

Integrated photonic multigrid solver for partial differential equations

Wolfram Pernice

wolfram.pernice@kip.uni-heidelberg.de

Heidelberg University <https://orcid.org/0000-0003-4569-4213>

Timoteo Lee

Heidelberg University <https://orcid.org/0009-0007-3236-8407>

Frank Brücknerhoff-Plückelmann

University of Heidelberg <https://orcid.org/0009-0009-7713-6735>

Jelle Dijkstra

Ruprecht-Karls-Univ Heidelberg

Jan Pawlowski

Heidelberg University

Article

Keywords:

Posted Date: September 2nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10339976/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Integrated photonic multigrid solver for partial differential equations

Timoteo Lee^{1,2*}, Frank Brückerohoff-Plückelmann^{1*}, Jelle Dijkstra¹, Jan M. Pawłowski², and Wolfram Pernice^{1†}

¹Kirchhoff-Institute for Physics, University of Heidelberg; Heidelberg, 69120, Germany.

²Institute for Theoretical Physics, University of Heidelberg; Heidelberg, 69120, Germany.

*These authors contributed equally.

Email: [†]wolfram.pernice@kip.uni-heidelberg.de

Solving partial differential equations underpins quantitative modelling across physics, engineering and the natural sciences, but high-accuracy simulations increasingly push conventional high-performance computers to their limits. Many of these problems reduce to large, sparse linear systems whose solution is limited not only by arithmetic cost, but also by the sequential and communication-bound nature of established digital algorithms. Current trends in general-purpose digital hardware are not well aligned with these properties, resulting in the utilization of only a small fraction of the peak performance. Specialised photonic processors offer ultralow-latency matrix-vector operations, yet analog noise has so far prevented optical solvers from reaching the accuracies required for demanding scientific simulations. Here, we show the synergy between multigrid methods and low-latency photonic matrix-vector multipliers. We develop a mixed-precision photonic multigrid framework that offloads the computationally demanding smoothing algorithm to the optical domain. Within this framework, we construct integrated photonic multigrid solvers based on a photonic crossbar array. The solvers converge with residual norms below 10^{-10} for Poisson and Schrödinger problems and reduces high-precision digital operations by 60% and 80%, respectively. Extending the approach with an adaptive multigrid scheme for lattice quantum chromodynamics, we find reductions of up to 97% for ill-conditioned systems with small quark masses, potentially improving computational speed and efficiency by more than an order of magnitude. Using the strengths of high-precision digital and low-latency photonic computing, the integrated photonic multigrid solvers present an exciting trajectory to increase computational power in the post-Moore's law era.

Introduction

Solving partial differential equations (PDEs) quickly and efficiently is fundamental to all natural sciences and engineering, ranging from fluid mechanics and atmospheric modelling to lattice quantum chromodynamics¹⁻⁴. Numerically approximating the solution on a discretized grid, as illustrated in Fig. 1a, allows to study systems without an analytical solution. However, high-accuracy PDE solutions require increasingly fine discretizations, resulting in linear systems with billions of unknowns. Solving them can account for up to 99% of the computational cost of large-scale simulations². In recent decades, continued development in semiconductor technology has enabled more powerful PDE solvers. As Moore’s law begins to slow down and with the end of Dennard’s scaling, innovations in hardware and algorithms are essential for future technological growth^{5,6}. On the algorithm side, mixed-precision methods are increasingly important tools for extracting performance from digital high-performance computing systems⁷. Here, software and algorithms are co-optimized to partially utilize low-precision operations reducing computational cost while maintaining high-accuracy solutions. However, this optimization only leverages a fraction of the maximum performance of high-performance computing systems due to the mismatch between the parallel nature of modern general-purpose digital hardware and the sequential nature of PDE solvers^{1,8}. In this context, new computing paradigms and hardware-algorithm co-design are crucial to sustain the needed performance growth^{5,6}.

Specialized analog processors provide an interesting alternative to modern general-purpose parallel digital hardware. By implementing multiplications and accumulations in the physical domain, their latency is greatly reduced^{9,10}, which is optimal for the sequential nature of PDE solvers⁸. Especially intriguing are optical matrix accelerators, as photons, unlike electrons, can transmit information without being subject to ohmic losses and capacitive charging¹¹⁻¹³. Optical accelerators excel at broadcast and weight architectures, which perform matrix multiplications in a single clock cycle, enabling speedups of more than two orders of magnitude in comparison to conventional graphics processing units⁹. Due to the optimal properties of ultra-low latency in-memory photonic processors for solving PDEs, multiple photonic PDE-solving approaches have been proposed in the literature, including data-driven machine learning approaches¹⁴, offloading matrix-vector multiplications from standard iterative solvers to the optical domain¹⁵, and directly incorporating the algorithms on the hardware¹⁶⁻¹⁸. However, none of these approaches currently achieves the accuracy required for many state-of-the-art applications, e.g., lattice quantum chromodynamics calculations. Co-optimizing hardware and algorithms through mixed-precision methods is an intriguing approach to both benefiting from ultra-low-latency computations and mitigating the inherently noisy nature of analog processors¹⁹⁻²². In this hybrid configuration, the analog accelerator performs most operations at low precision and the digital systems sparsely perform high-precision correction steps to ensure convergence with high accuracy, as illustrated in Fig. 1b. Although first hybrid approaches based

on the iterative refinement method demonstrate the potential of this approach for numerically easy problems with low condition number²⁰⁻²², it remains elusive to deploy efficient hybrid solvers in a wide range of large-scale problems.

Here, we develop a mixed-precision photonic multigrid (MPPMG) framework that offloads computationally expensive smoothing operations to the photonic domain. In this way, we exploit the low latency of photonic data processing while obtaining high-accuracy solutions for a wide range of large-scale calculations. We construct integrated photonic multigrid solvers by implementing a Richardson’s smoother on a photonic crossbar array for in-memory computation, benchmark the convergence for the Schrödinger and Poisson equations, and demonstrate the advantage of the MPPMG framework for numerically hard lattice quantum chromodynamics calculations with high condition numbers.

Mixed-Precision Photonic Multigrid Framework

To develop high-accuracy, low-latency hybrid photonic solvers, we exploit the mathematical properties of multigrid methods. These are considered the most efficient approach for a wide range of PDE problems³. Most multigrid methods consist of a smoother that reduces high-frequency error components, and transfer operators that connect different grid levels. By solving the transferred problem on coarser grids, the method efficiently removes the error across all frequencies. This idea can be made precise by looking at the decomposition of the error of the solution estimate into oscillatory components given by the eigenvectors of the matrix A . That is,

$$e_k := x - x_k = \sum_{n=0}^{V-1} c_n v_n, \quad (1)$$

where x_k is the current solution estimate, x is the real solution, v_n are the eigenvectors of A , and V is the dimension of the matrix. By first damping the high-frequency error components, which cannot be accurately represented on coarser grids, multigrid ensures that the remaining error can be transferred to the coarse levels without significant loss of information. The remaining low-frequency error components, which are numerically harder to eliminate, are then reduced by solving a considerably cheaper coarse-grid problem. The resulting error correction is subsequently transferred to the finer grids to improve the current solution estimate.

In the MPPMG framework, we offload the computationally demanding smoothing operations to the low-latency analog photonic domain and optimize the hybrid solvers for each problem by selecting which multigrid components are retained in the high-precision digital domain, as illustrated in Fig. 1c. In this configuration, the multigrid method remains robust to analog noise when used as a preconditioner for standard solvers, provided that the photonic smoother sufficiently attenuates the relevant high-frequency error components. In other words, we transform noise-sensitive hard problems into developing photonic smoothers, which is considerably easier. We provide a more

detailed analysis of the framework and optimization in the Supplementary Information. Finally, we highlight that any analog PDE solver with smoothing properties can benefit from the MPPMG approach, opening the possibility to explore innovative hardware-algorithm solutions while achieving high-accuracy solutions.

Low-Latency Richardson Smoother

We implement a Richardson smoother based on the integrated photonic processor array shown in Fig. 2a. It consists of a 9x3 crossbar array storing the matrix weights for in-memory computing, high-speed electro absorption modulators (EAMs) to convert the input vectors to the optical domain, and photodetectors (PDs) to readout the results of the matrix vector multiplications. The photonic processor performs the computations by manipulating the intensity of the optical light pulses rather than relying on logical gates. Two pulses are accumulated via superposition, and the matrix weights are stored in the optical transmission of the computational memory. Thus, a full matrix-vector multiplication reduces to an optical transmission measurement, and the lower bound of the latency is given by the pulse propagation time. As the optical bandwidth of the integrated photonic circuit spans several THz, the pulse shape does not deteriorate during propagation. To harness the ultra-low latency of photonic data processing, the system is interfaced by high-speed digital-to-analog converters (DACs) and analog-to-analog converters (ADCs) operating at 4 GS/s and 2 GS/s, respectively²³. Since the information is directly stored in the optical intensity, noise sources, like thermal noise in the readout chain, directly impact the precision of the photonic computation. Fig. 2b shows the measured result vs the ideal result for 100 random matrix vector multiplications. The accuracy is similar to fixed-point operations with 4 bits.

The Richardson smoother naturally maps onto the photonic processor. A single Richardson's iteration is defined as:

$$x_{k+1} = x_k + \omega(b - Ax_k), \quad (2)$$

where x_k is the current solution estimate and ω is a tunable parameter which we set to the inverse of the approximate spectral radius of A , also see the Supplementary Information. The computational cost and the latency of one iteration are dominated by the matrix-vector multiplication (MVM) of the matrix A describing the discretized PDE and the current estimate x_k . Since A is constant, we store it within in the computational memory, whereas the current estimate x_k is encoded via the high-speed DACs and EAMs. Similarly, we obtain the result of the MVM from the high-speed PDs and ADCs. In Fig. 2c, we test the photonic smoother on a manually chosen initial estimate x_0 of the solution with an error containing all frequency components. While the analog noise impacts the smoother, we obtain a large attenuation of the high-frequency components, necessary for reaching high-accuracy solutions.

High-Accuracy Photonic Multigrid Solvers

Algorithms for solving Poisson-like equations, also known as Poisson solvers, find applications in astrophysics, chemistry, mechanics, electromagnetics, statistics, and image processing^{24,25}. Here, we construct a photonic Poisson solver within the MPPMG framework and calculate the electric field of a parallel-plate capacitor (PC), described by the Poisson equation

$$\nabla^2 u(y, z) = f(y, z) = \frac{\rho}{\epsilon}, \quad (3)$$

which we discretize by using the finite-difference method (FDM). Then, the linear equations have the shape

$$-4u_{i,j} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1} = h^2 f_{i,j}, \quad (4)$$

where we choose $h = 1$ for simplicity. To solve the parallel capacitor problem, we need to use the correct boundary conditions²⁵. This leads to $f_{i,j} = 0$, except at the two plates where we have $f_{i,j} = 2$ and $f_{i,j} = -2$ for the respective plates, and to a matrix with a condition number of $\kappa = 467$. The condition number, also see Eq. (11), determines the numerical difficulty of the problem. Typical hybrid solvers based on the iterative refinement method reported in literature²⁰⁻²² solved linear systems with condition numbers $\kappa \in [1, 100]$. Finally, we obtain the electrical field by calculating the gradient of the potential u .

We solve this problem using an MPPMG solver that consists of a conjugate gradient (CG) solver with a smoothed aggregation algebraic multigrid preconditioner with two levels combined with an outer iterative refinement (IR) loop. The photonic multigrid solver reaches the accuracy of purely digital systems with a residual $|r| = 10^{-10}$, which is orders of magnitude better than the 85-98% accuracy achieved by purely photonic solvers¹⁴⁻¹⁸, also see Fig. 3a&b. At the same time, the photonic solver reduces the digital double-precision operations by 60% compared to a purely digital solver, shown in Fig. 3c. We use a digital twin of the photonic processor, also see Supplementary Information, to investigate the dependency of the performance gain on the accuracy. The emulated processor with 8-bits DACs and ADCs, and Gaussian noise, explained in the Methods Section, behaves similarly as the photonic system, whereas considerably reducing the noise can potentially decrease the number of digital operations by 80%, since the solver converges with less iterations.

In addition to linear systems, the MPPMG framework also supports eigenvalue solvers. Eigenvalue problems are central to many scientific and engineering fields, from quantum chemistry and materials science to imaging, data mining, and structural analysis. The locally optimal block preconditioned conjugate gradient (LOBPCG) method, which requires a good preconditioner to work efficiently, is widely used for these problems²⁶⁻³⁴. Hence, the MPPMG eigensolver consists of an LOBPCG solver with a smoothed aggregation algebraic multigrid preconditioner with three levels. We test this on the quantum quartic anharmonic oscillator (QQAQO), which does not have analytical solutions,

and which is described by the eigenvalue problem

$$\frac{1}{2} \left(-\frac{\partial^2}{\partial y^2} + y^2 + 2\mu y^4 \right) \varphi(y) = E\varphi(y). \quad (5)$$

We use again FDM to discretize y in the domain $[-3,3]$ with 120 points separated by a distance of $h = 0.05$, and we get the matrix

$$H_{ij} = \frac{1}{2} \left(\frac{2\delta_{i,j} - \delta_{i,j+1} - \delta_{i,j-1}}{h^2} + (y_j^2 + 2\mu y_j^4) \delta_{i,j} \right), \quad (6)$$

with a condition number of $\kappa = 971$ for $\mu = 0.25$. The hybrid photonic eigensolver uses 80% less double-precision operations than the digital eigensolver for the QQAO to converge to the same accuracy defined by $|r| = 10^{-10}$, see Fig. 3d-3f. In this way, photonic multigrid solvers deploying low-latency photonic matrix-vector multipliers can considerably accelerate the solution of a wide range of large-scale problems without sacrificing accuracy.

Accelerating Lattice Quantum Chromodynamics Calculations

Lattice quantum chromodynamics (LQCD) describes the interactions between quarks and gluons, as illustrated in Fig. 4a. LQCD calculations rank among the most demanding applications in computational science, routinely utilizing a substantial portion of global supercomputing resources¹. Future scientific objectives will require computing platforms that deliver at least an order-of-magnitude increase in performance beyond current exascale systems, highlighting the importance of jointly optimizing hardware architectures and numerical algorithms³⁵. Here, we extend the MPPMG framework with an adaptive scheme that enables the solver to effectively address the phenomenon of critical slowing down, where the condition number κ diverges as the physical quark mass approaches physically realistic values^{1,36}. The adaptive part can be seen as the training phase of the algorithm, where it learns to construct a multigrid solver able to deal with the critical slowing down.

Using the low-latency photonic Richardson smoother within the adaptive MPPMG framework creates a hybrid solver that efficiently deals with LQCD calculations with small physical quark masses, where we encounter linear systems with high condition numbers. To investigate the potential of these solvers, we deploy the emulated adaptive MPPMG solver on calculations with the Wilson-Dirac operator of the 2D U(1) Schwinger model, which is a widely used benchmark for developing and testing new LQCD solvers³⁶. We solve 100 linear systems using a similar MPPMG solver as the Poisson problem, with the addition of the adaptive part, which consists of the photonic smoother and the IR method, also see Supplementary Information. The Wilson-Dirac matrix has the form

$$D_{x,y}[U] = -\frac{1}{2} \sum_{\mu=1}^2 [(1 - \gamma_\mu) U_x^\mu \delta_{x+\hat{\mu},y} + (1 + \gamma_\mu) U_{x-\hat{\mu}}^{\mu\dagger} \delta_{x-\hat{\mu},y}] + (2 + m) \delta_{x,y}, \quad (7)$$

where γ_μ are the Euclidean gamma matrices, U are the gauge link configurations, and m is the bare quark mass. We solve the normal equation with $DD^\dagger$ as the matrix A , where both matrix-vector multiplications are processed individually with the digital noisy operations described in the Methods Section. In Fig. 4b, we compare the MPPMG approach with hybrid approaches based on the iterative refinement method. The MPPMG solver is not affected by the critical slowing down, while the performance of the hybrid IR-solver collapses for small physical quark masses. In other words, the MPPMG approach completely outperforms current hybrid approaches based on the iterative refinement method on problems with high condition numbers. Note that the physical quark mass becomes zero at a negative value of the bare quark mass parameter because the two quantities are not directly equivalent in the Wilson discretization. In Fig. 4c, we solve linear systems with $m = -0.063$ for 10 different gauge links configurations with an average high condition number of $\kappa = 4 \times 10^5$, each with 10 different random RHS with an accuracy defined by $|r| = 10^{-10}$. More details are included in the Methods Section and in the Supplementary Information. To estimate the performance gains of the proposed solver for a single RHS, we compare the performance of the standard solver with that of the emulated hybrid MPPMG solvers, whose computational cost is the combined cost of the adaptive part and solving a single RHS. The hybrid solvers reduce the double-precision digital operations by 87-97%, which is equivalent to a maximum performance gain both in terms of speed and energy consumption of 8-30, according to Eq. (10). Finally, we highlight that these adaptive multigrid approaches are successful in several settings, but are challenging to deploy efficiently due to the considerable cost of the adaptive procedure³⁷, as demonstrated in the Supplementary Information. By offloading most of the adaptive part to the ultra-low latency optical domain, the applicability of this robust “black-box” solver could be extended to different numerically hard large-scale calculations with high condition number κ . Hence, enhancing current high-performance computing systems with application-specific photonic processors could open new exciting paths for algorithmic development not possible with current digital hardware.

Discussion

The mixed-precision photonic multigrid framework developed here provides a flexible route to harness low-latency photonic computation for high-accuracy scientific computing. We demonstrate the approach for eigenvalue calculations through the Schrödinger equation and linear systems through the Poisson equation and lattice quantum chromodynamics (LQCD) calculations, where we deploy an adaptive scheme for the latter. The key idea is to offload the computationally demanding smoothing operations to the photonic domain, while operations in the high-precision digital domain ensure high-accuracy solutions. This separation allows the solver to reach residual norms below 10^{-10} , which is several orders of magnitude beyond the accuracy reported for fully photonic solvers^{14–18}. The finite precision and accuracy during the photonic smoothing process therefore mainly affects the convergence rate and the digital workload, rather than the final attainable accuracy. With the integrated photonic matrix processor used here, whose accuracy is comparable to 4-bit fixed-point arithmetic, the number of high-precision digital operations is reduced by 60-80% for the

Poisson and Schrödinger equations and by up to 97% for the LQCD calculations. This is equivalent to one order-of-magnitude gain in performance, with additional gains expected for larger systems¹⁹ and from the multi-precision extensions discussed in the Supplementary Information.

The compatibility with LQCD calculations is particularly significant because it places photonic acceleration in a regime that is both scientifically important and numerically demanding. The systems studied here reach average condition numbers up to order 10^5 , far beyond the comparatively benign problems with condition numbers of order 1 to 100 addressed in recent hybrid photonic solver demonstrations^{20–22}. LQCD has long driven the development of specialized computing architectures, reflecting the exceptional cost of these simulations and their importance for quantitative studies of the strong interaction^{19,38–40}. At the same time, further acceleration on modern general-purpose parallel hardware is increasingly limited by communication costs, low arithmetic intensity and the sequential nature of the algorithms^{1,41}. Our results show that MPPMG solvers can address this difficult class of sparse, ill-conditioned systems without sacrificing the accuracy required for scientific simulation.

The MPPMG framework is not restricted to the specific photonic processor used in this work. Any analog platform that can provide an effective smoothing operation can, in principle, be embedded into the framework. This makes the approach a natural interface between numerical analysis and emerging analog hardware. While featuring less bandwidth than its photonic counterpart, electronic analog in-memory computing is attractive for dense operators. Due to the smaller component size, more than four million analog in-memory computing cells can be integrated on-chip with an overall computation accuracy of 3-4 bits⁴². In photonics, scaling to larger matrices will likely require weight encoding across multiple degrees of freedom¹². Free-space photonic processors offer a promising route by exploiting an additional spatial dimension for parallelism and weight storage⁴³. As these systems grow, their efficiency can improve, pointing towards matrix processing with less than one photon per multiply and accumulate operation⁴⁴. In contrast, it remains difficult to efficiently use very low-precision digital tensor cores for sparse matrix-vector multiplications⁴⁵, which highlights the opportunity for physical processors whose native strength is low-latency linear computation.

Photonic multigrid solvers have the potential to provide transformative performance gains for large-scale calculations. Its strength lies in using the photonic processor where it is most effective, for repeated low-latency smoothing operations, while retaining digital computation for high-precision steps. In this role, low-precision optical matrix processing becomes a useful high-accuracy computational primitive rather than an isolated analog approximation. The integrated photonic multigrid solver therefore combines photonic bandwidth, analog parallelism and digital reliability within a single numerical architecture. This provides a scalable path towards accelerating partial differential equations and related sparse scientific computing problems without sacrificing the robustness expected from established digital solvers.

References

1. Clark, M. A. *et al.* Accelerating lattice qcd multigrid on gpus using fine-grained parallelization. In *SC '16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 795–806 (2016).
2. Isotton, G., Frigo, M., Spiezia, N. & Janna, C. Chronos: A general purpose classical amg solver for high performance computing. *SIAM Journal on Scientific Computing* **43**, C335–C357 (2021). URL <https://doi.org/10.1137/21M1398586>. <https://doi.org/10.1137/21M1398586>.
3. Stüben, K., Ruge, J. W., Clees, T. & Gries, S. *Algebraic Multigrid: From Academia to Industry*, 83–119 (Springer International Publishing, Cham, 2017). URL https://doi.org/10.1007/978-3-319-62458-7_5.
4. Mitchell, L. & Müller, E. H. High level implementation of geometric multigrid solvers for finite element problems: Applications in atmospheric modelling. *Journal of Computational Physics* **327**, 1–18 (2016). URL <https://www.sciencedirect.com/science/article/pii/S0021999116304582>.
5. Shalf, J. The future of computing beyond moore’s law. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **378**, 20190061 (2020). URL <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2019.0061>. <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2019.0061>.
6. Hennessy, J. L. & Patterson, D. A. A new golden age for computer architecture. *Commun. ACM* **62**, 48–60 (2019). URL <https://doi.org/10.1145/3282307>.
7. Kashi, A. *et al.* Mixed-precision numerics in scientific applications: survey and perspectives. *The Journal of Supercomputing* **82**, 287 (2026). URL <https://doi.org/10.1007/s11227-026-08264-4>.
8. Dongarra, J., Heroux, M. A. & Luszczek, P. High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems. *The International Journal of High Performance Computing Applications* **30**, 3–10 (2016). URL <https://doi.org/10.1177/1094342015593158>. <https://doi.org/10.1177/1094342015593158>.
9. Hua, S. *et al.* An integrated large-scale photonic accelerator with ultralow latency. *Nature* **640**, 361–367 (2025). URL <https://doi.org/10.1038/s41586-025-08786-6>.
10. Langenegger, J. *et al.* In-memory factorization of holographic perceptual representations. *Nature Nanotechnology* **18**, 479–485 (2023). URL <https://doi.org/10.1038/s41565-023-01357-8>.

11. Shastri, B. J. *et al.* Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics* **15**, 102–114 (2021). URL <https://doi.org/10.1038/s41566-020-00754-y>.
12. Bente, I. *et al.* The potential of multidimensional photonic computing. *Nature Reviews Physics* **7**, 439–450 (2025). URL <https://doi.org/10.1038/s42254-025-00843-3>.
13. Ahmed, S. R. *et al.* Universal photonic artificial intelligence acceleration. *Nature* **640**, 368–374 (2025). URL <https://doi.org/10.1038/s41586-025-08854-x>.
14. Tang, Y. *et al.* Optical neural engine for solving scientific partial differential equations. *Nature Communications* **16**, 4603 (2025). URL <https://doi.org/10.1038/s41467-025-59847-3>.
15. Yuan, H. *et al.* Microcomb-driven photonic chip for solving partial differential equations. *Advanced Photonics* **7**, 016007 (2025). URL <https://doi.org/10.1117/1.AP.7.1.016007>.
16. Tradonsky, C., Wolf, O., Vaknin, T., Gluck, D. & Furman, D. Solving partial differential equations on an analog, optical platform. invited paper. In *Proceedings of the 22nd ACM International Conference on Computing Frontiers: Workshops and Special Sessions*, CF ’25 Companion, 130–134 (Association for Computing Machinery, New York, NY, USA, 2025). URL <https://doi.org/10.1145/3706594.3729406>.
17. Ye, J. *et al.* Reconfigurable application-specific photonic integrated circuit for solving partial differential equations. *Nanophotonics* **13**, 2231–2239 (2024). URL <https://onlinelibrary.wiley.com/doi/abs/10.1515/nanoph-2023-0732>. <https://onlinelibrary.wiley.com/doi/pdf/10.1515/nanoph-2023-0732>.
18. Chen, M. *et al.* I/o-efficient iterative matrix inversion with photonic integrated circuits. *Nature Communications* **15**, 5926 (2024). URL <https://doi.org/10.1038/s41467-024-50302-3>.
19. Attanasio, F. *et al.* Speeding up fermionic lattice calculations with photonic accelerated inverters. *Computer Physics Communications* **317**, 109825 (2025). URL <https://www.sciencedirect.com/science/article/pii/S0010465525003273>.
20. Zhu, Z., Klein, A. B., Li, G. & Pang, S. Fixed-point iterative linear inverse solver with extended precision. *Scientific Reports* **13**, 5198 (2023). URL <https://doi.org/10.1038/s41598-023-32338-5>.
21. Zhu, Z. *et al.* Coherent general-purpose photonic matrix processor. *ACS Photonics* **11**, 1189–1196 (2024). URL <https://doi.org/10.1021/acsp Photonics.3c01694>. <https://doi.org/10.1021/acsp Photonics.3c01694>.

22. Mannocci, P., Larelli, G., Bonomi, M. & Ielmini, D. Achieving high precision in analog in-memory computing systems. *npj Unconventional Computing* **3**, 1 (2026). URL <https://doi.org/10.1038/s44335-025-00044-2>.
23. Meyer, L. *et al.* Deep neural network inference on an integrated, reconfigurable photonic tensor processor. *Nature Communications* **17**, 3396 (2026). URL <https://doi.org/10.1038/s41467-026-71599-2>.
24. Gholami, A., Malhotra, D., Sundar, H. & Biros, G. Fft, fmm, or multigrid? a comparative study of state-of-the-art poisson solvers for uniform and nonuniform grids in the unit cube. *SIAM Journal on Scientific Computing* **38**, C280–C306 (2016). URL <https://doi.org/10.1137/15M1010798>. <https://doi.org/10.1137/15M1010798>.
25. Zaman, M. A. Numerical solution of the poisson equation using finite difference matrix operators. *Electronics* **11** (2022). URL <https://www.mdpi.com/2079-9292/11/15/2365>.
26. Knyazev, A. V. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM Journal on Scientific Computing* **23**, 517–541 (2001). URL <https://doi.org/10.1137/S1064827500366124>. <https://doi.org/10.1137/S1064827500366124>.
27. Duersch, J. A., Shao, M., Yang, C. & Gu, M. A robust and efficient implementation of lobpcg. *SIAM Journal on Scientific Computing* **40**, C655–C676 (2018). URL <https://doi.org/10.1137/17M1129830>. <https://doi.org/10.1137/17M1129830>.
28. Acer, S., Boman, E. G., Glusa, C. A. & Rajamanickam, S. Sphynx: A parallel multi-gpu graph partitioner for distributed-memory systems. *Parallel Computing* **106**, 102769 (2021). URL <https://www.sciencedirect.com/science/article/pii/S0167819121000272>.
29. Yamada, S., Imamura, T., Kano, T. & Machida, M. High-performance computing for exact numerical approaches to quantum many-body problems on the earth simulator. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC '06*, 47–es (Association for Computing Machinery, New York, NY, USA, 2006). URL <https://doi.org/10.1145/1188455.1188504>.
30. Yamada, S., Imamura, T. & Machida, M. High performance parallel lobpcg method for large hamiltonian derived from hubbard model on multi-gpu systems. In Panda, D. K. & Sullivan, M. (eds.) *Supercomputing Frontiers*, 1–19 (Springer International Publishing, Cham, 2022).
31. Rakhuba, M. & Oseledets, I. Calculating vibrational spectra of molecules using tensor train decomposition. *The Journal of Chemical Physics* **145**, 124101 (2016). URL <https://doi.org/10.1063/1.4962420>. <https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.4962420/15517781/124101.1-online.pdf>.

32. Shao, M. *et al.* Accelerating nuclear configuration interaction calculations through a preconditioned block iterative eigensolver. *Computer Physics Communications* **222**, 1–13 (2018). URL <https://www.sciencedirect.com/science/article/pii/S0010465517302904>.
33. Fialko, S. & Żegleń, F. Block subspace projection pcg method for solution of natural vibration problem in structural analysis. In *2016 Federated Conference on Computer Science and Information Systems (FedCSIS)*, 669–672 (2016).
34. Knyazev, A. V. & Neymeyr, K. Efficient solution of symmetric eigenvalue problems using multigrid preconditioners in the locally optimal block conjugate gradient method. *Electron. Trans. Numer. Anal.* **15**, 38–55 (2003).
35. Boyle, P. *et al.* Lattice qcd and the computational frontier (2022). URL <https://arxiv.org/abs/2204.00039>. 2204.00039.
36. Brannick, J., Brower, R. C., Clark, M. A., Osborn, J. C. & Rebbi, C. Adaptive multigrid algorithm for lattice qcd. *Phys. Rev. Lett.* **100**, 041601 (2008). URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.041601>.
37. Xu, J. & Zikatanov, L. Algebraic multigrid methods. *Acta Numerica* **26**, 591–721 (2017).
38. Nakamura, Y. *et al.* Lattice qcd applications on qspace. *Procedia Computer Science* **4**, 841–851 (2011). URL <https://www.sciencedirect.com/science/article/pii/S1877050911001475>. Proceedings of the International Conference on Computational Science, ICCS 2011.
39. Belletti, F. *et al.* The apenext project. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **559**, 90–94 (2006). URL <https://www.sciencedirect.com/science/article/pii/S0168900205022382>. Proceedings of the X International Workshop on Advanced Computing and Analysis Techniques in Physics Research.
40. Joó, B. *et al.* Status and future perspectives for lattice gauge theory calculations to the exascale and beyond. *The European Physical Journal A* **55**, 199 (2019). URL <https://doi.org/10.1140/epja/i2019-12919-7>.
41. Clark, K., Howarth, D., Tu, J., Wagner, M. & Weinberg, E. Maximizing the bang per bit. In *Proceedings of The 39th International Symposium on Lattice Field Theory — PoS(LATTICE2022)*, LATTICE2022, 338 (Sissa Medialab, 2023). URL <http://dx.doi.org/10.22323/1.430.0338>.
42. Le Gallo, M. *et al.* A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference. *Nature Electronics* **6**, 680–693 (2023). URL <https://doi.org/10.1038/s41928-023-01010-1>.

- 407 43. Kalinin, K. P. *et al.* Analog optical computer for ai inference and combinatorial optimization.
408 *Nature* **645**, 354–361 (2025). URL [https://doi.org/10.1038/s41586-025-09430-](https://doi.org/10.1038/s41586-025-09430-z)
409 [z](https://doi.org/10.1038/s41586-025-09430-z).
- 410 44. Wang, T. *et al.* An optical neural network using less than 1 photon per multiplication. *Nature*
411 *Communications* **13**, 123 (2022). URL [https://doi.org/10.1038/s41467-021-](https://doi.org/10.1038/s41467-021-27774-8)
412 [27774-8](https://doi.org/10.1038/s41467-021-27774-8).
- 413 45. Zhang, L., Huang, J., Di, S., Matsuoka, S. & Wahib, M. Can tensor cores benefit memory-bound
414 kernels? (no!). In *Proceedings of the 17th Workshop on General Purpose Processing Using*
415 *GPU, GPGPU '25*, 28–34 (Association for Computing Machinery, New York, NY, USA, 2025).
416 URL <https://doi.org/10.1145/3725798.3725803>.
- 417 46. Davis, T. A. & Hu, Y. The university of florida sparse matrix collection. *ACM Trans. Math.*
418 *Softw.* **38** (2011). URL <https://doi.org/10.1145/2049662.2049663>.
- 419 47. Maris, P., Yang, C., Oryspayev, D. & Cook, B. Accelerating an iterative eigensolver for
420 nuclear structure configuration interaction calculations on gpus using openacc. *Journal of*
421 *Computational Science* **59**, 101554 (2022). URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1877750321002076)
422 [science/article/pii/S1877750321002076](https://www.sciencedirect.com/science/article/pii/S1877750321002076).
- 423 48. Bouwmeester, H., Dougherty, A. & Knyazev, A. V. Nonsymmetric preconditioning
424 for conjugate gradient and steepest descent methods1. *Procedia Computer Science* **51**,
425 276–285 (2015). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S1877050915010492)
426 [pii/S1877050915010492](https://www.sciencedirect.com/science/article/pii/S1877050915010492). International Conference On Computational Science, ICCS
427 2015.
- 428 49. Bell, N., Olson, L. N., Schroder, J. & Southworth, B. PyAMG: Algebraic multigrid solvers
429 in python. *Journal of Open Source Software* **8**, 5495 (2023). URL [https://doi.org/](https://doi.org/10.21105/joss.05495)
430 [10.21105/joss.05495](https://doi.org/10.21105/joss.05495).




Figure 1: Hybrid Partial Differential Equation Solver. **a)** Many complex theoretical models, for example, describing the physics of nucleons in the framework of lattice quantum chromodynamics (LQCD), can not be solved analytically. Discretizing the system enables numerical approximation but leads to large, sparse linear systems, like the LQCD matrix found on the SuiteSparse matrix collection⁴⁶, which are hard to solve, even on the current generation of supercomputers³⁵. **b)** We develop hybrid, mixed-precision partial differential equation solvers that leverage both the advantages of high-precision digital computation and low-latency analog computation. The majority of the computation is offloaded to the analog domain to improve the time-to-solution, while a few high-precision digital operations ensure high-accuracy convergence. **c)** We use multigrid methods, transferring the problem to multiple coarser discretization levels. At each level, we eliminate high-frequency error components using an analog smoother to transfer the problem to the coarser grids. At the coarsest level, described by the smallest linear system, we cheaply solve the problem and use this as an error correction to improve the solution estimate on the finer grids. Using this hybrid multigrid solver as a preconditioner, we drastically reduce the number of digital operations while maintaining high accuracy.




Figure 2: Ultra-low Latency Photonic Smoother. **a)** The photonic smoother deploys a crossbar array for in-memory computation. Electro-absorption modulators (EAMs) encode the input vector z and store the matrix weights A while photodiodes (PDs) read out the result y of the matrix-vector multiplication. The underlying crossbar routing mesh optically distributes the input vector to the individual matrix cells and accumulates the respective weighted inputs. The analog photonic computation itself avoids pipelining at optical data rates, enabling ultra-low latency computation. **b)** The photonic circuit encodes and processes information directly in the amplitude of light pulses rather than deploying an abstract binary encoding. Thus, physical noise sources directly impact the accuracy of the computation. For 100 random matrix-vector multiplications, the photonic processor exhibits an L_2 relative error of 7.7%, comparable to 4-bit digital fixed-point arithmetic, see Supplementary Information. **c)** The initial error has different frequency components, as described in Eq. (1), and the task of the smoother is to dampen the high-frequency ones. While the analog noise affects the output of the photonic smoother compared to the same smoother in the digital domain, we still obtain the required attenuation of the high-frequency error components to efficiently reach high-accuracy solutions.




Figure 3: High-Accuracy Photonic Multigrid Methods. **a) and d)** We use the ultra-low latency photonic smoother in the proposed mixed-precision photonic multigrid (MPPMG) framework to calculate the electric field of a parallel capacitor and the first eigenpairs $(\varphi_n, \mathcal{E}_n)$ of the quantum quartic anharmonic oscillator. **b) and e)** While the solution accuracy of fully analog solvers is inherently constrained by analog noise²⁰, the hybrid MPPMG solver compensates the noise with a sparse use of high-precision digital matrix-vector multiplications. Thus, the hybrid systems converge with a residual norm below $|r| = 10^{-10}$, equal to the accuracy reachable by fully digital solvers. **c) and f)** We benchmark the same solver on the digital domain with 64-bit precision (baseline), on the hybrid system with the photonic smoother (exp.), and on the digital twins of the photonic smoother with effectively 4 bits (emul.) and exactly 8 bits (ideal). To achieve an accuracy of $|r| = 10^{-10}$, the hybrid system requires 60% and 80% less high-precision digital operations than the baseline, respectively. The digital twin behaves similarly as to the photonic processor, and increasing the accuracy of the photonic smoother to 8 bit could reduce the number of high-precision digital operations by 80% and 85%, respectively.




Figure 4: Photonic Lattice QCD Accelerator. **a)** Quantum chromodynamics models the fundamental physics governing nucleons. After discretizing the system in space and time, we obtain a matrix D describing the dynamics of quarks on a gluon background g represented by the gauge links configurations U . As the bare quark mass m is tuned to the physically interesting regime, the condition number in Eq. (11) diverges. We deploy a hybrid adaptive scheme based on the photonic smoother to obtain a multigrid preconditioner, which is crucial for numerically hard problems with a high condition number, especially arising for realistic values of the quark masses. **b)** The bare quark mass m in Eq. (7) is adjusted to change the condition number of the matrix A , which determines the numerical difficulty of the linear systems to be solved. As expected, the performance of the hybrid solver (emulated) based on the IR method deteriorates quickly as we approach physically interesting values of the bare quark mass^{20,22}, while the MPPMG solver (emulated) remains robust for linear systems with high condition numbers. **c)** The standard method for solving a single RHS is the conjugate gradient method. The MPPMG approach, now with the adaptive scheme, considerably reduces the required number of digital operations while maintaining high-accuracy solutions with $|r| = 10^{-10}$ for $m = -0.063$. More information can be found in the Supplementary Information.

Acknowledgments

We thank Jochen Stuhmann, from Illustrato, for his assistance with the illustrations. This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) by the Collaborative Research Centre SFB 1225 - 273811115 (ISOQUANT) and Germany's Excellence Strategy EXC 2181/1 - 390900948 (the Heidelberg STRUCTURES Excellence Cluster). We further acknowledge funding for this work from the European Union's Horizon 2020 Research and Innovation Program, Darpa NAPSAC and ERC AdG PICNIC.

Author contributions: Conceptualization: TL, FB, JP, WP

Methodology: TL, FB, JD

Investigation: TL, FB

Visualization: TL, FB, WP

Funding acquisition: JP, WP

Project administration: JP, WP

Supervision: FB, JP, WP

Writing - original draft: TL

Writing - review & editing: All authors

Competing interests: The authors declare that they have no competing interests.

Data and materials availability: All data are available in the main text or the supplementary materials.

Methods

Experimental setup

We use a setup that includes an amplified spontaneous emission light source (Agilent 83438A), an incoherent optical 9x3 crossbar array fabricated by imec on a silicon-on-insulator (SOI) platform integrating electro-absorption modulators for input modulation and weighting and photodiodes for readout. We control the full system by a SoC-FPGA (Xilinx ZCU 216) featuring 4 GSPS DACs and 2 GSPS ADCs for fast optical processing²³. The weights for in-memory computing are programmed via a slower multi-channel DAC. We round the matrix values to 4-bit precision to reduce the amount of slow weight changes, and the input values to 8-bit precision. Due to the intrinsic optical noise of the ASE light source, we repeat the calculations 256 times and take the average as the result of the calculations. With this setup, we achieve errors comparable to digital 4-bit precision fixed-point calculations. Using a noise-free coherent light source reduces the need for averaging²³.

Digital Noisy Fixed-Point Operations

We use digital noisy fixed-point operations to mimic the analog nature of our experimental setup for the LQCD calculation by rescaling the input and matrix values to the range $[-1,1]$ with 8-bit discretization, adding Gaussian noise with a standard deviation of 0.045, performing the matrix-vector multiplications, and rescaling the output accordingly.

Residual

We use the norm of the residual, defined as

$$r := b - Ax, \quad (8)$$

as a metric of accuracy, which is common in lattice QCD calculations³⁶. This is related to the error defined in Eq. (1) as

$$Ae = r. \quad (9)$$

Estimation of the performance gain

We introduce a metric to evaluate the performance of the proposed solver. The goal of an efficient hybrid photonic solver is to offload most of the calculations to the optical domain. We estimate the maximum possible performance gain by assuming that the speed and energy efficiency of specialized hardware based on PICs are infinitely better than those of their digital counterparts. This is expressed by the quantity

$$G = \frac{N_d^{(S)}}{N_d^{(MP)}}, \quad (10)$$

where $N_d^{(S)}$ and $N_d^{(MP)}$ are the number of double-precision operations required by the double-precision solver and the mixed-precision solver, respectively. This would imply that current technical challenges, like data movement, are solved. We note that this value represents the upper limit of the performance gain of the proposed mixed-precision photonic solver. Additionally, in Fig. 3 and Fig. 4, we make the approximation that the MVMs with the matrix A are the only operations relevant to the computation cost, and that the cost of the solver in the coarsest grid is negligible, which is the default approximation in literature^{36,47}.

Condition number

The number of iterations required by an iterative linear solver typically increases with the condition number of the system matrix, generally defined as

$$\kappa(A) = \|A\| \|A^{-1}\|. \quad (11)$$

For symmetric matrices, this is simplified to

$$\kappa(A) = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}, \quad (12)$$

where $\lambda_{\min}$ and $\lambda_{\max}$ are the smallest and largest eigenvalues of A . Hybrid approaches based on the iterative refinement method struggle with problems with high condition numbers^{20,22}. In the LQCD problem, we tune the bare quark mass parameter of the model in Eq. (7) such that we approach the limit where the physical quark mass goes to zero, which is equivalent to $\lambda_{\min}$ going to zero, leading to the divergence of the condition number.

Iterative refinement Method

Intuitively, the iterative refinement (IR) method resets the solver's accumulation of rounding error, thereby improving numerical stability. The steps of this method are

1. Choose initial guess x .
2. Calculate residual $r = b - Ax$.
3. Approximate the solution of $Ae = r$ with some iterations of the solver.
4. Update $x \leftarrow x + e$.
5. Repeat 2 to 4 until convergence with the desired accuracy.

Details on the multigrid implementation

Due to empirical observations, we choose to perform asymmetrical smoothing for the mixed-precision solvers as in the work by Bouwmeester et al.⁴⁸, where no post-smoothing is done. For a fair comparison, we perform four photonic pre-smoothing steps for the hybrid solvers, and two pre-smoothing and two post-smoothing steps for the digital solvers, keeping the total number of smoothing steps per multigrid iteration equal. More details can be found in the Supplementary Information.

Lattice Quantum Chromodynamics

We test our solver on a benchmark similar to Brannick et al.³⁶. We provide further necessary data to reproduce the results in Fig. 4. We use the quench approximation to generate U(1) gauge link configurations on a 128 x 128 lattice with the action

$$S[U] = \beta \sum_{n \in \Lambda} \sum_{\mu < \nu} \text{Re tr} [\mathbb{1} - U_{\mu\nu}(n)], \quad (13)$$

with periodic boundary conditions and with $\beta = 6$. We generate the gauge links configurations U using a Metropolis-Hastings algorithm.

Adaptive Multigrid Method

The adaptive part is important to compute the space spanned by the low modes of the matrix, which are particularly challenging in lattice QCD calculations. We calculate these low modes by sampling

520 a random Gaussian initial guess x , approximately solving $Ax = 0$ by applying the smoother in
521 Eq. (2) many times to x , saving the result, and repeating this process 8 times in total. We block
522 the lattice into 4×4 blocks for building the coarse matrix. For the implementation, we build the
523 aggregate operator based on the 4×4 blocks and use the calculated low-modes and the PyAMG
524 library⁴⁹ to construct a smoothed aggregation multigrid solver.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SIsuubmit.pdf](#)