BSMBench: A HPC Benchmark for Beyond the Standard Model Lattice Physics

Ed Bennett, Biagio Lucini (Swansea U.), Kirk Jordan (IBM), Agostino Patella (CERN), Claudio Pica (CP³-Origins), Antonio Rago (Plymouth U.)

ABSTRACT: *Beyond the Standard Model (BSM) Lattice Physics is a growing area of computational theoretical physics encompassing extensions to/modifications of lattice QCD, which requires increasingly powerful machines. Depending on the theory under investigation, it can place greater demands on either the communications or the compute performance of a multi-node environment, relative to lattice QCD. Lattice QCD has been used to benchmark machines for many years. To allow for more accurate analysis of machines' suitability for a particular theory, as well as a more general analysis of machine's performance than a QCD benchmark would give, we present a new tool to benchmark the performance of some BSM theories, using a method close to that of Lüscher (2002), but based on the HiRep code of Pica, et al. (2009, 2010). Three regimes are probed; one QCD-like regime balancing demands on communications and compute power, and two emphasising each of those over the other. Some initial benchmark statistics are included for clusters, Blue Gene/P, and Blue Gene/Q machines.*

What is Lattice QCD?

- **Q**uantum **C**hromo**d**ynamics is a Quantum Field Theory (QFT) of the strong interaction, which binds atomic nuclei together
- It cannot be solved analytically
- A numerical solution is possible after discretising spacetime onto a four-dimensional lattice of discrete points.

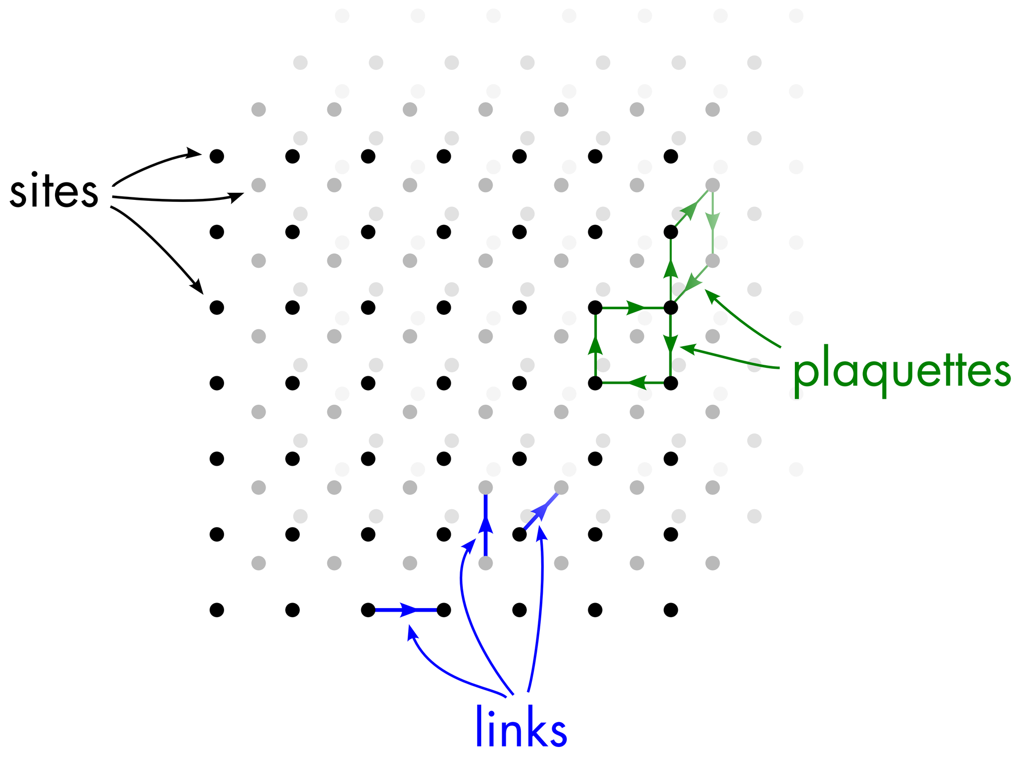

What is BSM?

- **B**eyond the **S**tandard **M**odel physics seeks to generalise QCD to explain new physics
- Fundamental variables are:
  - an $SU(N)$ matrix on each link (QCD fixes $N = 3$)
  - a vector (spinor field) comprising either $4N$ or $4(N^2 - 1)$ values on each site; QCD has the former

Lattice Computations

- Lattice volume (number of sites) V
  - Parallelisation splits lattice into parts, one per process
- Key quantity is the Dirac operator, a large, sparse matrix acting on all spinor variables
- Main computations involve inverting this matrix on selected vectors

Why is HPC necessary?

- Smallest lattices will fit on a desktop (4^4 lattice)
- Current research looks at 128^4 lattices (and beyond) – on a desktop this would require ~150GB RAM, ~10^4 CPU-years per data point
- BSM code can require more power again
- This is using state-of-the-art numerical techniques to minimise the amount of compute power necessary

Benchmarks

- QCD is already used to benchmark supercomputers [1]
- QCD codes place roughly equal demands on communications and compute speed
- BSM codes vary this
- BSM-derived benchmark allows more flexible testing of a machine's characteristics
  - Testing of machine's suitability for a given theory
  - More general communications/compute comparisons

HiRep

- Chroma (described in [2]), the most common publicly-available lattice code, only deals with QCD-like theories
- HiRep (described in [3], [4]) has been developed to be more flexible
  - Allows varying of N, type of spinor field
- Forms a complete, state-of-the-art suite for lattice BSM

Results

sqnorm test

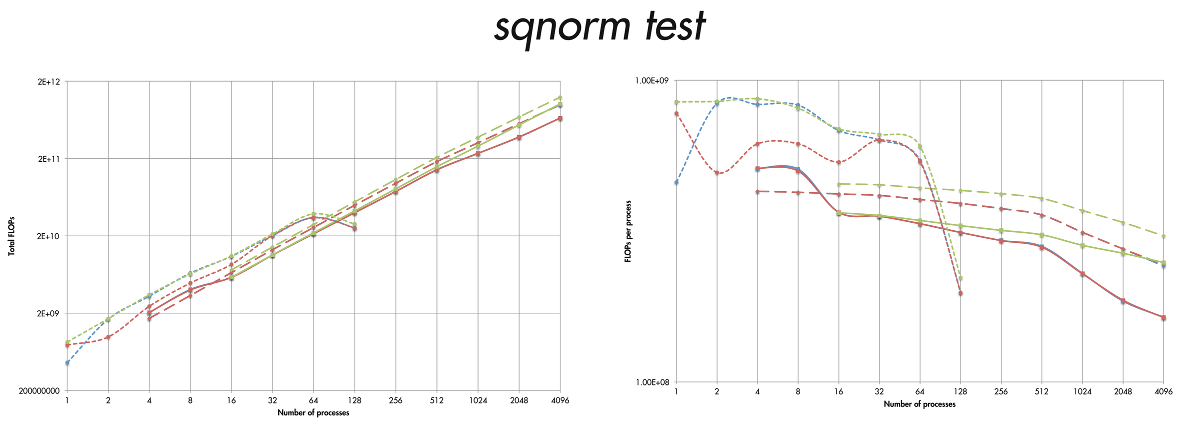

Square norm test demands less of inter-node communications; shows approximate scaling on all machines

Dphi test

Dirac operator application is more physically demonstrative; drop in performance on non-Infiniband cluster highlights higher demands placed on inter-node communications

Blue Gene Comparison

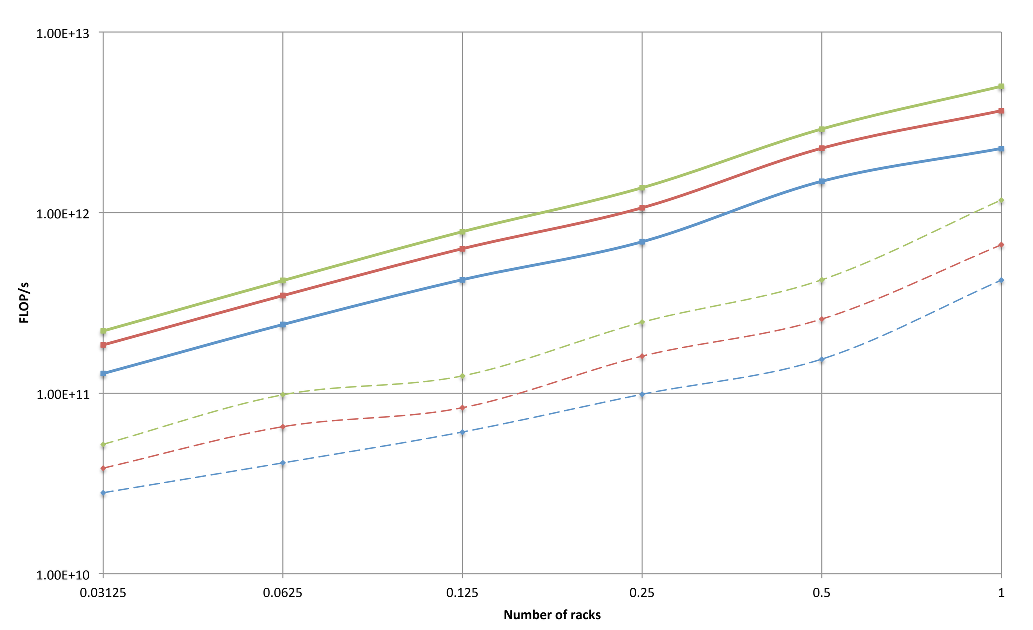

Rack-for-rack comparison between 128×64^3 lattice on Blue Gene/Q and equivalent data from a 64×32^3 lattice on Blue Gene/P

Processor Subdivision on Blue Gene/Q

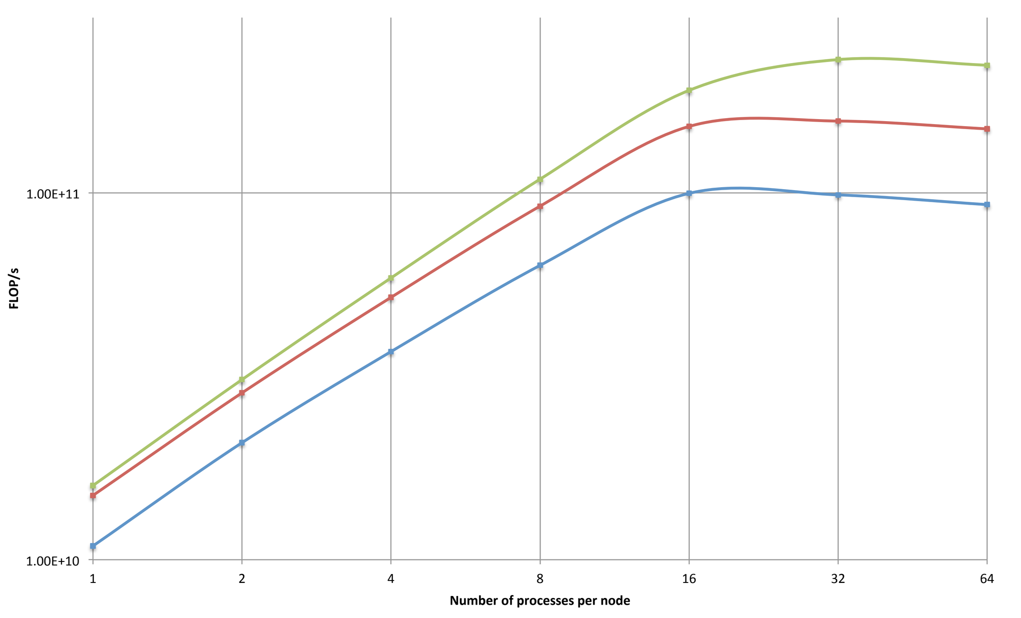

Blue Gene/Q supports four hardware threads per core, with 16 cores per node. Multithreading gives no benefit to the comms or balanced theories, but does give a small benefit to the compute theory

Comms test — Balance test — Compute test

Blue Gene/P — Blue Gene/Q — x86 cluster

Benchmark Strategy

Based on that of Lüscher (described in [5])

- Consistency check of arithmetic
  - Not used for performance analysis
  - Omitted for small machines
- Three operations tested for a given period of time
  - Spinor field square norm (sqnorm)
  - Spinor field multiply-add (mad) (not shown)
  - Dirac operator application (Dphi)
- FLOPs counted, performance reported
- Three regimes tested:
  - "Comms" – communications-intensive $SU(2)$ theory, 2 adjoint fermions
  - "Balance" – QCD-like $SU(3)$ theory, 2 fundamental fermions
  - "Compute" – computationally-intensive $SU(6)$ theory, 2 fundamental fermions
- Single lattice size of 64×32^3 used to allow direct comparison between machine sizes
- Based on the HiRep code

Conclusions

- Three machines tested: Blue Gene/P, Blue Gene/Q, and an x86 cluster without Infiniband.
- Square norm less demanding of communications; speed scales well beyond node size on cluster
- Dphi test is communications-intensive:
  - Speed scales well on Blue Gene
  - Speed drops off once exceeding node size of non-Infiniband cluster
- BGQ 4–5 times faster than BGP per rack
- Multithreading cores can have modest benefit to some theories

github.com/blucini/BSMBench

Architectures

Three different machines were tested

- Blue Gene/P:
  - 4 cores/node, 1 thread/core, 32-bit PowerPC CPU at 850MHz, 1 MPI process per core
  - Nodes connected by high-speed 3D torus
  - 4GB RAM, 8kB L2, 4MB L3 cache per node
  - IBM XL C compiler 9.0 for Blue Gene
- Blue Gene/Q:
  - 16 compute cores/node, 4 threads/core, 64-bit PowerPC CPU at 1.6GHz, 1 MPI process per core
  - Nodes connected by high-speed 5D torus
  - 16GB RAM, 32MB L2 cache per node
  - IBM XL C compiler 12.0 for Blue Gene
- x86 cluster:
  - 16 cores/node, 1 thread/core, 2 × Opteron 6128 CPU/node, 1 MPI process per core
  - Nodes connected by 1GigE (no Infiniband)
  - 96GB RAM, 8MB L2, 24MB L3 cache per node
  - GCC 4.1.2

References

1: P. Vranas, G. Bhanot, M. Blumrich, D. Chen, A. Gara, P. Heidelberger, V. Salapura, J. Sexton, *The BlueGene/L supercomputer and quantum ChromoDynamics*, Proceedings of the 2006 ACM/IEEE conference on Supercomputing (SC06) (2006)

2: R. G. Edwards (LHPC Collaboration), B. Joó (UKQCD Collaboration), *The Chroma Software System for Lattice QCD*, Proceedings of the 22nd International Symposium for Lattice Field Theory (Lattice2004), Nucl. Phys B1 (Proc. Suppl) 40 (2005) 832

3: L. Del Debbio, A. Patella, C. Pica, *Higher representations on the lattice: Numerical simulations. SU(2) with adjoint fermions*, Phys. Rev. D81 (2010) 094503

4: L. Del Debbio, B. Lucini, A. Patella, C. Pica, A. Rago, *Conformal versus confining scenario in SU(2) with adjoint fermions*, Phys. Rev. D80 (2009) 074507

5: M. Lüscher, *Lattice QCD on PCs?*, Nucl. Phys. B (Proc. Suppl.) 106 (2002) 21

Swansea University Prifysgol Abertawe — Research with Plymouth University — CP³ - Origins Particle Physics & Cosmology — UKQCD collaboration — CERN — IBM