Ontario Combinatorics Workshop Contributed Talks

May 9th and 10th 2026

Please note that all talks are 20 minutes in length with 5 minutes for questions. There is an additional 5 minutes for transition.

Saturday May 9^{th} Morning Session – WLH 210

Alexander Clow, Simon Fraser University 11:00-11:25am

Title: Large Degenerate Subgraphs in Degenerate Graphs

Abstract: Given a graph G and a non-negative integer d let $\alpha_d(G)$ be the order of a largest induced d-degenerate subgraph of G. We prove that for any pair of non-negative integers $k > d$, if G is a k-degenerate graph, then $\alpha_d(G) \geq \max\{\frac{(d+1)n}{k+d+1}, n - \alpha_{k-d-1}(G)\}$. For k-degenerate graphs this improves a more general lower bound of Alon, Kahn, and Seymour. By modifying our arguments we also obtain an improved lower bounds on $\alpha_d(G)$ for graphs of bounded genus. We will conclude by considering the open problem of determining $\alpha_d(k) := \inf_{G \text{ with degeneracy } k} \frac{\alpha_d(G)}{|V(G)|}$.

This is joint work with Sean Kim, and Ladislav Stacho.

Taite LaGrange, University of Waterloo 11:30-11:55am

Title: Classes of graphs with sub-linear twin-width

Abstract: Twin-width is a graph and matrix parameter introduced in 2021 by Bonnet, Kim, Thomassé, and Watrigant as essentially a measure of the 'error' between vertex neighbourhoods over a series of vertex contractions. This talk covers some graph classes with unbounded twin-width. We present a tool for obtaining twin-width bounds in general by contracting a graph based on a partition by distinct neighbourhoods. For a graph G on n vertices, we show that if such a partition exists, then the twin-width of G is at worst sub-linear with respect to n. We use this to obtain an upper bound on the twin-width of interval graphs and of graphs with bounded VC dimension. The latter implies that hereditary classes have sub-linear twin-width if and only if they have bounded VC dimension. This is joint work with Therese Biedl and Sophie Spirkl.

Saturday May 9^{th} Morning Session – WLH 212

Benjamin Mustico, Clarkson University 11:00-11:25am

Title: Uniform Mixing in Continuous-Time Quantum Walks on Graphs

Abstract: Given a graph G with adjacency matrix A, the continuous-time quantum walk on G is given by the time-varying unitary matrix $U(t) = exp(-itA)$. We say that G has uniform mixing at time t if $U(t)$ is a flat matrix (whose entries have the same absolute value). Ahmadi et al. (2002) showed that a complete graph on n vertices has uniform mixing if and only if $n = 2, 3, 4$. In this work, we prove that for any complete graph, there is a signing (assignment of complex weights to edges) so that the signed complete graph has uniform mixing. As a further corollary, we show that our signing yields a speedup of uniform mixing time for an infinite family of Hamming graphs. These results illustrate the power of signings (or chirality) for quantum transport on graphs. This is joint work with Luke Levine, Gabe Tucker, Hanmeng Zhan, and Christino Tamon, and is supported by NSF grant OSI-2427020. Note that this is anticipated to be a joint talk with Jessy J. Mesapam

Jessy Jacob Mesapam, Clarkson University 11:30-11:55am

Title: Average Mixing in Chiral Quantum Walks

Abstract: Given a graph G with adjacency matrix A, the continuous-time quantum walk on G is given by $U(t) = exp(-itA)$. The mixing matrix $M(t)$ of this quantum walk is defined as the Schur product of $U(t)$ with its conjugate. The (a, b)-entry of $M(t)$ represents the probability of observing the quantum walk at b at time t when starting at a. The average mixing matrix (AMM) is the Cesaro limit of $M(t)$ as time t tends to infinity. This matrix encodes the long-term behavior of the quantum walk. We say that a graph has average uniform mixing if its AMM matrix is the normalized all-one matrix. Godsil (2013) proved that no graph has average uniform mixing except for K_2. In contrast, we show that there are infinite families of oriented graphs with average uniform mixing. This is joint work with Luke Levine, Benjamin Mustico, Christino Tamon, Gabriel Tucker and Hanmeng Zhan, and is supported by NSF grant OSI-2427020.

Saturday May 9^{th} Afternoon Session – WLH 210

Adrian Lee, University of Guelph 2:00-2:25pm

Title: A report on the results of the ongoing investigation of the Weak Closure Algorithm (WCA)

Abstract: The non-Hamiltonian cycle decision problem is co-NP complete for which the existence of a polynomial time algorithm is unknown. In 2017, Gismondi and Lee developed a heuristic algorithm ($O(n^8)$) called the weak closure algorithm (WCA) that investigated this problem using a total of 461 snarks that are known to be non-Hamiltonian. This algorithm successfully decided these snarks to be non-Hamiltonian. In this presentation, a follow up report of an ongoing investigation of the WCA testing a total of 459,986 3-regular graphs with vertex counts ranging from 4 to 18; and also 153,000 out of 510,489 3-regular graphs with 20 vertices is presented. All non-Hamiltonian graphs were verified non-Hamiltonian, and, no Hamiltonian graph was decided non-Hamiltonian (Hamiltonian graphs cause the WCA to exit Undecided) via Maple 2023. All tested graphs were generated by GENREG, an open-source software developed by Meringer in 1999.

MacKenzie Carr, Toronto Metropolitan University 2:30-2:55pm

Title: Reconstructing C_4-free graphs from their digital convexity

Abstract: Fomin, Kratochvíl, Lokshtanov, Mancini, and Telle showed that every C_4-free graph is reconstructible from the multiset of closed neighborhoods. We strengthen their result proving that every C_4-free graph is reconstructible from the *set* of closed neighborhoods. A subset S of vertices in a graph G is digitally convex if, for every $v \notin S$, there is a private neighbor of v. We establish that reconstruction from digitally convex sets is equivalent to reconstruction from the set of closed neighborhoods, thereby extending the work of Lafrance, Oellermann, and Pressey by showing that all C_4-free graphs, and hence all graphs of girth at least five, are reconstructible from their digitally convex sets.

Saturday May 9th Afternoon Session – WLH 212

Matteus Froese, University of Waterloo 2:00-2:25pm

Title: Poset Path Splitting

Abstract: In 2001, Hoffman introduced an analogue of a covering space for a ranked, locally finite poset. The author proved an existence and uniqueness theorem for universal covers in this framework and explored a number of examples. We extend this work to the case of finite, but not necessarily ranked, posets. We give a characterization of the covering spaces in terms of paths of the original poset. We use this characterization to show that the set of all covering spaces naturally forms a supersolvable lattice and explore some of its properties. In particular, we classify when it is semidistributive and extremal. Finally, we give an analogue of a deck transformation and show that the automorphism group of a poset is intimately related to the automorphism group of its covering space lattice. This is joint work with Ian George.

Hin Lon Lao, York University 2:30-2:55pm

Title: Spotting Circuits in Bipartite Presentations of Transversal Matroids

Abstract: A natural way to present a transversal matroid is by drawing a bipartite graph. We are going to see when given a bipartite presentation of a transversal matroid, how we can spot circuits of the transversal matroid by looking at the bipartite graph. Also, we spot edges which can be removed without affecting the transversal matroid it presents when the corresponding induced subgraphs of circuits intersect in a nice way in the bipartite graph. Moreover, we can apply our result to to make certain communication networks simpler.

Sunday May 10^{th} Morning Session – WLH 210

Teddy Mishura, Toronto Metropolitan University 11:00-11:25am

Title: Cooling Generalized Hypercubes

Abstract: Graph cooling was introduced as an analogue to the well known game of graph burning. Both of these games are round based processes where fires spread from burning vertices to adjacent vertices and the player (the Arsonist) lights a vertex on fire every round. In graph burning, the player must burn all the vertices of G as quickly as possible, whereas in cooling, the player must burn all the vertices of G as slowly as possible. The cooling number of the hypercube graph Q_n, written $CL(Q_n)$, was recently shown to be exactly n.

In this talk, I analyze the cooling number of two families of graphs that can be viewed as the extension of the hypercube graph Q_n — Cartesian products of complete graphs K_m and Cartesian products of path graphs P_k. In the former case, I present an explicit form for the cooling number, and in the latter case, I present asymptotically tight bounds. I end the talk with a brief discussion on open problems.

Based on joint work with Anthony Bonato, MacKenzie Carr, Caleb Jones, and Trent Marbach.

Meagan Mann, Queen's University 11:30-11:55am

Title: Predicting The Cop Number Using Machine Learning

Abstract: Cops and Robbers is a pursuit evasion game played on a graph, first introduced independently by Quilliot and Nowakowski and Winkler over four decades ago. A main interest in recent the literature is identifying the cop number of graph families. The cop number of a graph, $c(G)$, is defined as the minimum number of cops required to guarantee capture of the robber. Determining the cop number is computationally difficult and exact algorithms for this are typically restricted to small graph families. This paper investigates whether classical machine learning methods and graph neural networks can accurately predict a graph's cop number from its structural properties and identify which properties most strongly influence this prediction. Of the classical machine learning models, tree-based models achieve high accuracy in prediction despite class imbalance, whereas graph neural networks achieve comparable results without explicit feature engineering. The interpretability analysis shows that the most predictive features are related to node connectivity, clustering, clique structure, and width parameters, which aligns with known theoretical results. Our findings suggest that machine learning approaches can be used in complement with existing cop number algorithms by offering scalable approximations where computation is infeasible.

Sunday May 10^{th} Morning Session – WLH 212

Maryam Mohammadi Yekta, University of Waterloo 11:00-11:25am

Title: A lower bound for the coefficients of denormalized Lorentzian Laurent series

Abstract: We introduce and study the class of denormalized Lorentzian (DL) Laurent series. We prove lower bounds for the coefficients of such Laurent series using Gurvits' capacity function. As a result, we are able to obtain an already known capacity bounds for the number of contingency tables, a new lower bound for the number of integral flows of any acyclic directed graph, and a lower bound for the dimension of weight spaces of parabolic Verma modules.

Blake Shirman, York University 11:30-11:55am

Title: Encoding lattice paths as walks on the path graph

Abstract: It is a fundamental result of algebraic graph theory that the uv-entry of the m^{th} power of the adjacency matrix of a graph is equal to the number of m-step walks from vertex u to vertex v. One could certainly conceive of a bounded, oriented portion of the (gridded) integer lattice as a directed graph and compute select entries of powers of its adjacency matrix to count the number of lattice paths over such a region. The orientation would be necessary to ensure that these walks are indeed paths (no repeated vertices) in the graph theoretic sense. For instance, all $\{N, E\}$-lattice paths of certain length could be drawn as walks on some oriented rectangular region of the integer lattice, with each vertical arc (directed edge) pointing up 'north' and each horizontal arc pointing right 'east.' However, computing powers of the adjacency matrix is cumbersome for such a large graph, since an n by n lattice grid would have an $(n+1)^2$ by $(n+1)^2$ adjacency matrix. Luckily, walks on such a grid can be encoded as walks on the path graph with $2n+1$ vertices. Such an undirected graph has a diagonalizable adjacency matrix with a known spectral decomposition. Of course, $\{N, E\}$-lattice paths are easily counted by a combinatorial argument, but we may equate the resulting binomial coefficient to a trigonometric polynomial (sourced from the spectral decomposition of the path graph). We can apply the same process to Dyck paths, which are famously enumerated by the Catalan numbers (a central binomial coefficient). Equating the two results gives us a family of identities between the real and imaginary parts of various roots of unity. The utility of our method becomes apparent with height-restricted Dyck paths, for which combinatorial arguments become more cumbersome. The path graph encoding shines, giving us as close to a 'closed formula' as the problem might allow.

Sunday May 10^{th} Afternoon Session – WLH 210

Mary Rose Jerade, University of Ottawa 3:00-3:25pm

Title: The Honeymoon Oberwolfach Problem: A Recursive Approach

Abstract: The *Honeymoon Oberwolfach Problem* (HOP) is a variant of the Oberwolfach Problem, which was first introduced by Mateja Šajna. The problem asks whether it is possible to seat n couples attending a conference for $2(n-1)$ consecutive nights at ℓ round tables of sizes $2m_1, 2m_2, \ldots, 2m_\ell$ such that $m_i \geq 2$ for all $i = 1, \ldots, \ell$, $\sum_{i=1}^{\ell} m_i = 2n$, and each participant must be seated next to their spouse every night, but next to each other person exactly once. This problem is denoted by $HOP(2m_1, 2m_2, \ldots, 2m_\ell)$.

This seating arrangement problem smoothly translates into a graph theory problem. It is equivalent to finding a cycle decomposition of the graph $K_{2n} + (2n-3)I$ (that is, the complete graph K_{2n} with $2n-3$ additional copies of a fixed 1-factor I) into 2-factors such that each 2-factor consists of disjoint I-alternating cycles of lengths $2m_1, \ldots, 2m_\ell$. So far, solutions to HOP have been established for various cases, among which: for all $n \leq 20$; when all tables are of the same size; when the table sizes are divisible by 8; and when n is odd and the Oberwolfach Problem $OP(m_1, \ldots, m_\ell)$ has a solution.

In this talk, we present a recursive approach to the problem. In particular, we show that if $m_1, \ldots, m_\ell$ are positive even integers, t is an integer such that $\sum_{i=1}^{\ell} m_i < t$, and $HOP(2m_1, \ldots, 2m_\ell)$ has a solution, then $HOP(2m_1, \ldots, 2m_\ell, 2t)$ has a solution as well. Moreover, we will briefly touch on the case when $m_1, \ldots, m_\ell$ are not all even.

This is joint work with Mateja Šajna.

Masoomeh Akbari, University of Ottawa 3:30-3:55pm

Title: On the generalized honeymoon Oberwolfach problem

Abstract: The Honeymoon Oberwolfach Problem (HOP), introduced by Šajna, is a recent variant of the classic Oberwolfach Problem. This problem asks whether it is possible to seat $2m_1 + 2m_2 + \cdots + 2m_t = 2n$ participants, consisting of n newlywed couples, at t round tables of sizes $2m_1, 2m_2, \ldots, 2m_t$ for $2n-2$ nights, so that each participant sits next to their spouse every night and next to every other participant exactly once. This problem is denoted by $\mathrm{HOP}(2m_1, 2m_2, \ldots, 2m_t)$. Jerade, Lepine, and Šajna have studied the HOP and resolved several important cases.

We generalized the HOP by allowing tables of size two, relaxing the previous restriction that tables must have a minimum size of four. In the generalized HOP, we aim to seat the $2n$ participants at s tables of size 2 and t round tables of sizes $2m_1, 2m_2, \ldots, 2m_t$, where $2n = 2s + 2m_1 + 2m_2 + \cdots + 2m_t$ and $m_i \geq 2$, while preserving the adjacency conditions of the HOP. We denote this problem by $\mathrm{HOP}(2^{\langle s \rangle}, 2m_1, \ldots, 2m_t)$.

In this talk, we will present a general approach to this problem, as well as recent solutions to several cases.

Sunday May 10^{th} Afternoon Session – WLH 212

Prangya Parida, University of Ottawa 3:00-3:25pm

Title: One Sequence to Rule them All

Abstract: A *binary reflected Gray code* (BRGC), denoted by $B(n)$, is an ordering of the 2^n binary n-tuples such that any two consecutive tuples differ in exactly one bit. This object can also be generated using a greedy approach (introduced by Williams): start the list with the zero n-tuple and then repeatedly create a new tuple by flipping the rightmost possible bit. It is also well known that the position of the bit that changes at each step in the generation of $B(n)$ is given by the *ruler sequence* (OEIS A001511). This observation further leads to a loopless generation of $B(n)$ (due to Bitner, Ehrlich, and Reingold), where each successive tuple can be obtained in $\mathcal{O}(1)$ time.

Mixed-radix reflected Gray codes and *mixed-radix modular Gray codes* are natural extensions of the BRGC, where the base of each position is not necessarily binary. They are extensively discussed in The Art of Computer Programming, Volume 4A by Knuth; their respective loopless algorithms can be found as Algorithm H and Exercise 77 in Section 7.2.1.1.

In this talk, we discuss that both of these Gray codes built on a mixed-radix base are guided by the same change sequence, which allows them to be generated in parallel. We show that modular Gray codes can also be generated using a greedy cyclic increment approach, generalizing the greedy method for generating the BRGC. Furthermore, we present a new family of modular Gray codes that can be constructed starting from any tuple w, and then generating the next word using the greedy cyclic increment approach that is lexicographically greater than or equal to w. We show that although this order is not a suffix of the full modular Gray code, its change sequence is a suffix of the change sequence of the full modular Gray code for the given base.

This is joint work with Lucia Moura, Brett Stevens, and Aaron Williams.

Xinyue Fan, University of Waterloo 3:30-3:55pm

Title: Deletion-only Induced Saturation on Trees

Abstract: Let H be a graph. A graph G is H*-induced-saturated* if G is H-free – meaning that no induced subgraph of G is isomorphic to H – and adding or removing any edge in G leaves a graph that is no longer H-free. It is natural to ask: For which graphs H does there exist an H-induced-saturated graph G? A full answer remains out of reach; indeed, it remains open whether H-induced-saturated graphs exist for every tree H. In this talk, we consider a weaker variant that only concerns about removing edges, and specifically look at non-complete trees with two leaves at distance at most three where k-uniform hypergraphs play important roles in the proof. This follows from an even stronger result for graphs with two leaves satisfying certain technical conditions.

This is joint work with Sahab Hajebi, Sepehr Hajebi, and Sophie Spirkl.