QuICS Workshop on Quantum Information and Computer Science

Center for Quantum Information and Computer Science

College Park Marriott Hotel and Conference Center
3501 University Blvd, East Hyattsville Maryland 20783

All speaker sessions will be held in room 2110/2111/2112.
The reception and poster session will be held in room 2100/2101/2102.

Monday, March 31

8:30	AM	Breakfast
9:00	AM	Opening Remarks
9:05	AM	Session 1: Tutorial and Quantum Algorithms I - Seth Lloyd, Andrew Childs
10:25	AM	Morning Break
10:45	AM	Session 2: Quantum Algorithms II - Sean Hallgren, Seth Lloyd
12:05	PM	Lunch
1:20	PM	Panel Discussion
2:00	PM	Session 3: Quantum Computation - Barbara Terhal, Daniel Gottesman
3:20	PM	Afternoon Break
3:45	PM	Session 4: Black Holes - Scott Aaronson, Samuel Braunstein
5:05	PM	Poster Session & Reception

Tuesday, April 1

8:30	AM	Breakfast
9:00	AM	Session 5: Quantum Information Theory - Graeme Smith, Debbie Leung
10:20	AM	Morning Break
10:45	AM	Session 6: Hamiltonian Complexity - Mario Szegedy, Fernando Brandão
12:05	PM	Lunch
1:10	PM	Session 7: Quantum Cryptography - Renato Renner, Nicolas Gisin
2:30	PM	Afternoon Break
2:55	PM	Session 8: Experimental Tests of Quantum Mechanics - Manny Knill
3:35	PM	End*

***Note:** Scott Aaronson will be giving the Physics Colloquium at the University of Maryland on "Quantum Computing and the Limits of the Efficiently Computable" at 4:00pm on Tuesday April, 1 in the John S. Toll Physics Building, Room 1412.

Session 1: Tutorial and Quantum Algorithms I

Seth Lloyd, MIT
Monday, 9:05AM.

Quantum information and computer science tutorial

Andrew Childs, University of Waterloo
Monday, 9:45AM.

Exponential improvement in precision for simulating sparse Hamiltonians

We provide a quantum algorithm for simulating the dynamics of sparse Hamiltonians with complexity sublogarithmic in the inverse error, an exponential improvement over previous methods. Our algorithm is based on a significantly improved simulation of the fractional-query model using discrete quantum queries, showing that fractional queries are not much more powerful than discrete ones even for very small error. We significantly simplify the analysis of this conversion, avoiding the need for a complex fault correction procedure. Our simplification relies on a new form of "oblivious amplitude amplification" that can be applied even though the reflection about the input state is unavailable. We also prove new lower bounds showing that our algorithms are optimal as a function of the error.

Based on joint work with Dominic Berry, Richard Cleve, Robin Kothari, and Rolando Somma.

Session 2: Quantum Algorithms II

Sean Hallgren, Pennsylvania State University
Monday, 10:45AM.

A quantum algorithm for computing the unit group of an arbitrary degree number field

Computing the group of units in a field of algebraic numbers is one of the central tasks of computational algebraic number theory. It is believed to be hard classically, which is of interest for cryptography. In the quantum setting, efficient algorithms were previously known for fields of constant degree. We give a quantum algorithm that is polynomial in the degree of the field and the logarithm of its discriminant. This is achieved by combining three new results. The first is a classical algorithm for computing a basis for certain ideal lattices with doubly exponentially large generators. The second shows that a Gaussian weighted superposition of lattice points, with an appropriate encoding, can be used to provide a unique representation of a real-valued lattice. The third is an extension of the hidden subgroup problem to continuous groups and a quantum algorithm for solving the HSP over $\mathbb{R}^n$.

Joint work with Kirsten Eisentraeger, Alexei Kitaev, and Fang Song.

Seth Lloyd, MIT
Monday, 11:25 AM.

Big quantum data

Machine learning techniques for handling big data frequently rely on linear manipulations of large vectors. Quantum computers are good at linear manipulations of large vectors. This talk presents a suite of quantum algorithms for big data analysis. The algorithms are exponentially faster than their classical counterparts.

Panel Discussion

Monday, 1:20 PM.

A panel discussion will be given by Sam Braunstein, Andrew Childs, Nicolas Gisin, Seth Lloyd, and Graeme Smith at 1:20 PM on Monday. The discussion may cover a number of topics, such as research directions that the QuICS center could pursue, activities that the center could organize, and research topics that we as a community would like to see funded. Jacob Taylor will moderate.

Session 3: Quantum Computation

Barbara Terhal, RWTH Aachen
Monday, 2:00 PM.

Space-time circuit-to-Hamiltonian construction and its applications

The circuit-to-Hamiltonian construction translates dynamics (a quantum circuit and its output) into statics (the groundstate of a circuit Hamiltonian) by explicitly defining a quantum register for a clock. The standard Feynman-Kitaev construction uses one global clock for all qubits while we consider a different construction in which a clock is assigned to each interacting qubit. This makes it possible to capture the spatio-temporal structure of the original quantum circuit into features of the circuit Hamiltonian. The construction is inspired by the original two-dimensional interacting fermionic model (see `http://journals.aps.org/pra/abstract/10.1103/PhysRevA.63.040302`). We prove that for one-dimensional quantum circuits the gap of the circuit Hamiltonian is appropriately lower-bounded, partially using results on mixing times of Markov chains, so that the applications of this construction for QMA (and partially for quantum adiabatic computation) go through. For one-dimensional quantum circuits, the dynamics generated by the circuit Hamiltonian corresponds to diffusion of a string around the torus.

See the paper at `http://arxiv.org/abs/1311.6101`.

Daniel Gottesman, Perimeter Institute
Monday, 2:40 PM.

Fault-tolerant quantum computation with constant overhead

The threshold theorem for fault tolerance tells us that it is possible to build arbitrarily large reliable quantum computers provided the error rate per physical gate or time step is below some threshold value. Most research on the threshold theorem so far has gone into optimizing the tolerable error rate under various assumptions, with other considerations being secondary. However, for the foreseeable future, the number of qubits may be an even greater restriction than error rates. The overhead, the ratio of physical qubits to logical qubits, determines how expensive (in qubits) a fault-tolerant computation is. Earlier results on fault tolerance used a large overhead which grows (albeit slowly) with the size of the computation. I show that it is possible in principle to do fault-tolerant quantum computation with low overhead, and with the overhead constant in the size of the computation. The result depends on recent progress on quantum low-density parity check codes.

Session 4: Black Holes

Scott Aaronson, MIT

The cryptographic hardness of decoding Hawking radiation
Monday, 3:45 PM.

The "firewall paradox" of Almheiri et al. is an argument that several reasonable-seeming assumptions about black-hole physics lead to an absurdity when combined. In a striking paper last year, Harlow and Hayden proposed that part of the resolution might be that, in order to "decode" a black hole's Hawking radiation in a way that leads to the paradox, the infalling observer Alice would need to perform a quantum computation that takes far longer than the life of the black hole. In particular, they proved that a natural abstraction of Alice's decoding task is hard for the complexity class QSZK (Quantum Statistical Zero Knowledge). We give a different, arguably-stronger hardness argument for Alice's decoding task, based on a more "standard" cryptographic assumption. Namely we show that it is at least as hard as inverting a one-way function.

Samuel Braunstein, University of York
Monday, 4:25 PM.

What quantum information tells us about black holes

Session 5: Quantum Information Theory

Graeme Smith, IBM Research
Tuesday, 9:00 AM.

Bound entangled states with a private key and their classical counterpart

Entanglement is a fundamental resource for quantum information processing. In its pure form, it allows quantum teleportation and sharing classical secrets. Realistic quantum states are noisy and their usefulness is only partially understood. Bound-entangled states are central to this question–they have no distillable entanglement, yet sometimes still have a private classical key. We present a construction of bound-entangled states with a private key based on classical probability distributions. From this emerge states possessing a new classical analogue of bound entanglement, distinct from the long-sought bound information. We also find states of smaller dimensions and higher key rates than previously known. Our construction has implications for classical cryptography: we show that existing protocols are insufficient for extracting private key from our distributions due to their "bound-entangled" nature. We propose a simple extension of existing protocols that can extract a key from them.

Debbie Leung, University of Waterloo
Tuesday, 9:40 AM.

Maximal privacy without coherence

A coherently transmitted quantum state is inherently private. Remarkably, coherent quantum communication is not a prerequisite for privacy: there are quantum channels that are too noisy to transmit any quantum information reliably that can nevertheless send private classical information. Here, we ask how much private classical information a channel can transmit if it has little quantum capacity. We present a class of channels N_d with input dimension d^2, quantum capacity $Q(N_d) \leq 1$, and private capacity $P(N_d) = \log d$. These channels asymptotically saturate an interesting inequality $P(N) \leq (\log d_A + Q(N))/2$ for any channel N with input dimension d_A, and capture the essence of privacy stripped of the confounding influence of coherence.

Joint work with Ke Li, Graeme Smith, and John Smolin

Session 6: Hamiltonian Complexity

Mario Szegedy, Rutgers University
Tuesday, 10:45 AM.

The area law of quantum physics does generalize to cut-law

We show that a proposed natural generalization of the conjectured area law of quantum physics is false. The area law states that if we have a domain in a d-dimensional grid representing a spin-system, then the entanglement entropy between the the the sub-systems (1) in the domain and the one that is (2) outside of it, is upper bounded by the area of the surface. We prove that a generalization of this conjecture for arbitrary interaction graphs, where "surface area" is replaced with "cut size" does not hold.

Joint with Zeph Landau, Daniel Nagaj and Umesh Vazirani.

Fernando Brandão, University College London
Tuesday, 11:25 AM.

Limitations for quantum PCPs

An interesting current open problem in quantum complexity theory is the quantum PCP conjecture. In analogy with the PCP theorem, the conjecture states that it is "quantum NP"-hard to tell whether a quantum constraint satisfaction problem (aka a local quantum Hamiltonian) is satisfiable or far from satisfiable, with a constant fraction of the constraints being violated in any assignment.

In this talk I will discuss limitations for quantum PCPs due to one of the most distinguishing features of quantum entanglement: its monogamous character. The monogamy of entanglement is the principle that the more entangled a system is with another one, the less entangled it can be with anything else. I will show how a quantitative understanding of entanglement monogamy leads both to limitations on the parameters that a potential quantum analogue of the PCP theorem might have, and to potential approaches to proving such an analogue (for instance by attempting to quantize the main steps of Dinur's proof of the PCP theorem).

The talk will be based on joint work with Aram Harrow (`arXiv:1310.0017`).

Session 7: Quantum Cryptography

Renato Renner, ETH Zurich
Tuesday, 1:10 PM.

Randomness amplification

Suppose you are in charge of drawing lottery numbers, and you are given a mechanical lottery machine. Clearly, you want the lottery numbers to be completely unpredictable. However, the lottery machine could have been in use since a long time, and it may be that someone (unknown to you) analysed it and found that certain numbers occur more likely than others. What can you do about it?

A famous result by Santha and Vazirani asserts that the quality of a source of randomness (such as our lottery machine) cannot be improved by any classical post-processing. However, the situation looks different in a quantum world. There exist entanglement-based "randomness amplification" protocols that can turn arbitrarily weak randomness into virtually perfect (and hence unpredictable) random bits. Crucially, these do not rely on the completeness of quantum theory and can therefore be used in the context of Bell-type experiments, or for device-independent information processing, for instance. In this talk, I will give an overview of the basics as well as recent research on randomness amplification.

Nicolas Gisin, University of Geneva
Tuesday, 1:50 PM.

Quantum communication: QKD, teleportation into a solid-state quantum memory, and large entanglement

- Commercial QKD system
- Why QKD?
- Longer distances: networks based on trusted nodes
- Quantum memories for quantum repeaters and networks
- Quantum hacking and decoy-detectors
- Large entanglement

Session 8: Experimental Tests of Quantum Mechanics

Emanuel Knill, NIST-Boulder
Tuesday, 2:55 PM.

Certifying violations of local realism

Many applications of quantum systems require measurements that verify the presence of sufficiently strong quantum correlations. The probability of the following unwanted event must be extremely small: The event where the correlations are not sufficiently strong but one is nevertheless convinced that they exist. Important examples of quantum correlation occur in experiments showing violations of Bell's inequalities, which are thought to invalidate local realism. This is a review of how such violations are quantified and robustly certified, with or without predetermined Bell's inequalities.

Posters

Joshua Bienfang, NIST

Bell-test beacon: verifiably random numbers for all to see

Randomness is an increasingly important resource in modern-day communications and information processing. As the diversity of applications of random-numbers has grown, so too have the requirements on the devices generating these numbers. In particular, applications such as cyber-security, gambling, and electronic commerce rely critically on the notion that the output of the random number generator is unknown and unpredictable. Actually, providing such assurance represents a profound and fundamental question for information theory. Indeed, while there are a wide variety of physically-based means to generate apparently random sequences, proving that the output of such a system could not have been known or predicted by a nefarious third-party operator is, in fact, impossible when the mechanism generating the bits is based solely on classical measurements. Proving (or disproving) correlations of this type is at the foundation of fundamental tests of the theory of quantum mechanics, known as Bell tests, that were developed specifically to address the notable objection that probabilities and uncertainty should not be physical properties. Loophole-free Bell tests are now recognized as a statement about whether or not the result of a measurement could have been predicted by any physically realistic means. This view represents a paradigm shift in how randomness can be verified. We present the design and applications of a public source of randomness that broadcasts strings of random bits generated by a Bell-test, NIST's random number beacon.

Sayan Choudhury, Cornell

Absence of the twisted superfluid state in a mean-field model of bosons on a honeycomb lattice

Motivated by recent observations [Soltan-Panahi et al., Nat. Phys. 8, 71 (2012)], we study the stability of a Bose-Einstein condensate within a spin-dependent honeycomb lattice towards forming a "twisted superfluid" state. Our exhaustive numerical search fails to find this phase, pointing to possible non-mean-field physics.

Daniela Frauchiger, Institute for Theoretical Physics, ETH Zurich, Switzerland

True randomness from realistic quantum devices

With Renato Renner, and Matthias Troyer.

True randomness is needed for applications such as cryptography, lottery and numerical simulations. As it turns out, many of the Random Number Generators (RNGs) used in practice are insufficient. Surprisingly, there seems to be no consensus on how to define randomness or measure its quality. What is even worse, is that the standard statistical tests used in practice, only analyse whether the output of an RNG is uniformly distributed, but they cannot verify that it is unpredictable. However, for many applications (such as the drawing of lottery numbers) the feature of unpredictability is crucial.

We propose a definition of true randomness that considers the generating process instead of the statistical properties of its output. This allows to capture the idea that it should be independent of all pre-existing information. Quantum systems are particularly interesting for random number generation, because their unpredictably can be proved based on physical principles.

In practical implementations of Quantum RNGs (QRNGs) the quantum system can never be perfectly isolated from the environment and is therefore always subject to noise. Because we cannot control the corresponding degrees of freedom, we cannot but sure either, that they are truly random. Here we propose a framework that allows to model any QRNG and determine the necessary amount of post-processing to turn its raw output into true randomness. The approach is complete in the sense that the bounds for the extractable randomness remain valid, if the model is replaced by another compatible model.

Michael Jarret, University of Maryland, College Park

Adiabatic optimization and the fundamental gap theorem

With Stephen Jordan.

By the adiabatic theorem, the runtime of an adiabatic optimization algorithm is upper bounded by $O(1/\gamma^3)$, where γ is the minimal eigenvalue gap. Thus, the analysis of the complexity of adiabatic algorithms reduces to the problem of bounding the eigenvalue gaps of the associated Hamiltonians. Despite some notable successes, this has historically proven to be a very difficult problem. In this work, rather than directly bounding the gaps of Hamiltonians for proposed adiabatic algorithms, we develop a collection of tools for proving lower bounds on eigenvalue gaps. These also may be of independent interest, such as in condensed-matter physics and spectral graph theory.

Dvir Kafri, Joint Quantum Institute, University of Maryland, College Park

A noise inequality for entanglement generation

Long-range interactions produce correlations between spatially separated systems, and thus are a means of communication. Here we present a method of confirming that a given observed interaction communicates quantum information. We show that for a Hamiltonian with linear couplings between canonical variables, disallowing quantum communication necessarily leads to excess noise in the systems conjugate variables. Analogously to Bell's inequality for quantum states, this gives a straightforward local test verifying the ability of an interaction to entangle systems, and we propose a simple quantum optics experiment demonstrating these ideas. We then present an underlying quantum circuit model for the generation of long-range, Markovian dynamics, which serves as a conceptual aid and motivates the proof of our claim.

Bill Kaminsky

A diagram expansion for calculating the moments of the eigenenergy distribution in Hamiltonians for adiabatic quantum optimization algorithms

We present a simple way to calculate the moments of the eigenenergy distribution in 2-local qubit Hamiltonians. Its first result is that random 2-local qubit Hamiltonians have asymptotically Gaussian eigenenergy distributions in the limit of an infinite number of qubits. This result is in sharp contrast to the much-fatter-tailed semicircle distribution associated with the completely nonlocal Hamiltonians considered in standard Wigner-Dyson random matrix theory, and it explains why adiabatic optimization algorithms based on 2-local qubit Hamiltonians are not immediately doomed to failure by having to fit an exponentially growing number of eigenvalues in an at most quadratically growing spectral range. Indeed, assuming the usual $O(N)$ extensive scaling of the range between the ground state eigenenergy and the modal eigenenergy, the eigenenergies' standard deviation will be just $O[\sqrt(N)]$. This in turn implies that far down in the Gaussian tails there may be $O(1)$ gaps on average between the lowest O(1) eigenenergies of a random 2-local Hamiltonian.

We then show that this tantalizing scenario obtains not only in such fully random 2-local Hamiltonians, but also in ones explicitly pertinent for adiabatic quantum optimization algorithms. In particular, we show that it obtains when the standard transverse-field Ising model adiabatic algorithm treats almost any instance of Max Independent Set drawn from the set of Erdős-Rényi random graphs with a subquadratic number of edges. Such Erdős-Rényi random instances have resisted all classical algorithmic attempts to produce a better than 50% approximation of their Max Independent Sets in a time growing slower than a superpolynomial $N^{\log N}$. We explain how the Gaussian scenario implies there is a chance the standard transverse-field Ising model adiabatic algorithm can do much better in approximating them, and if so, it may even maintain its advantage in the face of the practical restriction of having to operate at a temperature that cannot shrink as the instance size grows.

Shelby Kimmel, MIT

The quantum query complexity of read-many formulas

With Andrew Childs and Robin Kothari.

The quantum query complexity of evaluating any read-once formula with n black-box input bits is $\Theta(n^{1/2})$. However, the corresponding problem for read-many formulas (i.e., formulas in which the inputs have fanout) is not well understood. Although the optimal read-once formula evaluation algorithm can be applied to any formula, it can be suboptimal if the inputs have large fanout. We give an algorithm for evaluating any formula with n inputs, size S, and G gates using $O(\min\{n, S^{1/2}, n^{1/2}G^{1/4}\})$ quantum queries. Furthermore, we show that this algorithm is optimal, since for any n, S, G there exists a formula with n inputs, size at most S, and at most G gates that requires $\Omega(\min\{n, S^{1/2}, n^{1/2}G^{1/4}\})$ queries. We also show that the algorithm remains nearly optimal for circuits of any particular depth $k = 3$, and we give a linear-size circuit of depth 2 that requires $\tilde{\Omega}(n^{5/9})$ queries. Applications of these results include a $\tilde{\Omega}(n^{19/18})$ lower bound for Boolean matrix product verification, a nearly tight characterization of the quantum query complexity of evaluating constant-depth circuits with bounded fanout, new formula gate count lower bounds for several functions including parity, and a construction of an AC0 circuit of linear size that can only be evaluated by a formula with $\Omega(n^{2-\epsilon})$ gates.

Paulina S. Kuo, NIST Information Technology Laboratory

Quantum communications research at NIST Information Technology Laboratory

With Oliver Slattery, Lijun Ma, Yong-Su Kim, Alan Mink, and Xiao Tang.

We describe experimental quantum communications research in the Information Technology Laboratory at NIST. We are interested in hybrid quantum networks, which utilize quantum frequency conversion to interface between sources, detectors, qubits, and quantum memories that operate at different wavelengths. We study building blocks of these networks such as photon pair sources using spontaneous parametric downconversion and high-efficiency optical frequency conversion for frequency translation and improved photon detection.

Alejandra Maldonado, University of Concepción

Quantum discord underlies the optimal scheme for modifying the overlap between two states

The processes for decreasing probabilistically and increasing deterministically the overlap between two quantum states must be assisted by an auxiliary quantum system. The probabilistic scheme is implemented through a unitary reduction process, while the deterministic one is accomplished only with a joint unitary transformation. In both schemes a unitary operation is responsible for coupling the two involved systems thus introducing correlation between them. In this work we study the quantum correlations required to achieve the probabilistic scheme with optimal probability, and we characterize the quantum correlation involved in the deterministic procedure. We find a that both schemes can be accomplished with quantum discord and without entanglement. In addition, we propose a physical implementation of the probabilistic scheme with twin photons generated in the process of spontaneous parametric down-conversion. The states are encoded in the polarization of a single photon and the auxiliary system becomes its propagation path; the second photon is used for heralded detection.

Omar Shehab, University of Maryland, Baltimore County

Locality of different clock Hamiltonians for unitary evolution of Deutsch's algorithm

In this poster we derive three different clock Hamiltonians of the unitary evolution of the circuit of Deutsch's algorithm following Feynman, Kitaev and McClean-Parkhill-Aspuru-Guzik schemes respectively. Then, we compare the spectrum and locality of the Hamiltonians.

Borzu Toloui, Haverford College and Harvard University

Space-optimal quantum algorithms for simulating many-fermion systems in quantum chemistry

With Peter J. Love.

Simulating the time evolution of molecules is of great value in quantum chemistry. Numerical methods based on classical computational techniques are common but the cost of classical simulation of quantum systems grows exponentially with the system size and the required precision. Quantum computers, on the other hand, will be capable of simulating the evolution of other quantum systems efficiently. Simulating fermionic systems in chemistry, in particular, is one of the important applications of the new quantum simulation methods. Quantum simulation of fermion dynamics, when combined with appropriate state preparation techniques using quantum phase-estimation algorithms, allows for efficient calculation of correlation functions, electronic energies and reaction rates, yielding an exponential advantage over known classical numerical methods.

Here, we propose a new class of quantum algorithms for simulating fermionic systems based on oracle calls. We use the full Configuration-Interaction (CI) matrix representation for the first time where we take advantage of the sparsity of the Hamiltonian in the CI-basis. In our scheme, the number of oracle calls scales linearly with the sparsity of the Hamiltonian and quadratically with the maximum number of orbitals. Moreover, our method is optimal in the number of qubits necessary for the simulation. It can be shown that no other scheme with fewer number of qubits can be devised for this class of problems.
