

Evolving the COLA software library

Waseem Kamleh^{a,*}

*^aCentre for the Subatomic Structure of Matter, Department of Physics,
The University of Adelaide, SA 5005, Australia*

E-mail: waseem.kamleh@adelaide.edu.au

COLA is a software library for lattice QCD, written in a combination of modern Fortran and C/C++. Intel and NVIDIA have dominated the HPC domain in the years leading up to the exascale era, but the status quo has changed with the arrival of Frontier and other AMD-based systems in the supercomputing Top 500. Setonix is a next generation HPE Cray EX system hosted at the Pawsey Supercomputing Centre in Perth, Australia. Setonix features AMD EPYC CPUs and AMD Instinct GPUs. This report describes some of my experiences in evolving COLA to adapt to the current hardware landscape.

*The 39th International Symposium on Lattice Field Theory (Lattice2022),
8-13 August, 2022
Bonn, Germany*

*Speaker

The first supercomputer in the world to achieve a HPL benchmark greater than 1 exaflop/s was [Frontier](#), based at the [Oak Ridge National Laboratory](#). Frontier is a HPE Cray EX system featuring AMD EPYC “Trento” cores and Radeon Instinct MI250X accelerators. Topping both the [June](#) and [November Top 500](#) lists in 2022 with an R_{\max} score of 1.102 exaflop/s, Frontier displaced the Fugaku supercomputer at RIKEN, which features a bespoke chipset based on the ARM architecture. The previous crown holder was Summit, another Oak Ridge entry, but with NVIDIA accelerators and an IBM Power system. The diversity of hardware architectures on leading systems represents one of the challenges of high performance computing. Research groups typically have allocations on a variety of machines, and with the recent resurgence of AMD the range of target architectures has only widened. It is not just platform portability for scientific codes that is essential, but also performance portability.

The challenges of portability in the context of scientific computing are not new of course, but there are aspects of the contemporary architecture ecosystem that are distinct. Historically, there have always been a range of CPU chipsets deployed at HPC facilities. Performant scientific codes are typically written in C/C++ or Fortran. These languages are not tied to a specific vendor, and it is reasonable to expect that the compilers for these languages are available on any system. In this sense, platform portability for CPU codes presents a relatively low barrier for code development. For performance portability, due to tight language restrictions around aliasing, Fortran compilers developed by hardware vendors such as Intel and Cray have traditionally been very successful at generating optimised code. With regard C/C++, the use of architecture specific intrinsic has often been required to get the most benefit from specific processor features (such as vectorisation).

COLA is a custom in-house code that I began developing in Fortran as a graduate student. The code is more or less in a constant state of change, either evolving to adapt to new challenges or expanding to add new features. Key algorithms are solvers for linear systems and eigenmodes [1], and gauge field generation with Hybrid Monte Carlo [2]. The latter features a number of variants such as the RHMC algorithm [3] and a selection of filtering techniques [4–6]. Specific physics features are tailored to the CSSM lattice research programme, for which the COLA software library has formed the computational foundation for some time [7–64].

The fastest machine in Australia in the November 2022 Top 500 is Setonix, ranked at #15 with a HPL score of 27 petaflop/s and hosted at the Pawsey Supercomputing Centre in Perth. Similar to Frontier, Setonix is based on the AMD EPYC platform and also features the Radeon MI250X accelerators. Building upon the energy efficiency of these accelerators, [Setonix ranks at #4 in the Green 500 list](#) based on performance per Watt, 2 places higher than Frontier at #6. As the world moves towards a carbon-neutral future, there has been a developing focus on optimising the energy efficiency [65] of traditionally power hungry high performance computing systems.

To accompany the release of this novel architecture to the Australian supercomputing scene, the [Pawsey Centre for Extreme Scale Readiness \(PaCER\)](#) scheme was created. Several projects were chosen to partner with Pawsey to optimise codes and workflows for the next generation of supercomputers. The CSSM is partnered with the PaCER scheme via one of these projects, *Emergent phenomena revealed in subatomic matter*. User community initiatives with similar goals have accompanied the launch of other recent AMD-based systems.

As the Fortran components of COLA do not use any vendor or architecture specific features, adapting the software to the AMD EPYC platform was fairly straightforward. The enforcement of




Figure 1: Performance of the COLA fermion matrix on the AMD EPYC “Milan” CPU with the CrayEX, GCC, and AOCC Fortran compilers, measured relative to the Xeon “Cascade Lake” performance with the Intel Fortran compiler.

standard compliance does vary between the different compilers, requiring some minor changes for code which has primarily evolved on the Intel platform. There are three programming environments available on Setonix, namely Cray EX, GNU GCC, and AOCC (which is based on LLVM). Figure 1 shows the relative performance of the fermion matrix code under these three compilers on a single dual-socket “Milan” CPU node with a total of 128 cores, as compared to the Intel compiler on a dual-socket Xeon “Skylake” node with 48 cores. As would be expected given the effort that they put into their Fortran compiler optimisation, Cray performs the best of the three AMD programming environments, closely followed by the GNU compiler. At the time of writing, modern Fortran support within the AOCC programming environment should be considered as preliminary.

GPU acceleration for the COLA fermion matrix inverter was first introduced via NVIDIA CUDA C/C++ around the time that the Fermi architecture was released. This mixed-language approach persists to this day. As the CPU aspects of the code are written in modern Fortran, I utilise the interoperability provided by the intrinsic *ISO C Binding* module to interface with the GPU-accelerated routines that are implemented in CUDA C/C++. A consequence of this approach is that for many of the key algorithms in COLA there exist two independent implementations – one in Fortran for CPUs and one in CUDA for GPUs. The ability to cross-check the two implementations as a form of validation has proved very beneficial during code development. Significant amounts of utility code reuse are also realised, with the Fortran code that handles the reading of runtime parameters, the input and output of data, and the initialisation of the MPI process topology being common to both the CPU and GPU implementations.

Lattice codes have a relatively low flops/byte ratio and as such are typically memory bandwidth limited. The intrinsic geometric parallelism of the lattice maps naturally onto the massively parallel nature of the GPU platform. Whilst on a CPU a subvolume of the lattice would typically be




Figure 2: Performance of the COLA GPU-accelerated conjugate gradient inverter on various device architectures, measured relative to the Tesla K40 performance. Results are shown for various NVIDIA Tesla cards with generations ranging from Kepler to Ampere, and the AMD MI100 accelerator.

mapped to a single core, on a GPU each lattice site is mapped to a single thread. The latency hiding capabilities of the GPU execution model coupled with the high computational power means that opportunities to perform additional computation in order to reduce traffic from global memory generally result in an overall speedup of the code. Similarly, the use of mixed-precision techniques also results in a significant benefit [66]. Historically, the amount of device memory available was relatively small compared to the size of a fermion field. The available device memory on GPU accelerators has increased significantly with successive generations, such that limiting the number of vectors stored is less of a concern than it once was (similarly for register pressure).

Figure 2 shows the relative performance of the GPU-accelerated COLA conjugate gradient inverter on a variety of GPU architectures, using the NVIDIA Tesla K40 as a reference. Starting with Kepler, NVIDIA results are provided for generations up to Ampere with the A100. There is a single data point for AMD Instinct, the MI100. A mixed-precision solver is used for the benchmark, with 32-bit precision for the inner iterations and 64-bit precision for the outer solver.

To run on the MI100 architecture the CUDA components of COLA were converted to AMD’s Heterogeneous Interface for Portability (HIP). The HIP SDK provides scripts that automated much of the process of converting from CUDA code, although some manual effort was required to complete the conversion. The HIP compiler can target both NVIDIA and AMD accelerators, and on the Volta platform at least seemed to provide equivalent performance to *nvcc*. At the time of writing, the MI250X was not yet available to the author for benchmarking. Setonix will be fully launched in early 2023, when the MI250X partition becomes available to users. While platform portability has been demonstrated for the COLA software, it will be interesting to see what can be achieved in terms of performance on the MI250X.

It seems that one of the aims of HIP is to mirror the functionality provided by CUDA, though

this approach necessarily means there is a delay between a CUDA feature release and the appearance of the HIP equivalent. The dominance of CUDA for GPU-accelerated codes can be attributed to the fact that NVIDIA was the first vendor to successfully bring to market devices that targeted the HPC community. Arguably, their success would not have been possible without the large amount of effort put into developing the CUDA programming environment. Many researchers have implemented their code in CUDA due to a combination of the rapid maturity it achieved relative to other programming models (such as OpenCL), and the near-monopoly NVIDIA has had on accelerated HPC systems for many years.

AMD has had recent success at the hardware level with the launch of high profile HPC systems such as Frontier in the US, Lumi in Europe, and Setonix in Australia. They must now ensure that their programming environment rapidly achieves maturity in order to sustain momentum. It is also interesting to note that Intel is attempting to (re-)enter into the accelerator space with the Xe HPC platform.

This again raises the issue of (platform and performance) portability for accelerated computing. CPUs can generally be targeted by Fortran and C/C++ code in a platform-independent manner whilst maintaining performance (of course, platform-specific optimisations can always improve upon these). For accelerators, it is a different story. Vendor-specific programming environments are typically required to get the best performance. Developing divergent branches of the same code adds significant overhead to an activity that is already at a premium in the academic research environment, where time well-spent must necessarily translate into the *de facto* currency of the field.

In an ideal world we would treat the accelerator space in much the same way as we treat the traditional compute space. That is, through the establishment of platform-independent programming models with agreed upon standards that compiler providers implement to target their respective hardware platforms. Arguably, the natural way for this to proceed would be to have vendor-agnostic accelerated programming extensions to C/C++ and Fortran. Fortran has included intrinsic parallel computing features since the 1990s, and accelerator-based programming would seem like a natural extension. Language standards tend to evolve fairly slowly however, so in the short term we must look elsewhere. There are vendor-led candidates for cross-platform heterogeneous programming such as NVIDIA's OpenACC, AMD HIP, and Intel's DPC++ for oneAPI. Of course, these can be expected to perform well on the vendor's respective hardware, but support for (and performance on) the competitors hardware is not guaranteed. There are also a number of open candidates for heterogeneous programming such as OpenCL, SYCL, and Kokkos. The extent to which the various candidates above provide performance portability is being investigated [67–71]. Future work will explore this question in the context of continuing to evolve the COLA software library.

Acknowledgements

The author is supported by the Pawsey Supercomputing Centre through the Pawsey Centre for Extreme Scale Readiness (PaCER) program. This work is supported by the Australian Research Council through Grants No. DP190102215 and DP210103706.

References

- [1] T. Kalkreuter and H. Simma, *An Accelerated conjugate gradient algorithm to compute low lying eigenvalues: A Study for the Dirac operator in SU(2) lattice QCD*, *Comput. Phys. Commun.* **93** (1996) 33 [[hep-lat/9507023](#)].
- [2] S. Duane, A.D. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid Monte Carlo*, *Phys. Lett. B* **195** (1987) 216.
- [3] A.D. Kennedy, I. Horvath and S. Sint, *A New exact method for dynamical fermion computations with nonlocal actions*, *Nucl. Phys. B Proc. Suppl.* **73** (1999) 834 [[hep-lat/9809092](#)].
- [4] W. Kamleh and M. Peardon, *Polynomial Filtered HMC: An Algorithm for lattice QCD with dynamical quarks*, *Comput. Phys. Commun.* **183** (2012) 1993 [[1106.5625](#)].
- [5] T. Haar, W. Kamleh, J. Zanotti and Y. Nakamura, *Applying polynomial filtering to mass preconditioned Hybrid Monte Carlo*, *Comput. Phys. Commun.* **215** (2017) 113 [[1609.02652](#)].
- [6] T. Haar, W. Kamleh, J. Zanotti and Y. Nakamura, *Single flavour optimisations to Hybrid Monte Carlo*, *Comput. Phys. Commun.* **238** (2019) 111 [[1806.04350](#)].
- [7] D.-J. Kusterer, J. Hedditch, W. Kamleh, D.B. Leinweber and A.G. Williams, *Low lying eigenmodes of the Wilson-Dirac operator and correlations with topological objects*, *Nucl. Phys. B* **628** (2002) 253 [[hep-lat/0111029](#)].
- [8] W. Kamleh, D.H. Adams, D.B. Leinweber and A.G. Williams, *Accelerated overlap fermions*, *Phys. Rev. D* **66** (2002) 014501 [[hep-lat/0112041](#)].
- [9] W. Kamleh, D.B. Leinweber and A.G. Williams, *Hybrid Monte Carlo with fat link fermion actions*, *Phys. Rev. D* **70** (2004) 014502 [[hep-lat/0403019](#)].
- [10] S. Boinepalli, W. Kamleh, D.B. Leinweber, A.G. Williams and J.M. Zanotti, *Improved chiral properties of FLIC fermions*, *Phys. Lett. B* **616** (2005) 196 [[hep-lat/0405026](#)].
- [11] W. Kamleh, P.O. Bowman, D.B. Leinweber, A.G. Williams and J. Zhang, *The fat link irrelevant clover overlap quark propagator*, *Phys. Rev. D* **71** (2005) 094507 [[hep-lat/0412022](#)].
- [12] B.G. Lasscock, J.N. Hedditch, W. Kamleh, D.B. Leinweber, W. Melnitchouk, A.W. Thomas et al., *Search for the pentaquark resonance signature in lattice QCD*, *Phys. Rev. D* **72** (2005) 014502 [[hep-lat/0503008](#)].
- [13] J.N. Hedditch, W. Kamleh, B.G. Lasscock, D.B. Leinweber, A.G. Williams and J.M. Zanotti, *I-+ exotic meson at light quark masses*, *Phys. Rev. D* **72** (2005) 114507 [[hep-lat/0509106](#)].

- [14] S. Boinpolli, P.O. Bowman, U.M. Heller, W. Kamleh, J.N. Hedditch, B.G. Lasscock et al., *Some recent lattice QCD results from the CSSM*, *Int. J. Mod. Phys. A* **22** (2007) 5053.
- [15] J.N. Hedditch, W. Kamleh, B.G. Lasscock, D.B. Leinweber, A.G. Williams and J.M. Zanotti, *Pseudoscalar and vector meson form-factors from lattice QCD*, *Phys. Rev. D* **75** (2007) 094504 [[hep-lat/0703014](#)].
- [16] W. Kamleh, P.O. Bowman, D.B. Leinweber, A.G. Williams and J. Zhang, *Unquenching effects in the quark and gluon propagator*, *Phys. Rev. D* **76** (2007) 094501 [[0705.4129](#)].
- [17] B.G. Lasscock, J.N. Hedditch, W. Kamleh, D.B. Leinweber, W. Melnitchouk, A.G. Williams et al., *Even parity excitations of the nucleon in lattice QCD*, *Phys. Rev. D* **76** (2007) 054510 [[0705.0861](#)].
- [18] W. Kamleh, B. Lasscock, D.B. Leinweber and A.G. Williams, *Scaling analysis of FLIC fermion actions*, *Phys. Rev. D* **77** (2008) 014507 [[0709.1531](#)].
- [19] M.S. Mahbub, A. O. Cais, W. Kamleh, B.G. Lasscock, D.B. Leinweber and A.G. Williams, *Isolating Excited States of the Nucleon in Lattice QCD*, *Phys. Rev. D* **80** (2009) 054507 [[0905.3616](#)].
- [20] M.S. Mahbub, A.O. Cais, W. Kamleh, B.G. Lasscock, D.B. Leinweber and A.G. Williams, *Isolating the Roper Resonance in Lattice QCD*, *Phys. Lett. B* **679** (2009) 418 [[0906.5433](#)].
- [21] M.S. Mahbub, A.O. Cais, W. Kamleh, D.B. Leinweber and A.G. Williams, *Positive-parity Excited-states of the Nucleon in Quenched Lattice QCD*, *Phys. Rev. D* **82** (2010) 094504 [[1004.5455](#)].
- [22] M.S. Mahbub, W. Kamleh, D.B. Leinweber, A. O Cais and A.G. Williams, *Ordering of Spin- $\frac{1}{2}$ Excitations of the Nucleon in Lattice QCD*, *Phys. Lett. B* **693** (2010) 351 [[1007.4871](#)].
- [23] CSSM LATTICE collaboration, *Roper Resonance in 2+1 Flavor QCD*, *Phys. Lett. B* **707** (2012) 389 [[1011.5724](#)].
- [24] D.S. Roberts, P.O. Bowman, W. Kamleh and D.B. Leinweber, *Wave Functions of the Proton Ground State in the Presence of a Uniform Background Magnetic Field in Lattice QCD*, *Phys. Rev. D* **83** (2011) 094504 [[1011.1975](#)].
- [25] B.J. Menadue, W. Kamleh, D.B. Leinweber and M.S. Mahbub, *Isolating the $\Lambda(1405)$ in Lattice QCD*, *Phys. Rev. Lett.* **108** (2012) 112001 [[1109.6716](#)].
- [26] E.-A. O'Malley, W. Kamleh, D. Leinweber and P. Moran, *SU(3) centre vortices underpin confinement and dynamical chiral symmetry breaking*, *Phys. Rev. D* **86** (2012) 054503 [[1112.2490](#)].
- [27] D.S. Roberts, W. Kamleh, D.B. Leinweber, M.S. Mahbub and B.J. Menadue, *Accessing High Momentum States In Lattice QCD*, *Phys. Rev. D* **86** (2012) 074504 [[1206.5891](#)].

- [28] M.S. Mahbub, W. Kamleh, D.B. Leinweber, P.J. Moran and A.G. Williams, *Low-lying Odd-parity States of the Nucleon in Lattice QCD*, *Phys. Rev. D* **87** (2013) 011501 [1209.0240].
- [29] B.J. Owen, J. Dragos, W. Kamleh, D.B. Leinweber, M.S. Mahbub, B.J. Menadue et al., *Variational Approach to the Calculation of g_A* , *Phys. Lett. B* **723** (2013) 217 [1212.4668].
- [30] M.S. Mahbub, W. Kamleh, D.B. Leinweber, P.J. Moran and A.G. Williams, *Structure and Flow of the Nucleon Eigenstates in Lattice QCD*, *Phys. Rev. D* **87** (2013) 094506 [1302.2987].
- [31] F.M. Stokes, W. Kamleh, D.B. Leinweber, M.S. Mahbub, B.J. Menadue and B.J. Owen, *Parity-expanded variational analysis for nonzero momentum*, *Phys. Rev. D* **92** (2015) 114506 [1302.4152].
- [32] D.S. Roberts, W. Kamleh and D.B. Leinweber, *Wave Function of the Roper from Lattice QCD*, *Phys. Lett. B* **725** (2013) 164 [1304.0325].
- [33] A. Trewartha, W. Kamleh, D. Leinweber and D.S. Roberts, *Quark Propagation in the Instantons of Lattice QCD*, *Phys. Rev. D* **88** (2013) 034501 [1306.3283].
- [34] T. Primer, W. Kamleh, D. Leinweber and M. Burkardt, *Magnetic properties of the nucleon in a uniform background field*, *Phys. Rev. D* **89** (2014) 034508 [1307.1509].
- [35] M.S. Mahbub, W. Kamleh, D.B. Leinweber and A.G. Williams, *Searching for low-lying multi-particle thresholds in lattice spectroscopy*, *Annals Phys.* **342** (2014) 270 [1310.6803].
- [36] D.S. Roberts, W. Kamleh and D.B. Leinweber, *Nucleon Excited State Wave Functions from Lattice QCD*, *Phys. Rev. D* **89** (2014) 074501 [1311.6626].
- [37] F.M. Stokes, W. Kamleh and D.B. Leinweber, *Visualizations of coherent center domains in local Polyakov loops*, *Annals Phys.* **348** (2014) 341 [1312.0991].
- [38] S.D. Thomas, W. Kamleh and D.B. Leinweber, *Instanton contributions to the low-lying hadron mass spectrum*, *Phys. Rev. D* **92** (2015) 094515 [1410.7105].
- [39] J.M.M. Hall, W. Kamleh, D.B. Leinweber, B.J. Menadue, B.J. Owen, A.W. Thomas et al., *Lattice QCD Evidence that the $\Lambda(1405)$ Resonance is an Antikaon-Nucleon Molecule*, *Phys. Rev. Lett.* **114** (2015) 132002 [1411.3402].
- [40] B. Owen, W. Kamleh, D. Leinweber, B. Menadue and S. Mahbub, *Light Meson Form Factors at near Physical Masses*, *Phys. Rev. D* **91** (2015) 074503 [1501.02561].
- [41] A.L. Kiratidis, W. Kamleh, D.B. Leinweber and B.J. Owen, *Lattice baryon spectroscopy with multi-particle interpolators*, *Phys. Rev. D* **91** (2015) 094509 [1501.07667].
- [42] A. Trewartha, W. Kamleh and D. Leinweber, *Evidence that centre vortices underpin dynamical chiral symmetry breaking in $SU(3)$ gauge theory*, *Phys. Lett. B* **747** (2015) 373 [1502.06753].

- [43] B.J. Owen, W. Kamleh, D.B. Leinweber, M.S. Mahbub and B.J. Menadue, *Transition of $\rho \rightarrow \pi\gamma$ in lattice QCD*, *Phys. Rev. D* **92** (2015) 034513 [1505.02876].
- [44] A. Trewartha, W. Kamleh and D. Leinweber, *Connection between center vortices and instantons through gauge-field smoothing*, *Phys. Rev. D* **92** (2015) 074507 [1509.05518].
- [45] Z.-W. Liu, W. Kamleh, D.B. Leinweber, F.M. Stokes, A.W. Thomas and J.-J. Wu, *Hamiltonian effective field theory study of the $N^*(1535)$ resonance in lattice QCD*, *Phys. Rev. Lett.* **116** (2016) 082004 [1512.00140].
- [46] J. Dragos, R. Horsley, W. Kamleh, D.B. Leinweber, Y. Nakamura, P.E.L. Rakow et al., *Nucleon matrix elements using the variational method in lattice QCD*, *Phys. Rev. D* **94** (2016) 074505 [1606.03195].
- [47] Z.-W. Liu, W. Kamleh, D.B. Leinweber, F.M. Stokes, A.W. Thomas and J.-J. Wu, *Hamiltonian effective field theory study of the $N^*(1440)$ resonance in lattice QCD*, *Phys. Rev. D* **95** (2017) 034034 [1607.04536].
- [48] A.L. Kiratidis, W. Kamleh, D.B. Leinweber, Z.-W. Liu, F.M. Stokes and A.W. Thomas, *Search for low-lying lattice QCD eigenstates in the Roper regime*, *Phys. Rev. D* **95** (2017) 074507 [1608.03051].
- [49] J.M.M. Hall, W. Kamleh, D.B. Leinweber, B.J. Menadue, B.J. Owen and A.W. Thomas, *Light-quark contributions to the magnetic form factor of the $\Lambda(1405)$* , *Phys. Rev. D* **95** (2017) 054510 [1612.07477].
- [50] A. Trewartha, W. Kamleh and D. Leinweber, *Centre vortex removal restores chiral symmetry*, *J. Phys. G* **44** (2017) 125002 [1708.06789].
- [51] R. Bignell, J. Hall, W. Kamleh, D. Leinweber and M. Burkardt, *Neutron magnetic polarizability with Landau mode operators*, *Phys. Rev. D* **98** (2018) 034504 [1804.06574].
- [52] J.C. Biddle, W. Kamleh and D.B. Leinweber, *Gluon propagator on a center-vortex background*, *Phys. Rev. D* **98** (2018) 094504 [1806.04305].
- [53] J.J. Wu, W. Kamleh, D.t. Leinweber, R.D. Young and J.M. Zanotti, *Accessing high-momentum nucleons with dilute stochastic sources*, *J. Phys. G* **45** (2018) 125102 [1807.09429].
- [54] F.M. Stokes, W. Kamleh and D.B. Leinweber, *Opposite-Parity Contaminations in Lattice Nucleon Form Factors*, *Phys. Rev. D* **99** (2019) 074506 [1809.11002].
- [55] F.M. Stokes, W. Kamleh and D.B. Leinweber, *Elastic Form Factors of Nucleon Excitations in Lattice QCD*, *Phys. Rev. D* **102** (2020) 014507 [1907.00177].
- [56] R. Bignell, W. Kamleh and D. Leinweber, *Pion in a uniform background magnetic field with clover fermions*, *Phys. Rev. D* **100** (2019) 114518 [1910.14244].

- [57] A. Virgili, W. Kamleh and D. Leinweber, *Role of chiral symmetry in the nucleon excitation spectrum*, *Phys. Rev. D* **101** (2020) 074504 [1910.13782].
- [58] J.C. Biddle, W. Kamleh and D.B. Leinweber, *Visualization of center vortex structure*, *Phys. Rev. D* **102** (2020) 034504 [1912.09531].
- [59] R. Bignell, W. Kamleh and D. Leinweber, *Magnetic polarizability of the nucleon using a Laplacian mode projection*, *Phys. Rev. D* **101** (2020) 094502 [2002.07915].
- [60] R. Bignell, W. Kamleh and D. Leinweber, *Pion magnetic polarisability using the background field method*, *Phys. Lett. B* **811** (2020) 135853 [2005.10453].
- [61] CSSM/QCDSF/UKQCD collaboration, *State mixing and masses of the π_0 , η and η' mesons from $nf=1+1+1$ lattice QCD+QED*, *Phys. Rev. D* **104** (2021) 114514 [2110.11533].
- [62] A. Virgili, W. Kamleh and D. Leinweber, *Smoothing algorithms for projected center-vortex gauge fields*, *Phys. Rev. D* **106** (2022) 014505 [2203.09764].
- [63] J.C. Biddle, W. Kamleh and D.B. Leinweber, *Static quark potential from center vortices in the presence of dynamical fermions*, *Phys. Rev. D* **106** (2022) 054505 [2206.00844].
- [64] J.C. Biddle, W. Kamleh and D.B. Leinweber, *Impact of dynamical fermions on the center vortex gluon propagator*, *Phys. Rev. D* **106** (2022) 014506 [2206.02320].
- [65] C.D. Pietrantonio, C. Harris and M. Cytowski, *Energy-based accounting model for heterogeneous supercomputers*, *CoRR* **abs/2110.09987** (2021) [2110.09987].
- [66] M.A. Clark, R. Babich, K. Barros, R.C. Brower and C. Rebbi, *Solving Lattice QCD systems of equations using mixed precision solvers on GPUs*, *Comput. Phys. Commun.* **181** (2010) 1517 [0911.3191].
- [67] M. Bach, V. Lindenstruth, O. Philipsen and C. Pinke, *Lattice QCD based on OpenCL*, *Comput. Phys. Commun.* **184** (2013) 2042 [1209.5942].
- [68] P.A. Boyle, M.A. Clark, C. DeTar, M. Lin, V. Rana and A.V. Avilés-Casco, *Performance Portability Strategies for Grid C++ Expression Templates*, *EPJ Web Conf.* **175** (2018) 09006 [1710.09409].
- [69] C. Bonati, E. Calore, M. D'Elia, M. Mesiti, F. Negro, F. Sanfilippo et al., *Portable multi-node LQCD Monte Carlo simulations using OpenACC*, *Int. J. Mod. Phys. C* **29** (2018) 1850010 [1801.01473].
- [70] S. Cali, W. Detmold, G. Korcyl, P. Korcyl and P. Shanahan, *Implementation of the conjugate gradient algorithm for heterogeneous systems*, *PoS LATTICE2021* (2022) 507 [2111.14958].
- [71] A. Yamaguchi, P. Boyle, G. Cossu, G. Filaci, C. Lehner and A. Portelli, *Grid: OneCode and FourAPIs*, *PoS LATTICE2021* (2022) 035 [2203.06777].