

Book of Accepted Talks

Quantum Computing Theory in Practice (QCTiP)

April 23–25, 2025
Berlin

List of Presenters

Miriam Backens

Inria

Title: An algebraic interpretation of Pauli flow, leading to faster flow-finding algorithms

Pascal Baßler

Heinrich-Heine University Düsseldorf

Title: General, efficient, and robust Hamiltonian engineering

Lucas Berent

Technical University of Munich

Title: Localized statistics decoding: A parallel decoding algorithm for quantum low-density parity-check codes

Boris Bourdoncle

Quandela

Title: Minimizing resource overhead in fusion-based quantum computation using hybrid spin-photon devices

Zhenyu Cai

University of Oxford

Title: Quantum Error Mitigation for Sampling Algorithms

Laura Caune

Riverlane

Title: Demonstrating real-time and low-latency quantum error correction with superconducting qubits

Tim Chan

University of Oxford

Title: Snowflake: A Distributed Streaming Decoder

Giulio Croagnaletti

University of Trieste

Title: Estimates of loss function concentration in noisy parametrized quantum circuits

Theo Desselain

Quandela

Title: Enhanced Fault-tolerance in Photonic Quantum Computing: Floquet Code Outperforms Surface Code in Tailored Architecture

Dhruv Devulapalli

University of Maryland, College Park

Title: Quantum Routing and Entanglement Dynamics through Bottlenecks

Werner Dobrautz

Helmholtz Zentrum Dresden-Rossendorf

Title: Toward Real Chemical Accuracy on Current Quantum Hardware Through the Transcorrelated Method

Joao F. Doriguello

HUN-REN Alfréd Rényi Institute of Mathematics

Title: On the practicality of quantum sieving algorithms for the shortest vector problem

Timo Eckstein

FAU Erlangen-Nürnberg

Title: Shot-noise reduction for lattice Hamiltonians

Laura Clinton

Phasercraft

Title: Quantum Phase Estimation without Controlled Unitaries

Soumik Ghosh

University of Chicago

Title: The Hardness of Learning Quantum Circuits and its Cryptographic Applications

Jakob Günther

University of Copenhagen

Title: Phase estimation with partially randomized product formulas

Yaroslav Herasymenko

CWI

Title: Solving Free Fermion Problems on a Quantum Computer

Google Quantum AI

Google Quantum AI

Title: Quantum error correction below the surface code threshold

Marcel Hinsche

Freie Universität Berlin

Title: Single-copy stabilizer testing

Hong-Ye Hu

Harvard University

Title: Demonstration of Robust and Efficient Quantum Property Learning with Shallow Shadows

Ali Javadi-Abhari

IBM

Title: Suppressing Correlated Noise in Quantum Computers via Context-Aware Compiling

Oriel Kiss

CERN

Title: Early Fault-Tolerant Quantum Algorithms in Practice: Application to Ground-State Energy Estimation

Seok-Hyung Lee

The University of Sydney

Title: Low-overhead magic state distillation with color codes

Yuan Liu

North Carolina State University

Title: Unification of Finite Symmetries in Simulation of Many-body Systems on Quantum Computers

Alec Eickbusch

Google

Title: Demonstrating dynamic surface codes

Antonio Anna Mele

Freie Universität Berlin

Title: Tomography of bosonic systems and optimal estimates of the trace distance between Gaussian states

Nikolai Miklin

Hamburg University of Technology

Title: SPAM-free sound certification of quantum gates via quantum system quizzing

Daniel Miller

Freie Universität Berlin

Title: Experimental measurement and a physical interpretation of quantum shadow enumerators

Daniel Mills

Quantinuum

Title: On-Chip Verified Quantum Computation with an Ion-Trap Quantum Processing Unit

Frederik Nathan

NNF Quantum Computing Programme, University of Copenhagen

Title: Self-correcting GKP qubit in a superconducting circuit with an oscillating voltage bias

Yiyuan Chen

QuSoft / University of Amsterdam

Title: Trotter error and gate complexity of the SYK and sparse SYK models

Stefano Polla

Leiden University

Title: Error mitigation and circuit division for early fault-tolerant quantum phase estimation

Ricard Puig

École Polytechnique Fédérale de Lausanne

Title: Guarantees and limitations for warm starts and iterative methods in variational quantum computing

Javier Robledo Moreno

IBM Quantum

Title: Chemistry Beyond Exact Solutions on a Quantum-Centric Supercomputer*

Raul A Santos

Phasetcraft

Title: Approximating dynamical correlation functions with constant depth quantum circuits

Hasan Sayginel

University College London

Title: Fault-Tolerant Logical Clifford Gates from Code Automorphisms

Alexander Schuckert

University of Maryland

Title: Fermion-qubit fault-tolerant quantum computing

Alireza Seif

IBM Quantum

Title: Entanglement-enhanced learning of quantum processes at scale

Javier Robledo Moreno

IBM

Title: Sample-based Krylov Quantum Diagonalization*

Noah John Shutty

Google Quantum AI

Title: Optimization by Decoded Quantum Interferometry

Fedor Simkovic

IQM Quantum Computers

Title: Contextual Subspace Auxiliary-Field Quantum Monte Carlo: Improved bias with reduced quantum resources

Vincent Steffan

IQM Germany

Title: Logical Operators and Fold-Transversal Gates of Bivariate Bicycle Codes

Christopher Vairogs

University of Illinois Urbana-Champaign

Title: Localizing multipartite entanglement with local and global measurements

Katherine Van Kirk

Harvard University

Title: Derandomized shallow shadows: Efficient Pauli learning with bounded-depth circuits

Xin Wang

The Hong Kong University of Science and Technology (Guangzhou)

Title: Quantum Algorithm for Reversing Unknown Unitary Evolutions

Wenjun Yu

The University of Hong Kong

Title: Observable-Driven Speed-ups in Quantum Simulations

Jeffery Yu

University of Maryland

Title: Efficient preparation of Dicke states

You Zhou

Fudan University

Title: Auxiliary-free replica shadow estimation

Nadish de Silva

Simon Fraser University

Title: Classifying the gates of the Clifford hierarchy

Štěpán Šmíd

Imperial College London

Title: Polynomial Time Quantum Gibbs Sampling for Fermi-Hubbard Model at any Temperature

*These two submissions will be presented as part of a merged talk.

Abstracts

An algebraic interpretation of Pauli flow, leading to faster flow-finding algorithms

Presenter: Miriam Backens (Inria)

Authors: Piotr Mitosek, Miriam Backens

The one-way model of quantum computing is used not just to implement quantum computation but also as a tool for compilation and optimisation. One-way computations are driven by successive adaptive measurements of a graph state. 'Flow structures' witness whether a computation can be performed deterministically overall, and these structures can be found in polynomial time. However, existing definitions are lengthy and complex, hindering the ability to work with flow. We simplify these definitions by providing a new algebraic interpretation of Pauli flow, the most general flow structure. We define two matrices arising from the graph's adjacency matrix: the flow-demand matrix M and the order-demand matrix N . We show Pauli flow exists if and only if there is a right inverse C of M such that the product NC describes a directed acyclic graph. Using our algebraic interpretation, we obtain $O(n^3)$ algorithms for finding Pauli flow, improving previous bounds. We also introduce a first lower bound for the flow-finding problem by linking it to matrix invertibility over \mathbb{F}_2 , thus showing our algorithms are optimal, barring fundamental progress in matrix multiplication algorithms. The better understanding of Pauli flow and faster flow-finding algorithms will have applications in compilation, fault tolerance, and other fields.

General, efficient, and robust Hamiltonian engineering

Presenter: Pascal Baßler (Heinrich-Heine University Düsseldorf)

Authors: Pascal Baßler, Markus Heinrich, Martin Kliesch

We propose an efficient and robust scheme to engineer arbitrary many-body Hamiltonians by interleaving the free evolution of a fixed system Hamiltonian with layers of single-qubit Pauli or Clifford gates. These sequences are constructed by solving a linear program (LP) which minimizes the total evolution time. The target Hamiltonians that can be engineered by our method are only limited by the locality of the Pauli terms in the system Hamiltonian. We then construct an efficient relaxation of this LP with a classical runtime that depends on the number of Pauli terms in the Pauli decomposition of the system Hamiltonian and is thus a low-degree polynomial in the number of qubits in practice. With our method, it is also possible to engineer Hamiltonians if only partial knowledge of the system Hamiltonian is available, which can be used to cancel unknown unwanted Pauli terms. We show the classical efficiency of our method by engineering an arbitrary two-body Hamiltonian on a 2D square lattice with 225 qubits in only 60 seconds. We make our scheme robust against finite pulse time effects by applying average Hamiltonian theory. Moreover, we can combine our methods with arbitrary robust composite pulses, making them robust against many experimental error sources. We provide numerical simulations for a device with 2D lattice interactions as well as for a trapped-ion device with all-to-all interactions, demonstrating the benefits of our methods in simulating Ising and Heisenberg Hamiltonians. Moreover, we address dominant error sources in practical applications.

Localized statistics decoding: A parallel decoding algorithm for quantum low-density parity-check codes

Presenter: Lucas Berent (Technical University of Munich)

Authors: Timo Hillmann, Lucas Berent, Armando O. Quintavalle, Jens Eisert, Robert Wille, Joschka Roffe

Quantum low-density parity-check codes are a promising candidate for fault-tolerant quantum computing with considerably reduced overhead compared to the surface code. However, the lack of a practical decoding algorithm remains a barrier to their implementation. In this work, we introduce localized statistics decoding, a reliability-guided inversion decoder that is highly parallelizable and applicable to arbitrary quantum low-density parity-check codes. Our approach employs a parallel matrix factorization strategy, which we call on-the-fly elimination, to identify, validate, and solve local decoding regions on the decoding graph. Through numerical simulations, we show that localized statistics decoding matches the performance of state-of-the-art decoders while reducing the runtime complexity for operation in the sub-threshold regime. Importantly, our decoder is more amenable to implementation on specialized hardware, positioning it as a promising candidate for decoding real time syndromes from experiments.

Minimizing resource overhead in fusion-based quantum computation using hybrid spin-photon devices

Presenter: Boris Bourdoncle (Quandela)

Authors: Stephen Wein, Timothée Goubault de Brugière, Luka Music, Pascale Senellart, Boris Bourdoncle, Shane Mansfield

We present three schemes for constructing a (2,2)-Shor-encoded 6-ring photonic resource state for fusion-based quantum computing, each relying on a different type of photon source. We benchmark these architectures by analyzing their ability to achieve the best-known loss tolerance threshold for fusion-based quantum computation using the target resource state. More precisely, we estimate their minimum hardware requirements for fault-tolerant quantum computation in terms of the number of photon sources to achieve on-demand generation of resource states with a desired generation period. Notably, we find that a group of 12 deterministic single-photon sources containing a single matter qubit degree of freedom can produce the target resource state near-deterministically by exploiting entangling gates that are repeated until success. The approach is fully modular, eliminates the need for lossy large-scale multiplexing, and reduces the overhead for resource-state generation by several orders of magnitude compared to architectures using heralded single-photon sources and probabilistic linear-optical entangling gates. Our work shows that the use of deterministic single-photon sources embedding a qubit substantially shortens the path toward fault-tolerant photonic quantum computation.

Quantum Error Mitigation for Sampling Algorithms

Presenter: Zhenyu Cai (University of Oxford)

Authors: Kecheng Liu, Zhenyu Cai

Recent experimental breakthroughs have signalled the imminent arrival of the early fault-tolerant era. However, for a considerable period in the foreseeable future, relying solely on quantum error correction for full error suppression will remain extremely challenging due to its substantial hardware overhead. Additional help from quantum error mitigation (QEM) is essential for bridging this gap towards achieving quantum advantage. The application of QEM has so far been restricted to expectation value estimation, leaving its extension to sampling-based algorithms – which is expected to play a pivotal role in the early fault-tolerant era – an unresolved challenge. In this work, we present a framework for applying any QEM techniques to obtain the error-mitigated output distribution, showing that this incurs no greater cost than estimating a single observable. We also devised a way to sample from this distribution and constructed an explicit scheme for applying any QEM methods to quantum phase estimation, which can be generalised to other sampling algorithms. Numerical experiments were conducted to validate the efficacy of these methods. We believe our methods significantly broaden the scope of QEM, extending its applicability to most algorithms of practical interest and forming a crucial step towards realising quantum advantage.

Demonstrating real-time and low-latency quantum error correction with superconducting qubits

Presenter: Laura Caune (Riverlane)

Authors: Laura Caune, Luka Skoric, Nick S. Blunt, Archibald Ruban, Jimmy McDaniel, Joseph A. Valery, Andrew D. Patterson, Alexander V. Gramolin, Joonas Majaniemi, Kenton M. Barnes, Tomasz Bialas, Okan Buğdaycı, Ophelia Crawford, György P. Gehér, Hari Krovi, Elisha Matekole, Canberk Topal, Stefano Poletto, Michael Bryant, Kalan Snyder, Neil I. Gillespie, Glenn Jones, Kauser Johar, Earl T. Campbell, Alexander D. Hill

Quantum error correction (QEC) will be essential to achieve the accuracy needed for quantum computers to realise their full potential. The field has seen promising progress with demonstrations of early QEC and real-time decoded experiments. As quantum computers advance towards demonstrating a universal fault-tolerant logical gate set, implementing scalable and low-latency real-time decoding will be crucial to prevent the backlog problem and maintaining a fast logical clock rate. Here, we demonstrate low-latency feedback with a scalable FPGA decoder integrated into the control system of a superconducting quantum processor. We perform an 8-qubit stability experiment with up to 25 decoding rounds and a mean decoding time per round below $1 \mu\text{s}$, showing that we avoid the backlog problem even on superconducting hardware with the strictest speed requirements. We observe logical error suppression as the number of decoding rounds is increased. We also implement and time a fast-feedback experiment demonstrating a decoding response time of $9.6 \mu\text{s}$ for a total of 9 measurement rounds. The decoder throughput and latency developed in this work, combined with continued device improvements, unlock the next generation of experiments that go beyond purely keeping logical qubits alive and into demonstrating building blocks of fault-tolerant computation.

Snowflake: A Distributed Streaming Decoder

Presenter: Tim Chan (University of Oxford)

Authors: Tim Chan

We design Snowflake, a quantum error correction decoder that, for the surface code under circuit-level noise, is roughly 25% more accurate than the Union-Find decoder, with a better mean runtime scaling: subquadratic as opposed to cubic in the code distance. Our decoder runs in a streaming fashion and has a local implementation. In designing Snowflake, we propose a new method for general stream decoding that eliminates the processing overhead due to window overlap in existing windowing methods.

Estimates of loss function concentration in noisy parametrized quantum circuits

Presenter: Giulio Crognaletti (University of Trieste)

Authors: Giulio Crognaletti, Michele Grossi, Angelo Bassi

Variational quantum computing offers a powerful framework with applications across diverse fields such as quantum chemistry, machine learning, and optimization. However, its scalability is hindered by the exponential concentration of the loss function, known as the barren plateau problem. While significant progress has been made in understanding barren plateaus in unitary and noisy circuits independently, their combined effects remain poorly understood due to limitations in standard Lie algebraic methods. In this work, we introduce a new analytical formulation based on non-negative matrix theory that enables precise calculation of the variance in deep noisy circuits, unveiling the complex interplay between unitary circuits and noise. In particular, we show the emergence of a noise-induced absorption mechanism, a phenomenon that cannot arise in the purely reversible context of unitary quantum computing. Nevertheless, general lower bounds on the variance of deep circuits can still be established, constraining arbitrary circuits to effectively mimic the behaviour of shallow ones. When applied in the noisy setting, this allows us to establish a deep connection between the noise resilience of parameterized circuits and the potential to enhance their expressive power through smart initialization strategies. Theoretical developments are supported by numerical examples and related applications.

Enhanced Fault-tolerance in Photonic Quantum Computing: Floquet Code Outperforms Surface Code in Tailored Architecture

Presenter: Theo Dessertaine (Quandela)

Authors: Paul Hilaire, Theo Dessertaine, Boris Bourdoncle, Aurélie Denys, Grégoire De Gliniasty, Gerard Valenti-Rojas, Shane Mansfield

Fault-tolerant quantum computing is crucial for realizing large-scale quantum computation, and the interplay between hardware architecture and quantum error-correcting codes is a key consideration. We present a comparative study of two quantum error-correcting codes, the surface code and the honeycomb Floquet code, implemented on variants of the spin-optical quantum computing architecture, enabling a direct comparison of the codes using consistent noise models. Our results demonstrate that the honeycomb Floquet code significantly outperforms the surface code in this setting. Notably, we achieve a photon loss threshold of 6.4% for the honeycomb Floquet code implementation, to our knowledge the highest reported for photonic platforms to date without large-scale multiplexing. This finding is particularly significant given that photon loss is the primary source of errors in photon-mediated quantum computing.

Quantum Routing and Entanglement Dynamics through Bottlenecks

Presenter: Dhruv Devulapalli (University of Maryland, College Park)

Authors: Dhruv Devulapalli, Chao Yin, Andrew Y. Guo, Eddie Schoute, Adam Ehrenberg, Andrew M. Childs, Alexey V. Gorshkov, Andrew Lucas

To implement arbitrary quantum interactions in architectures with restricted geometries, one may simulate all-to-all connectivity by routing quantum information. We consider the entanglement dynamics and routing between two regions only connected through an intermediate region with few qubits that forms a bottleneck. In such systems, existing results give only a trivial constant lower bound on the routing time, or the minimum time taken to perform an arbitrary permutation. We significantly improve this lower bound. For any system with two regions with N_L, N_R qubits respectively, coupled only through an intermediate region of N_C qubits, for any $\delta > 0$ we show a lower bound of $\Omega(N_R^{1-\delta}/\sqrt{N_L N_C})$ on the routing time. We also prove an upper bound on the average amount of bipartite entanglement that can be generated in time T by an architecture-respecting Hamiltonian in systems constrained by vertex bottleneck, improving the scaling in the system size from $O(N_L T)$ to $O(\sqrt{N_L} T)$. When applied to the star graph, we obtain an $\Omega(\sqrt{N^{1-\delta}})$ lower bound on the routing time and on the time to prepare $N/2$ Bell pairs between the vertices. We also show a Hamiltonian routing protocol in systems of free fermions on the star graph with a speedup over gate-based routing.

Toward Real Chemical Accuracy on Current Quantum Hardware Through the Transcorrelated Method

Presenter: Werner Dobrautz (Helmholtz Zentrum Dresden-Rossendorf)

Authors: Werner Dobrautz, Igor O. Sokolov, Ke Liao, Pablo López Ríos, Martin Rahm, Ali Alavi, Ivano Tavernelli

Quantum computing is emerging as a new computational paradigm with the potential to transform several research fields including quantum chemistry. However, current hardware limitations (including limited coherence times, gate infidelities, and connectivity) hamper the implementation of most quantum algorithms and call for more noise-resilient solutions. We propose an explicitly correlated Ansatz based on the transcorrelated (TC) approach to target these major roadblocks directly. This method transfers, without any approximation, correlations from the wave function directly into the Hamiltonian, thus reducing the resources needed to achieve accurate results with noisy quantum devices. We show that the TC approach allows for shallower circuits and improves the convergence toward the complete basis set limit, providing energies within chemical accuracy to experiment with smaller basis sets and, thus, fewer qubits. We demonstrate our method by computing bond lengths, dissociation energies, and vibrational frequencies close to experimental results for the hydrogen dimer and lithium hydride using two and four qubits, respectively. To demonstrate our approach's current and near-term potential, we perform hardware experiments, where our results confirm that the TC method paves the way toward accurate quantum chemistry calculations already on today's quantum hardware.

On the practicality of quantum sieving algorithms for the shortest vector problem

Presenter: Joao F. Doriguello (HUN-REN Alfréd Rényi Institute of Mathematics)

Authors: Joao F. Doriguello, George Giapitzakis, Alessandro Luongo, Aditya Morolia

One of the candidates of post-quantum cryptography is lattice-based cryptography. Its cryptographic security is based on the hardness of lattice problems like the shortest vector problem (SVP), which asks to find the shortest non-zero vector in an integer lattice. Asymptotic quantum speedups for solving SVP are known and rely on Grover's search. However, to assess the security against these Grover-like quantum speedups, it is necessary to carry out a precise resource estimation. In this work, we perform a careful analysis on the resources required to implement several sieving algorithms aided by Grover's search for dimensions of cryptographic interests. We take into account fixed-point quantum arithmetic operations, non-asymptotic Grover's search, QRAM, different physical architectures, and quantum error correction. We find that even under optimistic assumptions, the best sieving algorithms require $\approx 10^{13}$ physical qubits and $\approx 10^{31}$ years to solve SVP on a 400-dimensional lattice, roughly the dimension for minimally secure post-quantum cryptographic standards. We estimate that a 6-GHz-clock-rate single-core classical computer would take the same time to solve the same problem. We conclude that there is currently little to no quantum speedup in the dimensions of cryptographic interest and significant breakthroughs in theoretical protocols and hardware development are needed.

Shot-noise reduction for lattice Hamiltonians

Presenter: Timo Eckstein (FAU Erlangen-Nürnberg)

Authors: Timo Eckstein, Refik Mansuroglu, Stefan Wolf, Ludwig Nützel, Stephan Tasler, Martin Kliesch, Michael J. Hartmann

Efficiently estimating energy expectation values of lattice Hamiltonians on quantum computers is a serious challenge, where established techniques can require excessive sample numbers. Here we introduce geometric partitioning as a scalable alternative. It splits the Hamiltonian into subsystems that extend over multiple lattice sites, for which transformations between their local eigenbasis and the computational basis can be efficiently found. This allows us to reduce the number of measurements as we sample from a more concentrated distribution without diagonalizing the problem. For systems in an energy eigenstate, we prove a lower bound on the sampling number improvement over the “naive” mutually commuting local operator grouping, which grows with the considered subsystem size, consistently showing an advantage for our geometric partitioning strategy. Notably, our lower bounds do not decrease but increase for more correlated states (Theorem 1). For states that are weakly isotropically perturbed around an eigenstate, we show how the sampling number improvement translates to imperfect eigenstate improvements, namely measuring close to the true eigenbasis already for smaller perturbations (Theorem 2). We illustrate our findings on multiple two-dimensional lattice models incl. the transverse field XY- and Ising model as well as the Fermi Hubbard model.

Quantum Phase Estimation without Controlled Unitaries

Presenter: Laura Clinton (Phasercraft)

Authors: Laura Clinton, Toby Cubitt, Raul Garcia-Patron, Ashley Montanaro, Stasja Stanisic, Maarten Stroeks

In this work we demonstrate the use of adapted classical phase retrieval algorithms to perform control-free quantum phase estimation. We eliminate the costly controlled time evolution and Hadamard test commonly required to access the complex time-series needed to reconstruct the spectrum. This significant reduction of the number of coherent controlled-operations lowers the circuit depth and considerably simplifies the implementation of statistical quantum phase estimation in near-term devices. This seemingly impossible task can be achieved by extending the problem that one wishes to solve to one with a larger set of input signals while exploiting natural constraints on the signal and/or the spectrum. We leverage well-established algorithms that are widely used in the context of classical signal processing, demonstrating two complementary methods to do this, vectorial phase retrieval and two-dimensional phase retrieval. We numerically investigate the feasibility of both approaches for estimating the spectrum of the Fermi-Hubbard model and discuss their resilience to inherent statistical noise.

The Hardness of Learning Quantum Circuits and its Cryptographic Applications

Presenter: Soumik Ghosh (University of Chicago)

Authors: Bill Fefferman, Soumik Ghosh, Makrand Sinha, Henry Yuen

We show that concrete hardness assumptions about learning or cloning the output state of a random quantum circuit can be used as the foundation for secure quantum cryptography. In particular, under these assumptions we construct secure one-way state generators, digital signature schemes, quantum bit commitments, and private key encryption schemes. We also discuss evidence for these hardness assumptions by analyzing the best-known quantum learning algorithms, as well as proving black-box lower bounds for cloning and learning given state preparation oracles. Our random circuit-based constructions provide concrete instantiations of quantum cryptographic primitives whose security do not depend on the existence of one-way functions. The use of random circuits in our constructions also opens the door to NISQ-friendly quantum cryptography. We discuss noise tolerant versions of our OWSG and digital signature constructions which can potentially be implementable on a noisy quantum computer. On the other hand, they are still secure against noiseless quantum adversaries, raising the intriguing possibility of a useful implementation of an end-to-end cryptographic protocol by a near-term quantum computer. Finally, our explorations suggest that the rich interconnections between learning theory and cryptography in classical theoretical computer science also extend to the quantum setting.

Phase estimation with partially randomized product formulas

Presenter: Jakob Günther (University of Copenhagen)

Authors: Jakob Günther, Freek Witteveen, Aram Harrow, Matthias Christandl, Marek Miller, Alexander Schmidhuber, Li Liu

Quantum phase estimation (QPE) combined with Hamiltonian simulation is the most promising algorithmic framework to computing ground state energies on quantum computers. Its main computational overhead derives from the Hamiltonian simulation subroutine. In this paper we use randomization to speed up this process, in particular product formulas. We observe that single-ancilla QPE methods combine well with error guarantees obtained by randomized Hamiltonian simulation. We also propose partially randomized Hamiltonian simulation methods in which some terms are kept deterministically and others are randomly sampled. This is useful if the Hamiltonian has a modest number of important term, and a long tail of small terms. We evaluate these methods numerically on examples from quantum chemistry, where we have Hamiltonians with tails of small terms, and find significant reductions in cost and circuit depth.

Solving Free Fermion Problems on a Quantum Computer

Presenter: Yaroslav Herasymenko (CWI)

Authors: Maarten Stroeks, Daan Lenterman, Barbara Terhal, Yaroslav Herasymenko

Simulating noninteracting fermion systems is a basic computational tool in many-body physics. In absence of translational symmetries, modeling free fermions on N modes usually requires $\text{poly}(N)$ computational resources. While often moderate, these costs can be prohibitive in practice when large systems are considered. We present several noninteracting fermion problems that can be solved by a quantum algorithm with exponentially-improved, $\text{poly log}(N)$ cost. We point out that the simulation of free-fermion dynamics belongs to the BQP-hard complexity class, implying that our discovered exponential speedup is robust. The key technique in our algorithm is the block-encoding of the correlation matrix into a unitary. We demonstrate how such a unitary can be efficiently realized as a quantum circuit, in the context of dynamics and thermal states of tight-binding Hamiltonians. The special cases of disordered and inhomogeneous lattices, as well as large non-lattice graphs, are presented in detail. Finally, we show that our simulation algorithm generalizes to other promising targets, including free boson systems.

Quantum error correction below the surface code threshold

Presenter: Google Quantum AI (Google Quantum AI)

Authors: Google Quantum AI

Quantum error correction provides a path to reach practical quantum computing by combining multiple physical qubits into a logical qubit. In this work, we present two surface code memories operating below threshold: a distance-7 code and a distance-5 code integrated with a real-time decoder. The logical error rate of our larger quantum memory is suppressed by a factor of $\Lambda = 2.14 \pm 0.02$ when increasing the code distance by two, culminating in a 101-qubit distance-7 code with $0.143\% \pm 0.003\%$ error per cycle of error correction. This logical memory is also beyond break-even, exceeding its best physical qubit's lifetime by a factor of 2.4 ± 0.3 . We maintain below-threshold performance when decoding in real time, achieving an average decoder latency of $63 \mu\text{s}$ at distance-5 up to a million cycles, with a cycle time of $1.1 \mu\text{s}$. To probe the limits of our error-correction performance, we run repetition codes up to distance-29 and find that logical performance is limited by rare correlated error events occurring approximately once every hour, or 3×10^9 cycles. Our results present device performance that, if scaled, could realize the operational requirements of large scale fault-tolerant quantum algorithms.

Single-copy stabilizer testing

Presenter: Marcel Hinsche (Freie Universität Berlin)

Authors: Marcel Hinsche, Jonas Helsen

We consider the problem of testing whether an unknown n -qubit quantum state is a stabilizer state, with only single-copy access. We give an algorithm solving this problem using $O(n)$ copies, and conversely prove that $\Omega(\sqrt{n})$ copies are required for any algorithm. The main observation behind our algorithm is that when repeatedly measuring in a randomly chosen stabilizer basis, stabilizer states are the most likely among the set of all pure states to exhibit linear dependencies in measurement outcomes. Our algorithm is designed to probe deviations from this extremal behavior. For the lower bound, we first reduce stabilizer testing to the task of distinguishing random stabilizer states from the maximally mixed state. We then argue that, without loss of generality, it is sufficient to consider measurement strategies that a) lie in the commutant of the tensor action of the Clifford group and b) satisfy a Positive Partial Transpose (PPT) condition. By leveraging these constraints, together with novel results on the partial transposes of the generators of the Clifford commutant, we derive the lower bound on the sample complexity.

Demonstration of Robust and Efficient Quantum Property Learning with Shallow Shadows

Presenter: Hong-Ye Hu (Harvard University)

Authors: Hong-Ye Hu, Andi Gu, Swarnadeep Majumder, Hang Ren, Yipei Zhang, Derek Wang, Yi-Zhuang You, Zlatko Minev, Susanne Yelin, Alireza Seif

Randomized measurements, or classical shadows, enable predicting many properties of arbitrary quantum states using few measurements. While random single-qubit measurements are experimentally friendly and suitable for learning low-weight Pauli observables, they perform poorly for nonlocal observables. Prepending a shallow random quantum circuit before measurements maintains this experimental friendliness, but also has favorable sample complexities for observables beyond low-weight Paulis, including high-weight Paulis and global low-rank properties such as fidelity. However, in realistic scenarios, quantum noise accumulated with each additional layer of the shallow circuit biases the results. To address these challenges, we propose the robust shallow shadows protocol. Our protocol uses Bayesian inference to learn the experimentally relevant noise model and mitigate it in postprocessing. This mitigation introduces a bias-variance trade-off: correcting for noise-induced bias comes at the cost of a larger estimator variance. Despite this increased variance, as we demonstrate on a superconducting quantum processor, our protocol correctly recovers state properties such as expectation values, fidelity, and entanglement entropy, while maintaining a lower sample complexity compared to the random single qubit measurement scheme. We also theoretically analyze the effects of noise on sample complexity and show how the optimal choice of the shallow shadow depth varies with noise strength.

Suppressing Correlated Noise in Quantum Computers via Context-Aware Compiling

Presenter: Ali Javadi-Abhari (IBM)

Authors: Alireza Seif, Haoran Liao, Vinay Tripathi, Kevin Krsulich, Moein Malekakhlagh, Mirko Amico, Petar Jurcevic, Ali Javadi-Abhari

Coherent errors, especially those that occur in correlation among a set of qubits, are detrimental for large-scale quantum computing. Correlations in noise can occur as a result of spatial and temporal configurations of instructions executing on the quantum processor. In this paper, we perform a detailed experimental characterization of many of these error sources, and theoretically connect them to the physics of superconducting qubits and gate operations. Equipped with this knowledge, we devise compiler strategies to suppress these errors using dynamical decoupling or error compensation into the rest of the circuit. Importantly, these strategies are successful when the context at each layer of computation is taken into account: how qubits are connected, what crosstalk terms exist on the device, and what gates or idle periods occur in that layer. Our context-aware compiler thus suppresses some dominant sources of error, making further error mitigation or error correction substantially less expensive. For example, our experiments show an increase of 18.5% in layer fidelity for a candidate 10-qubit circuit layer compared to context-unaware suppression. Owing to the exponential nature of error mitigation, this error suppression translates to several orders of magnitude reduction of sampling overhead for a circuit consisting of a moderate number of layers.

Early Fault-Tolerant Quantum Algorithms in Practice: Application to Ground-State Energy Estimation

Presenter: Oriel Kiss (CERN)

Authors: Oriel Kiss, Utkarsh Azad, Borja Requena, Alessandro Roggero, David Wakeham, Juan Miguel Arrazola

We investigate the feasibility of early fault-tolerant quantum algorithms with a focus on ground-state energy estimation problems. In particular, we examine the computation of the cumulative distribution function (CDF) of the spectral measure of a Hamiltonian and the identification of its discontinuities. Scaling these methods to larger system sizes reveals three key challenges: the smoothness of the CDF for large supports, the lack of tight lower bounds on the overlap with the true ground state, and the difficulty of preparing high-quality initial states. We propose a signal processing approach for finding the inflection point of the CDF. We advocate for improving classical estimates by targeting the low-energy support of the initial state. Additionally, we provide quantitative resource estimates, demonstrating a quadratic improvement in the maximum number of samples required to detect a specified change in the CDF compared to standard approaches. Our numerical experiments, conducted on a 26-qubit fully connected Heisenberg model, leverage a truncated DMRG initial state with low bond dimension. The results show that the predictions are accurate while requiring several orders of magnitude fewer samples than theoretical estimates suggest. Therefore CDF-based quantum algorithms are a practical and resource-efficient alternative to QPE in resource-constrained scenarios.

Low-overhead magic state distillation with color codes

Presenter: Seok-Hyung Lee (The University of Sydney)

Authors: Seok-Hyung Lee, Felix Thomsen, Nicholas Fazio, Benjamin J. Brown, Stephen D. Bartlett

Fault-tolerant implementation of non-Clifford gates is a major challenge for achieving universal fault-tolerant quantum computing with quantum error-correcting codes. Magic state distillation is the most well-studied method for this but requires significant resources. Hence, it is crucial to tailor and optimize magic state distillation for specific codes from both logical- and physical-level perspectives. In this work, we perform such optimization for two-dimensional color codes, which are promising due to their higher encoding rates compared to surface codes, transversal implementation of Clifford gates, and efficient lattice surgery. We propose two distillation schemes based on the 15-to-1 distillation circuit and lattice surgery, which differ in their methods for handling faulty rotations. Our first scheme uses faulty T-measurement, offering resource efficiency when the target infidelity is above a certain threshold ($\sim 35p^3$ for physical error rate p). To achieve lower infidelities while maintaining resource efficiency, our second scheme exploits a distillation-free fault-tolerant magic state preparation protocol, achieving significantly lower infidelities (e.g., $\sim 10^{-19}$ for $p = 10^{-4}$) than the first scheme. Notably, our schemes outperform the best existing magic state distillation methods for color codes by up to about two orders of magnitude in resource costs for a given achievable target infidelity.

Unification of Finite Symmetries in Simulation of Many-body Systems on Quantum Computers

Presenter: Yuan Liu (North Carolina State University)

Authors: Victor M. Bastidas, Nathan Fitzpatrick, K. J. Joven, Zane M. Rossi, Shariful Islam, Troy Van Voorhis, Isaac L. Chuang, Yuan Liu

Symmetry is fundamental in the description and simulation of quantum systems. Leveraging symmetries in classical simulations of many-body quantum systems often results in an exponential overhead due to the exponentially growing size of some symmetry groups as the number of particles increases. Quantum computers hold the promise of achieving exponential speedup in simulating quantum many-body systems; however, a general method for utilizing symmetries in quantum simulations has not yet been established. In this work, we present a unified framework for incorporating symmetry groups into the simulation of many-body systems on quantum computers. The core of our approach lies in the development of efficient quantum circuits for symmetry-adapted projection onto irreducible representations of a group or pairs of commuting groups. We provide resource estimations for common groups, including the cyclic and permutation groups. Our algorithms demonstrate the capability to prepare coherent superpositions of symmetry-adapted states and to perform quantum evolution across a wide range of models in condensed matter physics and ab initio quantum chemistry. In addition, we present a discussion of major open problems regarding the use of symmetries in digital quantum simulations of many-body systems, paving the way for future systematic investigations into leveraging symmetries for exponential quantum advantage.

Demonstrating dynamic surface codes

Presenter: Alec Eickbusch (Google)

Authors: Matthew McEwen, Alec Eickbusch, Alexis Morvan

QEC is traditionally implemented by repeatedly applying static syndrome checks, permitting correction of logical information. Recently, the development of time-dynamic approaches to error correction has uncovered new codes and new code implementations. In this work, we experimentally demonstrate three time-dynamic implementations of the surface code, each offering a unique solution to hardware design challenges and introducing flexibility in surface code realization. First, we embed the surface code on a hexagonal lattice, reducing the necessary couplings per qubit from four to three. Second, we *walk* a surface code, swapping the role of data and measure qubits each round, achieving error correction with built-in removal of accumulated non-computational errors. Finally, we realize the surface code using iSWAP gates instead of the traditional CNOT, extending the set of viable gates for error correction without additional overhead. We measure the error suppression factor when scaling from distance-3 to distance-5 codes, achieving state of the art error suppression for each. With detailed error budgeting, we explore their performance trade-offs and implications for hardware design. This work demonstrates that dynamic circuit approaches satisfy the demands for fault-tolerance and opens new alternative avenues for scalable hardware design

Tomography of bosonic systems and optimal estimates of the trace distance between Gaussian states

Presenter: Antonio Anna Mele (Freie Universität Berlin)

Authors: Lennart Bittel, Jens Eisert, Vittorio Giovannetti, Ludovico Lami, Lorenzo Leone, Antonio Anna Mele, Francesco Anna Mele, Salvatore Francesco Emanuele Oliviero, Salvatore Tirone

Quantum state tomography is a fundamental and practical task in quantum information. Here, we analyse the ultimate achievable performance of tomography of continuous-variable systems. We prove that tomography of these systems is extremely inefficient in terms of time resources, much more so than tomography of finite-dimensional systems: not only does the minimum number of state copies needed for tomography scale exponentially with the number of modes, but it also exhibits a dramatic scaling with the trace-distance error, even for low-energy states, in stark contrast with the finite-dimensional case. On a more positive note, we prove that tomography of Gaussian states is efficient. To accomplish this, we fully solve a fundamental question for the field of continuous-variable quantum information: if we know with a certain error the first and second moments of an unknown Gaussian state, what is the resulting trace-distance error that we make on the state? We do this by finding optimal bounds on the trace distance between two Gaussian states in terms of the norm distance of their first moments and covariance matrices. Lastly, we demonstrate that tomography of non-Gaussian states prepared through Gaussian unitaries and a few local non-Gaussian evolutions is efficient and experimentally feasible.

SPAM-free sound certification of quantum gates via quantum system quizzing

Presenter: Nikolai Miklin (Hamburg University of Technology)

Authors: Nikolai Miklin, Jan Nöller, Martin Kliesch, Mariami Gachechiladze

The rapid advancement of quantum hardware necessitates the development of reliable methods to certify its correct functioning. However, existing certification tests often fall short: they either rely on flawless state preparation and measurement, or fail to provide soundness guarantees, meaning that they do not ensure the correct implementation of the target operation by a quantum device. We introduce an approach, which we call quantum system quizzing, for certification of quantum gates in a practical server-user scenario, where a classical user tests the results of exact quantum computations performed by a quantum server. Importantly, this approach is free from state preparation and measurement (SPAM) errors. For a wide range of relevant gates, including a gate set universal for quantum computation, we demonstrate that our approach offers soundness guarantees based solely on the dimension assumption. Additionally, for a highly-relevant single-qubit phase gate - which corresponds experimentally to a $\pi/2$ -pulse - we prove that the method's sample complexity scales inverse-linearly relative to the average gate infidelity. By combining the SPAM-error-free and sound notion of certification with practical applicability, our approach paves the way for promising research into efficient and reliable certification methods for quantum computation.

Experimental measurement and a physical interpretation of quantum shadow enumerators

Presenter: Daniel Miller (Freie Universität Berlin)

Authors: Daniel Miller, Kyan Levi, Lukas Postler, Alex Steiner, Lennart Bittel, Gregory A.L. White, Yifan Tang, Eric J. Kuehnke, Antonio A. Mele, Sumeet Khatri, Lorenzo Leone, Jose Carrasco, Christian D. Marcinia, Ivan Pogorelov, Milena Guevara-Bertsch, Robert Freund, Rainer Blatt, Philipp Schindler, Thomas Monz, Martin Ringbauer, Jens Eisert

Throughout its history, the theory of quantum error correction has heavily benefited from translating classical concepts into the quantum setting. In particular, classical notions of weight enumerators, which relate to the performance of an error-correcting code, and MacWilliams' identity, which helps to compute enumerators, have been generalized to the quantum case. In this work, we establish a distinct relationship between the theoretical machinery of quantum weight enumerators and a seemingly unrelated physics experiment: we prove that Rains' quantum shadow enumerators—a powerful mathematical tool—arise as probabilities of observing fixed numbers of triplets in a Bell sampling experiment. This insight allows us to develop here a rigorous framework for the direct measurement of quantum weight enumerators, thus enabling experimental and theoretical studies of the entanglement structure of any quantum error-correcting code or state under investigation. On top of that, we derive concrete sample complexity bounds and physically-motivated robustness guarantees against unavoidable experimental imperfections. Finally, we experimentally demonstrate the feasibility of directly measuring weight enumerators on a trapped-ion quantum computer. Our experimental findings are in good agreement with theoretical predictions and illuminate how entanglement theory and quantum error correction crossfertilize each other once Bell sampling experiments are combined with the theoretical machinery of quantum weight enumerators.

On-Chip Verified Quantum Computation with an Ion-Trap Quantum Processing Unit

Presenter: Daniel Mills (Quantinuum)

Authors: Daniel Mills, Cica Gustiani, Dominik Leichtle, Jonathan Miller, Ross Grassie, Elham Kashefi

We present and experimentally demonstrate a novel approach to verifying and benchmarking quantum computing, implementing it on an ion-trap quantum computer. Unlike previous information-theoretically secure verification protocols, which typically require quantum communication between Client and Server, our approach is implemented entirely on-chip. This eliminates the need for a quantum client and significantly enhances practicality. We perform tomography to justify the additionally required assumption that the noise is independent of the secret used to prepare the Server’s single-qubit states. We quantify the soundness error which may be caused by residual secret dependencies. We demonstrate our protocol on the 20-qubit Quantinuum H1-1 ion-trap quantum processing unit, using qubit measurements and resets to construct measurement patterns with up to 52 vertices. To our knowledge, these are the largest verified measurement-based quantum computations performed to date. Our results pave the way for more accessible and efficient verification and benchmarking strategies in near-term quantum devices, enabling robust performance assessment without the added cost of external quantum infrastructure.

Self-correcting GKP qubit in a superconducting circuit with an oscillating voltage bias

Presenter: Frederik Nathan (NNF Quantum Computing Programme, University of Copenhagen)

Authors: Frederik Nathan, Max Geier

We propose a simple circuit architecture for a dissipatively error corrected Gottesman-Kitaev-Preskill (GKP) qubit. The device consists of a electromagnetic resonator with impedance $h/2e^2 \approx 12.91\text{k}\Omega$ connected to a Josephson junction with a voltage bias oscillating at twice the resonator frequency. For large drive amplitudes, the circuit is effectively described by the GKP stabilizer Hamiltonian, whose low-energy subspace forms the code space for a qubit protected against phase-space local noise. The GKP states in the codespace can be dissipatively stabilized and error corrected by coupling the resonator to a bath through a bandpass filter; a resulting side-band cooling effect stabilizes the system in the GKP code space, dissipatively correcting it against both bit and phase flip errors. Simulations show that this dissipative error correction can enhance coherence time by factor ~ 1000 with NbN-based junctions, for operating temperatures in the $\sim 100\text{mK}$ range. The scheme can be used to stabilize both square- and hexagonal-lattice GKP codes. Finally, a Josephson current based readout scheme, and dissipatively corrected single-qubit Clifford gates are proposed.

Trotter error and gate complexity of the SYK and sparse SYK models

Presenter: Yiyuan Chen (QuSoft / University of Amsterdam)

Authors: Yiyuan Chen, Jonas Helsen, Maris Ozols

The Sachdev–Ye–Kitaev (SYK) model is a prominent model of strongly interacting fermions that serves as a toy model of quantum gravity and black hole physics. In this work, we study the Trotter error and gate complexity of the quantum simulation of the SYK model using Lie–Trotter–Suzuki formulas. Building on recent results by Chen and Brandão [CB24] – in particular their uniform smoothing technique for random matrix polynomials –, we derive bounds on the first- and higher-order Trotter error of the SYK model, and subsequently find near-optimal gate complexities for simulating these models using Lie–Trotter–Suzuki formulas. For the k -local SYK model on n Majorana fermions, our gate complexity estimates for the first-order Lie–Trotter–Suzuki formula scales with $O(n^{k+\frac{5}{2}}t^2)$ for even k and $O(n^{k+3}t^2)$ for odd k , and the gate complexity of simulations using higher-order formulas scales with $O(n^{k+\frac{1}{2}}t)$ for even k and $O(n^{k+1}t)$ for odd k . Given that the SYK model has $\Theta(n^k)$ terms, these estimates are close to optimal. These gate complexities can be further improved upon in the context of simulating the time-evolution of an arbitrary fixed input state $|\psi\rangle$, leading to a $O(n^2)$ -reduction in gate complexity for first-order formulas and $O(\sqrt{n})$ -reduction for higher-order formulas.

Error mitigation and circuit division for early fault-tolerant quantum phase estimation

Presenter: Stefano Polla (Leiden University)

Authors: Alicja Dutkiewicz, Stefano Polla, Maximilian Scheurer, William J. Huggins, Christian Gogolin, Thomas O’Brien

As fully fault-tolerant quantum computers capable of solving useful problems remain a future goal, we anticipate an era of “early fault tolerance” allowing for limited error correction. We propose a framework for designing early fault-tolerant algorithms by trading between error correction overhead and residual logical noise, and apply it to quantum phase estimation (QPE). We develop a quantum-Fourier-transform (QFT)-based QPE technique that is robust to global depolarising noise and outperforms the previous state of the art at low and moderate noise rates. We further develop a data processing technique, Explicitly Unbiased Maximum Likelihood Estimation (EUMLE), allowing us to mitigate arbitrary error on QFT-based QPE schemes in a consistent, asymptotically normal way. This extends quantum error mitigation techniques beyond expectation value estimation, which was labeled an open problem for the field. Applying this scheme to the ground state problem of the two-dimensional Hubbard model and various molecular Hamiltonians, we find we can roughly halve the number of physical qubits with a $\approx 10\times$ wall-clock time overhead, but further reduction causes a steep runtime increase. This work provides an end-to-end analysis of early fault-tolerance cost reductions and space-time trade-offs and identifies which areas can be improved in the future.

Guarantees and limitations for warm starts and iterative methods in variational quantum computing

Presenter: Ricard Puig (École Polytechnique Fédérale de Lausanne)

Authors: Ricard Puig, Hela Mhiri, Marc Drudis, Sacha Lerch, Manuel S. Rudolph, Thiparat Chotibut, Supanut Thanasilp, Zoë Holmes

Barren plateaus are fundamentally a statement about quantum loss landscapes on average but there can exist patches of barren plateau landscapes with substantial gradients. This has motivated the study of warm starts whereby the algorithm is cleverly initialized closer to a minimum. Numerical studies indicate that these methods may be promising. In parallel, analytic studies have proven that small angle initializations, whereby the parameterized quantum circuit is initialized in a small region typically around identity or a Clifford, can exhibit non-exponentially vanishing gradients. However, a good solution may be far from this region and thus these methods can (in full generality) only work on a vanishing fraction of problem instances. In this joint submission we present general analysis of warm starts for physically-motivated ansatze and iterative training strategies. Our work thus suggests that while there are hopes to be able to warm-start variational quantum algorithms, any initialization strategy that cannot get increasingly close to the region of attraction with increasing problem size is likely to prove challenging to train.

Chemistry Beyond Exact Solutions on a Quantum-Centric Supercomputer

Presenter: Javier Robledo Moreno (IBM Quantum)

Authors: Javier Robledo Moreno, Mario Motta, Holger Haas, Ali Javadi-Abhari, Petar Jurcevic, William Kirby, Simon Martiel, Kunal Sharma, Sandeep Sharma, Tomonori Shirakawa, Iskandar Situdikov, Rong-Yang Sun, Kevin J. Sung, Maika Takita, Minh C. Tran, Seiji Yunoki, Antonio Mezzacapo

A universal quantum computer can be used as a simulator capable of predicting properties of diverse quantum systems. Electronic structure problems in chemistry are practical use cases around the hundred-qubit mark, and appear to be promising application since current quantum processors have reached these sizes. However, mapping these use cases onto quantum computers yields deep circuits, and for pre-fault-tolerant quantum processors, the large number of measurements to estimate molecular energies leads to prohibitive runtimes. As a result, realistic chemistry is out of reach of current quantum computers in isolation. Here we present quantum computations of chemistry that go beyond problem sizes amenable to current state-of-the-art exact diagonalization methods. Our results are obtained in a quantum-centric supercomputing architecture, using classical resources to assist an IBM Heron quantum processor. We simulate the N₂ triple bond breaking in a correlation-consistent basis set (cc-pVDZ), using 58 qubits and 5168 (1756 2-qubit) quantum gates, and the active-space electronic structure of [2Fe–2S] and [4Fe–4S] clusters, using 45 and 77 qubits respectively, with 3140 (1070) and 10540 (3560) quantum gates.

This submission will be presented as part of a merged talk together with ‘Sample-based Krylov Quantum Diagonalization’.

Approximating dynamical correlation functions with constant depth quantum circuits

Presenter: Raul A Santos (Phasecraft)

Authors: Raul A Santos, Reinis Irmejs

Dynamical correlation functions (DFCs) are important signals that characterize the properties of quantum systems. They are obtained by time-evolving a perturbation of an eigenstate of the system, typically the ground-state. In this work, we study approximations of these correlation functions that do not require time-dynamics. We show that having access to a circuit that prepares an eigenstate of the Hamiltonian, it is possible to approximate the DFCs up to exponential accuracy in the complex frequency domain, on a strip above the real line. We achieve this by exploiting their continued fraction representation as functions of frequency ω , where the level k approximant can be obtained by measuring a weight $O(k)$ operator on the eigenstate of interest. In the complex ω plane, this approach allows to determine approximations to DCFs with accuracy that increases exponentially with k . We analyse two algorithms to generate the continuous fraction representation in scalar or matrix form, starting from one or many initial operators. We prove that these algorithms generate an exponentially accurate approximation of the dynamical correlation functions on a region sufficiently far away from the real frequency axis. We present numerical evidence of these theoretical results through simulations of small lattice systems.

Fault-Tolerant Logical Clifford Gates from Code Automorphisms

Presenter: Hasan Sayginel (University College London)

Authors: Hasan Sayginel, Stergios Koutsoumpas, Mark Webster, Abhishek Rajput, Dan E. Browne

We study the implementation of fault-tolerant logical Clifford gates on stabilizer quantum error correcting codes based on their symmetries. Our approach is to map the stabilizer code to a binary linear code, compute its automorphism group, and impose constraints based on the Clifford operators permitted. We provide a rigorous formulation of the method for finding automorphisms of stabilizer codes and generalize ZX-dualities to non-CSS codes. We provide a Python package implementing our algorithms which uses the computational algebra system MAGMA. Our algorithms map automorphism group generators to physical circuits, calculate Pauli corrections based on the destabilizers of the code, and determine their logical action. We discuss the fault tolerance of the circuits and include examples of gates through automorphisms for the $[[4, 2, 2]]$ and perfect $[[5, 1, 3]]$ codes, bivariate bicycle codes, and the best known distance codes.

Fermion-qubit fault-tolerant quantum computing

Presenter: Alexander Schuckert (University of Maryland)

Authors: Alexander Schuckert, Eleanor Crane, Alexey V. Gorshkov, Mohammad Hafezi, Michael J. Gullans

Simulating the dynamics of electrons is one of the most promising applications of fault-tolerant quantum computers. However, mapping time evolution under fermionic Hamiltonians to qubit gates is challenging. We introduce fermion-qubit fault-tolerant quantum computing, a framework which removes this overhead altogether. Using native fermionic operations we first construct a repetition code which corrects phase errors only. We then engineer a fermionic color code which corrects for both phase and loss errors. We show how to realize a universal fermionic gate set in this code, including transversal Clifford gates. Interfacing with qubit color codes we realize qubit-fermion fault-tolerant computation, which allows for qubit-controlled fermionic time evolution, a crucial subroutine in state-of-the-art quantum algorithms for simulating fermions. We show how our framework can be implemented in neutral atoms, overcoming the apparent inability of neutral atoms to implement non-number-conserving gates by introducing a neutral-atom braiding gate using photodissociation of bosonic molecules. As an application, we consider the fermionic fast Fourier transform, an important subroutine for simulating crystalline materials, finding an exponential improvement in circuit depth from $O(N)$ to $O(\log(N))$ with respect to lattice site number N and a linear improvement from $O(N^2)$ to $O(N \log(N))$ in Clifford gate complexity compared to state-of-the-art qubit-only approaches.

Entanglement-enhanced learning of quantum processes at scale

Presenter: Alireza Seif (IBM Quantum)

Authors: Alireza Seif, Senrui Chen, Swarnadeep Majumder, Haoran Liao, Derek S Wang, Moein Malekakhlagh, Ali Javadi-Abhari, Liang Jiang, Zlatko K Minev

Learning unknown processes affecting a quantum system reveals underlying physical mechanisms and enables suppression, mitigation, and correction of unwanted effects. Generally, learning quantum processes requires exponentially many measurements. We show how entanglement with an ideal auxiliary quantum memory can provide an exponential advantage in learning certain quantum processes. We discuss practical limitations of entanglement-enhanced protocols, and show, both theoretically and experimentally, that even in the presence of noise, entanglement with auxiliary quantum memory combined with error mitigation considerably enhances the learning of quantum processes.

Sample-based Krylov Quantum Diagonalization

Presenter: Javier Robledo Moreno (IBM)

Authors: Jeffery Yu, Joseph Iosue, Luke Bertels, Daniel Claudino, Bryce Fuller, Peter Groszkowski, Travis S Humble, Petar Jurcevic, William Kirby, Thomas A. Maier, Mario Motta, Bibek Pokharel, Javier Robledo Moreno, Alireza Seif, Amir Shehata, Kevin J. Sung, Minh C. Tran, Vinay Tripathi, Antonio Mezzacapo, Kunal Sharma

Approximating the ground state of many-body systems is a key computational bottleneck underlying important applications in physics and chemistry. It has long been viewed as a promising application for quantum computers. However, the most widely known quantum algorithm for ground state approximation, quantum phase estimation, is out of reach for near-term and early fault-tolerant quantum devices due to its high circuit-depth requirements. Subspace algorithms represent alternatives for ground-state problems, which are feasible for pre-fault-tolerant quantum computers. Here, we introduce a new quantum algorithm which combines the strengths of existing ideas on quantum subspaces: a classical diagonalization based on quantum samples, and subspaces constructed with quantum Krylov states. First, we prove that our algorithm converges in polynomial time under the working assumptions of Krylov quantum diagonalization and a sparse ground state. Then, we carry out the largest ground-state quantum simulation of the single-impurity Anderson model on a system with 41 bath sites, using 85 qubits and up to $6 \cdot 10^3$ two-qubit gates on a Heron quantum processor, showing excellent agreement with density matrix renormalization group calculations. Lastly, we show numerical investigations of spin and chemistry Hamiltonians, which indicate that our method can outperform existing Krylov quantum diagonalization in the presence of shot noise.

This submission will be presented as part of a merged talk together with ‘Chemistry Beyond Exact Solutions on a Quantum-Centric Supercomputer’.

Optimization by Decoded Quantum Interferometry

Presenter: Noah John Shutty (Google Quantum AI)

Authors: Stephen Jordan, Noah John Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei Isakov, Ryan Babbush

We introduce Decoded Quantum Interferometry (DQI), a quantum algorithm for reducing classical optimization problems to classical decoding problems by exploiting structure in the Fourier spectrum of the objective function. DQI reduces sparse max-XORSAT to decoding LDPC codes, which can be achieved using powerful classical algorithms such as Belief Propagation (BP). As an initial benchmark, we compare DQI using belief propagation decoding against classical optimization via simulated annealing. In this setting we present evidence that, for a certain family of max-XORSAT instances, DQI with BP decoding achieves a better approximation ratio on average than simulated annealing, although not better than specialized classical algorithms tailored to those instances. We also analyze a combinatorial optimization problem corresponding to finding polynomials that intersect the maximum number of points. There, DQI efficiently achieves a better approximation ratio than any polynomial-time classical algorithm known to us, thus realizing an apparent exponential quantum speedup. Finally, we show that the problem defined by Yamakawa and Zhandry in order to prove an exponential separation between quantum and classical query complexity is a special case of the optimization problem efficiently solved by DQI.

Contextual Subspace Auxiliary-Field Quantum Monte Carlo: Improved bias with reduced quantum resources

Presenter: Fedor Simkovic (IQM Quantum Computers)

Authors: Matthew Kiser, Matthias Beuerle, Fedor Simkovic

Using trial wavefunctions prepared on quantum devices to reduce the bias of auxiliary-field quantum Monte Carlo (QC-AFQMC) has established itself as a promising hybrid approach to the simulation of strongly correlated many body systems. Here, we further reduce the required quantum resources by decomposing the trial wavefunction into classical and quantum parts, respectively treated by classical and quantum devices, within the contextual subspace projection formalism. Importantly, we show that our algorithm is compatible with the recently developed matchgate shadow protocol for efficient overlap calculation in QC-AFQMC. Investigating the nitrogen dimer and the reductive decomposition of ethylene carbonate in lithium-based batteries, we observe that our method outperforms a number of established algorithm for ground state energy computations, while reaching chemical accuracy with less than half of the original number of qubits.

Logical Operators and Fold-Transversal Gates of Bivariate Bicycle Codes

Presenter: Vincent Steffan (IQM Germany)

Authors: Vincent Steffan, Jens Niklas Eberhardt

Quantum low-density parity-check (qLDPC) codes offer a promising route to scalable fault-tolerant quantum computation with constant overhead. Recent advancements have shown that qLDPC codes can outperform the quantum memory capability of surface codes even with near-term hardware. The question of how to implement logical gates fault-tolerantly for these codes is still open. We present new examples of high-rate bivariate bicycle (BB) codes with enhanced symmetry properties. These codes feature explicit nice bases of logical operators (similar to toric codes) and support fold-transversal Clifford gates without overhead. As examples, we construct $[[98, 6, 12]]$ and $[[162, 8, 12]]$ BB codes which admit interesting fault-tolerant Clifford gates. Our work also lays the mathematical foundations for explicit bases of logical operators and fold-transversal gates in quantum two-block and group algebra codes, which might be of independent interest.

Localizing multipartite entanglement with local and global measurements

Presenter: Christopher Vairogs (University of Illinois Urbana-Champaign)

Authors: Christopher Vairogs, Samihr Hermes, Felix Leditzky

We study the task of localizing multipartite entanglement in pure quantum states onto a subsystem by measuring the remaining systems. To this end, we fix a multipartite entanglement measure and consider two quantities: the multipartite entanglement of assistance (MEA), defined as the entanglement measure averaged over the post-measurement states and maximized over arbitrary measurements; and the localizable multipartite entanglement (LME), defined in the same way but restricted to only local single-system measurements. First, we prove easily computable upper and lower bounds on MEA and LME. Using these bounds, we investigate the typical behavior of entanglement localization via concentration inequalities for the MEA evaluated on Haar-random states. We then turn our attention to protocols that transform graph states. Our bounds produce a criterion in terms of a matrix equation that in many cases efficiently deduces what transformations between graph states are possible, a generally challenging computational problem. We generalize this analysis to weighted graph states and show that our entanglement localization framework certifies the near-optimality of recently discussed local-measurement protocols to transform uniformly weighted line graph states into GHZ states. Finally, we demonstrate how our MEA and LME quantities can be used to detect phase transitions in transversal field Ising models.

Derandomized shallow shadows: Efficient Pauli learning with bounded-depth circuits

Presenter: Katherine Van Kirk (Harvard University)

Authors: Katherine Van Kirk, Christian Kokail, Jonathan Kunjummen, Hong-Ye Hu, Yanting Teng, Madelyn Cain, Jacob Taylor, Suzanne Yelin, Hannes Pichler, Mikhail Lukin

Efficiently estimating large numbers of non-commuting observables is an important subroutine of many quantum science tasks. We present the derandomized shallow shadow (DSS) algorithm for efficiently learning a large set of non-commuting observables, using shallow circuits to rotate into measurement bases. Exploiting tensor network techniques to ensure polynomial scaling of classical resources, our algorithm outputs a set of shallow measurement circuits that approximately minimizes the sample complexity of estimating a given set of Pauli strings. We numerically demonstrate systematic improvement, in comparison with state-of-the-art techniques, for energy estimation of quantum chemistry benchmarks and verification of quantum many-body systems, and we observe DSS's performance consistently improves as one allows deeper measurement circuits. These results indicate that in addition to being an efficient, low-depth, stand-alone algorithm, DSS can also benefit many larger quantum algorithms requiring estimation of multiple non-commuting observables.

Quantum Algorithm for Reversing Unknown Unitary Evolutions

Presenter: Xin Wang (The Hong Kong University of Science and Technology (Guangzhou))

Authors: Yu-Ao Chen, Yin Mo, Yingjian Liu, Lei Zhang, Xin Wang

Reversing an unknown quantum evolution is of central importance to quantum information processing and fundamental physics, yet it remains a formidable challenge as conventional methods necessitate an infinite number of queries to fully characterize the quantum process. Here we introduce the Quantum Unitary Reversal Algorithm (QURA), a deterministic and exact approach to universally reverse arbitrary unknown unitary transformations using $\mathcal{O}(d^2)$ calls of the unitary, where d is the system dimension. Our quantum algorithm resolves a fundamental problem of time-reversal simulations for closed quantum systems by confirming the feasibility of reversing any unitary evolution without knowing the exact process. The algorithm also provides the construction of a key oracle for unitary inversion in many quantum algorithm frameworks, such as quantum singular value transformation. It notably reveals a sharp boundary between the quantum and classical computing realms and unveils a quadratic quantum advantage in computational complexity for this foundational task.

Observable-Driven Speed-ups in Quantum Simulations

Presenter: Wenjun Yu (The University of Hong Kong)

Authors: Wenjun Yu, Jue Xu, Qi Zhao

As quantum technology advances, quantum simulation becomes increasingly promising, with significant implications for quantum many-body physics and quantum chemistry. Despite being one of the most accessible simulation methods, the product formula encounters challenges due to the pessimistic gate count estimation. In this work, we elucidate how observable knowledge can accelerate quantum simulations. By focusing on specific families of observables, we reduce product-formula simulation errors and gate counts in both short-time and arbitrary-time scenarios. For short-time simulations, we deliberately design and tailor product formulas to achieve size-independent errors for local and certain global observables. In arbitrary-time simulations, we reveal that Pauli-summation structured observables generally reduce average errors. Specifically, we obtain quadratic error reductions proportional to the number of summands for observables with evenly distributed Pauli coefficients. Our advanced error analyses, supported by numerical studies, indicate improved gate count estimation. We anticipate that the explored speed-ups can pave the way for efficiently realizing quantum simulations and demonstrating advantages on near-term quantum devices.

Efficient preparation of Dicke states

Presenter: Jeffery Yu (University of Maryland)

Authors: Jeffery Yu, Yuxin Wang, Sean Muleady, Nathan Schine, Alexey Gorshkov, Andrew Childs

We present an algorithm utilizing mid-circuit measurement and feedback that prepares Dicke states with polylogarithmically many ancillas and polylogarithmic depth. Our algorithm uses only global mid-circuit projective measurements and adaptively chosen global rotations. This improves over prior work that was only efficient for Dicke states of low weight, or was not efficient in both depth and width. Our algorithm can also naturally be implemented in a cavity QED context using polylogarithmic time, zero ancillas, and atom-photon coupling scaling with the square root of the system size.

Auxiliary-free replica shadow estimation

Presenter: You Zhou (Fudan University)

Authors: Qing Liu, Zihao Li, Xiao Yuan, Huangjun Zhu, You Zhou

Efficiently measuring nonlinear properties is a significant yet challenging task from quantum information processing to many-body physics. Current methodologies often suffer from an exponential sampling cost or require auxiliary qubits and deep quantum circuits. To address these limitations, we propose an efficient auxiliary-free replica shadow (AFRS) framework, which leverages the power of the joint entangling operation on a few input replicas while integrating the mindset of shadow estimation. We rigorously prove that AFRS can offer exponential improvements in estimation accuracy compared with the conventional shadow method, and facilitate the simultaneous estimation of various nonlinear properties, unlike the destructive swap test. Additionally, we introduce an advanced local-AFRS variant tailored to estimating local observables with constant-depth quantum circuits, significantly simplifying the experimental implementation. Our work paves the way for the application of AFRS on near-term quantum hardware, opening new avenues for efficient and practical quantum measurements.

Classifying the gates of the Clifford hierarchy

Presenter: Nadish de Silva (Simon Fraser University)

Authors: Imin Chen, Nadish de Silva, Oscar Lautsch

The Clifford hierarchy is a nested sequence of sets of quantum gates that can be fault-tolerantly performed using gate teleportation within standard quantum error correction schemes. The groups of Pauli and Clifford gates constitute the first and second ‘levels’, respectively. Non-Clifford gates from the third level or higher are necessary for achieving fault-tolerant universal quantum computation. Since it was defined twenty-five years ago by Gottesman-Chuang, two questions have been studied by numerous researchers. First, precisely which gates constitute the Clifford hierarchy? Second, which subset of the hierarchy gates admit efficient gate teleportation protocols? We completely solve both questions in the case of the Clifford hierarchy for gates of one qubit or one qudit. We express every such hierarchy gate uniquely as a product of three simple gates. We find that all such hierarchy gates admit efficient gate teleportation. Our decomposition of Clifford gates as a unique product of three simple gates is of broader applicability. We then consider the more complex case of two-qudit gates and focus on third-level gates. We prove that every two-qudit third-level gate admits efficient efficient gate teleportation. Our proof leverages tools of algebraic geometry which can be applied more widely within quantum information.

Polynomial Time Quantum Gibbs Sampling for Fermi-Hubbard Model at any Temperature

Presenter: Štěpán Šmíd (Imperial College London)

Authors: Štěpán Šmíd, Richard Meister, Mario Berta, Roberto Bondesan

Recently, there have been several advancements in quantum algorithms for Gibbs sampling. These algorithms simulate the dynamics generated by an artificial Lindbladian, which is meticulously constructed to obey a detailed-balance condition with the Gibbs state of interest, ensuring it is a stationary point of the evolution, while simultaneously having efficiently implementable time steps. The overall complexity then depends primarily on the mixing time of the Lindbladian, which can vary drastically. In this work, we calculate the spectral gap of the Lindbladian for free fermions using third quantisation, and then prove a constant gap of the perturbed Lindbladian corresponding to interacting fermions up to some maximal coupling strength. This is achieved by using theorems about stability of the gap for lattice fermions. Our methods apply at any constant temperature and independently of the system size. The gap then provides an upper bound on the mixing time, proving that the purified Gibbs state of weakly interacting (quasi)-local fermionic systems of any dimension can be prepared in quasi-cubic time. We provide exact numerical simulations for small system sizes supporting the theory and also identify different suitable jump operators and filter functions for the sought-after regime of intermediate coupling in the Fermi-Hubbard model.