The Thermal Quark Hadron Transition in Lattice QCD with two Quark Flavours

Owe Philipsen[1], **Matthias Bach**[2], **Volker Lindenstruth**[2], **and Christopher Pinke**[1]

[1] Institut für Theoretische Physik, Goethe Universität Frankfurt, Germany
E-mail: {philipsen, pinke}@th.physik.uni-frankfurt.de

[2] Institut für Informatik, Goethe Universität Frankfurt, Germany
and
Frankfurt Institute for Advanced Studies, Frankfurt, Germany
E-mail: bach@compeng.uni-frankfurt.de

We investigate the thermal transition from nuclear matter to a quark gluon plasma by simulations of lattice QCD with two quark species using the Wilson twisted mass formulation. A new code based on OpenCL is presented, which can be used on either CPU or GPU of any vendor. To optimise efficiency, we thermalise our Monte Carlo simulations using highly parallelised code on CPU machines, whereas production runs are done using multiple Monte Carlo chains on a GPU cluster, fitting an entire lattice on each GPU with zero overhead for communication. We discuss preliminary results for the thermal transition obtained in this way.

1 Introduction

The fundamental theory of the strong interactions governing nuclear and subnuclear forces is Quantum Chromodynamics (QCD). Its fundamental degrees of freedom are quarks and gluons, which combine into numerous tightly bound states, the hadrons, among them the familiar nucleons. A key feature of the theory is asymptotic freedom, according to which the coupling strength depends on the energy scale of a scattering process. For energies below 1 GeV, the coupling is too large to allow for weak coupling approximations and analytic predictions. On the other hand, at large temperatures the average energy per particle is large and the theory enters a weak coupling regime, where the quarks and gluons form a plasma rather than bound states. The transition from the hadronic to the plasma regime takes place at temperatures of about 170 MeV and still belongs to strong coupling physics. However, the theory can be reformulated on a space-time lattice, whereupon it is amenable to Monte Carlo simulations.

In this contribution we present a study of the thermal QCD transition with two mass-degenerate quark species. The main interest is in the critical temperature and the nature of the transition. In the limit of massless quarks, the theory has a chiral symmetry under mixing of the quark species. This symmetry is broken spontaneously, so that the lightest particles, the pions, are exactly massless. At finite temperature a non-analytic transition takes place in which this symmetry gets restored. Since massless QCD cannot be simulated, the order of this phase transition is not known to date. On the other hand, for finite quark mass the chiral symmetry is explicitly broken and the pions are massive. The phase transition then gets weakened to an analytic crossover. A first order chiral transition disappears gradually in a $Z(2)$ critical point, whereas a second order transition disappears immediately for non-zero mass, Fig. 1.

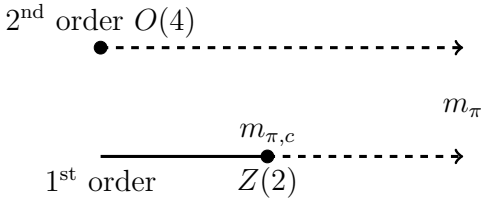

Figure 1. Possible scenarios for the chiral phase transition as a function of pion mass[4]. In the chiral (massless) limit there must be a true transition of first or second order, whereas for finite masses the transition is merely a smooth crossover. The boundary between the two corresponds to a critical point.

2 Twisted Mass Lattice QCD

Consider euclidean spacetime discretised on a hypercube with lattice spacing a. We denote the spatial and temporal extent of the system with N_σ and N_τ, respectively. The QCD action $\mathcal{S}_{\text{QCD}}$ is then replaced by a lattice version afflicted by discretisation errors,

$$\mathcal{S}_{\text{LQCD}} = \mathcal{S}_{\text{QCD}} + a\mathcal{S}_1 + a^2\mathcal{S}_2 + \dots \ , \tag{1}$$

and continuum physics can be obtained in the limit $a \to 0$.

The central object in statistical physics is the partition function $\mathcal{Z}$ of the system, and on the lattice, an expectation value of some observable A reads:

$$\begin{aligned} \langle A \rangle &= \mathcal{Z}^{-1} \int \mathcal{D}U\mathcal{D}\psi\mathcal{D}\overline{\psi} A \exp\left\{-S_{\text{LQCD}}[U]\right\} \\ &= \mathcal{Z}^{-1} \int \mathcal{D}U A \det D[U] \exp\left\{-S_{\text{gauge}}[U]\right\} \ . \end{aligned} \tag{2}$$

Here, ψ and U denote the fermion and gluon fields, respectively. The latter are represented by so-called links on the lattice. If one identifies $S_{\text{LQCD}} = \beta H$, the exponential in the first line is the Boltzmann factor. The fermion fields ψ can be integrated out exactly due to their Grassmann nature, and the resulting determinant of the fermion matrix D is expressed in terms of (*pseudo fermions*) ϕ,

$$\det D[U] \sim \int \mathcal{D}\phi^\dagger \ \mathcal{D}\phi \exp\left\{-\phi^\dagger D^{-1}[U]\phi\right\} \ , \tag{3}$$

yielding the effective action $S_{\text{eff}}[U,\phi] = S_{\text{gauge}}[U] + \phi^\dagger D^{-1}[U]\phi$. Importance sampling methods are used to evaluate this high-dimensional integral. Using the Boltzmann-weight $p[U,\phi] = \exp\left\{-S_{\text{eff}}[U,\phi]\right\}$ as probability measure, an ensemble of N gauge configurations $\{U_m\}$ is generated. Then, $\langle A \rangle$ may be approximated by

$$\langle A \rangle \approx \frac{1}{N} \sum_m A[U_m] \ . \tag{4}$$

The standard simulation algorithm to generate QCD gauge configurations is the *Hybrid Monte-Carlo (HMC)* algorithm, where the effective action is embedded in a fictitious classical system evolved over a time τ according to the hamiltonian equations of motion. Since the numerical integration is not exact, a Metropolis step is carried out in the end, thus ensuring detailed balance. For details see Ref. 6. In order to invert the high-dimensional,

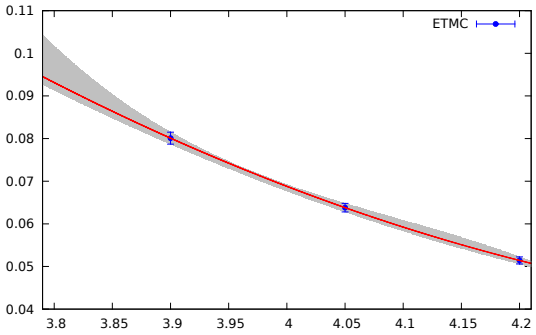

Figure 2. Lattice spacing a as a function of lattice coupling β[4].

sparse fermion matrix D, iterative Krylov space methods are used, i.e. D^{-1} is determined by equations like

$$D\phi = \psi \Rightarrow \phi = D^{-1}\psi \ . \tag{5}$$

This is the most cost-intensive part of a simulation and it is crucial to have a well tuned implementation of the derivative term $\not{D}$.

In our studies[4,5] we employ the so-called two flavour twisted mass Wilson fermions[6]. Their fermion matrix reads:

$$\begin{aligned} D^{\pm}_{\text{tm}} &= (1 \pm 2\mathrm{i}a\kappa\mu\gamma_5)\,\delta_{xy}\delta_{\alpha\beta}\delta_{ab} - \frac{\kappa}{2}\sum_{\mu}(1-\gamma_{\pm\mu})_{\alpha\beta}\,U_{\pm\mu}(x)_{ab}\delta_{n+\vec{\mu},y} \\ &= M^{\pm}_{\text{diag}} + \not{D}\ , \end{aligned} \tag{6}$$

with $\gamma_{-\mu} = -\gamma_\mu$ and $U_{-\mu}(x) = U_\mu(x-\vec{\mu})^\dagger$. γ_μ denotes Dirac matrices and a, b, α, β are colour and spinor indices, respectively. The sign in $M^{\pm}_{\text{diag}}$ corresponds to "up" and "down" quarks. In the gauge sector, the tree-level Symanzik improved Wilson action is used,

$$S_{\text{tlsym}} = \frac{\beta}{N_c}\sum_x \left(c_0 \sum_{\mu,\nu>\mu} \{1 - \text{Re Tr}(P_{\mu\nu}(x))\} + c_1 \sum_{\mu,\nu} \{1 - \text{Re Tr}(R_{\mu\nu}(x))\} \right) \ . \tag{7}$$

Here, $P_{\mu\nu}(x)$ and $R_{\mu\nu}(x)$ denote path-ordered plaquette and rectangle products of link variables. The parameters are the lattice coupling $\beta = 6/g^2$, $c_0 = 1 - 8c_1$ and $c_1 = 1/12$. For particular values $\kappa = \kappa_c(\beta)$, corresponding to "maximal twist", the $\mathcal{O}(a)$ discretisation effects vanish[6] and the quark mass is solely determined by μ.

A finite temperature T can be introduced by identifying

$$T = (a(\beta)N_\tau)^{-1} \ . \tag{8}$$

Thus, a scan in temperature equals a scan in β, and $a(\beta)$ is needed to set the scale. For this, we interpolate $T = 0$ data by the ETM collaboration[8] as in Fig. 2. ETMC also provides a formula from chiral perturbation theory to estimate $m_\pi(\mu)$ at maximal twist. The general phase structure of Twisted Mass fermions has been investigated in Ref. 5. The critical temperature and order of the chiral transition was studied in Ref. 4. Simulations

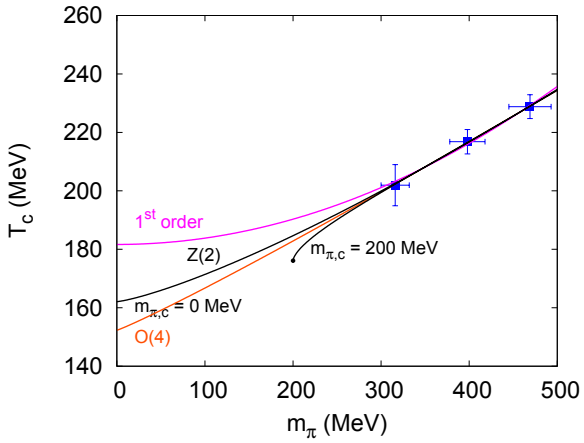

Figure 3. Extrapolation of the transition temperature to the chiral limit for various critical exponents[4].

were performed at three different pion masses ($300 \lesssim m_\pi \lesssim 500$ MeV) on lattices of size $32^3 \times 12$ and $32^3 \times 10$. The resulting critical temperatures are shown in Fig. 3. Extrapolations to the chiral limit for different orders of the transition with their associated scaling behaviour and critical exponents,

$$T_c(m_\pi) = T_c(0) + A m_\pi^{2/(\beta\delta)} \,, \tag{9}$$

are also shown. However, the combinations of exponents that go into the fit are numerically very similar, and much lower pion masses are needed to clarify the situation. This is exceedingly difficult, as it also implies larger lattices to fit the large correlation length of light pions on the lattice, while the numerical costs for the HMC scales like $V^{5/4}$ and m_π^{-6}.

3 LQCD Using OpenCL

In recent years, *Graphics Processing Units* (GPUs) have become an integral part of many modern computing clusters and are used in many LQCD applications. Tab. 1 shows an overview of available GPUs and CPUs. GPUs surpass CPUs in peak performance as well

	Chip	Peak SP [GFLOPS]	Peak DP [GFLOPS]	Peak BW [GB/s]
AMD Radeon HD 5870	Cypress	2720	544	154
AMD FirePro S10000	Tahiti	6820	1700	480
NVIDIA Tesla M2090	Fermi	1331	665	177
NVIDIA Tesla K20	Kepler	3520	1170	208
AMD Opteron 6278	Interlagos	307	154	51.2
Intel Xeon E5-2690	Sandy Bridge EP	371	186	51.2

Table 1. Theoretical peak performance of current GPUs and CPUs. SP and DP denote single and double precision, respectively. BW denotes bandwidth. Note that the AMD S10000 is a dual GPU.

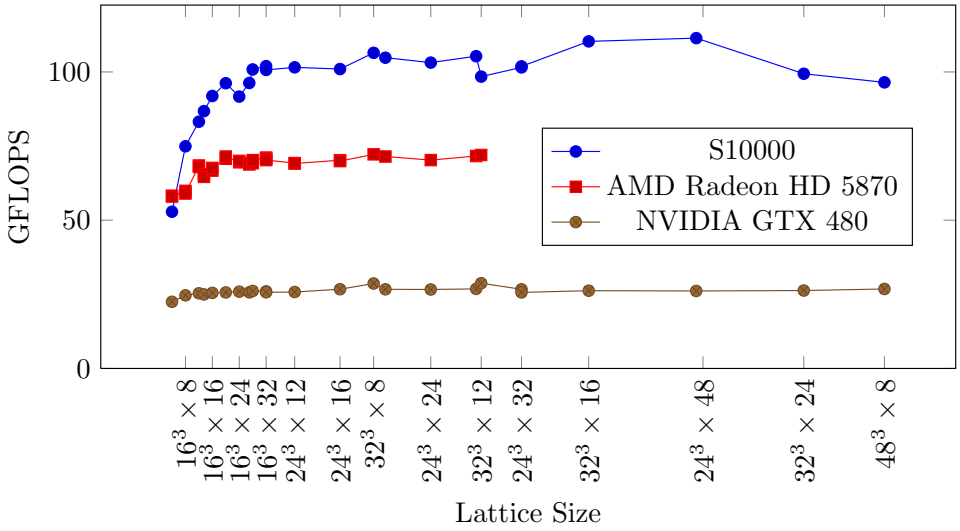

Figure 4. Performance of $\not{D}$ for various lattices sizes[3].

as in memory bandwidth, however, one also notes the drop in performance when going from single to double precision on the GPU. Current LQCD applications utilising GPUs are predominantly written using NVIDIA CUDA[9], and many routines are publicly available[10]. However, CUDA is only applicable to NVIDIA hardware. Currently we have access to two compute clusters with GPUs. One is the LOEWE-CSC (University of Frankfurt)[7], consisting of nodes with two 12-core AMD Magny-Cours CPUs and one AMD Radeon HD 5870 GPU. The other is the SANAM supercomputer (GSI Darmstadt), which has two AMD S10000 and two Intel Xeon E5-2650 CPUs per node.

For these architectures we developed a new HMC for twisted mass Wilson fermions[2,1,3]: CL^{2}QCD. It is based on OpenCL[11], an open standard for parallel computing that provides an alternative to the vendor-bound CUDA. All operations are carried out in double precision. Fig. 4 demonstrates excellent performance of the $\not{D}$ for the GPUs used in LOEWE-CSC and SANAM. 70 and 100 GFLOPS are achieved, respectively, over a wide range of lattice sizes. Also shown are results on NVIDIA GTX 680. Due to lack of optimisation the performance is poor, but it demonstrates the platform independence of OpenCL, which can also be run on CPUs. The limitation of applicability is mostly given by the GPU memory. For example, the AMD Radeon HD 5870 has only 1 GByte of memory which limits its performance when the entire lattice is put on one card.

The good performance of the $\not{D}$ carries over to the full HMC, Fig. 5. It shows three different setups corresponding to different pion masses[1] executed with CL^{2}QCD compared to a reference code[6] run on the CPUs of one node in LOEWE-CSC. The performance shows a speedup of two for the AMD Radeon HD 5870, and a speedup of four for the newer AMD S1000 with respect to the reference code. This means that one GPU is able to perform the HMC algorithm much more efficiently than two CPU nodes. In addition, the acquisition costs of a GPU is typically lower than those of a server CPU.

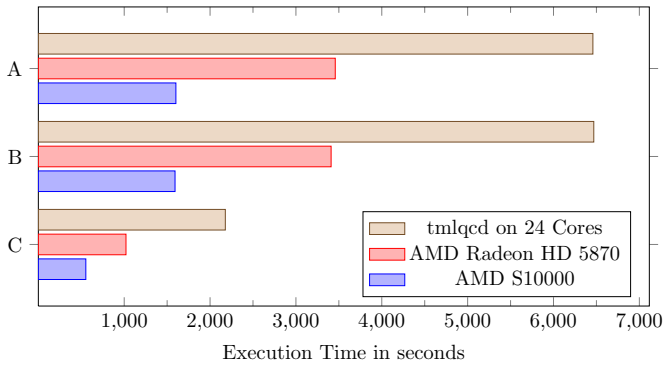

Figure 5. HMC performance compared to reference code tmlqcd[6] for different setups on a $24^3 * 8$ lattice[3].

4 First Results for $m_\pi \approx 270$ MeV

To further improve the results obtained in Ref. 4, we started simulations at a smaller pion mass of around 270 MeV on $32^3 \times 12$ lattices. The extrapolations of Fig. 3 predicts a critical β near 3.85.

We thermalised HMC chains using the tmlqcd CPU code[6] in highly parallelised fashion on LOEWE-CSC and JUGENE in Jülich[14]. On the other hand, the computer SANAM can fit the entire lattice in the memory of each GPU. We thus used this machine for production in "pedestrian parallelism": From each thermalised chain we started new chains at the same parameters, separated by sufficiently many trajectories to rule out autocorrelations. In order to reduce wall-time, we have also started to implement *Multi-GPU* usage in CL^{2}QCD, too[3] and are using it in current runs. However, this is only expected to be efficient once lattices are too large to fit into the memory of one unit.

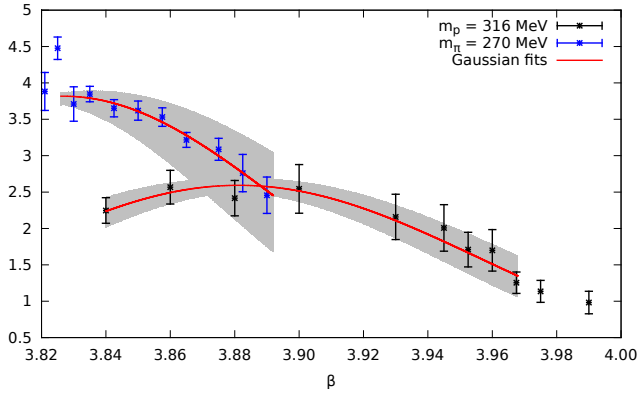

Figure 6. The susceptibility of the chiral condensate, $\sigma^2_{\bar{\psi}\psi}$. The left curve corresponds to the preliminary $m_\pi \approx 270$ MeV results, the right one to the previous $m_\pi \approx 316$ MeV.

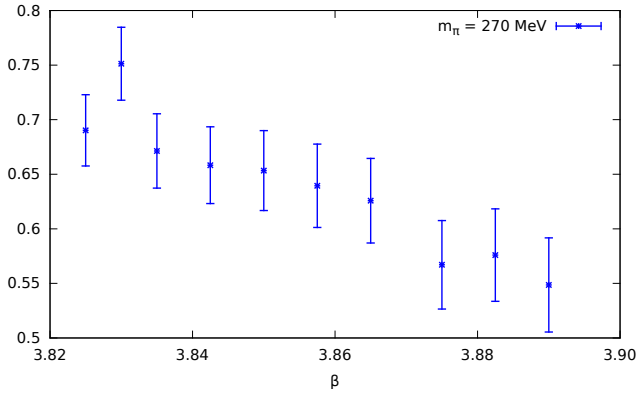

Figure 7. Preliminary results for the renormalised chiral condensate, $\langle\bar{\psi}\psi\rangle_{ren}$, for the lightest pion.

In this way we could effectively gather statistics of $\mathcal{O}(60k)$ trajectories for the 10 β values around the estimated β_c. The order parameter for the chiral transition is the chiral condensate $\langle\overline{\psi}\psi\rangle$. A peak in its susceptibility,

$$\sigma^2_{\overline{\psi}\psi} = V/T\left(\langle\overline{\psi}\psi^2\rangle - \langle\overline{\psi}\psi\rangle^2\right) \ , \tag{10}$$

signals maximal fluctuations and thus the location of the transition. This is shown in Fig. 6, corresponding to the lowest mass point in Fig. 3. Unfortunately, the new data do not show a pronounced peak, only a plateau around $\beta \approx 3.85$. On the other hand, we note that simulations for $\beta < 3.83$ become increasingly unreliable in the current setup, since the lattice spacing is only known for $\beta \geq 3.9$ and thus the pion mass cannot be held fixed reliably for smaller values. Clearly, additional simulations are necessary in order to clarify this issue.

Fig. 7 shows the renormalised chiral condensate[4],

$$\langle\overline{\psi}\psi\rangle_{ren} = \frac{\langle\overline{\psi}\psi\rangle(T,\mu) - \langle\overline{\psi}\psi\rangle(0,\mu) + \langle\overline{\psi}\psi\rangle(0,0)}{\langle\overline{\psi}\psi\rangle(0,0)} \ , \tag{11}$$

which may serve as an indicator for the transition. Again, the signal for an inflection point around $\beta \approx 3.85$ is very weak, but consistent with that of the susceptibility.

5 Conclusions

We have developed a lattice QCD code for twisted mass Wilson fermions based on OpenCL[1], CL^{2}QCD, which is able to utilise GPUs and CPUs in a vendor-independent way. It shows excellent performance on various generations of AMD GPUs. An HMC can be shown to be up to four times as fast as a reference code running on the CPUs of a whole LOEWE-CSC node. Currently, we are refining Multi-GPU usage[3].

We have used CL^{2}QCD on SANAM together with highly parallel CPU code on JUQUEEN to investigate the chiral transition in two-flavour QCD at finite temperature. Unfortunately, the current setup does not allow to investigate the full temperature range

needed for the transition at the lightest pion mass. Thus, enlarging the temporal extent N_τ is the next step. In addition, an analysis of the generated data of the Polyakovloop regarding the deconfinement transition is under way.

Acknowledgements

O. P. and C. P. are supported by the German BMBF grant *FAIR theory: the QCD phase diagram at vanishing and finite baryon density*, 06MS9150. M. B., O. P, and C. P. are supported by the Helmholtz International Center for FAIR within the LOEWE program of the State of Hesse. M. B. and C. P. are supported by the GSI Helmholtzzentrum für Schwerionenforschung. Some of the calculations have been performed on LOEWE-CSC[7], SANAM and JUGENE[14]. The authors thank the administrating teams for all the support.

References

1. M. Bach, V. Lindenstruth, O. Philipsen and C. Pinke, Comput. Phys. Commun. **184** (2013) 2042 [arXiv:1209.5942 [hep-lat]].
2. O. Philipsen, C. Pinke, C. Schafer, L. Zeidlewicz and M. Bach, PoS LATTICE **2011** (2011) 044 [arXiv:1112.5280 [hep-lat]].
3. M. Bach, V. Lindenstruth, O. Philipsen and C. Pinke, PoS LATTICE **2013** (2013). `http://www.lattice2013.uni-mainz.de/presentations/4G/Bach.pdf`
4. F. Burger, E.-M. Ilgenfritz, M. Kirchner, M. P. Lombardo, M. Muller-Preussker, O. Philipsen, C. Urbach and L. Zeidlewicz, Phys. Rev. D **87** (2013) 074508 [arXiv:1102.4530 [hep-lat]].
5. E.-M. Ilgenfritz, K. Jansen, M. P. Lombardo, M. Muller-Preussker, M. Petschlies, O. Philipsen and L. Zeidlewicz, Phys. Rev. D **80** (2009) 094502 [arXiv:0905.3112 [hep-lat]].
6. K. Jansen and C. Urbach, Comput. Phys. Commun. **180** (2009) 2717 [arXiv:0905.3331 [hep-lat]].
7. M. Bach, M. Kretz, V. Lindenstruth and D. Rohr, Computer Science - Research and Development **26** (2011) 3-4. `http://dx.doi.org/10.1007/s00450-011-0161-5`
8. R. Baron *et al.* [ETM Collaboration], JHEP **1008** (2010) 097 [arXiv:0911.5061 [hep-lat]].
9. NVIDIA, NVIDIA CUDA C Programming Guide. `http://developer.nvidia.com/nvidia-gpu-computing-documentation`
10. `http://lattice.github.io/quda/`
11. The OpenCL Specification, `http://www.khronos.org/registry/cl/`
12. `http://www.top500.org/`
13. `http://www.green500.org/`
14. `http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUGENE/JUGENE_node.html`