2 Invited talks

On Edge Removal, Errors, and Equivalence

Michelle Effros

California Institute of Technology, Pasadena, California, USA

We consider a simple question about network coding: How large is the impact of any single edge on the capacity of a network of noiseless, capacitated links? For example, if rate vector $(R_1, ..., R_k)$ is achievable in a k-unicast network $\mathcal{N}$, is rate vector $(R_1 - \delta, ..., R_k - \delta)$ achievable in the network $\mathcal{N}_\delta$ that results by removing a single edge of capacity δ from network $\mathcal{N}$? This simple question lies at the heart of a surprising number of mysteries, a few of which will be described in this talk.

California Institute of Technology
Department of Electrical Engineering, MC 136-93
1200 E. California Blvd.
Pasadena, California, 91125 USA
effros@caltech.edu

Probabilistic construction of t-designs over finite fields

Shachar Lovett

University of California, San Diego

(Joint work with Arman Fazeli (UC San Diego), Greg Kuperberg (UC Davis), Ron Peled (Tel Aviv Uni.) and Alex Vardy (UC San Diego))

A $t-(n,k,\lambda)$ design over a finite field $\mathbb{F}_q$ is a collection of k-dimensional subspaces of $\mathbb{F}_q^n$, called blocks, such that each t-dimensional subspace is contained in exactly λ blocks. Such t-designs over $\mathbb{F}_q$ are the q-analogs of conventional combinatorial designs. Constructions of nontrivial t-designs over finite fields are currently known only for $t \leq 3$.

We show by a probabilistic argument that nontrivial t-designs over $\mathbb{F}_q$ exist for any t and q, provided that $k > 12t$ and n is large enough. This is based on a new probabilistic framework, which was also used to establish similar results for other combinatorial structures. I will describe both the general probabilistic framework and its application for designs over finite fields.

References

[1] G. Kuperberg, S. Lovett, and R. Peled, *Probabilistic existence of rigid combinatorial structures*, in STOC'12—Proceedings of the 2012 ACM Symposium on Theory of Computing, ACM, New York, 2012, pp. 1091–1105.

[2] A. Fazeli, S. Lovett, and A. Vardy, *Nontrivial t-designs over finite fields exist for all t*, (Submitted, arXiv:1306.2088 (2013)).

University of California, 9500 Gilman Drive, La Jolla CA 92037, San Diego
slovett@cse.ucsd.edu

Coding for distributed storage systems via rank-metric codes

Natalia Silberstein

Technion - Israel Institute of Technology, Israel

(Joint work with Ankit Singh Rawat, O. Ozan Koyluoglu, and Sriram Vishwanath)

In distributed storage systems (DSS) data is reliably stored over a network of nodes in such a way that a user can retrieve the stored data even if some nodes fail. To achieve such a resilience against node failures DSS introduce data redundancy based on different coding techniques. For example, erasure codes are widely used in such systems. When a single node fails, the system reconstructs the data stored in the failed node to keep the required level of redundancy. This process of data reconstruction for a failed node is called node repair process. During a node repair process the node which is added to the system to replace the failed node downloads data from a set of appropriate and accessible nodes.

There are two important goals that guide the design of codes for DSS: reducing the repair bandwidth, i.e. the amount of data downloaded from the nodes during the node repair process, and achieving locality, i.e. reducing the number of nodes participating in the node repair process. These goals underpin the design of two families of codes for DSS called regenerating codes and locally repairable codes.

In this talk, we first will present a construction of a new family of optimal locally repairable codes based on maximum rank distance (MRD) codes. Second, we will discuss hybrid codes which for a given locality parameters minimize repair bandwidth. These codes are based on a combination of locally repairable codes with regenerating codes. Finally, we will consider DSS containing nodes with adversarial errors. The key challenge in such systems is the propagation of erroneous data from a single corrupted node to the rest of the system during a node repair process. We will present a novel coding scheme for DSS which provides resilience against adversarial errors. This coding scheme also makes use of MRD codes.

Department of Computer Science, Technion - Israel Institute of Technology, Haifa 32000, Israel.
natalys@cs.technion.ac.il

Matrix Codes over Finite Fields and Rings

Danilo Silva

Federal University of Santa Catarina, Brazil

(Joint work with Roberto Nóbrega, Chen Feng, Frank Kschischang and Bartolomeu Uchôa-Filho)

It is well-known that the end-to-end communication over a packet network employing random linear network coding can be suitably modeled by a matrix channel over a finite field subject to possible rank deficiency and/or additive rank errors. Previous work has shown that the capability of a matrix code to guarantee successful decoding in such a scenario is completely characterized by the rank distance between matrices (if the channel is coherent) or the injection distance between rowspaces (if the channel is noncoherent). Optimal or nearly-optimal codes are known in both situations. Recently, it has been shown that a form of wireless network coding based on nested lattices induces, in many important cases, a matrix channel over a more general finite ring. This talk reviews some of the existing work on the finite field case and describes our recent efforts in generalizing these results for finite chain rings.

Department of Electrical Engineering, Federal University of Santa Catarina, Brazil
danilo@eel.ufsc.br

Kronecker algebraic space-time MIMO network code

M. A. Vázquez-Castro

Universitat Autònoma de Barcelona

Networking at the lower layers over error-prone wireless networks is con- sidered. Building upon the seminal work by Kötter and Médard in [1], the aim is an algebraic approach to single-source network coding over finite fields. Two fundamentally different assumptions to those in [1] are identified. First, the network coding coefficients are obtained at the physical layer according to some optimization (e.g. as in [2]). Second, the case of access networks is considered with a MIMO (Multiple Input Multiple Output) topology. Our proposed algebraic approach consists of considering the MIMO topology as an encoding function of the input data with the given network coding coefficients acting as a space-time erasure channel [3]. Such encoding function is derived and turns out to algebraically decouple the spatial encoding domain (topology) from the temporal encoding domain (error- correction). Moreover, it has a Kronecker structure which accounts for the natural symmetry and redundancy of wireless access networks. Finally, the model yields an unified treatment of unicast and multicast as weighted algebraic projections of the topologically encoded data. The min-cut max-flow theorem derived for upper layers takes on here a different interpretation.

References

[1] M. Médard and R. Kötter, *An algebraic approach to network coding,* IEEE/ACM Trans. on Networking, 11 (2003), pp. 782–795.

[2] B. Nazer and M. Gastpar, *Compute-and-forward: Harnessing interference through structured codes*, IEEE Transactions on Information Theory, 57 (2011), pp. 6463–6486.

[3] M. A. Vázquez-Castro, *The kronecker space-time network code*, in IEEE Int. Conference on Ultra Modern Telecommunications and Control Systems, 2013.

Universitat Autònoma de Barcelona
Angeles.Vazquez@uab.es

Steiner systems, t-designs and large sets of t-designs over finite fields

Alfred Wassermann

University of Bayreuth

Let $\mathcal{V}$ be a v-set (i.e. a set with v elements) whose elements are called *points*. A t-(v, k, λ) *design* is a collection of k-subsets (called *blocks*) of $\mathcal{V}$ with the property that any t-subset of $\mathcal{V}$ is contained in exactly λ blocks. A t-$(v, k, 1)$ design is also called *Steiner system*.

The notion of t-designs and Steiner systems has been extended to vector spaces over finite fields by Cameron and Delsarte in the 1970s: Now, $\mathcal{V}$ is a v-dimensional vector space over a finite field $\mathbb{F}_q$. A t-$(v, k, \lambda; q)$ design is a collection of k-dimensional subspaces of $\mathcal{V}$ (called blocks) with the property that any t-dimensional subspace of $\mathcal{V}$ is contained in exactly λ blocks.

Recently, due to the work of Kötter and Kschischang on random network codes [3], the interest in these so called q-analogs of designs has much increased. In the setting of network coding, q-analogs of Steiner systems are optimal constant dimension subspace codes. However, Metsch (1999) conjectured that Steiner systems over finite vector spaces do not exist for $t \geq 2$.

In [2] the first q-analog of Steiner systems for $t \geq 2$, and in [1] the first large sets of q-analogs of designs for $t \geq 2$ have been constructed. In my talk I will explain the constructions of the Steiner systems and the large sets and give a survey on other recent results on q-analogs of designs.

This is a joint work with M. Braun, T. Etzion, A. Kohnert, P. R. J. Östergård, and A. Vardy.

References

[1] M. Braun, A. Kohnert, P. R. J. Östergård, and A. Wassermann, *Large Sets of t-Designs over Finite Fields*, ArXiv e-prints, May 2013.

[2] M. Braun, T. Etzion, P. R. J. Östergård, A. Vardy, and A. Wassermann, *Existence of q-Analogs of Steiner Systems*, ArXiv e-prints, April 2013.

[3] R. Koetter and F. Kschischang, *Coding for errors and erasures in random network coding*, IEEE Transactions on Information Theory, 54 (2008), pp. 3579–3591.

Mathematisches Institut, Universität Bayreuth, 95440 Bayreuth, Germany
alfred.wassermann@uni-bayreuth.de