Applications of Matroid Theory and Combinatorial Optimization to Information and Coding Theory

Navin Kashyap (Queen's University),
Emina Soljanin (Alcatel-Lucent Bell Labs)
Pascal Vontobel (Hewlett-Packard Laboratories)

August 2–7, 2009

The aim of this workshop was to bring together experts and students from pure and applied mathematics, computer science, and engineering, who are working on related problems in the areas of matroid theory, combinatorial optimization, coding theory, secret sharing, network coding, and information inequalities. The goal was to foster exchange of mathematical ideas and tools that can help tackle some of the open problems of central importance in coding theory, secret sharing, and network coding, and at the same time, to get pure mathematicians and computer scientists to be interested in the kind of problems that arise in these applied fields.

1 Introduction

Matroids are structures that abstract certain fundamental properties of dependence common to graphs and vector spaces. The theory of matroids has its origins in graph theory and linear algebra, and its most successful applications in the past have been in the areas of combinatorial optimization and network theory. Recently, however, there has been a flurry of new applications of this theory in the fields of information and coding theory.

It is only natural to expect matroid theory to have an influence on the theory of error-correcting codes, as matrices over finite fields are objects of fundamental importance in both these areas of mathematics. Indeed, as far back as 1976, Greene [7] (re-)derived the MacWilliams identities — which relate the Hamming weight enumerators of a linear code and its dual — as special cases of an identity for the Tutte polynomial of a matroid. However, aside from such use of tools from matroid theory to re-derive results in coding theory that had already been proved by other means, each field has had surprisingly little impact on the other, until very recently.

Matroid-theoretic methods are now starting to play an important role in the understanding of decoding algorithms for error-correcting codes. In a parallel and largely unrelated development, ideas from matroid theory are also finding other novel applications within the broader realm of information theory. Specifically, they are being applied to explore the fundamental limits of secret sharing schemes and network coding, and also to gain an understanding of information inequalities. We outline some of these recent developments next.

2 Background and Recent Developments

Our workshop covered four major areas within the realm of information theory — coding theory, secret sharing, network coding, and information inequalities — which have seen a recent influx of ideas from matroid theory and combinatorial optimization. We briefly discuss the applications of such ideas in each of these areas in turn.

2.1 Coding Theory

The serious study of (channel) coding theory started with Shannon's monumental 1948 paper [9]. Shannon stated the result that reliable communication is possible at rates up to channel capacity, meaning that for any desired symbol or block error probability there exists a channel code and a decoding algorithm that can achieve this symbol or block error probability as long as the rate of the channel code is smaller than the channel capacity. On the other hand, Shannon showed that if the rate is larger than capacity, the symbol and the block error probability must be bounded away from zero.

Unfortunately, the proof of the above achievability result is nonconstructive, meaning that it shows *only the existence* of such channel codes and decoding algorithms. Therefore, since the appearance of Shannon's theorem, the quest has been on to find codes with practical encoding and decoding algorithms that fulfill Shannon's promise.

The codes and decoding schemes that people have come up with can broadly be classified into two classes: "traditional schemes" and "modern schemes." In "traditional schemes," codes were proposed that have some desirable properties like large minimum Hamming distance (a typical example of such codes being the Reed-Solomon codes). However, given a code, it was usually unclear how to decode it efficiently. Often it took quite some time until such a decoding algorithm was found (e.g., the Berlekamp-Massey decoding algorithm for Reed-Solomon codes), if at all. In "modern schemes," the situation is reversed: given an iterative decoding algorithm like the sum-product algorithm, the question is what codes work well together with such an iterative decoding algorithm.

"Modern schemes" took off with the seminal paper by Berrou, Glavieu, and Thitimajshima in 1993 on turbo coding schemes [2]. Actually, codes and decoding algorithm in the spirit of "modern schemes" were already described in the early 1960s by Gallager in his Ph.D. thesis [5]. However, these schemes were, besides the work by Zyablov, Pinkser, and Tanner in the 1970s and 1980s, largely forgotten until the mid-1990s. Only then people started to appreciate Gallager's revolutionary approach to coding theory.

Gallager proposed to define codes in terms of graphs. Such graphs are now known as Tanner graphs: they are bipartite graphs where one class of vertices corresponds to codeword symbols and where the other class of vertices corresponds to parity-checks that are imposed on the adjacent codeword symbols. Decoding is then based on repeatedly sending messages with estimates about the value of the codeword symbols along edges, and to locally process these messages at vertices in order to produce new messages that are again sent along the edges. Especially for sparse Tanner graphs the resulting decoding algorithms have very low implementation complexity.

In the last fifteen years, Tanner graphs and iterative decoding algorithms have been generalized to factor graphs and algorithms operating on them, and many connections to techniques in statistical mechanics, graphical models, artificial intelligence, and combinatorial optimization were uncovered. The workshop talk by Kashyap (see Section 3.1) surveyed the connection between complexity measures for graphical models for a code and the treewidth (and other width parameters) of the associated matroid. On the other hand, the workshop talks by Wainwright and Vontobel (see Section 3.1) emphasized the connections between message-passing iterative decoding of codes and certain techniques from combinatorial optimization. In particular, they discussed the linear programming decoder by Feldman, Wainwright, and Karger [4], which is a low-complexity relaxation of an integer linear programming formulation of the maximum likelihood decoder. This linear programming decoder (and its variations) has paved the way for the use of tools from combinatorial optimization and matroids in the design and analysis of decoding algorithms.

2.2 Secret Sharing

The second major application of matroid-theoretic ideas that we mention here is with respect to secret-sharing schemes. A secret-sharing scheme is a method to distribute shares of a secret value among a certain number of participants such that *qualified* subsets of participants (e.g., subsets of a certain size) can recover the secret from their joint shares, but *unqualified* subsets of participants can obtain no information whatsoever about the secret by pooling together their shares. Secret-sharing schemes were originally motivated by the problem of secure storage of cryptographic keys, but have since found numerous other applications in cryptography and distributed computing.

It is not difficult to show that in a secret-sharing scheme, the size of each of the shares cannot be smaller than the size (information content) of the secret value. An *ideal* secret-sharing scheme is one in which all shares have the same size as the secret value. More generally, the *information rate* of a secret-sharing scheme is the ratio of the size of the secret to the maximum share size.

In a secret-sharing scheme, the collection of qualified subsets of participants is called the *access structure* of the scheme. It is known that for any monotone increasing collection, Γ, of subsets of a finite set, one can define a secret-sharing scheme with access structure Γ. The *information rate* $\rho(\Gamma)$ is defined to be the supremum of information rates among all secret-sharing schemes having access structure Γ. Γ is said to be an *ideal access structure* if it admits an ideal secret-sharing scheme.

Brickell and Davenport [3] began a line of work relating ideal secret-sharing schemes to matroids. They showed that any ideal access structure is induced by a matroid in a very specific sense. However, it is also known that not every matroid gives rise to an ideal access structure; for example, the access structures induced by the Vámos matroid are not ideal. Characterizing the matroids that give rise to ideal access structures has remained an open problem.

There has been some very recent work on computing the information rates of non-ideal access structures using polymatroid techniques, linear programming, and non-Shannon information inequalities. For example, it has been shown that for any access structure Γ induced by the Vámos matroid, $\rho(\Gamma) \leq 19/21$, which shows that such access structures are far from being ideal. This, and other related results, were surveyed in the workshop talks by Padró and Beimel (see Section 3.3).

Secret-sharing schemes have also been received some recent attention in the quantum domain, a topic covered in the workshop talk by Sarvepalli (see Section 3.3).

2.3 Network Coding

Another novel application of matroid theory and combinatorial optimization within the realm of information theory is in the area of network coding [10]. Network coding is an elegant technique introduced at the turn of the millennium to improve network throughput and performance. Since then, it has attracted significant interest from diverse scientific communities of engineers, computer scientists, and mathematicians in both academia and industry. This workshop explored connections between network coding and combinatorial optimization, matroids, and non-Shannon inequalities.

The area started when the simple but far reaching observation was made that in communication networks, (unlike in their transportation or fluid counterparts), data streams that are separately produced and consumed do not necessarily need to be kept disjoint while they are transported throughout the network [1]. (At the network layer, for example, nodes can perform binary addition of independent bit-streams.) Schemes that employ processing at network nodes of incoming independent data (as opposed to only forwarding) are referred to as network coding. Naturally, the throughput achievable by network coding is in general higher than what can be achieved by schemes that allow only forwarding. Certain standard problems in combinatorial optimization have been crucial in understanding the potential benefits of network coding. Charikar and Agraval as well as Chekuri, Fragouli, and Soljanin characterized the benefits for certain traffic scenarios and throughput measures, as discussed by Chekuri in his workshop talk (see Section 3.4).

Mathematically, data streams carried by network edges are treated as sequences of symbols which are elements over some finite field. Network nodes map the incoming multiple data streams into a single stream in a possibly different way for each of its outgoing edges. The goal is to choose these maps in way that will allow the intended receivers to recover the original information. In the simplest case of network multi-cast (one in which the source aims at communicating the same information to a set of receivers), it is sufficient that the nodes forward linear combinations of the incoming symbols. The coefficients in these linear combinations can even be chosen uniformly at random from a sufficiently large finite field. In more complex traffic scenarios, such linear network coding is not sufficient, and matroids have been instrumental in demonstrating this fact. In a series of recent papers, Dougherty, Freiling, and Zeger carried out an exploration of the fundamental limits of network coding. They used matroids to systematically construct various networks that demonstrated, for example, the insufficiency of linear network coding and the inachievability of network coding capacity. A survey of these results was given by Sprintson in his workshop talk (see Section 3.4).

Finally, network coding problems give certain operational meaning to non-Shannon information inequalities. Raymond Yeung, one of the inventors/pioneers in both areas believes that implications of non-Shannon-type inequalities in information theory will be finally understood in the context of network coding. He declared in his talk that “Every constraint on the entropy function is useful in some multi-source network coding problems!” These and other applications of non-Shannon information inequalities, as well as the fundamentals, were addressed in a separate session of the workshop.

2.4 Information inequalities

As mentioned above, non-Shannon information inequalities play a key role in computing the information rates of non-ideal secret-sharing access structures. Furthermore, the results of Dougherty *et al.* in the context of network coding also make heavy use of these inequalities. We briefly give some background on information inequalities here. The workshop talks of Yeung, Matúš, and Dougherty (see Section 3.5) contain a more comprehensive survey of this topic.

Constraints on the entropy function are sometimes referred to as the laws of information theory. It has been known for a long time that the entropy function must satisfy the polymatroid inequalities (non-negativity, monotonicity, and submodularity), and indeed, that these are equivalent to the non-negativity of the Shannon information measures. Inequalities that are implied by the polymatroid inequalities are referred to as *Shannon-type inequalities*. Until recently, Shannon-type inequalities were the only known linear constraints on the entropy function.

A *non-Shannon-type inequality* is a constraint on the entropy function which is not implied by the polymatroid inequalities. In the late 1990s, the discovery of a few such inequalities, starting with the Zhang-Yeung inequality [11], revealed that Shannon-type inequalities alone do not constitute a complete set of constraints on the entropy function.

Linear information inequalities correspond to the supporting hyperplanes of the closed convex cone $\overline{\Gamma}^*_N$ obtained by taking the closure of the set of entropy vectors defined by N random variables. By virtue of the fact that entropy vectors satisfy the polymatroid inequalities, the cone $\overline{\Gamma}^*_N$ is a subset of the closed convex cone Γ_N defined by the polymatroid inequalities. It is a fact that $\overline{\Gamma}^*_N = \Gamma_N$ for $N \leq 3$, but $\overline{\Gamma}^*_N \subsetneq \Gamma_N$ for $N \geq 4$. In fact, Matúš has shown that for $N \geq 4$, $\overline{\Gamma}^*_N$ is not even polyhedral, *i.e.*, it cannot be characterized by finitely many linear inequalities. This means that there are infinitely many distinct non-Shannon inequalities satisfied by entropy vectors defined by $N \geq 4$ random variables. Matúš’s study of the cones $\overline{\Gamma}^*_N$ also involves the use of matroid methods in a non-trivial way.

3 Presentation Highlights

We provide here brief descriptions of the talks presented at the workshop. Slides from most of the talks are available online at `http://robson.birs.ca/∼09w5103/`.

James Oxley kicked off the workshop with a tutorial on matroid theory. His talk introduced the most common ways to define matroids and then presented some fundamental examples, some basic constructions, and some of the main theorems of the subject. A more thorough introduction to matroids is contained in the survey paper “What is a matroid?” available at `http://www.math.lsu.edu/∼oxley/survey4.pdf`.

3.1 Coding theory

Oxley’s tutorial was followed by a day-long session, consisting of four talks, focusing on the use of matroid theory and combinatorial optimization in coding theory.

Navin Kashyap gave an overview of the applications of matroid methods to the study of graphical models for codes, and to the analysis of decoding methods such as the sum-product algorithm and linear-programming decoding. Among the topics covered were the use of code/matroid decomposition techniques, and various “width” parameters (treewidth, branchwidth) associated with graphs and matroids, in the analysis of graphical models and decoding algorithms for linear codes.

Martin Wainwright’s talk began with an overview of the various hierarchies of linear programming (LP) relaxations, as well as related conic programming relaxations (e.g., SOCP and SDP), that can be applied to a given integer program. It then went to cover their on-going applications in coding theory and other areas of applied mathematics, and the connection between such hierarchies and the hypergraph defined by the underlying integer program. He also described some links between these relaxations, and various types of "message-passing" algorithms that are widely used in communication theory as well as many other domains (e.g., statistical physics, computer vision, machine learning, computational biology).

Pascal Vontobel focused on pseudo-codewords, i.e., certain non-zero fractional vectors that play an important role in the performance characterization of iterative message-passing decoders as well as linear programming decoding. This is in contrast to classical coding theory where decoding algorithms are mostly characterized by non-zero codewords. The talk gave an overview of results about pseudo-codewords and their influence on message-passing iterative decoding and linear programming decoding. The topics that were covered included: pseudo-codewords for cycle codes and their relationship to the graph zeta function; pseudo-codewords for finite-geometry-based codes; pseudo-codewords obtained by canonical completion, and how they upper bound the performance of linear programming decoding; the influence of redundant rows in the parity-check matrix on the set of pseudo-codewords; the relationship of pseudo-codewords to other concepts like stopping sets, near-codewords, trapping sets, and absorbing sets.

The final talk in the coding theory session was given by **Thomas Britz**, who presented a brief overview on what is presently known about the support and weight connections between coding and matroid theory, and gave applications of these connections to coding and graph theory. The newest results included an interesting variation of the Tutte polynomial as well as an interesting but ever-evolving dual identity.

3.2 The Matroid Minors Project

The morning of the second day of the workshop was devoted to the Matroid Minors Project of Jim Geelen, Bert Gerards, and Geoff Whittle. This project aims to extend the results and techniques of the Graph Minors Project of Robertson and Seymour (see *e.g.* [8]) to matrices and matroids. One of the main goals of this theory is to describe precisely the structure of minor-closed classes of matroids representable over finite fields. This requires a peculiar synthesis of graphs, topology, connectivity, and algebra. In addition to proving several long-standing conjectures in the area, the structure theory is expected to help find efficient algorithms for a general class of problems on matrices and graphs [6].

Bert Gerards presented an overview of the structure theorem (whose proof has just recently been completed by Geelen, Gerards and Whittle) for minor-closed classes of binary matroids. This theorem is a major milestone of the Matroid Minors Project. One important implication of this theorem is that every minor-closed class of binary matroids is characterized by a finite set of excluded minors.

Jim Geelen followed Gerards' talk by surveying some of the applications of the binary matroids structure theorem. It follows from the theorem that there exists an $O(n^7)$ algorithm for testing an n-element binary matroid for the presence of a fixed minor. An application pertinent to coding theory is the interesting result that proper minor-closed families of binary linear codes cannot be asymptotically good. Geelen further presented some open problems concerning minor-closed classes of binary matroids.

3.3 Secret sharing

The theme for the afternoon session on the second day was secret-sharing schemes. In a secret-sharing scheme, a secret value is distributed into shares among a set of participants in such a way that the qualified subsets of participants can recover the secret value, while the non-qualified ones do not obtain any information about it. In this situation, the size of every share is at least the size of the secret. If all shares have the same size as the secret, which is the best possible situation, the scheme is said to be ideal. Only a few access structures admit an ideal secret sharing scheme. In general, one is interested in finding schemes with optimal share length for every given access structure. This is a difficult problem that has attracted the attention of many researchers.

Carles Padró discussed several methods to find upper and lower bounds on the share length. He presented the most important results and techniques that have been obtained about this open problem from combinatorics, specially from the use of matroids and polymatroids. He also discussed some combinatorial techniques to construct efficient linear secret sharing schemes.

Amos Beimel, in a talk based on joint works with Noam Livne, Carles Padró, and Ilan Orlov, presented the use of non-Shannon information inequalities for proving lower bounds on the size of shares in secret-sharing schemes. He described two results:

1. A proof, using non-Shannon information inequalities, of lower bounds on the size of the shares in every secret-sharing scheme realizing an access structure induced by the Vámos matroid. This is the first result showing the existence of an access structure induced by a matroid which is not nearly ideal.
2. A proof of the fact that all the information inequalities known to date cannot yield a lower bound of $\Omega(n)$ on the share size.

Pradeep Kiran Sarvepalli talked about the applications of matroids quantum secret sharing, which deals with the problem of distribution of a quantum state among n players so that only authorized players can reconstruct the secret. He presented the first steps toward a matroidal characterization of quantum secret-sharing schemes. This characterization allows one to construct efficient schemes from self-dual matroids that are coordinatizable over a finite field. In the process, he also provided a connection between a class of quantum stabilizer codes and secret-sharing schemes.

Sarvepalli also briefly surveyed the use of matroids in quantum computation and quantum cryptography. He reviewed a recent work by Shepherd and Bremner which claims that even restricted models of quantum computation, such as those consisting of abelian gates, give rise to probability distributions that cannot be sampled efficiently by a classical computer. He sketched their arguments that use the theory of binary matroids to substantiate their claim.

Sarvepalli also considered an open problem related to the classification of a class of quantum states called the stabilizer states. A restricted version is to classify the equivalence classes of a subclass of stabilizer states (namely, the CSS states) under the action of the local unitary group and a subgroup of the local unitary group, called the local Clifford group. Specifically, the problem is to find necessary and sufficient conditions for when a CSS stabilizer state has distinct equivalence classes. Sarvepalli showed that CSS stabilizer states whose equivalence classes are distinct must arise from binary matroids which are neither graphic nor co-graphic. In doing so, he arrived at a class of minor-closed matroids whose excluded minors have not yet been characterized.

3.4 Network Coding

Network coding was the theme for the third day of the workshop, when a tutorial and two survey talks were given, followed by a presentation of an open problem. Network coding was also discussed on the two following days in connection with non-Shannon inequalities, some recent results in wireless networks, and general hardness to find a network coding scheme that achieves, or approximately achieves, capacity.

Emina Soljanin gave a tutorial talk on coding for network multicast (simultaneously transmitting the same information to multiple receivers in the network). She explained sufficient and necessary conditions that the network has to satisfy to be able to support the multicast at a certain rate. For the case of unicast (when only one receiver at the time uses the network), such conditions have been known for the past fifty years, and, clearly, we must require that they hold for each receiver participating in the multicast. The fascinating fact that the main network coding theorem brings is that the conditions necessary and sufficient for unicast at a certain rate to each receiver are also necessary and sufficient for multicast at the same rate, provided the intermediate network nodes are allowed to combine and process different information streams.

Chandra Chekuri surveyed results that seek to understand the potential benefit that network coding offers over more traditional and simpler transmission schemes such as store and forward routing. This was examined by asking the following question: what is the maximum ratio (over all networks) between the rate achievable via network coding and via routing? He restricted his attention to the wireline setting. This question has been answered to a large extent in the multicast setting in both undirected and directed graphs. In the multiple unicast setting, the benefit is known to be very large in some directed graph instances while the case of undirected graphs is wide open. Combinatorial optimization plays an important role in understanding this question. Steiner-tree packings and integrality gaps of linear programming relaxations for Steiner trees are the key tools in the multicast setting. Multicommodity flow-cut gaps play a role in the multiple unicast setting.

In his talk, **Alex Sprintson** gave an extensive survey of connections between matroid theory and network coding. He presented two ways of constructing new classes of coding networks from matroids. These constructions are instrumental for establishing several important properties of coding networks, such as insufficiency of scalar and vector linear network coding and inachievability of network coding capacity. He also explained the recently introduced problem of index coding, and pointed out its role as an intermediate step from a given matroid to a network whose dependency relations satisfy the given matroidal constraints. He presented recent results in this research area and outlined directions for future work.

The final talk of the session was given by **Randall Dougherty**, who outlined an approach that, if two proof-holes in it can be filled or worked around, will yield a proof that the solvability problem for network coding is undecidable. The idea was to try to represent groups satisfying or not satisfying identities as networks, in order to reduce Rhodes' problem on finite groups to the network coding solvability problem.

3.5 Information inequalities

The penultimate day of the workshop was the last "themed" day, the theme this time being information inequalities. Information inequalities are inequalities that must be satisfied by entropies of random variables.

Raymond Yeung's tutorial talk gave the necessary background on information inequalities. It is well-known that the entropy function must satisfy the polymatroidal axioms. All information inequalities implied by the polymatroidal axioms are called Shannon-type inequalities. In 1998, Zhang and Yeung discovered a non-Shannon-type inequality, an information inequality that is independent of the polymatroidal axioms. Since then, many more such inequalities have been found, and connections between the entropy function and a number of fields in information science, mathematics, and physics have been established. Yeung gave several examples of such connections to the fields of probability theory, network coding, combinatorics, group theory, Kolmogorov complexity, matrix theory, and quantum information theory.

Frantisek Matúš considered the problem of characterizing the closed cone, $\overline{\Gamma}^*_N$, formed by taking the closure of the set of entropic points on N random variables. He showed that this cone is not polyhedral, meaning that it cannot be characterized by finitely many linear inequalities, if and only if $N \geq 4$. He also discussed the problem of determining which matroids lie within $\overline{\Gamma}^*_N$, and mentioned that it remains an open problem to identify the excluded minors for this class of "almost entropic" matroids.

The third talk of this session was given by **Andreas Winter**, and was mainly based on a joint paper with N. Linden on quantum (van Neumann) entropy inequalities. Pippenger has initiated the generalization of the programme to find all the "laws of information theory" to quantum entropy. The standard quantum information inequalities derive from strong subadditivity (SSA), which corresponds to the third polymatroidal axiom. SSA of the von Neumann entropy, proved in 1973 by Lieb and Ruskai (who was present at the workshop), is a cornerstone of quantum information theory. All other known inequalities for entropies of quantum systems may be derived from it. In his talk, Winter proved a new inequality for the von Neumann entropy which is independent of strong subadditivity: it is an inequality which is true for any four party quantum state, provided that it satisfies three linear relations (constraints) on the entropies of certain reduced states. He also discussed the possibility of finding an unconstrained inequality (work with N. Linden and B. Ibinson).

Randall Dougherty gave his second talk of the workshop in this session, this talk being on non-Shannon-type information inequalities and linear rank inequalities. He first gave an alternate proof of Zhang and Yeung's non-Shannon-type inequality in four random variables. Zhang and Yeung's original proof used the technique of adding two auxiliary variables with special properties and then applying Shannon-type inequalities to the enlarged list. Dougherty presented a derivation of this inequality by adding just one auxiliary variable. He then used the same basic technique of adding auxiliary variables to give many other non-Shannon inequalities in four variables (which, surprisingly, are all of the same general form). He also derived rules for generating new non-Shannon inequalities from old ones, which can be applied iteratively to generate infinite families of inequalities such as the one used by Matúš to show that the cone $\overline{\Gamma}^*_4$ is not polyhedral.

Dougherty further showed how a variant of this approach (using a different sort of auxiliary variable) allowed one to derive inequalities which always hold for ranks of linear subspaces, but need not hold for entropies of random variables. It is known that the Ingleton inequality and the Shannon inequalities give a complete list of the rank inequalities for four variables (subspaces). Dougherty derived a list of 24 additional inequalities in five variables which, together with the Shannon inequalities and instances of the Ingleton inequality, are complete for rank inequalities on five subspaces. He also gave general many-variable families of rank inequalities.

3.6 Short talk sessions

The remaining sessions of the workshop consisted of short talks on several different topics related to the overall theme of the workshop.

Alex Grant presented his work with Terence Chan on quasi-uniform codes and their applications. Quasi-uniform random variables have probability distributions that are uniform over their support. They are of fundamental interest because a linear information inequality is valid if and only if it is satisfied by all quasi-uniform random variables. In his talk, Grant investigated properties of codes induced by quasi-uniform random variables. He proved that quasi-uniform codes (which include linear codes as a special case) are distance-invariant and that Greene's Theorem holds in the setting of quasi-uniform codes. He also showed that almost affine codes are a special case of quasi-uniform codes in the sense that quasi-uniform codes are induced by entropic polymatroids while almost affine codes are induced by entropic matroids. Applications of quasi-uniform codes in error correction and secret sharing were also given.

Serap Savari presented a combinatorial study of linear deterministic relay networks. This network model has gained popularity in the last few years as a means of studying the flow of information over wireless communication networks. This model considers layered directed graphs, and a node in the graph receives a linear transformation of the signals transmitted to it by neighbouring nodes. There is recent work extending the celebrated max-flow/min-cut theorem of Ford and Fulkerson to this model. This result was first established by a randomized transmission scheme over large blocks of transmitted signals. In joint work with S. Tabatabaei-Yazdi, Savari demonstrated the same result with a simple, deterministic, polynomial-time algorithm which takes as input a single transmitted signal instead of a long block of signals. Their capacity-achieving transmission scheme requires the extension of a one-dimensional Rado-Hall transversal theorem on the independent subsets of columns of a column-partitioned matrix into a two-dimensional variation for block matrices. The rank function arising from the study of cuts in their model has an important difference from the rank functions considered in the literature on matroids in that it is submodular but not monotone.

Eimear Byrne presented upper bounds for a particular model of error-correcting codes for coherent network coding. Versions of the Singleton, sphere-packing, and Gilbert-Varshamov bounds for this model were previously given by Yang and Yeung. In her talk, Byrne showed how to extend the classical Plotkin and Elias bounds for the same model.

The final session of the workshop began with a talk by **Dillon Mayhew** on the excluded minors for real-representable matroids. Rota conjectured that if F is a finite field, then there is only a finite number of minor-minimal matroids that are not F-representable. Such matroids are called excluded minors for F-representability. Rota's conjecture contrasts with the long-established fact that there are infinitely many excluded minors for representability over the real numbers. Geelen (2008) conjectured a much stronger fact: if M is any real-representable matroid, then there is an excluded minor, N, for real-representability, such that N contains M as a minor. Mayhew presented a proof of Geelen's conjecture (joint work with Mike Newman and Geoff Whittle).

Michael Langberg discussed the algorithmic complexity of network coding, focusing on how ”hard” it is to find a network coding scheme that achieves, or approximately achieves, capacity. He gave proofs of the fact that deciding whether or not a given instance of a network coding problem (acyclic network plus communication requirements) has scalar linear capacity of 1 is NP-complete. He further showed that it is "hard" (in the sense of being reducible to an open problem in graph colouring) to find a scalar linear code that enables communication with any constant factor of capacity. The same hardness result extends to the problem of finding a vector linear code of a fixed dimension.

Olgica Milenkovic gave a talk which approached the problem of compressive sensing via matroid theory. Compressive sensing is a new sampling technique for sparse signals that has the potential to significantly reduce the complexity of many data acquisition techniques. Most compressive sensing reconstruction techniques are still prohibitively time-consuming, narrowing the scope of practical applications of this method. Milenkovic, in joint work with Wei Dai and Vin Pham Hoa, proposed a new method for compressive sensing signal reconstruction of logarithmic complexity that combines iterative decoding methods with greedy subspace pursuit algorithms. The performance of the method depends on certain characteristics of support weight enumerators of the codes used for constructing the sensing matrix, which can be described via matroid theory.

The final talk of the workshop was given by **Alexander Barg** on the subject of linear codes in the ordered Hamming space. As is well known, the weight distribution of MDS codes in the Hamming metric can be recovered easily from the rank function of a uniform matroid. No such association has been established for the ordered Hamming space (the Niederreiter-Rosenbloom-Tsfasman space), although the weight distribution of MDS codes is also easily found. The question becomes more challenging when one considers codes with distance even one less than the MDS distance. Barg presented his work with Punarbasu Purkayastha which computes such weight distributions for an arbitrary poset metric and characterizes distributions of points in the unit cube that arise from near-MDS codes in the ordered metric.

4 Outcome of the Meeting

The workshop achieved its stated goal of encouraging interactions between researchers from several different disciplines, for whom there is currently no other forum (conference or workshop) that could serve as a natural meeting point. As a result, the workshop was extremely well received by all the participants, making it an unqualified success. Here we list some of the feedback that we received from the participants.

Thanks for organizing a beautiful workshop. I enjoyed my time during the workshop days no less than a fantastic weekend of hikes. *— Alexander Barg (University of Maryland, College Park)*

Once again, many thanks for the invitation to Banff – it was a very enlightening workshop.
— Eimear Byrne (University College Dublin)

It was very nice to be in Banff, thank you once again for the invitation.
— František Matúš (Institute of Information Theory and Automation Prague)

Thanks for your role in the workshop, it was very educational.
— Dillon Mayhew (Victoria University of Wellington)

Thanks a lot for the invitation to Banff – was a great workshop! Just continue organizing more of these.
— Olgica Milenkovic (University of Illinois, Urbana-Champaign)

Thanks again for the great workshop! *— Michael Langberg (The Open University of Israel)*

Thanks very much for putting together such an interesting workshop. I have enjoyed it very much indeed and am very glad I was invited to speak. *— James Oxley (Louisiana State University)*

Thanks for organizing an interesting and stimulating workshop. I also want to thank you for the opportunity to present at the workshop. I personally benefited a lot from the workshop especially in the sense of gaining a big picture of the associations between various fields. I was glad to have had some useful discussions with some of the workshop participants. *— Pradeep Kiran Sarvepalli (University of British Columbia)*

Thanks again for all of your work in organizing the workshop. *— Serap Savari (Texas A&M University)*

Thank you so much for letting me participate in this workshop. I learned a lot and enjoyed it very much.
— Beth Ruskai (Tufts University)

Best workshop that I've attended for quite a while. *— Alex Vardy (University of California, San Diego)*

Thanks again for organizing the workshop. *— Martin Wainwright (University of California, Berkeley)*

It was indeed a very nice meeting, I learnt a lot of new maths, and enjoyed myself very much!
— Andreas Winter (University of Bristol)

References

[1] R. Ahlswede, N. Cai, S.-Y.R. Li, and R.W. Yeung, Network information flow, *IEEE Trans. Inform. Theory*, **46** (2005), 1204–1216.

[2] C. Berrou, A. Glavieux, and P. Thitimajshima, Near Shannon limit error-correcting coding and decoding: turbo-codes (1), *Proc. IEEE Int. Conf. Communications*, Geneva, Switzerland, (1993), pp. 1064–1070.

[3] E.F. Brickell and D.M. Davenport, On the classification of ideal secret sharing schemes, *J. Cryptol.*, **4** (1991), 123–134.

[4] J. Feldman, M.J. Wainwright and D.R. Karger, Using linear programming to decode binary linear codes, *IEEE Trans. Inform. Theory*, **51** (2005), 954–972.

[5] R. G. Gallager, *Low-density parity-check codes*, M.I.T. Press, Cambridge, MA, (1963).

[6] J. Geelen, B. Gerards, and G. Whittle, Towards a matroid-minor structure theory. In *Combinatorics, Complexity and Chance. A Tribute to Dominic Welsh (G. Grimmett and C. McDiarmid, eds)*, Oxford University Press, 2007.

[7] C. Greene, Weight enumeration and the geometry of linear codes, *Studia Appl. Math.*, **55** (1976), 119–128.

[8] L. Lovász, Graph minor theory, *Bull. Amer. Math. Soc.*, **43** (2006), 75–86.

[9] C. E. Shannon, “A mathematical theory of communication,” Bell Syst. Tech. J., **27** (1948), pp. 379–423, 623–656.

[10] R.W. Yeung, *Information Theory and Network Coding*, Springer, 2008.

[11] Z. Zhang and R.W. Yeung, On characterization of entropy function via information inequalities, *IEEE Trans. Inform. Theory*, **44** (1998), 1440–1452.