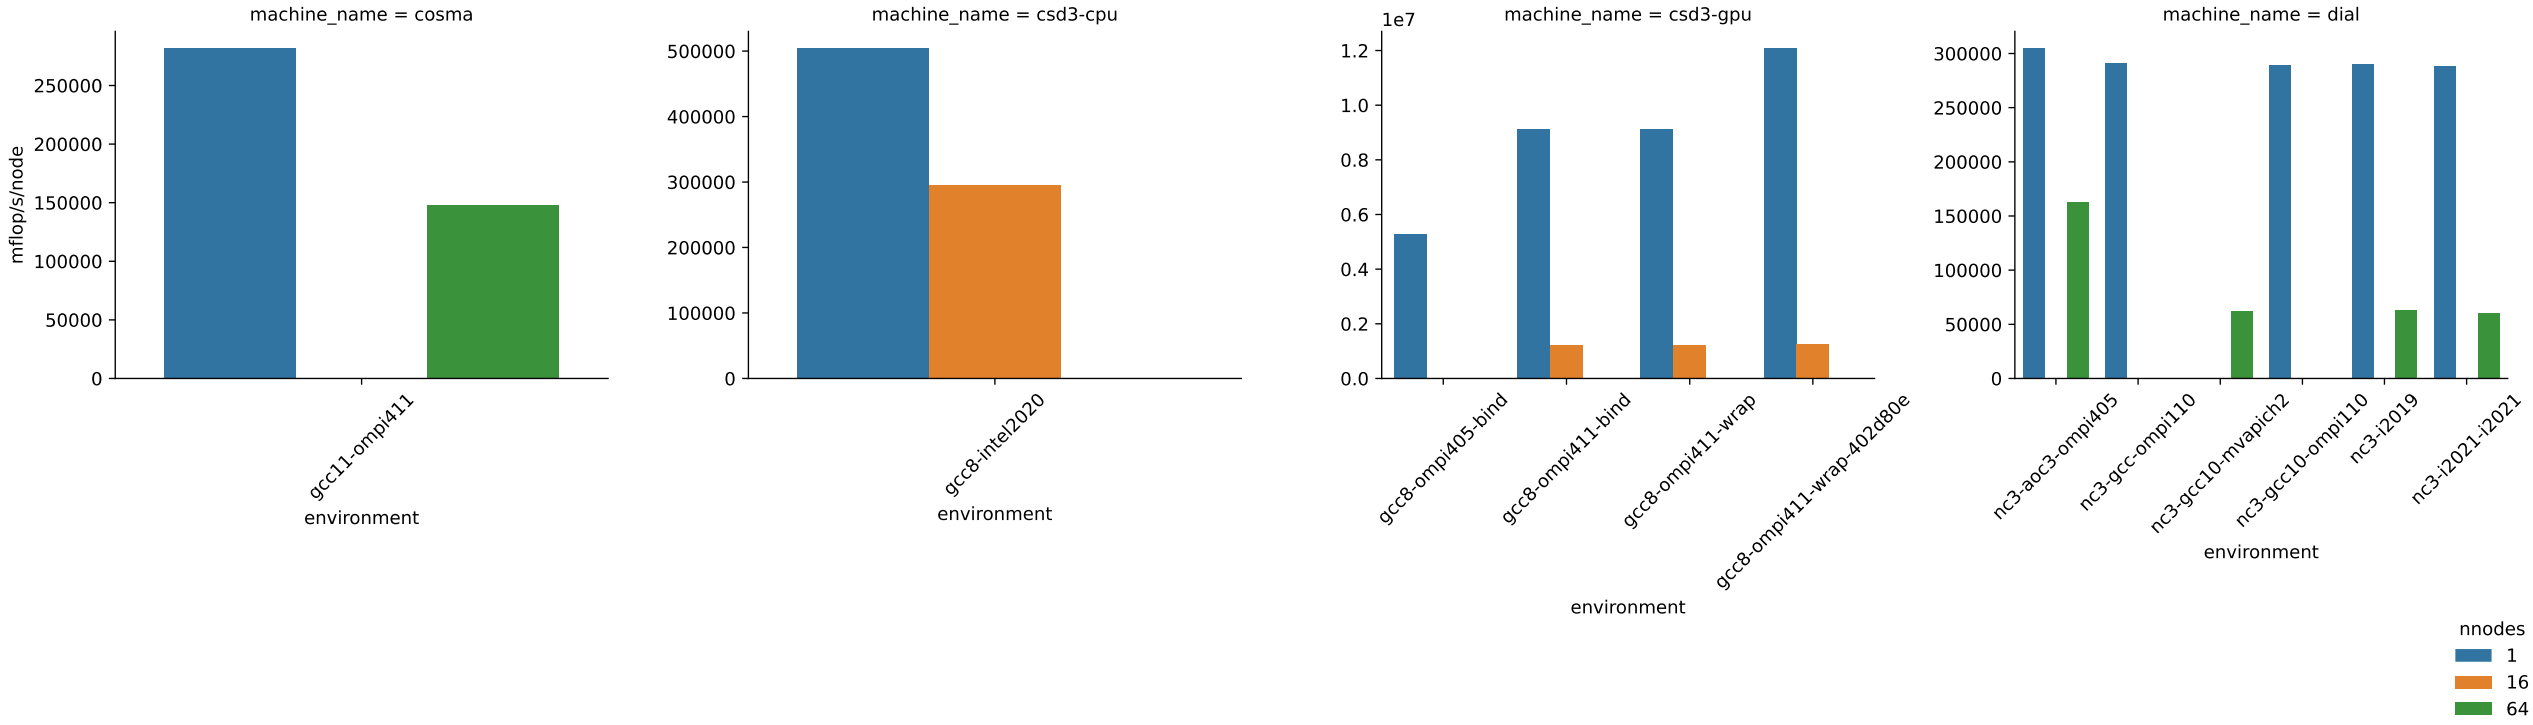

Higher is better

Scaling for SOMBRERO

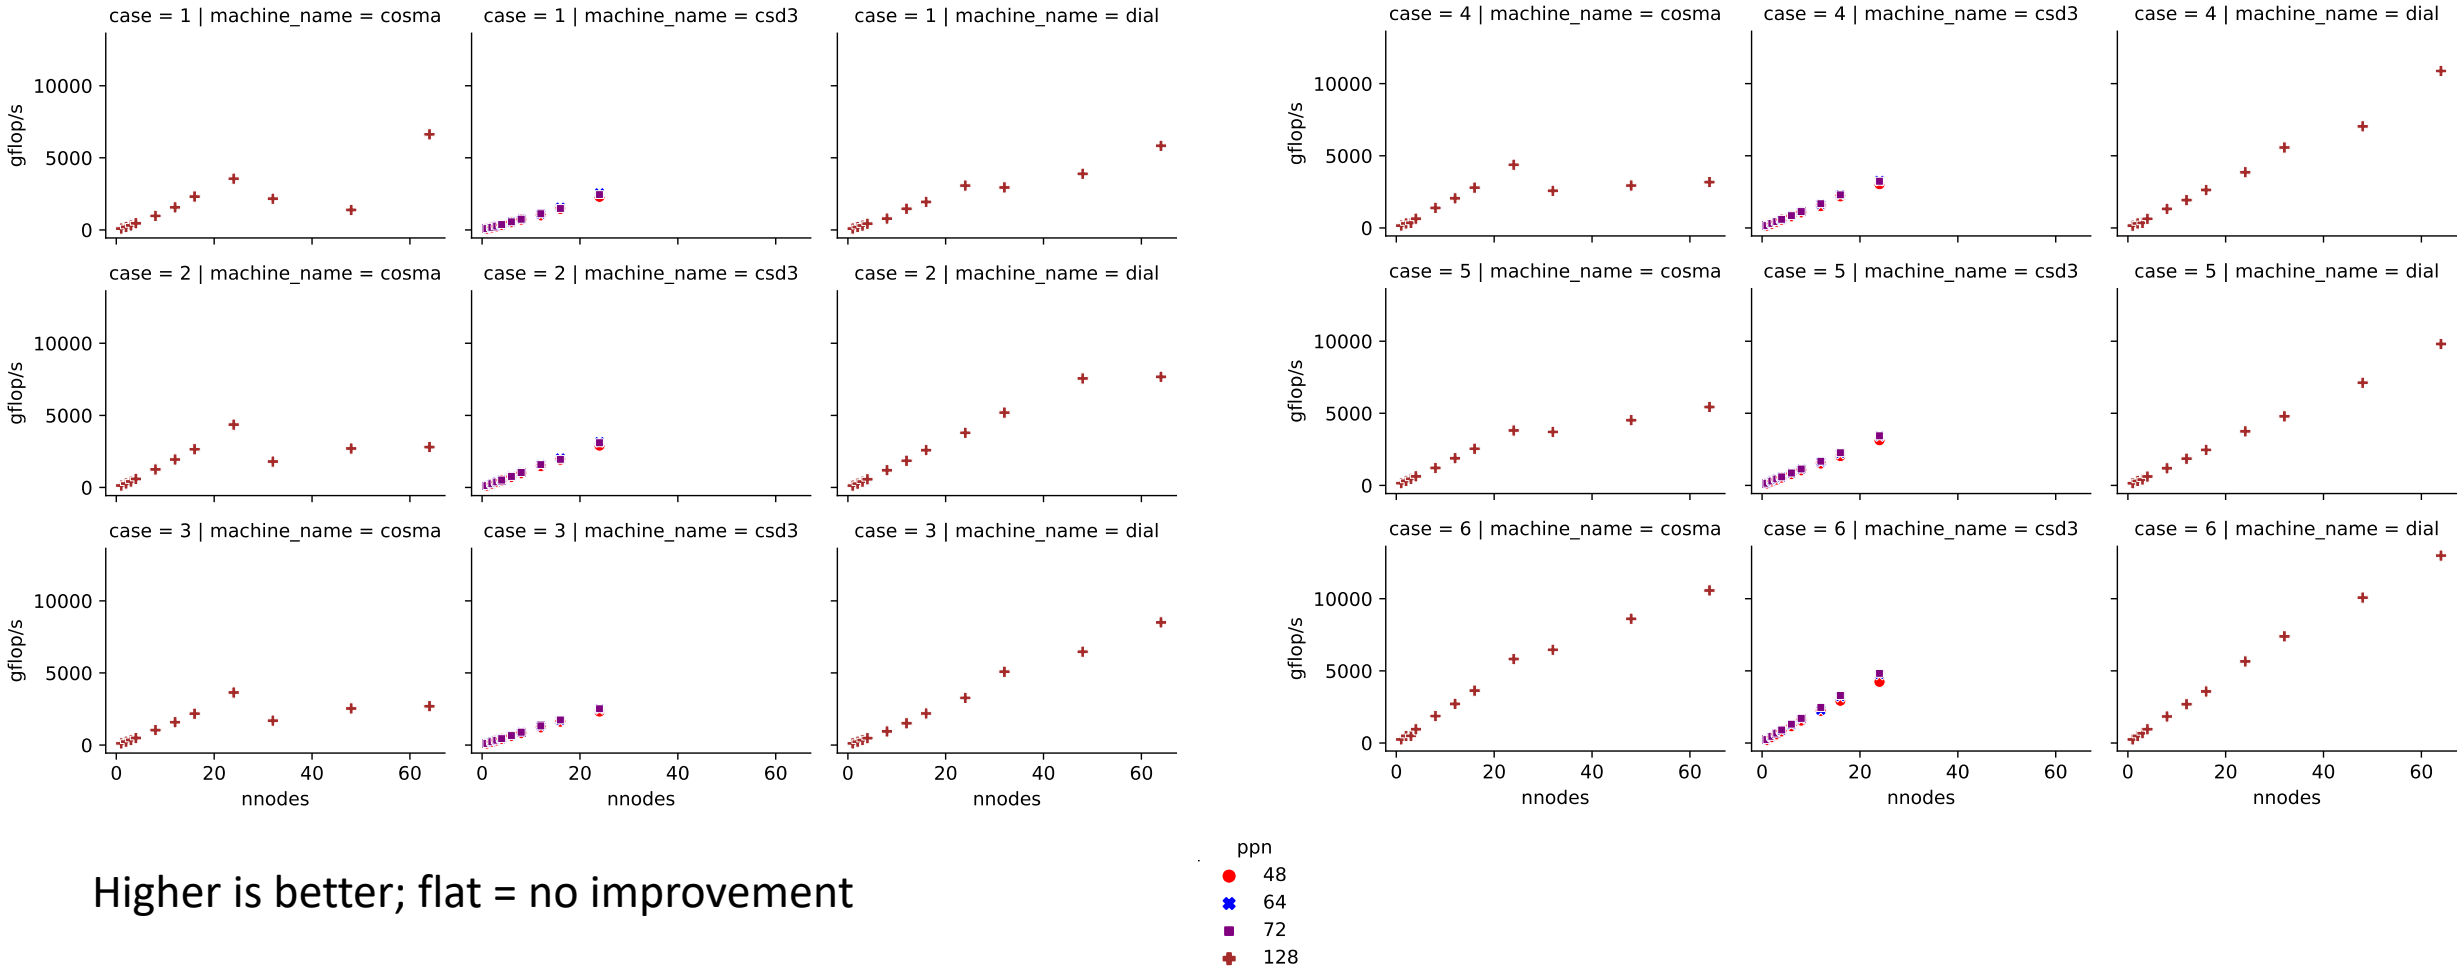

Higher is better; flat = no improvement

Scaling for Grid

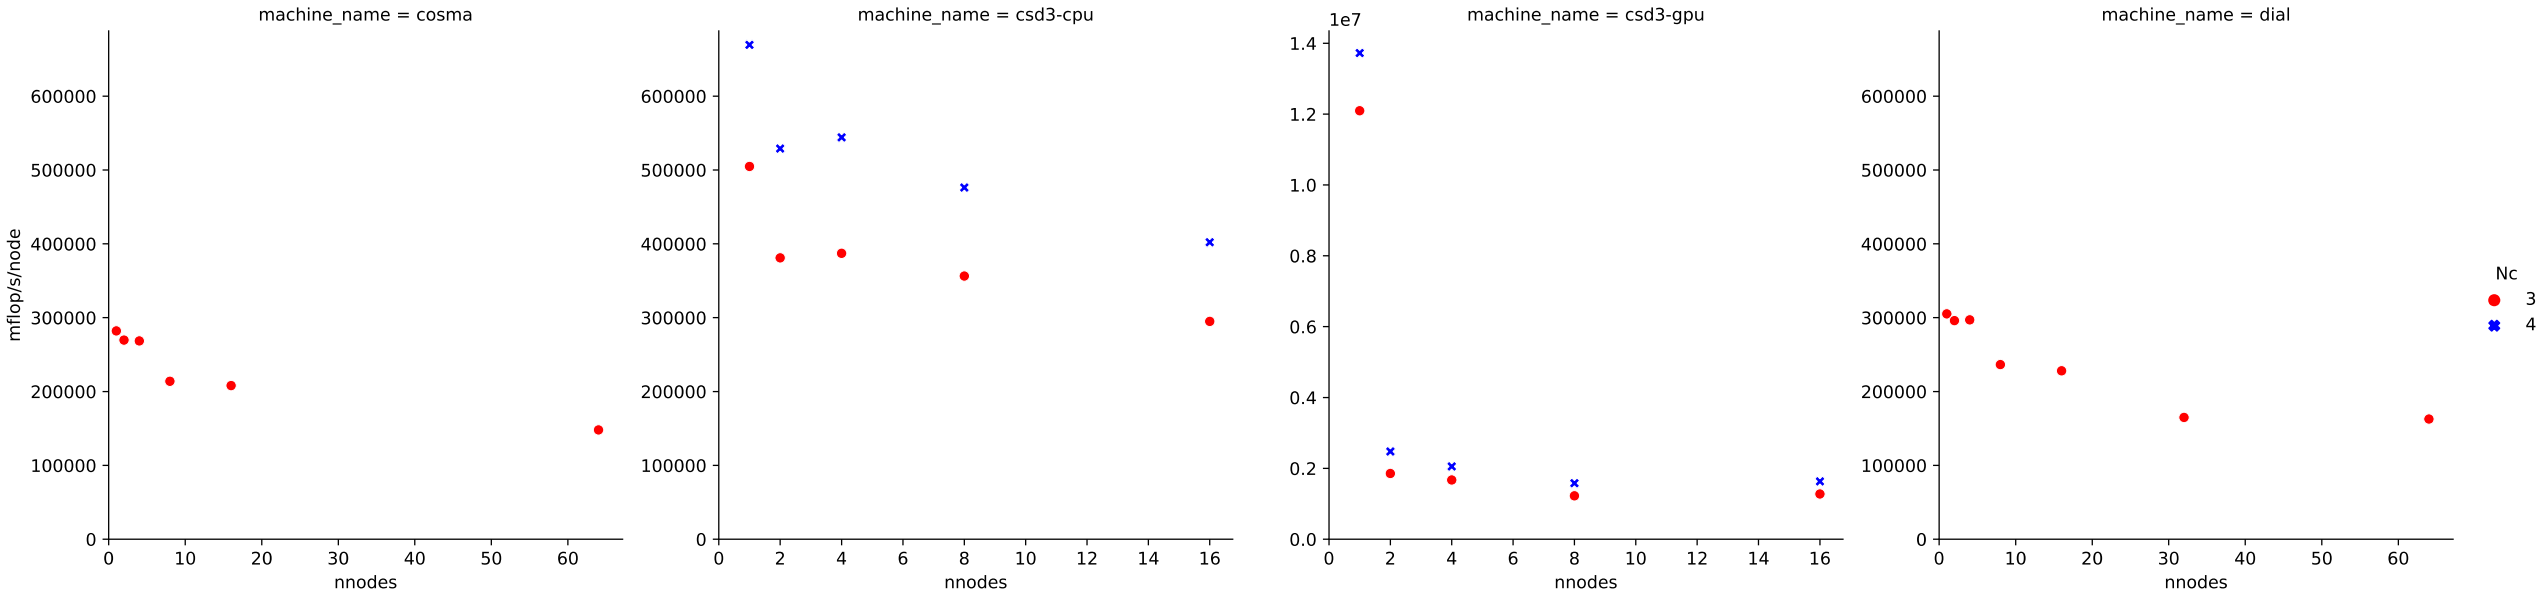

Higher is better; flat = perfect scaling

ExCALIBUR Benchmarking

Michele Mesiti, EXALAT

- Quasi-automated benchmarking suite based on Spack and ReFrame
- SOMBRERO adapted into this framework

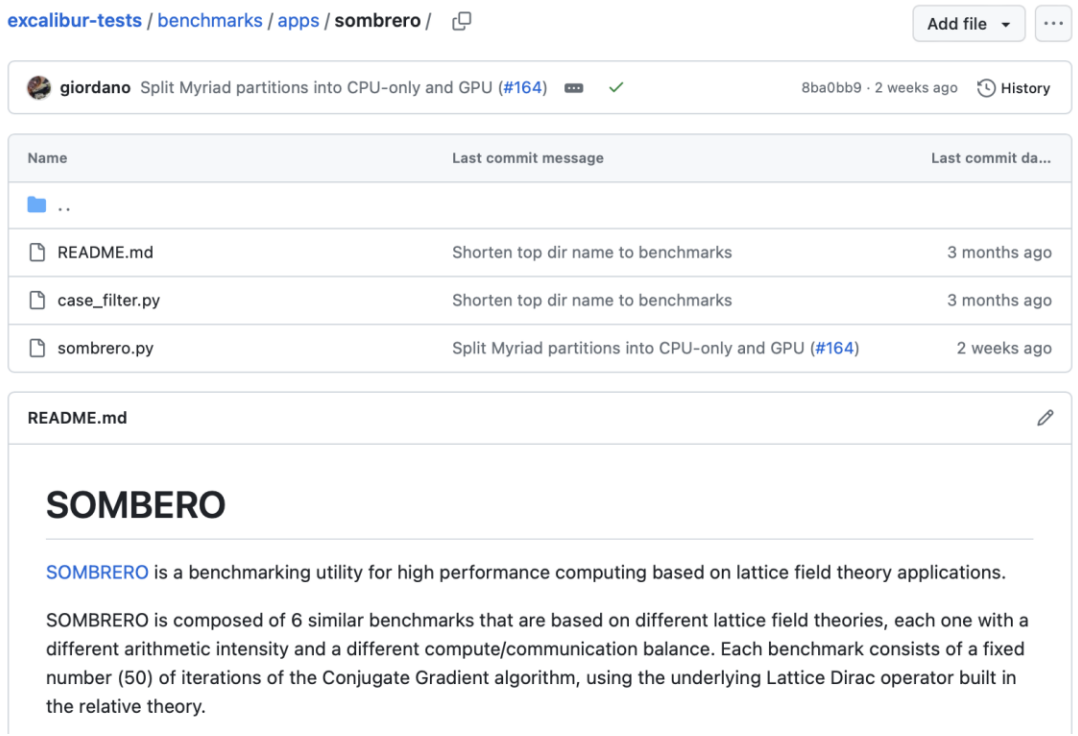

Continuous Benchmark

Andy Sunderland, EXALAT

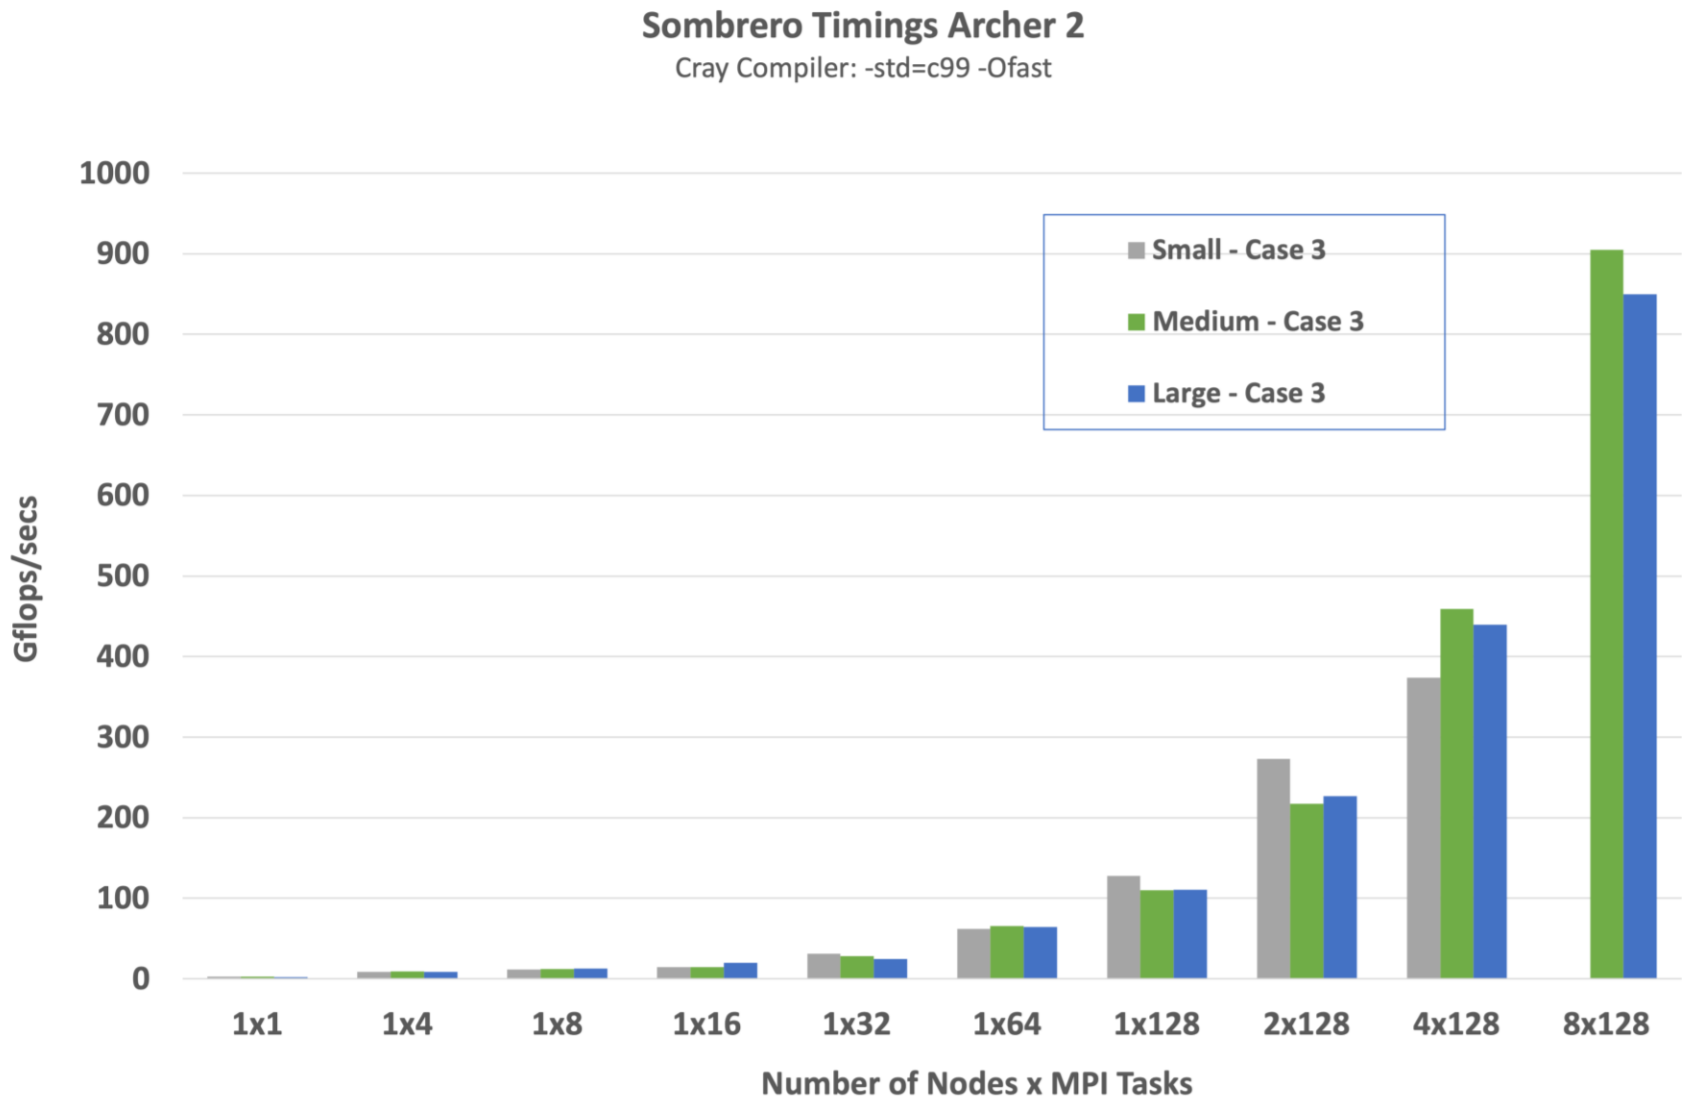

Other outcomes

- Lattice benchmarks (SOMBRERO and predecessors + Grid) have also
  - Identified performance issues on new machine setups
  - Identified compiler bugs on new architectures
  - Contributed to procurement decisions and acceptance tests
  - Been used by vendors internally

Tasks for ExaTEPP

- HiRep development has moved on
  - CUDA branch under active development
  - Shoplifter no longer compatible; needs updating
- Upcoming DiRAC uplift
  - Repeat previous technical commissioning exercise
  - Ideally now with SOMBRERO on GPU as well
- ExCALIBUR test beds
  - Some are compatible with Grid and/or HiRep, so want benchmarking