Performance of the Eurora cluster for Lattice QCD calculations

Mario Schröck, Silvano Simula
INFN Roma Tre

Motivation

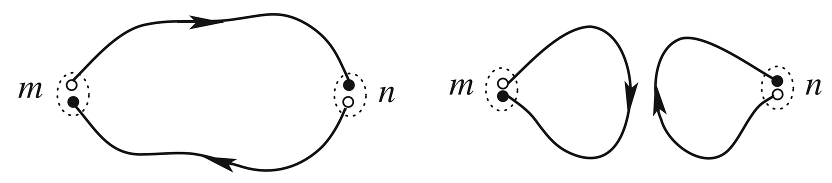

- inclusion of disconnected quark loops in lattice QCD requires *ab initio* the complete inversion of a rank $\geq$ one million matrix
- clever algorithms lower this to 100-1000 solutions of linear equation systems per gaugefield configuration
- costs of the inversions still highly dominate the post gaugefield generation analysis

$\Rightarrow$ adopt modern hardware to accelerate the inversions!

The Eurora cluster @ Cineca

- Eurora - Eurotech Aurora HPC 10-20
- Xeon E5-2687W 8C 3.100GHz
- Infiniband QDR
- 64 NVIDIA Tesla K20 (Kepler)
- 64 Intel Xeon Phi (MIC)
- Linpack Performance (R_{max}): 100.9 TFlop/s [1]
- Theoretical Peak (R_{peak}): 175.667 TFlop/s [1]
- No. 1 of the Green500 List in June 2013 [2]

Code package I: Quda

- "QCD on CUDA" [3, 4] `http://lattice.github.com/quda`
- Effort started at Boston University in 2008, now in wide use as the GPU solver backend for Chroma, MILC, and various other codes.
- Various solvers for several discretizations, including multi-GPU support and domain-decomposed (Schwarz) preconditioners.

Code package II: cuLGT

- "CUDA Lattice Gauge Theory" [5] `http://www.cuLGT.com`
- Evolved since 2010, developed in Graz and Tübingen.
- Main focus lies on lattice gauge fixing (Coulomb, Landau and maximally Abelian gauge) but a very general, object oriented infrastructure for lattice QCD calculations on GPUs is offered.

Single GPU performance (Tesla K20)

- lattice size 32^4
- Quda: twisted mass conjugent gradient inverter (nondegenerate doublet of quark flavours)
- cuLGT: Landau gauge fixing with the overrelaxation algorithm

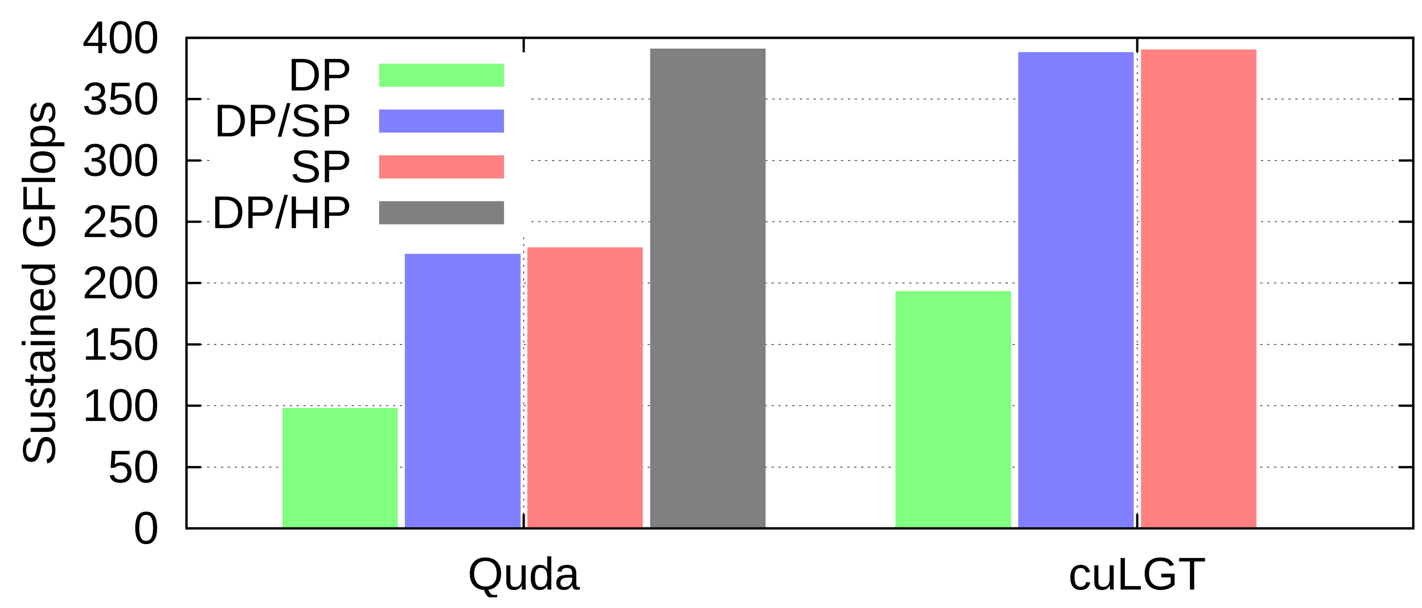

Figure 1: The different colors correspond to double precision (DP), mixed double/single precision (DP/SP), single precision (SP) and mixed double/half precision (DP/HP) (only Quda).

Comparison to IBM BlueGene/Q

- Fermi @ Cineca: BlueGene/Q (10.240 nodes with 16 cores each)
- Processor Type: IBM PowerA2, 1.6 GHz
- setup which is in production at Roma Tre:
  - $32^3 \times 64$ lattice: twisted mass inverter with mixed double/single precision and SSE vector instructions
  - **128 Fermi nodes (2048 CPU cores)**
- compare to:
  - same lattice, equivalent inverter (Quda) with mixed double/half precision
  - **one Eurora node (two GPUs)**
- ratio of average time per inverter iteration:

$$\frac{t_{\text{Fermi}}}{t_{\text{Eurora}}} = \frac{0.009898s}{0.017016s} = 0.58$$

Multi GPU performance

- Quda twisted mass inverter with one flavour of quarks
- double/half mixed precision and 12 parameter reconstruction

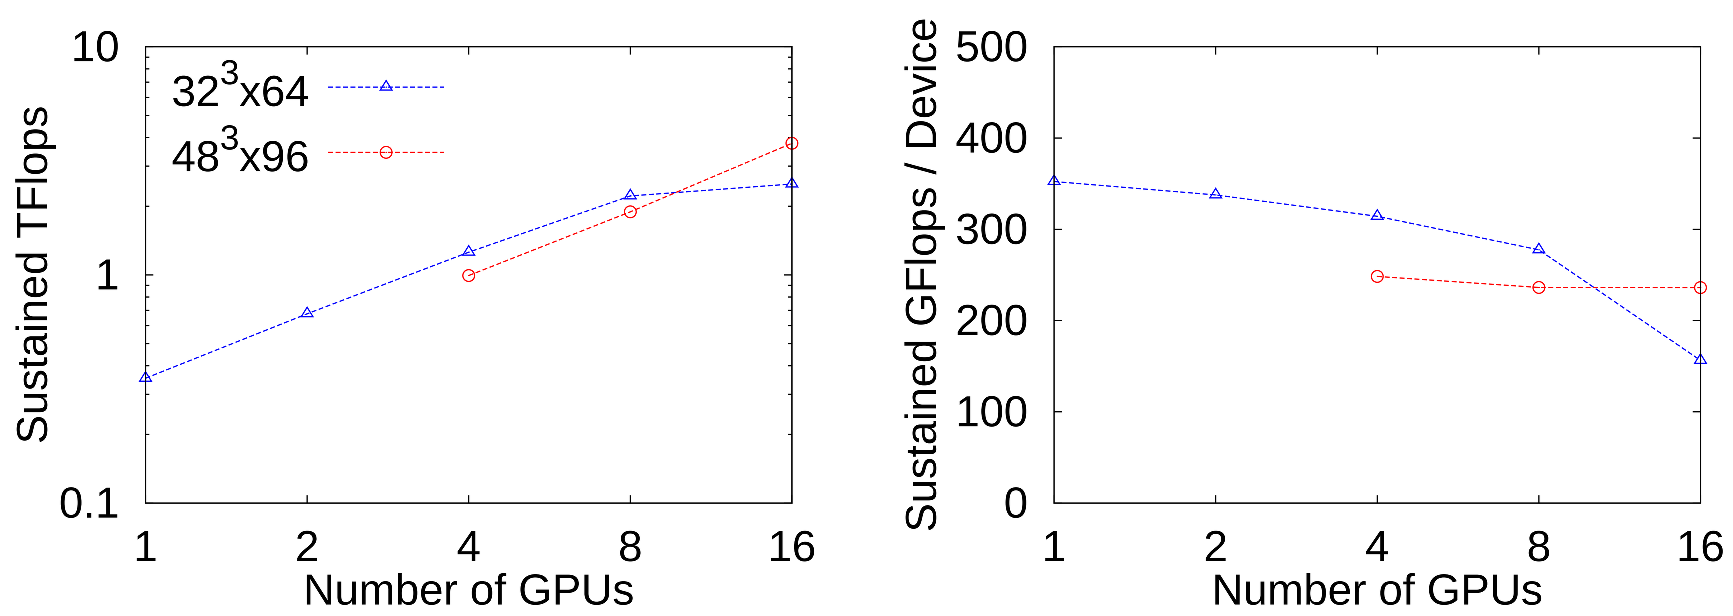

Figure 2: Strong scaling for two different lattice sizes. On the left the total performance and on the right the performance per GPU.

Conclusions

- GPUs can efficiently accelerate the main bulk of QCD simulations (large matrix inversions)
- highly optimized software packages are available for many lattice QCD problems (Quda, cuLGT, QDP-JIT and more)
- a single node of Eurora offers the same magnitude of performance as 128 nodes of the Fermi BlueGene/Q cluster for the inversions of lattice sizes currently in production at Roma Tre
- weak point is the size of the memory: use as many nodes/GPUs as necessary to accommodate the problem at hand
- our production setup ($32^3 \times 64$ and $48^3 \times 96$, respectively) requires the use of up to eight nodes (16 GPUs) to match the memory needs
- outlook: we will greatly profit from the larger memory of the Tesla K40 (fewer GPUs necessary to fit the problem).

References

[1] `http://www.top500.org/system/178077`.

[2] `http://www.green500.org/lists/green201306`.

[3] M. Clark *et al*, *Comput.Phys.Commun.* **181** (2010) 1517. `arXiv:0911.3191`.

[4] A. Strelchenko *et al*, *PoS LATTICE* (2013) 415. `arXiv:1311.4462`.

[5] M. Schröck and H. Vogt, *Comput.Phys.Commun.* **184** (2013) 1907. `arXiv:1212.5221`.