Booklet of abstracts

Thomas Vidick

California Institute of Technology

Tsirelson's problem and MIP*=RE

Boris Tsirelson in 1993 implicitly posed "Tsirelson's Problem", a question about the possible equivalence between two different ways of modeling locality, and hence entanglement, in quantum mechanics. Tsirelson's Problem gained prominence through work of Fritz, Navascues et al., and Ozawa a decade ago that establishes its equivalence to the famous "Connes' Embedding Problem" in the theory of von Neumann algebras. Recently we gave a negative answer to Tsirelson's Problem and Connes' Embedding Problem by proving a seemingly stronger result in quantum complexity theory. This result is summarized in the equation MIP* = RE between two complexity classes. In the talk I will present and motivate Tsirelson's problem, and outline its connection to Connes' Embedding Problem. I will then explain the connection to quantum complexity theory and show how ideas developed in the past two decades in the study of classical and quantum interactive proof systems led to the characterization (which I will explain) MIP* = RE and the negative resolution of Tsirelson's Problem. Based on joint work with Ji, Natarajan, Wright and Yuen available at arXiv:2001.04383.

Joonho Lee, Dominic Berry, Craig Gidney, William Huggins, Jarrod McClean, Nathan Wiebe and Ryan Babbush

Columbia University | Macquarie University | Google | Google Research | Google | University of Washington | Google

Efficient quantum computation of chemistry through tensor hypercontraction

We show how to achieve the highest efficiency yet for simulations with arbitrary basis sets by using a representation of the Coulomb operator known as tensor hypercontraction (THC). We use THC to express the Coulomb operator in a non-orthogonal basis, which we are able to block encode by separately rotating each term with angles that are obtained via QROM. Our algorithm has the best complexity scaling for an arbitrary basis, as well as the best complexity for the specific case of FeMoCo. By optimising the surface code resources, we show that FeMoCo can be simulated using about 4 million physical qubits and 3.5 days of runtime, assuming 1 -µs cycle times and physical gate error rates no worse than 0.1%.

Peter Brown, Hamza Fawzi and Omar Fawzi

ENS Lyon | University of Cambridge | ENS Lyon

New quantum Rényi divergences and their application to device-independent cryptography and quantum Shannon theory

In the analysis of quantum information processing tasks, the choice of distance measure between states or channels often plays a crucial role. This submission introduces new quantum Rényi divergences for states and channels that are based on a convex optimization program involving the matrix geometric mean. These divergences have mathematical and computational properties that make them applicable to a wide variety of problems. We use these Rényi divergences to obtain semidefinite programming lower bounds on the key rates for device-independent cryptography, and in particular we find a new bound on the minimal detection efficiency required to perform device-independent quantum key distribution without additional noisy preprocessing. Furthermore, we give several applications to quantum Shannon theory, in particular proving that adaptive strategies do not help in the strong converse regime for quantum channel discrimination and obtaining improved bounds for quantum capacities.

Anne Broadbent and Alex Bredariol Grilo

University of Ottawa | LIP6, CNRS/Sorbonne Université

QMA-hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge

We provide several advances to the understanding of the class of Quantum Merlin-Arthur proof systems (QMA), the quantum analogue of NP. Our central contribution is proving a longstanding conjecture that the Consistency of Local Density Matrices (CLDM) problem is QMA-hard under Karp reductions. The input of CLDM consists of local reduced density matrices on sets of at most k qubits, and the problem asks if there is an n-qubit global quantum state that is locally consistent with all of the k-qubit local density matrices. The containment of CLDM in QMA and the QMA-hardness under Turing reductions were proved by Liu [APPROX-RANDOM 2006]. Liu also conjectured that CLDM is QMA-hard under Karp reductions, which is desirable for applications, and we finally prove this conjecture. We establish this result using the techniques of simulatable codes of Grilo, Slofstra, and Yuen [FOCS 2019], simplifying their proofs and tailoring them to the context of QMA. In order to develop applications of CLDM, we propose a framework that we call locally simulatable proofs for QMA: this provides QMA proofs that can be efficiently verified by probing only k qubits and, furthermore, the reduced density matrix of any k-qubit subsystem of a good witness can be computed in polynomial time, independently of the witness. Within this framework, we show several advances in zero-knowledge in the quantum setting. We show for the first time a commit-and-open computational zero-knowledge proof system for all of QMA, as a quantum analogue of a ``sigma'' protocol. We then define a Proof of Quantum Knowledge, which guarantees that a prover is effectively in possession of a quantum witness in an interactive proof, and show that our zero-knowledge proof system satisfies this definition.

Ramis Movassagh and Yingkai Ouyang

IBM Quantum Research, IBM Research, MIT-IBM A.I. Lab | University of Sheffield

Constructing quantum codes from any classical code and their embedding in ground space of local Hamiltonians

We introduce a framework for constructing a quantum error correcting code from *any* classical error correcting code. This includes CSS codes and goes beyond the stabilizer formalism to allow quantum codes to be constructed from classical codes that are not necessarily linear or self-orthogonal. We give an algorithm that explicitly constructs quantum codes with linear distance and constant rate from classical codes with a linear distance and rate. As illustrations for small size codes, we obtain Steane's 7-qubit code uniquely from Hamming's [7,4,3] code, and obtain other error detecting quantum codes from other explicit classical codes of length 4 and 6. Motivated by quantum LDPC codes and the use of physics to protect quantum information, we introduce a new 2-local frustration free quantum spin chain Hamiltonian whose ground space we analytically characterize completely. By mapping classical codewords to basis states of the ground space, we utilize our framework to demonstrate that the ground space contains explicit quantum codes with linear distance. This side-steps the Bravyi-Terhal no-go theorem because our work allows for more general quantum codes beyond the stabilizer and/or linear codes. This model may be called an example of *subspace* quantum LDPC codes with linear distance.

Kyungjoo Noh, Stefano Pirandola and Liang Jiang

AWS Center for Quantum Computing | University of York | University of Chicago

Enhanced energy-constrained quantum communication over bosonic Gaussian channels using multi-channel strategies

Quantum communication is an important branch of quantum information science, promising unconditional security to classical communication and providing the building block of a future large-scale quantum network. Noise in realistic quantum communication channels imposes fundamental limits on the communication rates of various quantum communication tasks. It is therefore crucial to identify or bound the quantum capacities of a quantum channel. Here, we consider Gaussian channels that model energy loss and thermal noise errors in realistic optical and microwave communication channels and study their various quantum capacities in the energy-constrained scenario. We provide improved lower bounds to various energy-constrained quantum capacities of these fundamental channels and show that higher communication rates can be attained than previously believed. Specifically, we show that one can boost the transmission rates of quantum information and private classical information by using a correlated multi-mode thermal state instead of the single-mode thermal state of the same energy.

Philippe Faist, Mischa Woods, Victor V. Albert, Joseph M. Renes, Jens Eisert and John Preskill

Freie Universität Berlin | ETH Zurich | National Institute of Standards and Technology, Gaithersburg, MD | ETH Zurich | FU Berlin | Caltech

Bipartite energy-time uncertainty relation for quantum metrology with noise

Noise in quantum metrology reduces the sensitivity to which one can determine an unknown parameter in the evolution of a quantum state, such as time. Here, we consider a probe system prepared in a pure state that evolves according to a given Hamiltonian. We study the resulting local sensitivity of the probe to time after the application of a given noise channel. We show that the decrease in sensitivity due to the noise is equal to the sensitivity that the environment gains with respect to the energy of the probe. We obtain necessary and sufficient conditions for when the probe does not suffer any sensitivity loss; these conditions are analogous to, but weaker than, the Knill-Laflamme quantum error correction conditions. New upper bounds on the sensitivity of the noisy probe are obtained via our uncertainty relation, by applying known sensitivity lower bounds on the environment's system. Our time-energy uncertainty relation also generalizes to any two arbitrary parameters whose evolutions are generated by Hermitian operators. This uncertainty relation asserts a general trade-off between the sensitivities that two parties can achieve for any two respective parameters of a single quantum system, in terms of the commutator of the associated generators. We consider applications to strongly interacting many-body probes. We find probe states for general interaction graphs of Ising and Heisenberg interactions that are robust to any single located error. For a 1D spin chain with nearest-neighbor interactions subject to amplitude damping noise on each site, we verify numerically that our probe state does not lose any sensitivity to first order in the noise parameter.

Gorjan Alagic, Andrew Childs, Andrea Coladangelo, Alex Bredariol Grilo, Shih-Han Hung, Thomas Vidick and Tina Zhang // Nir Bitansky and Omri Shmueli

QuICS | University of Maryland | UC Berkeley | LIP6, CNRS/Sorbonne Université | QuICS | Caltech | Caltech // Tel Aviv University | Tel Aviv University

Non-interactive Zero-knowledge Protocols for QMA and Post-quantum Zero-knowledge in Constant Rounds

A non-interactive zero-knowledge (NIZK) proof system for a language L in NP allows a prover (who is provided with an instance x and a witness w) to compute a classical certificate for the claim that x is in L, with the following properties: 1) the protocol can be verified efficiently, and 2) the protocol does not reveal any information about w, besides the fact that it exists (i.e., that x is in L). While NIZKs are known to be impossible in the plain model (i.e., with no additional trusted resource), they are well studied in alternative models and have seen widespread application in classical cryptography. Given the importance of NIZKs, and more generally zero-knowledge protocols, in classical cryptography, there has been a recent effort to achieve such protocols for QMA, a natural quantum analog of NP. However, all previous results only achieved interactive protocols, limiting their cryptographic use. Moreover, they all rely on quantum communication between the prover and the verifier, which may be difficult to achieve. In this submission, we present two NIZK protocols for QMA in the Common Reference String (CRS) model, with additional offline setup. Both protocols are achieved through the homomorphic computation of classical NIZKs for NP, and rely on the hardness of the Learning With Errors problem. However, each of them then combines this core idea with different (seemingly incomparable) techniques: 1) our first protocol makes use of quantum teleportation and quantum communication in an offline setup phase, with a classical online phase; our second protocol leverages techniques for classical verification of quantum computations, and is the only known NIZK for QMA to be completely classical, as well as reusable, meaning that a single setup allows to prove many theorems. Security of the latter is in the Quantum Random Oracle model. // We construct the first constant-round zero-knowledge classical argument for NP secure against quantum attacks. We assume the existence of Quantum Fully Homomorphic Encryption and other standard primitives, known based on the Learning with Errors Assumption for quantum algorithms. As a corollary, we also obtain the first constant-round zero-knowledge quantum argument for QMA. At the heart of our protocol is a new no-cloning non-black-box simulation technique.

Zvika Brakerski and Henry Yuen

Weizmann Institute of Science | University of Toronto

Quantum Garbled Circuits

We present a garbling scheme for quantum circuits, thus achieving a decomposable randomized encoding scheme for quantum computation. Specifically, we show how to compute an encoding of a given quantum circuit and quantum input, from which it is possible to derive the output of the computation and nothing else. In the classical setting, garbled circuits (and randomized encodings in general) are a versatile cryptographic tool with many applications such as secure multiparty computation, delegated computation, depth-reduction of cryptographic primitives, complexity lower-bounds, and more. However, a quantum analogue for garbling general circuits was not known prior to this work. We hope that our quantum randomized encoding scheme can similarly be useful for applications in quantum computing and cryptography. To illustrate the usefulness of quantum randomized encoding, we use it to design a conceptually-simple zero-knowledge (ZK) proof system for the complexity class QMA. Our protocol has the so-called Σ format with a single-bit challenge, and allows the inputs to be delayed to the last round. The only previously-known ZK Σ-protocol for QMA is due to Broadbent and Grilo (FOCS 2020), which does not have the aforementioned properties.

Honghao Fu, Carl Miller and William Slofstra

QUICS, University of Maryland | QUICS, University of Maryland, and National Institute of Standards and Technology | University of Waterloo

The membership problem of constant-sized quantum correlations is undecidable

When two spatially separated parties make measurements on an unknown entangled quantum state, what correlations can they achieve? How difficult is it to determine whether a given correlation “is quantum”? This question is central to problems in quantum communication and computation. Previous work has shown that the general membership problem for quantum correlations is computationally undecidable. In the current work we show something stronger: there is a family of constant-sized correlations --- that is, correlations for which the number of measurements and number of measurement outcomes are fixed --- such that solving the quantum membership problem for this family is computationally impossible. Intuitively, our result means that the undecidability that arises in understanding Bell experiments is innate, and is not dependent on varying the number of measurements in the experiment. This places strong constraints on the types of descriptions that can be given for quantum correlation sets. Our proof is based on a combination of techniques from quantum self-testing and from undecidability results of the third author for linear system nonlocal games.

Vikesh Siddhu

Carnegie Mellon University

Leaking information to gain entanglement via log-singularities

Entanglement lies at the root of quantum theory. It is a remarkable resource that is generally believed to diminish when entangled systems interact with their environment. On the contrary, we find that engaging a system with its environment increases its ability to retain entanglement. The maximum rate of retaining entanglement is given by the quantum capacity. We counter-intuitively boost the quantum capacity of a channel by leaking almost all quantum information to the channel's environment. This result also has surprising implication for quantum key distribution: maximum rates for key distribution are boosted by allowing leakage of information to the eavesdropping environment. These boosts exploit two-letter level non-additivity in the channel's coherent information. The resulting non-additivity has a far larger magnitude and a qualitatively wider extent than previously known. This wide extent is proven using a new insight into the von-Neumann entropy which shows that logarithmic singularities in the entropy are a source of quantum capacity and non-additivity. These singularities can be used further to show a new type of non-additivity displayed by a zero quantum capacity qubit amplitude damping channel used in parallel with a simple qutrit channel.

Laura Mančinska, Jitendra Prakash and Christopher Schafhauser

University of Copenhagen | University of Copenhagen | University of Nebraska-Lincoln

Constant-sized robust self-tests for states and measurements of unbounded dimensions

We consider correlations, $p_{n,x}$, arising from measuring a maximally entangled state using n measurements with two outcomes each, constructed from n projections that add up to identity. We show that the correlations $p_{n,x}$ robustly self-test the underlying states and measurements. To achieve this, we lift the group-theoretic Gowers-Hatami based approach for proving robust self-tests to a more natural algebraic framework. A key step is to obtain an analogue of the Gowers-Hatami theorem allowing to perturb an ``approximate" representation of the relevant algebra to an exact one. For n=4, the correlations $p_{n,x}$ self-test the maximally entangled state of every odd dimension as well as 2-outcome projective measurements of arbitrarily high rank. The only other family of constant-size self-tests for strategies of unbounded dimension is due to Fu (QIP 2020) who presents such self-tests for an infinite family of maximally entangled states with *even* local dimension. Therefore, we are the first to exhibit a constant-size self-test for measurements of unbounded dimension as well as all maximally entangled states with odd local dimension. In addition, correlations $p_{4,x}$ represent the first self-tests for measurements of rank higher than one.

Marius Junge and Nicholas Laracuente

University of Illinois | University of Chicago

Multivariate Trace Inequalities, Recovery, and p-Fidelity Beyond Tracial Settings

Trace inequalities are powerful techniques in studying quantum entropy. The physics of quantum field theory and holography nonetheless motivate entropy inequalities in von Neumann algebras that lack a useful notion of a trace. Haagerup and Kosaki L_p spaces enable re-expressing trace inequalities in scenarios that start with a non-tracial von Neumann algebra, and we show how the generalized Araki-Lieb-Thirring and Golden-Thompson inequalities from (Sutter, Berta & Tomamichel 2017) port to general von Neumann algebras. Using an approximation method of Haagerup, we prove the tightened recovery map correction to data processing for relative entropy in (potentially non-tracial) von Neumann algebras. We also prove a p-fidelity version. Furthermore, we prove that non-decrease of relative entropy is equivalent to the existence of an L_1-isometry implementing the channel on both input states.

Edward Farhi, Jeffrey Goldstone, Sam Gutmann and Leo Zhou

Google | Massachusetts Institute of Technology | not applicable | Harvard University

The Quantum Approximate Optimization Algorithm and the Sherrington-Kirkpatrick Model at Infinite Size

The Quantum Approximate Optimization Algorithm (QAOA) is a general-purpose algorithm for combinatorial optimization problems whose performance can only improve with the number of layers p. While QAOA holds promise as an algorithm that can be run on near-term quantum computers, its computational power has not been fully explored. In this work, we study the QAOA applied to the Sherrington-Kirkpatrick (SK) model, which can be understood as energy minimization of n spins with all-to-all random signed couplings. There is a recent classical algorithm by Montanari that, assuming a widely believed conjecture, can be tailored to efficiently find an approximate solution for a typical instance of the SK model to within (1-epsilon) times the ground state energy. We can only hope to match its performance with the QAOA. Our main result is a novel technique that allows us to evaluate the typical-instance energy of the QAOA applied to the SK model. We produce a formula for the expected value of the energy, as a function of the 2p QAOA parameters, in the infinite size limit that can be evaluated on a computer with $O(16^p)$ complexity. We evaluate the formula up to p=12, and find that the QAOA at p=11 outperforms the standard semidefinite programming algorithm. Moreover, we show concentration: With probability tending to one as n goes to infinity, measurements of the QAOA will produce strings whose energies concentrate at our calculated value. As an algorithm running on a quantum computer, there is no need to search for optimal parameters on an instance-by-instance basis since we can determine them in advance. What we have here is a new framework for analyzing the QAOA, and our techniques can be of broad interest for evaluating its performance on more general problems where classical algorithms may fail.

William Kirby, Andrew Tranter and Peter Love

Tufts University | Cambridge Quantum Computing | Tufts University

Exploiting Contextuality in Variational Quantum Eigensolvers

We describe how contextuality may be used to advantage in variational quantum eigensolvers. Contextuality is a characteristic feature of quantum mechanics, and identifying contextuality in quantum algorithms provides a means for distinguishing them from their classical counterparts. We first describe how contextuality may be identified in variational quantum eigensolvers (VQEs), which are a leading algorithm for noisy intermediate-scale quantum computers. We then show how to construct a classical phase-space model for any noncontextual Hamiltonian, which provides a classical simulation algorithm for noncontextual VQE and allows us to prove that the noncontextual Hamiltonian problem is only NP-complete, rather than QMA-complete. Finally, we describe an approximation method called contextual subspace VQE that permits us to partition a general Hamiltonian into a noncontextual part and a contextual part, and estimate its ground state energy using a technique that combines classical simulation of the noncontextual part with quantum simulation of the contextual part. By using more quantum resources (in qubits and simulated terms of the Hamiltonian), we can increase the accuracy of the approximation. We tested contextual subspace VQE on electronic structure Hamiltonians, and found that to reach chemical accuracy in most cases it requires fewer qubits and simulated terms than standard VQE.

James Bartusek, Andrea Coladangelo, Dakshita Khurana and Fermi Ma

UC Berkeley | UC Berkeley | University of Illinois, Urbana Champaign | Princeton University and NTT Research

On the Round Complexity of Two-Party Quantum Computation

We investigate the round complexity of maliciously-secure two-party quantum computation (2PQC) with setup, and obtain the following results: - A three-message protocol (two-message if only one party receives output) in the common random string (CRS) model assuming classical two-message oblivious transfer (OT) with post-quantum malicious security. This round complexity is optimal for the sequential communication setting. Under the additional assumption of reusable malicious designated-verifier non-interactive zero-knowledge (MDV-NIZK) arguments for NP, our techniques give an MDV-NIZK for QMA. Each of the assumptions mentioned above is known from the quantum hardness of learning with errors (QLWE). - A protocol with two simultaneous rounds of communication, in a quantum preprocessing model, assuming sub-exponential QLWE. In fact, we construct a three-round protocol in the CRS model with only two rounds of online communication, which implies the above result. Along the way, we develop a new delayed simulation technique that we call “simulation via teleportation,” which may be useful in other settings. In addition, we perform a preliminary investigation into barriers and possible approaches for two round 2PQC in the CRS model, including an impossibility result for a natural class of simulators, and a proof-of-concept construction from a strong form of quantum virtual black-box (VBB) obfuscation. Prior to our work, maliciously-secure 2PQC required round complexity linear in the size of the quantum circuit.

Alexander Dalzell, Nicholas Hunter-Jones and Fernando Brandao

California Institute of Technology | Perimeter Institute for Theoretical Physics | AWS Center for Quantum Computing

Random quantum circuits anti-concentrate in log depth

We consider quantum circuits consisting of randomly chosen two-local gates and study the number of gates needed for the distribution over measurement outcomes for typical circuit instances to be anti-concentrated, roughly meaning that the probability mass is not too concentrated on a small number of measurement outcomes. Understanding the conditions for anti-concentration is important for determining which quantum circuits are difficult to simulate classically, as anti-concentration has been in some cases an ingredient of mathematical arguments that simulation is hard and in other cases a necessary condition for easy simulation. Our definition of anti-concentration is that the expected collision probability, that is, the probability that two independently drawn outcomes will agree, is only a constant factor larger than if the distribution were uniform. We show that when the 2-local gates are each drawn from the Haar measure (or any two-design), at least Ω(n log(n)) gates (and thus Ω(log(n)) circuit depth) are needed for this condition to be met on an n qudit circuit. In both the case where the gates are nearest-neighbor on a 1D ring and the case where gates are long-range, we show O(n log(n)) gates are also sufficient, and we precisely compute the optimal constant prefactor for the n log(n). The technique we employ relies upon a mapping from the expected collision probability to the partition function of an Ising-like classical statistical mechanical model, which we manage to bound using stochastic and combinatorial techniques.

Adam Bouland, Bill Fefferman, Zeph Landau and Yunchao Liu

UC Berkeley | University of Chicago | UC Berkeley | University of California, Berkeley

Noise and the frontier of quantum supremacy

Understanding the power of random quantum circuit sampling experiments has emerged as one of the most pressing topics in the near-term quantum era. In this work we make progress toward bridging the major remaining gaps between theory and experiment, incorporating the effects of experimental imperfections into the theoretical hardness arguments. We do this first by proving that computing the output probability of an m-gate random quantum circuit to within additive imprecision $2^{(-O(m^{1+\varepsilon}))}$ is #P-hard for any $\varepsilon>0$, an exponential improvement over the prior hardness results of Bouland et al. and Movassagh which were resistant to imprecision $2^{(-O(m^3))}$. This improvement very nearly reaches the threshold $(2^{(-O(m))}/$ poly(m)) sufficient to establish the hardness of sampling for constant-depth random quantum circuits. To prove this result we introduce new error reduction techniques for polynomial interpolation, as well as a new robust Berlekamp-Welch argument over the Reals which may be of independent interest. Second we show that these results are still true in the presence of a constant rate of noise, so long as the noise rate is below the error detection threshold. That is, even though random circuits with a constant noise rate converge rapidly to the maximally mixed state, the (exponentially) small deviations in their output probabilities away from uniformity remain difficult to compute. Interestingly, we then show that our two main results are in tension with one another, and the latter result implies the former result is essentially optimal with respect to additive imprecision error, even with substantial generalizations of our techniques.

Dante Bencivenga, Xining Chen and Peter Høyer

University of Calgary | University of Calgary | University of Calgary

Quantum sampling in Markov chains

We consider the problem of sampling from a target distribution in a Markov chain P. The random walk starts in a state drawn from the stationary distribution of P, walks according to P, and eventually stops once some predefined conditions are satisfied, generating a desired target distribution τ. We present a quantum algorithm that generates the target distribution τ using quadratically fewer steps than the optimal random walk. Our algorithm uses a generalization of controlled quantum walks. We introduce a quantum analogue of the exit frequencies of a random walk and use it to prove a relationship between sampling in random and quantum walks. Our framework yields simple and natural proofs. Applications of the framework include: A quantum rejection sampling algorithm achieving a quadratic speedup. A proof that controlled quantum walks can emulate generalized quantum interpolated walks. A proof that the extended hitting time is a measure of the complexity of a sampling problem. Efficient sampling is an important computational task used in e.g. simulations based on stochastic processes. This is the first quantum algorithm that achieves a quadratic speed-up for sampling from general probability distributions over states of a Markov chain starting from the stationary distribution.

Michael Beverland, Aleksander Kubica and Krysta M. Svore

Microsoft | AWS Center for Quantum Computing | Microsoft

The cost of universality: A comparative study of the overhead of state distillation and code switching with color codes

Estimating the reducing overhead of existing fault tolerance schemes is a crucial step toward realizing scalable quantum computers. Many of the most promising schemes are based upon two-dimensional (2D) topological codes such as the surface and color codes. In these schemes, universal computation is typically achieved using readily implementable Clifford operations along with a less convenient and more costly implementation of the T gate. In our work, we compare the cost of fault-tolerantly implementing the T-gate in 2D color codes using two leading approaches: state distillation and code switching to a 3D color code. We report that state distillation is more resource-efficient than code switching, in terms of both qubit overhead and space-time overhead. In particular, we find a T gate threshold via code switching of 0.07(1)% under circuit noise, almost an order of magnitude below that for distillation with 2D color codes. To arrive at this result, we provide and implement a simplified end-to-end recipe for code switching, detailing each step and providing important optimization considerations. We not only find numerical overhead estimates of this code switching protocol, but also lower bound various conceivable improvements. We also optimize the 2D color code for circuit noise yielding it's largest threshold to date 0.37(1)%, and adapt and optimize the restriction decoder and find a threshold of 0.80(5)% for the 3D color code with perfect measurements under Z noise. We foresee that this analysis will influence the choice of which FT schemes and which salable hardware designs should be pursued in future.

Alexis Schotte, Guanyu Zhu, Lander Burgelman and Frank Verstraete

Ghent University | IBM T.J. Watson Research Center | Ghent University | University of Ghent

Quantum error correction thresholds for the universal Fibonacci Turaev-Viro code

We consider a two-dimensional quantum memory of qubits on a torus encoding an extended Fibonacci string-net model, and construct error correction strategies when those qubits are subjected to depolarizing noise. In the case of a fixed-rate sampling noise model, we find an error correcting threshold of 4.75% with a clustering decoder. Using the concept of tube algebras, we construct a set of measurements and of quantum gates which map arbitrary qubit errors to the Turaev-Viro subspace. Tensor network techniques then allow to quantitatively study the action of Pauli noise on that subspace. We perform Monte-Carlo simulations of the Fibonacci code, and compare the performance of several decoders. To the best of our knowledge, this is the first time that a threshold has been calculated for a two-dimensional error correcting code in which universal quantum computation can be performed in its code space.

David Aasen, Daniel Bulmash, Abhinav Prem, Kevin Slagle and Dominic Williamson

Microsoft Station Q / KITP @ UCSB | University of Maryland, | Princeton University | California Institute of Technology | Stanford University

Topological Defect Networks for Fractons of all Types

Fracton phases exhibit striking behavior which appears to render them beyond the standard topological quantum field theory (TQFT) paradigm for classifying gapped quantum matter. Here, we explore fracton phases from the perspective of defect TQFTs and show that topological defect networks-networks of topological defects embedded in stratified 3+1D TQFTs-provide a unified framework for describing various types of gapped fracton phases. In this picture, the sub-dimensional excitations characteristic of fractonic matter are a consequence of mobility restrictions imposed by the defect network. We conjecture that all gapped phases, including fracton phases, admit a topological defect network description and support this claim by explicitly providing such a construction for many well-known fracton models, including the X-Cube and Haah's B code. To highlight the generality of our framework, we also provide a defect network construction of a novel fracton phase hosting non-Abelian fractons. As a byproduct of this construction, we obtain a generalized membrane-net description for fractonic ground states as well as an argument that our conjecture implies no topological fracton phases exist in 2+1D gapped systems. Our work also sheds light on new techniques for constructing higher order gapped boundaries of 3+1D TQFTs.

Cambyse Rouzé, Ivan Bardet, Ángela Capel and Daniel Stilck França

Technische Universität München | INRIA Paris | Technische Universität München | University of Copenhagen

On the entropic convergence of quantum Gibbs samplers

Given a uniform, frustration-free family of local Lindbladians defined on a quantum lattice spin system in any spatial dimension, we prove a strong exponential convergence in relative entropy of the system to equilibrium under a condition of spatial mixing of the stationary Gibbs states and the rapid decay of the relative entropy on finite-size blocks. Our result leads to the first examples of the positivity of the modified logarithmic Sobolev inequality for quantum lattice spin systems independently of the system size. Moreover, we show that our notion of spatial mixing is a consequence of the recent quantum generalization of Dobrushin and Shlosman's complete analyticity of the free-energy at equilibrium. The latter typically holds above a critical temperature T_c. Our results have wide applications in quantum information processing. As an illustration, we discuss three of them: first, using techniques of quantum optimal transport, we show that a quantum annealer subject to a finite range classical noise will output an energy close to that of the fixed point after constant annealing time. Second, we prove a finite blocklength refinement of the quantum Stein lemma for the task of asymmetric discrimination of two Gibbs states of commuting Hamiltonians satisfying our conditions. In the same setting, our results imply the existence of a local quantum circuit of logarithmic depth to prepare Gibbs states of a class of commuting Hamiltonians.

Tomotaka Kuwahara, Alvaro Alhambra and Anurag Anshu

RIKEN Center for Advanced Intelligence Project | Max Planck Institute for Quantum Optics | University of California, Berkeley

Improved thermal area law and quasi-linear time algorithm for quantum Gibbs states

One of the most fundamental problems in quantum many-body physics is the characterization of correlations among thermal states. Of particular relevance is the thermal area law, which justifies the tensor network approximations to thermal states with a bond dimension growing polynomially with the system size. In the regime of sufficiently low temperatures, which is particularly important for practical applications, the existing techniques do not yield optimal bounds. Here, we propose a new thermal area law that holds for generic many-body systems on lattices. We improve the temperature dependence from the original $O(\beta)$ to $\tilde{O}(\beta^{2/3})$, thereby suggesting diffusive propagation of entanglement by imaginary time evolution. This qualitatively differs from the real-time evolution which usually induces linear growth of entanglement. We also prove analogous bounds for the Rényi entanglement of purification and the entanglement of formation. Our analysis is based on a polynomial approximation to the exponential function which provides a relationship between the imaginary-time evolution and random walks. Moreover, for one-dimensional (1D) systems with n spins, we prove that the Gibbs state is well-approximated by a matrix product operator with a sublinear bond dimension of $\exp(\sqrt{\tilde{O}(\beta \log(n))})$. This allows us to rigorously establish, for the first time, a quasi-linear time classical algorithm for constructing an MPS representation of 1D quantum Gibbs states at arbitrary temperatures of $\beta=o(\log(n))$. Our new technical ingredient is a block decomposition of the Gibbs state, that bears resemblance to the decomposition of real-time evolution given by Haah et al., FOCS'18.

Daniel Ranard, Freek Witteveen and Michael Walter

Stanford University | University of Amsterdam | University of Amsterdam

Classifying unitary dynamics with approximate light cones in one dimension

Unitary dynamics with a strict causal cone (or "light cone") have been studied extensively, under the name of locality preserving unitaries (LPUs) or quantum cellular automata. In particular, LPUs in one dimension have been completely classified by an index theory. Physical systems often exhibit only approximate causal cones; Hamiltonian evolutions on the lattice satisfy Lieb-Robinson bounds rather than strict locality. This motivates us to study approximately locality preserving unitaries (ALPUs). We show that the index theory is robust and completely extends to one-dimensional ALPUs. As a consequence, we achieve a converse to the Lieb-Robinson bounds: any ALPU of index zero can be exactly generated by some time-dependent, quasi-local Hamiltonian in constant time. For the special case of finite chains with open boundaries, any unitary satisfying the Lieb-Robinson bound may be generated by such a Hamiltonian. We also discuss some results on the stability of operator algebras which may be of independent interest.

Sam Roberts and Dominic Williamson

PsiQuantum | Stanford University

3-Fermion topological quantum computation

We present a scheme for universal topological quantum computation based on Clifford complete braiding and fusion of symmetry defects in the 3-Fermion anyon theory, supplemented with magic state injection. We formulate a fault-tolerant measurement-based realisation of this computational scheme on the lattice using ground states of the Walker--Wang model for the 3-Fermion anyon theory with symmetry defects. The Walker--Wang measurement-based topological quantum computation paradigm that we introduce provides a general construction of computational resource states with thermally stable symmetry-protected topological order. We also demonstrate how symmetry defects of the 3-Fermion anyon theory can be realized in a 2D subsystem code due to Bombin -- pointing to an alternative implementation of our 3-Fermion defect computation scheme via code deformations.

Bowen Shi, Kohtaro Kato and Isaac Kim

University of California San Diego | Osaka University | The University of Sydney

Entanglement bootstrap program

We introduce the entanglement bootstrap program, a powerful new approach to study topologically ordered quantum many-body systems. In this program, we posit that local reduced density matrices of the underlying system obey a set of simple constraints that are motivated by the entanglement entropy calculations in the literature. We show that these constraints lead to a highly nontrivial set of identities that can be interpreted as the basic axioms of the emergent theory describing the topological charges in such systems. The surprising nature of our work lies in the fact that these basic axioms of the emergent theory, which are typically assumed in the literature, can actually be derived from a simple entanglement property of the ground state. Furthermore, we apply this line of reasoning to systems with gapped domain walls and derive a hitherto unknown set of topological charges and identities that those topological charges need to satisfy. Thus, our work establishes a deep connection between entanglement and the exotic behaviors of topological charges in topologically ordered systems.

Kianna Wan, Soonwon Choi, Isaac Kim, Noah Shutty and Patrick Hayden

Stanford University | University of California, Berkeley | The University of Sydney | Stanford University | Stanford University

Fault-tolerant qubit from a constant number of components

With gate error rates in multiple technologies now below the threshold required for fault-tolerant quantum computation, the major remaining obstacle to useful quantum computation is scaling, a challenge greatly amplified by the huge overhead imposed by quantum error correction itself. We propose a fault-tolerant quantum computing scheme that can nonetheless be assembled from a small number of experimental components, potentially dramatically reducing the engineering challenges associated with building a large-scale fault-tolerant quantum computer. Our scheme has a threshold of 0.39% for depolarising noise, assuming that memory errors are negligible. In the presence of memory errors, the logical error rate decays exponentially with $\sqrt{\mathrm{T}/\tau}$, where T is the memory coherence time and τ is the timescale for elementary gates. Our approach is based on a novel procedure for fault-tolerantly preparing three-dimensional cluster states using a single actively controlled qubit and a pair of delay lines. Although a circuit-level error may propagate to a high-weight error, the effect of this error on the prepared state is always equivalent to that of a constant-weight error. We describe how the requisite gates can be implemented using existing technologies in quantum photonic and phononic systems. With continued improvements in only a few components, we expect these systems to be promising candidates for demonstrating fault-tolerant quantum computation with a comparatively modest experimental effort.

William Kretschmer

University of Texas at Austin

The Quantum Supremacy Tsirelson Inequality

A leading proposal for verifying near-term quantum supremacy experiments on noisy random quantum circuits is linear cross-entropy benchmarking. For a quantum circuit C on n qubits and a sample $z \in \{0,1\}^n$, the benchmark involves computing $|\langle z|C|0^n\rangle|^2$, i.e. the probability of measuring z from the output distribution of C on the all zeros input. Under a strong conjecture about the classical hardness of estimating output probabilities of quantum circuits, no polynomial-time classical algorithm given C can output a string z such that $|\langle z|C|0^n\rangle|^2$ is substantially larger than $1/2^n$ (Aaronson and Gunn, 2019). On the other hand, for a random quantum circuit C, sampling z from the output distribution of C achieves $|\langle z|C|0^n\rangle|^2 \approx 2/2^n$ on average (Arute et al., 2019). In analogy with the Tsirelson inequality from quantum nonlocal correlations, we ask: can a polynomial-time quantum algorithm do substantially better than $2/2^n$? We study this question in the query (or black box) model, where the quantum algorithm is given oracle access to C. We show that, for any $\varepsilon \geq 1/\mathrm{poly}(n)$, outputting a sample z such that $|\langle z|C|0^n\rangle|^2 \geq (2+\varepsilon)/2^n$ on average requires at least $\Omega(2^{n/4}/\mathrm{poly}(n))$ queries to C, but not more than $O(2^{n/3})$ queries to C, if C is either a Haar-random n-qubit unitary, or a canonical state preparation oracle for a Haar-random n-qubit state. We also show that when C samples from the Fourier distribution of a random Boolean function, the naive algorithm that samples from C is the optimal 1-query algorithm for maximizing $|\langle z|C|0^n\rangle|^2$ on average.

Justin Holmgren and Ran Raz

NTT Research | Princeton University

A Parallel Repetition Theorem for the GHZ Game

We prove that parallel repetition of the (3-player) GHZ game reduces the value of the game polynomially fast to 0. That is, the value of the GHZ game repeated in parallel t times is at most $t^{-\Omega(1)}$. Previously, only a bound of $\approx 1/\alpha(t)$, where α is the inverse Ackermann function, was known (Verbitsky '96). The GHZ game was recently identified by Dinur, Harsha, Venkat and Yuen as a multi-player game where all existing techniques for proving strong bounds on the value of the parallel repetition of the game fail. Indeed, to prove our result we use a completely new proof technique. Dinur, Harsha, Venkat and Yuen speculated that progress on bounding the value of the parallel repetition of the GHZ game may lead to further progress on the general question of parallel repetition of multi-player games. They suggested that the strong correlations present in the GHZ question distribution represent the ``hardest instance'' of the multi-player parallel repetition problem. Another motivation for studying the parallel repetition of the GHZ game comes from the field of quantum information. The GHZ game, first introduced by Greenberger, Horne and Zeilinger, is a central game in the study of quantum entanglement and has been studied in numerous works. For example, it is used for testing quantum entanglement and for device-independent quantum cryptography. In such applications a game is typically repeated to reduce the probability of error, and hence bounds on the value of the parallel repetition of the game may be useful.

Satvik Singh and Ion Nechita

Department of Physical Sciences, Indian Institute of Science Education and Research, Mohali | Université de Toulouse

The PPT-squared conjecture holds for all Choi type maps

In the rapidly developing field of Quantum technologies, the task of entanglement distribution between two parties occupies a central stage in many important protocols. However, as the distance between the two parties increases, the error probability in any transmission channel gets larger, resulting in degradation of the quality of the distributed entanglement. To overcome this problem, quantum repeater devices are used. The basic idea of such a device is to split up the long transmission channel into shorter manageable segments, each of which can be provided with high fidelity entangled states. Then, the well-known entanglement swapping technique can be used to transfer the entanglement from the intermediate segments to the ends of the long channel. A key conjecture in this regard was proposed by M. Christandl, which states that all PPT entangled states are useless from the perspective of repeater devices, since the swapping of entanglement in such states inevitably leads to a separable state. The conjecture admits an equivalent formulation in terms of linear maps, where it amounts to saying that the composition of any two PPT maps (these are the maps which are both completely positive and completely copositive) is entanglement-breaking. In the present work, we prove that this conjecture holds for all linear maps which are covariant under the diagonal unitary group's action. Many salient examples like the Choi-type maps, Schur multipliers, Classical maps, etc. lie in this class. Our proof relies on a generalization of the matrix-theoretic notion of factor width for pairwise completely positive matrices, as well as on our previous characterization of the aforementioned class of maps. Hence, in a nutshell, our research proves the unsuitability of a large class of states from the perspective of repeater protocols and thus significantly contributes to the solution of a long-standing open problem in quantum information theory.

Pablo Bonilla Ataides, David Tuckett, Stephen Bartlett, Steven Flammia and Benjamin Brown

University of Sydney | University of Sydney | The University of Sydney | AWS | University of Sydney

The XZZX surface code

We show that a variant of the surface code---the XZZX code---offers remarkable performance for fault-tolerant quantum computation. The error threshold of this code matches what can be achieved with random codes (hashing) for *every* single-qubit Pauli noise channel; it is the first explicit code shown to have this universal property. We present numerical evidence that the threshold even exceeds this hashing bound for an experimentally relevant range of noise parameters. Focusing on the common situation where qubit dephasing is the dominant noise, we show that this code has a practical, high-performance decoder and surpasses all previously known thresholds in the realistic setting where syndrome measurements are unreliable. We go on to demonstrate the favorable sub-threshold resource scaling that can be obtained by specializing a code to exploit structure in the noise. We show that it is possible to maintain all of these advantages when we perform fault-tolerant quantum computation. We finally suggest some small-scale experiments that could exploit noise bias to reduce qubit overhead in two-dimensional architectures. The complete version of this paper can be found at https://arxiv.org/abs/2009.07851.

Daniel Ranard and Xiao-Liang Qi

Stanford University | Stanford University

Emergent classicality in multipartite channels

In a quantum measurement process, classical information about the measured system spreads throughout the environment. Meanwhile, quantum information about the system becomes inaccessible to local observers. Here we prove a result about quantum channels indicating that an aspect of this phenomenon is completely general. We show that for any evolution of the system and environment, for everywhere in the environment excluding an O(1)-sized region we call the "quantum Markov blanket," any locally accessible information about the system must be approximately classical, i.e. obtainable from some fixed measurement. The result strengthens the earlier result of arXiv:1310.8640 in which the excluded region was allowed to grow with total environment size. It may also be seen as a new consequence of the principles of no-cloning or monogamy of entanglement. Our proof offers a constructive optimization procedure for determining the "quantum Markov blanket" region, as well as the effective measurement induced by the evolution. Alternatively, under channel-state duality, our result characterizes the marginals of multipartite states.

Alexander Müller-Hermes and Matthias Christandl

Claude Bernard University of Lyon 1 | University of Copenhagen

Fault-tolerant coding for quantum communication

Designing encoding and decoding circuits to reliably send messages over many uses of a noisy channel is a central problem in communication theory. When studying the optimal transmission rates achievable with asymptotically vanishing error it is usually assumed that these circuits can be implemented using noise-free gates. While this assumption is satisfied for classical machines in many scenarios, it is not expected to be satisfied in the near term future for quantum machines where decoherence leads to faults in the quantum gates. As a result, fundamental questions regarding the practical relevance of quantum channel coding remain open. By combining techniques from fault-tolerant quantum computation with techniques from quantum communication, we initiate the study of these questions. We introduce fault-tolerant versions of quantum capacities quantifying the optimal communication rates achievable with asymptotically vanishing total error when the encoding and decoding circuits are affected by gate errors with small probability. Our main results are threshold theorems for the classical and quantum capacity: For every quantum channel T and every $\varepsilon>0$ there exists a threshold $p(\varepsilon,T)$ for the gate error probability below which rates larger than $C-\varepsilon$ are fault-tolerantly achievable with vanishing overall communication error, where C denotes the usual capacity. Our results are not only relevant in communication over large distances, but also on-chip, where distant parts of a quantum computer might need to communicate under higher levels of noise than affecting the local gates.

Minh Tran, Abhinav Deshpande, Andrew Guo, Andrew Lucas and Alexey Gorshkov

QuICS/University of Maryland | QuICS/JQI/University of Maryland | QuICS/JQI/University of Maryland | University of Colorado Boulder | Joint Quantum Institute (JQI)

Optimal State Transfer and Entanglement Generation in Power-law Interacting Systems

We present an optimal protocol for encoding an unknown qubit state into a multiqubit Greenberger-Horne-Zeilinger-like state and, consequently, transferring quantum information in large systems exhibiting power-law ($1/r^\alpha$) interactions. For all power-law exponents α between d and 2d+1, where d is the dimension of the system, the protocol yields a polynomial speedup for $\alpha>2d$ and a superpolynomial speedup for $\alpha\leq 2d$, compared to the state of the art. For all $\alpha>d$, the protocol saturates the Lieb-Robinson bounds (up to subpolynomial corrections), thereby establishing the optimality of the protocol and the tightness of the bounds in this regime. The protocol has a wide range of applications, including in quantum sensing, quantum computing, and preparation of topologically ordered states.

Gorjan Alagic, Prabhanjan Ananth, Zvika Brakerski, Yfke Dulek, Rolando La Placa and Christian Schaffner

QuICS | UC Santa Barbara | Weizmann Institute of Science | QuSoft, University of Amsterdam | MIT | QuSoft, University of Amsterdam

Secure Software Leasing and Implications to Quantum Copy-Protection and Obfuscation

In quantum copy-protection, an adversary who is given a quantum state computing a function f cannot produce two (possibly entangled) quantum states that each individually compute f. No constructions for copy-protection are known in the plain model. We consider a weaker notion, secure software leasing (SSL), where it is only impossible to produce two copies that can both compute f using the honest evaluation algorithm. We show the following: (1) SSL is possible for a subclass of evasive functions, assuming the existence of post-quantum indistinguishability obfuscators and hardness of LWE; (2) SSL is impossible in general, assuming hardness of LWE. The second statement has important implications for existing quantum-cryptographic notions: in particular, it implies the impossibility of quantum copy-protection for arbitrary unlearnable functions, and impossibility of quantum virtual-black-box obfuscation of classical circuits.

James Bartusek, Andrea Coladangelo, Dakshita Khurana and Fermi Ma // Alex Bredariol Grilo, Huijia Lin, Fang Song and Vinod Vaikuntanathan

LIP6, CNRS/Sorbonne Université | University of Washington | Portland State University | MIT // UC Berkeley | UC Berkeley | University of Illinois, Urbana Champaign | Princeton University and NTT Research

Secure Computation is in MiniQCrypt

MiniQCrypt is a world where quantum-secure one-way functions exist, and quantum communication is possible. We construct an oblivious transfer (OT) protocol in MiniQCrypt that achieves simulation-security against malicious quantum polynomial-time adversaries, building on the foundational work of Bennett, Brassard, Crépeau and Skubiszewska (CRYPTO 1991). Combining the OT protocol with prior works, we obtain secure two-party and multi-party computation protocols also in MiniQCrypt. This is in contrast to the classical world, where it is widely believed that one-way functions alone do not give us OT.

Greg Kuperberg

Professor of Mathematics and Computer Science, UC Davis, USA

The Hidden Subgroup Problem for Infinite Groups

We consider the hidden subgroup problem (HSP) for infinite groups, beyond the celebrated original cases established by Shor and Kitaev. We prove that HSP is NP-hard in the rational numbers $\mathbb{Q}$ under addition, as well as for normal subgroups of a non-abelian free group F_k. We can show that HSP in the lattice $\mathbb{Z}^k$ is uSVP-hard with unary encoding of vectors. On the other hand, HSP in $\mathbb{Z}^k$ with standard binary encoding of vectors can be solved in BQP, uniformly in the dimension, generalizing the Shor-Kitaev algorithm to infinite-index hidden subgroups. HSP in any fixed, finitely generated, virtually abelian subgroup can also be solved in subexponential time using established quantum algorithms for the abelian hidden shift problem.

Chenyi Zhang, Jiaqi Leng and Tongyang Li

Tsinghua University | University of Maryland | Massachusetts Institute of Technology

Quantum algorithms for escaping from saddle points

We initiate the study of quantum algorithms for escaping from saddle points with provable guarantee. Given a function $f: \mathbb{R}^n \to \mathbb{R}$, our quantum algorithm outputs an ε-approximate local minimum using $\tilde{O}(\log^2 n/\varepsilon^{1.75})$ queries to the quantum evaluation oracle (i.e., the zeroth-order oracle). Compared to the classical state-of-the-art algorithm by Jin et al. with $\tilde{O}(\log^6 n/\varepsilon^{1.75})$ queries to the gradient oracle (i.e., the first-order oracle), our quantum algorithm is polynomially better in terms of n and matches its complexity in terms of $1/\varepsilon$. Our quantum algorithm is built upon two techniques: First, we replace the classical perturbations in gradient descent methods by simulating quantum wave equations, which constitutes the polynomial speedup in n for escaping from saddle points. Second, we show how to use a quantum gradient computation algorithm due to Jordan to replace the classical gradient queries in nonconvex optimization by quantum evaluation queries with the same complexity, extending the same result from convex optimization due to van Apeldoorn et al. and Chakrabarti et al. Finally, we also perform numerical experiments that support our quantum speedup.

Bill Fefferman and Zachary Remscrim // Uma Girish, Ran Raz and Wei Zhan

University of Chicago | University of Chicago // Princeton University | Princeton University | Princeton University

Eliminating Intermediate Measurements in Space-Bounded Quantum Computation // Quantum Logspace Algorithm for Powering Matrices with Bounded Norm

A foundational result in the theory of quantum computation, known as the ``principle of safe storage,'' shows that it is always possible to take a quantum circuit and produce an equivalent circuit that makes all measurements at the end of the computation. While this procedure is time efficient, meaning that it does not introduce a large overhead in the number of gates, it uses extra ancillary qubits, and so is not generally space efficient. It is quite natural to ask whether it is possible to eliminate intermediate measurements without increasing the number of ancillary qubits. We give an affirmative answer to this question by exhibiting a procedure to eliminate all intermediate measurements that is simultaneously space efficient and time efficient. In particular, this shows that the definition of a space-bounded quantum complexity class is robust to allowing or forbidding intermediate measurements. A key component of our approach, which may be of independent interest, involves showing that the well-conditioned versions of many standard linear-algebraic problems may be solved by a quantum computer in less space than seems possible by a classical computer. We give a quantum logspace algorithm for powering contraction matrices, that is, matrices with spectral norm at most 1. The algorithm gets as an input an arbitrary $n \times n$ contraction matrix A, and a parameter $T \leq poly(n)$ and outputs the entries of A^T, up to (arbitrary) polynomially small additive error. The algorithm applies only unitary operators, without intermediate measurements. We show various implications and applications of this result: First, we use this algorithm to show that the class of quantum logspace algorithms with only quantum memory and with intermediate measurements is equivalent to the class of quantum logspace algorithms with only quantum memory without intermediate measurements. This shows that the deferred-measurement principle, a fundamental principle of quantum computing, applies also for quantum logspace algorithms (without classical memory). More generally, we give a quantum algorithm with space $O(S + \log T)$ that takes as an input the description of a quantum algorithm with quantum space S and time T, with intermediate measurements (without classical memory), and simulates it unitarily with polynomially small error, without intermediate measurements. Since unitary transformations are reversible (while measurements are irreversible) an interesting aspect of this result is that it shows that any quantum logspace algorithm (without classical memory) can be simulated by a reversible quantum logspace algorithm. This proves a quantum analogue of the result of Lange, McKenzie and Tapp that deterministic logspace is equal to reversible logspace. Finally, we use our results to show non-trivial classical simulations of quantum logspace learning algorithms.

Aleksander Kubica and Rafał Demkowicz-Dobrzański

AWS Center for Quantum Computing | Faculty of Physics, University of Warsaw

Using Metrological Bounds in Quantum Error Correction

We present a simple proof of the approximate Eastin-Knill theorem, which connects the quality of a quantum error-correcting code (QECC) with its ability to achieve a universal set of transversal logical gates. Our derivation employs powerful bounds on the quantum Fisher information in generic quantum metrological protocols to characterize the QECC performance measured in terms of the worst-case entanglement fidelity. The theorem is applicable to a large class of decoherence models, including erasure and depolarizing noise. Our approach is unorthodox, as instead of following the established path of utilizing QECCs to mitigate noise in quantum metrological protocols, we apply methods of quantum metrology to explore the limitations of QECCs.

Sara Bartolucci, Patrick Birchall, Hector Bombin, Hugo Cable, Chris Dawson, Mercedes Gimeno-Segovia, Eric Johnston, Konrad Kieling, Naomi Nickerson, Mihir Pant, Fernando Pastawski, Terry Rudolph and Chris Sparrow

PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp. | PsiQuantum Corp.

Fusion-based quantum computation

We introduce fusion based quantum computing (FBQC) - a model of universal quantum computation in which entangling measurements, called fusions, are performed on the qubits of small constant sized entangled resource states. We introduce a stabilizer formalism for analyzing fault tolerance and computation in these schemes. This framework naturally captures the error structure that arises in certain physical systems, such as linear optics. FBQC can offer significant architectural simplifications, enabling hardware made up of many identical modules, and reducing classical processing requirements. We present two pedagogical examples of fault-tolerant schemes constructed in this framework and numerically evaluate their threshold under a general error model with measurement erasure and error, and a linear-optical error model with fusion failure and photon loss. In these schemes, fusion failures, as well as errors, are directly dealt with by the quantum error correction protocol. We find that tailoring the fault-tolerance framework to the physical system allows the scheme to have a higher threshold than schemes reported in literature. We present a ballistic scheme which can tolerate a 10.3% probability of suffering photon loss in each fusion.

Oscar Higgott and Nikolas Breuckmann

University College London | University College London

Subsystem codes with high thresholds by gauge fixing and reduced qubit overhead

We introduce a technique that uses gauge fixing to significantly improve the quantum error correcting performance of subsystem codes. By changing the order in which check operators are measured, valuable additional information can be gained, and we introduce a new method for decoding which uses this information to improve performance. Applied to the subsystem toric code with three-qubit check operators, we increase the threshold under circuit-level depolarising noise from 0.67% to 0.81%. The threshold increases further under a circuit-level noise model with small finite bias, up to 2.22% for infinite bias. Furthermore, we construct families of finite-rate subsystem LDPC codes with three-qubit check operators and optimal-depth parity-check measurement schedules. To the best of our knowledge, these finite-rate subsystem codes outperform all known codes at circuit-level depolarising error rates as high as 0.2%, where they have a qubit overhead that is 4.3× lower than the most efficient version of the surface code and 5.1× lower than the subsystem toric code. Their threshold and pseudo-threshold exceeds 0.42% for circuit-level depolarising noise, increasing to 2.4% under infinite bias using gauge fixing.

Burak Sahinoglu and Rolando Somma

Los Alamos National Laboratory | Los Alamos National Laboratory

Hamiltonian simulation in the low energy subspace

We study the problem of simulating the dynamics of spin systems when the initial state is supported on a subspace of low energy of a Hamiltonian H. This is a central problem in physics with vast applications in many-body systems and beyond, where the interesting physics takes place in the low-energy sector. We analyze error bounds induced by product formulas that approximate the evolution operator and show that these bounds depend on an effective low-energy norm of H. We find improvements over the best previous complexities of product formulas that apply to the general case, and these improvements are more significant for long evolution times that scale with the system size and/or small approximation errors. To obtain these improvements, we prove novel exponentially-decaying upper bounds on the leakage to high-energy subspaces due to the product formula. Our results provide a path to a systematic study of Hamiltonian simulation at low energies, which will be required to push quantum simulation closer to reality.

Hakop Pashayan, Oliver Reardon-Smith, Kamil Korzekwa and Stephen Bartlett

Institute for Quantum Computing, University of Waterloo ; Perimeter Institute | Jagiellonian University | Jagiellonian University | The University of Sydney

Fast estimation of outcome probabilities for quantum circuits

We present two classical algorithms for the simulation of universal quantum circuits on n qubits constructed from c instances of Clifford gates and t arbitrary-angle Z-rotation gates such as T gates. Our algorithms complement each other by performing best in different parameter regimes. The Estimate algorithm produces an additive precision estimate of the Born rule probability of a chosen measurement outcome with the only source of run-time inefficiency being a linear dependence on the stabilizer extent (which scales like $\approx 1.17^t$ for T gates). Our algorithm is state-of-the-art for this task: as an example, in approximately 25 hours (on a standard desktop computer), we estimated the Born rule probability to within an additive error of 0.03, for a 50 qubit, 60 non-Clifford gate quantum circuit with more than 2000 Clifford gates. The Compute algorithm calculates the probability of a chosen measurement outcome to machine precision with run-time $O(2^{(t-r)}(t-r)t)$ where r is an efficiently computable, circuit-specific quantity. With high probability, r is very close to $\min\{t, n-w\}$ for random circuits with many Clifford gates, where w is the number of measured qubits. Compute can be effective in surprisingly challenging parameter regimes, e.g., we can randomly sample Clifford+T circuits with $n = 55$, $w = 5$, $c = 10^5$ and $t = 80$, T-gates, and then compute the Born rule probability with a run-time consistently less than 104 seconds using a single core of a standard desktop computer. We provide a C+Python implementation of our algorithms.

Matthew Coudron and Nolan Coble

University of Maryland Computer Science/NIST | University of Maryland Mathematics

Quasi-polynomial Time Approximation of Output Probabilities of Constant-depth, Geometrically-local Quantum Circuits

We present a classical algorithm that, for any geometrically-local, constant-depth quantum circuit C, and any bit string $x \in \{0,1\}^n$, can compute the quantity $< 0|^{\otimes n} C|x >^2$ to within any inverse-polynomial additive error in quasi-polynomial time. It is known that it is $\# P$-hard to compute this same quantity to within 2^{-n^2} additive error. The previous best known algorithm for this problem used $O(2^{n^{\frac{1}{3}}} poly(1/\varepsilon))$ time to compute probabilities to within additive error ϵ \cite{BGM19}. Notably, the [BGM19] paper included an elegant polynomial time algorithm for the same estimation task with 2D circuits, which makes a novel use of 1D Matrix Product States carefully tailored to the 2D geometry of the circuit in question. Surprisingly, it is not clear that it is possible to extend this use of MPS to address the case of 3D circuits in polynomial time. This raises a natural question as to whether the computational complexity of the 3D problem might be drastically higher than that of the 2D problem. In this work we address this question by exhibiting a quasi-polynomial time algorithm for the 3D case. In order to surpass the technical barriers encountered by previously known techniques we are forced to pursue a novel approach: Constructing a recursive sub-division of the given 3D circuit using carefully designed block-encodings.

Runyao Duan, Baidu

Director of Institute for Quantum Computing, Baidu Research

Quantum Computing at Baidu

Quantum computing is believed to be the heart of the next-generation computing technology. Our mission at Baidu Research is to be a world-class Quantum Artificial Intelligence (AI) research strength, and to continuously integrate relevant quantum technologies into Baidu's core business. We have general interests in Quantum Information Science with three research priority topics: Quantum AI, Quantum Algorithm, and Quantum Architecture - altogether form the QAAA research plan. As a result of this plan, we develop the Baidu Quantum Platform (BQP), which currently contains three essential products: 1) Quantum Leaf, a cloud-native quantum computing platform; 2) Paddle Quantum, a quantum machine learning toolkit developed based on Baidu's deep learning platform PaddlePaddle; 3) Quanlse, a cloud-based platform for quantum control. With the help of BQP and its further development, we strive to explore more possibilities of quantum technology, to build a sustainable quantum ecosystem, and finally to achieve our vision that "Everyone Can Quantum".

Bruno Taketani, IQM

IQM

Co-designing quantum computers at IQM

Quantum computers are set to revolutionise many fields. However, a full-fledged quantum computer is still a long term goal. IQM is invested in reaching quantum advantage faster via application specific, co-design quantum computers. In this talk, we will present recent updates on the technological development at IQM Quantum Computers in developing scalable superconducting quantum processors. We will discuss our implementation of digital-analog quantum algorithms as a promising route to improve circuit depth. On the broader scope, we will also present our latest results in qubit and quantum processor development.

Dr. Johannes Klepsch, BMW

Product Owner Quantum Computing, BMW AG

Application of quantum computing in the increasingly complex automotive value chain

The complexity is increasing rapidly in many areas of the automotive industry. The design of an automobile involves many different engineering disciplines and various tradeoffs. The vehicle software and hardware system comprise millions of lines of code. Complex processes, such as manufacturing, logistics, distribution and sales, need to be handled. In all these domains, myriads of optimization problems arise. Quantum computing-based optimization approaches promise to overcome some of the inherent scalability limitations of classical systems. This presentation investigates quantum computing applications across the automotive value chain and evaluates their suitability for quantum-based methods.

Havi Carel

Professor of Philosophy, University of Bristol

How to promote social justice in academia

This talk will provide an introduction and training to participants on equality, diversity and inclusion in academia, including microaggressions and implicit bias. We will learn about the processes and factors that are barriers to equality, diversity and inclusion and discuss some ways of mitigating these in the context of academia.

Anurag Anshu and Chinmay Nirkhe

University of California, Berkeley | University of California, Berkeley

Circuit lower bounds for low-energy states of code Hamiltonians

The No Low-energy Trivial States (NLTS) conjecture of Freedman and Hastings (Quantum Information and Computation, 2014) -- which posits the existence of a local Hamiltonian with a super-constant circuit lower bound on the complexity of all low-energy states -- identifies a fundamental obstacle to the resolution of the quantum PCP conjecture. In this work, we provide new techniques based on entropic and local indistinguishability arguments that prove circuit lower bounds for all the low-energy states of local Hamiltonians arising from quantum error-correcting codes. For local Hamiltonians arising from nearly linear-rate and polynomial-distance LDPC stabilizer codes, we prove super-constant circuit lower bounds for the complexity of all states of energy $o(n)$ (which can be viewed as an almost linear NLTS theorem). Such codes are known to exist and are not necessarily locally-testable, a property previously suspected to be essential for the NLTS conjecture. Curiously, such codes can also be constructed on a two-dimensional lattice, showing that low-depth states cannot accurately approximate the ground-energy in physically relevant systems.

Srinivasan Arunachalam, Alex B. Grilo, Tom Gur, Igor C. Oliveira and Aarthi Sundaram

IBM T. J. Watson Research Center | Sorbonne Universite, CNRS, LIP6 | University of Warwick | University of Warwick | Microsoft Quantum

Quantum Learning Algorithms Imply Circuit Lower Bounds

We establish the first general connection between the design of quantum algorithms and circuit lower bounds. Specifically, let C be a class of polynomial-size concepts, and suppose that C can be learned in the PAC model under the uniform distribution with membership queries, and with error $\frac{1}{2} - c$ by a time T quantum algorithm. We prove that if $(c^2.T) \ll \frac{2^n}{n}$, then BQE is not contained in C, where BQE = BQTIME$[2^{O(n)}]$ is an exponential-time analogue of BQP. This result is optimal in both c and T since it is not hard to learn any class C of functions in (classical) time $T = 2^n$ (with no error) or in quantum time $T = \mathrm{ploy}(n)$ with error at most $\frac{1}{2} - \omega(2^{\frac{-n}{2}})$ via Fourier sampling. In other words, even a marginal improvement on these generic learning algorithms would lead to major consequences in complexity theory. Our proof builds on several works in learning theory, pseudorandomness, and computational complexity, and crucially, on a connection between non-trivial classical learning algorithms and circuit lower bounds established by Oliveira and Santhanam (CCC 2017). Extending their approach to quantum learning algorithms turns out to create significant challenges. To achieve that, we show among other results how pseudorandom generators imply learning-to-lower-bound connections in a generic fashion, construct the first conditional pseudorandom generator secure against uniform quantum computations, and extend the local list-decoding algorithm of Impagliazzo, Jaiswal, Kabanets and Wigderson (SICOMP 2010) to quantum circuits via a delicate analysis. We believe that these contributions are of independent interest and might find other applications.

Dorit Aharonov, Jordan Cotler and Xiao-Liang Qi

The Hebrew University of Jerusalem | Harvard University | Stanford University

Quantum Algorithmic Measurement

Can quantum computational tools enhance the precision and efficiency of physical experiments? Promising examples are known, but a systematic treatment and comprehensive framework are missing. We introduce Quantum Algorithmic Measurements (QUALMs) to enable the study of quantum measurements and experiments from the perspective of computational complexity and communication complexity. The measurement process is described, in its utmost generality, by a many-round quantum interaction protocol between the experimental system and a full-fledged quantum computer. The QUALM complexity is quantified by the number of elementary operations performed by the quantum computer, including its coupling to the experimental system. We study how the QUALM complexity depends on the type of allowed access the quantum computer has to the experimental system: local-local, incoherent, coherent, adaptive, etc. We provide the first example of a measurement "task" for which the coherent QUALM complexity is exponentially better than the incoherent one, even if the latter is adaptive; this implies that using entanglement between different systems in experiments may lead to exponential savings in resources. We extend our results to derive a similar exponential advantage for a physically motivated measurement task which determines the symmetry class of the time evolution operator for a quantum many-body system. Many open questions are raised towards better understanding how quantum computational tools can be applied in experimental physics. A major question is whether an exponential advantage in QUALM complexity can be achieved in the NISQ era; an equally important one is to design new, efficient quantum algorithmic measurements based on our framework, perhaps relying on ideas from quantum algorithms.

Tony Metger, Yfke Dulek, Andrea Coladangelo, Rotem Arnon-Friedman and Thomas Vidick

ETH Zurich | QuSoft, University of Amsterdam | UC Berkeley | Weizmann Institute of Science | Caltech

Device-independent protocols from computational assumptions

Device-independent protocols use untrusted quantum devices to achieve a cryptographic task. Such protocols are typically based on Bell inequalities and require the assumption that the quantum device is composed of separated non-communicating components. In this submission, we present protocols for self-testing and device-independent quantum key distribution (DIQKD) that are secure even if the components of the quantum device can exchange arbitrary quantum communication. Instead, we assume that the device cannot break a standard post-quantum cryptographic assumption. Importantly, the computational assumption only needs to hold during the protocol execution and only applies to the (adversarially prepared) device in possession of the (classical) user, while the adversary herself remains unbounded. The output of the protocol, e.g. secret keys in the case of DIQKD, is information-theoretically secure. For our self-testing protocol, we build on a recently introduced cryptographic tool (Brakerski et al., FOCS 2018; Mahadev, FOCS 2018) to show that a classical user can enforce a bipartite structure on the Hilbert space of a black-box quantum device, and certify that the device has prepared and measured a state that is entangled with respect to this bipartite structure. Using our self-testing protocol as a building block, we construct a protocol for DIQKD that leverages the computational assumption to produce information-theoretically secure keys. The security proof of our DIQKD protocol uses the self-testing theorem in a black-box way. Our self-testing theorem thus also serves as a first step towards a more general translation procedure for standard device-independent protocols to the setting of computationally bounded (but freely communicating) devices.

Srijita Kundu and Ernest Y.-Z. Tan // Anne Broadbent and Rabib Islam

National University of Singapore | ETH Zurich // University of Ottawa | University of Ottawa

Composably secure device-independent encryption with certified deletion // Quantum encryption with certified deletion

We study the task of encryption with certified deletion (ECD) introduced by Broadbent and Islam (2019), but in a device-independent setting: we show that it is possible to achieve this task even when the honest parties do not trust their quantum devices. Moreover, we define security for the ECD task in a composable manner and show that our ECD protocol achieves composable security. Our protocol is based on device-independent quantum key distribution (DIQKD), and in particular the parallel DIQKD protocol based on the magic square non-local game, given by Jain, Miller and Shi (2017). To achieve certified deletion, we use a property of the magic square game observed by Fu and Miller (2017), namely that a two-round variant of the game can be used to certify deletion of a single random bit. In order to achieve certified deletion security for arbitrarily long messages from this property, we prove a parallel repetition theorem for two-round non-local games, which may be of independent interest.

Given a ciphertext, is it possible to prove the deletion of the underlying plaintext? Since classical ciphertexts can be copied, clearly such a feat is impossible using classical information alone. In stark contrast to this, we show that quantum encodings enable certified deletion. More precisely, we show that it is possible to encrypt classical data into a quantum ciphertext such that the recipient of the ciphertext can produce a classical string which proves to the originator that the recipient has relinquished any chance of recovering the plaintext should the decryption key be revealed. Our scheme is feasible with current quantum technology: the honest parties only require quantum devices for single-qubit preparation and measurements; the scheme is also robust against noise in these devices. Furthermore, we provide an analysis that is suitable in the finite-key regime.

Andrea Coladangelo, Christian Majenz and Alexander Poremba // Anne Broadbent, Stacey Jeffery, Sébastien Lord, Supartha Podder and Aarthi Sundaram

UC Berkeley | CWI and QuSoft | Caltech // University of Ottawa | QuSoft and CWI | University of Ottawa | University of Ottawa | Microsoft Quantum

Quantum Copy-Protection of Compute-and-Compare Programs in the Quantum Random Oracle Model // Secure Software Leasing Without Assumptions

Copy-protection allows a software distributor to encode a program in such a way that it can be evaluated on any input, yet it cannot be ``pirated'' -- a notion that is impossible to achieve in a classical setting. Aaronson (CCC 2009) initiated the formal study of quantum copy-protection schemes, and speculated that quantum cryptography could offer a solution to the problem thanks to the quantum no-cloning theorem. In this work, we introduce a quantum copy-protection scheme for a large class of evasive functions known as ``compute-and-compare programs'' -- a more expressive generalization of point functions. A compute-and-compare program CC$[f, y]$ is specified by a function f and a string y within its range: on input x, CC$[f, y]$ outputs 1, if $f(x) = y$, and 0 otherwise. We prove that our scheme achieves non-trivial security against fully malicious adversaries in the quantum random oracle model (QROM), which makes it the first copy-protection scheme to enjoy any level of provable security in a standard cryptographic model. As a complementary result, we show that the same scheme fulfils a weaker notion of software protection, called ``secure software leasing'', introduced very recently by Ananth and La Placa (eprint 2020), with a standard security bound in the QROM, i.e. guaranteeing negligible adversarial advantage. // Quantum cryptography is known for enabling functionalities that are unattainable using classical information alone. Recently, Secure Software Leasing (SSL) has emerged as one of these areas of interest. Given a target circuit C from a circuit class, SSL produces an encoding of C which enables the evaluation of C, and also enables a verify procedure, by which the originator of the software becomes convinced that the software is returned --- meaning that the recipient has relinquished the possibility of any further use of the software. Clearly, such functionality is unachievable using classical information alone, since it is impossible to prevent a user from keeping a copy of the software. Recent results have shown the achievability of SSL using quantum information for a class of functions called compute-and-compare (these are a generalization of the well-known point functions). These prior works, however, all make use of setup or computational assumptions. Here, we show that SSL is achievable for compute-and-compare circuits without any assumptions. Our technique is a generic reduction from any quantum message authentication code to such an SSL scheme. Along the way, we also show that point functions can be copy-protected without any assumptions, for a security definition that involves one honest and one malicious evaluator.

Abhinav Deshpande, Alexey Gorshkov and Bill Fefferman

QuICS/JQI/University of Maryland | Joint Quantum Institute (JQI) | University of Chicago

The importance of the spectral gap in estimating ground-state energies

The field of quantum Hamiltonian complexity lies at the intersection of quantum many-body physics and computational complexity theory, with deep implications to both fields. The main object of study is the LocalHamiltonian problem, which is concerned with estimating the ground-state energy of a local Hamiltonian and is complete for the class QMA, a quantum generalization of the class NP. A major challenge in the field is to understand the complexity of the LocalHamiltonian problem in more physically natural parameter regimes. One crucial parameter in understanding the ground space of any Hamiltonian in many-body physics is the spectral gap, which is the difference between the smallest two eigenvalues. Despite its importance in quantum many-body physics, the role played by the spectral gap in the complexity of the LocalHamiltonian is less well-understood. In this work, we make progress on this question by considering the precise regime, in which one estimates the ground-state energy to within inverse exponential precision. Computing ground-state energies precisely is a task that is important for quantum chemistry and quantum many-body physics. In the setting of inverse-exponential precision, there is a surprising result that the complexity of LocalHamiltonian is magnified from QMA to PSPACE, the class of problems solvable in polynomial space. We clarify the reason behind this boost in complexity. Specifically, we show that the full complexity of the high precision case only comes about when the spectral gap is exponentially small. As a consequence of the proof techniques developed to show our results, we uncover important implications for the representability and circuit complexity of ground states of local Hamiltonians, the theory of uniqueness of quantum witnesses, and techniques for the amplification of quantum witnesses in the presence of postselection.

Anurag Anshu, Srinivasan Arunachalam, Tomotaka Kuwahara and Mehdi Soleimanifar

University of California, Berkeley | IBM T. J. Watson Research Center | RIKEN Center for Advanced Intelligence Project | MIT

Sample-efficient learning of quantum many-body systems

We study the problem of learning the Hamiltonian of a quantum many-body system given samples from its Gibbs (thermal) state. The classical analog of this problem, known as learning graphical models or Boltzmann machines, is a well-studied question in machine learning and statistics. In this work, we give the first sample-efficient algorithm for the quantum Hamiltonian learning problem. In particular, we prove that polynomially many samples in the number of particles (qudits) are necessary and sufficient for learning the parameters of a spatially local Hamiltonian in l2-norm. Our main contribution is in establishing the strong convexity of the log-partition function of quantum many-body systems, which along with the maximum entropy estimation yields our sample-efficient algorithm. Classically, the strong convexity for partition functions follows from the Markov property of Gibbs distributions. This is, however, known to be violated in its exact form in the quantum case. We introduce several new ideas to obtain an unconditional result that avoids relying on the Markov property of quantum systems, at the cost of a slightly weaker bound. In particular, we prove a lower bound on the variance of quasi-local operators with respect to the Gibbs state, which might be of independent interest. Our work paves the way toward a more rigorous application of machine learning techniques to quantum many-body problems.

Sisi Zhou and Liang Jiang

Yale University | University of Chicago

Asymptotic theory of quantum channel estimation

The quantum Fisher information (QFI), as a function of quantum states, measures the amount of information that a quantum state carries about an unknown parameter. The (entanglement-assisted) QFI of a quantum channel is defined to be the maximum QFI of the output state assuming an entangled input state over a single probe and an ancilla. In quantum metrology, people are interested in calculating the QFI of N identical copies of a quantum channel when $N\to\infty$, which we call the asymptotic QFI. It was known that the asymptotic QFI grows either linearly or quadratically with N. Here we obtain a simple criterion that determines whether the scaling is linear or quadratic. In both cases, the asymptotic QFI and a quantum error correction protocol to achieve it are solvable via a semidefinite program. When the scaling is quadratic, the Heisenberg limit, a feature of noiseless quantum channels, is recovered. When the scaling is linear, we show the asymptotic QFI is still in general larger than N times the single-channel QFI and furthermore, sequential estimation strategies provide no advantage over parallel ones. For details, see arXiv: 2003.10559.

Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao and Avishay Tal

The University of Texas at Austin | University of Waterloo | Microsoft | Northwestern University | University of California at Berkeley

Degree vs. Approximate Degree and Quantum Implications of Huang's Sensitivity Theorem

Based on the recent breakthrough of Huang (2019), we show that for any total Boolean function f, deg(f) = O(∼deg(f)^2): The degree of f is at most quadratic in the approximate degree of f. This is optimal as witnessed by the OR function. D(f) = O(Q(f)^4): The deterministic query complexity of f is at most quartic in the quantum query complexity of f. This matches the known separation (up to log factors) due to Ambainis, Balodis, Belovs, Lee, Santha, and Smotrovs (2017). We apply these results to resolve the quantum analogue of the Aanderaa--Karp--Rosenberg conjecture. We show that if f is a nontrivial monotone graph property of an n-vertex graph specified by its adjacency matrix, then Q(f)=Ω(n), which is also optimal. We also show that the approximate degree of any read-once formula on n variables is Θ($\sqrt{n}$).

Matthew Hastings, Jeongwan Haah and Ryan O'Donnell

Microsoft | Microsoft | Carnegie Mellon University

Fiber Bundle Codes: Breaking the $N^{1/2}$polylog(N) Barrier for Quantum LDPC Codes

We present a quantum LDPC code family that has distance $\Omega(N^{3/5}/\text{polylog}(N))$ and $\Theta(N^{3/5})$ logical qubits, up to polylogs, where N is the code length. This is the first quantum LDPC code construction which achieves distance greater than $N^{1/2}$polylog(N). The construction is based on generalizing the homological product of codes to a fiber bundle.

Laura Mančinska and David Roberson

University of Copenhagen | Technical University of Denmark

Quantum isomorphism is equivalent to equal homomorphism counts from planar graphs

Over 50 years ago, Lovasz proved that two graphs are isomorphic if and only if they admit the same number of homomorphisms from any graph. Other equivalence relations on graphs, such as cospectrality or fractional isomorphism, can be characterized by equality of homomorphism counts from an appropriately chosen class of graphs. Dvorak [J. Graph Theory 2010] showed that taking this class to be the graphs of treewidth at most k yields a tractable relaxation of graph isomorphism known as k-dimensional Weisfeiler-Leman equivalence. Together with a famous result of Cai, Furer, and Immerman [FOCS 1989], this shows that homomorphism counts from graphs of bounded treewidth do not determine a graph up to isomorphism. Dell, Grohe, and Rattan [ICALP 2018] raised the questions of whether homomorphism counts from planar graphs determine a graph up to isomorphism, and what is the complexity of the resulting relation. We answer the former in the negative by showing that the resulting relation is equivalent to the so-called quantum isomorphism [Mancinska et al, ICALP 2017]. Using this equivalence, we further resolve the latter question, showing that testing whether two graphs have the same number of homomorphisms from any planar graph is, surprisingly, an undecidable problem, and moreover is complete for the class coRE (the complement of recursively enumerable problems). Quantum isomorphism is defined in terms of a one-round, two-prover interactive proof system in which quantum provers, who are allowed to share entanglement, attempt to convince the verifier that the graphs are isomorphic. Our combinatorial proof leverages the quantum automorphism group of a graph, a notion from noncommutative mathematics.

Giacomo De Palma, Milad Marvian, Dario Trevisan and Seth Lloyd

MIT | University of New Mexico | University of Pisa | MIT

The quantum Wasserstein distance of order 1

We propose a generalization of the Wasserstein distance of order 1 to the quantum states of n qudits. The proposal recovers the Hamming distance for the vectors of the canonical basis, and more generally the classical Wasserstein distance for quantum states diagonal in the canonical basis. The proposed distance is invariant with respect to permutations of the qudits and unitary operations acting on one qudit and is additive with respect to the tensor product. Our main result is a continuity bound for the von Neumann entropy with respect to the proposed distance, which significantly strengthens the best continuity bound with respect to the trace distance. We also propose a generalization of the Lipschitz constant to quantum observables. The notion of quantum Lipschitz constant allows us to compute the proposed distance with a semidefinite program. We prove a quantum version of Marton's transportation inequality and a quantum Gaussian concentration inequality for the spectrum of quantum Lipschitz observables. Moreover, we derive bounds on the contraction coefficients of shallow quantum circuits and of the tensor product of one-qudit quantum channels with respect to the proposed distance. We discuss other possible applications in quantum machine learning, quantum Shannon theory, and quantum many-body systems.

Simon Becker, Nilanjana Datta, Ludovico Lami and Cambyse Rouzé

University of Cambridge | University of Cambridge | Ulm University | Technische Universität München

Energy-constrained discrimination of unitaries, quantum speed limits and a Gaussian Solovay-Kitaev theorem

We investigate the energy-constrained (EC) diamond norm distance between unitary channels acting on possibly infinite-dimensional quantum systems, and establish a number of results. Firstly, we prove that optimal EC discrimination between two unitary channels does not require the use of any entanglement. Extending a result by Acin, we also show that a finite number of parallel queries suffices to achieve zero error discrimination even in this EC setting. Secondly, we employ EC diamond norms to study a novel type of quantum speed limits, which apply to pairs of quantum dynamical semigroups. We expect these results to be relevant for benchmarking internal dynamics of quantum devices. Thirdly, we establish a version of the Solovay-Kitaev theorem that applies to the group of Gaussian unitaries over a finite number of modes, with the approximation error being measured with respect to the EC diamond norm relative to the photon number Hamiltonian.

Gergely Bunth, Christopher Perry, Péter Vrana and Albert H. Werner

Budapest University of Technology and Economics | QMath - University of Copenhagen | Budapest University of Technology and Economics | QMath - University of Copenhagen

The semiring of dichotomies and asymptotic relative submajorization

We study quantum dichotomies and the resource theory of asymmetric distinguishability using a generalization of Strassen's theorem on preordered semirings. We find that an asymptotic variant of relative submajorization, defined on unnormalized dichotomies, is characterized by real-valued monotones that are multiplicative under the tensor product and additive under the direct sum. These strong constraints allow us to classify and explicitly describe all such monotones, leading to a rate formula expressed as an optimization involving sandwiched Renyi divergences. As an application we give a new derivation of the strong converse error exponent in quantum hypothesis testing.

Alexander Sherstov, Andrey Storozhenko and Pei Wu

University of California, Los Angeles | University of California, Los Angeles | University of California, Los Angeles

An Optimal Separation of Randomized and Quantum Query Complexity

We prove that for every decision tree, the absolute values of the Fourier coefficients of given order $\ell \leq 1$ sum to at most $c^{\ell}\sqrt{\binom{d}{\ell}(1+\log n)^{\ell-1}}$, where n is the number of variables, d is the tree depth, and c>0 is an absolute constant. This bound is essentially tight and settles a conjecture due to Tal (arxiv 2019; FOCS 2020). The bounds prior to our work degraded rapidly with ℓ becoming trivial already at $\ell = \sqrt{d}$. As an application, we obtain, for every integer $k \geq 1$, a partial Boolean function on n bits that has bounded-error quantum query complexity at most $\lceil k/2 \rceil$ and randomized query complexity $\widetilde{\Omega}\left(n^{1-1/k}\right)$. This separation of bounded-error quantum versus randomized query complexity is best possible, by the results of Aaronson and Ambainis (STOC 2015). Prior to our work, the best known separation was polynomially weaker: $O(1)$ versus $\Omega\left(n^{\frac{2}{3}-\epsilon}\right)$ for any $\epsilon > 0$ (Tal, FOCS 2020). As another application, we obtain an essentially optimal separation of $O(\log n)$ versus $\Omega(n^{1-\epsilon})$ for bounded-error quantum versus randomized communication complexity, for any $\epsilon > 0$. The best previous separation was polynomially weaker: of $O(\log n)$ versus $\Omega\left(n^{\frac{2}{3}-\epsilon}\right)$ (implicit in Tal, FOCS 2020).

Makrand Sinha and Nikhil Bansal

CWI Amsterdam | CWI Amsterdam and TU Eindhoven

k-Forrelation Optimally Separates Quantum and Classical Query Complexity

Aaronson and Ambainis (SICOMP `18) showed that any partial function on N bits that can be computed with an advantage δ over a random guess by making q quantum queries, can also be computed classically with an advantage $\delta/2$ by a randomized decision tree making $O_q\left(N^{1-\frac{1}{2q}}\delta^{-2}\right)$ queries. Moreover, they conjectured the k-Forrelation problem --- a partial function that can be computed with $q = \lceil k/2 \rceil$ quantum queries --- to be a suitable candidate for exhibiting such an extremal separation. We prove their conjecture by showing a tight lower bound of $\widetilde{\Omega}\left(N^{1-1/k}\right)$ for the randomized query complexity of k-Forrelation, where the advantage $\delta = 2^{-O(k)}$. By standard amplification arguments, this gives an explicit partial function that exhibits an $O_\epsilon(1)$ vs $\Omega(N^{1-\epsilon})$ separation between bounded-error quantum and randomized query complexities, where $\epsilon > 0$ can be made arbitrarily small. Our proof also gives the same bound for the closely related but non-explicit k-Rorrelation function introduced by Tal (FOCS `20). Our techniques rely on classical Gaussian tools, in particular, Gaussian interpolation and Gaussian integration by parts, and in fact, give a more general statement. We show that to prove lower bounds for k-Forrelation against a family of functions, it suffices to bound the ℓ_1-weight of the Fourier coefficients between levels k and (k-1)k. We also prove new interpolation and integration by parts identities that might be of independent interest in the context of rounding high-dimensional Gaussian vectors.

Anurag Anshu, Shalev Ben-David and Srijita Kundu

University of California, Berkeley | University of Waterloo | National University of Singapore

On Query-to-Communication Lifting of Quantum Adversaries

We investigate query-to-communication lifting theorems for models related to the quantum adversary bounds. Our results are as follows: 1. We show that the classical adversary bound lifts to a lower bound on randomized communication complexity with a constant-sized gadget. We also show that the classical adversary bound is a strictly stronger lower bound technique than the previously-lifted measure known as critical block sensitivity, making our lifting theorem one of the strongest lifting theorems for randomized communication complexity using a constant-sized gadget. 2. Turning to quantum models, we show a connection between lifting theorems for quantum adversary bounds and secure 2-party quantum computation in a certain "honest-but-curious" model. Under the assumption that such secure 2-party computation is impossible, we show that a simplified version of the positive-weight adversary bound lifts to a quantum communication lower bound using a constant-sized gadget. We also give an unconditional lifting theorem which lower bounds bounded-round quantum communication protocols. 3. Finally, we give some new results in query complexity. We show that the classical adversary and the positive-weight quantum adversary are quadratically related. We also show that the positive-weight quantum adversary is never larger than the square of the approximate degree. Both relations hold even for partial functions.

Mirjam Weilenmann, Edgar A. Aguilar and Miguel Navascués

IQOQI Vienna | IQOQI Vienna | IQOQI Vienna

Quantum Preparation Games

A preparation game is a task whereby a player sequentially sends a number of quantum states to a referee, who probes each of them and announces the measurement result. The measurement setting in each round, as well as the final score of the game, are decided by the referee based on the past history of settings and measurement outcomes. Many experimental tasks in quantum information, such as entanglement quantification or magic state detection, can be cast as preparation games. In this paper, we introduce general methods to design n-round preparation games, with tight bounds on the average game scores achievable by players subject to constraints on their preparation devices. We illustrate our results by devising new adaptive measurement protocols for entanglement detection and quantification. Surprisingly, we find that the standard procedure in entanglement detection, namely, estimating n times the average value of a given entanglement witness, is in general sub-optimal for detecting the entanglement of a specific quantum state. On the contrary, there exist n-round experimental scenarios where detecting the entanglement of a known state optimally requires adaptive measurement schemes.

Harry Buhrman, Noah Linden, Laura Mančinska, Ashley Montanaro and Maris Ozols

Centrum Wiskunde & Informatica | University of Bristol | University of Copenhagen | PhaseCraft Ltd and University of Bristol | University of Amsterdam

Quantum majority and other Boolean functions with quantum inputs

Majority vote is a basic method for amplifying correct outcomes that is widely used in computer science and beyond. It can, for example, be used to amplify the correctness of a quantum device whose output is classical. However, when the output of a device is a quantum state, it is not apriori clear how to implement an analogous *quantum* majority vote. To this end, we consider an extension of majority vote to quantum inputs and outputs: given a product state of the form $|\phi_1\rangle\otimes|\phi_2\rangle\otimes\cdots\otimes|\phi_n\rangle$, where each qubit $|\phi_i\rangle$ is in one of two orthogonal states $|\psi_0\rangle$ or $|\psi_1\rangle$, output the majority state $|\psi_0\rangle$ or $|\psi_1\rangle$. We provide an optimal algorithm for this problem that achieves worst-case fidelity of $1/2 + \Theta\left(1/\sqrt{n}\right)$. Under the promise that at least 2/3 of the qubits are in the majority state, the fidelity increases to of $1 - \Theta(1/n)$ and approaches one in the limit. More generally, we initiate the study of covariant and symmetric Boolean functions $f\colon \{0,1^n\} \to \{0,1\}$ with quantum inputs and outputs. We provide a simple linear program of size roughly n/2 for computing the optimal worst-case fidelity and show that a generalization of our algorithm is optimal for computing f. Our algorithm has complexity $O(n^4 \log n)$ where n is the number of qubits.

Zuzana Gavorova, Matan Seidel and Yonathan Touati

The Hebrew University of Jerusalem | Tel Aviv University | The Hebrew University of Jerusalem

Topological obstructions to implementing controlled unknown unitaries

Is a quantum algorithm capable of implementing an if-clause? Given a black-box subroutine, a d-dimensional unitary U from U(d), a quantum if-clause would correspond to applying it to an input qudit if and only if the value of a control qubit is 1. It was previously shown by Thompson et al. and Araujo et al. that implementations using single access to the oracle U are impossible. Our main result is a strong generalization of this impossibility result: we prove that there is no unitary oracle algorithm implementing $control_\phi(U) = |0\rangle\langle 0| \otimes I + e^{i\phi(U)}|1\rangle\langle 1| \otimes U$, for some U-dependent phase $\phi(U)$, even if allowed any finite number of calls to U and its inverse, and even if required to only approximate the desired operator $control_\phi(U)$. Even further, there is no such postselection oracle algorithm, i.e. a unitary oracle algorithm followed by a binary success/fail measurement, that reports success and implements $control_\phi(U)$ with a nonzero probability for each U in U(d). Our proof relies on topological arguments which can be viewed as a modification of the Borsuk-Ulam theorem. Combining the topological arguments with the algorithm of Dong et al. in fact leads to an interesting dichotomy: implementing $control_\phi(U^m)$ in our model is possible if and only if the integer m is a multiple of the oracle's dimension d. Our impossibility no longer holds if the model is relaxed, either by dropping the worst-case requirement that the algorithm works for all U from U(d), or, remarkably, by allowing measurements more general than a binary postselection measurement. We observe that for both relaxations, inefficient algorithms exist, and it remains open whether efficient ones do.

Atul Singh Arora, Jérémie Roland and Chrysoula Vlachou

Université libre de Bruxelles (ULB) and Caltech | Université libre de Bruxelles (ULB) | Université libre de Bruxelles (ULB)

Analytic quantum weak coin flipping protocols with arbitrarily small bias

Weak coin flipping (WCF) is a fundamental cryptographic primitive for two-party secure computation, where two distrustful parties need to remotely establish a shared random bit whilst having opposite preferred outcomes. It is the strongest known primitive with arbitrarily close to perfect security quantumly while classically, its security is completely compromised (unless one makes further assumptions, such as computational hardness). A WCF protocol is said to have bias ϵ if neither party can force their preferred outcome with probability greater than $1/2 + \epsilon$. Classical WCF protocols are shown to have bias 1/2, i.e., a cheating party can always force their preferred outcome. On the other hand, there exist quantum WCF protocols with arbitrarily small bias, as Mochon showed in his seminal work in 2007 [arXiv:0711.4114]. In particular, he proved the existence of a family of WCF protocols approaching bias $\epsilon(k)=1/(4k + 2)$ for arbitrarily large k and proposed a protocol with bias 1/6. Last year, Arora, Roland and Weis presented a protocol with bias 1/10 and to go below this bias, they designed an algorithm that numerically constructs unitary matrices corresponding to WCF protocols with arbitrarily small bias [STOC'19, p.205-216]. In this work, we present new techniques which yield a fully analytical construction of WCF protocols with bias arbitrarily close to zero, thus achieving a solution that has been missing for more than a decade. Furthermore, our new techniques lead to a simplified proof of existence of WCF protocols by circumventing the non-constructive part of Mochon's proof. As an example, we illustrate the construction of a WCF protocol with bias 1/14.

Anurag Anshu, Aram Harrow and Mehdi Soleimanifar

University of California, Berkeley | Massachusetts Institute of Technology | MIT

From communication complexity to an entanglement spread area law in the ground state of gapped local Hamiltonians

In this work, we make a connection between two seemingly different problems. The first problem involves characterizing the properties of entanglement in the ground state of gapped local Hamiltonians, which is a central topic in quantum many-body physics. The second problem is on the quantum communication complexity of testing bipartite states with EPR assistance, a well-known question in quantum information theory. We construct a communication protocol for testing (or measuring) the ground state and use its communication complexity to reveal a new structural property for the ground state entanglement. This property, known as the entanglement spread, roughly measures the log of the ratio between the largest and the smallest Schmidt coefficients across a bipartite cut in the ground state. Our main result shows that gapped ground states possess limited entanglement spread across any cut, exhibiting an "area law" behavior. Our result applies to any interaction graph with an improved bound for the special case of lattices. This entanglement spread area law includes interaction graphs constructed in [AHL+14] that violate a generalized area law for the entanglement entropy. Our construction also provides evidence for a conjecture in physics by Li and Haldane on the entanglement spectrum of lattice Hamiltonians [LH08]. On the technical side, we use recent advances in Hamiltonian simulation algorithms along with the quantum phase estimation to give a new construction for an approximate ground space projector (AGSP) over arbitrary interaction graphs, which might be of independent interest.

Noah Shutty, Mary Wootters and Patrick Hayden

Stanford University | Stanford University | Stanford University

Tight Limits on Nonlocality from Nontrivial Communication Complexity

It has long been known that the existence of certain superquantum nonlocal correlations would cause communication complexity to collapse. The absurdity of a world in which any function could be evaluated by two players with a constant amount of communication in turn provides a tantalizing way to distinguish quantum mechanics from incorrect theories of physics; the statement “communication complexity is nontrivial” has even been conjectured to be a concise information-theoretic axiom for characterizing quantum mechanics. We directly address the viability of that perspective with two results. First, we exhibit a nonlocal game such that communication complexity collapses in any physical theory whose maximal winning probability exceeds the quantum value. Second, we consider the venerable CHSH game that initiated this line of inquiry. In that case, the quantum value is about 0.85 but it is known that a winning probability of approximately 0.91 would collapse communication complexity. We show that the 0.91 result is the best possible using a large class of proof strategies, suggesting that the communication complexity axiom is insufficient for characterizing CHSH correlations. Both results build on new insights about reliable classical computation. The first exploits our formalization of an equivalence between amplification and reliable computation, while the second follows from a rigorous determination of the threshold for reliable computation with formulas of noise-free XOR gates and ε-noisy AND gates.

Jonas Haferkamp, Felipe Montealegre-Mora, Markus Heinrich, Jens Eisert, David Gross and Ingo Roth

FU Berlin | Universität Köln | Universität Köln | FU Berlin | Universität Köln | FU Berlin

Efficient unitary designs with a system-size independent number of non-Clifford gates

Many quantum information protocols require the implementation of random unitaries. Because it takes exponential resources to produce Haar-random unitaries drawn from the full n-qubit group, one often resorts to t-designs. Unitary t-designs mimic the Haar-measure up to t-th moments. It is known that Clifford operations can implement at most 3-designs. In this work, we quantify the non-Clifford resources required to break this barrier. We find that it suffices to inject $O(t^4 \log^2(t) \log(1/\varepsilon))$ many non-Clifford gates into a polynomial-depth random Clifford circuit to obtain an ε-approximate t-design. Strikingly, the number of non-Clifford gates required is independent of the system size -- asymptotically, the density of non-Clifford gates is allowed to tend to zero. We also derive novel bounds on the convergence time of random Clifford circuits to the t-th moment of the uniform distribution on the Clifford group. Our proofs exploit a recently developed variant of Schur-Weyl duality for the Clifford group, as well as bounds on restricted spectral gaps of averaging operators.

Michał Oszmaniec, Adam Sawicki and Michal Horodecki

Center for Theoretical Physics, Polish Academy of Sciences | Center for Theoretical Physics, Polish Academy of Sciences | International Centre for Theory of Quantum Technologies, University of Gdansk

Epsilon-nets, unitary designs and random quantum circuits

Epsilon-nets and approximate unitary t-designs are natural notions that capture properties of unitary operations relevant for numerous applications in quantum information and quantum computing. The former constitute subsets of unitary channels that are epsilon-close to any unitary channel in the diamond norm. The latter are ensembles of unitaries that (approximately) recover Haar averages of polynomials in entries of unitary channels up to order t. In this work we systematically study quantitative connections between these two notions. Specifically, we prove that, for a fixed dimension d of the Hilbert space, unitaries constituting δ-approximate t-expanders form ϵ-nets for $t \simeq \frac{d^{5/2}}{\epsilon}$ and $\delta \simeq \left(\frac{\epsilon^{3/2}}{d}\right)^{d^2}$. We also show that ϵ-nets can be used to construct δ-approximate unitary t-designs for $\delta \simeq \epsilon t$, where the notion of approximation is based on the diamond norm. Finally, we prove that the degree of an exact unitary t-design necessary to obtain an ϵ-net must grow at least fast as $1/\epsilon$ (for fixed dimension) and not slower than d^2 (for fixed epsilon). This shows near optimality of our result connecting t-designs and ϵ-nets. We further apply our findings in conjunction with the recent results of Varju 2013 in the context of quantum computing. First, we show that that approximate t-designs can be generated by shallow random circuits formed from a set of universal two-qudit gates in the parallel and sequential local architectures considered in Brandao-Harrow-Horodecki 2016. Importantly, our gate sets need not to be symmetric (i.e. contains gates together with their inverses) or consist of gates with algebraic entries. Second, we consider a problem of compilation of quantum gates and prove a non-constructive version of the Solovay-Kitaev theorem for general universal gate sets. Our main technical contribution is a new construction of efficient polynomial approximations to the Dirac delta in the space of quantum channels, which can be of independent interest.

Michael Gullans, Stefan Krastanov, David Huse, Liang Jiang and Steven Flammia

Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park | John A. Paulson School of Engineering and Applied Sciences, Harvard University | Department of Physics, Princeton University | University of Chicago | AWS

Quantum coding with low-depth random circuits

Random quantum circuits have played a central role in establishing the computational advantages of near-term quantum computers over their conventional counterparts. Here, we use ensembles of low-depth random circuits with local connectivity in D spatial dimensions to generate quantum error-correcting codes. For random stabilizer codes and the erasure channel, we find strong evidence that a depth $O(\log N)$ random circuit is necessary and sufficient to converge (with high probability) to zero failure probability for any finite amount below the channel capacity for any D. Previous results on random circuits have only shown that $O(N^{1/D})$ depth suffices or that $O(\log^3 N)$ depth suffices for all-to-all connectivity. We then study the critical behavior of the erasure threshold in the so-called moderate deviation limit, where both the failure probability and the distance to the channel capacity converge to zero with N. We find that the requisite depth scales like $O(\log N)$ only for dimensions $D \geq 2$, and that random circuits require $O(\sqrt{N})$ depth for $D = 1$. Finally, we introduce an "expurgation" algorithm that uses quantum measurements to remove logical operators that cause the code to fail by turning them into either additional stabilizers or into gauge operators in a subsystem code. With such targeted measurements, we can achieve sub-logarithmic depth in $D \geq 2$ spatial dimensions below capacity without increasing the maximum weight of the check operators. We find that for any rate beneath the capacity, high-performing codes with thousands of logical qubits are achievable with depth 4-8 expurgated random circuits in $D = 2$ dimensions. These results indicate that finite-rate quantum codes are practically relevant for near-term devices and may significantly reduce the resource requirements to achieve fault tolerance for near-term applications.

Piotr Czarnik, Andrew Arrasmith, Patrick Coles and Lukasz Cincio

Los Alamos National Laboratory | Los Alamos National Lab | Los Alamos National Laboratory | Los Alamos National Laboratory

Error mitigation with Clifford quantum-circuit data

Achieving near-term quantum advantage will require accurate estimation of quantum observables despite significant hardware noise. For this purpose, we propose a novel, scalable error-mitigation method that applies to gate-based quantum computers. The method generates training data $\{X_i^{noisy}, X_i^{exact}\}$ via quantum circuits composed largely of Clifford gates, which can be efficiently simulated classically, where X_i^{noisy} and X_i^{exact} are noisy and noiseless observables respectively. Fitting a linear ansatz to this data then allows for the prediction of noise-free observables for arbitrary circuits. We analyze the performance of our method versus the number of qubits, circuit depth, and number of non-Clifford gates. We obtain an order-of-magnitude error reduction for a ground-state energy problem on 16 qubits in an IBMQ quantum computer and on a 64-qubit noisy simulator.

Guang Hao Low, Vera von Burg, Thomas Haner, Damian Steiger, Markus Reiher, Martin Roetteler and Matthias Troyer

Microsoft | ETH Zurich | Microsoft | Microsoft | ETH Zurich | Microsoft | Microsoft

Quantum computing enhanced computational catalysis

The quantum computation of electronic energies can break the curse of dimensionality that plagues many-particle quantum mechanics. It is for this reason that a universal quantum computer has the potential to fundamentally change computational chemistry and materials science, areas in which strong electron correlations present severe hurdles for traditional electronic structure methods. Here, we present a state-of-the-art analysis of accurate energy measurements on a quantum computer for computational catalysis, using improved quantum algorithms with more than an order of magnitude improvement over the best previous algorithms. As a prototypical example of local catalytic chemical reactivity we consider the case of a ruthenium catalyst that can bind, activate, and transform carbon dioxide to the high-value chemical methanol. We aim at accurate resource estimates for the quantum computing steps required for assessing the electronic energy of key intermediates and transition states of its catalytic cycle. In particular, we present new quantum algorithms for double-factorized representations of the four-index integrals that can significantly reduce the computational cost over previous algorithms, and we discuss the challenges of increasing active space sizes to accurately deal with dynamical correlations. We address the requirements for future quantum hardware in order to make a universal quantum computer a successful and reliable tool for quantum computing enhanced computational materials science and chemistry, and identify open questions for further research.

Isaac Kim, Eugene Tang and John Preskill

The University of Sydney | California Institute of Technology | Caltech

The ghost in the radiation: Robust encodings of the black hole interior

We reconsider the black hole firewall puzzle, emphasizing that quantum error-correction, computational complexity, and pseudorandomness are crucial concepts for understanding the black hole interior. We assume that the Hawking radiation emitted by an old black hole is pseudorandom, meaning that it cannot be distinguished from a perfectly thermal state by any efficient quantum computation acting on the radiation alone. We then infer the existence of a subspace of the radiation system which we interpret as an encoding of the black hole interior. This encoded interior is entangled with the late outgoing Hawking quanta emitted by the old black hole, and is inaccessible to computationally bounded observers who are outside the black hole. Specifically, efficient operations acting on the radiation, those with quantum computational complexity polynomial in the entropy of the remaining black hole, commute with a complete set of logical operators acting on the encoded interior, up to corrections which are exponentially small in the entropy. Thus, under our pseudorandomness assumption, the black hole interior is well protected from exterior observers as long as the remaining black hole is macroscopic. On the other hand, if the radiation is not pseudorandom, an exterior observer may be able to create a firewall by applying a polynomial-time quantum computation to the radiation.

Cody Jones and Ryan Babbush, Google

Google | Google

Quantum Computer Science at Google

This talk will give an update regarding Google's plans in quantum computing. We will highlight some recent results from our team and discuss some example problems where engagement with the quantum computer science community endemic to QIP will be important as the field advances. The first part of our talk will focus on quantum error-correction and the second part of our talk will focus on quantum algorithms.

Sarah Sheldon, IBM

IBM

Demonstrating the capabilities of state-of-the-art quantum systems

Development of quantum hardware has accelerated in recent years and quantum systems with tens of qubits and error rates in the 10-2 range are now reliable and accessible. The trajectory of this development begs the question: how can we make near term systems useful without quantum error correction? The answer is that we need new and diverse advancements in theory, software, and hardware to allow us to approach ever harder problems. This talk will focus in particular on using classical resources to augment the capabilities of today's quantum hardware, whether through classical post processing techniques or classical computations within a quantum experiment. I will discuss recent work performed on IBM Quantum systems to demonstrate error mitigation, benchmarking techniques for measuring progress, and strategies for extending the range of accessible applications.

Yuval Sanders, Dominic Berry, Pedro Costa, Louis Tessler, Nathan Wiebe, Craig Gidney, Hartmut Neven and Ryan Babbush

Macquarie University | Macquarie University | Macquarie University | Macquarie University | University of Washington | Google | Google | Google

Compilation of Fault-Tolerant Quantum Heuristics for Combinatorial Optimization

We compile explicit circuits and evaluate the computational cost for heuristic-based quantum algorithms for combinatorial optimization. We consider several variants of quantum-accelerated simulated annealing as well as adiabatic algorithms, quantum-enhanced population transfer, the quantum approximate optimization algorithm, and other approaches. We provide novel methods for executing the bottleneck subroutines for these heuristics, and our methods can easily be applied to other algorithms where numerical performance matters. We estimate how quickly the subroutines could be executed on a modestly sized superconducting-qubit-based quantum computer with surface code error correction. We conclude that quadratic speedups for heuristic-based quantum optimization algorithms are insufficient for early quantum computers to beat present day classical computers.

Daniel Stilck França and Raul Garcia-Patron

University of Copenhagen | University of Edinburgh

Limitations of optimization algorithms on noisy quantum devices

Recent technological developments have focused the interest of the quantum computing community on investigating how near-term devices could outperform classical computers for practical applications. A central question that remains open is whether their noise can be overcome or it fundamentally restricts any potential quantum advantage. We present a transparent way of comparing classical algorithms to quantum ones running on near-term quantum devices for a large family of problems that include optimization problems and approximations to the ground state energy of Hamiltonians. Our approach is based on the combination of entropic inequalities that determine how fast the quantum computation state converges to the fixed point of the noise model, together with established classical methods of Gibbs state sampling. The approach is extremely versatile and allows for its application to a large variety of problems, noise models and quantum computing architectures. We use our results to provide estimates for a variety of problems and architectures that have been the focus of recent experiments, such as quantum annealers, variational quantum eigensolvers, and quantum approximate optimization. The bounds we obtain indicate that substantial quantum advantages are unlikely for classical optimization unless the current noise rates are decreased by orders of magnitude or the topology of the problem matches that of the device. This is the case even if the number of qubits increases substantially. We reach similar but less stringent conclusions for quantum Hamiltonian problems.

Kun Fang and Zi-Wen Liu // Bartosz Regula and Ryuji Takagi

University of Waterloo | Perimeter Institute for Theoretical Physics // Nanyang Technological University | Nanyang Technological University

No-go theorems and limitations for quantum resource purification

The manipulation of quantum “resources” such as entanglement and coherence lies at the heart of quantum science and technology, empowering potential advantages over classical methods. In practice, a particularly important kind of manipulation is to “purify” the quantum resources, since they are inevitably contaminated by noises and thus often lost their power or become unreliable for direct usage. In these two works, we establish a theory of the universal limitations on the accuracy and efficiency of resource purification tasks which apply to any well-behaved resource theory, for both state (static) and channel (dynamical) resources. The general results bring new insights and imply various forms of fundamental limits to a broad range of problems of great theoretical and practical importance, including magic state distillation and fault tolerant quantum computing, quantum error correction, quantum Shannon theory, and quantum circuit synthesis. //

We establish universal limitations on the manipulation of quantum channel resources under the most general transformation protocols. Focusing in particular on the class of distillation tasks -- which can be understood either as the purification of noisy channels into unitary ones, or the extraction of state-based resources from channels -- we develop fundamental restrictions on the error necessarily incurred in such transformations. Our results are applicable to the study of general quantum resources under any physical manipulation scheme, which includes general adaptive protocols with or without a definite causal order. We introduce comprehensive lower bounds for the overhead of any distillation protocol in terms of required channel uses, imposing strong limitations on the practical efficiency and cost of channel resource manipulation. In the asymptotic setting, our results yield broadly applicable strong converse bounds for the rates of distillation. As a special case, our methods apply to the manipulation of quantum states, in which case they significantly improve on and extend previous approaches. We demonstrate our results through explicit applications to quantum communication, where we recover in particular a number of strong converse bounds for the quantum capacity of channels assisted by different classes of operations, as well as to fault-tolerant quantum computation, where we obtain improved bounds for the overhead cost of magic state distillation and gate synthesis.

Piotr Kopszak, Marek Mozrzymas, Michal Studzinski and Michal Horodecki

University of Wroclaw | University of Wroclaw | University of Gdansk | International Centre for Theory of Quantum Technologies, University of Gdansk

Multi-port teleportation schemes

We introduce and discuss a novel multi-port based teleportation schemes performing transmission of a number of unknown quantum states or one composite system in one go. We fully characterize the probabilistic and deterministic case by presenting expressions for the average probability of success and entanglement fidelity in both non-optimal and optimal variant. We also deliver explicit forms of the measurements and the resource state exploited by parties to perform the process. To obtain our results, i.e. explicit expressions for the performance of the new schemes, we deliver novel mathematical tools concerning representation theory of the algebra of partially transposed permutation operators, where the transposition acts on more than one subsystem. Additionally, the optimal values of the entanglement fidelity and probability success emerge from formulated and solved primal and dual semidefinite problems, which due to existing symmetries and delivered mathematical tools could be solved analytically. Next, we have applied the obtained formulas for the performance of multi-port based teleportation schemes to get a qualitative improvement of asymptotic "teleportation capacities" of multi-port based teleportation schemes over the pre-existing port-based teleportation schemes.

Yuxiang Yang, Renato Renner and Giulio Chiribella

Institute for Theoretical Physics, ETH Zurich | Institute for Theoretical Physics, ETH Zurich | The University of Hong Kong

Optimal universal programming of unitary gates

A universal quantum processor is a device that takes as input a (quantum) program, containing an encoding of an arbitrary unitary gate, and a (quantum) data register, on which the encoded gate is applied. While no perfect universal quantum processor can exist, approximate processors have been proposed in the past two decades. A fundamental open question is how the size of the smallest quantum program scales with the approximation error. Here we answer the question, by proving a bound on the size of the program and designing a concrete protocol that attains the bound in the asymptotic limit. Our result is based on a connection between optimal programming and the Heisenberg limit of quantum metrology, and establishes an asymptotic equivalence between the tasks of programming, learning, and estimating unitary gates.

Tanmay Singal, Filip Maciejewski and Michał Oszmaniec

Center for Theoretical Physics, Polish Academy of Sciences | Center for Theoretical Physics, Polish Academy of Sciences | Center for Theoretical Physics, Polish Academy of Sciences

Implementation of quantum measurements using classical resources and only a single ancillary qubit

It is imperative to minimize resources needed to implement quantum operations on existing near-term quantum devices. With this in mind, we propose a scheme to implement an arbitrary general quantum measurement, also known as Positive Operator Valued Measures (POVM) in dimension d using only classical resources and a single ancillary qubit. The proposed method is based on probabilistic implementation of d outcome measurements which is followed by postselection on some of the received outcomes. This is an extension of an earlier work which required dichotomic measurements, no additional ancillary qubits, and whose success probability scaled like $\frac{1}{d}$. The success probability of our scheme depends on the operator norms of the coarse grained POVM effects. Significantly, we show that for typical Haar random rank-one POVMs with at most d^2 outcomes, the success probability of our simulation scheme does not go to zero with the dimension of the system. We conjecture that this is true for all POVMs in dimension d. This is supported by numerical computations showing constant success probability for SIC-POVMs and (non symmetric) IC-POVMs in dimensions up to 323. Additionally, for the gate noise model used in the recent demonstration of quantum computational advantage by Google, we prove that for typical Haar random POVMs noise compounding in circuits required by our scheme is substantially lower than in the scheme that directly uses Naimark's dilation theorem.

Amit Behera and Or Sattath

Ben-Gurion University of the Negev, Israel | Ben-Gurion University of the Negev, Israel

Almost Public Quantum Coins

In a quantum money scheme, a bank can issue money that users cannot counterfeit. Similar to bills of paper money, most quantum money schemes assign a unique serial number to each money state, thus potentially compromising the privacy of the users of quantum money. However in a quantum coins scheme, just like the traditional currency coin scheme, all the money states are exact copies of each other, providing a better level of privacy for the users. A quantum money scheme can be private, i.e., only the bank can verify the money states, or public, meaning anyone can verify. In this work, we propose a way to lift any private quantum coin scheme -- which is known to exist based on the existence of one-way functions by Ji, Liu, and Song (CRYPTO'18) -- to a scheme that closely resembles a public quantum coin scheme. Verification of a new coin is done by comparing it to the coins the user already possesses, by using a projector on to the symmetric subspace. No public coin scheme was known prior to this work. It is also the first construction that is close to a public quantum money scheme and is provably secure based on standard assumptions. The lifting technique when instantiated with the private quantum coins scheme by Mosca and Stebila (2010) gives rise to the first construction that is close to an inefficient unconditionally secure public quantum money scheme.

Pierre Fraigniaud, Francois Le Gall, Harumichi Nishimura and Ami Paz

CNRS and Universite de Paris | Nagoya University | Nagoya University | Universitat Wien

Distributed Quantum Proofs for Replicated Data

This paper tackles the issue of checking that all copies of a large data set replicated at several nodes of a network are identical. The fact that the replicas may be located at distant nodes prevents the system from verifying their equality locally, i.e., by having each node consult only nodes in its vicinity. On the other hand, it remains possible to assign certificates to the nodes, so that verifying the consistency of the replicas can be achieved locally. However, we show that, as the replicated data is large, classical certification mechanisms, including distributed Merlin-Arthur protocols, cannot guarantee good completeness and soundness simultaneously, unless they use very large certificates. The main result of this paper is a distributed quantum Merlin-Arthur protocol enabling the nodes to collectively check the consistency of the replicas, based on small certificates, and in a single round of message exchange between neighbors, with short messages. In particular, the certificate-size is logarithmic in the size of the data set, which gives an exponential advantage over classical certification mechanisms. We propose yet another usage of a fundamental quantum primitive, called the SWAP test, in order to show our main result.

Alex Bredariol Grilo, Kathrin Hövelmanns, Andreas Hülsing and Christian Majenz

LIP6, CNRS/Sorbonne Université | Ruhr-Universität Bochum | Eindhoven University of Technology | CWI and QuSoft

Tight adaptive reprogramming in the Quantum Random Oracle Model

The random oracle model (ROM) enjoys widespread popularity, mostly because it tends to allow for tight and conceptually simple proofs where provable security in the standard model is elusive or costly. While being the adequate replacement of the ROM in the post-quantum security setting, the quantum-accessible random oracle model (QROM) has thus far failed to provide these advantages in many settings. In this work, we focus on adaptive reprogrammability, a feature of the ROM enabling tight and simple proofs in many settings. We show that the straightforward quantum-accessible generalization of adaptive reprogramming is feasible by proving a bound on the adversarial advantage in distinguishing whether a random oracle has been reprogrammed or not. We show that our bound is tight by providing a matching attack. We go on to demonstrate that our technique recovers the mentioned advantages of the ROM in three QROM applications: 1) We give a tighter proof of security of the message compression routine as used by XMSS. 2) We show that the standard ROM proof of chosen-message security for Fiat-Shamir signatures can be lifted to the QROM, straightforwardly, achieving a tighter reduction than previously known. 3) We give the first QROM proof of security against fault injection and nonce attacks for the hedged Fiat-Shamir transform.

Arne Heimendahl, Markus Heinrich and David Gross

University of Cologne | Universität Köln | Universität Köln

The axiomatic and the operational approach to resource theories of magic do not coincide

Stabiliser operations occupy a prominent role in the theory of fault-tolerant quantum computing. They are defined operationally: by the use of Clifford gates, Pauli measurements and classical control. Within the stabiliser formalism, these operations can be efficiently simulated on a classical computer, a result which is known as the Gottesman-Knill theorem. However, an additional supply of magic states is enough to promote them to a universal, fault-tolerant model for quantum computing. To quantify the needed resources in terms of magic states, a resource theory of magic has been developed during the last years. Stabiliser operations (SO) are considered free within this theory, however they are not the most general class of free operations. From an axiomatic point of view, these are the completely stabiliser-preserving (CSP) channels, defined as those that preserve the convex hull of stabiliser states. It has been an open problem to decide whether these two definitions lead to the same class of operations. In this work, we answer this question in the negative, by constructing an explicit counter-example. This indicates that recently proposed stabiliser-based simulation techniques of CSP maps might be strictly more powerful than Gottesman-Knill-like methods. The result is analogous to a well-known fact in entanglement theory, namely that there is a gap between the class of local operations and classical communication (LOCC) and the class of separable channels. Along the way, we develop a number of auxiliary techniques which allow us to better characterise the set of CSP channels.

Guillaume Aubrun, Ludovico Lami, Carlos Palazuelos and Martin Plávala

Université Lyon 1 | Ulm University | Departamento de Análisis Matemático y Matemática Aplicada, Universidad Complutense de Madrid, E-28040 Madrid, Spain | University of Siegen

Entangleability of cones

We prove that two non-classical general probabilistic theories must give rise to entanglement, either at the level of states or at the level of measurements, when combined. This reveals a deep connection between a local phenomenon (non-classicality, or the existence of superpositions) and a global one (entanglement), and raises the latter to a generically non-classical rather than merely quantum phenomenon, in a precise mathematical sense. Instrumental in our proof is the solution of a long-standing conjecture by Barker.

Benjamin Yadin, Benjamin Morris and Gerardo Adesso

University of Oxford | University of Nottingham | University of Nottingham

Mixing indistinguishable systems leads to a quantum Gibbs paradox

The classical Gibbs paradox concerns the entropy change upon mixing two gases. Whether an observer assigns an entropy increase to the process depends on their ability to distinguish the gases. A resolution is that an `ignorant' observer, who cannot distinguish the gases, has no way of extracting work by mixing them. Moving the thought experiment into the quantum realm, we reveal new and surprising behaviour: the ignorant observer can extract work from mixing different gases, even if the gases cannot be directly distinguished. Moreover, in the macroscopic limit, the quantum case diverges from the classical ideal gas: as much work can be extracted as if the gases were fully distinguishable. We show that the ignorant observer assigns more microstates to the system than found by naive counting in semiclassical statistical mechanics. This demonstrates the importance of accounting for the level of knowledge of an observer, and its implications for genuinely quantum modifications to thermodynamics.

Hsin-Yuan Huang, Richard Kueng, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, Jarrod McClean and John Preskill

Caltech | Johannes Kepler University Linz | Google Research | Google Research | Google | Google Research | Google | Google | Caltech

Fundamental aspects of solving quantum problems with machine learning

Machine learning (ML) provides the potential to solve challenging quantum many-body problems in physics and chemistry. Yet, this prospect has not been fully justified. In this work, we establish rigorous results to understand the power of classical ML and the potential for quantum advantage in an important example application: predicting outcomes of quantum mechanical processes. We prove that for achieving a small average prediction error, one can always design a classical ML model whose sample complexity is comparable to the best quantum ML model (up to a small polynomial factor). Regarding computational complexity, we show that the class of problems that can be solved by efficient classical ML models with access to sampled data is strictly larger than BPP. Hence, classical ML models may be able to solve some challenging quantum problems after training from data obtained in physical experiments. As a concrete example, we prove that a simple, classical ML model can efficiently learn to predict ground state representations that approximate expectation values of local observables up to a small, constant error. This holds for any smooth family of gapped local Hamiltonians in a finite spatial dimension.

Shalev Ben-David, Andrew M. Childs, András Gilyén, William Kretschmer, Supartha Podder and Daochen Wang

University of Waterloo | University of Maryland | California Institute of Technology | University of Texas at Austin | University of Ottawa | University of Maryland

Symmetries, graph properties, and quantum speedups

Aaronson and Ambainis (2009) and Chailloux (2018) showed that fully symmetric (partial) functions do not admit exponential quantum query speedups. This raises a natural question: how symmetric must a function be before it cannot exhibit a large quantum speedup? In this work, we prove that hypergraph symmetries in the adjacency matrix model allow at most a polynomial separation between randomized and quantum query complexities. We also show that, remarkably, permutation groups constructed out of these symmetries are essentially the only permutation groups that prevent super-polynomial quantum speedups. We prove this by fully characterizing the primitive permutation groups that allow super-polynomial quantum speedups. In contrast, in the adjacency list model for bounded-degree graphs—where graph symmetry is manifested differently—we exhibit a property testing problem that shows an exponential quantum speedup. These results resolve open questions posed by Ambainis, Childs, and Liu (2010) and Montanaro and de Wolf (2013).

Simon Apers, Troy Lee and Ronald de Wolf

ULB - CWI | University of Technology, Sydney | QuSoft, CWI and University of Amsterdam (Amsterdam)

Quantum speedups for graph sparsification, graph cut problems and Laplacian solving

Graph sparsification underlies a large number of algorithms for graph cut problems and solving Laplacian systems. In its strongest form, "spectral sparsification" reduces the number of edges to near-linear in the number of nodes, while approximately preserving the cut and spectral structure of the graph. We give a quantum algorithm that outputs a classical description of an 98ε-spectral sparsifier of a weighted graph with n vertices and m edges. It has time complexity $\tilde{O}(\sqrt{mn}/\varepsilon)$ in the adjacency array model and $\tilde{O}(n^{3/2}/\varepsilon)$ in the adjacency matrix model. These bounds are tight up to polylogarithmic factors, and improve on the optimal classical complexities $\Omega(m)$ and $\Omega(n^2)$, respectively. Using classical algorithms on the obtained sparsifier yields immediate quantum speedups for approximately solving Laplacian systems and for approximating a range of graph cut problems in essentially the same complexity. As a significantly more involved application we show how to speed up the quantum query complexity for computing exactly the edge connectivity of simple graphs. We show upper bounds on the query complexity of $\tilde{O}(\sqrt{mn})$ and $\tilde{O}(n^{3/2})$ in the adjacency matrix and adjacency array models, respectively. The upper bound for the adjacency matrix model is tight up to logarithmic factors, while for the adjacency array model the best quantum query lower bound we know is $\Omega(n)$.

Ashley Montanaro and Changpeng Shao // Troy Lee, Miklos Santha and Shengyu Zhang

PhaseCraft Ltd and University of Bristol | School of Mathematics, University of Bristol // University of Technology, Sydney | CNRS, IRIF, Université de Paris and CQT Singapore | Tencent Quantum Laboratory

Quantum algorithms for learning graphs // Quantum algorithms for graph problems with cut queries

We study the problem of learning an unknown graph provided via an oracle using a quantum algorithm. We consider three query models. In the first model ("OR queries"), the oracle returns whether a given subset of the vertices contains any edges. In the second ("parity queries"), the oracle returns the parity of the number of edges in a subset. In the third model, we are given copies of the graph state corresponding to the graph. We give quantum algorithms that achieve speedups over the best possible classical algorithms in the OR and parity query models, for some families of graphs, and give quantum algorithms in the graph state model whose complexity is similar to the parity query model. For some parameter regimes, the speedups can be exponential in the parity query model. On the other hand, without any promise on the graph, no speedup is possible in the OR query model. A main technique we use is the quantum algorithm for solving the combinatorial group testing problem, for which a query-efficient quantum algorithm was given by Belovs. Here we additionally give a time-efficient quantum algorithm for this problem, based on the algorithm of Ambainis et al. for a "gapped" version of the group testing problem. We also give simple time-efficient quantum algorithms based on Fourier sampling and amplitude amplification for learning the exact-half and majority functions, which almost match the optimal complexity of Belovs' algorithms. //

Let G be an n-vertex graph with m edges. When asked a subset S of vertices, a cut query on G returns the number of edges of G that have exactly one endpoint in S. We show that there is a bounded-error quantum algorithm that determines all connected components of G after making $O(\log(n)^6)$ many cut queries. In contrast, it follows from results in communication complexity that any randomized algorithm even just to decide whether the graph is connected or not must make at least $\Omega(n/\log(n))$ many cut queries. We further show that with $O(\log(n)^8)$ many cut queries a quantum algorithm can with high probability output a spanning forest for G. En route to proving these results, we design quantum algorithms for learning a graph using cut queries. We show that a quantum algorithm can learn a graph with maximum degree d after $O(d\log(n)^2)$ many cut queries, and can learn a general graph with $O(\sqrt{m}\log(n)^{3/2})$ many cut queries. These two upper bounds are tight up to the poly-logarithmic factors, and compare to $\Omega(dn)$ and $\Omega(m/\log(n))$ lower bounds on the number of cut queries needed by a randomized algorithm for the same problems, respectively. The key ingredients in our results are the Bernstein-Vazirani algorithm, approximate counting with "OR queries", and learning sparse vectors from inner products as in compressed sensing.

Ankit Garg, Robin Kothari, Praneeth Netrapalli and Suhail Sherif

Microsoft | Microsoft | Microsoft | Tata Institute of Fundamental Research

No quantum speedup over gradient descent for non-smooth convex optimization

We study the first-order convex optimization problem, where we have black-box access to a (not necessarily smooth) function $f: \mathbb{R}^n \to \mathbb{R}$ and its (sub)gradient. Our goal is to find an eps-approximate minimum of f starting from a point that is distance at most R from the true minimum. If f is G-Lipschitz, then the classic gradient descent algorithm solves this problem with $O((GR/\epsilon)^2)$ queries. Importantly, the number of queries is independent of the dimension n and gradient descent is optimal in this regard: No deterministic or randomized algorithm can achieve better complexity that is still independent of the dimension n. In this paper we reprove the matching randomized lower bound using a simpler argument than previous lower bounds. We then show that although the function family used in the lower bound is hard for randomized algorithms, it can be solved quadratically faster using quantum queries. We then show an improved lower bound against quantum algorithms using a different set of instances and establish our main result that in general even quantum algorithms need $\Omega((GR/\epsilon)^2)$ queries to solve the problem. Hence there is no quantum speedup over gradient descent for black-box first-order convex optimization without further assumptions on the function family. In our second result, we consider the case of smooth functions. Here the optimal classical algorithm is not gradient descent, but accelerated gradient descent, which is also known to be optimal among all classical (randomized) algorithms. We show a matching quantum lower bound, showing that there is no quantum speedup over accelerated gradient descent either.

Jintai Ding, Vlad Gheorghiu, András Gilyén, Sean Hallgren and Jianqiang Li

University of Cincinnati, OH, USA | softwareQ Inc. / Institute for Quantum Computing / Dept.\ of Combinatorics \& Optimization, University of Waterloo | California Institute of Technology | Department of Computer Science and Engineering, Pennsylvania State University | Department of Computer Science and Engineering, Pennsylvania State University

Limitations of the Macaulay matrix approach for using the HHL algorithm to solve multivariate polynomial systems

Recently Chen and Gao proposed a new quantum algorithm for Boolean polynomial system solving, motivated by the cryptanalysis of some post-quantum cryptosystems. The key idea of their approach is to apply a Quantum Linear System (QLS) algorithm to a Macaulay linear system over $\mathcal{C}$, which is derived from the Boolean polynomial system. The efficiency of their algorithm depends on the condition number of the Macaulay matrix. In this paper, we give a strong lower bound on the condition number as a function of the Hamming weight of the solution. We describe a Grover-based exhaustive search algorithm that always outperforms their algorithm. Then, we improve upon Chen and Gao's algorithm by introducing the Boolean Macaulay linear system over $\mathcal{C}$ by reducing the original Macaulay linear system. This improved algorithm could potentially significantly outperform the brute-force algorithm, when the Hamming weight of the solution is logarithmic in the number of variables. Furthermore, we provide a simple and more elementary proof of correctness for our improved algorithm using a reduction employing the Valiant-Vazirani affine hashing method, and also extend the result to polynomial systems over $\mathcal{F}_q$ improving on subsequent work by Chen, Gao and Yuan. We also suggest a new approach for extracting the solution of the Boolean polynomial system via a generalization of the quantum coupon collector problem.

Jay Gambetta

IBM Fellow and VP of Quantum Computing, IBM Quantum

Challenges and Directions of Quantum Information Technology

In the past decade, the quantum computing community has expanded from a small, research-focused group of physicists, engineers and mathematicians to a large and interdisciplinary field that includes experts from all domains including industry, government, and academia. As a result, we have seen accelerated progress toward understanding the scope of quantum computing, pushing its hardware technology, developing applications, and advancing error correction protocols. In this talk, I would like to share my five-year vision for advancing quantum information technology, including how we will push the limits and what the key challenges to realizing frictionless quantum computing are—that is, quantum computing integrated seamlessly with high performance computing resources.

John Napp, Rolando La Placa, Alexander Dalzell, Fernando Brandão and Aram Harrow

Massachusetts Institute of Technology | MIT | California Institute of Technology | California Institute of Technology and Amazon Web Services | Massachusetts Institute of Technology

Efficient classical simulation of random shallow 2D quantum circuits

Random quantum circuits are commonly viewed as hard to simulate classically. In some regimes this has been formally conjectured, and there had been no evidence against the more general possibility that for circuits with uniformly random gates, approximate simulation of typical instances is almost as hard as exact simulation. We prove that this is not the case by exhibiting a shallow circuit family with uniformly random gates that cannot be efficiently classically simulated near-exactly under standard hardness assumptions, but can be simulated approximately for all but a superpolynomially small fraction of circuit instances in time linear in the number of qubits and gates. We furthermore conjecture that sufficiently shallow random circuits are efficiently simulable more generally. To this end, we propose and analyze two simulation algorithms. Implementing one of our algorithms numerically, we give strong evidence that it is efficient both asymptotically and, in some cases, in practice. To argue analytically for efficiency, we reduce the simulation of 2D shallow random circuits to the simulation of a form of 1D dynamics consisting of alternating rounds of random local unitaries and weak measurements -- a type of process that has generally been observed to undergo a phase transition from an efficient-to-simulate regime to an inefficient-to-simulate regime as measurement strength is varied. Using a mapping from quantum circuits to statistical mechanical models, we give evidence that a similar computational phase transition occurs for our algorithms as parameters of the circuit architecture like the local Hilbert space dimension and circuit depth are varied.

Matthew B. Hastings, Umesh Vazirani and András Gilyén

Station Q, Microsoft Research, and Microsoft Quantum | Berkeley Quantum Information & Computation Center, UC Berkeley | California Institute of Technology

(Sub)Exponential advantage of adiabatic quantum computation with no sign problem

We demonstrate the possibility of (sub)exponential quantum speedup via a quantum algorithm that follows an adiabatic path of a gapped sparse Hamiltonian with no sign problem. This is in sharp contrast with frustration-free stoquastic Hamiltonians, where no such speedup is possible as shown by Bravyi and Terhal (2008). The Hamiltonian that exhibits this speed-up comes from the adjacency matrix of an undirected graph, and we can view the adiabatic evolution as an efficient $\mathcal{O}(\mathrm{poly}(n))$-time quantum algorithm for finding a specific "EXIT" vertex in the graph given the "ENTRANCE" vertex. On the other hand we show that if the graph is given via an adjacency-list oracle, there is no classical algorithm that finds the "EXIT" with probability greater than $e^{\frac{-n}{\delta}}$ using at most $e^{\frac{n}{\delta}}$ queries for $\delta = \frac{1}{5} - o(1)$. Our construction of the graph is somewhat similar to the ``welded-trees'' construction of Childs et al. (2003), but uses additional ideas for simultaneously achieving a spectral gap and a short adiabatic path.

Joel Klassen and Charles Derby

Phasecraft Ltd. | Phasecraft Ltd. and University College London

A Compact Fermion to Qubit Mapping

Mappings between fermions and qubits are valuable constructions in physics. To date only a handful exist. In addition to revealing dualities between fermionic and spin systems, such mappings are indispensable in any simulation of fermionic physics on quantum computers. The number of qubits required per fermionic mode, and the locality of mapped fermionic operators strongly impact the cost of such simulations. Local fermionic encodings are a class of fermion to qubit mappings that map local fermionic operators to local qubit operators. We present a novel local fermionic encoding -- called the compact encoding -- which outperforms all previous local fermionic encodings in both the qubit to mode ratio, and the locality of mapped operators. We demonstrate how this encoding may be applied to any uniform tiling of degree 4 or less, for example square and hexagonal lattices, and to a 3d cubic lattice. In order to characterize the encoding on various lattices we clarify the group theoretic structure underlying the design of such encodings. We also illuminate an elegant relationship between the compact encoding and the toric code, and show how the compact encoding may be understood as condensing the fermionic excitations of the toric code into its code space.

Adrian Chapman and Steven Flammia

The University of Sydney | AWS

Characterization of solvable spin models via graph invariants

Exactly solvable models are essential in physics. For many-body spin-$1/2$ systems, an important class of such models consists of those that can be mapped to free fermions hopping on a graph. We provide a complete characterization of models which can be solved this way. Specifically, we reduce the problem of recognizing such spin models to the graph-theoretic problem of recognizing line graphs, which has been solved optimally. A corollary of our result is a complete set of constant-sized commutation structures that constitute the obstructions to a free-fermion solution. We find that symmetries are tightly constrained in these models. Pauli symmetries correspond to either: (i) cycles on the fermion hopping graph, (ii) the fermion parity operator, or (iii) logically encoded qubits. Clifford symmetries within one of these symmetry sectors, with three exceptions, must be symmetries of the free-fermion model itself. We demonstrate how several exact free-fermion solutions from the literature fit into our formalism and give an explicit example of a new model previously unknown to be solvable by free fermions.

Adam Bene Watts and J. William Helton

MIT | UCSD

3XOR Games with Perfect Commuting Operator Strategies Have Perfect Tensor Product Strategies and are Decidable in Polynomial Time

We consider 3XOR games with perfect commuting operator strategies. Given any 3XOR game, we show existence of a perfect commuting operator strategy for the game can be decided in polynomial time. Previously this problem was not known to be decidable. Our proof leads to a construction, showing a 3XOR game has a perfect commuting operator strategy iff it has a perfect tensor product strategy using a 3 qubit (8 dimensional) GHZ state. This shows that for perfect 3XOR games the advantage of a quantum strategy over a classical strategy (defined by the quantum-classical bias ratio) is bounded. This is in contrast to the general 3XOR case where the optimal quantum strategies can require high dimensional states and there is no bound on the quantum advantage. To prove these results, we first show equivalence between deciding the value of an XOR game and solving an instance of the subgroup membership problem on a class of right angled Coxeter groups. We then show, in a proof that consumes most of this paper, that the instances of this problem corresponding to 3XOR games can be solved in polynomial time.

Yuan Su, Hsin-Yuan Huang and Earl Campbell

QuICS/University of Maryland | Caltech | AWS Center for Quantum Computing

Nearly tight Trotterization of interacting electrons

We consider simulating quantum systems on digital quantum computers. We show that the performance of quantum simulation can be improved by simultaneously exploiting the commutativity of Hamiltonian, the sparsity of interactions, and the prior knowledge of initial state. We achieve this using Trotterization for a class of correlated electrons that encompasses various physical systems, including the plane-wave-basis electronic structure and the Fermi-Hubbard model. We estimate the simulation error by taking the transition amplitude of nested commutators of Hamiltonian terms within the η-electron manifold. We develop multiple techniques for bounding the transition amplitude and the expectation of general fermionic operators, which may be of independent interest. We show that it suffices to use $\mathcal{O}(\frac{n^{\frac{5}{3}}}{\eta^{\frac{2}{3}}} + n^{4/3}\eta^{2/3})$ gates to simulate electronic structure in the plane-wave basis with n spin orbitals and η electrons up to a negligible factor, improving the best previous result in second quantization while outperforming the first-quantized simulation when $\eta = \mathcal{O}(\eta^2)$. We also obtain an improvement for simulating the Fermi-Hubbard model. We construct concrete examples for which our bounds are almost saturated, giving a nearly tight Trotterization of correlated electrons.

Leo Zhou and Dorit Aharonov

Harvard University | The Hebrew University of Jerusalem

Strongly Universal Hamiltonian Simulators

A universal family of Hamiltonians can be used to simulate any local Hamiltonian by encoding its full spectrum as the low-energy subspace of a Hamiltonian from the family. Many spin-lattice model Hamiltonians---such as Heisenberg or XY interaction on the 2D square lattice---are known to be universal. However, the known encodings can be very inefficient, requiring interaction strengths that scales exponentially with system size if the original Hamiltonian have complex, possibly all-to-all connectivity. In this work, we provide an efficient construction by which these universal families are in fact ``strongly'' universal; this means that the required interaction strengths as well as all other resources scale polynomially, regardless of the connectivity of the original Hamiltonian. This exponential improvement over previous constructions based on perturbative gadgets is achieved by combining the tools of quantum phase-estimation algorithm and circuit-to-Hamiltonian transformation in a non-perturbative way that only incurs polynomial overhead. Furthermore, we show that 1D Hamiltonians with nearest-neighbor interaction of 8-dimensional particles on a line are also strongly universal Hamiltonian simulators. Our results demonstrate that analog quantum simulation of general Hamiltonians can be made efficient for all target local Hamiltonians; this has potential application for future quantum technologies.

Lin Lin and Yu Tong

University of California, Berkeley | University of California, Berkeley

Near-optimal ground state preparation

Preparing the ground state of a given Hamiltonian and estimating its ground energy are important but computationally hard tasks. However, given some additional information, these problems can be solved efficiently on a quantum computer. We assume that an initial state with non-trivial overlap with the ground state can be efficiently prepared, and the spectral gap between the ground energy and the first excited energy is bounded from below. With these assumptions we design an algorithm that prepares the ground state when an upper bound of the ground energy is known, whose runtime has a logarithmic dependence on the inverse error. When such an upper bound is not known, we propose a hybrid quantum-classical algorithm to estimate the ground energy, where the dependence of the number of queries to the initial state on the desired precision is exponentially improved compared to the current state-of-the-art algorithm proposed in [Ge et al. 2019]. These two algorithms can then be combined to prepare a ground state without knowing an upper bound of the ground energy. We also prove that our algorithms reach the complexity lower bounds by applying it to the unstructured search problem and the quantum approximate counting problem.

Dmitri Maslov, Jin-Sung Kim, Sergey Bravyi, Theodore J. Yoder and Sarah Sheldon

IBM Thomas J. Watson Research Center | IBM Almaden Research Center | IBM Thomas J. Watson Research Center | IBM Thomas J. Watson Research Center | IBM Almaden Research Center

Quantum advantage for computations with limited space

Quantum computations promise the ability to solve problems intractable in the classical setting. Restricting the types of computations considered often allows to establish a provable theoretical advantage by quantum computations, and later demonstrate it experimentally. In this paper, we consider space-restricted computations, where input is a read-only memory and only one (qu)bit can be computed on. We show that n-bit symmetric Boolean functions can be implemented exactly through the use of quantum signal processing as restricted space quantum computations using $O(n^2)$ gates, but some of them may only be evaluated with probability $\frac{1}{2} + O(n/2^{n/2})$ by analogously defined classical computations. We experimentally demonstrate computations of 3-, 4-, 5-, and 6-bit symmetric Boolean functions by quantum circuits, leveraging custom two-qubit gates, with algorithmic success probability exceeding the best possible classically. This establishes and experimentally verifies a different kind of quantum advantage---one where quantum scrap space is more valuable than analogous classical space---and calls for an in-depth exploration of space-time tradeoffs in quantum circuits.

Nathan Ju, Daniel Grier and Luke Schaeffer

University of Illinois, Urbana-Champaign | University of Waterloo | University of Waterloo

Interactive quantum advantage with noisy, shallow Clifford circuits

Recent work by Bravyi et al. constructs a relation problem that a noisy constant-depth quantum circuit (QNC^0) can solve with near certainty (probability $1 - o(1)$), but that any bounded fan-in constant-depth classical circuit (NC^0) fails with some constant probability. We show that this robustness to noise can be achieved in the other low-depth quantum/classical circuit separations in this area. In particular, we show a general strategy for adding noise tolerance to the interactive protocols of Grier and Schaeffer. As a consequence, we obtain an unconditional separation between noisy QNC^0 circuits and $AC^0[p]$ circuits for all primes $p \geq 2$, and a conditional separation between noisy QNC^0 circuits and log-space classical machines under a plausible complexity-theoretic conjecture. A key component of this reduction is showing average-case hardness for the classical simulation tasks---that is, showing that a classical simulation of the quantum interactive task is still powerful even if it is allowed to err on some constant fraction of inputs. We show that is possible even for quantum tasks which are parity-L-hard to simulate. To do this, we borrow techniques from randomized encodings used in cryptography.

David Gosset, Daniel Grier, Alex Kerzner and Luke Schaeffer

University of Waterloo | University of Waterloo | University of Waterloo | University of Waterloo

Fast simulation of planar Clifford circuits

A general quantum circuit can be simulated in exponential time on a classical computer. If it has a planar layout, then a tensor-network contraction algorithm due to Markov and Shi has a runtime exponential in the square root of its size, or more generally exponential in the treewidth of the underlying graph. Separately, Gottesman and Knill showed that if all gates are restricted to be Clifford, then there is a polynomial time simulation. We combine these two ideas and show that treewidth and planarity can be exploited to improve Clifford circuit simulation. Our main result is a classical algorithm with runtime scaling asymptotically as $n^{\frac{\omega}{2}} < n^{1.19}$ which samples from the output distribution obtained by measuring all n qubits of a planar graph state in given Pauli bases. Here ω is the matrix multiplication exponent. We also provide a classical algorithm with the same asymptotic runtime which samples from the output distribution of any constant-depth Clifford circuit in a planar geometry. Our work improves known classical algorithms with cubic runtime. A key ingredient is a mapping which, given a tree decomposition of some graph G, produces a Clifford circuit with a structure that mirrors the tree decomposition and which emulates measurement of the quantum graph state corresponding to G. We provide a classical simulation of this circuit with the runtime stated above for planar graphs and otherwise $nt^{\omega-1}$ where t is the width of the tree decomposition. Our algorithm incorporates two subroutines which may be of independent interest. The first is a matrix-multiplication-time version of the Gottesman-Knill simulation of multi-qubit measurement on stabilizer states. The second is a new classical algorithm for solving symmetric linear systems over $\mathbb{F}_2$ in a planar geometry.