Mathematical Frontiers in Quantum Information Theory
(Org: **Frédéric Dupuis** (Université de Montréal), **Benjamin Lovitz** (Concordia University) and/et **Arthur Mehta** (University of Ottawa))

AVANTIKA AGARWAL, Institute for Quantum Computing, University of Waterloo
Quantum Capacity of Pauli Channels

The quantum capacity captures the value of a quantum channel for transmitting quantum information, establishing fundamental limits on quantum communication. In spite of its central role in quantum information, the quantum capacity of most channels is unknown. In our recent work, we report significant increases in the capacity thresholds of the depolarizing channel. This is the first improvement in 18 years, giving a bigger increase beyond the hashing bound than all previous improvements combined. Our starting point is the representation theoretic framework recently proposed by Bhalerao and Leditzky (2025) to compute coherent information for special permutation invariant states. We generalize their framework to the full symmetric subspace, which allow us to optimize coherent information over rank two states in that space. A representation theoretic calculation shows that exponentially many Kraus operators of the channel annihilate the symmetric space, corresponding to a massive decrease in environment entropy for states on the symmetric space compared to the maximally mixed state. This explains the enhanced coherent information as a manifestation of degeneracy for the resulting codes.

XAVIER COITEUX-ROY, University of Calgary
Unconditional quantum advantage at the Razborov–Smolensky frontier

Constant-depth quantum circuits can solve relation problems that their classical counterparts provably cannot solve. Until now, the strongest complexity-theoretic separation for QNC^0-circuits was against AC^0, the class of constant-depth, polynomial-size classical circuits composed of unbounded fan-in AND and OR gates, as well as NOT gates. The classical lower bounds in all known demonstrations of that separation rely on random restrictions and switching lemmas. Here we extend the separation to $\mathrm{AC}^0[p]$, the class obtained by augmenting AC^0 with unbounded fan-in MOD_p gates. For each prime p, we give a relation problem that QNC^0-circuits over a fixed finite gate set can solve with probability at least $\frac{99}{100}$, but that cannot be solved by $\mathrm{AC}^0[p]$-circuits with probability $\frac{2}{3}$. The validity of a proposed solution can nevertheless be checked in $\mathrm{AC}^0[p]$. The classical hardness is proven by deriving a contradiction with the fundamental lower bound by Razborov–Smolensky that MOD_6 cannot be approximated by $\mathrm{AC}^0[p]$-circuits for any prime p. Since no lower bounds comparable to those of Razborov–Smolensky are known for even slightly stronger models such as $\mathrm{AC}^0[6]$, extending this unconditional quantum advantage to such models will require a breakthrough in classical circuit complexity.

NADISH DE SILVA, Simon Fraser University
The generalised semi-Clifford conjecture is false

The Clifford hierarchy is a nested sequence of sets of quantum gates that can be fault-tolerantly performed using gate teleportation within standard quantum error correction schemes. The importance of these gates has motivated numerous studies of their structure. Zeng-Chen-Chuang conjectured in 2007 that all hierarchy gates are generalised semi-Clifford, i.e. take the form $C_1 \Pi D C_2$ for Clifford gates C_1, C_2, a permutation gate Π, and a diagonal gate D; Beigi-Shor proved in 2008 that this holds for all third-level gates.

We construct a five-qubit gate that is in the fifth level of the Clifford hierarchy but is not generalised semi-Clifford. Rather than simply present and verify our counterexample to the generalised semi-Clifford conjecture, we show how its form can be deduced. Our counterexample also demonstrates that the Clifford hierarchy is not closed under inverses.

(Joint work with Oscar Lautsch.)

HARM DERKSEN, Northeastern University
Bi-Lipschitz Invariant Theory and Quantum Channels

Given an affine variety in a high dimensional Euclidean space, a random projection onto a subspace of smaller dimension is often preserves the norm up to a small perturbation and is bi-Lipschitz with small distortion. This has applications in Bi-Lipschitz Invariant Theory and counterexamples for the additivity of quantum channels. I will discuss the problem of finding such projections with bounds for the distortion without using randomization.

HONGHAO FU, Concordia University
Nonlocal Games in the High-Noise Regime: Optimal Quantum Values and Rigidity

Self-testing is a key characteristic of certain nonlocal games, which allow one to uniquely determine the underlying quantum state and measurement operators used by the players, based solely on their observed input-output correlations. Motivated by the limitations of current quantum devices, we study self-testing in the high-noise regime, where the two players are restricted to sharing many copies of a noisy entangled state with an arbitrary constant noise rate. In this setting, many existing self-tests fail to certify any nontrivial structure. We first characterize the maximal winning probabilities of the CHSH game, the Magic Square game, and the 2-out-of-n versions of these two games as functions of the noise rate. These results enable the construction of device-independent protocols for estimating the noise rate. Building on this analysis, we show that these games can self-test one, two, and n pairs of anticommuting Pauli operators, respectively. These are the first known self-tests that are robust in the high-noise regime and remain sound even when the players' measurements are noisy. Our proofs rely on Sum-of-Squares (SoS) decompositions and Pauli analysis techniques developed in the contexts of quantum proof systems and quantum learning theory.

OLIVIER LALONDE, University of Waterloo
An exponential separation between entanglement-assisted and unassisted one-way quantum communication

A longstanding question in quantum communication complexity is whether some task can be accomplished with a small amount of communication in the presence of entanglement, yet require much more quantum communication in the absence of entanglement. Separations of this nature were previously known for relational problems and, in the simultaneous message passing model, for partial functions. But it has remained unresolved whether any such separation exists for a total Boolean function. We resolve this question with an exponential separation in the one-way setting: we exhibit a family of total Boolean functions that can be computed with bits of one-way classical communication given prior entanglement, but that require qubits of one-way quantum communication without entanglement. Our function is a special case of the subgroup membership problem, first studied in the communication setting by Aaronson, Le Gall, Russell, and Tani.

KIERAN MASTEL, University of Ottawa
Reductions between entangled constraint satisfaction problems

The constraint satisfaction problem (CSP) dichotomy theorem of Bulatov and Zhuk states that, for each constraint language over a finite domain, the corresponding CSP is either in P or NP-complete. CSP instances can be formulated as nonlocal games with entangled players in which one player receives a constraint and the other a variable. Building on MIP* = RE, recent work shows that many classically NP-complete CSPs become RE-hard, and hence undecidable, in this setting, while many classically tractable CSPs remain in P. Motivated by the search for an entangled CSP dichotomy, I will discuss recent work on reductions between entangled CSP languages and their role in transferring complexity results.

MAXENCE MAYRAND, Université de Sherbrooke
The complex geometry of GKP codes

GKP codes, introduced by Gottesman, Kitaev and Preskill in 2001, encode finite-dimensional logical spaces into quantum harmonic oscillators, with code states given by joint eigenstates of commuting translations in phase space. I will explain

how the theory of complex abelian varieties, which originated in the work of Abel, Jacobi and Riemann and has remained a very active area of research ever since, provides a rigorous mathematical framework for studying them. As a result, the tools and expertise developed for abelian varieties offer new ways of constructing and understanding GKP codes. In turn, natural questions about GKP codes, such as finding codes with many easily implementable logical gates or with optimal robustness to noise, become concrete problems on which mathematicians and physicists can work together. This talk is based on joint work with Baptiste Royer, arXiv:2605.28784.

CONNOR PADDOCK, University of Calgary

From matrix algebras to Cuntz algebras: self-testing commuting operator correlations

An amalgamation of the matrix algebras $M_3(\mathbb{C})$ and $M_2(\mathbb{C})$, obtained by identifying a minimal projection in each, produces a corner isomorphic to the Cuntz algebra $\mathcal{O}_2$. A recent theorem of Sam Harris shows that just four prescribed moments uniquely determine a state on $\mathcal{O}_2 \otimes_{\max} \mathcal{O}_2$. In joint work with Simon Schmidt, we use these operator-algebraic ideas to construct a finite two-player nonlocal game whose optimal commuting operator correlation uniquely determines the underlying state. Remarkably, this state encodes an exact entanglement embezzlement protocol which cannot be realized in the local tensor-product model, even when the local spaces are allowed to be infinite dimensional (despite being approximable by finite-dimensional strategies). As a corollary, our nonlocal game gives the first genuine example of a commuting operator self-test. I will focus on the operator-algebraic structure behind the construction, particularly how finite-dimensional matrix algebras give rise to Cuntz algebras and how uniqueness of a state on a corner extends to the full algebra.

WILLIAM SLOFSTRA, University of Waterloo

DAVE TOUCHETTE, Université de Sherbrooke

Type-constrained de Finetti reductions

For many information processing tasks, de Finetti-style theorems can often simplify the analysis in worst-case input scenarios for which the task exhibits some permutation-invariance symmetry, as they can allow for a reduction from an analysis on worst-case inputs to that of i.i.d. inputs. If further information is available on the inputs, it might be advantageous to reflect this information in the de Finetti reduction.

In this presentation, we focus on a form of such constraint, based on the type of the input. This allows us to obtain a conceptually simple proof of a new de Finetti reduction for classical probability distributions, derived from elementary properties from the method of types.

NATHANIEL JOHNSTON (MOUNT ALLISON UNIVERSITY, Mount Allison University

PPT states of almost maximal Schmidt number

We construct PPT states on $\mathbb{C}^m \otimes \mathbb{C}^n$ that have Schmidt number asymptotically approaching the smaller local dimension. More specifically, we construct a PPT state with Schmidt number at least

$$\left\lceil \frac{m+n-\sqrt{(m-n)^2+4(m+n-1)}}{2} \right\rceil.$$

In the case of equal local dimensions ($m = n$), this becomes $n - \lfloor\sqrt{2n-1}\rfloor$, far exceeding previous constructions, which achieved $n/2 + O(1)$. In the case of unequal local dimensions, our result shows that there exists a PPT state on $\mathbb{C}^n \otimes \mathbb{C}^{3n-4}$ with Schmidt number at least $n-1$.

RUTH WANG, University of Waterloo
Peak state transfer in continuous quantum walks

Quantum spin networks offer a promising way to transfer quantum information between different parts of a quantum computer, but how reliably can a quantum state be transferred between two sites? We study this question using continuous-time quantum walks generated by adjacency or Laplacian Hamiltonians. While the desirable perfect state transfer (PST) is not always possible, the transfer probability may still attain its upper bound. This is known as peak state transfer. In this talk, we discuss peak state transfer and its behaviour under graph operations. We also explore some elementary properties of the bounding matrix, which provides the entrywise time-independent upper bounds on transition amplitudes between any pair of vertices in a graph. This work is a product of the PIMS First Year Interest Groups (FYIG), joint with Joy Cooper, Homer De Vera, Hermie Monterde, and Alireza Talebpour.

PEIXUE WU, Syracuse University
Classical communication over quantum channels: does entanglement help?

A fundamental question in quantum information theory is whether entangling the inputs across multiple uses of a quantum channel can increase the rate at which it transmits classical information. An old conjecture dating back nearly three decades ago asserted that it cannot. In 2009, Hastings dis proved this conjecture by proving that there exist channels for which entangled inputs strictly increase the communication rate. His argument, however, is probabilistic: it shows that suitably chosen random channels exhibit this effect, but it does not identify any specific one. In this talk, I will discuss how to construct an explicit construction of a channel with such an entanglement advantage. On the other hand, we show that for phase-covariant and all qubit-to-qubit channels, there is no entanglement advantage.

ALICE ZHENG, NIST / UMD
Computing Boolean properties of quantum sequences

Quantum memory is a scarce and costly resource, yet little is known about which learning tasks remain feasible under severe memory constraints. We study the problem of computing global properties of quantum sequences when quantum systems must be measured individually, without storing or jointly processing them. In our setting, a bit string $x \in \{0,1\}^n$ is encoded into an n-qubit product state $|\psi_{x_1}\rangle \otimes \cdots \otimes |\psi_{x_n}\rangle$, and the goal is to infer $f(x) \in \{0,1\}$ from measurements of this quantum encoding. We consider a simple local strategy, which we call the greedy strategy, that applies the same optimal single-system measurement independently to each subsystem and then infers $f(x)$ from the outcomes. Our main result gives a complete characterization of when the greedy strategy is optimal: it achieves the same maximum success probability as an unrestricted global measurement if and only if the target Boolean function is affine (in all but finitely many cases). We establish a universal performance guarantee for general Boolean functions, showing that the success probability of the greedy strategy is always at least the square of the optimal global success probability, in direct analogy with the Barnum-Knill bound for the pretty good measurement. These results demonstrate that even under extreme memory constraints, simple local measurement strategies can remain provably competitive for learning global properties of quantum sequences.

CUNLU ZHOU, Université de Sherbrooke
Modular-Annihilator Parent Hamiltonians for Purified Gibbs States

Purified Gibbs states provide a bridge between finite-temperature physics, dissipative dynamics, and ground-state preparation. In this talk, I will present an exact finite sum-of-squares (SoS) construction of their parent Hamiltonians and the associated Lindbladian based on modular annihilators. Given a finite set of Hermitian generators, the corresponding modular annihilators yield a frustration-free SoS representation without continuous time integrals or an explicit decomposition into Bohr-frequency sectors. The purified Gibbs state remains a common zero mode while the freedom to choose and combine the generators can be used to optimize the spectral properties of the parent Hamiltonian. For interacting systems, where the modularly dressed generators are not available in closed form, we introduce a Krylov–Lanczos approximation scheme and bound the

resulting ground-state error in terms of the modular approximation error and the parent Hamiltonian gap. Numerical results illustrate the accuracy of the interacting construction depends on temperature, interaction strength, and Krylov dimension. arXiv:2608.30272.