

Scaling SU(2) to 1000 GPUs using HiRep

Sofie Martins,^{a,b,*} Erik Kjellgren,^a Emiliano Molinaro,^a Claudio Pica^{a,b} and Antonio Rago^{a,b}

^aUniversity of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark

^bhQTC, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark

E-mail: martinss@imada.sdu.dk

HiRep allows flexible simulations of higher representations of Wilson Fermions with various actions and gauge groups and a range of inverters and integrators. This is particularly important for enabling evaluations of observables relevant to phenomenological inputs for Beyond-the-Standard-Model physics from lattice field theory. We present progress on the GPU porting of available features, especially in terms of scaling to large jobs on AMD GPUs.

The 41st International Symposium on Lattice Field Theory (LATTICE2024)
28 July - 3 August 2024
Liverpool, UK

*Speaker

1. Motivation

Modern supercomputers reach unprecedented peak bandwidths and computational throughput using Graphics Processing Units (GPUs). Exploiting these peak performances to simulate BSM physics allows the generation of high-precision, non-perturbative predictions for experimental input. While there is support for models with higher representations from numbers of colors of 2 to 5 in Grid [1, 2] using Staggered or Wilson Fermions, HiRep [3, 4] is currently the highest performing option supporting Wilson Fermions with arbitrary numbers of flavors and colors in $SU(N)$ gauge groups with fermions in the fundamental and selected higher representations for CPUs and GPUs. We aim to minimize the time needed for state-of-the-art lattice calculations with unprecedented statistics, supporting possibly full machine runs without loss of efficiency. While the GPU version was developed for NVIDIA GPUs, we reached reasonable efficiency also on AMD machines, such as LUMI-G.

The software is available at

<https://github.com/claudiopica/HiRep>

2. Execution structure

2.1 Wilson-Dirac Operator

Implementing the Wilson-Dirac operator for multi-GPU runs depends on the execution of different kernels. First, a kernel calculates the inner computations, that is, the computations that do not depend on communicated information from other GPUs. Then, there is a boundary calculation that computes all terms only after communicated information is received. The communications rely on executing a kernel that synchronizes data from the local lattice to a send buffer, followed by a memory transfer.

The efficiency of this kernel depends on the organization of these different kernels and memory operations. They can be analyzed by using profiler tracing data, either from NVIDIA Nsight (CUDA) or rocprof (ROCm). Below, we show traces for LUMI-Gs AMD MI250X GCDs and UCloud DeiC Interactive HPCs NVIDIA H100 GPUs.

In these figures, times are given relative to an arbitrarily chosen starting point chosen as a representative example of the sequence of kernel executions. There is inevitable variability from one Dirac execution to another. The example we show is typical and not the best one we found in the tracing data. Any timing reported here is up to the precision of the profiler's sampling.

Our ability to achieve the peak capabilities of the device is measured in terms of memory bandwidth instead of computational throughput, as lattice simulations are memory-bound.

In Fig. 1, we mainly want to see that the send buffer synchronizations and the Dirac boundary execution are sufficiently parallelized to minimize the overhead. Further, the communications have to be in parallel to the inner Dirac application, and they mainly are. In principle, the split up even and odd parts of the Dirac inner point computations can also be streamed in parallel; in practice, this will not work because the inner Dirac operator kernel uses the GPU to its full extent. Additionally, this parallelization is impossible in the even-odd preconditioned version of the Dirac operator used mainly in the HMC, so this improvement would be of little practical relevance.




Figure 1: ROCm profiles for single and multiple Dirac operator executions on LUMI-G in a 4 GCD/single-node job on a 48^4 local lattice parallelized in two dimensions. The vertical direction is arranged first by kernel type and then by queue ID, showing the parallelization of different kernel types in different streams. In this case, the profile on the right shows only one of the four MPI processes.

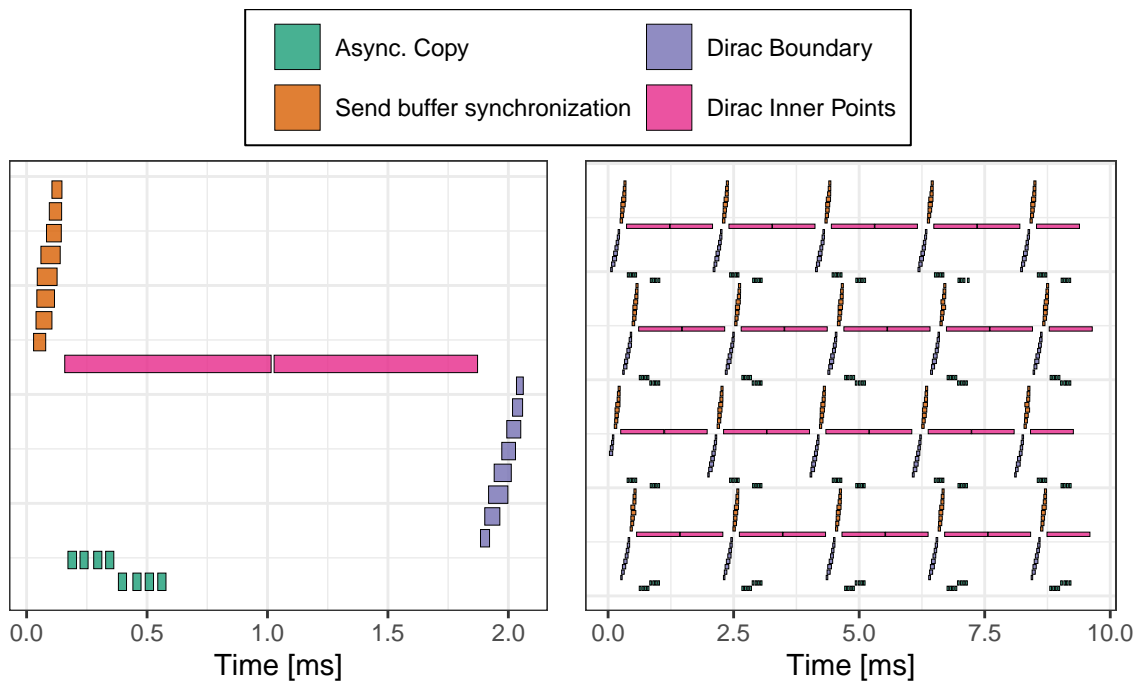


Figure 2: Profile of repeated Dirac operator applications on four NVIDIA H100 GPUs using a 48^4 local lattice parallelized in two dimensions. The tracing plot shows on the left side a single application of the Dirac operator as part of this profile and on the right side the repeated Dirac executions on all four MPI processes. The vertical direction is arranged first by process, then kernel type, and finally by stream ID. In contrast to the ROCm profiles, this shows both the MPI parallelization at the node level and the software-level streaming structure. The ROCm profiles only show a single MPI process and the queue ID at the hardware level.

For comparison, we show in Fig. 2 the execution structure of a 4-GPU NVIDIA H100 node. The structure in execution is similar, but the node achieves slightly better results in executing the communications in parallel to the heavy inner Dirac kernel, and the overall execution times are significantly reduced, mainly due to the higher theoretical peak bandwidth. Note, that the two profiles are different in displaying the parallelism. While the NVIDIA Nsight trace contains information about all processes and the stream to which the software submitted the kernel, the ROCm traces only show a single process and the hardware level queue that the kernel ends up being scheduled to, regardless of which stream the kernel was selected for.

2.1.1 Blocking and non-blocking communications

All communications in HiRep are executed from a different POSIX thread. This means we can either call blocking or non-blocking communications in parallel to the inner Dirac execution. While non-blocking communications have the advantage that they do not pile up requests, sending all requests at once can enable a well-configured network to manage the requests better. Which setup is more efficient needs to be tested on an individual basis.

2.2 Clover improvement

Clover improvement of the Wilson-Dirac operator requires the application of another diagonal term. While the parallelization of this term is trivial, the site-by-site operations require more memory and computation. Therefore, the applications were optimized by precomputing certain fields, such as the LDL-factorization and exponential of the clover term, respectively, in such a way that they need only to be recomputed whenever the gauge field is updated.

The results in Fig. 3 show that the application of the additional clover term only takes little additional time. The extra time needed is approximately the same irrespective of the specific improvement, either regular clover improvement (left) [5] or exponential clover improvement (right) [6]. Due to the mathematical structure of the two improvements, the inverse of the regular clover improvement kernel takes an approximate factor of 2 longer than the kernel shown here for this local lattice of 48^4 . In the exponential clover, the diagonal part of the Dirac operator and its inverse take the same time to execute. This is particularly useful when using Hasenbusch acceleration [7].

3. Weak and strong scaling

We must stress that a good scaling up to 1000 GPUs only works for weak scaling. Here, we observe a close-to-optimal scaling. For a single GCD execution, we are achieving around 82% (1.34 TB/s) of the theoretical peak bandwidth of the MI250X GCD, given by 1.6384 TB/s. For larger jobs, we are achieving less because of the overhead caused by separately calling the boundary kernels and the buffer synchronization. If the inner Dirac kernel execution is becoming too fast to mask the communications, additional overhead is expected causing an issue with strong scaling.

Fig. 4 shows the scaling behavior on LUMI-G, with corresponding data in Tab.1. While the weak scaling is near perfect, achieving a peak bandwidth of 679 TB/s, jobs that contain over 128 GPUs are showing decreasing efficiency when scaling strongly. This is likely because the inner Dirac computations are becoming too fast to mask the communications even on a state-of-the-art network. While this is an issue for the efficiency of jobs, it is still possible to scale to relatively

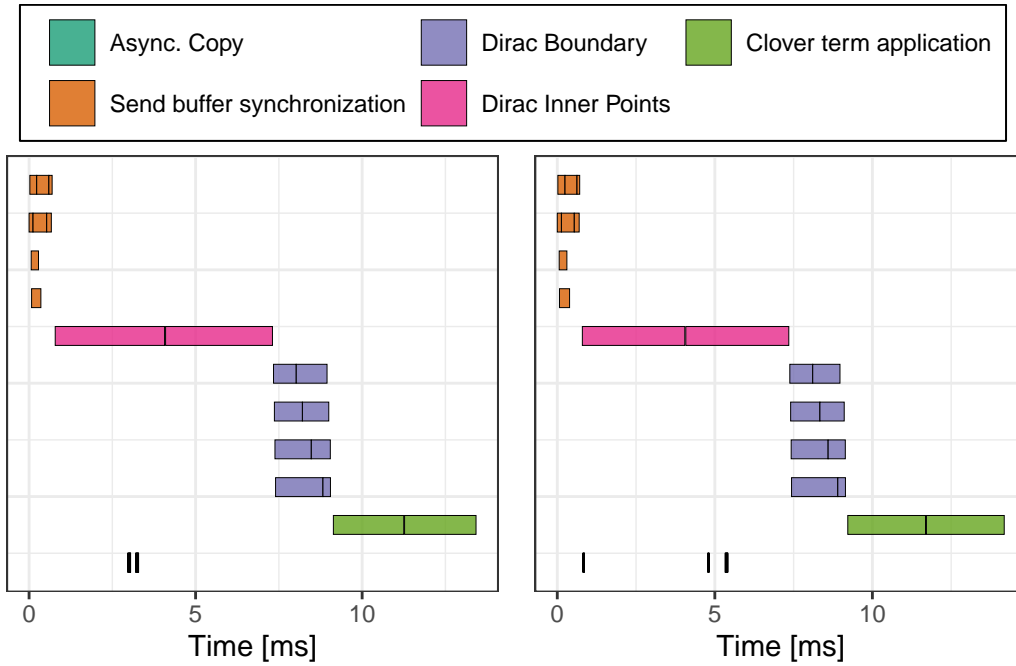


Figure 3: A single Dirac operator application including the diagonal clover term executed on AMD MI250X GCDs on LUMI-G using a 48^4 local lattice with two parallelized dimensions, both for regular clover improvement (left) and exponential clover improvement (right). The vertical direction is organized first by kernel type and second by queue ID.




Figure 4: Weak and strong scaling from 16 to 1024 GCDs on LUMI-G with all four dimensions parallelized on a local lattice of 48^4 (weak scaling) and global lattice of 128^4 (strong scaling).

Table 1: Strong and weak scaling data for the unimproved Wilson-Dirac operator in terms of bandwidth reached on LUMI-G using AMD MI250X GCDs

T/a	L_1/a	L_2/a	L_3/a	GCDs	Bandw. [TB/s]	[%]
96	96	96	96	16	12.93176	49%
192	96	96	96	32	23.00843	44%
384	96	96	96	64	42.18804	40%
768	96	96	96	128	85.74323	41%
1536	96	96	96	256	166.92260	40%
3072	96	96	96	512	331.60242	40%
6144	96	96	96	1024	679.34508	40%
128	128	128	128	32	21.38658	41%
128	128	128	128	64	36.94729	35%
128	128	128	128	128	71.44381	34%
128	128	128	128	256	96.84180	23%
128	128	128	128	512	168.57047	20%
128	128	128	128	1024	262.96682	16%

small local lattices and scale jobs to achieve unparalleled execution times. So far, for $SU(2)$ gauge groups, we observe that local lattices smaller than 16^4 should be avoided, but for larger numbers of colors, this needs to be tested individually.

4. Performance reached in comparison

In Fig. 5, we compare the peak bandwidths reached on two different architectures. While we exceed the peak theoretical bandwidth on the NVIDIA-A100 on Tursa due to L2 reuse, we have yet to reach peak performance on the MI250X on LUMI-G. The impact of the overhead is also heavier in the case of LUMI-G. However, the scaling of the code for multi-node calculations is close to perfect and reaches 40% of the peak performance of 1024 GCDs on LUMI-G.

5. Conclusion and outlook

We can reach 82% (1.34 TB/s) of the peak performance on AMD MI250X cards for a code originally developed for CUDA. The code scales up to 1024 GCDs, reaching 679 TB/s. This increases the possibilities for the BSM physics lattice community to reach high precisions for predictions from the lattice.

Acknowledgments

This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement №813942. Testing, development, and benchmarking of this software was possible using resources on LUMI-G provided




Figure 5: Comparison of bandwidths per GPU (Tursa) or GCD (LUMI-G) reached for single and multi-node jobs. The red dashed line denotes the peak theoretical bandwidth of the respective GPU/GCD.

by the Danish eInfrastructure Consortium under grant application number DeiC-SDU-N5-2024055 and NVIDIA V100, A100, and H100 nodes provided by the UCloud DeiC Interactive HPC system managed by the eScience Center at the University of Southern Denmark. This work used the DiRAC Extreme Scaling service Tursa at the University of Edinburgh, managed by the Edinburgh Parallel Computing Centre on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC service at Edinburgh was funded by BEIS, UKRI and STFC capital funding and STFC operations grants. DiRAC is part of the UKRI Digital Research Infrastructure.

References

- [1] P.A. Boyle, G. Cossu, A. Yamaguchi and A. Portelli, *Grid: A next generation data parallel C++ QCD library*, *PoS LATTICE2015* (2016) 023.
- [2] E. Bennett et al., *Symplectic lattice gauge theories in the grid framework: Approaching the conformal window*, *Phys. Rev. D* **108** (2023) 094508 [2306.11649].

- [3] L. Del Debbio, A. Patella and C. Pica, *Higher representations on the lattice: Numerical simulations. $SU(2)$ with adjoint fermions*, *Phys. Rev. D* **81** (2010) 094503 [0805.2058].
- [4] S. Martins, E. Kjellgren, E. Molinaro, C. Pica and A. Rago, *GPU-accelerated Higher Representations of Wilson Fermions with HiRep*, *PoS EuroPLEx2023* (2024) 035 [2405.19294].
- [5] B. Sheikholeslami and R. Wohlert, *Improved Continuum Limit Lattice Action for QCD with Wilson Fermions*, *Nucl. Phys. B* **259** (1985) 572.
- [6] A. Francis, P. Fritsch, M. Lüscher and A. Rago, *Master-field simulations of $O(a)$ -improved lattice QCD: Algorithms, stability and exactness*, *Comput. Phys. Commun.* **255** (2020) 107355 [1911.04533].
- [7] M. Hasenbusch, *Speeding up the hybrid Monte Carlo algorithm for dynamical fermions*, *Phys. Lett. B* **519** (2001) 177 [hep-lat/0107019].