QCTiP

Oxford | 2026

Quantum Computing Theory in Practice

Book of Abstracts





Mathematical Institute, University of Oxford

20-22 April 2026

Hosts

UNIVERSITY OF OXFORD

UK Research and Innovation

Co-Host

National Quantum Computing Centre


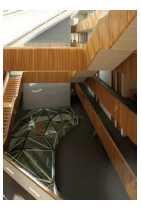

Contents

About

Quantum Computing Theory in Practice (QCTiP)

We are witnessing impressive progress in quantum hardware development and ongoing theoretical advancements are bringing practical applications of this hardware closer to reality. Quantum Computing Theory in Practice (QCTiP) aims to bring together the academic community and industry representatives to foster discussions on how to unlock the full potential of quantum computers.

Steering Committee

- Ophelia Crawford (Riverlane)
- Elham Kashefi (CNRS & University of Edinburgh)
- Jens Eisert (Freie Universität Berlin)
- Noah Linden (University of Bristol)
- Ashley Montanaro (University of Bristol & Phasecraft)

Local Organisers

- Bálint Koczor (Mathematical Institute, University of Oxford)
- Aleks Kissinger (Computer Science, University of Oxford)

Programme Committee Chair

- Richard Kueng (Johannes Kepler University Linz)

Programme Committee

- Antonio Anna Mele (Freie Universität Berlin)
- Francesco Anna Mele (Scuola Normale Superiore di Pisa)
- Leandro Aolita (Technology Innovation Institute Abu Dhabi)
- Janos Asboth (Budapest University of Technology)
- Lucas Berent (Iceberg Quantum)

- Vjosa Blakaj (University of Copenhagen)
- Paolo Braccia (Los Alamos National Lab)
- Michael Bremner (UTS Sydney)
- Dan Browne (University College London)
- Daniel Cavalcanti (Algorithmiq)
- Ieva Cepaite (Phasecraft)
- Jonathan Conrad (EPFL Lausanne)
- Arjan Cornelissen (Simons Institute)
- Abhinav Deshpande (IBM Research)
- Andrew Doherty (University of Sydney)
- Arkopal Dutt (IBM Quantum)
- Andreas Elben (Paul Scherrer Institute)
- Francisco Escudero Guiterrez (CWI Amsterdam)
- Paul Faehrmann (Freie Universität Berlin)
- Philippe Faist (Freie Universität Berlin)
- Marco Fanizza (Inria)
- Diego Garcia Martin (Los Alamos National Lab)
- Linnea Grans-Samuelsson (University of Oxford)
- Sander Gribling (Tilburg University)
- Dmitry Grinko (QuSoft Amsterdam)
- Shouzhen Gu (Yale)
- Jonas Haferkamp (Saarland University)
- Robin Harper (University of Sydney)
- Tobias Haug (Technology Innovation Institute Abu Dhabi)
- Jonas Helsen (CWI Amsterdam)
- Nicholas Hunter-Jones (UT Austin)
- Sofiene Jerbi (Freie Universität Berlin)
- Yifan Jia (University of Copenhagen)
- Michael Kastoryano (AWS and University of Copenhagen)
- Robbie King (Google Quantum AI and Caltech)
- Johannes Knoerzer (ETH Zürich)
- Aleks Kubica (Yale)
- Anthony Leverrier (INRIA)
- Yunchao Liu (IBM Research)
- Peter Love (Tufts University)
- Angelo Lucia (Politecnico di Milano)
- Sam McArdle (AWS)
- Alex Meiburg (Perimeter Institute)
- Alex Nietner (Freie Universität Berlin)
- Harold Ollivier (Inria)
- Michal Oszmaniec (CTP PAS Warsaw)
- Aadil Oufkir (RWTH Aachen)

- Adrian Perez-Salinas (ETH Zürich)
- Lirandë Pira (National University of Singapore)
- Alejandro Posas-Kerstjens (Universite de Geneve)
- Sergi Ramos Calderer (National University of Singapore)
- Joschka Roffe (University of Edinburgh)
- Gregory Rosenthal (University of Toronto)
- Mark Steudtner (PsiQuantum)
- Daniel Stilck-Franca (University of Copenhagen)
- David Sutter (IBM Quantum)
- Jordi Tura (Leiden University)
- Albert Werner (University of Copenhagen)
- Mischa Woods (Inria)
- James Wootton (MothQuantum)
- Leonardo Zambrano (ICFO)
- Zoltan Zimboras (Algorithmiq and University of Helsinki)
- Sebastian Zur (IRIF, CNRS)

Acknowledgements

The Organising Committee thank Viet Tran (Johannes Kepler University Linz, Austria), Po-Wei Huang and Abhishek Rajput (University of Oxford) for their help in managing submissions and scheduling the scientific programme. Many thanks also to Po-Wei Huang for creating the book of abstracts.

The organisers acknowledge the generous support of the following confirmed sponsors and academic partners.

Hosts

University of Oxford

UK Research and Innovation

Co-Host

National Quantum Computing Centre

Main Sponsors

Quantinuum

nu Quantum

Phasecraft

Gala Dinner Sponsor

Quantum Motion

General Sponsors

Google Quantum AI

river lane

Academic Partners

Oxford Mathematics

University of Oxford Department of Computer Science

University of Oxford Oxford Quantum Institute
Supporting Oxford's quantum research

QCi3 HUB

Keynote Talks

Multi-qubit Toffoli with exponentially fewer T gates

Robin Kothari

Google Quantum AI

Prior work of Beverland et al. has shown that any exact Clifford+T implementation of the n-qubit Toffoli gate must use at least n T gates. Here we show how to get away with exponentially fewer T gates, at the cost of incurring a tiny $1/\mathrm{poly}(n)$ error that can be neglected in most practical situations. More precisely, the n-qubit Toffoli gate can be implemented to within error ϵ in the diamond distance by a randomly chosen Clifford+T circuit with at most $O(\log(1/\epsilon))$ T gates. We also give a matching $\Omega(\log(1/\epsilon))$ lower bound that establishes optimality, and we show that any purely unitary implementation achieving even constant error must use $\Omega(n)$ T gates. We also extend our sampling technique to implement other Boolean functions. Finally, we describe upper and lower bounds on the T-count of Boolean functions in terms of non-adaptive parity decision tree complexity and its randomized analogue.

Pipes, Loops and Snakes: Quantum computing with semiconductor devices

Simon Benjamin [1, 2]

[1] University of Oxford
[2] Quantum Motion

Semiconductor foundries can create chips within which qubits, rather than bits, are stored and processed. Transistor-like structures at the 100nm scale can each hold an electron whose spin represents a qubit. This is an appealing approach for scalability: a QPU with millions of qubits needs only to occupy a couple of square millimetres of silicon. Recent progress with fast and high-fidelity shuttling of electrons has further increased excitement for this paradigm because it profoundly increases the connectivity. I'll briefly introduce these concepts and the hardware progress to-date, before focusing on the potential architectures for full scale fault tolerant QC.

In the looped pipeline approach, multiple qubits move in closed cycles – remarkably, this allows an effective 3D architecture even though the physical device is 2D. This can be very helpful in routing information and for supporting otherwise-impossible codes. A more radical paradigm is "snakes on a plane" which reformats logical qubits into 1D entities that can then move freely over a 2D lattice. This creates a vulnerability that can be tackled using a powerful resource called decoder confidence (a.k.a. soft information) that has multiple applications. All these ideas are compatible with both single-electron physical qubits and with dual-spin encoding, the latter bringing the advantage of "erasure conversion".

Contributed Talks on Monday

Algorithms I – Optimization

A Rigorous Quantum Framework for Inequality-Constrained and Multi-Objective Binary Optimization

Sebastian Egginger[1], Kristina Kirova[1], Sonja Bruckner[2], Stefan Hillmich[2], Richard Kueng[1]

[1] Johannes Kepler University Linz
[2] Software Competence Center Hagenberg

Encoding combinatorial optimization problems into physically meaningful Hamiltonians with tractable energy landscapes forms the foundation of quantum optimization. Numerous works have studied such efficient encodings for the class of Quadratic Unconstrained Binary Optimization (QUBO) problems. However, many real-world tasks are constrained, and handling equality and, in particular, inequality constraints on quantum computers remains a major challenge. In this work, we show that including inequality constraints is equivalent to solving a multi-objective optimization. This insight motivates the Multi-Objective Quantum Approximation (MOQA) framework, which approximates the maximum via p-norms and comes with rigorous performance guarantees. MOQA operates directly at the Hamiltonian level and is compatible with, but not restricted to, ground-state solvers such as quantum adiabatic annealing, the Quantum Approximate Optimization Algorithm (QAOA), or imaginary-time evolution. Moreover, it is not limited to quadratic functions.

A scalable quantum-enhanced greedy algorithm for maximum independent set problems

Elisabeth Wybo, Jami Rönkkö, Olli Hirviniemi, Jernej Rudi Finzgar, Martin Leib

IQM Quantum Computers

We investigate a hybrid quantum-classical algorithm for solving the Maximum Independent Set (MIS) problem on regular graphs, combining the Quantum Approximate Optimization Algorithm (QAOA) with a minimal degree classical greedy algorithm. The method leverages pre-computed QAOA angles, derived from depth-p QAOA circuits on regular trees, to compute local expectation values and inform sequential greedy decisions that progressively build an independent set. This hybrid approach maintains shallow quantum circuit and avoids instance-specific parameter training, making it well-suited for implementation on current quantum hardware: we have implemented the algorithm on a 20 qubit IQM superconducting device to find independent sets in graphs with thousands of nodes. We perform tensor network simulations to evaluate the performance of the algorithm beyond the reach of current quantum hardware and compare to established classical heuristics. Our results show that even at low depth ($p = 4$), the quantum-enhanced greedy method significantly outperforms purely classical greedy baselines as well as more sophisticated approximation algorithms. The modular structure of the algorithm and relatively low quantum resource requirements make it a compelling candidate for scalable, hybrid optimization in the NISQ era and beyond.

Verifiable Quantum Advantage via Optimized DQI Circuits

Tanuj Khattar, Noah Shutty, Craig Gidney, Adam Zalcman, Noureldin Yosri, Dmitri Maslov, Ryan Babbush, Stephen P. Jordan

Google Quantum AI

Recently, a quantum algorithm called Decoded Quantum Interferometry (DQI) was introduced that achieves an apparent exponential speedup for Optimal Polynomial Intersection (OPI) problem, which has previously been studied in the contexts of cryptography and error correcting codes. However, this left open the question of how many logical gates and logical qubits would be needed to solve a classically intractable instance of OPI. Here, we develop optimized implementations of DQI which greatly reduce its resource requirements. We establish that DQI for OPI is the first known candidate for verifiable quantum advantage with optimal asymptotic speedup: solving instances with classical hardness $O(2^N)$ requires only $\widetilde{O}(N)$ quantum gates, matching the theoretical lower bound. To realize this, we overcome the primary bottleneck of reversible Reed-Solomon decoding by introducing novel quantum circuits for the Extended Euclidean Algorithm (EEA) that reduce the leading-order space complexity to the theoretical minimum of $2nb$ qubits. These improvements are broadly applicable, including to Shor's algorithm for the discrete logarithm. We analyze OPI over binary extension fields $\mathrm{GF}(2^b)$, assess hardness against new classical attacks, and identify resilient instances. Our resource estimates show that classically intractable OPI instances (requiring $> 10^{23}$ classical trials) can be solved with approximately 5.72 million Toffoli gates. This is roughly 1000 times fewer gates than required for factoring RSA-2048 and, remarkably, is also less than the leading interactive protocol for computational proof of quantumness, positioning DQI as a compelling candidate for practical, verifiable quantum advantage.

QEC I – Decoders

Interpretability of Neural Network Decoders for Fault-Tolerant Quantum Error Correction

Lukas Bödeker, Luc Kusters, Markus Müller

Forschungszentrum Jülich

Neural-network (NN) decoders are increasingly used in quantum error correction (QEC) due to their flexibility and strong empirical performance, yet their internal decision logic typically remains opaque. This lack of interpretability poses challenges for fault-tolerant quantum computing, where reliability and physical consistency are essential. We introduce a systematic interpretability framework for NN-based QEC decoders based on Shapley value attribution. By adapting feature-attribution methods from machine learning to the decoding setting, we quantify how individual syndrome and flag-qubit measurements contribute to specific correction decisions, both per instance and in aggregate, and resolve their temporal influence in recurrent architectures. We demonstrate the approach in a case study of a recurrent neural decoder for flag-based fault-tolerant error correction with the Steane [[7,1,3]] code. The analysis shows that the trained network assigns high importance to flag outcomes precisely when correlated errors are likely, and distinguishes transient measurement faults from persistent data errors through structured temporal attributions. At the same time, the method uncovers specific misattributions linked to elevated logical error rates in targeted noise regimes, revealing concrete circuit-level failure mechanisms that are not visible in aggregate performance metrics. Beyond explanation, interpretability serves as a diagnostic and design tool, guiding architectural refinements and providing an additional convergence criterion during training. Our results establish feature-attribution analysis as a practical method to validate, debug, and improve neural decoders, supporting their trustworthy integration into fault-tolerant quantum computing architectures.

Colour Codes Reach Surface Code Performance using Vibe Decoding

***Stergios Koutsioumpas*[1], *Tamas Noszko*[1], *Hasan Sayginel*[2], *Mark Webster*[2], *Joschka Roffe*[1]**

[1] University of Edinburgh
[2] University College London

Two-dimensional quantum colour codes hold significant promise for quantum error correction, offering advantages such as planar connectivity and low-overhead logical gates. Despite their theoretical appeal, the practical deployment of these codes faces challenges due to complex decoding requirements compared to surface codes. We introduce vibe decoding which, for the first time, brings colour code performance on par with the surface code under practical polynomial-time decoding with parallel sub-routines that can be implemented on specialised hardware. We will discuss the accuracy and speed of Vibe decoding on multiple settings and future work.

Surprisingly useful local decoders for topological codes

Nathaniel Selub

University of California, Berkeley

We show that quantum error correction in a wide class of stabilizer codes can be performed using simple autonomous local dynamics, with no reliance on global control. We present a framework for implementing this approach using anisotropic cellular automaton rules. We give rigorous threshold proofs, resolve some old questions about the functioning of cellular automaton decoders, and show that our schemes achieve performance competitive with state-of-the-art global decoders. Our schemes apply to a broad range of topological codes, from the surface code to various models of fractons.

Algorithms II – Quantum Chemistry

A comprehensive framework to simulate real-time chemical dynamics on a fault-tolerant quantum computer

***Karthik Seetharam*[1], *Matteo Lostaglio*[1], *Burak Sahinoglu*[1], *Sam Pallister*[1], *Felipe Jornada*[2]**

[1] PsiQuantum
[2] Stanford University

We present a comprehensive end-to-end framework for simulating the real-time dynamics of chemical systems on a fault-tolerant quantum computer, incorporating both electronic and nuclear quantum degrees of freedom. An all-particle simulation is nominally efficient on a quantum computer, but practically infeasible. Hence, central to our approach is the construction of a first-quantized plane-wave algorithm making use of *pseudoions*. The latter consolidate chemically inactive electrons and the nucleus into a single effective dynamical ionic entity, extending the well-established concept of pseudopotentials in quantum chemistry to a two-body interaction. We explicitly describe efficient quantum circuits for initial state preparation across all degrees of freedom, as well as for block-encoding the Hamiltonian describing interacting pseudoions and chemically active electrons, by leveraging recent advances in quantum rejection sampling to optimize the implementations. To extract useful chemical information, we first design molecular fingerprints by combining density-functional calculations with machine learning techniques, and subsequently validate them through surrogate classical molecular dynamics simulations. These fingerprints are then coherently encoded on a quantum computer for efficient molecular identification via amplitude estimation. We provide an extensive analysis of the cost of running the algorithm on a fault-tolerant quantum computer for several chemically interesting systems. As an illustration, simulating the interaction between NH_3 and BF_3 (a 40-particle system) requires 808 logical qubits to encode the problem, and approximately 10^{11} Toffoli gates per femtosecond of time evolution. Our results establish a foundation for further quantum algorithm development targeting chemical and material dynamics.

Faster Quantum Chemistry Simulations on a Quantum Computer with Improved Tensor Factorization and Active Volume Compilation

***William Pol*[1], *Sukin Sim*[1], *Mark Steudtner*[1], *Athena Caesura*[1], *Cristian Cortes*[1], *Gian-Luca R. Anselmetti*[2], *Matthias Degroote*[2], *Nikolaj Moll*[2], *Raffaele Santagati*[2], *Michael Streif*[2], *Christofer S. Tautermann*[2]**

[1] PsiQuantum

[2] Boehringer Ingelheim

Electronic structure calculations of molecular systems are among the most promising applications for fault-tolerant quantum computing (FTQC) in quantum chemistry and drug design. However, while recent algorithmic advancements such as qubitization and tensor hypercontraction (THC) have significantly reduced the complexity of such calculations, they do not yet achieve computational runtimes short enough to be practical for industrially relevant use cases. In this work, we combine several advances to electronic structure calculation for molecular systems, resulting in a 2-orders-of-magnitude speedup of estimated runtimes over prior-art algorithms run on comparable quantum devices. One of these advances is a novel framework for block-invariant symmetry-shifted tensor hypercontraction, with which we achieve the tightest Hamiltonian factorizations reported to date. We compile our algorithm for an active volume (AV) architecture, a technical layout that has recently been proposed for fusion-based photonic quantum hardware. AV compilation contributes towards a lower runtime of our computation by eliminating overheads stemming from connectivity issues in the underlying surface code. We present a detailed benchmark of our approach, focusing primarily on the computationally challenging benchmark molecule P450. Leveraging a number of hardware trade-offs in interleaving-based photonic FTQC, we estimate runtimes for the electronic structure calculation of P450 as a function of the device footprint.

Fullqubit alchemist: Quantum algorithm for alchemical free energy calculations

***Po-Wei Huang*[1,2], *Gregory Boyd*[2], *Gian-Luca R. Anselmetti*[3], *Matthias Degroote*[3], *Nikolaj Moll*[3], *Raffaele Santagati*[3], *Michael Streif*[3], *Benjamin Ries*[3], *Daniel Marti-Dafcik*[2], *Hamza Jnane*[1,2], *Sophia Simon*[4], *Nathan Wiebe*[4], *Thomas R. Bromley*[2], *Bálint Koczor*[1,2]**

[1] University of Oxford
[2] Quantum Motion
[3] Boehringer Ingelheim
[4] University of Toronto

Accurately computing the free energies of biological processes is a cornerstone of computer-aided drug design, but it is a daunting task. The need to sample vast conformational spaces and account for entropic contributions makes the estimation of binding free energies very expensive. While classical methods, such as thermodynamic integration and alchemical free energy calculations, have significantly contributed to reducing computational costs, they still face limitations in terms of efficiency and scalability. We tackle this through a quantum algorithm for the estimation of free energy differences by adapting the existing Liouvillian approach and introducing several key algorithmic improvements. We directly implement the Liouvillian operator and provide an efficient description of electronic forces acting on both nuclear and electronic particles on the quantum ground state potential energy surface. This leads to super-polynomial runtime scaling improvements in the precision of our Liouvillian simulation approach and quadratic improvements in the scaling with the number of particles relative to prior quantum algorithms. Second, our algorithm calculates free energy differences via a fully quantum implementation of thermodynamic integration and alchemy, thereby foregoing expensive entropy estimation subroutines used in prior works. Our results open new avenues towards the application of quantum computers in drug discovery.

Testing and Verification

Composable Verification in the Circuit-Model via Magic-Blindness

Sami Abdul-Sater[1,2,3]***, Harold Ollivier***[1,2,3]

[1] ENS-PSL
[2] INRIA
[3] CNRS

As quantum computing machines move towards the utility regime, it is essential that users are able to verify their delegated quantum computations with security guarantees that are (i) robust to noise (ii) composable with other secure protocols and (iii) exponentially stronger as the number of resources dedicated to security increases. Previous works that achieve these guarantees are expressed in the Measurement-Based Quantum Computation (MBQC) model and benefit from a modular framework of verification protocols. This leaves architectures based on the circuit model—in particular those using the Magic State Injection (MSI)—with fewer options to verify their computations or with the need to compile their circuits in MBQC which leads to overheads. This paper introduces a family of noise robust, composable and efficient verification protocols for Clifford + MSI circuits that are secure against arbitrary malicious behavior. This family contains the verification protocol of Broadbent (2018, ToC), extends its security guarantees while also bridging the modularity gap between protocols for MBQC and those for the circuit model, and reducing quantum communication costs. As a result, it opens the prospect of rapid implementation tailored to near-term quantum devices. Our technique is based on a refined notion of blindness, called magic-blindness, which hides only the injected magic states—the sole source of non-Clifford computational power. This enables verification by randomly interleaving computation rounds with classically simulable, magic-free test rounds, leading to a trap-based framework for circuit verification. As a result, circuit-based quantum verification attains the same level of security and robustness previously known only in MBQC. It also reduces the quantum communication cost as transmitted qubits are required only at the locations of state injection.

Clifford testing: algorithms and lower bounds

Marcel Hinsche[1], Zongbo Bao[2], Philippe van Dordrecht[2], Jens Eisert[1], Jop Briët[2], Jonas Helsen[2]

[1] Freie Universität Berlin
[2] QuSoft and Centrum Wiskunde & Informatica (CWI)

We consider the problem of Clifford testing, which asks whether a black-box n-qubit unitary is a Clifford unitary or at least ε-far from every Clifford unitary. We give the first 4-query Clifford tester, which decides this problem with probability $\mathrm{poly}(\varepsilon)$. This contrasts with the minimum of 6 copies required for the closely-related task of stabilizer testing. We show that our tester is tolerant, by adapting techniques from tolerant stabilizer testing to our setting. In doing so, we settle in the positive a conjecture of Bu, Gu and Jaffe, by proving a polynomial inverse theorem for a non-commutative Gowers 3-uniformity norm. We also consider the restricted setting of single-copy access, where we give an $O(n)$-query Clifford tester that requires no auxiliary memory qubits or adaptivity. We complement this with a lower bound, proving that any such, potentially adaptive, single-copy algorithm needs at least $\Omega(n^{1/4})$ queries. To obtain our results, we leverage the structure of the commutant of the Clifford group, obtaining several technical statements that may be of independent interest.

Is it Gaussian? Testing bosonic quantum states

***Filippo Girardi*[1], *Freek Witteveen*[2], *Francesco Anna Mele*[1], *Lennart Bittel*[3], *Salvatore Francesco Emanuele Oliviero*[1], *David Gross*[4], *Michael Walter*[5]**

[1] Scuola Normale Superiore di Pisa
[2] QuSoft and Centrum Wiskunde & Informatica
[3] Freie Universität Berlin
[4] University of Cologne
[5] LMU Munich

Gaussian states are widely regarded as one of the most relevant classes of continuous-variable (CV) quantum states, as they naturally arise in physical systems and play a key role in quantum technologies. This motivates a fundamental question: given copies of an unknown CV state, how can we efficiently test whether it is Gaussian? We address this problem from the perspective of representation theory and quantum learning theory, characterizing the sample complexity of Gaussianity testing as a function of the number of modes. For pure states, we prove that just a constant number of copies is sufficient to decide whether the state is exactly Gaussian. We then extend this to the tolerant setting, showing that a polynomial number of copies suffices to distinguish states that are close to Gaussian from those that are far. In contrast, we establish that testing Gaussianity of general mixed states necessarily requires exponentially many copies, thereby identifying a fundamental limitation in testing CV systems. Our approach relies on rotation-invariant symmetries of Gaussian states together with the recently introduced toolbox of CV trace-distance bounds.

Learning I

Cloning is as Hard as Learning for Stabilizer States

***Nikhil Bansal*[1], *Matthias C. Caro*[1], *Gaurav Mahajan*[2]**

[1] University of Warwick
[2] Yale University

The impossibility of simultaneously cloning non-orthogonal states lies at the foundations of quantum theory. Even when allowing for approximation errors, cloning an arbitrary unknown pure state requires as many initial copies as needed to fully learn the state. Rather than arbitrary unknown states, modern quantum learning theory often considers structured classes of states and exploits such structure to develop learning algorithms that outperform general-state tomography. This raises the question: How do the sample complexities of learning and cloning relate for such structured classes? We answer this question for an important class of states. Namely, for n-qubit stabilizer states, we show that the optimal sample complexity of cloning is $\Theta(n)$. Thus, also for this structured class of states, cloning is as hard as learning. To prove these results, we use representation-theoretic tools in the recently proposed Abelian State Hidden Subgroup framework and a new structured version of the recently introduced random purification channel to relate stabilizer state cloning to a variant of the sample amplification problem for probability distributions that was recently introduced in classical learning theory. This allows us to obtain our cloning lower bounds by proving new sample amplification lower bounds for classes of distributions with an underlying linear structure. Our results provide a more fine-grained perspective on No-Cloning theorems, opening up connections from foundations to quantum learning theory and quantum cryptography.

Learning T-conjugated stabilizers: The multiple-squares dihedral StateHSP

***Gideon Lee*[1], *Jonathan A Gross*[2], *Masaya Fukami*[2], *Zhang Jiang*[2]**

[1] University of Chicago
[2] Google Quantum AI

The state hidden subgroup problem (StateHSP) is a recent generalization of the hidden subgroup problem. We present an algorithm that solves the non-abelian StateHSP over N copies of the dihedral group of order 8 (the symmetries of a square). This algorithm is of interest for learning non-Pauli stabilizers, as well as related symmetries relevant for the problem of Hamiltonian spectroscopy. Our algorithm is polynomial in the number of samples and computational time, and requires only constant depth circuits. This result extends previous work on the abelian StateHSP and, as a special case, provides a solution for the ordinary hidden subgroup problem on this specific non-abelian group.

Optimal randomized measurements for a family of non-linear quantum properties

Zhenyu Du[1], Yifan Tang[2], Andreas Elben[3], Ingo Roth[4], Jens Eisert[2], Zhenhuan Liu[1]

[1] Tsinghua University
[2] Freie Universität Berlin
[3] Paul Scherrer Institute
[4] Technology Innovation Institute

Quantum learning encounters fundamental challenges when estimating non-linear properties, owing to the inherent linearity of quantum mechanics. Although recent advances in single-copy randomized measurement protocols have achieved optimal sample complexity for specific tasks like state purity estimation, generalizing these protocols to estimate broader classes of non-linear properties without sacrificing optimality remains an open problem. In this work, we introduce the observable-driven randomized measurement (ORM) protocol enabling the estimation of $\mathrm{Tr}(O\rho^2)$ for an arbitrary observable O—an essential quantity in quantum computing and many-body physics. We establish an upper bound for ORM's sample complexity and show its optimality for observables with a large trace-norm, including Pauli and local observables, closing a gap in the literature. For these observables, ORM admits an efficient implementation with Clifford circuits. Numerical experiments validate that ORM requires substantially fewer state samples to achieve the same precision compared to classical shadows. Additionally, we introduce a braiding randomized measurement protocol for multiple low-rank non-linear observables, reducing circuit complexities in practical applications.

ALGORITHMS III – GENERAL

State-to-Hamiltonian conversion with a few copies

Kaito Wada, Jumpei Kato, Hiroyuki Harada, Naoki Yamamoto

Keio University

Density matrix exponentiation (DME) is a general procedure that converts an unknown quantum state into the Hamiltonian evolution. This enables state-dependent operations and can reveal nontrivial properties of the state, among other applications, without full tomography. However, it has been proven that for any physical process, the DME requires $\Theta(1/\varepsilon)$ state copies in error ε. In this work, we go beyond the lower bound and propose a procedure called the *virtual* DME that achieves $\tilde{O}(\log(1/\varepsilon))$ or $\tilde{O}(1)$ state copies, by using non-physical processes. Using the virtual DME in place of its conventional counterpart realizes a general-purpose quantum algorithm for property estimation, that achieves *exponential* circuit-depth reductions over existing protocols across tasks including quantum principal component analysis, quantum emulator, calculation of nonlinear functions such as entropy, and linear system solver with quantum precomputation. In such quantum algorithms, the non-physical process for virtual DME can be effectively simulated via simple classical post-processing while retaining a near-unity measurement overhead. We numerically verify this small constant overhead together with the exponential reduction of copy count in the quantum principal component analysis task. The number of state copies used in our algorithm essentially saturates the theoretical lower bound we proved.

Randomized Truncation for Quantum State Preparation and Series-Truncated Algorithms

***Yue Wang*[1], *Xiao-Ming Zhang*[2], *Xiao Yuan*[3], *Qi Zhao*[1]**

[1] The University of Hong Kong
[2] School of Physics, South China Normal University
[3] Peking University

Quantum algorithms promise transformative speedups, yet their practical impact is often limited by prohibitively deep circuits, with two recurring sources of depth: costly preparation of structured input states and the execution of subroutines based on truncated series approximations. We introduce a unified, resource-efficient paradigm that uses classical randomness to improve accuracy–depth trade-offs. For realistic states with hierarchical amplitude structure, we propose a randomized state-preparation protocol that probabilistically amplifies small amplitudes using ensembles of low-complexity circuits, reducing the number of amplitudes that must be encoded and achieving up to 99% reductions in CNOT and T-gate counts in simulations on LiH wavefunctions and power-law decay states. Target quantum algorithms, we develop Randomized Truncated Series, a generic acceleration principle for any quantum algorithm built from truncated series, which quadratically suppresses truncation error while enabling continuous tuning of the effective truncation order via random mixing of shallow circuits. Together, these results broaden the regime where end-to-end quantum advantage is feasible on near-term and early fault-tolerant hardware.

Quantum algorithms through graph composition

Arjan Cornelissen

Simons Institute

In this work, we unify several quantum algorithmic frameworks for boolean functions that are based on the quantum adversary bound. First, we show that the st-connectivity framework subsumes the (adaptive/extended) learning graph framework, and the weighted-decision-tree framework. Additionally, we show that every randomized algorithm can be turned into an st-connectivity problem as well, with the same complexity up to constants. This situates the st-connectivity framework in between randomized and quantum algorithms, indicating that it's an intermediate computational model between classical and quantum. We also introduce a generalization of the st-connectivity framework, the graph composition framework, and show that it subsumes part of the quantum divide and conquer framework, and it is itself subsumed by the multidimensional quantum walk framework. Second, we investigate these frameworks' power by investigating the most efficient algorithms they can produce, in terms of the number of queries to the input. We show that the weighted-decision-tree framework's power is polynomially related to deterministic query complexity, showing that the quantum speed-ups that can be obtained with this framework are at most quadratic. Finally, we turn our attention to time-efficient implementations of the algorithms constructed through the st-connectivity and graph composition frameworks. To that end, we convert instances to the two-subspace phase estimation framework, and we show how we can implement these as transducers. This has the added benefit of removing the effective spectral gap lemma from the construction, significantly simplifying the analysis. We showcase the techniques developed in this work to give improved algorithms for various string search problems.

Contributed Talks on Tuesday

Many-Body Quantum Physics and Information

Computational complexity of Berry phase estimation in topological phases of matter

***Ryu Hayakawa*[1,2], *Kazuki Sakamoto*[3], *Chusei Kiumi*[4]**

[1] Yukawa Institute for Theoretical Physics
[2] The Hakubi Center, Kyoto University
[3] Graduate School of Engineering Science, The University of Osaka
[4] Center for Quantum Information and Quantum Biology, The University of Osaka

The Berry phase is a fundamental quantity for classifying topological phases of matter. We present a new quantum algorithm for Berry phase estimation (BPE) that is both more general than previously known approaches and comes with a rigorous polynomial-time performance guarantee. Moreover, we provide a new circuit-to-Hamiltonian construction that results in a closed loop of parameterized Hamiltonians. Building on these, we prove that a BPE formulation is BQP-complete when given a guiding state with large overlap with the ground state. This shows the first complexity-theoretic evidence of an exponential quantum speedup for quantum-computational approaches to studying topological phases of matter. We also establish several complexity-theoretic results for BPE, including dUQMA-completeness, $\mathrm{P}^{\mathrm{dUQMA}[\log]}$-hardness, and containment in $\mathrm{P}^{\mathrm{PGQMA}[\log]}$, depending on the BPE setting. Here, dUQMA is a variant of the unique-witness class UQMA that we introduce and remarkably, this dUQMA-complete BPE variant appears to be the first natural problem known to lie in UQMA ∩ co-UQMA.

Correcting and extending Trotterized quantum many-body dynamics

Gian Gentinetta, Friederike Metz, Giuseppe Carleo

EPFL

A complex but important challenge in understanding quantum mechanical phenomena is the simulation of quantum many-body dynamics. Although quantum computers offer significant potential to accelerate these simulations, their practical application is currently limited by noise and restricted scalability. In this work, we address these problems by introducing a hybrid ansatz combining the strengths of quantum and classical computational methods. Using Trotterization, we propose to evolve an initial state on the quantum computer according to a simplified Hamiltonian, focusing on terms that are difficult to simulate classically. A classical model then corrects the simulation by including the terms omitted in the quantum circuit. While the classical ansatz is optimized during the time evolution, the quantum circuit has no variational parameters. Derivatives can thus be calculated purely classically, avoiding challenges arising in the optimization of parameterized quantum circuits. We demonstrate three applications of this hybrid method. First, our approach allows us to avoid SWAP gates in the quantum circuit by restricting the quantum part of the ansatz to hardware-efficient terms of the Hamiltonian. Second, we can mitigate errors arising from the Trotterization of the time evolution unitary. Finally, we can extend the system size while keeping the number of qubits in the quantum circuit constant by including additional degrees of freedom in the classical ansatz.

Rapid Mixing of Quantum Gibbs Samplers for Weakly-Interacting Quantum Systems

***Štěpán Šmíd*[1], *Richard Meister*[1], *Mario Berta*[2], *Roberto Bondesan*[1]**

[1] Imperial College London
[2] RWTH Aachen University

Dissipative quantum algorithms for state preparation in many-body systems are increasingly recognised as promising candidates for achieving large quantum advantages in application-relevant tasks. Recent advances in algorithmic, detailed-balance Lindbladians enable the efficient simulation of open-system dynamics converging towards desired target states. However, the overall complexity of such schemes is governed by system-size dependent mixing times. In this work, we analyse algorithmic Lindbladians for Gibbs state preparation and prove that they exhibit rapid mixing, i.e., convergence in time poly-logarithmic in the system size. We first establish this for non-interacting spin systems, free fermions, and free bosons, and then show that these rapid mixing results are stable under perturbations, covering weakly interacting qudits and perturbed non-hopping fermions. Further, we adapt the techniques from separable qudits to the fermionic setting and prove rapid mixing of the strongly-interacting regime of the Fermi-Hubbard model. Our results constitute the first efficient mixing bounds for non-commuting qudit models and bosonic systems at arbitrary temperatures. Compared to prior spectral-gap-based results for fermions, we achieve exponentially faster mixing, further featuring explicit constants on the maximal allowed interaction strength. This not only improves the overall polynomial runtime for quantum Gibbs state preparation, but also enhances robustness against noise. Our analysis relies on oscillator norm techniques from mathematical physics, where we introduce tailored variants adapted to specific Lindbladians — an innovation that we expect to significantly broaden the scope of these methods.

Compilation and Resource Estimation

Characterizing Space Requirements for Quantum Computations via Signaling Conditions

***Kosuke Matsui*[1,2], *Jun-Yi Wu*[2,3], *Hayata Yamasaki*[1], *Min-Hsiu Hsieh*[2], *Mio Murao*[1]**

[1] University of Tokyo
[2] Hon Hai Research Institute
[3] Tamkang University

Scaling up the number of qubits available on quantum processors remains technically demanding; it is therefore crucial to clarify the number of qubits required to execute a quantum computation. When circuit compilation techniques such as mid-circuit measurements and delayed input preparation are permitted, the qubit requirement of a given quantum computation can be smaller than that implied by its naive description. However, no general method was known for characterizing how much such reductions can be achieved. In this work, we characterize lower and upper bounds on the number of qubits required to implement a given quantum computation in terms of the causal structure of the corresponding quantum instrument. We further show that these lower and upper bounds coincide for quantum computations used in entanglement distillation protocols, and thereby obtain optimal space requirements for several well-known entanglement distillation protocols.

Compiling Quantum Regular Language States

***Armando Bellante*[1,2], *Reinis Irmejs*[1,2], *Marta Florido-Llinàs*[1,2], *María Cea Fernández*[1,2], *Marianna Crupi*[1,2], *Matthew Kiser*[3,4], *J. Ignacio Cirac*[1,2]**

[1] Max Planck Institute of Quantum Optics
[2] Munich Center for Quantum Science and Technology (MCQST)
[3] TUM School of Natural Sciences
[4] IQM Quantum Computers

State preparation compilers for quantum computers typically sit at two extremes: general-purpose routines that treat the target as an opaque amplitude vector, and bespoke constructions for a handful of well-known state families. We ask whether a compiler can instead accept simple, structure-aware specifications while providing predictable resource guarantees. We answer this by designing and implementing a quantum state-preparation compiler for regular language states (RLS): uniform superpositions over bitstrings accepted by a regular description, and their complements. Users describe the target state via (i) a finite set of bitstrings, (ii) a regular expression, or (iii) a deterministic finite automaton (DFA), optionally with a complement flag. By translating the input to a DFA, minimizing it, and mapping it to an optimal matrix product state (MPS), the compiler obtains an intermediate representation (IR) that exposes and compresses hidden structure. The efficient DFA representation and minimization offloads expensive linear algebra computation in exchange of simpler automata manipulations. The combination of the regular-language frontend and this IR gives concise specifications not only for RLS but also for their complements that might otherwise require exponentially large state descriptions. This enables state preparation of an RLS or its complement with the same asymptotic resources and compile time, which to our knowledge is not supported by existing compilers. We outline two hardware-aware backends: SeqRLSP, which yields linear-depth, ancilla-free circuits for linear nearest-neighbor architectures via sequential generation, and TreeRLSP, which achieves logarithmic depth on all-to-all connectivity via a tree tensor network. On the theory side, we prove circuit-depth and gate-count bounds that scale with the system size and the maximal Schmidt rank of the target state, and we give compile-time bounds that expose the benefit of the initial DFA representation. We implement the full pipeline and evaluate it on Dicke and W states, random uniform superpositions, and complement states, comparing against general-purpose, sparse-state, and specialized baselines.

The FLuid Allocation of Surface code Qubits (FLASQ) cost model for early fault-tolerant quantum algorithms

William J. Huggins, Tanuj Khattar, Amanda Xu, Matthew Harrigan, Christopher Kang, Guang Hao Low, Austin Fowler, Nicholas C. Rubin, Ryan Babbush

Google Quantum AI

Holistic resource estimates are essential for guiding the development of fault-tolerant quantum algorithms and the computers they will run on. This is particularly true when we focus on highly constrained early fault-tolerant devices. Many attempts to optimize algorithms for early fault-tolerance focus on simple metrics, such as the circuit depth or T-count. These metrics fail to capture critical overheads, such as the spacetime cost of Clifford operations and routing, or miss the key optimizations. We propose the FLuid Allocation of Surface code Qubits (FLASQ) cost model, tailored for architectures that use a two-dimensional lattice of locally connected qubits to implement the two-dimensional surface code. FLASQ abstracts away the complexity of routing by assuming that ancilla space and time can be fluidly rearranged, allowing for the tractable estimation of spacetime volume while still capturing important details neglected by simpler approaches. At the same time, it enforces constraints imposed by the circuit's measurement depth and the processor's reaction time. We apply FLASQ to analyze the cost of a standard two-dimensional lattice model simulation, finding that modern advances (such as magic state cultivation and the combination of quantum error correction and mitigation) reduce both the time and space required for this task by an order of magnitude compared with previous estimates. We also analyze the Hamming weight phasing approach to synthesizing parallel rotations, revealing that despite its low T-count, the overhead from imposing a 2D layout and from its use of additional ancilla qubits will make it challenging to benefit from in early fault-tolerance. We hope that the FLASQ cost model will help to better align early fault-tolerant algorithmic design with actual hardware realization costs without demanding excessive knowledge of quantum error correction from quantum algorithmists.

Algorithms IV – Dynamics

Quantum algorithms for general nonlinear dynamics based on the Carleman embedding

***David Jennings*[1], *Kamil Korzekwa*[1], *Matteo Lostaglio*[1], *Andrew T Sornborger*[2], *Yigit Subasi*[2], *Guoming Wang*[1]**

[1] PsiQuantum
[2] Los Alamos National Laboratory

Important nonlinear dynamics, such as those found in plasma and fluid systems, are typically hard to simulate on classical computers. In a recent breakthrough [Liu et al., PNAS 2021], the first efficient quantum algorithm for solving nonlinear differential equations was constructed, based on a single condition $R < 1$, where R characterizes the ratio of nonlinearity to dissipation. This result, however, is limited to the class of purely dissipative systems with negative log-norm, which excludes application to many important problems. In this work, we correct technical issues with this and other prior analysis, and substantially extend the scope of nonlinear dynamical systems that can be efficiently simulated on a quantum computer in a number of ways. Firstly, we extend the existing results from purely dissipative systems to a much broader class of stable systems, and show that every quadratic Lyapunov observable for the linearized system corresponds to an independent R-number criterion for the convergence of the Carleman scheme. Secondly, we extend our stable system results to physically relevant settings where conserved polynomial quantities exist. Finally, we provide extensive results for the class of non-resonant systems. With this, we are able to show that efficient quantum algorithms exist for a much wider class of nonlinear systems than previously known, and prove the BQP-completeness of nonlinear oscillator problems of exponential size. In our analysis, we also obtain several results related to the Poincaré-Dulac theorem and diagonalization of the Carleman matrix, which could be of independent interest.

An end-to-end quantum algorithm for nonlinear fluid dynamics with bounded quantum advantage

David Jennings[1], Kamil Korzekwa[1], Matteo Lostaglio[1], Paul Mannix[1], Richard Ashworth[2], Emanuele Marsili[2], Stephen Rolston[2]

[1] PsiQuantum
[2] Airbus

Computational fluid dynamics (CFD) is a cornerstone of classical scientific computing, and there is growing interest in whether quantum computers can accelerate such simulations. In this work, we first show that the existing proposals for fault-tolerant quantum algorithms for CFD suffer from a range of severe bottlenecks that negate conjectured quantum advantages. We then develop a novel algorithm for the incompressible lattice Boltzmann equation that circumvents these obstacles, and provide a detailed analysis of our algorithm, including all potential sources of algorithmic complexity, as well as gate count estimates. We find that for an end-to-end problem, a modest quantum advantage may be preserved for selected observables in the high-error-tolerance regime. We lower bound the Reynolds number scaling of our quantum algorithm in dimension D at Kolmogorov microscale resolution with $O(\mathrm{Re}^{\frac{3}{4}(1+\frac{D}{2})} \times q_M)$, where q_M is a multiplicative overhead for data extraction with $q_M = O(\mathrm{Re}^{\frac{3}{8}})$ for the drag force. This upper bounds the scaling improvement over classical algorithms by $O(\mathrm{Re}^{\frac{3D}{8}})$. However, our numerical investigations suggest a lower speedup, with a scaling estimate of $O(\mathrm{Re}^{1.936} \times q_M)$ for $D = 2$. Finally, to support our theoretical analysis, we provide a classical numerical study illustrating the accuracy, complexity, and convergence of the algorithm for representative incompressible-flow cases, including the driven Taylor–Green vortex, the lid-driven cavity flow, and the flow past a cylinder. Our results give robust evidence that small, but nontrivial, quantum advantages can be achieved in the context of CFD, and motivate the need for additional rigorous end-to-end quantum algorithm development.

Optimized quantum algorithms for simulating the Schwinger effect

Angus Kan, Jessica Lemieux, Burak Şahinoğlu, Olga Okrut

PsiQuantum

The Schwinger model, which describes lattice quantum electrodynamics in $1+1$ space-time dimensions, provides a valuable framework to investigate fundamental aspects of quantum field theory, and a stepping stone towards non-Abelian gauge theories. Specifically, it enables the study of physically relevant dynamical processes, such as the nonperturbative particle-antiparticle pair production, known as the Schwinger effect. In this work, we analyze the quantum computational resource requirements associated with simulating the Schwinger effect under two distinct scenarios: (1) a quench process, where the initial state is a simple product state of a non-interacting theory and then interactions are turned on at time $t = 0$, and (2) a splitting (or scattering) process where two Gaussian states, peaked at given initial momenta, are shot away from (or towards) each other. We explore different physical regimes in which the Schwinger effect is expected to be observable. These regimes are characterized by initial momenta and coupling strengths, as well as simulation parameters such as lattice size and electric-field cutoffs. Leveraging known rigorous bounds for electric-field cutoffs, we find that a reliable simulation of the Schwinger effect is provably possible at high cutoff scales. Furthermore, we provide optimized circuit implementations of both the second-order Trotter formula and an interaction-picture algorithm based on the Dyson series to implement the time evolution. Our detailed resource estimates show the regimes in which the interaction-picture approach outperforms the Trotter approach, and vice versa. The improved theoretical error bounds, optimized quantum circuit designs, and explicitly compiled subroutines developed in this study are broadly applicable to simulations of other lattice models in high-energy physics and beyond.

QEC II - Topological Codes

Rethinking Lattice Surgery Compilation: Diverse Topological Codes and Movable Logical Qubits

***Laura S. Herzog*[1], *Lucas Berent*[2], *Aleksander Kubica*[3], *Robert Wille*[1]**

[1] Technical University Munich
[2] Iceberg Quantum
[3] Yale University

Fault-tolerant quantum computation (FTQC) requires compiling logical quantum circuits encoded using a quantum error-correcting code into physical operations tailored to specific hardware architectures. Current lattice surgery based compilation techniques are limited due to certain paradigmatic assumptions that are widely regarded as standard. In our works we identify and address two of these limiting paradigms. First, prior work has *predominantly focused on the surface code*, even though other topological codes offer certain advantages. Second, compilation schemes usually follow a *place-and-route paradigm* where logical qubits remain fixed in space throughout the computation. We initiate a more flexible line of work that goes beyond both aforementioned paradigms.

Transversal surface-code game with reconfigurable qubits

***Shinichi Sunami*[1,2], *Akihisa Goban*[1], *Hayata Yamasaki*[1,3]**

[1] NanoQT
[2] University of Oxford
[3] University of Tokyo

Physical platforms supporting qubit reconfigurability have enabled novel approaches to realizing fault-tolerant quantum computation (FTQC), offering a promising alternative to lattice surgery through transversal gates. Yet a critical gap remains between FTQC theory and its experimental realization: in reconfigurable-qubit platforms, system-size-dependent qubit transport prevents meeting the key requirement that syndrome extraction be performed frequently enough to keep error accumulation below a threshold. Here, focusing on surface-code implementation, we close this gap by formulating the transversal-gate FTQC protocol as a "transversal surface-code game" realizable with demonstrated technological capabilities. The game's rules governing mode transitions of code cells ensure compatibility with the threshold theorem at scale and enable extensive rule-based protocol optimization. Furthermore, resource estimation for neutral-atom platforms shows that this approach enhances the space-time performance of FTQC with reconfigurable qubits. These results establish a solid foundation bridging FTQC theory and experiment, charting a compelling path toward scalable quantum computation with reconfigurable qubits.

Logical gates on Floquet codes via folds and twists

Alexandra E. Moylett*, *Bhargavi Jonnadula

Nu Quantum Ltd.

Floquet codes have recently emerged as a new family of error-correcting codes, and have drawn significant interest across both theoretical and practical quantum computing. A central open question has been how to implement logical operations on these codes. In this work, we show how two techniques from static quantum error-correcting codes can also be implemented on Floquet codes. First, we present a way of implementing fold-transversal operations on Floquet codes in order to yield logical Hadamard and S gates. And second, we present a way of implementing logical CNOT gates on Floquet codes via Dehn twists. We discuss the requirements for these techniques, and show that they are applicable to a wide family of Floquet codes defined on colour code lattices. Through numerical benchmarking of the logical operations on the CCS Floquet code, we establish a logical-gate threshold of 0.25-0.35% and verify sub-threshold exponential error suppression. Our results show that these logical operations are robust, featuring a performance that is close to the baseline set by a quantum memory benchmark. Finally, we explain in detail how to implement logical gates on Floquet codes by operating on the embedded codes.

Learning II

Shedding light on classical shadows: learning photonic quantum states

***Hugo Thomas*[1,2,3], *Ulysse Chabaud*[3], *Pierre-Emmanuel Emeriau*[1]**

[1] Quandela
[2] LIP6
[3] DIENS

Efficient learning of quantum state properties is both a fundamental and practical problem in quantum information theory. Classical shadows have emerged as an efficient method for estimating properties of unknown quantum states, with rigorous statistical guarantees, by performing randomized measurement on a few number of copies. With the advent of photonic technologies, formulating efficient learning algorithms for such platforms comes out as a natural problem. Here, we introduce a classical shadow protocol for learning photonic quantum states via randomized passive linear optical transformations and photon-number measurement. We show that this scheme is efficient for a large class of observables of interest. We experimentally demonstrate our findings on both a twelve-mode and a twenty-four-mode photonic integrated quantum processing unit. Our protocol allows for scalable learning of a wide range of photonic state properties and paves the way to applying the already rich variety of applications of classical shadows to photonic platforms.

Polynomial Speed-Up in Photonic Neural Networks via Adaptive State Injection

***Léo Monbroussou*[1,2], *Eliott Z. Mamon*[3], *Hugo Thomas*[3,4,5], *Verena Yacoub*[3], *Elham Kashefi*[1,5]**

[1] School of Informatics, University of Edinburgh
[2] Terra Quantum AG
[3] LIP6 (La Sorbonne Université)
[4] Quandela
[5] INRIA
[6] CNRS

Quantum Machine Learning (QML) has become a promising area for real world applications of quantum computers, and near-term methods and their scalability are still important research topics. A consequent amount of efforts has been put into understanding how to avoid Barren Plateaus (BPs), a vanishing gradient phenomenon that prevents the variational algorithms from being trained efficiently. In particular, evidence has recently been shown that the structures that allow us to avoid BP seem to allow classical simulation techniques. In addition, other important questions must be tackled to design near-term quantum algorithms that may offer an advantage. How to ensure that the performance of the algorithms will scale with input size, and how to compare classical and quantum algorithms on different figures of merit for a same use case? Recent works have proposed to use subspace preserving quantum circuits to mimic classical neural network architectures. By restricting the Hilbert space to a subspace of polynomial size with respect to the number of qubits, such architectures are likely to avoid BPs. This comes at the cost of *classical simulability* that is, a classical method can perform the same computation in polynomial time. In this work, we propose a paradigm shift: focusing on subspace-preserving methods that aim for a practical polynomial advantage. In particular, we propose to use linear optical circuits that are intrinsically subspace preserving as they conserve the number of particles during the computation. We believe that this approach could be sufficient to create useful QML applications as the generation of Fock states with few particles can be extremely high.

Do you know what q-means?

Alessandro Luongo[1], João Doriguello[2], Ewin Tang[3], Arjan Cornelissen[4]

[1] Center for Quantum Technologies, NUS
[2] HUN-REN Alfréd Rényi Institute of Mathematics
[3] Department of EECS, UC Berkeley
[4] Simons Institute, UC Berkeley

Clustering is one of the most important tools for analysis of large datasets, and perhaps the most popular clustering algorithm is Lloyd's algorithm for k-means. This algorithm takes n vectors $V = [v_1, \ldots, v_n] \in \mathbb{R}^{d \times n}$ and outputs k centroids $c_1, \ldots, c_k \in \mathbb{R}^d$; these partition the vectors into clusters based on which centroid is closest to a particular vector. We present a classical ε-k-means algorithm that performs an approximate version of one iteration of Lloyd's algorithm with time complexity $\widetilde{O}\big(\frac{\|V\|_F^2}{n}\frac{k^2 d}{\varepsilon^2}(k + \log n)\big)$, exponentially improving the dependence on the data size n and matching that of the "q-means" quantum algorithm originally proposed by Kerenidis, Landman, Luongo, and Prakash (NeurIPS'19). Moreover, we propose an improved q-means quantum algorithm with time complexity $\widetilde{O}\big(\frac{\|V\|_F}{\sqrt{n}}\frac{k^{3/2} d}{\varepsilon}(\sqrt{k} + \sqrt{d})(\sqrt{k} + \log n)\big)$ that quadratically improves the runtime of our classical ε-k-means algorithm in several parameters. Our quantum algorithm does not rely on quantum linear algebra primitives of prior work, but instead only uses QRAM to prepare simple states based on the current iteration's clusters and multivariate quantum mean estimation. Our upper bounds are complemented with classical and quantum query lower bounds, showing that our algorithms are optimal in most parameters. Finally, we conduct numerical experiments that evidence the substantially improved runtime our classical algorithm over the standard Lloyd's algorithm, thus being one of the first cases of a practical dequantised algorithm.

Benchmarking and Classical Simulation

Limitations on Measurement-free Fault-tolerant Protocols using Clifford Circuits

***Jon Nelson*[1], *Joel Rajakumar*[1], *Dominik Hangleiter*[1,2], *Michael Gullans*[1,3]**

[1] University of Maryland
[2] ETH Zürich
[3] NIST

When designing fault-tolerant protocols for quantum computers, it is important to understand which ingredients are necessary and which can be safely excluded. This is especially true in the so-called NISQ era, where near-term hardware lacks the ability to easily implement the full suite of quantum operations. For instance, non-Clifford gates, mid-circuit resets, and intermediate measurements can all be challenging to execute on near-term hardware. In this work, we examine the interplay between these resources, with the aim of understanding how they contribute to fault-tolerance. We show that in many regimes, we can rigorously rule out forms of measurement-free fault-tolerance in all Clifford circuits. Further, in some regimes not covered by our no-go results, we also provide evidence that random Clifford (or non-Clifford) circuits provide useful fault-tolerance properties.

A magic criterion (almost) as nice as PPT, with applications in distillation and detection

***Zhenhuan Liu*[1], *Tobias Haug*[2], *Qi Ye*[1], *Zi-Wen Liu*[1], *Ingo Roth*[2]**

[1] Tsinghua University
[2] Technology Innovation Institute

We introduce a mixed-state magic criterion, the Triangle Criterion, which plays a role for magic analogous to the Positive Partial Transposition (PPT) criterion for entanglement: it combines strong detection capability, a clear geometric interpretation, and an operational link to magic distillation. Using this criterion, we uncover several new features of multi-qubit magic distillation and detection. We prove that genuinely multi-qubit magic distillation protocols are strictly more powerful than all single-qubit schemes by showing that the Triangle Criterion is not stable under tensor products, in sharp contrast to the PPT criterion. Moreover, we show that, with overwhelming probability, multi-qubit magic states with relatively low rank cannot be distilled by any single-qubit distillation protocol. We derive an upper bound on the minimal purity of magic states, which is conjectured to be tight with both numerical and constructive evidences. Using this minimal-purity result, we predict the existence of unfaithful magic states, namely states that cannot be detected by any fidelity-based magic witness, and reveal fundamental limitations of mixed-state magic detection in any single-copy scheme.

Fail fast: techniques to probe rare events in quantum error correction

Michael Beverland, Malcolm Carroll, Andrew W. Cross, Ted Yoder

IBM

The ultimate goal of quantum error correction is to create logical qubits with very low error rates (e.g. 1e-12) and assemble them into large-scale quantum computers capable of performing many (e.g. billions) of logical gates on many (e.g. thousands) of logical qubits. However, it is necessarily difficult to directly assess the performance of such high-quality logical qubits using standard Monte Carlo sampling because logical failure events become very rare. Building on existing approaches to this problem, we develop three complementary techniques to characterize the rare-event regime for general quantum low-density parity-check (qLDPC) codes under circuit noise. (I) We propose a well-motivated, low-parameter ansatz for the failure spectrum (the fraction of fault sets of each size that fail) that empirically fits all the QEC systems we studied and predicts logical error rates at all physical error rates. (II) We find min-weight logical operators of syndrome measurement circuits and exactly compute the number of min-weight failing configurations. (III) We generalize the splitting method to qLDPC codes using multi-seeded Metropolis sampling to improve convergence for systems with many inequivalent logical operators. We apply these tools to distance-6, -12, and -18 bivariate bicycle codes under circuit noise, observing strong low-error-rate performance with the recently proposed Relay decoder but also considerable scope for further improvement.

Contributed Talks on Wednesday

Algorithms V – Hamiltonian Simulation

Quantum singular value transformation without block encodings

Shantanav Chakraborty[1], Soumyabrata Hazra[1], Tongyang Li[2], Changpeng Shao[3], Xinzhao Wang[2], Yuxin Zhang[3]

[1] IIIT Hyderabad
[2] Peking University
[3] Academy of Mathematics and Systems Science, Chinese Academy of Sciences

We develop new algorithms for Quantum Singular Value Transformation (QSVT), a unifying framework that encapsulates most known quantum algorithms and serves as the foundation for new ones. Existing implementations of QSVT rely on block encoding, incurring an intrinsic $O(\log L)$ ancilla overhead and circuit depth $\widetilde{O}(Ld\lambda)$ for polynomial transformations of a Hamiltonian $H = \sum_{k=1}^{L} H_k$, where d is the polynomial degree and $\lambda = \sum_k \|H_k\|$. We introduce a simple yet powerful approach that utilizes only basic Hamiltonian simulation techniques, namely, Trotter methods, to: (i) eliminate the need for block encoding, (ii) reduce the ancilla overhead to only a single qubit, and (iii) still maintain near-optimal complexity. Our method achieves a circuit depth of $\widetilde{O}(L(d\lambda_{\text{comm}})^{1+o(1)})$, without requiring any complicated multi-qubit controlled gates. Moreover, λ_{comm} depends on the nested commutators of the terms of H and can be substantially smaller than λ for many physically relevant Hamiltonians, a feature absent in standard QSVT. To achieve these results, we make use of Richardson extrapolation in a novel way, systematically eliminating errors in any interleaved sequence of arbitrary unitaries and Hamiltonian evolution operators, thereby establishing a general framework that encompasses QSVT but is more broadly applicable. We further design two randomized algorithms for QSVT in settings with only sampling access to the Hamiltonian terms. The first is a direct randomization of standard QSVT, while the second integrates qDRIFT within our interleaved-circuit architecture. Both achieve a complexity quadratic in d, which we establish as a lower bound for any randomized method implementing polynomial transformations in this model. Finally, as applications, we develop end-to-end quantum algorithms for solving linear systems and estimating ground state properties of Hamiltonians, both achieving near-optimal complexity without relying on oracular access. Overall, our results establish a new framework for quantum algorithms, significantly reducing hardware overhead while maintaining near-optimal performance, with implications for both near-term and fault-tolerant quantum computing.

Trotter error mitigation by error profiling with shallow quantum circuit

***Sangjin Lee*[1], *Youngseok Kim*[2], *Seung-Woo Lee*[3]**

[1] Korea Institute for Advanced Study (KIAS)
[2] Ohio State University
[3] Postech

Understanding the dynamics of quantum systems is crucial in many areas of physics, but simulating many-body systems presents significant challenges due to the large Hilbert space to navigate and the exponential growth of computational overhead. Quantum computers offer a promising platform to overcome these challenges, particularly for simulating the time evolution with Hamiltonians. Trotterization is a widely used approach among available algorithms in this regard, and well suited for near-term quantum devices. However, it introduces algorithmic Trotter errors due to the non-commutativity of Hamiltonian components. Several techniques such as multi-product formulas have been developed to mitigate Trotter errors, but often require deep quantum circuits, which can introduce additional physical errors. In this work, we propose a resource-efficient scheme to reduce the algorithmic Trotter error with relatively shallow circuit depth. We develop a profiling method by introducing an auxiliary parameter to estimate the error effects in expectation values, enabling significant error suppression with a fixed number of Trotter steps. Our approach offers an efficient way of quantum simulation on near-term quantum processors with shallow circuits.

Trotterization, Operator Scrambling, and Entanglement

Tianfeng Feng, Yue Cao, Qi Zhao

The University of Hong Kong

Operator scrambling governs the delocalization of quantum information in many-body systems and is closely tied to quantum chaos and information spreading. In this work, we reveal a fundamental connection between operator scrambling and the reliability of quantum simulations. We show that the Trotter error in simulating operator dynamics is bounded by the degree of operator scrambling, providing the most refined analysis of Trotter errors in operator dynamics so far. We then derive entanglement-based bounds showing that, for sufficiently entangled states, Trotter errors scale with the normalized Frobenius norms of both the observable and the error operator, which may give a double quadratic reduction in system-size dependence and operator-structure dependence. This indicates that, for general physical observables, the upper bound on the simulation error may be reduced from $O(N)$ to $\bar{O}(1)$ in systems undergoing entanglement growth and thermalization. Even when the state entanglement remains small, operator-induced entanglement can still suppress errors. Our results clarify the fundamental connection of Trotterization, scrambling, and entanglement, and provide refined guarantees for quantum simulation of many-body dynamics.

QEC III – Fault Tolerance and Compilation

Check-weight-constrained quantum codes: Bounds and examples

***Lily Wang*[1], *Andy Zeyi Liu*[2], *Ray Li*[3], *Aleksander Kubica*[2], *Shouzhen (Bailey) Gu*[2]**

[1] University of Michigan
[2] Yale University
[3] Santa Clara University

Quantum low-density parity-check (qLDPC) codes can be implemented by measuring only low-weight checks, making them compatible with noisy quantum hardware and central to the quest to build noise-resilient quantum computers. A fundamental open question is how constraints on check weight limit the achievable parameters of qLDPC codes. Here, we study stabilizer and subsystem codes with constrained check weight, combining analytical arguments with numerical optimization to establish strong upper bounds on their parameters. We show that stabilizer codes with checks of weight at most three cannot have nontrivial distance. We also prove tight tradeoffs between rate and distance for broad families of CSS stabilizer and subsystem codes with checks of weight at most four and two, respectively. Notably, our bounds are applicable to general qLDPC codes, as they rely only on check-weight constraints without assuming geometric locality or special graph connectivity. In the finite-size regime, we derive numerical upper bounds using linear programming techniques and identify explicit code constructions that approach these limits, delineating the landscape of practically relevant qLDPC codes with tens or hundreds of physical qubits.

Fault-tolerant transformations of spacetime codes

***Arthur Pesah*[1], *Austin K. Daniel*[2], *Ilan Tzitrin*[2], *Michael Vasmer*[3]**

[1] University College London
[2] Xanadu
[3] Inria Paris

Recent advances in quantum error-correction (QEC) have shown that it is often beneficial to understand fault-tolerance as a dynamical process—a circuit with redundant measurements that help correct errors—rather than as a static code equipped with a syndrome extraction circuit. Spacetime codes have emerged as a natural framework to understand error correction at the circuit level while leveraging the traditional QEC toolbox. Here, we introduce a framework based on chain complexes and chain maps to model spacetime codes and transformations between them. We show that stabilizer codes, quantum circuits, and decoding problems can all be described using chain complexes, and that the equivalence of two spacetime codes can be characterized by specific maps between chain complexes, the fault-tolerant maps, that preserve the number of encoded qubits, fault distance, and minimum-weight decoding problem. As an application of this framework, we extend the foliated cluster state construction from stabilizer codes to any spacetime code, showing that any Clifford circuit can be transformed into a measurement-based protocol with the same fault-tolerant properties. To this protocol, we associate a chain complex which encodes the underlying decoding problem, generalizing previous cluster state complex constructions. Our method enables the construction of cluster states from non-CSS, subsystem, and Floquet codes, as well as from logical Clifford operations on a given code.

Automated Compilation Including Dropouts: Tolerating Defective Components in Stabiliser Codes

Stasiu Wolanski[1,2]

[1] Riverlane
[2] University College London

Utility-scale solid-state quantum devices will need to fabricate quantum devices at scale using imperfect processes. By introducing tolerance to fabrication defects into the design of the quantum devices, we can improve the yield of usable quantum chips and lower the cost of useful systems. Automated Compilation Including Dropouts (ACID) is a framework that works in the ancilla-free (or 'middle-out') paradigm, to generate syndrome extraction circuits for general stabiliser codes in the presence of defective couplers or qubits. In the ancilla-free paradigm, we do not designate particular qubits as measurement ancillas, instead measuring stabilisers using any of the data qubits in their support. This approach leads to a great deal of flexibility in how syndrome extraction circuits can be implemented. ACID works by constructing and solving an optimisation problem within the ancilla-free paradigm to find a short syndrome extraction circuit. Applied to the surface code, ACID produces syndrome-extraction circuits of depth between 1× (no overhead) and 1.5× relative to the depth of defect-free circuits. I demonstrate the broad applicability of ACID by compiling syndrome extraction circuits for bivariate bicycle codes and the colour code. For these circuits, we incur a circuit-depth overhead of between 1× (no overhead) and 2.5× relative to defect-free circuits. I believe this work is the first to simulate both of these families of codes in the presence of fabrication defects.

Fermionic Systems

The Quantum Paldus Transform: Efficient Circuits with Applications

***Jedrzej Burkat*[1], *Nathan Fitzpatrick*[2]**

[1] University of Cambridge
[2] Quantinuum

We present the Quantum Paldus Transform: an efficient quantum algorithm for block-diagonalising fermionic, spin-free Hamiltonians in the second quantisation. Our algorithm implements an isometry between the occupation number basis of a fermionic Fock space of $2d$ modes, and the Gelfand-Tsetlin (GT) states spanning irreducible representations of the group $U(d) \times SU(2)$. The latter forms a basis indexed by well-defined values of total particle number N, global spin S, spin projection M, and $U(d)$ GT patterns. This realises the antisymmetric unitary-unitary duality discovered by Howe and developed into the Unitary Group Approach (UGA) for computational chemistry by Paldus and Shavitt in the 1970s. The Paldus transform lends tools from the UGA readily applicable to quantum computational chemistry, leading to maximally sparse representations of spin-free Hamiltonians, efficient preparation of Configuration State Functions, and a direct interpretation of quantum chemistry reduced density matrix elements in terms of $SU(2)$ angular momentum coupling. The transform also enables the encoding of quantum information into novel Decoherence-Free Subsystems for use in communication and error mitigation. Our work can be seen as a generalisation of the quantum Schur transform for the second quantisation, made tractable by the Pauli exclusion principle. Alongside self-contained derivations of the underlying dualities we provide fault-tolerant circuit compilation methods with full gate counts for the Paldus transform, resulting in $\mathcal{O}(d^3)$ Toffoli complexity, where a transform on 50 spatial orbitals would require a modest 5500 Toffoli gates. This paves the way for significant advancements in quantum simulation on quantum computers enabled by the UGA paradigm.

Near-Term Fermionic Simulation with Subspace Noise Tailored Quantum Error Mitigation

Miha Papič, Manuel G. Algaba, Emiliano Godinez-Ramirez, Inés de Vega, Adrian Auer, Fedor Šimkovic IV, Alessio Calzona

IQM Quantum Computers

Quantum error mitigation (QEM) has emerged as a powerful tool for the extraction of useful quantum information from quantum devices. Here, we introduce the Subspace Noise Tailoring (SNT) algorithm, which efficiently combines the cheap cost of Symmetry Verification (SV) and low bias of Probabilistic Error Cancellation (PEC) QEM techniques. We study the performance of our method by simulating the Trotterized time evolution of the spin-1/2 Fermi-Hubbard model (FHM) using a variety of local fermion-to-qubit encodings, which define a computational subspace through a set of stabilizers, the measurement of which can be used to post-select noisy quantum data. We study different combinations of QEM and encodings and uncover a rich state diagram of optimal combinations, depending on the hardware performance, system size and available shot budget. We then demonstrate how SNT extends the reach of current noisy quantum computers in terms of the number of fermionic lattice sites and the number of Trotter steps, and quantify the required hardware performance beyond which a noisy device may outperform classical computational methods.

Matchgate synthesis via Clifford matchgates and T gates

***Berta Casas*[1], *Paolo Braccia*[2], *Élie Gouzien*[3], *M. Cerezo*[2], *Diego Garcia Martin*[2]**

[1] Barcelona Supercomputing Center
[2] Los Alamos National Lab
[3] Alice & Bob

Matchgate unitaries are ubiquitous in quantum computation due to their relation to non-interacting fermions and because they can be used to benchmark quantum computers. Implementing such unitaries on fault-tolerant devices requires first compiling them into a discrete universal gate set, typically Clifford+T. Here, we propose a different approach for their synthesis: compile matchgate unitaries using only matchgate gates. To this end, we first show that the matchgate-Clifford group (the intersection of the matchgate and Clifford groups) plus the $\overline{T}$ gate (a T unitary up to a phase) is universal for the matchgate group. Our approach leverages the connection between n-qubit matchgate circuits and the standard representation of $\mathbb{SO}(2n)$, which reduces the compilation from $2^n \times 2^n$ unitaries to $2n \times 2n$ ones, thus reducing exponentially the size of the target matrix. Moreover, we rigorously show that this scheme is efficient, as an approximation error $\varepsilon_{\mathbb{SO}(2n)}$ incurred in this smaller-dimensional representation translates at most into an $\mathcal{O}(n\,\varepsilon_{\mathbb{SO}(2n)})$ error in the exponentially large unitary. In addition, we study the exact version of the matchgate synthesis problem, and we prove that all matchgate unitaries U such that $U \otimes U^*$ has entries in the ring $\mathbb{Z}\left[1/\sqrt{2}, i\right]$ can be exactly synthesized by a finite sequence of gates from the matchgate-Clifford+$\overline{T}$ set, without ancillas. We then use this insight to map optimal exact matchgate synthesis to Boolean satisfiability, and compile the circuits that diagonalize the free-fermionic XX Hamiltonian on $n = 4, 8$ qubits.

Quantum Control and Emulation

Quantum Optimal Control with Geodesic Pulse Engineering

***Dylan Lewis*[1], *Roeland Wiersema*[2], *Sougato Bose*[3]**

[1] Imperial College London
[2] Flatiron Institute
[3] University College London

Designing multi-qubit quantum logic gates with experimental constraints is an important problem in quantum computing. Here, we develop a new quantum optimal control algorithm for finding unitary transformations with constraints on the Hamiltonian. The algorithm, geodesic pulse engineering (GEOPE), uses differential programming and geodesics on the Riemannian manifold of $\mathrm{SU}(2^n)$ for n qubits. We demonstrate significant improvements over the widely used gradient-based method, GRAPE, for designing multi-qubit quantum gates. Instead of a local gradient descent, the parameter updates of GEOPE are designed to follow the geodesic to the target unitary as closely as possible. We present numerical results that show that our algorithm converges significantly faster than GRAPE for a range of gates and can find solutions that are not accessible to GRAPE in a reasonable amount of time. The strength of the method is illustrated with varied multi-qubit gates in 2D neutral Rydberg atom platforms.

Universal Dynamics with Globally Controlled Analog Quantum Simulators

***Hong-Ye Hu*[1], *Abigail McClain Gomez*[1], *Liyuan Chen*[1], *Aaron Trowbridge*[2], *Andy J. Goldschmidt*[3], *Zachary Manchester*[2], *Frederic T. Chong*[3], *Arthur Jaffe*[1], *Susanne F. Yelin*[1]**

[1] Harvard University
[2] Carnegie Mellon University
[3] University of Chicago

Analog quantum simulators with global control fields have emerged as powerful platforms for exploring complex quantum phenomena. Recent breakthroughs, such as the coherent control of thousands of atoms, highlight the growing potential for quantum applications at scale. Despite these advances, a fundamental theoretical question remains unresolved: to what extent can such systems realize universal quantum dynamics under global control? Here we establish a necessary and sufficient condition for universal quantum computation using only global pulse control, proving that a broad class of analog quantum simulators is, in fact, universal. We further extend this framework to fermionic and bosonic systems, including modern platforms such as ultracold atoms in optical superlattices. Moreover, we observe that analog simulators driven by random global pulses exhibit information scrambling comparable to random unitary circuits. In a dual-species neutral-atom array setup, the measurement outcomes anti-concentrate on a $\log N$ timescale despite the presence of only temporal randomness, opening opportunities for efficient randomness generation. To bridge theoretical possibility with experimental reality, we introduce *direct quantum optimal control*, a control framework that enables the synthesis of complex effective Hamiltonians while incorporating realistic hardware constraints. Using this approach, we experimentally engineer three-body interactions outside the blockade regime and demonstrate topological dynamics on a Rydberg-atom array. Experimental measurements reveal dynamical signatures of symmetry-protected-topological edge modes, confirming both the expressivity and feasibility of our method. Our work opens a new avenue for quantum simulation beyond native hardware Hamiltonians, enabling the engineering of effective multi-body interactions and advancing the frontier of quantum information processing with globally-controlled analog platforms.

A Fast and Frugal Gaussian Boson Sampling Emulator

***Tom Dodd*[1], *Javier Martínez-Cifuentes*[2], *Oliver Thomson Brown*[1], *Nicolás Quesada*[2], *Raúl García-Patrón*[1]**

[1] University of Edinburgh
[2] École Polytechnique de Montréal

If classical algorithms have been successful in reproducing the estimation of expectation values of observables of some recent quantum experiments, even using off-the-shelf computing resources, the simulation of sampling tasks, both qubit- and bosonic-based, are known to pose significantly harder challenges to classical simulations. The implementations of these algorithms required an extreme cost in terms of memory and computing operations, making them accessible to only a handful of supercomputers around the world and unreasonable to reproduce on a regular basis. In this work we demonstrate a classical simulation of imperfect GBS matching - and sometimes outperforming - the actual quantum experiments in a series of standardized tests, while requiring, for the first time, only a single CPU or GPU. Our algorithm samples from the joint probabilities of M-bit samples using a decomposition in terms of cumulants of increasing order and approximated marginals of the original distribution. Truncating the expression to a constant order K allows us to approximately capture all relevant statistics of the hardware experiment while making its running time scale as $O(M^K)$.

Algorithms VI – Optimization II

A quantum analogue of convex optimization

Eunou Lee

Korea Institute for Advanced Study (KIAS)

Convex optimization is the powerhouse behind the theory and practice of optimization. We introduce a quantum analogue of unconstrained convex optimization: find the minimum eigenvalue of a Schrödinger operator $h = -\Delta + V$ with convex potential $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ such that $V(x) \to \infty$ as $\|x\| \to \infty$. For this problem, we present an efficient quantum algorithm that computes the minimum eigenvalue of h up to error ϵ in polynomial time in n, $1/\epsilon$, and parameters that depend on V. Adiabatic evolution of the ground state is used as a key subroutine, which we analyze with novel techniques that allow us to focus on the low-energy space. We apply our algorithm to give the first known polynomial-time algorithm for finding the lowest frequency of an n-dimensional convex drum, or mathematically, the minimum eigenvalue of the Dirichlet Laplacian on an n-dimensional region that is defined by m linear constraints in polynomial time in n, m, $1/\epsilon$ and the radius R of a ball encompassing the region.

A measurement-driven quantum algorithm for SAT: Performance guarantees via spectral gaps and measurement parallelization

Franz J. Schreiber, Maximilian J. Kramer, Alexander Nietner, Jens Eisert

Dahlem Center for Complex Quantum Systems, Freie Universitaet Berlin, Germany

The Boolean satisfiability problem (SAT) is of central importance in both theory and practice. Yet, most provable guarantees for quantum algorithms rely exclusively on Grover-type methods that cap the possible advantage at only quadratic speed-ups, making the search for approaches that surpass this quadratic barrier a key challenge. In this light, this work presents a rigorous worst-case runtime analysis of a recently introduced measurement-driven quantum SAT solver. Importantly, this quantum algorithm does not exclusively rely on Grover-type methods and shows promising numerical performance. Our analysis establishes that the algorithm's runtime depends on an exponential trade-off between two key properties: the spectral gap of the associated Hamiltonian and the success probability of the driving measurements. We show that this trade-off can be systematically controlled by a tunable rotation angle. Beyond establishing a worst-case runtime expression, this work contributes significant algorithmic improvements. First, we develop a new readout routine that efficiently finds a solution even for instances with multiple satisfying assignments. Second, a measurement parallelization scheme, based on perfect hash families, is introduced. Third, we establish an amplitude-amplified version of the measurement-driven algorithm. Finally, we demonstrate the practical utility of our framework: By suitably scheduling the algorithm's parameters, we show that its runtime collapses from exponential to polynomial on a special class of SAT instances, consistent with their known classical tractability. A problem we leave open is to establish a non-trivial lower bound on the spectral gap as a function of the rotation angle. Resolving this directly translates into an improved worst-case runtime, potentially realizing a super-quadratic quantum advantage.

Grover's algorithm is an approximation of imaginary-time evolution

***Yudai Suzuki*[1], *Marek Gluza*[2], *Jeongrak Son*[2], *Bi Hong Tiang*[2], *Nelly H. Y. Ng*[2], *Zoë Holmes*[1]**

[1] EPFL

[2] Nanyang Technological University

We reveal the power of Grover's algorithm from thermodynamic and geometric perspectives by showing that it is a product formula approximation of imaginary-time evolution (ITE), a Riemannian gradient flow on the special unitary group. This viewpoint uncovers three key insights. First, we show that the ITE dynamics trace the shortest path between the initial and the solution states in complex projective space. Second, we prove that the geodesic length of ITE determines the query complexity of Grover's algorithm. This complexity notably aligns with the known optimal scaling for unstructured search. Lastly, utilizing the geodesic structure of ITE, we construct a quantum signal processing formulation for ITE without post-selection, and derive a new set of angles for the fixed-point search. These results collectively establish a deeper understanding of Grover's algorithm and suggest a potential role for thermodynamics and geometry in quantum algorithm design.

QEC IV – qLDPC Codes

Directional Codes: a new family of quantum LDPC codes on hexagonal- and square-grid connectivity hardware

Gyorgy Geher, David Byfield, Archibald Ruban

Riverlane

Utility-scale quantum computing requires quantum error correction (QEC) to protect quantum information against noise. Currently, superconducting hardware is a promising candidate for achieving fault tolerance due to its fast gate times and feasible scalability. However, it is often restricted to two-dimensional nearest-neighbour connectivity, which is thought to be incapable of accommodating high-rate quantum low-density parity-check (qLDPC) codes that promise to greatly reduce the number of physical qubits needed to encode logical qubits. In this paper we construct a new family of qLDPC codes, which we call "Directional Codes", that outperforms the rotated planar code (RPC) while naturally meeting the connectivity requirements of the widely adopted square-grid, and some even the sparser hexagonal-grid. The key idea is to utilise the iSWAP gate – a natural native gate for superconducting qubits – to construct circuits that measure the stabilisers of these qLDPC codes without the need for any long-range connections or an increased degree of connectivity. We numerically evaluate the performance of directional codes, encoding four, six and twelve logical qubits, using a common superconducting-inspired circuit-level Pauli noise model. We also compare them to the RPC and to the bivariate bicycle (BB) codes, currently the two most popular quantum LDPC code families. As a concrete example, directional codes outperform the RPC by achieving approximately the same logical error probability at physical error rate $p = 10^{-3}$ using only $18.75 - 45\%$ of the physical qubits at distance up to 10. Our discovery represents a breakthrough in QEC code design that suggests complex long-range, high-connectivity hardware may not be necessary for low-overhead fault-tolerant quantum computation.

High-performance syndrome extraction circuits for quantum codes

***Armands Strikis*[1], *Dan Browne*[2], *Michael Beverland*[3]**

[1] University of Oxford
[2] University College London
[3] IBM

We present a fast and effective framework for analysing and designing syndrome-extraction circuits (SECs). Our approach is based on left–right circuits, a general design for SECs which maintain low depth by staggering X and Z checks without interleaving gates. Initially proposed for specific classes of codes, we generalise this construction to arbitrary CSS codes and optimise the circuit structure to achieve low qubit idling time, large effective distance, and reduced minimum-weight failure mechanisms. A key component of our framework is the formal notion of *residual errors* and their associated distance metrics, which form lightweight tools for capturing error propagation and quantifying the potential harm of circuit-level errors. Applying our automated framework to diverse classes of codes, we observe consistent improvements in logical performance of up to an order of magnitude compared to existing single-ancilla SEC designs. We also use these tools to prove that no non-interleaving SEC can achieve circuit distance 12 for the gross code, and identify an explicit circuit that we conjecture achieves distance 11, exceeding previously known constructions.

Computing Efficiently in QLDPC Codes

***Alexander Malcolm*[1], *Andrew Glaudell*[1], *Patricio Fuentes*[1], *Daryus Chandra*[1], *Alexis Schotte*[1], *Colby DeLisle*[1], *Rafael Haenel*[1,2], *Amir Ebrahimi*[1], *Joschka Roffe*[1,3], *Armanda Quintavalle*[1,4], *Stefanie Beale*[1,5], *Nichola Lee-Hone*[1], *Stephanie Simmons*[1]**

[1] Photonic Inc.
[2] QuEra
[3] University of Edinburgh
[4] Freie Universität Berlin
[5] Nord Quantique

Quantum error correction is essential for building utility-scale quantum computers that outperform classical machines, yet leading approaches incur substantial physical qubit overhead. Quantum low-density parity check (QLDPC) codes offer a promising alternative by significantly reducing the number of physical qubits required per logical qubit. However, existing work on QLDPC codes has focused primarily on quantum memories, with no efficient method known for implementing arbitrary logical Clifford operations at low circuit depth. Here, we introduce a new family of QLDPC codes that enables efficient implementation of the full Clifford group via transversal operations, allowing any m-qubit Clifford operation to be executed in at most $O(m)$ syndrome extraction rounds. We run circuit-level simulations of depth-126 logical circuits to demonstrate the near-memory logical performance of these logical operations. In combination with known methods for implementing T gates, these results establish QLDPC codes as a viable route toward resource-efficient universal quantum computation. Further validation is provided via competitive QRE analysis of algorithms running on a SHYPS-based, distributed quantum computing architecture.