Algebraic Structure in Network Information Theory

Michael Gastpar (UC Berkeley and EPFL),
Frank Kschischang (University of Toronto)

August 14-19, 2011

Mathematics has always played an important role in the design and optimization of communication systems, particularly following Shannon's groundbreaking work in 1948 that gave theorems establishing fundamental limits on the rates of reliable communication over point-to-point channels (i.e., those involving a single transmitter-receiver pair). In order to develop fundamental bounds, concepts from probability and statistics and arguments involving averages taken over random code ensembles have been of key importance in most of the communications problems studied to date. Endowing these code ensembles with particular algebraic structure (e.g., the structure of a vector space) was not necessary to establish fundamental limits, and entered the stage only to allow for compact code descriptions and computationally-efficient encoding and decoding algorithms.

In *networks* (i.e., communication systems involving multiple transmitters or receivers or relay nodes) the situation becomes much more complicated, and a general framework that establishes fundamental limits is lacking. The current grand challenge in information theory is to devise communication strategies and architectures that optimally exploit communication networks. An emerging key insight—and the motivation for this workshop—is that in certain network information theory problems, endowing code ensembles with algebraic structure is a key necessity, not only for engineering convenience, but for the derivation of fundamental limits. The goal of the workshop was to bring together experts from information theory, coding theory, and algebra to shed light on this observation, with the goal of beginning to understand the type and extent of algebraic structure needed to extend Shannon's insights about point-to-point channels to the more general case of networks.

1 Overview of the Field

In his groundbreaking paper [23], Shannon determined the fundamental limits of reliable communication across a noisy channel and set forth the bit as the underlying unit of information. The mathematical framework used to prove these results involved statistical averages taken over ensembles of codebooks chosen at random, and thus forged a deep connection between information theory, probability, and statistics. Following Shannon, information theorists proving capacity theorems often resort to probabilistic arguments to prove the existence of good codes. As no particular algebraic structure is imposed, the resulting codebooks are generally regarded as being unstructured. In fact, Shannon's original random coding arguments do indeed determine optimal performance for all single-user noisy channels as well as for several multi-user networks.

In this classical context, algebraic structure is not a necessity, but it is a practical convenience, as such structure often allows for efficient encoding and decoding. The simplest examples are linear block codes, where the codewords form a finite-dimensional vector space over a finite field, allowing an exponentially-large number of codewords to be described as linear combinations of only a small number of basis codewords, and where decoding algorithms can exploit this algebraic structure for efficiency. In this classical context, it

seems that there may be a price to pay for this algebraic convenience: a folk theorem holds that, in general, codes with algebraic structure perform worse than the best unstructured codes. (This is in fact true for the case of a single noisy channel, where it is known that group codes are suboptimal except in the case of additive noise [1].)

More recently, it has become clear that this folk theorem may give the wrong insight in the network setting. That is, when several users communicate over noisy channels and interfere with each other (as occurs, for example, in wireless communications), requiring algebraic structure can be advantageous, even in the context of proving a capacity theorem. More precisely, proving the existence of a good algebraically-structured code becomes possible, even when arguments involving unstructured code ensembles fail. A first inkling of this appeared in the paper by Körner and Marton [12], which considers a very simple, non-standard distributed coding problem where the goal is to recover, not the original bits, but only their modulo-2 sum. For this problem, it was shown that linearly-structured codes attain optimal performance, while standard arguments involving ensembles without algebraic structure are not able to show that good codes exist.

Recent work by several groups has shown that arguments involving algebraically-structured codes for communicating bits across a network can result in proofs for the existence of codes having significant gains relative to what can be proved about codes drawn from unstructured ensembles, and specific examples (with references) will be given in the sequel. These gains appear to stem from the fact that algebraic structure enables users to perform distributed processing as if it were centralized.

Beyond these examples, algebraic structure is emerging as a key argument in several of the most important and challenging problems in network information theory, including the following:

Interference alignment: One of the long-standing open problems in information theory is the determination of the capacity of the so-called interference channel, wherein several transmitter-receiver pairs share the same communication channel. For many years, it was assumed that the maximum bit rate per user is inversely proportional to the number of active users. Surprisingly, each user can achieve half its interference-free bit rate, regardless of the number of interferers. The key is to carefully assign subspaces for transmission such that all signals (except the desired one) end up in a single subspace at each receiver [4]. Codes with algebraic structure have a crucial role, as recent work has shown [3, 21]. Algebraic structure also enters as an argument determining the maximum amount of interference alignment possible. An early account of this can be found in [8].

Distributed interference cancellation: Consider a power-limited transmitter and receiver that communicate over a noisy channel subject to interference. One of the celebrated results of information theory is that the effects of the interference can be completely removed as soon as the transmitter has prior knowledge of the interfering signal, regardless of its strength [6]. What if multiple users communicate to a single receiver over interfering links? When each transmitter has partial knowledge of the interference, cancellation of an additional interferer is possible, but only if algebraically structured codes are used. For an additive white Gaussian noise model, this involves lattice codes and has been studied in [22].

Computation over noisy channels: The traditional view of interference is that it is an obstacle to reliable communication. This is true if the objective of each receiver is to recover the message sent by a particular transmitter. However, if the receiver is only interested in a function of the messages, then interference can be harnessed to compute the function more efficiently, provided the transmitters use a structured code that is appropriately matched to the function of interest (e.g., see [20]).

Wireless "Network Coding": Even in simple networks, it is not sufficient to just route packets; instead, intermediate nodes may have to forward linear combinations of their incoming packets. This insight was found in [2] and has since sparked a great deal of research under the name of "network coding." If we consider said intermediate node, all it really needs to know is a particular linear combination of its incoming packets. Using codes with algebraic structure, we can enable just this, as advocated in [19], and further developed in recent work such as [18, 17]. Beyond this local perspective, codes with algebraic structure are also crucial at the end-to-end level, as recent work has shown [13].

Distributed quantization: It is increasingly the case that data is collected at several locations and then compressed for transmission across a network. The receiver may only be interested in a function of the sampled data and, in this case, it has been shown that gains are possible through the use of group and lattice codes [14, 15].

Secrecy: Algebraic structure enables several users to collude against adversaries and eavesdroppers directly over a wireless link, without any prior coordination (e.g., see [11]).

The emerging theme is that algebraically-structured codes can enable powerful new schemes, especially for wireless channel models, that have the potential to dramatically increase end-to-end bit rates. To date, most of the above research has used existing algebraic code constructions in a "plug-and-play" fashion. For discrete channels, these are mostly codes over finite fields, and for continuous-input channels, lattice codes [5]. Much work has focused on showing that lattice codes can approach the capacity of a single-user additive white Gaussian noise channel, beginning with [7] and culminating in several successful constructions [16, 24, 9], see [10] for a commentary. While these codes already give interesting results in networks, they are optimal only in certain limited special cases. One of the most interesting avenues for future research will be to develop new approaches to the construction of algebraically structured codes that are particularly well-suited for a given network communication scenario.

The research described above is the beginning of a new algebraic methodology for existence proofs in network information theory. To complement these methods, new tools for proving impossibility results will be needed as well. As it stands, impossibility results focus on which statistical dependencies can be established between different transmitters and receivers. Quite recently, new bounds have emerged for the interference channel that make use of results in additive combinatorics [8]. Much work is needed to bring algebraic structure to the forefront in impossibility results.

2 Recent Developments and Presentation Highlights

The presentation of recent developments was organized into five sessions, not necessarily mutually exclusive.

1. **Mathematical Foundations.** This session started with a tutorial lecture by Ram Zamir (Tel Aviv University) on the fundamentals of random lattices and how they relate to network information theory. Starting from the Minkowski-Hlawka ensemble, he gave an extensive coverage of the basic tools and ended with a series of intriguing starting points for discussion, which came under the heading of "anti-structure problems," namely, problems where algebraic arguments are not of key importance.

   Uri Erez (Tel Aviv University) presented a refined analysis of algebraic codes in network information theory, including a new result on error exponents on the Gaussian multiple-access channel.

   Sandeep Pradhan (University of Michigan, Ann Arbor) continued the exposition of the basic foundations with the discussion of nested linear constructions, going beyond the classical construction of nesting a "fine" linear code inside a "coarse" one.

   Nigel Boston (University of Wisconsin, Madison) presented on the mathematical foundations of convolutional codes.

   The session was completed by two short talks by graduate students about very recent work, namely Katie Morrison (University of Nebraska, Lincoln) on rank-metric codes, and Anna-Lena Trautmann (University of Zurich, Switzerland) on cyclic orbit codes.

2. **Algebraic Structure in Relay Networks.** This session started with a longer lecture of a certain tutorial character by Shlomo Shamai (Technion, Israel), diligently summarizing the state of the art of algebraic arguments as they appear in the recent literature, with a particular connection to the important emerging topic of interference management.

   This was followed by a sequence of shorter, more focused talks.

   Bobak Nazer (Boston University) started with a discussion of the "compute-and-forward" paradigm and then gave a thought-provoking presentation about the main missing ingredients and open problems

in the use of lattice coding in Gaussian networks. To underline the importance of his open problems as well as his Canadian heritage, he carefully chose names for each of the open problems, including the "icewine" problem, referring to the prizes offered for the solution of these problems.

Urs Niesen (Bell Labs, New Jersey) presented an improved "compute-and-forward" construction for Gaussian networks that performs significantly better in the limit as the noise process disappears (but requires a significant amount of channel state information at the coding devices in turn).

Krishna Narayanan (Texas A&M University) talked about practical code constructions for lattice codes, showing ways of explicitly going beyond binary codes.

In the second part of this session, the talks focused on questioning the ultimate need for algebraic structure.

Sae-Young Chung (Korea Advanced Institute of Science and Technology) started this part in a classical devil's advocate fashion. He considered one of the flagship results that are widely used to underline the need for algebraic structure, namely, the Gaussian two-way relay channel. For this channel, he showed that in the limit as the noise process disappears, almost the same performance can be attained by a code without any algebraic structure.

Gerhard Kramer (Technical University of Munich) presented a well-structured review of the emerging "noisy network coding" approach, again an approach that chooses not to deal with algebraic structure, as well as some intriguing extensions of this approach, in particular pertaining to what he referred to as "short messages."

Slawomir Stanczak (Technical University of Berlin) presented a complementary set of ideas about practical approaches to compute-and-forward (as well as an intriguing connection to Hilbert's 15th problem).

The final three talks considered in detail specific network problems for which it does not seem that algebraic structure is needed.

Natasha Devroye (University of Illinois, Chicago) considered the classical relay channel and showed that codes with algebraic structure can attain all results of codes without algebraic structure.

Liang Xie (University of Waterloo) presented list decoding ideas for the classical relay channel.

Ashish Khisti (University of Toronto) talked about delay-efficient coding for streaming applications.

This session was completed by several graduate student talks, covering recent developments. Chen Feng (University of Toronto) presented an algebraic approach to compute-and-forward. Matthew Nokleby (Rice University) presented ideas on "cooperative" compute-and-forward. Yiwei Song (University of Illinois at Chicago) presented further results on lattices in Gaussian relay networks (involving list decoding). Jiening Zhan (University of California, Berkeley) presented on algebraic structure and receiver architectures in networks.

3. **Algebraic Structure and Interference Alignment.** Algebraic structure of a slightly different kind is of key importance to another promising emerging direction in network information theory, namely, interference alignment. The third session of the workshop was devoted to this.

   Alex Dimakis (University of Southern California) started off from the real-world problem of distributed database repair, and showed how codes with algebraic structure play a key role in this problem.

   Viveck Cadambe (University of California, Irvine) posed the interference alignment problem as a problem of finding common invariant subspaces for a collection of (linear) operators. He elegantly expressed the problem as a canonical (though unsolved) question about tensors.

   Sriram Vishwanath (University of Texas, Austin) considered lattice codes for interference scenarios.

   Aylin Yener (Pennsylvania State University) expanded the horizon by considering the problem of communication under secrecy constraints, and showed how algebraic structure again plays a key role.

   This session was completed by a graduate student talk on very recent developments. Guy Bresler (University of California, Berkeley) presented new results on interference alignment for vector channels.

4. **Algebraic Structure and Network Coding.** This session covered the important topic of network coding and its many connections to algebraic structure.

   Joachim Rosenthal (University of Zurich) started the session with an introduction to Schubert calculus and its connection to rank-metric codes.

   Frederique Oggier (Nanyang Technological University, Singapore) discussed wiretap code design.

   Emanuele Viterbo (Monash University, Australia) presented results about network coding over finite rings, giving explicit constructions.

   Babak Hassibi (California Institute of Technology) discussed the important connections between network coding and matroid theory.

   Danilo Silva (Federal University of Santa Catarina, Brazil) presented the key ideas behind using algebraic codes to obtain error control for noncoherent network coding.

   Tracey Ho (California Institute of Technology) used an intriguing "zigzag" network example to illustrate fundamental algebraic questions in network coding.

   Alex Vardy (University of California, San Diego) then talked about a third class of codes of fundamental importance to the workshop, namely, subspace codes. He presented novel list decoding algorithms for these codes.

5. **Algebraic Structure and Source Coding.** The final session concerned the role of algebraic structure in distributed source coding.

   Aaron Wagner (Cornell University) presented fundamental results delineating those distributed source coding problem for which structure is important.

   Prakash Ishwar (Boston University) talked about interactive source coding.

   Mohammad Maddah-Ali (Bell Labs) presented new algebraic approaches to distributed source coding, giving crisp examples that showed the need for algebraic structure.

3 Open Problems

The workshop also allowed ample time for the discussion of open problems. For one, short breaks were scheduled between all talks to enable addressing detailed open problems arising in a narrower context. On top of this, dedicated open problems sessions were held in the evenings. These sessions happened both in the lecture halls as well as in the BIRS lounge and were very well attended.

One of the most heavily debated open question concerned the information expressions for the Gaussian networks. In particular, typical information-theoretic capacity results for Gaussian channels involve expressions of the type $\log(1 + \text{SNR})$. When using lattices for the "compute-and-forward" problem, in the two-terminal setting, it is possible to show a formula of the type $\log(\frac{1}{2} + \text{SNR})$. This might appear to be a small difference, but it has important ramifications when it comes to layering and multi-level codes. Therefore, the question is whether this difference is fundamental to the problem, and if it is not, what kinds of constructions might permit to improve performance. A fair amount of discussion was devoted to understanding this.

Another open problem, raised by Uri Erez, concerns the so-called "dirty paper" problem, where the received signal is of the form $Y^n = X^n + hS^n$, where h is a constant known only to the receiver and S^n is a sequence known only to the transmitter (in a non-causal fashion, at the beginning of time). Several approaches were developed during the session, but it remains to be determined whether they will work.

Bobak Nazer posed several questions pertaining to the possibility of doing "compute-and-forward" in a fast-fading environment, and Krishna Narayanan considered a simple two-way relay channel to understand the trade-offs between compute-and-forward and decode-and-forward, for which it seemed highly tempting that a final capacity result should be within reach.

Many other open problems, beyond this short sampling, were discussed, and we hope to see solutions in the literature shortly.

4 Concluding Remarks

The workshop achieved its main goal of providing a focused environment for in-depth discussion of this emerging research direction. None of the existing forums could provide a similarly dedicated setting. Deep technical discussions were enabled by the many high-quality research talks providing a steady stream of input, but perhaps even more importantly by several sessions dedicated to open problems. The BIRS setup provides a uniquely stimulating environment for this, including lecture halls of various sizes with ample blackboard space. While algebraic arguments provide network information-theoretic results that are out of reach of the classical random coding arguments, they can at present rarely be proved to be strictly optimal. This nagging fact provided ample open problems and many intriguing discussions. Another factor that contributed to the highly interactive environment was the fact that room and board were provided by BIRS, free of charge to participants. Therefore, many interesting discussions took place over breakfast, lunch and dinner and well into the night over coffee and other drinks in the well-equipped lounge at the BIRS facility. Holding a workshop at BIRS was a truly enjoyable experience for all involved.

References

[1] R. Ahlswede, "Group Codes do not Achieve Shannon's Channel Capacity for General Discrete Channels," *Annals Math. Stat.*, vol. 42, no. 1, pp. 224–240, Feb. 1971.

[2] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network Information Flow," *IEEE Trans. on Info. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[3] G. Bresler, A. Parekh and D. Tse, "The Approximate Capacity of the Many-to One and One-to-Many Gaussian Interference Channels," submitted to *IEEE Trans. on Inform. Theory*, Sep. 2008.

[4] V. Cadambe and S. Jafar, "Interference Alignment and Degrees of Freedom of the K-User Interference Channel," *IEEE Trans. on Info. Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.

[5] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, 3rd ed. New York: Springer-Verlag, 1999.

[6] M. Costa, "Writing on Dirty Paper," *IEEE Trans. on Info. Theory*, vol. 29, no. 3, pp. 439–441, May 1983.

[7] R. de Buda, "Some Optimal Codes Have Structure," *IEEE J. on Selected Areas in Commun.*, vol. 7, no. 6, pp. 893–899, Aug. 1989.

[8] R. Etkin and E. Ordentlich , "On the Degrees-of-Freedom of the K-User Gaussian Interference Channel," submitted to *IEEE Trans. on Info. Theory*, Jun. 2008. http://arxiv.org/abs/0901.1695

[9] U. Erez and R. Zamir, "Achieving 0.5 log(1+SNR) over the Additive White Gaussian Noise Channel with Lattice Encoding and Decoding", *IEEE Trans. on Info. Theory*, pp. 2293–2314, Oct. 2004.

[10] G. D. Forney, Jr., "On the role of MMSE estimation in approaching the information-theoretic limits of linear Gaussian channels: Shannon meets Wiener," *Proc. 2003 Allerton Conf.* (Monticello, IL), Oct. 2003.

[11] X. He and A. Yener, "Providing Secrecy With Structured Codes: Tools and Applications to Two-User Gaussian Channels," Submitted to *IEEE Trans. on Info. Theory*, Jul. 2009. http://arxiv.org/abs/0907.5388

[12] J. Korner and J. Marton, "How to Encode the Modulo-Two Sum of Binary Sources," *IEEE Trans. on Info. Theory*, vol. 25, no. 2, pp. 219–221, Mar. 1979.

[13] R. Koetter and F. R. Kschischang, "Coding for Errors and Erasures in Random Network Coding," *IEEE Trans. on Info. Theory*, vol. 54, no. 8, pp. 3579–3591, Aug. 2008.

[14] D. Krithivasan and S. Pradhan, "Lattices for Distributed Source Coding: Jointly Gaussian Sources and Reconstruction of a Linear Function," Submitted to *IEEE Trans. on Info. Theory*, Jul. 2007. http://arxiv.org/abs/0707.3461

[15] D. Krithivasan and S. Pradhan, "Distribued Source Coding using Abelian Group Codes," Submitted to *IEEE Trans. on Info. Theory*, Aug. 2008. http://arxiv.org/abs/0808.2659

[16] T. Linder, C. Schlegel, and K. Zeger, "Corrected Proof of de Buda's Theorem," *IEEE Trans. on Info. Theory*, vol. 39, no. 5, pp. 1735–1737, Sep. 1993.

[17] W. Nam, S.-Y. Chung, and Y. H. Lee, "Capacity Bounds for Two-Way Relay Channels," *2008 Int. Zurich Seminar on Commun. (IZS 2008)*, Zurich, Switzerland, Mar. 2008

[18] K. Narayanan, M. P. Wilson, and A. Sprintson, "Joint Physical Layer Coding and Network Coding for Bi-Directional Relaying," *45th Annual Allerton Conf.*, Monticello, IL, Sep. 2007.

[19] B. Nazer and M. Gastpar, Computing over Multiple-Access Channels with Connections to Wireless Network Coding, *Proc. of the IEEE Int. Symp. on Info. Theory (ISIT 2006)*, Seattle, WA, Jul. 2006.

[20] B. Nazer and M. Gastpar, "Computation over Multiple-Access Channels," *IEEE Trans. on Info. Theory*, vol. 53, no. 10, pp. 3498–3516, Oct. 2007.

[21] B. Nazer, M. Gastpar, S. Vishwanath, and S. Jafar, "Ergodic Interference Alignment," 2009 *IEEE Int. Symp. on Info. Theory*, Seoul, Korea, Jul. 2009. http://arxiv.org/abs/0901.4379

[22] T. Philosof, R. Zamir, and U. Erez and A. Khisti, "Lattice Strategies for the Dirty Multiple Access Channel," Submitted to *IEEE Trans. on Info. Theory*, Apr. 2009. http://arxiv.org/abs/0904.1892

[23] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623-656, 1948.

[24] R. Urbanke and B. Rimoldi, "Lattice Codes Can Achieve Capacity on the AWGN Channel," *IEEE Trans. on Info. Theory*, vol. 44, no. 1, pp. 273–278, Jan. 1998.