

Quantum Software Lab


ANNUAL REPORT

From Pathfinder to Turbocharger Engine

Making quantum advantage useful, trusted and deliverable

Quantum computing is moving from promise to delivery. The challenge is not only to build larger machines, but to ensure quantum computation delivers useful and deployable answers. This requires a full software layer that brings together algorithms, applications, architectures, high-performance computing, error correction, benchmarking and verification.

QSL was created to close the gap between algorithms and machine, to build the foundation of the software layer.



Since 2023, QSL has pioneered the journey from a vision for quantum software to a concrete agenda and new national capabilities.

QAP

The Quantum Advantage Pathfinder (QAP) programme created the methodology. It has allowed us to systematically understand real-world pain points, translate them into mathematical specifications, design quantum and quantum-inspired methods, implement on emerging hardware, and benchmark and verify results. It enabled QSL to become an anchor for testbed evaluation, application discovery, international coordination, skills and talent, verification and benchmarking, and a sector-facing quantum software strategy.

QATCH

Quantum Advantage TurboCHarger (QATCH) is a new programme that builds on, and scales from, QAP. QATCH is the delivery engine for the software layer, moving beyond the pathfinding regime to provide the tools, methods, and systems that will deliver significant economic and social impact in the UK.



QSL

Board Update

This annual update marks the Quantum Software Lab's third anniversary. Quantum technology has moved from laboratory to industry, and the UK has positioned itself at the forefront of that transition with a ten-year National Quantum Strategy. The United Nations designation of 2025 as the International Year of Quantum Science and Technology reflected a global recognition that the next chapter of quantum computing requires institutional capability in equal measure to hardware investment.

As history tells us, lasting value will not come from hardware alone. It belongs to those who champion the algorithms, software and applications layer: translating physical qubits into useful answers for chemistry, finance, life sciences, security, materials, manufacturing, AI, and energy.

Over the past year, the ecosystem has made clear its requirements for accelerated growth: international collaboration, applications excellence, the co-design of hardware and software, trusted benchmarking, investment in skills and talent, and structural delivery.


QSL has proved that it can meet these urgent demands. Underpinning our response to the needs of the ecosystem is an institutional model built meticulously over three years. QSL is home to faculty, researchers, students, software engineers and professional staff from across scientific disciplines; working together with hardware experts, domain specialists, and policymakers towards a single purpose and mission: the pursuit of quantum advantage. It lies at the very heart of who we are, the culture we have built, the motivation behind QATCH, and how we engage collectively with the NQCC and the wider quantum ecosystem.



**Quantum Software Lab
Board Members**

The QSL Methodology

QSL's core methodology is simple, but it is also hard to execute.



This is the Full-Stack Feedback Loop.

It starts with a real-world challenge that is translated into a mathematical object with measurable figures of merit. Algorithms and protocols are designed against that specification, and then implemented on available or emerging hardware using novel programming techniques. The results are benchmarked against the best classical baselines and verified against clear claims of correctness, performance and reliability throughout the process.

This final step is the most important. Evidence does not stop at a paper or a proof-of-concept demonstration. It feeds into the future design of the software layer, and has a direct impact on the hardware requirements, performance evaluation, and the development of killer applications.

This is how QSL transforms quantum advantage from an undefined slogan into real and evidenced impact.

Our Impact

This core methodology has enabled QSL to respond to the urgent needs of the UK quantum ecosystem, with key achievements over 3 years comprising:



130+

Conference Talks


Invited and contributed talks in major conferences: including QIP, POPL, PLDI, CRYPTO, FOCS, SODA, QCTIP, OPTICA, STOC, QPL, QCrypt, TQC, APS March Meeting



70+

Quantum Researchers

Full-time quantum researchers covering all aspects of the quantum software stack



170+

Publications

Including Nature, Nature Communications, Proceedings of the National Academy of Sciences, Physical Review Letters, PRX Quantum, and Quantum, Quantum Science and Technology



75+

Internal seminars

< QSL >
Quantum Software Lab



60+

Visitors

From 35+ institutions



55+

Active projects & Collaborations

With external partners (academic and industry)



40+

PhD Students

Graduated or in training through the QI CDT and other direct sponsorships via QCI3, NQCC and industries



55+

International Events & Trade Missions

Including Q2B (Paris, Silicon Valley, Tokyo), Quantum.Tech (UK, USA), Economist Commercialising Quantum, Quantum Korea, Quantum Innovation Summit Dubai, UK Quantum Showcase



QSL has delivered the needs of the ecosystem for international cooperation, application excellence, hardware-software co-design, trusted benchmarking, workforce training and structured delivery.

The story of Year 3 is therefore not a collection of separate achievements. It is the moment when an institutional model empowers real success.

Research Depth Across the Stack

QSL's research strength spans the full quantum software stack. QSL's third anniversary is an opportunity to reflect and celebrate the achievements of all of our researchers and students who have made significant contributions to QSL's success and advancing the field.



In **algorithms and quantum machine learning**, QSL develops methods for identifying where quantum resources can genuinely outperform classical ones, and where quantum inspired classical methods offer a better solution. Year 3 work has examined the conditions under which quantum generative models can be classically surrogated and where they cannot, and the trainability and expressivity of subspace-preserving and Hamming-weight preserving quantum neural architectures. Algorithmic and certification advances, including new sample bounds for trace estimation, quantum-enhanced versions of state-of-the-art Markov-Chain-Monte-Carlo based classical optimisation heuristics, and a Quantum Hamiltonian Certification framework, extend the rigour with which quantum advantage claims can be established.



In **simulation**, QSL connects quantum theory, chemistry, materials, physics and HPC to study regimes where hybrid quantum-classical workflows can make a difference. A Quantum Elastic Network Model framework, developed in partnership with industry collaborator, demonstrated exponentially advantaged subroutines for the dynamics of coupled-oscillator systems applied to graphene, and reached the semi-finals of the XPRIZE Quantum Applications. Programmable tools for analogue quantum simulation in HPC contexts, and local quench spectroscopy on digital quantum computers, extend the same translational arc into condensed-matter and high-energy regimes.


QSL's research strength is not concentrated in one topic. It spans the full quantum software stack.

In **programming languages, compilation and HPC integration**, QSL builds the abstractions needed to connect high-level applications to heterogeneous quantum and classical systems. Year 3 contributions include an equational quantum programming framework, efficient verification methods for quantum programs and fault-tolerant codes, and an API for OpenMP for quantum task offloading. An energy-efficient model for simulating quantum state vectors at scale was also developed jointly with the EPCC.



In **distributed and networked quantum computing**, QSL develops secure and verifiable architectures for multi-user and multi-node computation. This year we co-authored the design and demonstration of an operating system for executing applications on quantum network nodes, and contributed to an experimental multi-client verifiable blind quantum computation on a photonic Qline architecture. Companion theoretical work on hybrid authentication protocols and selectively blind quantum computation extends these foundations to advanced quantum networks.

In **error correction, resource estimation, benchmarking and verification**, QSL establishes the assurance layer that provides trust in claims of quantum advantage. This includes work on zero-trust verification to provide a secure approach for tasks aiming to achieve near-term quantum advantage, a framework for efficient certification of fault-tolerant computations, and verification-inspired benchmarking combining a heuristic-free protocol with a community-wide review of metrics and benchmarks. Other activities include entropy benchmarking of near-term devices, QEC decoding, fault-tolerant compilation, and resource estimation for verifiable quantum advantage workflows. Foundational results on the role of coherence for quantum computational advantage anchor these protocols in the deeper structure of quantum resource theory.



Partners and Collaborators

QSL aims to translate theory into application for the benefit of our economy and society.

This objective cannot be achieved without a vibrant ecosystem of partners and collaborators. QSL plays a vital role in the UK National Quantum Technologies Programme and landscape.

Working with our strategic partner, the **National Quantum Computing Centre (NQCC)**, QSL has played an integral role in the hardware testbed initiative, boosting end-user adoption through the SparQ programme, and co-designing software to work with the hardware. Further, QSL has worked with **National Physical Laboratory (NPL)** on benchmarking and standards, with the **QCi3 Hub** and the **IQN Hub** on hardware, integration and networked quantum systems, and with the **National Cyber Security Centre (NCSC)** on quantum-safe cybersecurity engagement. Through the **EPSRC Centre for Doctoral Training in Quantum Informatics**, QSL researchers work with Heriot-Watt University, University of Strathclyde, Oxford University and University College London to train the next generation of talent.



Beyond these national partners, QSL also has active collaborations with leading international industries.

- **Quantum / classical / HPC vendors:** Aegiq, Alice & Bob, Atom Computing, AWS, Fujitsu, IBM, Infleqtion, IonQ, Nvidia, Nu Quantum, ORCA Computing, Pasqal, Photonic Inc., Quandela, Quantinuum, Quantum Motion, Rigetti, SeeQC, Xanadu
- **Quantum software / control / algorithms vendors:** Algorithmiq, Cambridge Consultants, Cisco, Classiq, Input Output, JiJ, Phascraft, Q-CTRL, Qinara, Quantum Base Alpha, Riverlane, VeriQloud
- **End-users:** British Telecommunications (BT), CERN, HSBC, National Energy Systems Operator, NatWest, Nomura Investments, Rolls-Royce, Scottish Government, Standard Chartered
- **Standards and learned-society partners:** Institute of Physics, National Physical Laboratory, Financial Conduct Authority, Ofgem, Hartree Centre




QSL's international academic collaborations are strengthened through its role as a founding member of the **Quantum Software Alliance** (QSA), an independent international consortium formally launched on 10 February 2026 in Ghana at the closing ceremony of the United Nations International Year of Quantum Science and Technology. Bringing together 21 founding research institutions across the UK, Europe, North America, Asia and the Middle East, the Alliance aims to co-ordinate research agendas, inform decision makers, identify and secure new funding opportunities, and address the grand challenges of the field.



Year 3 research collaborations have included extended work with our international partners. QSL members have worked with the Free University Berlin and the Technical University of Munich on quantum error correction; with the Sapienza University of Rome on photonic quantum computing; with the INRIA QAT team in Paris on verifiable quantum computing; with the University of Tokyo on higher order operations for quantum protocols, and with the University of Twente on hardware security modules for secure delegated quantum cloud computing. Through the EPSRC Centre for Doctoral Training in Quantum Informatics, QSL researchers help to address talent development needs with 11 overseas research institutes including: QuSoft (Netherlands), Qmath (Denmark), Sydney Quantum Academy (Australia), the Institute for Quantum Computing (Canada), and Lawrence Berkeley National Laboratory (USA).



To bring together the wider ecosystem, QSL has led an ambitious programme of events and public outreach. The flagship is our **Quantum Fringe**, providing an open platform for workshops, conferences, hackathons, public lectures and community exchange.



Quantum Fringe 2026

Quantum Fringe 2026 features **14 events** taking place across the UK from **May to June 2026**.

PEOPLE: FULL-STACK ARCHITECTS

QSL's long-term output is not only papers, software, and demonstrators. It is people.



This human architecture matters. A national quantum strategy cannot be delivered by hardware procurement alone. It requires software, and above all people who understand the whole stack and can defend claims of utility with evidence.

Our environment fosters and develops full-stack quantum architects: researchers who can move between mathematical foundations, algorithms, applications, programming systems, HPC integration, error correction, verification, benchmarking, and hardware constraints.

This skills and talent model is visible across QSL. It starts with the EPSRC Centre for Doctoral Training in Quantum Informatics, which brings together universities and partners to inspire and train a new generation of quantum software researchers. QATCH will extend QSL's existing team of 70+ members, adding new tenure-track Quantum Fellows and several new research and joint innovation positions with external partners. Through these initiatives QSL is building a pipeline of future leaders that will transform theory into application and deliver the promise of quantum advantage.

DIRECTOR

Elham Kashefi — Director

FACULTY

Alex Cojocar — Faculty

Ava Khamseh — Faculty

Chris Dent — Faculty

Chris Heunen — Faculty

Julien Michel — Faculty

Joschka Roffe — Faculty

Lukas Szpruch — Faculty

Mina Doosti — Faculty

Myrto Arapinis — Faculty

Oliver Brown — Faculty

Petros Wallden — Faculty

Raul Garcia Patron Sanchez — Faculty

Steven Thomson — Faculty

SENIOR RESEARCHER

Brian Coyle — Senior Researcher (Testbed Technical Lead)

Caitlin Jones — Senior Researcher (Testbed Technical Lead)

Dominik Leichtle — Senior Researcher

Jinge Bao — Senior Researcher

Qisheng Wang — Senior Researcher (Former Member)

Ross Grassie — Senior Researcher (Innovation Fellow)

Stergios Koutsoumpas — Senior Researcher

Suchetana Goswami — Senior Researcher

William Waites — Senior Researcher, Software Developer

RESEARCH ASSOCIATE

Alex Rice — Research Associate

Andrew Jackson — Research Associate (Former Member)

Ben Symons — Research Associate (Former Member)

James Mills — Research Associate

Kenza Hammam — Research Associate

Kit Battarbee — Research Associate

Leo Monbroussou — Research Associate

Louis Lemonnier — Research Associate

Malin Altenmuller — Research Associate

Marine Demarty — Research Associate (Former PhD Student)

Nicolas Huertel — Research Associate

Pauline Besserve — Research Associate

Sara Metwalli — Research Associate

Sean Adamson — Research Associate

Wang Fang — Research Associate

RESEARCH ASSISTANT

David Halpenny — Research Assistant

Feroz Hassan — Research Assistant

Virag Umathe — Research Assistant, Software Developer

PHD STUDENT

Adi Sireesh – PhD Student
Chirag Wadhwa – PhD Student
Holly Farler – PhD Student
Jakub Adamski – PhD Student
Kano Bacho – PhD Student
Kengo Hirata – PhD Student
Kim Worrall – PhD Student
Kyle Campbell – PhD Student
Liam Veeder-Sweeney – PhD Student
Maria Bahna – PhD Student
Mario Herrero Gonzalez – PhD Student
María Gragera Garcés – PhD Student
Mateusz Meller – PhD Student
Oscar Miller – PhD Student
Roo Dunnill – PhD Student
Sam Winter – PhD Student
Sean Purcell – PhD Student
Sean Thrasher – PhD Student
Stuart Ferguson – PhD Student
Tamas Noszko – PhD Student
Tom Dodd – PhD Student
Tommy Williams – PhD Student
Weixi Zhang – PhD Student

SUPPORT TEAM

Craig Skeldon – Head of Business & Operations
Ramin Jafarzadegan – Project Manager
Laura Zhou – Event and Marketing Manager

QUANTUM INFORMATICS CDT

Jared de Bruin – CDT Manager, Quantum Informatics

Anahita Manchala – CDT Student
Aristomenis Zazanis – CDT Student
David Byfield – CDT Student
Emma Hughes – CDT Student
Gabriela Pinheiro Costa – CDT Student
Haytham McDowall-Rose – CDT Student
Owen Cundy – CDT Student
Peter Shanahan – CDT Student
Philip Ilono – CDT Student
Sophia Widmer – CDT Student
Victor Fischer – CDT Student



As we reflect on the last three years, what we are most proud of is the creation of an institutional model where the full quantum software stack is developed in one place: from mathematical foundations to applications, from programming systems to HPC integration, from error correction to verification, and from academic insight to national testbed evidence.

This is what the UK needs for the quantum mission decade.

QAP showed that the method works.

Year 3 showed that the institution can empower and accelerate progress.

QATCH now turns it into delivery.

*QSL is the UK's quantum software capability:
making quantum advantage useful, trusted and deliverable.*

RESEARCH AREAS

Our various activities are targeting scientific knowledge creation in several complementary domains: Foundation of Quantum Information, Programming and Modeling, Software Development, Algorithm and Protocol Design, Complexity Analysis, Error Correction-Mitigation, Verification, Benchmarking, etc.



APPLICATION

The QSL puts applications in front and centre, fuelling the development of the field in the direction that is most useful societally.



ARCHITECTURE, PROGRAMMING HPC INTEGRATION

We design a network architecture of quantumly enhanced hybrid devices, each possessing, alongside HPC units, a powerful Quantum Processor, that can quantumly access large amounts of data and perform computations that would have been impossible to perform with any classical computing machine.



QUANTUM CYBER SECURITY

Adopting a hybrid approach, which integrates quantum and classical elements, we explore various networking scenarios that span from the near-term post-quantum cryptography to the distant future of the quantum internet.



NOISE HANDLING: QUANTUM ERROR CORRECTION, MITIGATION AND VERIFICATION

Our approach encompasses a range of tools, including benchmarking, verification, and error characterization, mitigation and correction to assess and derive the best performance of different hardware platforms for specific applications.



QUANTUM ALGORITHMS AND MACHINE LEARNING

Quantum information processing is about finding algorithms and protocols that can solve problems more efficiently than their classical counterparts by storing and manipulating information within quantum systems.

Selected Research Outcomes



APPLICATION

Caitlin Jones, Using Spatially Informed Interpretable Quantum Neural Networks for Classifying Cancerous Mutations

Pauline Besserve, Towards large-scale correlated states preparation

Sean Thrasher, Adiabatic-Inspired Hybrid Quantum-Classical Methods for Molecular Ground State

Snehal Raj, Beyond Weisfeiler-Lehman Limitations with Scalable Quantum Graph Neural Networks

Stuart Ferguson, Quantum-enhanced Parallel Tempering

Stuart Ferguson, Quantum-enhanced Simulated Annealing

Stuart Ferguson (XPRIZE), Quantum Elastic Network Models and their Application to Graphene

Weixi Zhang, Dynamical Mean Field Theory with Interacting Bath

ARCHITECTURE, PROGRAMMING & HPC INTEGRATION

Maria Bahna, Exploring Performance and Power of CXL Memory and PCIe 5.0 NVMe for Memory-Hungry Workloads

Nicolas Huertel, Towards a generalisation of complete equational theories

Virag Umathe, QSL Software Stack

William Waites, Consulting the Unwritten Manual

QUANTUM CYBERSECURITY

Abbas Poshtvan, Selective Blindness: A new framework for Blind Quantum Computation

Kano Bacho, Eavesdropper-Blinded Remote State Preparation & Applications to Quantum Public-Key Encryption

Kit Battarbee, Authentication for Quantum Networks

Suchetana Goswami, Partial blindness for sub-universal photonic quantum computing schemes

NOISE HANDLING: QUANTUM ERROR CORRECTION, MITIGATION AND VERIFICATION

Bo Yang, Verifiable Blind Probabilistic Error Cancellation

Dominik Leichtle, Plugging Leaks in Fault-Tolerant Quantum Cryptography

James Mills, Logical Accreditation: A Framework for Efficient Certification of Fault-Tolerant Computations

Marine Demarty, Error-mitigation aware benchmarking strategy for quantum optimization problems

Stergios Koutsoumpas, Colour Codes Reach Surface-Code Performance with Vibe Decoding

Suchetana Goswami, Non-Markovian noise limits for sustaining entanglement in multipartite quantum states

Tom Dodd, A Fast and Frugal Gaussian Boson Sampling Simulator

QUANTUM ALGORITHMS AND MACHINE LEARNING

Brian Coyle, Adaptive directional gradients for parameterised quantum circuits

Chirag Wadhwa, Distributed Quantum Property Testing with Communication Constraints

David Halpenny & Adi, Improved algorithms for the Quantum Fourier Transform on prime-order groups

Elliott Mamon, Orbit dimensions in bosonic systems

Hugo Thomas, Shedding light on classical shadows: learning photonic quantum states

Jing Bao, On Estimating the Quantum Tsallis Relative Entropy

Kyle Campbell, PQC Pruning with Circuit Harmonic Matrices

Léo Monbroussou, Boson sampling beyond the dilute regime: second moments and anticoncentration

Mario Herrero Gonzalez, Exploiting the Gradient Flow to Characterize Parameterized Quantum Circuit Landscapes

Mario Herrero Gonzalez, Simulability in Quantum Generative Models: Conditions for Advantageous Quantum Generative Modelling

Sam Winter, Average Case Sampling Complexity of QAOA

Sean Adamson, Benincasa-Dowker-Glaser causal set actions by quantum counting

Tigran Sedrakyan, Private training in quantum machine learning

Using Spatially Informed Interpretable Quantum Neural Networks for Classifying Cancerous Mutations

➤ **Caitlin Jones, Senior Researcher, cjones18@ed.ac.uk**


Hugh Warden, Sara Metwalli, Tommy Williams, Siddharth Setlur, Luke Boulter, Sjoerd Beentjes, Mina Doosti, Petros Wallden, Elham Kashefi, Ava Khamseh

Problem


Different cancerous mutations alter cellular morphology and the structure of tissue in different ways. This means that cellular morphology can be used diagnostically to quickly predict these mutations – potentially leading to faster and cheaper testing for patients in a clinical setting.

The current state of the art in this area is to train large, weakly supervised, foundation models on large tissue sections [1]. Here we investigate if a quantum machine learning approach can better capture the relationship between these morphological properties and improve over pixel based classification approaches.

KRAS + P53 Mutation




CTNNB1 + P53 Mutation



Examples of imaged tumours. Nuclei are stained with haematoxylin in blue and cancerous cells are stained with DAB in brown


Solution



Identify Nuclei in Image

Measure Nucleus Morphology:

- Area
- Perimeter
- Eccentricity
- Minor Axis Length




Measure morphology of each cell

Measure mean morphology of cells in successive annuli around the cell. Append this information to the single cell measurements along with counts of cells in each annulus

Cancerous cells are biologically marked with a stain. Each tissue has a known combination of cancerous mutations and cancerous cells are filtered for using the staining

We use a quantum neural network (QNN) [2] to classify the cancer cell treating its surrounding environment as separate quantum subsystems. Cellular features extracted from each region are encoded in independent quantum registers, and a parameterized interaction Hamiltonian is used to model the relationships between the primary cell and its neighboring regions.



Unlike many QNNs that function as black-box models [3], the interaction strengths between regions provide interpretable information about the relative influence regions. These learned parameters may help identify which neighboring regions contribute most strongly to the classification outcomes and potentially reveal biologically relevant patterns of cell to cell communication.





Outcome

Our quantum model will be benchmarked against a state of the art classical model, using the same feature set. These benchmarks will be rigorously compared using constrained budget hyperparameter optimization to find the best performing configurations of each model.



Sample plot showing reconstructed coupling coefficients λ between different regions of the image, modeled by the Hamiltonian in the RHS box. Error bars are 95% CI over 15 runs of the QNN model. These initial results suggest that the innermost region is the most relevant for image classification.

The QNN used in this project is a first order Trotter expansion of the following Hamiltonian:

$$H(\theta) = \lambda_A H_A(\theta) + \lambda_{AB} H_{AB}(\theta) + \lambda_{AC} H_{AC}(\theta) + \lambda_{AD} H_{AD}(\theta) + H_A(\theta) + H_B(\theta) + H_C(\theta) + H_D(\theta)$$

Where each letter, A...D corresponds to features in a different region in the image.

$$H_A(\theta) = \sum_{i,j \in n_A} A_{ij}(\theta) \sigma_i^z \sigma_j^z - \sum_{i \in n_A} h_i \sigma_i^x$$

$$H_{AB}(\theta) = \sum_{i \in n_A, j \in n_B} B_{ij}(\theta) \sigma_i^z \sigma_j^z.$$

To ensure that the magnitude of the λ are comparable we can impose a restrains on interaction coefficients:

$$\sum_{i,j \in n_A} |A_{ij}(\theta)| + \sum_{i \in n_A} |h_i| = 1, \quad \sum_{i \in n_A, j \in n_B} |B_{ij}| = 1$$

References


1. Li, Qing, et al. "Progress and opportunities of foundation models in bioinformatics." Briefings in Bioinformatics 25.6 (2024): bbae548.
2. C. Kak, Quantum neural computing, Advances in imaging and electron physics 94, 259 (1995).
3. Pira, Lirandë, and Chris Ferrie. "On the interpretability of quantum neural networks." Quantum Machine Intelligence 6.2 (2024): 52.

Towards large-scale correlated states preparation

➤ Pauline Besserve, Research Associate, pbesserv@ed.ac.uk

Problem


- State preparation refers to the **preparation, on a quantum computer, of a target state**. It is a central step in condensed matter and chemistry applications, where typically one is interested in preparing the ground state of a given electronic structure Hamiltonian.
- State preparation usually relies on **deep quantum circuits**. Part of these circuits is required to generate the target physics, whereas part is a consequence of a **suboptimal representation**.
- **Sparser representations could be discovered by quantum computers**, shifting the expectations over quantum devices from quantitative to qualitative.
- **Natural Orbitals** (NOs), defined from one-body correlators (1-RDM), offer a way to get **compact state representations**. The sparsity gain is particularly striking for **ground states of impurity models**, which play a key role in solving the **Fermi-Hubbard model** via embedding.
- However, these orbitals are **state-dependent** and thus, stricto sensu, require to have knowledge of the target state.



Natural Orbitals are associated with sparse representations. These are easier to handle, whether quantumly or classically.

Solution and current open question

- Natural Orbitalization [1] is an algorithm consisting in **using approximate target state preparation to infer natural orbitals** and thereby enable shallower state preparation. This **iterative method** was observed to be robust to simulated depolarising noise.
- Natural Orbitals were also leveraged by IBM in the context of sample-based Krylov quantum diagonalization (SQKD [2]), which requires a sparse ground state representation.
- At scale, approximate state preparation circuit is likely to be very far from the exact target due to **capped depth**. Also, there is no known exact state preparation circuit for correlated states, calling for **heuristic circuits**.
- The question remains open whether such circuits can still **display relevant one-body information** to enable subsequent **high-accuracy compact state representation** of the target state.




Algorithmic workflow this work focuses on. The motivation of this work is to understand the requirements over approximate state preparation to enable the discovery of compact state representations which can be leveraged by QPUs and/or CPUs.




Experiment on IBM hardware

- Idea: **mimic degraded expressivity** by tweaking the (fixed) **Givens rotations angles** involved in the **known exact circuit** for uncorrelated ground state preparation $\rightarrow$ **controlled experiment**, comparison can be made against noisy execution of the exact circuit.
- Upon **measuring the correlation matrix (RDM)**, one gets an orbital rotation with which **the Hamiltonian is dressed**. Measurement of its expectation value over the **product state** reflects to which extent the enabled compact state preparation is jeopardized. Gained insights apply to **correlated states** as well.



Results

For **uncorrelated 20-qubit impurity model** with physically relevant parameters, randomly tweaking angles within a given magnitude:



We observe that, from an energetic criterion, **approximate Natural Orbitals are robust**. The evolution of the energy error matches that of **perturbation theory**. We also observe (not shown here) **differentiated effects** of altering one or the other type of angles in the Givens rotations [3].

Next steps

- Add quantum error mitigation: PEC, Hamming-weight based post-selection.
- Consider approximate state preparation of correlated states.
- Consider larger system sizes.

References


- [1] Compact fermionic quantum state preparation with a natural-orbitalizing variational quantum eigensolving scheme, P. Besserve, M. Ferrero, T. Ayral, arXiv:2406.14170.
- [2] Quantum-Centric Algorithm for Sample-Based Krylov Diagonalization, J. Yu, J. Robledo-Moreno, J.T. Iosue et al, arXiv:2501.09702.
- [3] Manuscript in preparation.

Adiabatic-inspired hybrid quantum-classical methods for molecular ground state preparation

➤ Sean Thrasher, PhD Student, s.thrasher@sms.ed.ac.uk
Joint work with Petros Wallden, Ioannis Kolotouros, Julien Michel

Problem

- Variational Quantum Algorithms (VQAs) could revolutionise drug discovery by providing a way to compute molecular ground states with current quantum devices.
- VQAs feature exponentially vanishing gradients (barren plateaus) and many local minima: a significant obstacle to unlocking quantum utility for drug discovery.
- Adiabatically-inspired Variational Quantum algorithms (AVQ)s give us a way to warm-start this algorithm, potentially helping us to avoid barren plateaus and false local minima.
- Several AVQ algorithms [1-3]: how do they compare on quantum chemistry problems?
- AVQ algorithms with the mean-field Hamiltonian as the initial Hamiltonian has not yet been explored.
- Finally, the AVQ algorithm only guarantees the angles chosen makes us remain at a stationary point of the energy.



Solution

• Framework

- We view each adiabatically-inspired method as “predictor-corrector methods” from homotopy methods.
- The Davidenko differential equation is utilised.

• Proposed G-AQC-PQC to tackle the following bottlenecks of AQC-PQC:

- Generalizes AQC-PQC to arbitrary adiabatic schedules.
- Reduces the computational cost while maintaining the positive-semi-definite constraint of AQC-PQC.

• Chemistry application


- VQE with Fock initialisation struggles with finding the ground state of BeH2 at large bond lengths, being a good testing ground for AVQ methods.

• Comparison

- Compared the performance of AVQ methods, varying: ansatz depth, initial Hamiltonian, bond-length, number of steps, and choice of ansätze.

$$H_{\text{Fock}} = \sum_{p \in A} \epsilon_p \left(\hat{a}_{p\alpha}^\dagger a_{p\alpha} + \hat{a}_{p\beta}^\dagger a_{p\beta} \right) + K_{\text{C,Fock}}$$

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} \langle pq|rs \rangle a_p^\dagger a_q^\dagger a_s a_r$$



Algorithm 1: Generic Predictor-Corrector Scheme in Our Framework

```
Input:  $\theta \leftarrow \theta^* \in \mathbb{R}^p$ ;  
 $t \leftarrow 0 \in \mathbb{R}$ ;  
 $h > 0$ ;  
while  $t < 1$  do  
  run a prediction subroutine with  $\theta$  as the initial  
  angle on  $H(t)$  to obtain a shift vector  $\epsilon$  such  
  that:  
     $E(\theta + \epsilon, t) \approx E_0(t + h)$ ,  
  where  $E_0(t)$  is the true ground state energy at  
  interpolation value  $t$ ;  
   $t \leftarrow t + h$ ;  
   $\theta' \leftarrow \theta + \epsilon$ ;  
  run a correction subroutine with  $\theta + \epsilon$  as the  
  initial angle on the Hamiltonian  $H(t)$  to obtain  
   $\theta'$  such that  
     $E(\theta', t) \approx E_0(t)$   
end
```






Outcomes

- Unified the AAVQE algorithm, the VAQC method, and AQC-PQC within a single framework, relating it to homotopy methods.
- Proposed a novel algorithm G-AQC-PQC, which improves AQC-PQC.
- Performed an empirical benchmark of adiabatically inspired approaches for chemistry testing the performance on the ground state of BeH₂ (describing using the STO-3G, CAS(4,4) level of theory) at different bond lengths

Future Work

- Perform a comparison of the methods under realistic noise models
- Consider alternative initial Hamiltonians such as the Dyall Hamiltonian



[1] M. Cerezo et al., "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.

[2] A. Garcia-Saez and J. Latorre, "Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers," *arXiv preprint arXiv:1806.02287*, 2018.

[3] S. M. Harwood et al., "Improving the variational quantum eigensolver using variational adiabatic quantum computing," *ACM Transactions on Quantum Computing*, vol. 3, no. 1, pp. 1–20, 2022.

[4] I. Kolotouros, I. Petrongonas, M. Prokop, and P. Wallden, "Simulating adiabatic quantum computing with parameterized quantum circuits," *Quantum Science and Technology*, 2024.

[5] S. Lee et al., "Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry," *Nature communications*, vol. 14, no. 1, p. 1952, 2023.

Beyond Weisfeiler–Leman with quantum graph neural networks




➤ Snehal Raj, PhD student, snehal.raj@lip6.fr
Joint work with Brian Coyle, Leo Monbroussou, Elham Kashefi

Problem

- Graph neural networks power modern chemistry and combinatorial-optimisation pipelines, but the standard message-passing recipe is blind to structural differences it cannot resolve.
- Classical methods that go beyond this limit store features on every k -subset of nodes, at a cost that grows as N^k and becomes infeasible for $k \geq 3$ on real graphs.
- Existing quantum graph models inherit the same blind spot and enforce symmetry by tying many parameters together, restricting what they can learn.
- Deep parametrised quantum circuits also suffer "barren plateaus": gradient signals vanish exponentially with system size, blocking training at scale.

Solution


The architecture is built from quantum gates that conserve particle number; acting on two registers. A node register of N qubits with j excitations sets how the model represents groups of nodes; an embedding register of D qubits with k excitations carries the trainable features. Varying j sets graph-distinguishing power; varying (D, k) sets feature capacity; the two knobs are independent. A single layer applies three operations: a loader that encodes the input graph, an equivariant adjacency step driven by edge weights, and a trainable evolution on the embedding register. A mixing layer then couples both registers while keeping the total excitation count fixed.





Outcome

Our quantum graph neural network delivers four properties together that no prior quantum graph model has combined: it tells apart graphs that standard message passing cannot, going beyond the Weisfeiler–Leman expressivity limit; it is exactly permutation-equivariant, so relabelling input nodes simply relabels the output; it trains reliably at scale, with provable polynomial gradient signals instead of vanishing ones; and it carries trained parameters across the classical-to-quantum boundary, opening a concrete path to quantum utility.



References

- [1] K. Xu et al., How Powerful are Graph Neural Networks? ICLR 2019.
- [2] C. Morris et al., Weisfeiler and Leman go neural. AAAI 2019.
- [3] H. Maron et al., Invariant and equivariant graph networks. ICLR 2019.

Quantum-enhanced Parallel Tempering

➤ Stuart Ferguson, PhD student, S.A.Ferguson-3@sms.ed.ac.uk
Petros Wallden

Problem - Sampling & Optimisation

Combinatorial optimisation is one of the most important types of computational problem that might benefit from quantum computing. It is also related to the harder problem of **sampling**, from a Boltzmann distribution for example:


$$\mu(s) = \frac{1}{Z} e^{-f(s)/T}$$

For certain problem functions, $f(s)$, it is very difficult to sample configurations, s , with appropriate probability $\mu(s)$. Highly parallelised classical methods such as **Parallel Tempering (PT)** exist, where Markov chains explore this distribution at different temperatures, T , at the same time. Here, we ask whether a quantum version of this algorithm is useful and if it can be parallelised.

Solution - Quantum-enhanced Parallel Tempering (QePT)


Classical

The Markov Chain Monte Carlo (MCMC) algorithm samples configurations from a Boltzmann distribution for a particular temperature, T . A random walk (Markov Chain) is performed, directed by the energy of the states. In high T , all s are sampled uniformly but when a chain is colder, only low energy configurations are preferred, and the chains becomes very localised, unable to move. Low T is thus much less efficient. One way to mitigate this is to allow chains of global (high T) chains to communicate and help low- T chains.



Quantum

Our approach relies on a recent quantum algorithm called Quantum-enhanced (Qe) MCMC. It suggests that you can find a quartic scaling advantage in sampling for some problems. This advantage is in reducing the thermalisation time of the algorithm - the Markov chain requires fewer steps. The obvious thing to do is to run PT where each chain is Quantum-enhanced!



Quantum-Classical integration


High T chains are classically trivial for many problems, so there is no point in using quantum there. In fact, there is a **law of diminishing returns in using too many quantum-enhanced markov chains**.

We suggest using quantum-enhancement on only the lowest chains (dotted purple). In this way we have an algorithm that is **highly parallelised between both quantum and classical** devices and we only require the exchange of bitstrings between them.



Outcome

Given this hybrid approach we have to consider how we actually run it in practise. For example, how do we know how long to run our algorithm for? How much quantum resource should we use? We simulate our algorithm for an optimisation problem for a spin glass Ising model. We always have 4 chains running in parallel, but vary the number of these chains that are quantum-enhanced.



Simulations

As the size of the problem (x axis) increases, the total effort (y-axis) required by any algorithm increases exponentially.

Quantum methods scale better.

Majority of advantage gained using only one quantum chain (green diamonds).

Since only a couple of low temperature chains are required to outperform the classical approach, our algorithm is expected to be useful in heterogeneous computing as it relies on close communication with quantum and classical hardware. We can parallelise different computation over multiple QPU, with only classical communication between them. This flexibility is rare in quantum algorithms!

Next up

Real hardware runs - current bottleneck in classical $\leftrightarrow$ quantum communication.

Update QeMCMC to state of the art.

Other algorithms - even more parallelisation?

References

- 1) David Layden et al., "Quantum-enhanced Markov Chain Monte Carlo", Nature 619.7969 (2023), pp. 282–287
- 2) Stuart Ferguson, Petros Wallden, "Methods for non-variational heuristic quantum optimisation". Arxiv:2602.01353

Quantum-enhanced Simulated Annealing

➤ Stuart Ferguson, PhD student, S.A.Ferguson-3@sms.ed.ac.uk
Petros Wallden

Problem - Combinatorial optimisation

Combinatorial optimisation is one of the most important types of computational problem that might benefit from quantum computing. The aim is to find the best “configuration”, s^* , out of all $s \in S$, according to a cost function, $f(s)$.

$$s^* = \operatorname{argmin}_{s \in S} (f(s))$$


We represent all configurations simply by bitstrings, however they may be routes in a pathfinding problem, or even a protein in drug design. Sadly, the quadratic speedup of Grover’s algorithm is unlikely to be fruitful when considering fault tolerant overheads, and variational algorithms suffer difficulties in trainability. **Here, we introduce new quantum optimisation heuristics that are not variational, and do not rely on Grover’s style speedup.**

Solution - Quantum-enhanced Simulated Annealing (QeSA)

Classical

The Markov Chain Monte Carlo (MCMC) algorithm samples configurations from a Boltzmann distribution for a particular temperature, T . In high T , all s are equally likely but at low temperatures, configurations with small $f(s)$ dominate.

One of the traditional classical approaches to optimisation is Simulated Annealing (SA). Basically, you “cool down” your system, and if done slowly enough your Markov chain finds optimal configurations and stays there.




Quantum

Our approach relies on a recent quantum algorithm called Quantum-enhanced (Qe) MCMC [1]. It suggests that you can find a quartic scaling advantage in this sampling procedure for some problems. The question then arises:

Can we use QeMCMC to speed up the comparatively easier problem of optimisation?

Our method is simple, we plug in QeMCMC into SA, to produce QeSA.







Outcome

Success Probability

In simulation, compared to classical SA, QeSA finds the optimal solution with much higher probability after a certain number of iterations. The **classical** approach gets stuck in local minima, which the **quantum** approach does not.



Optimal effort

For each method we find the optimal number of steps to take, balancing it with how many re-runs of the algorithm to do to find the optimal value with high probability. The **quantum method requires less "effort"** than the classical approach, as the size of the problem increases.

Next up

Quantum-enhanced Parallel Tempering - Parallel quantum & classical chains for Q-HPC.

Real hardware runs - current bottleneck in classical <-> quantum communication.

Apply recent QeMCMC improvements.


David Layden et al., "Quantum-enhanced Markov Chain Monte Carlo", Nature 619.7969 (2023), pp. 282–287
Stuart Ferguson, Petros Wallden, "Methods for non-variational heuristic quantum optimisation". Arxiv:2602.01353

Quantum Elastic Network Models and their Application to Graphene

➤ Ioannis Kolotouros, Adithya Sireesh, Stuart Ferguson, Sean Thrasher, Petros Wallden, and Julien Michel

Problem - Graphene

Recently, a quantum algorithm to simulate a system of coupled oscillators was described with the potential for exponential quantum advantage [1].



However, to use this advantage we must map interesting problems For example, graphene, to this simple system and carefully describe efficient quantum circuits for the oracles required by the construction of the algorithm. For many useful problems this is very difficult - or even impossible.

Solution - Connectivity Oracle

The connectivity oracle should find the index, $a(j, \ell)$, which corresponds to the ℓ^{th} neighbour of carbon atom at index j

$$|j\rangle |\ell\rangle \mapsto |j\rangle |a(j, \ell)\rangle$$

The lattice is split into diamond shaped unit cells, each containing two carbon atoms, one from each coloured sub-lattice. The coordinate triplet, (r, c, s) , uniquely identifies an atom.

r : row index


c : column index

s : position in sublattice (bit)

The neighbours of any red (A) node must be a blue (B) node. The three neighbours to any red node are:

1. The other occupant of the unit cell
2. Blue occupant of the upper left cell
3. Blue occupant of the upper right cell

Thus, we can construct an oracle that indicates the neighbours of a node by referencing only the coordinate triplet.



We store the coordinate in a register:

$$|j\rangle_{RCS}$$

Since we need to populate another register with the neighbours, we must calculate the row and column indices of the neighbours.

These “relative shifts”, δ_r, δ_c , are calculable from the table, while they can be used to create neighbour coordinates via quantum arithmetic:

$$|r + \delta_r, c + \delta_c, s \oplus 1\rangle_{R'C'S'}$$

$$|j\rangle_{RCS} |\ell\rangle |0\rangle_{R'C'S'} |0\rangle_F \xrightarrow{CC-U_{\delta}} |j\rangle_{RCS} |\delta_r, \delta_c, 1\rangle_{R'C'S'} |0\rangle_F \xrightarrow{U_{calc}} |j\rangle_{RCS} |r + \delta_r, c + \delta_c, s \oplus 1\rangle_{R'C'S'} |0\rangle_F$$

$|k\rangle$


Control Bits		Relative Shift Vector (δ_r, δ_c)		
r_0	s	$\ell = 0$	$\ell = 1$	$\ell = 2$
0	0	(0, 0)	(-1, 0)	(-1, +1)
0	1	(0, 0)	(+1, 0)	(+1, +1)
1	0	(0, 0)	(-1, -1)	(-1, 0)
1	1	(0, 0)	(+1, -1)	(+1, 0)




Outcome - Applications

Out of plane rippling effect

In nature, graphene does not stay perfectly flat, small perturbations create large oscillations. This is called the "out-of-plane rippling effect."



Heat Transfer



If you heat up one part of a graphene sheet, how does energy disperse through the material?

Doping & Imperfections

In practice, we are much more interested in graphene in the presence of deformities or deliberate doping, not pure graphene. We can model changes to the lattice structure by altering the connectivity matrix, which must be implemented by the oracle. **With QENM, we can simulate both of these phenomena, in the presence of doping, efficiently.** To simulate just 1 cm² of graphene would take more memory than available on the largest supercomputer in the world. Just ~160 fault tolerant qubits are required.

References

- [1] Ryan Babbush, et al. "Exponential quantum speedup in simulating coupled classical oscillators." Physical Review X 13.4 (2023): 041041.
- [2] Ioannis Kolotouros, Adithya Sireesh, Stuart Ferguson, Sean Thrasher, Petros Wallden, and Julien Michel, "Quantum Elastic Network Models and their Application to Graphene". Arxiv:2601.05161


Dynamical Mean Field Theory with Interacting Bath

➤ Weixi Zhang, PhD Student, weixi.zhang@ed.ac.uk

Joint work with Petros Wallden, Ivan Rungger, Yannic Rath

Problem


- Understanding and modelling the electronic structure of strongly correlated materials is essential for high-tech applications.
- A widely used method to understand and model real-world strongly correlated materials is Dynamical Mean Field Theory (DMFT) [1], which uses an effective impurity problem to model the original lattice.




- In standard DMFT, electron-electron interaction is only considered on the impurity site. However, if one wants to use quantum devices as impurity solvers, especially in an analogue setting, it can be hard to turn off completely or make perfectly arbitrary on-site interactions.

Solution

- We use the Maximally Localized Dynamical Embedding (MLDE) approach [2], which modifies DMFT by considering an interacting bath environment for the effective impurity problem.
- It minimally breaks the locality assumption needed for DMFT to be valid, while still giving accurate results compared to classically expensive exact solutions.
- MLDE still requires the important self-consistency relation between the original lattice problem and the impurity problem to be fulfilled. The key differences to standard DMFT are summarised schematically as follows.







Progress

(Right) With MLDE, we can show the Mott transition behaviour akin to what standard DMFT can.

(Bottom) A comparison of the self-energy obtained with MLDE against the one obtained with Continuous Time Quantum Monte Carlo (CTQMC) method, showing good agreement.



Next steps

Systematically benchmark MLDE against standard DMFT to look for advantages such as potential reduction in the number of sites needed.

Develop ways of measuring related quantities to execute MLDE on a quantum device.

References

- [1] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.* **68**, 13 (1996).
- [2] Lupo, C., Jamet, F., Tse, W.H.T. et al. Maximally localized dynamical quantum embedding for solving many-body correlated systems. *Nat Comput Sci* **1**, 410–420 (2021).
- [3] Paper in preparation for this work.

Work In Progress – Exploring Performance and Power of CXL Memory and PCIe 5.0 NVMe for Memory-Hungry Workloads


➤ Maria Bahna, PhD Student, maria.bahna@ed.ac.uk
David Fitzsimmons, Antonio Barbalace

Introduction


The recent convergence of storage and memory enabled by PCIe 5.0 NVMe SSDs and Compute eXpress Link (CXL) 2.0 memory expanders offers promising solutions to memory bottlenecks of memory hungry workloads.

Idea

Exploring how these new storage and memory technologies improve the performance and energy efficiency with an example memory hungry workload: quantum statevector simulations



Setup



Workload

```
initQuESTEnv();
Qureg qureg = createQureg(size);

startReading_power();
startTiming();

foreach qubit in qureg:
    applyHadamard(qubit);

stopTiming();
stopReading_power();

destroyQureg(qureg);
finalizeQuESTEnv();
```

Simulator [1]



- ❖ Relatively fast benchmark compared to other circuits
- ❖ All memory addresses are touched



Costs

SYSTEM	COMPONENTS	
Machine	8 x DDR5 DIMM	
Machine + Disk	4 x SSD	
Machine + CXL	4 x CXL	
Machine + full RAM	16 x DDR5 DIMM	
ARCHER2	4096 nodes	



PCIe 5.0 NVMe SSD

SnuQS Framework [2] (*not open source)
We reimplement the pipeline with QuEST[1] as a backend
vs
OS subsystem swap




CXL 2.0 MEMORY EXPANSION

Compute eXpress Link (CXL) is a new high-speed interconnect for memory expanders [3]
Weighted Interleaving Configurations
vs
Default Memory Allocation

CXL vs PCIe NVMe SSD vs Double RAM

Takeaway
CXL – within 30% when using half of the machine
DISK – within 50% on small number of cores
ISSUE latency



References

- [1] Jones, T., Brown, A., Bush, I. et al. QuEST and High Performance Simulation of Quantum Computers. Sci Rep 9, 10736 (2019). <https://doi.org/10.1038/s41598-019-47174-9>
- [2] Daeyoung Park, Heehoon Kim, Jinpyo Kim, Taehyun Kim, and Jaemin Lee. 2022. SnuQS: scaling quantum circuit simulation using storage devices. In Proceedings of the 36th ACM International Conference on Supercomputing (ICS '22). Association for Computing Machinery, New York, NY, USA, Article 6, 1–13. <https://doi.org/10.1145/3524059.3532375>
- [3] Sharma, D. D., Blankenship, R., & Berger, D. S. (2024). An Introduction to the Compute Express Link (CXL) Interconnect. arXiv [Cs.AR]. Retrieved from <http://arxiv.org/abs/2306.11227>

Towards a Generalisation of Complete Equational Theories*

➤ Nicolas Heurtel, Research Associate
With Chris Heunen, Robin Kaarsgaard, Louis Lemonnier and Kim Worrall

What are Complete Equational Theories?

An equational theory (ET) is a set of relations between terms. In quantum computing, these terms are often diagrams representing quantum gates, making the ET a set of diagrammatic equations between quantum circuits.

An ET is **complete** when every true equation can be derived within the theory. finding a complete equational theory (CET) was a long-standing open problem, only solved recently [1].

Since then, many CETs have been developed, simplified, and even proven to be minimal [2], i.e. when no axiom can be derived from the others.

While some of the works use similar methods, deriving CETs for unitaries are often specific to their chosen gate sets. The goal of this work is to provide a framework generalising such results.

Minimal **Complete** Equational Theories

For Quantum Circuits, up to control and structural equations[3]

$$\begin{array}{ll} \textcircled{2\pi} = \text{ } & (\text{p}2\pi) \qquad \textcircled{\alpha_1} \textcircled{\alpha_2} = \textcircled{\alpha_1 + \alpha_2} \quad (\text{p-p}) \quad \otimes \\ \text{---} \boxed{H} \text{---} \boxed{H} \text{---} = \text{---} & (\text{H}) \qquad \text{---} \times \text{---} = \text{---} \oplus \text{---} \quad (\text{s wap}) \\ \text{---} \boxed{H} \text{---} \boxed{Z(\alpha_1)} \text{---} \boxed{H} \text{---} \boxed{Z(\alpha_2)} \text{---} \boxed{H} \text{---} = \textcircled{\beta_0} \text{---} \boxed{Z(\beta_1)} \text{---} \boxed{H} \text{---} \boxed{Z(\beta_2)} \text{---} \boxed{H} \text{---} \boxed{Z(\beta_3)} \text{---} & (\text{E2}) \end{array}$$

For Matrix Theory [4]

$$\begin{array}{ll} \text{---} \boxed{2\pi} \text{---} = \text{---} & (\text{p}2\pi) \qquad \text{---} \boxed{\varphi_1} \text{---} \boxed{\varphi_1} \text{---} = \text{---} \boxed{\varphi_1 + \varphi_2} \text{---} \quad (\text{p-p}) \quad \oplus \\ \text{---} \times \text{---} = \text{---} \boxed{R_X(\pi)} \text{---} \begin{array}{c} \boxed{\frac{\pi}{2}} \\ \boxed{\frac{\pi}{2}} \end{array} \text{---} & (\text{s wap}) \qquad \text{---} \boxed{R_X(\alpha_1)} \text{---} \begin{array}{c} \boxed{\alpha_0} \\ \boxed{\alpha_2} \end{array} \text{---} \boxed{R_X(\alpha_3)} \text{---} = \text{---} \boxed{R_X(\beta_2)} \text{---} \begin{array}{c} \boxed{\beta_0} \\ \boxed{\beta_3} \end{array} \text{---} \quad (\text{E2}) \\ \text{---} \boxed{R_X(\delta_1)} \text{---} \boxed{R_X(\delta_2)} \text{---} \boxed{R_X(\delta_3)} \text{---} = \text{---} \boxed{R_X(\gamma_1)} \text{---} \boxed{R_X(\gamma_2)} \text{---} \boxed{R_X(\gamma_3)} \text{---} & (\text{E3}) \end{array}$$

Can we derive CETs for unitaries, regardless of the generators?
Why is (E3) not present in CETs for quantum circuits?




*Work in progress. Will be presented at LIQCS26. Current and future work involve arbitrary-dimensional generators.



Generic equations for generic generators

Let's consider a 2-wire block that is universal for 2×2 unitary matrices, up to phases on the right. Every unitary can be uniquely decomposed as a triangle composed of these 2-mode block. This form can be obtained with the oriented versions (from left to right) of (G-E2) and (G-E3). For the matrix theory of [4], this block would consist of one phase and one RX rotation.

$$\begin{array}{|c|} \hline \text{Two vertical bars on two wires} \\ \hline \end{array} = \begin{array}{|c|} \hline \text{One vertical bar on two wires} \\ \hline \end{array} \quad (\text{G-E2}) \quad \begin{array}{|c|} \hline \text{Three vertical bars on two wires} \\ \hline \end{array} = \begin{array}{|c|} \hline \text{Two vertical bars on two wires} \\ \hline \end{array} \quad (\text{G-E3})$$



Derivation of (E3) thanks to structural equations

$$\begin{array}{|c|} \hline \text{Circuit with two A blocks and two B blocks} \\ \hline \end{array} = \begin{array}{|c|} \hline \text{Circuit with two B blocks and two A blocks} \\ \hline \end{array} \quad (1) \quad \begin{array}{|c|} \hline \text{Circuit with A and B blocks} \\ \hline \end{array} \otimes \begin{array}{|c|} \hline \text{Circuit with A and B blocks} \\ \hline \end{array} = \begin{array}{|c|} \hline \text{Circuit with A and B blocks} \\ \hline \end{array} \otimes \begin{array}{|c|} \hline \text{Circuit with A and B blocks} \\ \hline \end{array} \quad (2)$$

Equation (1) is the structural rule that, along with (G-E2) and rules with phases, can derive (G-E3).

It is the direct-sum analogy of the bifunctionality of the Kronecker product, i.e. the Equation (2), which is native for quantum circuits.

Theorem. Consider a set of circuits C generated by arbitrary 2-dimensional generators and phases. If C is universal for unitary matrices, then an ET on C , complete for up to 2-dimensional unitaries and containing (1), is complete for unitaries of 4 dimensions and above.

References

- [1] A Complete Equational Theory for Quantum Circuits. Alexandre Clément, Nicolas Heurtel, Shane Mansfield, Simon Perdrix, Shane Mansfield.
- [2] Minimal Equational Theories for Quantum Circuits. Alexandre Clément, Noé Delorme, Simon Perdrix.
- [3] Diagrammatic Reasoning with Control as a Constructor, Applications to Quantum Circuits. Noé Delorme and Simon Perdrix
- [4] A Complete Graphical Language for Linear Optical Circuits with Finite-Photon-Number Sources and Detectors. Nicolas Heurtel.

QSL Software Stack.

➤ Virag Umathe, Research Assistant, (vumathe@ed.ac.uk)
Brian Coyle, Caitlin Jones, Mario Herrero Gonzalez,
Pauline Besserve, Sean Thrasher, Kenza Hammam

Problem

- With the rapidly evolving quantum computing landscape, it is challenging to ensure researchers can make the most of available quantum hardware. New programming languages, software frameworks, and hardware capabilities emerge continuously.
- There is no single standardised process for executing quantum programs across devices. The lack of a common abstraction layer also means that the code is less reproducible and documentations are irrelevant.
- Quantum hardware needs estimation because resources are restricted and often bounded by credits.
- Beyond the fragmentation of hardware stacks, Quantum research often demands running thousands of experiments sweeping over circuit parameters, noise models, optimisation hyperparameters, and hardware configurations.
- Compounding this, researchers are expected to maintain high-quality documentation.

Solution

What if there existed a single, unified flow that works seamlessly for both hardware-specific and agnostic projects, without requiring researchers to worry. We propose our very own software stack that sits between the researcher and the hardware, providing a consistent support for the tedious tasks. The stack composes

- QSL Backend,
- IBM resource estimator,
- Experimental Version Control System (EVCS)
- Auto documentation.

QSL Backend:

We are developing the QSL Backend, a testbed specific software layer that provides a unified interface for submitting and executing quantum programs across devices available through NQCC's testbed and the SPARQ program.

Resource Estimator


Estimates QPU runtime for six quantum algorithm classes across small, moderate, and large problem sizes on IBM Heron 156-qubit hardware, using circuit submissions as the atomic scaling unit.

EVCS:

Automatically captures, hashes, and logs all experiment hyperparameters, configurations, and results at runtime – ensuring every experiment is consistently tracked and reproducible without any manual overhead.

Auto Documentation

Researchers comment their code – the system generates structured documentation automatically, keeping all codebases consistently documented with minimal effort.






Outcome

The QSL Backend is implemented in Python and currently supports two hardware backends, Infleqtion's Superstaq and Qiskit. These two platforms were chosen as the starting point given their relevance to the current research activities.

The backend provides a single, consistent submission process through which researchers can submit quantum programs to backends without modifying their core implementation of the original frameworks.



The intention is not to prescribe which hardware platforms will be supported, but rather to grow the backend organically in response to the evolving needs of the researchers using it. As new devices become relevant to the work at QSL, support for them will be integrated accordingly.

The resource estimation checks the circuit count, applying IBM's formula $\text{QPU time} = 2 + (\text{device factor} \times \text{num_executions})$ to produce zero cost pre-execution runtime estimates. Validation on IBM Heron 156-qubit hardware yielded a predict-

-ed runtime of 63.4 minutes against an actual execution of 62.5 minutes, a deviation of under 1.5%. The estimator also accounts for per-circuit repetition delay ($\text{rep_delay} = 250\mu\text{s}$ default on IBM backends), incorporating it into the total execution time calculation alongside circuit length and shot count.

The EVCS is built around a JSON based logging system that automatically captures all experimental hyperparameters, configurations,

and results at the time of execution. Every experiment run through the system produces a structured log. Beyond logging, EVCS also incorporates a lightweight code restoration mechanism.

For documentation, Sphinx parser that traverses the codebase and automatically generates structured, up to date documentation directly from the source code. Rather than documentation being a task that researchers return to at the end of a project.

Consulting the Unwritten Manual:

Uncovering tacit quantum computing knowledge with AI surrogates

➤ William Waites, Senior Researcher, wwaites@ed.ac.uk
Vincent Danos, Pauline Besserve, Brian Coyle, Marine Demarty,
Kenza Hammam, Mario Herrero Gonzalez, Caitlin
Jones, Léo Monbroussou, Virag Umathe, Bo Yang, Elham Kashefi

Surfacing Tacit Knowledge

The quantum computing landscape is large and varied. Hardware modalities, encoding and manipulation of qubits (where the concept of qubit makes sense), handling of errors and noise, compilers whose output is not guaranteed to be identical across invocations, the absence of standardised machine-readable encoding of algorithms, all of these conspire to make the seemingly straightforward question of how best to run a given algorithm rather difficult to answer.

This information is not all readily available. Though much is published in papers and vendor documentation, crucial information is exchanged informally between postdocs and graduate students, not easily discoverable and rarely written down.


This project gathers this tacit information through interviews with practitioners and encodes it in a form where it can be used by AI agents for exploration, discovery and planning.

Literature Scanning and Curation


The knowledge base is grounded in a literature curation process. An interesting source is a regularly scheduled search of the arXiv preprint server to identify new publications that pertain to an application of quantum computing to some domain.

Hardware modality × application domain	Super-conducting	Trapped ions	Neutral atoms	Photonic	Annealing	Q-Inspired	NV center	Silicon spin	Cloud	Total
Platforms & Tooling	5	3	2	3	1			1	1	16
Quantum Machine Learning	4	5	3	2	1					15
Drug Discovery & Pharma	7	3	1	1	1	1				14
Optimization	4	3	3	1	2					13
Chemistry & Materials	7	3	1	1						12
Simulation & Physics	3	4		2						9
Logistics & Scheduling	1	3	1		3					8
Crypto & Security	4	2	1	1						8
Finance	3	1			1	1				6
Error Correction	2		3							5
Energy & Climate		1	2							3
Networking & Edge		1		1			1			3
Boson Sampling				1						1

The literature is triaged and evaluated by AI agents under human supervision and then classified and analysed in various ways and can be viewed by vendor, by modality, by specific application or domain or algorithm class. As we build up the corpus of reference material, we can produce an increasingly nuanced account of the relationship between applications of quantum computing and the machines that implement them.



Debate with AI Surrogates




An example debate topology. To obtain candidate answers to questions such as “which hardware might be best suited to run this particular algorithm?” We construct a network of AI agents. Each agent has access to a general knowledge database about quantum computing and may run on a specially trained model (see below). Each may also have access to specialist or confidential information according to its specific role.

The question is posed to vendor surrogates who try to make the case for suitability of their particular hardware. The original question, and cases are then put to an agent trained on the informal knowledge in QSL as elicited in interviews, a sceptical agent that believes nothing and a jester agent who tests assumptions and the foundations of the question itself. All of these opinions are then presented to a judge (or panel of judges) who decide if there is consensus. If there is, the answer is produced. If there is not, the process begins again.

To experiment with different debate topologies, and other applications of AI agents, we developed a concurrent language based on concepts from category theory with type checking and a verified compiler.

Fine tuning. Because we work with confidential information that ought not be shared with third parties including cloud AI providers – we are developing specialised language models that can be run locally with relatively modest resources. The first layer of specialisation is training in knowledge about quantum computing from the published literature including 99 papers from QSL. Further layers are trained based on transcripts of interviews with researchers at QSL where we attempt to capture the rich informal knowledge that is not normally formally written down. Future layer adapters might consist of vendor materials.




Selective Blindness: A new framework for Blind Quantum Computation

➤ Abbas Poshtvan, PhD Student, A.Poshtvan@ed.ac.uk

Joint work with Sasha Lapiha, Mina Doosti, Dominik Leichtle, Luka Music, Elham Kashefi

Problem

- Goal: Secure delegation of quantum computation; Enable a server to perform a computation without revealing the computation itself.
- Achievable with quantum communication which has a linear cost in the size of the computation [1].
- Can we reduce the communication overhead? Do we always need full masking?
- UBQC delegates one unitary from the all possible computations.
- All the MBQC graph on the server side is encrypted.



Solution



Our Solution [2]:

Change the functionality: Delegation of one unitary from a public set of unitaries (which represents the few possibilities remained after partial specification of the problem): $\{U_1, U_2, \dots, U_n\}$

Key observation: The only nodes which need quantum encryption are the ones that have different measurement angles in the set of unitaries (target nodes) + Nodes with non-clifford angles (corresponded to magic gates) which get X-correction from target nodes.


Example: One of two unitaries with one different measurement angle (the red node in the figure)

This approach pushes the hiding to its limit using the rules of MBQC.





Other attempts:

Saving on the architecture level [3]: Blind operations are performed through photonic interactions, The rest (including fault tolerant supporting operations) are performed on the neutral atom qubits.



Partial Blindness [4]: UBQC only on the privacy sensitive parts and using Quantum homomorphic encryption on the public parts of the computation.






Outcome and future directions

A new protocol that identifies the minimal and essential hiding (qubit communication) required to delegate a quantum computation from a public set taking into account all the possible differences between the possible computations.


This approach can bring saving on quantum communication in cases that only some sensitive parts of the computation need to be kept hidden, such as Federated machine learning and Grover search algorithm.

Cryptographic compilations: This protocol motivates new ways of compiling quantum circuits (or their corresponding MBQC patterns) such that:

1. T gates (equivalently non-clifford angles) are pushed to the front of the circuit
2. Differences between the nodes measurement angles of the possible computations are pushed to the end of the measurement pattern. An example:



Blind swap: The hiding approach can be used as a subroutine to develop methods for the cases that the difference between the possible computations is significant. In such cases, performing both computations and applying a swap in a blind way can be a better solution.



References

- [1] A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal blind quantum computation," in 2009 50th annual IEEE symposium on foundations of computer science, pp. 517–526, IEEE, 2009.
- [2] A. Poshtvan, O. Lapiha, M. Doosti, D. Leichtle, L. Music, and E. Kashefi, "Selectively blind quantum computation," arXiv preprint arXiv:2504.17612, 2025.
- [3] Gefen Baranes, Iria W. Wang, Francisco Machado, Aziza Suleymanzade, et. al. "Designing Fault tolerant blind quantum computation" arXiv preprint: 2505.21621, 2025
- [4] Youngkyung Lee, Doyoung Chung, "Partial blind quantum computation: a framework for selective circuit protection", Quantum Information Processing, 2025

Eavesdropper-Blind Remote State Preparation & Applications to Quantum Public-Key Encryption

➤ Kaniuar Bacho, PhD Student, k.bacho@sms.ed.ac.uk
Joint work with Alexandru Cojocaru, University of Edinburgh

Problem

Motivation:

- Remote state preparation (RSP) lets a classical verifier prepare quantum states on a quantum prover using only classical communication. It is the backbone of many quantum cryptographic protocols.
- Known RSP constructions typically rely on strong trapdoor assumptions, especially trapdoor claw-free functions.
- The usual blindness definition against a malicious prover may be stronger than what some applications need.

Central Question:


- Can a weaker, eavesdropper-security notion still enable useful cryptography?
- Can an RSP-type primitive be built from trapdoor-free assumptions?

Solution


Eavesdropper-Blind Remote State Preparation:

- Prepares root of unity states: $\left\{ |+\theta\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle) \mid \theta \in \frac{2\pi}{N}\mathbb{Z}_N \right\}$
- An eavesdropper given the honest transcript can't distinguish from a uniformly random angle.

(X,Y)-Plane Bloch Sphere (N = 8)



Two-Message Version:



Cryptographic Assumption:


- Specific One-Way Group Actions (OWGAs) or Trapdoor Claw-Free Functions (TCFs):

$$\begin{cases} V \text{ samples } x_0 \leftarrow X \text{ and } s, r \leftarrow G. \\ \text{msg}_V = (x_0, s \star x_0, r), \text{msg}_P = (z, g) \in \mathbb{Z}_p \times G \\ \theta = \langle s, z \cdot r + g \rangle \cdot \frac{2\pi}{p} \end{cases}$$


Application:

- Quantum Public-Key Encryption (QPKE) with classical public keys and quantum ciphertexts:

$$\begin{cases} \text{pk} = (x_0, s \star x_0, r), \text{sk} = s \\ \text{Enc}(\text{pk}, b \in \{0, 1\}) \approx \langle Z^b | +\theta \rangle, z, g \end{cases}$$



Outcome



References


- [1] Gilles Brassard and Moti Yung. One-way group actions. CRYPTO 1990.
- [2] Kaniuar Bacho, Alexander Kulpe, Giulio Malavolta, Simon Schmidt, and Michael Walter. Compiled nonlocal games from any trapdoor claw-free function. CRYPTO 2025.
- [3] James Bartusek and Dakshita Khurana. On the power of oblivious state preparation. CRYPTO 2025.
- [4] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. From the hardness of detecting superpositions to cryptography: Quantum public key encryption and commitments. EUROCRYPT 2023.

Authentication for Quantum Networks

➤ Christopher Battarbee, QSL, kit.battarbee@ed.ac.uk
with Myrto Arapinis, Mina Doosti, and Suchetana Goswami

Quantum cryptography requires classical authentication.

Quantum cryptography reduces to classical communication about outcomes obtained from carefully structured quantum interactions. Security guarantees are void if this classical communication is not properly secured. How can we ensure classical authentication?



Stronger security guarantees come at a cost.

Unconditional security is the strongest possible guarantee; forgeries occur with bounded probability...

We have secure constructions **assuming one-way functions (OWFs)** exist...

We have secure (and possibly quantum-resistant) constructions **assuming quantum-resistant trapdoor OWFs** exist...

more secure



more practical


...but is only possible if **single-use pre-shared keys** are pre-distributed

...and a **multi-use pre-shared key** suffices

...and these have no pre-shared key requirement. However, we need **one-time authentication** to ensure public keys are correct

The future quantum internet needs to be scalable.

Set-up requirements are important when many users may join a network. Public-key cryptography has a special property – **the cost to join a network is independent of the size of that network**, provided a Public Key Infrastructure (PKI) is available. Moreover, even if full unconditional security is lost, we obtain long-term **everlasting security**.



Partial blindness for sub-universal photonic quantum computing schemes

- Verena Yacoub, PostDoc, verena.yacoub@lip6.fr
 Suchetana Goswami, Senior Researcher, sgoswami@ed.ac.uk
 Joint work with Mina Doosti and Elham Kashefi

Problem

- **Objective:** Establishment of *quantum cloud computing* (QCC) that preserves privacy when processing sensitive information (such as, medical records etc.).
- This challenge can be addressed by the introduction of quantum internet, where the security is ascertained by **blind quantum computing** -> operational structuration!
- Our motivation: To obtain operational *blindness for the near-term photonic devices*.
- The photonic devices are generally deprived of quantum entanglement as a resource that is essential for blindness -> sub-universal -> hard to achieve full blindness -> **introduction of the notion of partial blindness!**
- Our interest is to explore the photonic setups and the metrics to quantify the amount of blindness and association of suitable quantification with particular photonic setup.
- Considering the photonic systems as *bosonic systems* instead of qubit encoding.
- We explore the fundamental resources required in these systems to achieve blindness and correctness in photonic systems.

Solution

Blind quantum computing with qubits

- Client has limited computational power -> wants the server to perform desired computation in a blind way.
- Client prepares initial quantum states and sends them to the server in a hidden way. Hiding qubit states:

$$|\psi_j\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta_j}|1\rangle)$$

$$\text{where, } \theta_j \in \{0, \frac{\pi}{4}, \dots, \frac{7\pi}{4}\}$$

- As the above angle is chosen uniformly at random by the client, the server essentially receives a maximally mixed state.
- Client also encodes the computation (measurement) to be performed by the server by encoding with one time padding:


$$\delta_j = \phi_j + \theta_j + r_j\pi$$

- The server receives the encrypted angle to perform the measurement and returns classical outcomes.
- To decrypt the computation, client classically apply XOR to the obtained outcome with the corresponding random bit $r_j \in \{0, 1\}$

Hiding the state in the photonic setting

To achieve hiding for the input state while the unitary is known: **Quantum Fully Homomorphic Encryption** [1, 2].

- Client uses random coherent displacements in optical phase space to encrypt the input state.
- Note that, the encryption has to commute with the linear optics evolution (unitary) for efficiently decryption after computation by the server.



Drawback: The unitary is not hidden!

Hiding the computation

- Let the computation is given by an unitary operation.
- To hide -> To multiply the unitary classically with a Haar random unitary.
- To decrypt the result of the computation -> the client performs a heterodyne measurement that allows the client to recover the desired state [3].

Drawback: Is the delegated computation hard enough?




Outlook

The **goals** are the following:

- Successful initial state hiding with minimal computational resource on Bosonic systems: FHE or by randomisation of creation operation on a multi-mode state.

$$|\psi\rangle = \otimes_{j=1}^m A_j(i)|0\rangle^{\otimes m} \text{ With, } A(i) = \mathbb{I} \text{ for } i = 0, \text{ and } A(i) = \hat{a}^\dagger \text{ for } i = 1 \quad \forall i \in \{0, 1\}$$

- Next, we want to integrate the previous frameworks into the full framework of delegation either using minimal quantum powers that permit hiding on the client's side or by inserting a Hardware secure module in the malicious server that does the hiding task and the client only communicates classically as in [4].
- Another direction: can adding nonlinear gates (e.g. Kerr gates) while not allowing full graph state entanglement help with the hiding property of the initial photonic system?



References


- [1] Ouyang, Yingkai, et al. "Homomorphic encryption of linear optics quantum computation on almost arbitrary states of light with asymptotically perfect security." Physical Review Research 2.1 (2020): 013332.
- [2] Rohde, Peter P., Joseph F. Fitzsimons, and Alexei Gilchrist. "Quantum walks with encrypted data." Physical review letters 109.15 (2012): 150501.
- [3] Upreti, Varun, and Ulysse Chabaud. "An efficient quantum state verification framework and its application to bosonic systems." arXiv preprint arXiv:2411.04688 (2024).
- [4] Ma, Yao, et al. "QEnclave-A practical solution for secure quantum cloud computing." npj Quantum Information 8.1 (2022): 128.

Verifiable Blind Probabilistic Error Cancellation

➤ Bo Yang, PhD candidate, Bo.Yang@lip6.fr
Elham Kashefi, Harold Ollivier

Problem


- Quantum error mitigation (QEM) is essential for suppressing hardware errors with moderate quantum overhead. [1]



- Its reliability often relies on modelling, calibration, and hardware-specific assumptions.

e.g.) Probabilistic error cancellation (PEC) [2]

- apply inverse of the noise via quantum-classical hybrid procedure




- Estimation bias relies heavily on
 - accuracy of the noise model
 - correctness of error cancellation gates

End-to-end reliability of QEM is unresolved

Solution

Cryptographic protocols for “Verifiable Blind Quantum Computation (VBQC)” [3] offer a principled route to end-to-end reliability.



For the Client with the ability to prepare and send single qubits, VBQC protocols provide


- perfect blindness of computation
- verifiability with exponentially small soundness error probability

However...

- Existing verification protocols are not actively robust to noise.
- QEM has remained outside cryptographic framework.


Fundamental gap between verification protocols and noise-robustness


We formalise the secure cryptographic resource for the ideal PEC procedure as “**Secure Delegated Probabilistic Error Cancellation (SDPEC)**” as below:



The proposed “**Verifiable Blind Probabilistic Error Cancellation (VBPEC)**” protocol below securely constructs SDPEC resource.

Test rounds: compute $\tilde{\Delta} = \|\tilde{p}_{PEC} - \tilde{p}_T\|_1$. If $\tilde{\Delta} \leq \epsilon_t$, Acc. Otherwise Abort.





Outcome: The first cryptographically secure QEM

	RVBQC [4]	VBOE [6]	VBPEC (this work)
Computation	(randomised) discrete output	continuous-valued estimator	continuous-valued estimator
Test	count of failed deterministic rounds	count of failed deterministic rounds	continuous-valued estimator
noise-robustness	inherent robustness only	inherent robustness only	inherent robustness, active error suppression


Our results

- Exponential composable security of VBPEC in the abstract cryptography framework
 - perfect blindness
 - verifiability with negligible error
- Active noise robustness of VBPEC by QEM
 - honest noise - mitigate with sampling cost
 - noise mismatch - inherently robust

Technical contributions

- Making PEC-based QEM blindly executable
- Overhead-free MBQC implementation of PEC
- Further generalising trap constructions in [5] through classical post-processing
- New proof technique for non-deterministic, continuous-valued test/computation rounds

Reliable end-to-end PEC execution pipeline
Powerful tool for recent interest in verifiable quantum advantage
The first formulation of QEM in the cryptographic framework
Flexible design and proof for a wider class of verification protocols
Enhancing practicality of cryptographic verification protocols in more realistic setups



References

- [1] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, *Reviews of Modern Physics* 95, 045005 (2023).
- [2] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Physical review letters* 119, 180509 (2017).
- [3] J. F. Fitzsimons and E. Kashefi, Unconditionally verifiable blind quantum computation, *Physical Review A* 96, 012303 (2017).
- [4] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, Verifying bqp computations on noisy devices with minimal overhead, *PRX Quantum* 2, 040302 (2021).
- [5] T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, Unifying quantum verification and error-detection: theory and tools for optimisations, *Quantum Science and Technology* 9, 035036 (2024).
- [6] B. Yang, E. Kashefi, and H. Ollivier, Verifiable blind observable estimation, *arXiv preprint*, arXiv:2510.08548 (2025).

Plugging Leaks in Fault-Tolerant Quantum Cryptography

➤ Dominik Leichtle, Senior Researcher, dominik.leichtle@ed.ac.uk
Joint work with Theodoros Kapourniotis, Luka Music, Harold Ollivier

The robustness problem


- Security analyses of most cryptographic protocols assume that honest parties behave perfectly
- However, quantum devices are naturally imperfect, this exposes the protocol to side-channel attacks

How do protocols behave if an honest party uses a real (imperfect) device?

We are interested in protocols with the following properties:

1. **Robust correctness:** the protocol performs correctly if all parties are behaving honestly, *but are potentially noisy**
2. **Robust security:** honest, but noisy* parties obtain security guarantees

* up to a reasonable amount



Plugging leaks using coherent privacy amplification

Faults, as a model for **secret-independent noise**, are a problem for correctness, but not for security




Figure: Remote State Preparation with secret-independent noise (is equivalent to normal RSP)

Leaks, as a model for **secret-dependent noise**, are a problem for security






Figure: Leaky Remote State Preparation

Coherent privacy amplification for single-qubit states




Step 1: entanglement



Step 2: measurement

Exponential leak suppression for output state

$$\theta = \sum_j (-1)^{t_j} \theta_j$$



Fault-and-leak-tolerant quantum protocols




Figure: simultaneous concatenation of error-correcting codes and leak-distillation

Theorem 1 (Threshold theorem) There exists a universal threshold $p_0 > 0$ such that for all circuits C with L locations, and C^k its level- k simulation with p_c -imperfect operations, the diamond distance δ between C^k and an ideal computation performing C on the logical space and producing a secret-independent syndrome is upper-bounded by

$$\delta \leq 2Lp_0 \left(\frac{p_c}{p_0} \right)^{2^k}.$$

Take-home message: Scalable fault-and-leak-tolerant quantum computation is possible

References

- [1] Dorit Aharonov, Michael Ben-Or, Elad Eban, Urmila Mahadev: Interactive Proofs for Quantum Computations (2017), arXiv:1704.04487.
- [2] Theodoros Kapourniotis, Dominik Leichtle, Luka Music, Harold Ollivier: Plugging Leaks in Fault-Tolerant Quantum Computation and Verification (2025), arXiv:2510.03227.

Logical accreditation: a framework for efficient certification of fault-tolerant computations

➤ James Mills, Adithya Sireesh, Dominik Leichtle, Joschka Roffe and Elham Kashefi

Problem Setting

With rapid progress in both quantum hardware and error correction, quantum computing is transitioning from the NISQ to the **early fault-tolerance era**¹.

Soon we will be able to run encoded computations that are neither classically simulable nor classically verifiable. E.g. problems in simulation and sampling.

Key challenge: How do we trust a quantum computer's output? Need an efficient method to certify correctness of logical quantum computations. This is especially necessary for problems that are **not classically verifiable**!



Logical Accreditation

We propose a new method for certifying logical (error corrected) computations, that fits within the accreditation framework:

1. Run a target circuit alongside multiple trap circuits on a fault-tolerant device.
2. Each trap is constructed to be affected by the same logical noise as the target.

Each logical trap circuit is designed such that it has:

- A known deterministic output in the presence of no logical errors.
- A non-zero bit string output with high probability in the presence of logical errors.




Fig.1 The fault-tolerant quantum computing stack including the logical accreditation layer.




Fig.2 Numerical simulation results of logical accreditation applied to IQP² circuits.



Main Result

Theorem (informal): The logical accreditation protocol provides an upper-bound on the total variational distance (TVD) between the experimental target circuit output distribution and the ideal output distribution of the form:

$$\frac{1}{2} \sum_{s \in \{0,1\}^n} |p_{\text{exp}}(s) - p_{\text{ideal}}(s)| \leq \gamma$$

The upper bound is experimentally estimated from the measured outputs of the logical trap circuits, with a statistical error that may be bounded using Hoeffding's concentration inequality.

Apart from the TVD, the protocol's output can be used to bound the total circuit error rate, second-order Renyi entropy and state infidelity.




Fig.3 Logical accreditation is performed with a quantum device and a classical processing unit (CPU).

References

¹Katabarwa, A., Gratsea, K., Caesura, A., & Johnson, P. D. (2024). Early fault-tolerant quantum computing. PRX quantum, 5(2), 020101.

²Shepherd, D., & Bremner, M. J. (2008). Instantaneous quantum computation. arXiv preprint arXiv:0809.0847.

Error-mitigation aware benchmarking strategy for quantum optimization problems

➤ Marine Demarty, Research assistant, marine.demarty@ed.ac.uk
Bo Yang, Kenza Hammam, and Pauline Besserve

Problem

- To assess the potential of quantum computers at achieving quantum advantage for real-world utility tasks, we need **principled and reliable application-level benchmarks**.
- Existing methods for quantum optimization [SFGP21, BGP25] do not fully apply to realistic quantum experiments, characterised by **finite-shot constraints** and the use of **quantum error mitigation** (QEM).
- No quantitative task-level criterion to determine **whether QEM is useful** within a **fixed shot budget**.

Can we develop a **task-level benchmarking method** that

- takes into account finite-shot statistics and quantum error mitigation;
- determines whether the deployment of QEM is advantageous under a fixed shot budget?

→ Focus on **probabilistic error cancellation** (PEC), a commonly used QEM method


Solution

Problem considered:

- $E_0 = \min_p \text{Tr}[\rho H]$
- Assume known lower and upper bounds E_- and E_+

Benchmarked quantum approach:

- Assume noiseless circuit yields E_0
- Estimation method: Pauli averaging ($\|H\|_2$: Pauli 2-norm) Estimate
- distribution: assumed Gaussian (mean E , variance σ^2)




PROBABILITY OF QUANTUM ADVANTAGE

$$P_{\text{success}} = \frac{1}{2} \left[\text{erf} \left(\frac{E_+ - E}{\sigma\sqrt{2}} \right) - \text{erf} \left(\frac{E_- - E}{\sigma\sqrt{2}} \right) \right]$$

Consider a D-layer quantum circuit under layer-wise global depolarizing noise with probability P

PEC	raw
Mean: $E = E_0 = (E_- + E_+)/2$ as E is unknown	Mean: $E = E_{\text{noisy}} = (1-P)^D (E_- + E_+)/2 + ((1-(1-P)^D)/d) \text{Tr}[H]$
Variance: $\sigma^2 = \beta (\ H\ _2 \gamma_{\text{tot}})^2 / N_{\text{shots}}$ [TSY23]	With dimension $d=2^n$ (n : number of qubits)
where $\beta=1$ for depolarizing and $\gamma_{\text{tot}}=f(n,D,P)$ [T21]	Variance: $\sigma^2 = (\ H\ _2) / N_{\text{shots}}$




Outcome

EXAMPLE QUANTUM OPTIMIZATION TASK

Estimating the ground state energy of the **2D Fermi-Hubbard Hamiltonian**

- Foundational model for strongly correlated fermions on a lattice
- Thoroughly studied → classical results to benchmark against
- Intractable for large system sizes → quantum advantage candidate



PARAMETERS

Square 8x8 lattice
($L=64$ sites)

$(U, t, \mu) = (8, 1, 3.75)$

Known bounds [VHdA91]

Quantum algorithm:
Hamiltonian Variational
Ansatz [WHT15]
($D=L$; $n=2L$)

Successful strategy

Highest P_{success}
and
 $P_{\text{success}} \geq 0.95$

APPLICATION: for a state-of-the-art quantum computer with all-to-all connectivity, $P \sim 4 \cdot 10^{-2}$, which is just at the successful quantum advantage boundary for large $N \sim 10^6$. A one-order-of-magnitude improvement $P \sim 4 \cdot 10^{-3}$ allows quantum advantage with a reasonable shot count $N_{\text{shots}} \sim 10^3$.

→ Given access to the **calibration data** of a quantum computer (gate fidelities) and a maximum **shot budget**, one can easily gain insight into (i) whether **quantum advantage is achievable** with high probability and, if so, (ii) whether to **deploy PEC**.

References

- [SFGP21] Stilck França, D., & García-Patrón, R. (2021). Limitations of optimization algorithms on noisy quantum devices. *Nature Physics*, 17(11), 1221-1222.
- [BGP25] Besserve, P., & García-Patrón, R. (2025). Probing quantum advantage for solving the Fermi-Hubbard model with entropy benchmarking. *arXiv preprint arXiv:2510.00930*.
- [T21] Takagi, R. (2021). Optimal resource cost for error mitigation. *Physical Review Research*, 3(3), 033178.
- [TSY23] Tsubouchi, K., Sagawa, T., & Yoshioka, N. (2023). Universal cost bound of quantum error mitigation based on quantum estimation theory. *Physical review letters*, 131(21), 210601.
- [VHdA91] Valenti, R., Hirschfeld, P. J., & d'Auriac, J. A. (1991). Rigorous lower bounds on the ground-state energy of correlated Fermi systems. *Physical Review B*, 44(8), 3995.
- [WHT15] Wecker, D., Hastings, M. B., & Troyer, M. (2015). Progress towards practical quantum variational algorithms. *Physical Review A*, 92(4), 042303.

Colour Codes Reach Surface-Code Performance with Vibe Decoding

➤ Stergios Koutsoumpas, Postdoctoral researcher, skoutsio@ed.ac.uk
Tamas Noszko, Hasan Sayginel, Mark Webster, Joschka Roffe

Problem

- Colour codes promise efficient logical computation: transversal Clifford gates and low-overhead magic-state cultivation make them appealing alternatives to surface codes.
- In practice they have lagged surface codes: existing decoders fail to match surface-code qubit footprints under realistic circuit-level noise.
- Standard belief propagation (BP) gets trapped on circuit-level Tanner graphs full of short cycles, leading to split beliefs and non-convergence.
- Matching-based decoders force the colour-code detector graph into a structure it does not naturally have, sacrificing accuracy.

Solution

VibeLSD: ensemble BP → LSD.

1. Ensemble of serial BPs


L parallel BP decoders, each with a random serial schedule π_i over error mechanisms. Asynchronous updates break the symmetries and short-cycle split-beliefs that defeat a single BP on circuit-level Tanner graphs.

2. Early stop & rank

Once M decoders converge the rest terminate; the most-likely correction is selected from the survivors using the detector-error-model prior.

3. LSD fallback

If none converge, normalised LLR vectors are averaged and passed to a localised statistics decoder (sparse-matrix cluster inversion).




Configurability

Offline setup is the generation of L random permutations – no training, no per-circuit hyperparameter tuning, no automorphism analysis.

Complexity & speed

Worst-case $O(n^3)$; sub-linear in practice. Min-sum BP and union-find LSD both have mature FPGA/ASIC mappings; GPU prototype uses one warp lane per ensemble member.



Outcome

Colour code footprint comparable to surface codes

VibeLSD is the first colour-code decoder to match



surface-code qubit footprints under realistic circuit-level noise.

Robust across regimes

Works across syndrome-extraction circuits, noise models, and physical error rates. Including experimental data from Lacroix et al. (Nature 645, 2025).

GPU throughput

~900× faster than the Tesseract decoder at $p = 10^{-3}$, with competitive accuracy; microsecond per round on superconducting-relevant distances.



References


- [1] S. Koutsoumpas, T. Nolzko, H. Sayginel, M. Webster, J. Roffe. Colour Codes Reach Surface Code Performance using Vibe Decoding. arXiv:2508.15743 (2025).
- [2] T. Hillmann et al. Localised statistics decoding. arXiv:2406.18655.
- [3] C. Gidney & C. Jones. Chromobius decoder for the colour code. arXiv:2312.08813.
- [4] O. Higgott & C. Gidney. PyMatching v2. Quantum 9, 1600 (2025).
- [5] G. Lacroix et al. Scaling and logic in the colour code on a superconducting quantum processor. Nature 645, 614–619 (2025).
- [6] L. R. Beni et al. Tesseract: a search-based decoder for QEC. arXiv:2503.10988.

Non-Markovian noise limits for sustaining entanglement in multiparty quantum states

➤ Suchetana Goswami, Senior Researcher, sgoswami@ed.ac.uk
Joint work with Ujjwal Sen

Problem

- Multipartite entanglement provides quantum advantages as a resource in many distributed network scenarios.
- Practical implementation involves interaction with noisy environment.
- If multi-qubit GHZ and W states are considered as resources, how do the noises affect the dynamics with time when distributing them across a network?
- Can information-backflow from environment to the system, help sustain the initial resource – Introduction of non-Markovian evolution of multipartite entangled system.



Solution

Initial quantum states

- N-qubit GHZ state:

$$|\psi\rangle_{GHZ}^N = \frac{1}{\sqrt{2}}(|0\rangle^{\otimes N} + |1\rangle^{\otimes N})$$

- N-qubit W state:

$$|\psi\rangle_W^N = \frac{1}{\sqrt{N}}(|0\rangle^{\otimes(N-1)}|1\rangle + |0\rangle^{\otimes(N-2)}|1\rangle|0\rangle + \dots + |0\rangle|1\rangle|0\rangle^{\otimes(N-2)} + |1\rangle|0\rangle^{\otimes(N-1)})$$

Proposition: For any N-qubit system, the GHZ and the W states are symmetric across all possible unique bipartitions.

Modelling of the environmental interaction

For a single qubit, the evolution of the state with time is governed by the time-local master equation,

$$\frac{d\rho}{d(\omega_0 t)} = \dot{\rho}(t) = -\frac{i}{\hbar\omega_0}[H(t), \rho(t)] + \frac{1}{\omega_0} \sum_k \gamma_k(t) \times \left(L_k(t)\rho(t)L_k^\dagger(t) - \frac{1}{2}\{L_k^\dagger(t)L_k(t), \rho(t)\} \right)$$

$H(t)$ is the Hamiltonian of the system, $L(t)$ are the Lindblad operators. Rate of dissipation: $\gamma_k(t)$

- The time dependency of the decay rate $\rightarrow$ memory effect in the system dynamics $\rightarrow$ induction of non-Markovianity in the system.

Noise models

- Dephasing noise:

$$\frac{H_D}{\hbar} = \omega_0 \sigma_z + \sum_i \omega_i a_i^\dagger a_i + \sum_i \sigma_z (g_i a_i + g_i^* a_i^\dagger)$$

$$\dot{\rho}(t) = \frac{\gamma(t)}{\omega_0} (\sigma_z \rho(t) \sigma_z - \rho(t))$$

$$\gamma_{T=0}(t, s) = \omega_c [1 + (\omega_c t)^2]^{-\frac{s}{2}} \Gamma[s] \sin(s \arctan(\omega_c t))$$


- Depolarising noise:

$$H_{DP}/\hbar = \omega_0 \sigma_z + \sum_i \omega_i a_i^\dagger a_i + \sum_{j=x,y,z} \sum_i \sigma_j (g_i a_i + g_i^* a_i^\dagger)$$

$$\dot{\rho}(t) = \sum_{j=x,y,z} \frac{\gamma_j(t)}{\omega_0} (\sigma_j \rho(t) \sigma_j - \rho(t))$$







$$\begin{aligned} \gamma_x(t) + \gamma_y(t) &\geq 0, \\ \gamma_y(t) + \gamma_z(t) &\geq 0, \quad \forall t \geq 0 \\ \gamma_z(t) + \gamma_x(t) &\geq 0 \end{aligned}$$

The above inequalities imply P-divisibility of the dynamics. Violation of any of the inequalities imply presence of non-Markovianity in the dynamics.



Outcome

- Non-Markovian dephasing preserves residual entanglement for longer time (Plot 1 for GHZ state).
- W states retain higher value entanglement than GHZ states under dephasing noise with increasing number of qubits. (Plots 2 and 3 for comparison)
- Removing non-Markovianity eliminates long-time saturation of residual entanglement.
- Under non-Markovian depolarising noise, entanglement revivals appear after collapse due to information backflow. (Plot 4 for GHZ state)
- W states exhibit some special features:
 - asymptotic non zero residual entanglement. (Plot 5)
 - even-odd behavior with qubit number. (Plot 6)



References

1. H.-P. Breuer and F. Petruccione, The Theory of Open Quantum Systems (Oxford University Press, New York, 2002).
2. P. Haikka, T. H. Johnson, and S. Maniscalco, Non-Markovianity of local dephasing channels and time-invariant discord, Phys. Rev. A **87**, 010103(R) (2013).
3. D. Chruscinski and S. Maniscalco, Degree of non-Markovianity of quantum evolution, Phys. Rev. Lett. **112**, 12040 ((2014).

A Fast and Frugal Gaussian Boson Sampling Emulator

➤ Tom Dodd, PhD Student, T.H.Dodd@sms.ed.ac.uk

Collaborators: Javier Martínez-Cifuentes, Oliver Thomson Brown, Nicolás Quesada, Raúl García-Patrón


Problem

- Sampling tasks pose a harder challenge to classical simulation than estimating expectation values: for ideal GBS, exact classical sampling is closely tied to complexity-theoretic hardness arguments, and generic simulation costs grow exponentially with the system size and mean photon count.
- Current photonic devices are not ideal: photon loss and other possible imperfections yield an imperfect experimental distribution, which can be benchmarked against a lossy “ground truth” Gaussian state reconstructed from the device transmission matrix and squeezing parameters.
- In threshold-detector GBS, each output mode gives a binary outcome (the absence or presence of photons), so the task is to sample from a joint probability distribution over M bits.
- The central question is whether the noisy distribution still contains high-order correlations that make it classically hard, or whether the relevant benchmark statistics are already captured by low-order cumulants.
- Previous large-scale noisy GBS simulations required extreme memory and compute costs, making routine reproduction impractical outside the use of a handful of supercomputing facilities.

Solution

- We sample from the joint distribution of M binary outcomes using a cumulant expansion combined with crafted marginal probability approximations.
- Starting from the lossy ground-truth covariance matrix, we first compute the $O(M^K)$ low-order Walsh-Hadamard coefficients $c(S)$ of the ground truth probability distribution up to a cutoff order K .
- We then compute the $O(M^K)$ cumulants associated with these coefficients, which we can use directly in our expansion.
- A cumulant $\kappa(S)$ of order k captures the novel k -mode correlation of the set of modes S not contained in lower-order cumulants:

$$\kappa(S) = \sum_{\pi \in \mathcal{P}(S)} (-1)^{|\pi|-1} (|\pi| - 1)! \prod_{B \in \pi} c(B)$$



- Samples are generated sequentially using the chain rule. For the bit at index n , the bias between an output of 0 and an output 1 is written as a sum of cumulants and marginals,


$$\Delta x_n = \sum_{R \subseteq [n-1]} \frac{1}{2^{|R|}} \chi_{R \cup \{n\}}(\bar{x}_{R \cup \{n\}}) \kappa(R \cup \{n\}) p(\bar{x}_{[n-1] \setminus R})$$

where $[n-1]$ denotes the set of all previous indices, $\chi_S(x_S)$ is the parity function of the bits with indices S , and $p(x_S)$ are the marginal probabilities

- The $R = \emptyset$ term gives the independent contribution, while higher-order terms add correlations with previously sampled bits.

Keeping every term gives an $O(2^M)$ scaling in the number of terms. Because photon loss and other noise suppresses higher-order cumulants, we truncate to the cutoff order: $|R| < K$.

- Even after truncation, we still have the marginal probabilities within the sum to compute, which retains the exponential hardness of the overall computation. Our “double-elision” algorithm is a scheme for approximating these required inner marginals using dynamic tables and a series of further truncated recurrent relations, giving us the desired $O(M^K)$ scaling for fixed cutoff order K .



Outcome


- All our samples were generated using the double-elision algorithm with cutoff $K = 5$. We compared them against the 144-mode Jiuzhang configurations' experimental samples using established validation tests: click cumulants, total click distributions and, where possible, click-count subset XEB scores.
- For low- and high-brightness Jiuzhang 2.0/3.0, our sampler's click cumulants match the ground truth more closely than those of the experimental samples.
- For low-brightness Jiuzhang 2.0, our sampler closely matches and often exceeds the XEB scores of the experimental samples.
- For high-brightness Jiuzhang 2.0 and coarse-grained Jiuzhang 3.0, XEB is too costly in the relevant click range, so validation uses total-click distributions.

• In all cases, the total-click distribution from our sampler overlaps the ground truth more accurately than the experimental samples.

• Our embarrassingly parallel implementation requires only a single CPU or GPU and 2 GB of RAM to store the 144-mode click cumulants up to the cutoff order, chosen to be $K = 5$. Each concurrent thread also uses 2 MB to cache intermediate marginals as part of executing the double-elision scheme.

• The main conclusion of this work is that GBS hardness requires the experiments to exhibit non-negligible, accurate, high-order cumulants, otherwise experiments may perform no better than our simulation method (which can be ran cheaply).

• We hope to extend this sampling method to the case of photon counting.



References

- T. Dodd, J. Martínez-Cifuentes, O. T. Brown, N. Quesada, R. García-Patrón, "A fast and frugal Gaussian Boson Sampling emulator," arXiv:2511.14923, 2025.
- C. Oh, M. Liu, Y. Alexeev, B. Fefferman, L. Jiang, "Classical algorithm for simulating experimental Gaussian boson sampling," Nature Physics 20, 1461–1469, 2024.
- N. Quesada, J. M. Arrazola, N. Killoran, "Gaussian boson sampling using threshold detectors," Phys. Rev. A 98, 062322, 2018.

QUIVER: Adaptive directional gradients for parameterised quantum circuits

➤ Brian Coyle - Senior Researcher, (b.coyle@ed.ac.uk)

Snehal Raj, Virag Umathe, El Amine Cherrat, Virag Umathe, and Elham Kashefi

Problem

- Training a parameterised quantum circuit requires gradient estimates at each step.
- The standard method (parameter-shift rule) costs $O(N)$ measurements per step which is prohibitive at $N > 5$ parameters.
- Cheaper alternatives like SPSA use only one random direction and become too noisy at scale.
- There is no principled way to trade off cost against accuracy: parameter-shift is locked to $O(N)$, SPSA to $O(1)$.
- This measurement bottleneck limits how large a model can be trained on near-term hardware.

Solution


Part 1: Forward gradients

Instead of estimating the full gradient, we aggregate V random directional derivatives per step, a **forward gradient** estimator. The per-step measurement cost is $O(V)$, not $O(N)$, so V can be set freely regardless of model size. At $V = 1$ this recovers SPSA; at $V = N$ it recovers the full gradient. In between, a modest $V = 10$ – 25 directions consistently matches the accuracy of the parameter-shift rule at a fraction of its shot budget, with the saving growing linearly with N .

$$\tilde{\nabla}^F f(\theta) = \frac{1}{V} \sum_{\ell=1}^V (\nabla_{v^\ell} f) v^\ell$$

Part 2: QUIVER

Rather than fixing V by hand, we derive the **optimal joint allocation** of directions V and shots-per-direction M that achieves a target gradient accuracy at minimum total measurement cost. The resulting closed-form rule, QUIVER, adapts both V and M online using running estimates of the gradient signal and noise level. We show that Rademacher random directions uniquely minimise the estimator variance, and that the resulting estimator is Cramér–Rao optimal, matching the parameter-shift rule in information-theoretic efficiency while costing a **constant factor fewer measurements in practice**.






$$V_t^* = \frac{(N - 1 + \alpha) \hat{g}_t^2}{\tau^2}, \quad M_t^* = \frac{N \hat{\sigma}_t^2}{\alpha \hat{g}_t^2}$$




Outcome

Forward gradients **beat** parameter-shift, SPSA and Random Coordinate Descent on 4 problems: Optimisation (QAOA), Quantum Simulation (VQE), Time Series Classification (ECG) and Image Classification (MNIST)



QUIVER **beats** adaptive measurement competitor strategies such as iCANS and gCANS.



References

- Gradients without Backpropagation, Baydin et al., arXiv:2202.08587 (2022)
- An Adaptive Optimizer for Measurement-Frugal Variational Algorithms, Kübler et al., Quantum 4, 263 (2020)
- Adaptive Shot Allocation for Fast Convergence in Variational Quantum Algorithms, Gu et al., arXiv:2108.10434 (2021)
- Quantum Circuit Learning, Mitarai et al., Phys. Rev. A 98, 032309 (2018)
- Evaluating Analytic Gradients on Quantum Hardware, Schuld et al., Phys. Rev. A 99, 032331 (2019)
- Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation, Spall, IEEE Trans. Autom. Control 37, 332 (1992)
- A Quantum Approximate Optimization Algorithm — Farhi, Goldstone & Gutmann, arXiv:1411.4028 (2014)

Distributed Quantum Property Testing with Communication Constraints


➤ Chirag Wadhwa, PhD Student, chirag.wadhwa@ed.ac.uk
With Mina Doosti and Ryan Sweke

Distributed Quantum State Certification

Our work focuses on the problem of state certification: Given copies of an unknown mixed state ρ , test whether this state is identical to a known hypothesis state σ or ε -far from it in trace distance.

The complexity of state certification is well-understood in the standard setting where a central algorithm possesses all copies, and $\Theta(d/\varepsilon^2)$ copies are optimal [1,2].

Instead, we consider solving this problem in a new, distributed setting, where m distributed nodes each possess a copy of the unknown state, and can communicate with a central node who must solve this problem.



The classical version of this problem has been studied previously [3,4], and results are also implied for the quantum problem when only classical communication is available [5]. We will thus consider the setting where each node can send q qubits, where q is small.

Techniques

Quantum Channels

We model the action of each distributed node N_i as a quantum channel Φ_i from a d -dimensional space to a 2^q -dimensional space.

Then, the central node receives states $\Phi_1(\rho), \dots, \Phi_m(\rho)$, and must use only this information to solve the certification task.

Public-vs-Private Randomness

We obtain separate complexities for this problem when the distributed nodes have shared randomness and when they do not.

In the first case, we can imagine these nodes use their shared randomness to pick the same random channels, while they are unable to do so in the second case.

Upper Bound Ideas:

We will show a choice of random channels Φ_U , such that with high probability, the output satisfies $\|\Phi_U(\rho) - \Phi_U(\sigma)\|_2 \geq \frac{1}{2} \sqrt{\frac{d_q}{d}} \cdot \|\rho - \sigma\|_2$.

Then, each distributed node applies this random channel to their copies and sends the output state to the central node, which then directly applies the algorithm of [2].

The construction of this channel is remarkably simple: apply a Haar-random unitary U , and then trace out all but the first q qubits.

Lower Bound Ideas:

We consider a point-vs-mixture distinguishing task, which any certification algorithm must be able to solve. The point is the maximally mixed state, while the mixture of alternatives is from [5].

To prove lower bounds, we want to show that no matter the choice of channels $\Phi_1, \dots, \Phi_m$, the output states are indistinguishable across the two cases, unless the number of copies is large. We do this by using the quantum Ingster-Suslina method [6] to bound the quantum χ^2 -divergence:

$$D_{\chi^2} \left(\mathbb{E}_{\rho \sim D} \left[\bigotimes_{i=1}^m \Phi_i(\rho) \right] \middle\| \bigotimes_{i=1}^m \Phi_i \left(\frac{\mathbb{1}}{d} \right) \right)$$



Results

Theorem 1 (Informal Upper Bound): With shared randomness, the complexity of certifying d -dimensional states with q -qubits of communication per node is at most $O(d^2/2^q \epsilon^2)$.

The above theorem implies that as q approaches $\log(d)$, i.e., the nodes can send over the fully copy, the complexity is $O(d/\epsilon^2)$, recovering the standard centralized bound.

Next, we show lower bounds on the copy complexity, under the assumption that each channel is mixedness-preserving.

Definition 1 (Mixedness-preserving channels): A quantum channel Φ from d_1 - to d_2 -dimensional states is said to be mixedness preserving if $\Phi(1/d_1) = 1/d_2$, i.e., the channel maps the maximally mixed state of dimension d_1 to the maximally mixed state of dimension d_2 .

Theorem 2 (Informal Public-Coin Lower Bound): With shared randomness, the complexity of certifying d -dimensional states with q -qubits of communication per node is at least $\Omega(d^2/2^q \epsilon^2)$, if each channel is mixedness-preserving.

Our lower bound shows that the complexity obtained in Theorem 1 is optimal under our mixedness-preserving assumption. We note that the algorithm we use to obtain Theorem 1 does indeed satisfy this assumption.

We also show that shared randomness is necessary to obtain this complexity, by proving a stronger lower bound when this resource is not available.

Theorem 3 (Informal Private-Coin Lower Bound): Without shared randomness, the complexity of certifying d -dimensional states with q -qubits of communication per node is at least $\Omega(d^3/4^q \epsilon^2)$, if each channel is mixedness-preserving.

Our lower bound shows that without shared randomness, the complexity of this problem is worse by at least a factor of $d/2^q$.

This presents a tradeoff between the resources of quantum communication and shared classical randomness for solving the task of state certification in our distributed setting.

References

- [1]: Quantum spectrum testing; Ryan O'Donnell and John Wright; STOC 2015.
- [2]: Quantum state certification; Costin Bădescu, Ryan O'Donnell, and John Wright; STOC 2019.
- [3]: Inference under information constraints I: Lower bounds from chi-square contraction; Jayadev Acharya, Clément Canonne, and Himanshu Tyagi; IEEE Transactions on Information Theory, vol. 66, no. 12, pp. 7835-7855, 2020.
- [4]: Inference under information constraints II: Communication constraints and shared randomness; Jayadev Acharya, Clément Canonne, and Himanshu Tyagi; IEEE Transactions on Information Theory, vol. 66, no. 12, pp. 7856-7877, 2020.
- [5]: Quantum state testing with restricted measurements; Yuhua Liu and Jayadev Acharya; arXiv:2408.17439.
- [6]: Instance-optimal quantum state certification with entangled measurements; Ryan O'Donnell and Chirag Wadhwa; STOC 2026 (To appear).

Improved algorithms for the Quantum Fourier Transform on prime-order groups

➤ David Halpenny, Research Assistant, david.halpenny@ed.ac.uk
Adithya Sireesh, Petros Wallden

Problem

Motivation

- **The Quantum Fourier Transform** is arguably the most famous and significant subroutine in quantum computing, playing a key role in almost all algorithms that achieve a super-polynomial quantum computational advantage.
- In recent years, the importance of the QFT on **prime-order groups** for algorithms that admit exponential quantum speed-ups has become apparent. [1,2]
- For example, the **Decoded Quantum Interferometry (DQI)** algorithm [1]. DQI's advantage on the OPI problem represents the first example of a provable exponential separation between quantum and classical approaches to an optimisation problem.
- This result uses the QFT on $\mathbb{Z}_p$

Objective

- New (qubit and gate efficient) algorithms for implementing the following transformation over qubits:

$$\forall j \in \mathbb{Z}_p; |j\rangle \rightarrow |\phi_p(j)\rangle = \frac{1}{\sqrt{p}} \sum_{k \in \mathbb{Z}_p} e^{2\pi i(jk/p)} |k\rangle$$

Solution

It is easy to do the transformation $|j\rangle \rightarrow |j\rangle |\phi_p(j)\rangle$

The hard part is reversibly uncomputing $|j\rangle$ i.e. we want to be able to do

$$|j\rangle |\phi_p(j)\rangle \rightarrow |0\rangle |\phi_p(j)\rangle$$

Key insight: Notice the encoded state j is hidden in the phase factor.

$$|\phi_p(j)\rangle = \frac{1}{\sqrt{p}} \sum_{k \in \mathbb{Z}_p} \exp\left(2\pi i \frac{jk}{p}\right) |k\rangle$$

More importantly, The state $|\phi_p(j)\rangle$ is an eigenstate for a carefully chosen shift unitary U i.e.

$$U|\phi(j)\rangle = e^{2\pi i j/p} |\phi(j)\rangle$$

By applying a tailored variant of **phase estimation** alongside using U , we should be able to effectively extract and uncompute j

Further operations are needed and involve doing coherent quantum arithmetic using subroutines such as (modular) addition, multiplication, division, rounding etc.




Outcome

We explore 3 variants of the

- ✓ Exact Algorithm
- ✓ Approximate Algorithm optimised for gate count
- ✓ Approximate Algorithm optimised for ancilla count

Below, we show the minimum number of ancillas required by previous methods, and our proposed methods to achieve a certain fidelity (0.95 and 0.99 respectively) for different prime orders

Minimum no. of ancilla qubits required to achieve fidelity thresholds for all primes in the range 0-4000



References

- [1] Stephen P. Jordan, Noah Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei V. Isakov, Tanuj Khattar, and Ryan Babbush. Optimization by Decoded Quantum Interferometry, May 2025
- [2] Takashi Yamakawa and Mark Zhandry. Verifiable Quantum Advantage without Structure. Journal of the ACM, 71(3):1–50, June 2024.
- [3] David Halpenny, Adithya Sireesh, and Petros Wallden. Improved algorithms for the Quantum Fourier Transform on prime-order groups. (Manuscript in preparation)

Orbit dimensions in bosonic systems

➤ Elliott Z. Mamon, PhD student, elliott.mamon@lip6.fr

Problem

- Bosonic (e.g. photonic) platforms are promising candidates for quantum computing.
- Both *Passive Linear (PLO)* and *Gaussian (GO)* quantum optics are however *sub-universal*. It is thus important to understand the limitations on the space of quantum states they can produce.
- Furthermore, the two settings differ on a theoretical level:
 - in PLO one usually considers a finite dimensional Hilbert space, with quantum states given in Fock representation (the *Discrete Variable (DV)* setting)
 - while GO acts on an infinite-dimensional Hilbert space, and one commonly describes quantum states in phase-space representation e.g. Wigner or stellar functions (the *Continuous Variable (CV)* setting)
- This calls for a common framework to quantify the expressivity in both settings at once.

Solution

The *manifold dimension* of the space of reachable states (called the orbit) can be easily calculated for various states of interest, via:

$$\begin{aligned} \dim(\text{Orb}_G(|\psi\rangle)) &= \text{rank}_{\mathbb{R}}(\{H_1|\psi\rangle, \dots, H_d|\psi\rangle\}), \\ \dim(\text{Orb}_G(\rho)) &= \text{rank}_{\mathbb{R}}(\{[H_1, \rho], \dots, [H_d, \rho]\}). \end{aligned}$$


with the following Hamiltonians of PLO or GO:

Quadratic Hamiltonians (Lie algebra bases)

$$\begin{aligned} e_{kl} &= \frac{1}{2}(a_k^\dagger a_l + a_l^\dagger a_k) & E_{kl} &= \frac{1}{2}(a_k^\dagger a_l - a_l^\dagger a_k) \\ N_k &= a_k^\dagger a_k \\ q_k &= \frac{a_k^\dagger + a_k}{\sqrt{2}} & p_k &= \frac{i(a_k^\dagger - a_k)}{\sqrt{2}} \\ s_k &= \frac{1}{2}(a_k^\dagger{}^2 + a_k^2) & S_k &= \frac{1}{2}(a_k^\dagger{}^2 - a_k^2) \\ r_{kl} &= \frac{1}{2}(a_k^\dagger a_l^\dagger + a_k a_l) & R_{kl} &= \frac{1}{2}(a_k^\dagger a_l^\dagger - a_k a_l). \end{aligned}$$

This formalism reveals topological obstructions preventing certain state or gate engineerings.

The formalism is **proven to be valid** in both discrete-variable and continuous-variable. It finds connections with quantum machine learning, metrology, resources theories of asymmetry and non-Gaussianity, symplectic rank, stellar hierarchy.



The same Hamiltonians become the following differential operators in the Wigner-function picture:

$$\begin{aligned} \widehat{\text{ad}}_{e_{kl}} &= i \left(\frac{q_l}{2} \partial_{p_k} - \frac{p_l}{2} \partial_{q_k} + \frac{q_k}{2} \partial_{p_l} - \frac{p_k}{2} \partial_{q_l} \right), \\ \widehat{\text{ad}}_{E_{kl}} &= i \left(-\frac{p_l}{2} \partial_{p_k} - \frac{q_l}{2} \partial_{q_k} + \frac{p_k}{2} \partial_{p_l} + \frac{q_k}{2} \partial_{q_l} \right), \\ \widehat{\text{ad}}_{N_k} &= i \left(\frac{q_l}{2} \partial_{p_k} + \frac{p_l}{2} \partial_{q_k} + \frac{q_k}{2} \partial_{p_l} + \frac{p_k}{2} \partial_{q_l} \right), \\ \widehat{\text{ad}}_{R_{kl}} &= i \left(\frac{p_l}{2} \partial_{p_k} - \frac{q_l}{2} \partial_{q_k} + \frac{p_k}{2} \partial_{p_l} - \frac{q_k}{2} \partial_{q_l} \right), \\ \widehat{\text{ad}}_{N_k} &= i (q_k \partial_{p_k} - p_k \partial_{q_k}), \\ \widehat{\text{ad}}_{S_k} &= i (q_k \partial_{p_k} + p_k \partial_{q_k}), \\ \widehat{\text{ad}}_{S_k} &= i (p_k \partial_{p_k} - q_k \partial_{q_k}), \\ \widehat{\text{ad}}_{q_k} &= i \partial_{p_k}, \\ \widehat{\text{ad}}_{p_k} &= -i \partial_{q_k}, \\ \widehat{\text{ad}}_{\mathbf{1}} &= 0. \end{aligned}$$

$$\dim(\text{Orb}_G(\rho)) = \text{rank}_{\mathbb{R}} \{ \widehat{\text{ad}}_{H_1} W_\rho, \dots, \widehat{\text{ad}}_{H_d} W_\rho \}$$



Outcome

Direct *geometrical meaning*: tells you something about “how many” states you can reach:

$$\text{localdim} \left(\begin{array}{c} \rho_{in} \\ \text{PLO} \\ \text{a specific PLO circuit} \end{array} \right) \leq \text{OrbDim} \left(\begin{array}{c} \rho_{in} \\ \text{PLO} \\ \text{whole PLO group} \end{array} \right)$$

$$\text{rank}_{\mathbb{R}}(\text{QFIM}_{\theta}) \leq \dim(\text{Orb}_G(|\psi_{in}\rangle\langle\psi_{in}|)) .$$

GO-orbit dimension is a *universal witness of non-Gaussianity* for multimode pure states:

Theorem:

For $m \geq 2$ modes:

$$\dim(\text{Orb}_{G_{GO}}(|\psi\rangle\langle\psi|)) \geq m(m+3),$$



with equality if and only if $|\psi\rangle$ is Gaussian.

Orbit dimensions are also *monotones for the resource theory of asymmetry*:

For every free CPTP map $\mathcal{E}$,

$$\dim(\text{Orb}_G(\mathcal{E}(\rho))) \leq \dim(\text{Orb}_G(\rho)).$$

Different proposed schemes to experimentally measure these orbit dimensions:



Example of two kinds of “two-mode cat states” that both have infinite stellar-rank but different orbit dimensions (hence not interconvertible):

$$f_a(z_1, z_2) = c(1 + e^{z_1} + e^{z_2}) \quad \dim(\text{Orb}_{G_{GO}}(|\psi_a\rangle)) = 15.$$

$$f_b(z_1, z_2) = c'(\cosh(\alpha z_1 z_2)) \quad \dim(\text{Orb}_{G_{GO}}(|\psi_b\rangle)) = 14.$$

References

Kok, Pieter, et al. “Linear Optical Quantum Computing with Photonic Qubits.” Reviews of Modern Physics, vol. 79, 2007

Adesso, Gerardo, et al. “Continuous Variable Quantum Information: Gaussian States and Beyond.” Open Systems & Information Dynamics, vol. 21, 2014

Volkoff, T. J., and Yiğit Subaşı. “Ancilla-Free Continuous-Variable SWAP Test.” Quantum, vol. 6, 2022

Shedding light on classical shadows: learning photonic quantum states

➤ Hugo THOMAS, PhD Student, hugo.thomas@lip6.fr
Pierre-Emmanuel EMERIAU, Ulysse CHABAUD

Problem


- Extracting meaningful informations from a photonic state is a challenging task, as it can require **many copies** of the input state, require an **expensive post-processing**, or simply the procedure is **unknown**.
- Such features of a photonic state play a role in **many applications**.

How to extract meaningful information of a photonic states from **few copies**, with an **efficient** classical post-processing?

Solution

Randomised measurements! We design a classical shadow protocol tailored to learning photonic states.

- Photonic shadow protocol mimics qubit protocol and is experimentally friendly:
 - Simple, linear-optical evolutions
 - Photon counting
- Efficient post-processing
 - Design a novel classical method for estimating expectation values of observables
 - Scales polynomially with the number of modes and photons, exponentially with the degree
 - Space efficient
- We provide guarantees for the number of copies and post-processing time required.
- We implemented a Julia package for the post-processing [1].







Outcome: Experimental validation on 2 QPUs [2]

QPU	Type	Application	Usage
Ascella		Mode correlators	Statistical benchmarking
12 modes, 6 photons	Reproduced experiments	Binned distribution	Validation of Boson Sampling
		Lie algebraic invariants	State Characterization
Belenos	First implementation of the	Energy estimation	Quantum chemistry
24 modes, 12 photons	protocol via shadows	Boson Sampling state learning	Learning

Ground energy estimation:



Boson Sampling state learning:



References

- [1] <https://github.com/thmhugo/LOShadows.j>
- [2] <https://cloud.quandela.com>

On Estimating the Quantum Tsallis Relative Entropy

➤ Jinge Bao, Senior Researcher, jinge.bao@ed.ac.uk
Joint work with Minbo Gao & Qisheng Wang

Problem

- We study estimation of the quantum α -Tsallis relative entropy

$$D_{Tsa,\alpha}(\rho||\sigma) = \frac{\text{tr}(\rho^\alpha \sigma^{1-\alpha}) - 1}{\alpha - 1}$$

between quantum states ρ, σ of rank r , for $\alpha \in (0,1)$.

- $D_{Tsa,\alpha}(\rho||\sigma)$ generalizes quantum α -Tsallis entropy

$$S_\alpha(\rho) = \frac{1 - \text{tr}(\rho^\alpha)}{\alpha - 1}$$

- $D_{Tsa,\alpha}(\rho||\sigma)$ reduces to von Neumann relative entropy as $\alpha \rightarrow 1$,

$$D(\rho||\sigma) = \text{tr}(\rho \log \rho) - \text{tr}(\rho \log \sigma)$$

- Quantum Hellinger distance as $\alpha = 1/2$,

$$d_{\text{Hell}}^2(\rho||\sigma) = 2(1 - \text{tr}(\sqrt{\rho}\sqrt{\sigma}))$$

- Query model (state-preparation circuits of ρ, σ)
- Sample model (copies of ρ, σ)

Techniques

Algorithmic tools

- **Block-encoding (BE):** embed matrices into left-top corner of unitary operators
- **Quantum Singular Value Transform (QSVT):** apply polynomial transformations to singular values of a BE.
- **Approximation of neg. monomial $x^{\alpha-1}$:** there exists polynomial p_- such that
 $|p_-(x)| \leq 1$ for $x \in [-1,1]$ and
 $\left|p_-(x) - \frac{\delta_-^{1-\alpha}}{2} x^{\alpha-1}\right| \leq \epsilon_-$ for $x \in [-1, -\delta_-] \cup [\delta_-, 1]$
- **Approximation of pos. monomial $x^{1-\alpha}$:** there exists polynomial p_+ such that
 $|p_+(x)| \leq 1$ for $x \in [-1,1]$ and
 $\left|p_+(x) - \frac{1}{2} x^{1-\alpha}\right| \leq \epsilon_+$ for $x \in [-1,1]$
- **Hadamard test:** given U as a $(1, a, 0)$ -BE of operator A and density operator ρ , return 0 with probability $\frac{1}{2} + \frac{1}{2} \text{Re}(\text{tr}(A\rho))$.
- **Key step:** estimate quantum affinity
 $A_\alpha(\rho||\sigma) = \text{tr}(\rho^\alpha \sigma^{1-\alpha}) = \text{tr}(\rho \cdot \rho^{\alpha-1} \sigma^{1-\alpha})$

Algorithm (estimate quantum affinity)

- Step 1: Build BEs U_A of ρ and U_B of σ via purified query access O_ρ, O_σ .
- Step 2: Apply QSVT to get $U_{p_1(A)}$, a BE of $p_1(\rho) \approx \rho^{\alpha-1}$.
- Step 3: Apply QSVT to get $U_{p_2(B)}$, a BE of $p_2(\sigma) \approx \sigma^{1-\alpha}$.
- Step 4: Multiply two BEs: $U_{p_1(A)p_2(B)}$ is the BE of $p_1(\rho)p_2(\sigma) \approx \rho^{\alpha-1}\sigma^{1-\alpha}$.
- Step 5: Estimate $\text{tr}(\rho \cdot p_1(\rho) \cdot p_2(\sigma))$ via Hadamard test.

Quantum sampler: convert query algorithms into sample-efficient algorithms with only polynomial overhead.

Computational Hardness

- **QSZK-hardness:** reduction from Quantum State Distinguishability problem (trace distance).
- **BQP-hardness:** reduction from pure state version of Quantum State Distinguishability problem (trace distance).



Results

Let $r = \max\{\text{rank}(\rho), \text{rank}(\sigma)\}$.

Query complexity

- Upper bound: $\tilde{O}(r^{1.5})$
- Lower bound: $\Omega(r^{\frac{1}{3}})$
- queries of state-preparation circuits of ρ and σ .

Sample complexity

- Upper bound: $\tilde{O}(r^{3.5})$
- Lower bound: $\Omega(r)$
- copies of ρ and σ .

Computational hardness

$TsallisQSD_\alpha[a, b]$ is

- QSZK-complete for $0 < b < 2\alpha(1 - \alpha)^4 a^4 < 32(1 - \alpha)^4 a^5$, and
- BQP-complete for $0 < b < a < 1/(1 - \alpha)$ at low rank.

$HellingerQSD[a, b]$ is

- QSZK-complete for $0 < \sqrt{2}b < a^4 < 1/4$, and
- BQP-complete for $0 < b < a < 1$ at low rank.

	Quantum ℓ_α Distance ($\alpha > 1$)	Trace Distance	Uhlmann Fidelity (Bures Distance)	Von Neumann Relative Entropy	Quantum α -Tsallis Relative Entropy ($0 < \alpha < 1$)
Query Complexity	$O(1)$	$O(r)$	$O(r)$	/	$O(r^{1.5})$
Sample Complexity		$O(r^2)$	$O(r^{5.5})$	$O(d^2)$	$O(r^{3.5})$
Hardness (Low-Rank)	BQP-hard	BQP-hard	BQP-hard	BQP-hard	BQP-hard
Hardness (General)		QSZK-hard	QSZK-hard	QSZK-hard	QSZK-hard

References

Gilyén, Su, Low & Wiebe. Quantum singular value transformation. (STOC 2019)
 Low & Chuang. Hamiltonian simulation by qubitization. (Quantum 2019)
 Rastegin. Bounds of the Pinsker and Fannes types on the Tsallis relative entropy. (MPAG 2013)
 Liu & Wang. On estimating the trace of quantum state powers. (SODA 2025)
 Aharonov, Jones & Landau. A polynomial quantum algorithm for approximating the Jones polynomial. (STOC 2006)
 Gilyén & Poremba. Improved quantum algorithms for fidelity estimation. (TQC 2024)
 Wang & Zhang. Quantum lower bounds by sample-to-query lifting (SICOMP 2025)
 Watrous. Limits on the power of quantum statistical zero-knowledge. (FOCS 2002)
 Watrous. Zero-knowledge against quantum attacks. (SICOMP 2009)
 Chakraborty, Fischer, Matsliah & de Wolf. New results on quantum property testing. (FSTTCS 2010)
 O'Donnell and Wright. Efficient quantum tomography. (CMP 2021)
 Anshu, Landau & Liu. Distributed quantum inner product estimation. (STOC 2022)

PQC Pruning with Circuit Harmonic Matrices

K.J.S.Campbell, L.DelDebbio, P.Wallden. "Circuit Harmonic Matrices: A Spectral Framework for Quantum Machine Learning." arXiv:2604.04292 (2026).

➤ Kyle J.S. Campbell, PhD Student, Kyle.J.S.Campbell@ed.ac.uk
Joint work with Luigi Del Debbio, Petros Wallden

Problem

How parameterised quantum circuits (PQCs) learn is difficult to understand analytically. We do know that for certain classes of PQCs, the architecture controls which Fourier modes they can represent, how strongly those modes are expressed, and how different modes interact during training [1,2,3,4].

This structure is usually found empirically, by training many circuits and comparing performance. This is costly, noisy, and gives limited insight into why one architecture works better than another and how we can choose such architectures a priori.

Key question: can we describe a PQC in a framework that reveals its learning structure before training and use this to inform our design choices?

Solution

The Circuit Harmonic Matrix C: a training-free representation of a parameterised quantum circuit, computed from its gate structure [1].

A PQC output can be written as a joint Fourier expansion in both the input x and trainable parameters θ :

$$f(x; \theta) = \sum_{\omega \in \Omega} a_{\omega}(\theta) e^{i\omega x} = \sum_{\omega \in \Omega} \sum_{k \in K} C_{\omega k} e^{i\omega x} e^{ik \cdot \theta}.$$

The rows of C correspond to input Fourier modes ω . The columns correspond to parameter harmonics k . In this way, C records which frequencies the circuit can access, how strongly they appear, and which parameter structures generate them.

Fourier covariance matrix: From C , we compute a covariance matrix

$$\Sigma_{\omega\omega'} = \sum_{k \neq 0} C_{\omega k} C_{\omega' k}^{\dagger}.$$

The diagonal entries measure how strongly each Fourier mode is represented. The off-diagonal entries show which modes are coupled through shared parameter harmonics. This gives a compact, data-free representation of the circuit's learning structure.

C Matrix Structure

	k_a	k_b	k_c	k_d	...
ω_0	$C_{\omega_0 k_a}$	$C_{\omega_0 k_b}$	0	0	...
ω_1	$C_{\omega_1 k_a}$	0	$C_{\omega_1 k_c}$	0	...
ω_2	0	$C_{\omega_2 k_b}$	0	0	...

ω_1, ω_2 share no k -column $\Rightarrow \text{Cov}[a_{\omega_1}, a_{\omega_2}] = 0$
 ω_0, ω_2 share k_b $\Rightarrow \text{Cov}[a_{\omega_0}, a_{\omega_2}] \neq 0$

Parameter importance score: We can use this object to understand how much each part of the circuit contributes to learning. For each parameter θ_i , we effectively remove its corresponding gate and compute the shift in the covariance matrix:

$$\Delta_i = \|\Sigma - \Sigma(\theta_i = 0)\|.$$




A large covariance shift Δ_i means that freezing the gate strongly changes the circuit's learnable structure. A small Δ_i means the gate has little effect on the Fourier covariance and is a candidate for pruning.

Circuit pruning rule: Rank gates by Δ_i , remove low- Δ_i gates, and keep the parameters that most shape the circuit's learnable Fourier structure whilst retaining circuit's learning ability.






Outcome

Slightly Prunable Circuit



Highly Prunable Circuit



Key findings: C-guided pruning via covariance matrices (A) successfully identifies redundant gates and preserves training ability.

- In the slightly prunable circuit, removing **low Δ_i gates** gradually increases final loss error (C), while removing **high Δ_i gates** causes rapid increases. This is explained in (B) where both pruning methods remove similar amounts of Fisher matrix eigenvalue mass (learnability of the circuit).
- In the highly prunable circuit, we can remove around 70% of the **low Δ_i gates** with negligible loss increase. This is seen in (B) where bottom pruning removes significantly less Fisher mass than top.

Why it matters: Smaller circuits mean reduced depth, fewer operations, lower exposure to decoherence and gate noise, and lower training and hardware costs.

Impact and applications: Potentially move QML to more near-term devices and allow architecture design without trial-and-error training.

References

- [1] K.J.S. Campbell, L. Del Debbio, P. Wallden, "Circuit Harmonic Matrices: A Spectral Framework for Quantum Machine Learning," arXiv:2604.04292 (2026).
- [3] Adrián Pérez-Salinas et al. In: Quantum 4 (2020). arXiv:1907.02085 [quant-ph].
- [4] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. In: Physical Review A 103.3 (2021).
- [5] Hela Mhiri et al. In: Quantum 9 (2025). arXiv:2403.09417 [quant-ph].

Boson sampling beyond the dilute regime: second moments and anti-concentration

➤ Hela Mhiri, Hugo Thomas, Léo Monbroussou (Imonbrou@ed.ac.uk), Ulysse Chabaud, Zoë Holmes, Elham Kashefi


Problem

- Boson sampling is a leading candidate for demonstrating quantum advantage in photonic systems, but its output statistics remain poorly understood beyond the dilute regime, where most previous analyses of second moments and anti-concentration have been carried out. This is a major limitation, since experimentally relevant photonic devices operate in collision-dominated regimes.
- Closed-form expressions for second moments of generic particle-number-preserving bosonic observables across arbitrary photon-mode scalings are hard to compute, but important for many applications (certification, benchmarking, classical simulation, variational algorithms).
- Anti-concentration of the output bosonic distribution is an important ingredient in hardness arguments for random sampling schemes. Until this work, it was still a conjecture beyond the dilute regime.

Solution

First, we develop a representation-theoretic framework for passive linear optics that yields closed-form expressions for second moments of expectation values of particle-number-preserving observables. Concretely, we express these moments in terms of Hilbert-Schmidt norms of projections onto irreducible components of the operator space and show that these norms admit compact analytical formulas from the underlying symmetry structure.

Second, we apply this framework to boson sampling. In particular, we derive a closed-form expression for the second moment of averaged output probabilities, corresponding to the normalized average outcome collision probability, and use its asymptotic scaling to prove anti-concentration beyond the dilute regime, including in the presence of photon collisions. This complements recent complexity-theoretic advances on boson sampling in the saturated regime and also yields quantities of independent interest for linear cross-entropy benchmarking.



Schematics of passive linear optics transformations from a representation theory perspective. a) Linear-optical circuit made of beam-splitters and phase-shifters on m spatial modes described by a unitary in $U(m)$, with single photons as input (left); action of a linear interferometer with n input photons through the photonic homomorphism φ giving rise to a block-diagonal decomposition of the Fock space under passive linear optics transformations (right). b) Decomposition of the operator space of particle-number-preserving observables into irreps under the action of the adjoint representation ω , same copy appearing in different particle sector have same color and connected via raising and lowering maps.




Outcome

- Second Moment:

$$\mathbb{E}_{U \sim U(m)} [f_U(\rho, O)^2] = \sum_{k=0}^n \frac{1}{d_k^{(n)}} \|P_k^{(n)}(\rho)\|_2^2 \|P_k^{(n)}(O)\|_2^2$$

- Iterative Projection for Purity Analysis based on representation theory, to calculate purity norms:




Purity norm of Fock states considering 20 modes and 10 particles.

- Proof of anti-concentration beyond the dilute regime:

$$\Pr_{\substack{U \sim U(m) \\ S \sim \Phi_m^n}} \left[p_U(S) \geq \frac{\alpha}{|\Phi_m^n|} \right] \geq \gamma(\alpha)$$

With $p_U(S)$ the probability of observing the outcome S . A natural quantity used for diagnosing anti-concentration is given by:

$$P_2(m, n) := |\Phi_m^n| \sum_{S \in \Phi_m^n} \mathbb{E}_{U \sim U(m)} [p_U(S)^2]$$



References


[Boson sampling beyond the dilute regime: second moments and anti-concentration](#), Hela Mhiri, Hugo Thomas, Léo Monbroussou, Ulysse Chabaud, Zoë Holmes, Elham Kashefi, arXiv preprint arXiv:2604.14323.

Exploiting the Gradient Flow to Characterize Parameterized Quantum Circuit Landscapes

➤ Mario Herrero González, mario.herrero@ed.ac.uk
Elham Kashefi, Mina Doosti

Problem

- We work with parameterized quantum type of losses $\mathcal{L}(\theta) = \text{Tr}[OU(\theta)\rho U^\dagger(\theta)]$
- Barren Plateaux (BPs), which describe vanishing variance of the loss $\text{Var}(\mathcal{L}) \sim 2^{-n}$, work primarily with random initialization where the loss landscape is treated uniformly, no information of the optimizer is taken into account. BP literature does not contemplate the fact that after a polynomial number of steps one may fall in a better section of the loss landscape.
- These results have affected negatively the reputation of variational methods, where for data relevant cases may be inaccurate.
- Our objective is then to show that more precise scaling behaviors can be obtained by harnessing the behavior of the landscape and show how different assumptions can render different results



Solution

We propose a more landscape-aware approach by introducing a new integral measure. This integral measure in its raw form follows the gradient flow which can be characterized by conserved quantities.

The variance, which is a proxy for trainability, is obtained via integrating over this measure. [See Fig 1. for a simple example of how different parameters affect the variance]

$$\text{Var}_{\Gamma_r}[\mathcal{L}] = \int_{\Gamma_r} (\mathcal{L}(\theta) - \mathbb{E}_{\Gamma_r}[\mathcal{L}])^2 d\theta$$

However, we cannot always compute the exact gradient flow, or obtain all of the conserved quantities as this would imply being able to solve exactly for large problem sizes.

To make evident the misinterpretation given by fully random measures we provide a modification to the gradient flow based measure obtained through the gradual randomization (or Haarization) of the solution to the conserved quantities which in the limit equates to the result obtained through random initialization. This partial randomization can ease the calculation of the gradient flow,

$$\mu_{\tau_H} = f(\mu_{\text{Haar}}, \mu_{\text{GF}})$$




Figure 1: For the 2 qubit case, we show **a)** the gradient flow paths for different initializations characterized by the ratio parameter h and different integration length paths r . **b)** The variance at different initialization points and h but fixed r . **c)** The variance throughout different initializations at fixed conserved quantity but different lengths of the integration path.



Outcome

The cases we care about are those of large number of qubits. We focus in the Brickwork case with single RY rotations and nearest-neighbor CZ, and global cost. Which is now to plateau no matter the depth of the circuit.


In contrast with Fig 1. we now study the effect of different ensembles of initial points within a vicinity. This makes the integral in Eq. 1 go from being in one dimension (throughout the conserved quantity) to two dimensional. This defines the pushforward measure.

$$\int_{\Theta} f(\theta) d\mu_t(\theta) = \int_{\Theta} f(\varphi_t(\theta)) d\mu_0(\theta)$$

Our results show that under different assumptions we can achieve different behaviors in the scaling. These different assumptions are implemented in the measure through a stochasticity or Brownian parameter τ_B depicting a Haarization procedure.

Theorem 1 (informal): Under the assumption where local path information is relevant, Barren plateaus are less prohibitive.

Figure 2 on the right shows how the Barren plateau phenomena is not as aggressive when taking into account local structure as in the Uniform integration case. In all values of depth D the cost function vanishes exponentially for the Haar uniform case, while we see that for different stochasticity strength values this does not follow.



References


- [1] Abide by the Law and Follow the Flow: Conservation Laws for Gradient Flows - Sibylle Marcotte, Rémi Gribonval, Gabriel Peyré
- [2] Cost Function Dependent Barren Plateaus in Shallow Parametrized Quantum Circuits - M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, Patrick J. Coles

Simulability in Quantum Generative Models: Conditions for Advantageous Quantum Generative Modelling

➤ **Mario Herrero González, PhD Student, mario.herrero@ed.ac.uk**
Joint work with Ross Grassie, Kieran McDowall, Sjoerd Beentjes, Ava Khamseh & Elham Kashefi


Problem

- Assess the potential quantum advantage of Quantum Generative Models.
- Parameterised Quantum Circuits (PQCs) are difficult to train. Although a certain family of them show sampling advantage over classical counterparts
- We surrogate the training from a k-order correlator perspective
- Then we assess if the inherent bias of the architecture gives us a representation advantage when sampling



Solution

We analyse the importance of k-order correlations. Some turn up to be essential to have an exact training




We note that there is a great dependence in the complexity of the data at hand. If it is not too complex we can use simple Tensor Networks or Pauli Propagation to achieve good values of correlators (extremes). If it is too complex (middle) we need to use more resources in the surrogate.



Outcome


Quantum generative models are kernel methods and its advantage can be assessed as such. The following conditions apply:

- Easy to sample from correlators
- Quantum Fourier distribution has polynomial size effective support
- The Classical Fourier distribution is aligned with that of the Quantum one



Moreover, when training classically and then deploying quantumly, different discrepancies arise, making the advantage at inference less likely.

To be able to raise our chances of having advantage at inference we can introduce different inductive biases. The first figure shows improvement in the bias for different datasets. The second figure shows regimes where we do not have systematic improvement vs one where we do.



Average Case Sampling Complexity of QAOA

➤ Sam Winter, PhD Student, (S.M.B.Winter@sms.ed.ac.uk)
Joint work with Petros Wallden

Problem

- The quantum approximate optimisation algorithm (QAOA) [1] has been applied to many industrially relevant NP-hard problems, and has shown a polynomial scaling advantage for particular use cases i.e. max k sat [2] and CVP [3].
- The success probability of finding the true optimum for these problems using the shallowest depth version of QAOA is not known for average case problems [4].
- Further, identifying optimal parameters for the QAOA typically requires expensive sampling. Parameter setting strategies are a key component of addressing the utility of QAOA, and understanding where scaling advantages come from. Ideally, there would be a fully classical routine that would allow the optimal angles to be determined before executing QAOA circuits on quantum hardware [4].
- In this work we address these problems by: identifying a global property of families of problems in the form of a polynomial, using this to construct an approximation of the success probability of QAOA, and bounding this value.

Solution

Binary optimisation problems mostly take the form:

$$f(x) = \sum_{R \subseteq [n]} J_R \prod_{i \in R} x_i$$

where J_R is a real number. The degree of f is the largest size of R for which there is a J_R not equal to zero. Determining the input that minimises f

$$z = \operatorname{argmin}_{x \in \mathcal{S}^n} f(x).$$




is hard to do. However, even without knowing the value, it is possible to calculate the average value of f over the bit strings that are at some hamming distance d from z . The expectation is a polynomial of hamming distance, while the distribution of output values of f are gaussian.

The figure below shows the distribution of output values of f on the right, and the average value with increasing hamming distance on the right. This polynomial allows the overlap with the ground state

$$\langle z | \psi(\beta, \gamma) \rangle = \frac{e^{i\gamma f(z)}}{\sqrt{2^n}} \sum_{d=0}^n \sum_{x \in \mathcal{S}(z,d)} c^{n-d} e^{i\gamma d} e^{i\gamma (f(x) - f(z))}$$

and the success probability by


$$P_s = \left| \frac{1}{\sqrt{2^n}} \sum_{d=0}^n s^d c^{n-d} \binom{n}{d} e^{i\frac{\pi}{2} + i\gamma \omega_d - d(\gamma \sigma_d)^2/2} \right|^2$$





Outcome

The consequences of this has lead to the ability to efficiently approximate the success probability for average case instances. The left figure shows the success probability for the cubic optimisation problem in 16 variables. The right hand side shows simulations of 4096 samples from noiseless hardware, while the right hand sides hows our guess. This demonstrates a capability to efficiently estimate the success probability before the execution QAOA, and a method to find the optimal parameters ahead of time.



This is then used to construct a function that estimates the success probability as a function of the two variational parameters. Further, we go on to show that the sampling complexity, with parameter dependant q and $\alpha \geq 0$

$$N_{samples} = \Omega(2^{n^{\alpha}}).$$

This result shows a sub-exponential complexity with single layer QAOA may be possible for average case problems.

The next stages of this research will cover:

- Analysing families of optimisation problems such as graph colouring and max-k-sat.
- Extending these results to higher depth versions of QAOA
- The applications to industry datasets, notably graph problems in real telecommunications network topologies.
- The comparison with real hardware samples.

References

- [1] Farhi, Edward, Jeffrey Goldstone, and Sam Gutmann. "A quantum approximate optimization algorithm." *arXiv preprint arXiv:1411.4028* (2014).
- [2] Boulebnane, Sami, and Ashley Montanaro. Solving Boolean satisfiability problems with the quantum approximate optimization algorithm. *PRX Quantum* 5.3 (2024): 030348.
- [3] Priestley, Ben, and Petros Wallden. "A practically scalable approach to the closest vector problem for sieving via qaoa with fixed angles." *Quantum Science and Technology* 11.2 (2026): 025018.
- [4] Blekos, Kostas, et al. "A review on quantum approximate optimization algorithm and its variants." *Physics Reports* 1068 (2024): 1-66.

Benincasa-Dowker-Glaser causal set actions by quantum counting

➤ Sean A. Adamson, Researcher, sean.adamson@ed.ac.uk
Joint work with Petros Wallden [1]

Problem

Causal set theory is an approach to quantum gravity in which spacetime is treated as a discrete collection of points.

- The Benincasa–Dowker–Glaser (BDG) action is a candidate discrete action defined on causal sets [2].
- Taking its continuum limit recovers the Einstein–Hilbert action of general relativity, from which Einstein's field equations follow.

- Classical algorithms that evaluate the BDG action scale as $O(n^3)$ in the number of spacetime points n .
- Simulating very large causal sets could prove valuable for advancing the phenomenology of the theory.

Could quantum computers provide a speedup in computing the BDG action of a given causal set?

Solution

Computing the BDG action reduces to counting order intervals of given cardinalities.




Figure 1: Order intervals of cardinality $k=1$ (left) and cardinality $k=7$ (right).

- We show classical random sampling provides only a constant speedup, at the cost of introducing an error $\epsilon > 0$.
- We find a linear-time quantum circuit (Fig. 2) of an oracle identifying order intervals of a given cardinality.
- This makes use of recent optimizations in compiling multi-controlled Toffoli gates [3].

- We develop a two-stage variant of quantum counting [4] to decrease the number of order intervals searched from $O(n^2)$ to $O(n)$.
- The running time of each oracle call is fast enough to yield an overall speedup, even when the causal set is given as classical data.

The algorithm proceeds as follows:

1. Prepare a uniform superposition

$$\frac{1}{n} \sum_{i,j=1}^n |\mathbf{r}_i\rangle |\mathbf{c}_j\rangle$$

where $\mathbf{r}$ and $\mathbf{c}$ describe the causal set (row i and column j of its adjacency matrix).

2. Perform a two-stage quantum counting variant, calling the oracle circuit (Fig. 2) as a subroutine.

- This estimates the count N_k of cardinality k order intervals.

3. Repeat steps 1 and 2 for different k and combine the counts to obtain the BDG action

$$\frac{1}{\hbar} S(C) = \frac{4}{\sqrt{6}} \left(\frac{l}{l_p} \right)^2 [n - N_0 + 9N_1 - 16N_2 + 8N_3]$$

Quantum resources required:

- Step 1 requires $\tilde{O}(n^2)$ time, $O(n)$ space.
- Step 2 requires $\tilde{O}(n^2\epsilon^{-1})$ time, $O(n)$ space.
- Step 3 requires constant time and space.




Figure 2: Quantum oracle circuit identifying cardinality k order intervals for a causal set with n elements.

Outcome

- Overall time is $\tilde{O}(n^2)$, improved from $O(n^3)$.
- Memory required is $O(n)$ qubits: still scales better than those relying on classical matrix multiplication.

Next steps:

- Quantum algorithms for other causal set quantities, e.g., spacelike distance.
- Adiabatic-inspired NISQ approach to counting with optimized scheduling [5].

Approach	Type	Time complexity	Space complexity
Naive classical	Classical	$O(n^3)$	$\tilde{O}(1)$
Random sampling	Classical	$O(n^3)$	$\tilde{O}(1)$
Strassen ($n \geq 1000$)	Classical	$O(n^{2.81})$	$\tilde{O}(n^2)$
Laser method (galactic)	Classical	$O(n^{2.38})$	$\tilde{O}(n^2)$
Quantum counting	Quantum	$\tilde{O}(n^2 \epsilon^{-1})$	$O(n)$

References

- [1] S. A. Adamson, P. Wallden, "Benincasa-Dowker-Glaser causal set actions by quantum counting", Phys. Rev. Research **8**, 023188 (May 2026).
- [2] D. M. T. Benincasa and F. Dowker, "Scalar curvature of a causal set", Phys. Rev. Lett. **104**, 181301 (May 2010).
- [3] B. Claudon et al., "Polylogarithmic-depth controlled-NOT gates without ancilla qubits" Nat. Commun. **15**, 5886 (2024).
- [4] S. Aaronson and P. Rall, "Quantum approximate counting, simplified", Symposium on Simplicity in Algorithms (SOSA), pp. 24–32 (2020).
- [5] S. A. Adamson and P. Wallden, "Adiabatic quantum unstructured search in parallel", arXiv:2502.08594 (2025).

Private training in quantum machine learning

➤ Tigran Sedrakyan, PhD Student, tigran.sedrakyan@lip6.fr
Joint work with Frédéric Grosshans and Elham Kashefi

Problem

- Training machine learning models on sensitive data needs formal privacy guarantees. These can be enforced at the input level, during training or at the inference stage. The common tool for training-time privacy is DP-SGD [1] based on differential privacy (DP) [2]: at every training iteration clip per-example gradients at C and add calibrated Gaussian noise of scale σC . DP-SGD trades accuracy for privacy due to the way it perturbs the training procedure. Two perturbation sources dominate: clipping bias [3], when per-example gradient norms exceed C and injected noise scaling with the available privacy budget defined by parameters (ϵ, δ) .
- QML models are not immune to privacy attacks [4] – yet the interplay between QML and classical DP-SGD remains poorly understood.
- In our work we consider two questions. First, when considering training-time perturbation, can we use native quantum noise to replace the calibrated DP noise in DP-SGD? Second, how can the clipping bias explain the observation that quantum models retain higher accuracy under the same privacy budget than matched classical baselines [5]?

Solution

Quantum noise propagation

Given a quantum classifier with traceless-Pauli logits, depolarizing channel of strength λ , shot count S , the per-example loss-gradient perturbation is approximately additive Gaussian:

$$\widehat{\partial_{\theta} \ell} \approx \partial_{\theta} \ell + \mathcal{N}(m, v)$$

where m and v depend on λ , S , but also on the training data. Thus despite the fact that quantum shot and hardware noise look Gaussian, the scale of the noise is data dependent and not suitable for differential privacy. Instead, the formal guarantee must come from classical DP-SGD [1] and we consider how quantum models may have an edge in this case.

Clipping-bias bound

For a circuit defined by $U(\theta) = \prod e^{-i\theta_j H_j}$ class observables O_k and cross-entropy loss at temperature τ , define G via the gate generators and observables


$$G = 2 \max_k \|O_k\| \left(\sum_j \|H_j\|^2 \right)^{1/2}$$

then per-example loss gradients and the clipping bias are bounded deterministically:

$$\|\partial_{\theta} \ell(x_i, \theta)\|_2 \leq \frac{2G}{\tau}$$

$$\|b_C(\theta)\|_2 \leq \left(\frac{2G}{\tau} - C \right)_+$$

indicating that quantum models may tolerate a smaller clipping threshold C without inducing clipping bias, which lowers σC and improves the privacy-utility trade-off for training-time perturbations.







Outcome


Synthetic dataset – clipping & DP-SGD

To study the effect of accuracy loss under clipping in an isolated environment, we perform a non-trivial 4-class classification task on a synthetic dataset. The classical model is a fully connected neural network, while the quantum model uses amplitude-encoded variational classifier with a matching number of parameters. Both classical and quantum models have similar accuracy without clipping, however the classical model has a lower accuracy under clipped training due to clipping bias. We perform the same experiment for private training sweeping across the clipping threshold and privacy budget, observing a similar picture.



Private image classification on real-world data

We also consider a more practical task of image classification, on an image dataset containing 64 x 96 grayscale images distributed among 4 classes. The classical model is a convolutional neural network, while the quantum model is an amplitude-encoded variational classifier with a matched number of parameters. Under DP-SGD with $\epsilon=10$, $\delta=10^{-5}$ and $C=1.0$ the quantum model achieves higher accuracy and stays closer to the non-private optimisation trajectory throughout training, indicating better performance due to the lower clipping bias under identical privacy budget.



References

- [1] Abadi et al., Deep Learning with Differential Privacy. CCS '16.
- [2] Dwork & Roth, The Algorithmic Foundations of Differential Privacy. NOW, 2014.
- [3] Chen, Wu, Hong, Understanding Gradient Clipping in Private SGD: a Geometric Perspective. NeurIPS '20.
- [4] Heredge et al., Characterizing Privacy in Quantum Machine Learning. npj Quantum Information, 2025.
- [5] Watkins, Chen, Yoo, Quantum Machine Learning with Differential Privacy. Scientific Reports 13, 2023.



www.quantumsoftwarelab.com



QuantumSoftwareLab@ed.ac.uk



Informatics Forum,
The University of Edinburgh,
10 Crichton St, Newington,
Edinburgh, EH8 9AB United Kingdom