

Landau gauge fixing on the lattice using GPU's

Nuno Cardoso

*CFTP, Departamento de Física, Instituto Superior Técnico, Universidade Técnica de Lisboa,
Avenida Rovisco Pais 1, 1049-001 Lisbon, Portugal*

E-mail: nuno.cardoso@ist.utl.pt

Paulo J. Silva*

*CFC, Departamento de Física, Faculdade de Ciências e Tecnologia, Universidade de Coimbra,
3004-516 Coimbra, Portugal*

E-mail: psilva@teor.fis.uc.pt

Orlando Oliveira

*CFC, Departamento de Física, Faculdade de Ciências e Tecnologia, Universidade de Coimbra,
3004-516 Coimbra, Portugal*

E-mail: orlando@teor.fis.uc.pt

Pedro Bicudo

*CFTP, Departamento de Física, Instituto Superior Técnico, Universidade Técnica de Lisboa,
Avenida Rovisco Pais 1, 1049-001 Lisbon, Portugal*

E-mail: bicudo@ist.utl.pt

In this work, we consider the GPU implementation of the steepest descent method with Fourier acceleration for Landau gauge fixing, using CUDA. The performance of the code in a Tesla C2070 GPU is compared with a parallel CPU implementation.

Xth Quark Confinement and the Hadron Spectrum

8-12 October 2012

TUM Campus Garching, Munich, Germany

*Speaker.

On the lattice, Landau gauge is defined through the maximization of the functional

$$F_U[g] = \frac{1}{N_d N_c V} \sum_x \sum_\mu \text{Re} [\text{Tr} (g(x) U_\mu(x) g^\dagger(x + \mu))] ,$$

where N_d is the dimension of the space-time, N_c is the dimension of the gauge group and V the lattice volume, on each gauge orbit. The functional $F_U[g]$ can be maximised using a steepest descent method [1, 2]. However, when the method is applied to large lattice volumes, it faces the problem of critical slowing down, which can be attenuated by Fourier acceleration.

The main goal of this work is to compare the difference in performance between GPU and CPU implementations of the Fourier accelerated Landau gauge fixing method. The MPI parallel version of the algorithm was implemented in C++, using the machinery provided by the Chroma library [3]; for the Fourier transforms, the code uses PFFT, a parallel FFT library written by Michael Pippig [4]. For the GPU code, we used version 4.1 of CUDA [5] – see also [6]; FFT are performed using the CUFFT library by NVIDIA [7].

For such a comparison, we use a NVIDIA Tesla C2070. The GPU code has been run using a 12 real number representation; furthermore, we used texture memory and we switched ECC off. The CPU code has been run in the Centaurus cluster, at Coimbra. In Centaurus, each node has 2 Intel Xeon E5620@2.4 GHz (quad core), with 24 GB of RAM, and it is equipped with a DDR Infiniband network.

In order to compare the performance of the two codes, we used a 32^4 lattice volume. The configurations have been generated using the Wilson gauge action, with three different values of β . The runs used $\alpha = 0.08$ and $\theta < 10^{-15}$.




Figure 1: Strong scaling CPU tested for a 32^4 lattice volume and comparison with the GPU for the best performance, 12 real number parametrization, ECC Off and using texture memory in double precision, [8]. In Centaurus, a cluster node means 8 computing cores.

In figure 1 we compare the performance of the parallel CPU code against the GPU result. The CPU code shows a good strong scaling behaviour, with a linear speed-up against the number of computing nodes. However, the GPU code was much faster: in order to reproduce the performance of the GPU code, one needs 256 CPU cores.

For more details on this work, please see [8, 9, 10].

Acknowledgments

We thank Bálint Joó and Michael Pippig for discussions about Chroma and PFFT libraries respectively. In particular, we thank Michael Pippig for extending his library for 4-dimensional FFTs.

This work was partly funded by the FCT contracts POCI/FP/81933/2007, CERN/FP/83582/2008, PTDC/FIS/100968/2008, CERN/FP/109327/2009, CERN/FP/116383/2010, CERN/FP/123612/2011, projects developed under initiative QREN financed by UE/FEDER through Programme COMPETE. Nuno Cardoso and Paulo Silva are supported by FCT under contracts SFRH/BD/44416/2008 and SFRH/BPD/40998/2007 respectively. We would like to thank NVIDIA Corporation for the hardware donation used in this work via Academic Partnership program.

References

- [1] C. Davies, G. Batrouni, G. Katz, A. S. Kronfeld, G. Lepage, et al., Fourier acceleration in lattice gauge theories. I. Landau gauge fixing, *Phys.Rev. D*37 (1988) 1581.
[doi:10.1103/PhysRevD.37.1581](https://doi.org/10.1103/PhysRevD.37.1581).
- [2] O. Oliveira, P. J. Silva, A global optimization method for Landau gauge fixing in lattice QCD, *Comp. Phys. Comm.* 158 (2004) 73–88. [arXiv:hep-lat/0309184](https://arxiv.org/abs/hep-lat/0309184),
[doi:10.1016/j.cpc.2003.12.001](https://doi.org/10.1016/j.cpc.2003.12.001).
- [3] R. G. Edwards, B. Joó, The Chroma software system for lattice QCD, *Nucl. Phys. Proc. Suppl.* 140 (2005) 832.
- [4] M. Pippig, PFFT - An extension of FFTW to Massively Parallel Architectures, Department of Mathematics, Chemnitz University of Technology, 09107 Chemnitz, Germany Preprint 06.
- [5] NVIDIA, NVIDIA CUDA C Programming Guide, version 4.1 Edition (2011).
- [6] N. Cardoso, P. Bicudo, Generating SU(Nc) pure gauge lattice QCD configurations on GPUs with CUDA, *Comput.Phys.Commun.* 184 (2013) 509–518. [arXiv:1112.4533](https://arxiv.org/abs/1112.4533),
[doi:10.1016/j.cpc.2012.10.002](https://doi.org/10.1016/j.cpc.2012.10.002).
- [7] NVIDIA, CUDA Toolkit 4.1 CUFFT Library (2012).
- [8] N. Cardoso, P. J. Silva, P. Bicudo, O. Oliveira, Landau Gauge Fixing on GPUs, *Computer Physics Communications* 184 (1) (2013) 124–129. [arXiv:1206.0675](https://arxiv.org/abs/1206.0675),
[doi:10.1016/j.cpc.2012.09.007](https://doi.org/10.1016/j.cpc.2012.09.007).
- [9] N. Cardoso, P. Bicudo, P. J. Silva, O. Oliveira, Landau Gauge Fixing on GPUs and String Tension, *Acta Phys.Polon.Supp.* 5 (2012) 1135–1141. [arXiv:1208.1186](https://arxiv.org/abs/1208.1186),
[doi:10.5506/APhysPolB.5.1135](https://doi.org/10.5506/APhysPolB.5.1135).
- [10] N. Cardoso, P. J. Silva, O. Oliveira, P. Bicudo, GPU implementation of a Landau gauge fixing algorithm, *PoS(Lattice 2012)* 245. [arXiv:1210.8018](https://arxiv.org/abs/1210.8018).