Report from Dagstuhl Seminar 16321

Coding Theory in the Time of Big Data

Edited by

Martin Bossert[1], Eimear Byrne[2], and Emina Soljanin[3]

1   **Universität Ulm, DE,** `martin.bossert@uni-ulm.de`
2   **University College Dublin, IE,** `ebyrne@ucd.ie`
3   **Rutgers University – Piscataway, US,** `emina.soljanin@rutgers.edu`

---- **Abstract** --

This report documents the program and the outcomes of Dagstuhl Seminar 16321 "Coding Theory in the Time of Big Data". The overarching technical theme was on how fundamentals of coding theory could be applied to data storage and transmission in the context of big data and conversely, on new problems in coding theory arising from such applications.

1   Executive Summary

Eimear Byrne
Martin Bossert
Emina Soljanin

The Dagstuhl Seminar 16321 *Coding Theory in the Time of Big Data*, held in August 7–12, 2016, was the third of a series of Dagstuhl seminars relating modern aspects of coding theory and its applications in computer science. The overarching technical theme was on how fundamentals of coding theory could be applied to data storage and transmission in the context of big data and conversely, on emerging topics in coding theory arising from such applications. In Dagstuhl Seminar 11461 the main topics discussed were list decoding, codes on graphs, network coding and the relations between them. The themes of distributed storage, network coding and polar codes were central to Dagstuhl Seminar 13351.

The conference was organised into six main working groups, as listed below:
1. Distributed Storage & Index Coding,
2. Private Information Retrieval for Storage Codes,
3. DNA-Based Storage,
4. Age & Delay of Information.
5. Code-Based Cryptography,
6. Rank-Metric Codes.

The amount of data that is being stored is scaling at a rapid pace making efficient data storage an important problem that inspires several lines of scientific research. During the seminar, several discussions were conducted on the theme of using classical and new techniques from coding theory to store/compute data efficiently in distributed storage systems. A number of open problems were identified, such as the design of codes with optimal repair bandwidth, fundamental trade-offs between storage & communication cost, applications to content distribution networks, connections between fundamental limits of storage/caching and the index coding problem and applications of coding theory for parallel computing. A theoretical framework and numerical simulation for the long term reliability of a distributed storage system were presented by Luby.

DNA-based storage was recently proposed to address new challenges to handle extremely high volume recording media to propose new compression methods for non-traditional data formats. Since DNA may be easily replicated and a massive amount of information stored reliably with minimal space requirements, it has enormous potential as a method of big data storage. This was the focus of the DNA working group. Problems such read and write cost, insertion and deletion errors arising in sequences, error reduction were discussed. Milenkovic gave an introductory talk describing several problems associated with whole genome, sequencing read, RNA-seq and ChiP-seq data compression, and outlined the first portable DNA-based rewritable and random access storage system.

Private information retrieval (PIR) enables a user to retrieve a data item from a database without disclosing the identity of the item retrieved, while the data itself may be public. The PIR working group considered this problem in the context of storage codes, in particular for dynamic coded storage and adversarial PIR, with some extensions to asynchronized systems, batch codes and private keyword search. Hollanti gave a tutorial overview of recent results in the area.

Age of information is a metric for status updating systems, where a monitor is interested in staying timely about the status of a source. The optimal updating strategy that minimizes the average age exists when the updating rate is constrained by limited network resources. Streaming source coding problems can be applied to the problem of age analysis. The main focus the Age & Delay working group was to introduce the age of information concept to participating coding theorists and explore potential age and delay problems in coding and storage. An adaptive arithmetic coding scheme was proposed as a potential solution to avoid huge decoding delay. Several possible delay problems in file downloading from multiple servers were discussed. Two PhD students, Zhong and Najm gave a tutorial overview of the topic.

Code-based crypto-systems are some of the very few that resist quantum-based attacks. In the case subfield subcodes such as the Goppa or Srivastava codes no successful attack is known yet. Moderate-density parity-check (MDPC) codes have been proposed for key size reduction in such crytposystems. The group identified open problems such as investigating other subfield subcodes and attacks on MDPC structured codes. An overview was presented by Bossert.

Rank-metric codes have applications in random network coding, coded-caching and in code-based cryptography. The working group focussed on maximum rank distance (MRD) codes, specifically their classifications and on algebraic methods for constructing and decoding families of them. New nontrivial classifications were obtained. Further research directions on the classification problem were identified such as adapting semi-field theory techniques and searches for codes with high symmetry. Given the known limitation of list decoders for Gabidulin codes, the group worked on adapting decoders for Gabidulin codes to recent

families of MRD codes. Sheekey presented recent results on MRD codes and described links to semifields.

A total of 44 researchers participated in the seminar across these working groups. In addition, several participants took the opportunity to collaborate with others on specific related projects. There were 16 talks in total, several related to storage of big data and others on topics such as maximum rank distance codes, chip-to-chip communication, the MDS conjecture, the SAGE computer algebra system, age of information, the edge removal problem, convolutional codes and network coding. Among the talks given were some tutorial presentations, aimed at introducing researchers to fundamentals of a related working group. The working groups focussed on identifying and addressing new and/or important open problems in the area. Age & Delay, PIR for storage codes and DNA-based storage were new topics to many participants and generated considerable interest.

2 Table of Contents

Working groups

Overview of Talks

3.1 An Explicit, Coupled-Layer Construction of a High-Rate MSR Code with Low Sub-Packetization Level, Small Field Size and All-Node Repair

P. Vijay Kumar (Indian Institute of Science -Bangalore, IN)

This talk presents an explicit construction for an $((n, k, d = n - 1), (\alpha, \beta))$ regenerating code over a finite field of size Q operating at the Minimum Storage Regeneration (MSR) point. The MSR code can be constructed to have rate k/n as close to 1 as desired, sub-packetization given by $r^{n/r}$, for $r = (n - k)$, field size no larger than n and where all code symbols can be repaired with the same minimum data download. The construction modifies a prior construction by Sasidharan et al. which required far larger field-size. A building block appearing in the construction is a scalar MDS code of block length n. The code has a simple layered structure with coupling across layers, that allows both node repair and data recovery to be carried out by making multiple calls to a decoder for the scalar MDS code. The construction can be extended to handle the case of $d < (n-1)$ under a mild restriction on the choice of helper nodes. While this work was carried out independently, there is considerable overlap with a prior construction by Ye and Barg. It is shown here that essentially the same architecture can be employed to construct MSR codes using vector binary MDS codes as building blocks in place of scalar MDS codes. The advantage here is that computations can now be carried out over a field of smaller size potentially even over the binary field as we demonstrate in an example.

3.2 A Coding Theory Inspired Approach to Computing Large Linear Transforms in Parallel/Distributed Systems

Viveck Cadambe (Pennsylvania State University – University Park, US)

In distributed and parallel computing, the performance of a computation is often limited by "stragglers", a usually small set of processors that slow down the entire computation. Job replication has been recently proposed as a method for overcoming the straggler effect. In this talk, we will describe a simple, coding-theory inspired method to compute matrix-vector multiplications that is robust to stragglers in distributed systems. In particular, our method

performs a small set of carefully designed redundant computations such that the desired matrix-vector multiplication can be obtained from any sufficiently large subset of processors.

3.3 Coding and Distributed Caching for Content Delivery

Alex Dimakis (University of Texas – Austin, US)

Smartphone and tablet proliferation is generating an enormous increase in the demand for multimedia content. Modern wireless networks cannot support this demand and its large projected growth. We explain how caching of popular content can play a fundamental role in addressing this problem and how several novel mathematical and algorithmic problems arise. We focus on the Femtocaching problem and the Coded Caching problem introduced by Maddah-Ali and Niesen and discuss how caching is very promising for giving gains that scale surprisingly well in the size of the wireless system. Unfortunately, we show that for these gains to appear, the cached files must be separated in a number of blocks that scales exponentially in the number of users and files. We show how this problem can be resolved if we modify the Maddah-Ali and Niesen scheme to place and deliver coded packets in a less optimistic way.

3.4 MDS Conjecture and the Projective Line

Iwan M. Duursma (University of Illinois – Urbana Champaign, US)

Simeon Ball famously solved the MDS conjecture for prime fields and together with Jan de Beule extended the proof to codes in characteristic p with $k \leq 2p - 2$. Ameera Chowdhury formulated the proof in a combinatorial setting of inclusion matrices. The proofs rely on Segre's tangent lemma and polynomial interpolation. We replace these two main ingredients with a new simple duality relation for the projective line, which greatly simplifies the overall proof.

3.5 The Missing Link: An Introduction to the Edge Removal Problem

Michelle Effros (CalTech – Pasadena, US)

In a world where massive networks carry massive quantities of data, it is surprising how little we know about even the most fundamental properties of the networks that we employ. This talk will introduce one such fundamental question: How much does a single link of some fixed capacity affect the capacity of the network in which it is employed? This simple question remains largely unsolved. Its solution turns out to be the key to many other fundamental questions whose solutions are also unknown. This talk will introduce the question and explore

both what is known about its solution and questions about reliability and data dependence to which it is linked.

3.6 Alphabet Size for Network Coding – Vectors Outperform Scalars

Tuvi Etzion (Technion – Haifa, IL)

A survey on the known results on the alphabet size for solutions of multicast network is given. Vector network coding solutions based on rank-metric codes and subspace codes are considered. The main result of this paper is that vector solutions can significantly reduce the required field size compared to the optimal scalar linear solution for the same multicast network. The multicast networks considered in this paper have one source with h messages and the vector solution is over a field of size q with vectors of length t. The achieved gap of the field size between the optimal scalar linear solution and the vector solution is $q^{(h-2)t^2/h+o(t)}$ for any $q \geq 2$ and any even $h \geq 4$. If $h \geq 5$ is odd, then the achieved gap of the field size is $q^{(h-3)t^2/(h-1)+o(t)}$. Previously, only a gap of constant size had been shown for networks with a very large number of messages.

These results imply the same gap of the field size between the optimal scalar linear and any scalar *nonlinear* network coding solution for multicast networks. For three messages, we also show an advantage of vector network coding, while for two messages the problem remains open. Several networks are considered, all of them generalizations and modifications of the well-known combination network. The vector network codes that are used as a solution for those networks are based on subspace codes and in particular subspace codes obtained from rank-metric codes. Some of these codes form a new family of subspace codes which poses a new interesting research problem. Finally, the exposition given in this paper suggests a sequence of related problems for future research.

3.7 An Algebraic Framework for Physical-Layer Network Coding

Elisa Gorla (Université de Neuchâtel, CH) and Alberto Ravagnani (Université de Neuchâtel, CH)

In this talk we will propose a new algebraic framework for physical-layer network coding. Our setup is based on nested-lattice-based physical-layer network coding as proposed by Nazer-Gastpar (2011), and on the algebraic approach proposed by Feng-Silva-Kschischang (2013) and Feng-Nobrega-Kschischang-Silva (2014). We will discuss the algebra which is used in our setup, and why our approach recovers and generalizes the previous ones.

3.8 Repairing Reed-Solomon Codes

Venkatesan Guruswami (Carnegie Mellon University – Pittsburgh, US)

A fundamental fact about polynomial interpolation is that k evaluations of a degree-$(k-1)$ polynomial f are sufficient to determine f. This is also necessary in a strong sense: given $k-1$ evaluations, we learn nothing about the value of f on any kth point. Motivated by the exact repair problem for Reed-Solomon codes in distributed storage systems that are ubiquitous in the time of big data, we study a variant of the polynomial interpolation problem: instead of querying entire evaluations of f (which are elements of a large field F) to recover an unknown evaluation, we are allowed to query only a few bits of evaluations.

We show that in this model, one can do significantly better than in the traditional setting, in terms of the amount of information required to determine the missing evaluation. More precisely, only $O(k)$ bits are necessary to recover the missing evaluation, and this result is optimal for linear methods.

3.9 Private Information Retrieval from Coded Data in Distributed Storage Systems

Camilla Hollanti (Aalto University, FI)

Private information retrieval (PIR) enables a user to retrieve a data item from a database without disclosing the identity of the item retrieved, while the data itself may be public. In this talk, I will give an introduction to PIR, concentrating in particular on recent advances in PIR from coded storage systems, where storage nodes may collude [1]. Some open problems will be shortly introduced, leaving further discussions to the PIR working group.

References

1 R. Tajeddine and S. El Rouayheb, "Private Information Retrieval from MDS Coded data in Distributed Storage Systems," *2016 IEEE International Symposium on Information Theory (ISIT)*, Barcelona, 2016, pp. 1411–1415. Extended version available at http://www.ece.iit.edu/~salim/PIRMDS.pdf.

3.10   Linear Systems and Convolutional Codes

Julia Lieb (Universität Würzburg, DE)

Some properties such as reachability or observability are of considerable interest for dealing with linear systems. If one focuses on systems defined over finite fields, it becomes possible to count the number of systems with a certain property and therefore, to estimate probabilities [1]. Using a criterion for the reachability and observability of interconnected systems [2], we extend these probability estimations to networks of systems, e.g. parallel connection. Since there is a correspondence between linear systems and convolutional codes [3], these considerations could be transferred to interconnected convolutional codes. Hereby, the reachability and observability of the system correlate with the minimality and non-catastrophicity of the code.

References
1   U. Helmke, J. Jordan, J. Lieb. "Probability estimates for reachability of linear systems defined over finite fields", Advances in Mathematics of Communications, Vol. 10, No. 1(2016), 63–78.
2   P. A. Fuhrmann, U. Helmke. "The Mathematics of Networks of Linear Systems," Springer, 2015.
3   J. Rosenthal, E.V. York. "BCH convolutional codes," IEEE Trans. Inform. Theory, Vol. 45, No. 6 (1999), 1833–1844.

3.11   A Mathematical Theory of Distributed Storage

Michael Luby (Qualcomm Inc. – San Diego, US)

We describe a natural and general model of distributed storage. A distributed storage system uses two fundamental mechanisms to ensure that stored objects remain recoverable as nodes fail and are replaced with new nodes: (1) storage overhead, i.e., the aggregate size of stored objects is less than the aggregate system storage capacity; (2) a repair mechanism, i.e., a repairer continually reads data stored at the nodes, performs computations on the read data, and writes the results of the computations back to the nodes.

We show lower bounds and upper bounds on trade-offs between storage overhead and repairer read rates (and write rates). The lower bounds are information-theoretic, i.e., we prove all repairers must operate above certain storage overhead/repairer read rate trade-offs. The upper bounds are algorithmic, i.e., we prove there is a repairer that operates below certain storage overhead/repairer read rate trade-offs. The lower and upper bound trade-offs are asymptotically equal as the storage overhead goes to zero.

3.12 DNA-Based Storage and Storing DNA

Olgica Milenkovic (University of Illinois – Urbana Champaign, US)

The surge of Big Data platforms has imposed new challenges to the storage community to identify extremely high volume recording media and to information theorists to propose new compression methods for nontraditional data formats. To address the first challenge, the new paradigm of DNA-based storage was recently proposed and implemented by a number of researchers. At the same time, to address the second challenge, several new initiatives for genomic read and RNA expression data compression were put forward by the National Institute of Health.

We describe several problems associated with whole genome, sequencing read, RNA-seq and ChiP-seq data compression, including parallel transform coding and large-alphabet source coding. We then proceed to outline the implementation of the first portable DNA-based rewritable and random access storage system. In this setting, we introduce new problems in prefix-synchronized, profile and Damerau-distance coding.

3.13 SageMath for Research and Teaching in Coding Theory

Johan Rosenkilde (Technical University of Denmark – Lyngby, DK)

SageMath is a powerful open-source computer-algebra system. In excess of being open-source – a clear advantage and "ethical" feature for researchers – SageMath has user-friendly interfaces that support experimentation, such as the Jupyter Notebook and SageMathCloud. For Coding Theory, Sagemath has in recent years become vastly more powerful, mainly due to the ACTIS project, which employed David Lucas full-time for two years as software developer. We give a succinct demonstration of selected capabilities that SageMath now offers to the working and teaching coding theorist.

3.14 Maximum Rank Distance Codes: Constructions, Classifications, and Applications

John Sheekey (University College Dublin, IE)

We survey the known constructions and classifications of MRD codes. We illustrate their links with algebraic structures known as semifields and quasifields, and present some open problems.

3.15 A Theory of Coding for Chip-to-Chip Communication

M. Amin Shokrollahi (EPFL – Lausanne, CH)

Modern electronic devices consist of a multitude of IC components: the processor, the memory, the RF modem and the baseband chip (in wireless devices), and the graphics processor are only some examples of components scattered throughout a device. The increase of the volume of digital data that needs to be accessed and processed by such devices calls for ever faster communication between these IC's. Faster communication, however, often translates to higher susceptibility to various types of noise, and inevitably to a higher power consumption in order to combat the noise. This increase in power consumption is, for the most part, far from linear, and cannot be easily compensated for by Moore's Law. In this talk I will give a short overview of problems encountered in chip-to-chip communication, and will advocate the use of novel coding techniques to solve those problems. I will also briefly talk about Kandou Bus, and some of the approaches the company is taking to design, implement, and market such solutions.

References
1   A. Abbasfar, "Generalized differential vector signaling," Proc. of the ICC, pp. 1-5, 2009.
2   A. Abbasfar, "Simplified receiver for use in communication systems," U.S. patent no. 8,159,375.
3   A. Amirkhany, "Multi-carrier signaling for high speed electrical links," Ph.D. Thesis, Stanford University, 2008.
4   A. Amirkhany, A. Abbasfar, V. Stojanovic, and M. A. Horowitz, "Practical limits of multi-tone signaling over high-speed backplane electrical links," Proc. Of the ICC, pp. 2693–2698, 2007.
5   A. Amirkhany, K. Kaviani, A. Abbasfar, F. Shuaeb, W. Beyene, C. Hoshino, C. Madden, K. Chang, and C. Yuan, "A 4.1pJ/b 16Gb/s Coded Differential Bidirectional Parallel Electrical Link," ISSCC 2012, pp. 138–140.
6   A. Bechtolsheim, "Moore's law and networking," North American Network Operator's Group (NANOG) meeting, 2012. (https://www.nanog.org/meetings/nanog55/presentations/Monday/Bechtolsheim.pdf)
7   D. M. Chiarulli, J. D. Bakos, J. R. Martin, and S. P. Levitan, "Area, power, and pin efficient bus transceiver using multi-bit-differential signaling," IEEE International Symposium on Circuits and Systems, pp. 1662–1665, 2005.
8   H. Cronie and A. Shokrollahi, "Orthogonal differential vector signaling," U.S. Patent application no. 12/784414.
9   H. Cronie, A. Shokrollahi, and A. Tajalli, "Methods and systems for noise resilient, pin-efficient and low-power communications with sparse signaling codes," U.S. Patent no. 8,649,445.
10 K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, N. Masuda, T.Takemoto, F. Yui, and T. Saito, "A 12.3-mW 12.5-Gb/s complete transceiver in 65-nm CMOS process," IEEE J. Solid State Circuits, 45, 2010.
11 K. Gharibdoust, A. Tajalli, and Y. Leblebici, "A 7.5 mW 7.5 Gb/s mixed NRZ/multi- tone serial- data transceiver for multi-drop memory interfaces in 40nm CMOS," ISSCC 2015, pp. 1–3, 2015.
12 M. Harwood et al., "A 12.5 Gb/s SerDes in 65nm CMOS using a Baud-Rate ADC with Digital Receiver Equalization and Clock Recovery," ISSCC 2007, pp. 436–439.

13 A. Healey and C. Morgan, "A comparison of 25 Gbps NRZ & PAM-4 Modulation used in legacy & premium backplane channels," DesignCon 2012.

14 A. Hormati, A. Shokrollahi, and R. Ulrich, "Method and apparatus for low power chip- to-chip communications with constrained ISI-ratio," U.S. Patent application no. 14/612241.

15 A. Hormati and A. Shokrollahi, "ISI tolerant signaling: a comparative study of PAM4 and ENRZ," DesignCon 2016.

16 A. Hormati, A. Tajalli, Ch. Walter, K. Gharibdoust, and A. Shokrollahi, "A Versatile Spectrum Shaping Scheme for Communicating Beyond Notches in Multi- Drop Interfaces," DesignCon 2016.

17 S. A. Ibrahim and B. Razavi, "Design requirements of 20-Gb/s serial links using multi-tone signaling," ISSCC 2009, pp. 1–4.

18 J. Lee, M. Chen, and H. Wang, "Design and Comparison of Three 20-Gb/s Backplane Transceivers for Duobinary, PAM4, and NRZ Data," JSSC, vol. 43, no. 9, 2008.

19 F. J. MacWilliams and N. J. A. Sloane. "The Theory of Error-Correcting Codes." North-Holland, 1988.

20 M. Mansuri, J. E. Jaussi, J. T. Kennedy, T. Hsueh, S. Shekhar, G. Balamurugan, F. O'Mahony, C. Roberts, R. Mooney, and B. Casper, "A scalable 0.128-to-1Tb/s 0.8- to-2.6pJ/b 64-lane parallel I/O in 32nm CMOS," ISSCC 2013, pp. 402–403.

21 A. Nazemi, Kangmin Hu, B. Catli, Delong Cui, U. Singh, T. He, Zhi Huang, Bo Zhang, A. Momtaz, and J. Cao, "A 36Gb/s PAM4 transmitter using an 8b 18GS/s DAC in 28nm CMOS," ISSCC 2015, pp. 58–60.

22 D. Oh, F. Ware, J.-H. Kim, A. Abbasfar, J. Wilson, L. Luo, R. Schmitt, and C. Yuan, "Pseudo- differential vector signaling for noise reduction in single-ended signaling systems," Designcon, 2009.

23 P. Orlik and H. Terao. Arrangements of Hyperplanes. Springer Verlag, 1992. Number 300 in Grundlehren der mathematischen Wissenschaften.

24 V. Parthasaraty, "PAM4 digital receiver performance and feasibility," IEEE 802.3bj meeting, 2012. (http://www.ieee802.org/3/bj/public/jan12/parthasarathy_01_0112.pdf)

25 D. V. Perino and J. B. Dillon, "Apparatus and method for multilevel signaling," U.S. Patent no. 6,005,895.

26 J. W. Poulton, S. Tell, and R. Palmer, "Multiwire differential signaling," University of North Carolina-Chapel Hill, 2003.

27 A. Shokrollahi, "Vector signaling codes with reduced receiver complexity," U.S. Patent application no. 14/313966.

28 A. Shokrollahi, "Vector signaling codes with high pin-efficiency and their applications to chip-to-chip communications and storage," U.S. Patent application no. 14/612252.

29 A. Shokrollahi and R. Ulrich, "Vector signaling codes with increased signal to noise characteristics," U.S. Patent application no. 62/015172.

30 A. Shokrollahi et al., "A Pin-Efficient 20.83 Gb/s/wire 0.94 pJ/bit Forwarded Clock CNRZ-5 coded Serial Link up to 12mm for MCM Packages in 28nmTechnology," ISSCC 2016.

31 A. Singh et al., "A pin- and power-efficient low-latency 8-to-12Gb/s/wire 8b8w- coded SerDes link for high-loss channels in 40nm technology," ISSCC 2014, pp. 442–443.

32 J. Poulton, W. J. Dally, , X. Chen, J. G. Eyles, Th. H. Greer, S. G. Tell, J. M. Wilson, and C.Th. Gray, "A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short- reach serial link in 28 nm CMOS for advanced packaging applications," JSSC, vol. 48, pp. 3206–3218, 2013.

33 D. Slepian, "Permutation modulation," Proc. IEEE, vol. 53, pp. 228–236, 1965.

34 D. Stauffer, J. Trinko-Mechler, M. A. Sorna, K. Dramstad, C. R. Ogilvie, A. Mohammad, and J. D. Rockrohr. "High Speed SerDes Devices and Applications." Springer Verlag, 2009.

35 V. M. Stojanovic, A. Amirkhany, and J. Zerbe, "Multi-tone system with oversampled precoders," U.S. Patent no. 7,817,743.

36 A. Tajalli, H. Cronie, and A. Shokrollahi, "Methods and circuits for efficient processing and detection of balanced codes," U.S. Patent no. 8,593,305.

37 Th. Toifl, C. Menolfi, M. Ruegg, R. Reutemann, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. L. Schmatz, "A 22-Gb/s PAM-4 receiver in 90-nm CMOS SOI technology," JSSC, vol. 41, pp. 954–965, 2006.

38 R. Ulrich, "Multilevel driver for high speed chip-to-chip communications," U.S. patent application no. 14/315,306.

39 G. A. Wiley, "Three phase and polarity encoded serial interface," U.S. Patent no. 8,472,551.

40 L. Yang and J. Armstrong, "Oversampling to reduce the effect of timing jitter on high speed OFDM systems," IEEE Communication Letters, vol. 14, pp. 196–198, 2010.

41 S. Zogopoulos and W. Namgoong, "High-Speed Single-Ended Parallel Link Based on Three-Level Differential Encoding," JSSC, Vol. 44, pp. 549–557, 2009.

3.16 Age of Information

Jing Zhong (Rutgers University – New Brunswick, US) and Elie Najm (EPFL – Lausanne, CH)

Age of information is a recently introduced timeliness metric for status updating systems, where a monitor is interested in staying timely about the status of the source which is connected to the monitor through some communication systems. In this talk, we will discuss the general methods to calculate the average age that can be applied to different service systems, and present several most recent results on age of information for a variety of queueing systems. The optimal updating strategy that minimizes the average age exists when the updating rate is constrained by limited network resources. We also connect age of information to coding theory by applying the age analysis to streaming source coding problems, where the receiver is required to decode the source message in a real-time fashion. Unlike the traditional source coding problem that focuses on the coding redundancy, the age-optimized source coding scheme balances the redundancy and higher moments of code lengths in order to minimize the queuing delay.

References

1 S. Kaul, R. D. Yates, and M. Gruteser, "Real-time status: How often should one update?" Proc. of IEEE INFOCOM, 2731–2735, 2012.

2 M. Costa, M. Codreanu and A. Ephremides, "Age of Information with Packet Management," IEEE Int. Symp. on Info. Theory (ISIT), 1583–1587, 2014.

3 E. Najm, R. Nasser, "Age of Information: the Gamma Awakening," IEEE Int. Symp. on Info. Theory (ISIT), 2574–2578, 2016.

4 J. Zhong, R. D. Yates, "Timeliness in Lossless Block Coding," Data Compression Conf. (DCC), 2016.

4   Working groups

4.1   Code-Based Cryptography

Martin Bossert (Universität Ulm, DE)

We discussed the status of coding-based cryptography. The McEliece and the Niederreiter cryptosystems are based on codes and remain the only methods. As of now these systems are some of the very few that resist quantum-based attacks. Lattice-based crypto is another such promising scheme. However, the usage of McEliece and Niederreiter with highly structured codes is broken. For example RS, RM, wild Goppa, LDPC, BCH, Gabidulin, and others cannot be used in these systems. The successful attacks are almost all based on the convolution theorem of the Fourier transform (Schur product, star product, etc.). However, the original Goppa code is not broken yet. It seems that for subfield subcodes like this Goppa code or for the Srivastava code no successful attack is known yet. Thus, an open problem is to check other subfield subcodes.

The concept of learning with errors is very similar to syndrome based methods and is viewed in the crypto community as provably secure. One recent approach is MDPC (moderate-density parity-check) codes with quasi-cyclic property for key size reduction (key size being one of the major criticisms of code-based cryptography). Possible attacks were discussed using the quasi-cyclicity and decoding by calculating low-weight codewords. Thus another open problem is to check possible attacks for MDPC codes. A further recent approach is to improve the permutation in the McEliece cryptosystem such that even structured codes can be used (not published yet).

A common argument against the use of code-based cryptography has been that signatures are not possible; however, this is not true since five methods are known and two of them are not broken yet ([1], [2]).

References
1   Nicolas T. Courtois, Matthieu Finiasz, and Nicolas Sendrier. "How to Achieve a McEliece-Based Digital Signature Scheme", pages 157–174. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.
2   G. Kabatianskii, E. Krouk, and B. Smeets. "A Digital Signature Scheme Based on Random Error-Correcting Codes", pages 161–167. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.

4.2   Distributed Storage & Index Coding

Viveck Cadambe (Pennsylvania State University – University Park, US)

In the big data era, the amount of data that is being stored is scaling at a rapid pace. This makes efficient data storage an important problem that inspires several lines of scientific research. During the seminar, several discussions were conducted on the theme of using classical and new techniques from coding theory to store/compute data efficiently in distributed storage systems. We summarize the discussions below.

1. Open Problems in Codes for Distributed Storage: Several open problems requiring design of new codes for storing data in distributed storage systems, and study of their information-theoretic fundamental limits were discussed. A list of specific open problems that are relevant to wide variety of applications was composed. Follow-up discussions included a summary of recent progress in some of the listed problems and related challenges. The topics discussed include:
   a. design of codes with good/optimal repair bandwidth, and fundamental trade-offs between storage cost and communication cost,
   b. fundamental limits of locality and availability of codes,
   c. formulations that are applicable to caching and content distribution networks, and
   d. connections between fundamental limits storage/caching and the index coding problem in network information theory.
2. A Mathematical Theory for Distributed Storage: Details related to a recently developed theoretical framework that studies the long term reliability of a distributed storage system were presented by its author (Michael Luby). The model, unlike others discussed in the working group, studies of the evolution of the storage system over time, estimating the likelihood of data loss. A numerical simulation of the mathematical model was also demonstrated. Follow up discussions included comparisons and connections between the exact and functional repair problem in coding theory and the presented framework.
3. Coding Theory Inspired Methods for Parallel/Distributed Computing: Mathematical details related to a coding theory inspired method for implementing linear transforms in parallel/distributed systems, were presented by its author (Viveck Cadambe). The utility of coding theory was to incorporate redundancy in the computations to make it robust to slow nodes (stragglers). The discussion also involved connections to classical topics in coding theory, such as the role of the field size of the operation and the decoding complexity.
4. Consistency Issues in Distributed Storage: In a brief discussion conducted jointly with the Private Information Retrieval for Storage Codes working group, consistency related issues in storage of multiple versions of data were discussed. A discussion of relevant applications and a related coding theoretic formulation ensued.

References

1 M. Ye and A. Barg, "Explicit constructions of high-rate MDS array codes with optimal repair bandwidth." arXiv:1604.00454 (2016).

2 B. Sasidharan, M. Vajha, and P. V. Kumar, "An Explicit Coupled-Layer Construction of a High-Rate MSR Code with Low Sub-Packetization Level, Small Field Size and All-Node Repair." http://arxiv.org/pdf/1607.07335.pdf

3 N. Raviv, N. Silberstein and T. Etzion, "Constructions of high-rate minimum storage regenerating codes over small fields," 2016 IEEE International Symposium on Information Theory (ISIT), Barcelona, 2016, pp. 61–65.

4 B. and P. V. Kumar, "High-rate regenerating codes through layering," 2013 IEEE International Symposium on Information Theory Proc, pp. 1611–1615.

5 Z. Wang, A.G. Dimakis, and J. Bruck, "Rebuilding for array codes in distributed storage systems," 2010 IEEE Globecom Workshops, 2010.

6 T. Westerbäck, R. Freij-Hollanti, and C. Hollanti, "Applications of polymatroid theory to distributed storage systems," 53rd Annual Allerton Conference on Communication, Control, and Computing, 2015.

7 T. Westerbäck, R. Freij-Hollanti, T. Ernvall and C. Hollanti, "On the Combinatorics of Locally Repairable Codes via Matroid Theory," in IEEE Transactions on Information Theory, vol. 62, no. 10, pp. 5296–5315, Oct. 2016.

8 M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," IEEE Transactions on Information Theory 60.5 (2014): 2856–2867.

9 K. Shanmugam, M. Ji, A. M. Tulino, J. Llorca, and A. G. Dimakis, "Finite length analysis of caching-aided coded multicasting." 52nd Annual Allerton Conference on Communication, Control, and Computing, 2014.

10 N. Golrezaei, A. Molisch, A. G. Dimakis, and G. Caire, "Femtocaching and device-to-device collaboration: A new architecture for wireless video distribution." IEEE Communications Magazine 51.4 (2013): 142–149.

11 A. Golovnev, O. Regev, and O. Weinstein, "The Minrank of Random Graphs." arXiv preprint arXiv:1607.04842 (2016).

12 Z. Wang and V. R. Cadambe, "Multi-Version Coding-An Information Theoretic Perspective of Consistent Distributed Storage." arXiv preprint arXiv:1506.00684 (2015).

13 K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, K. Ramchandran, "Speeding up distributed machine learning using codes." arXiv preprint arXiv:1512.02673 (2015).

14 S. Dutta, V. R. Cadambe, and P. Grover, "Short-Dot: Computing Large Linear Transforms Distributedly Using Coded Short Dot Products," to appear in Proceedings of The Thirtieth Annual Conference on Neural Information Processing Systems (NIPS), 2016.

15 I. Tamo and A. Barg, "A family of optimal locally recoverable codes," IEEE Transactions on Information Theory 60.8 (2014): 4661–4676.

16 V. Guruswami and M. Wootters, "Repairing Reed-Solomon Codes," Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, 2016.

4.3 Private Information Retrieval for Storage Codes

Camilla Hollanti (Aalto University, FI)

The PIR working group discussed various topics related to (dynamic) coded storage and (adversarial) PIR, with some extensions to asynchronized systems and private keyword search. Recent results were reviewed related to matroids (Hollanti), batch codes (Skachek), and PIR, and this part of the discussion is likely to lead to a joint publication (Hollanti and Skachek, jointly with Thomas Westerback and Ragnar Freij-Hollanti).

Discussions have been continued via email (e.g., between Etzion and Hollanti), as well as during some visits that were prompted during this seminar (e.g., Eimear Byrne is visiting Hollanti and Greferath at Aalto University Sep.-Dec. 2016). During this seminar and working group discussions, open Ph.D. positions related to coded storage and PIR in the ANTA group at Aalto University were advertised, and a new student was found (from University of Neuchatel, M.Sc. advised by Elisa Gorla) and who started at Aalto University in October 2016.

4.4   DNA-Based Storage

Olgica Milenkovic (University of Illinois – Urbana Champaign, US)

The DNA working group focused on the problem of using DNA as information storage media. In this setting, information is encoded into strings over the alphabet of nucleobase symbols $\{A, T, G, C\}$. Because DNA may be easily replicated and a massive amount of information stored reliably with minimal space requirements, it has enormous potential as a method of storage in the time of big data. A current obstacle for practical implementations is that it remains costly to write information to and read from strands of DNA, though this cost is going down with time.

Recently, technologies for sequencing (reading) strands of DNA have improved dramatically, and will hopefully continue to become more reliable and accessible. In particular, nanopore technologies allow researchers to sequence longer strands of DNA than was previously possible. DNA is read as it passes through a nano-scale hole called a nanopore. Any given strand of DNA is read multiple times, until there is sufficient data for decoding. The errors introduced by this sequencing method include insertion and deletion errors: while being read, a strand may oscillate, doubling the sequence of read symbols back on itself. That is, if $\overline{s} \in \mathbb{F}_4^\ell$ is a sequence of ℓ symbols, and $\overline{s}^R$ is the same sequence reversed, then we may have insertion error patterns of the form $\overline{s}t\overline{s}^R$, where $t \in \mathbb{F}_4$ is a transition symbol before the reversal of the strand. Thus, when designing codes for use in this application, we seek to avoid sequences which contain both a sequence of length ℓ, for some ℓ, and its reversal. Other constraints to consider include having an appropriate balance of bases, creating unique addresses which identify each strand of DNA, and avoiding particular subsequences which are biologically difficult to synthesize (write). After discussing the background of the problem, we briefly discussed several relevant open problems, which are given below.

- How may sequences exist in $\mathbb{F}_4^n$ with the property that there are no "oscillating" subsequences of length ℓ? That is, how many sequences avoid the pattern $\overline{s}t\overline{s}^R$ for $\overline{s} \in \mathbb{F}_4^\ell$?
- Can we mathematically justify methods for error reduction that seem to work in practice? For example: if we balance the nucleobase content in subblocks of a given length, we can avoid deletions when sequencing. However, this is an observed phenomenon, rather than a mathematically-developed strategy.

4.5   Rank-Metric Codes

John Sheekey (University College Dublin, IE)

Rank-metric codes are codes using matrices over a field as codewords, with the distance function determined by the rank of the difference of two matrices. The topic is experiencing a resurgence as of late, due in part to its applications in random network coding, and potential applications in code-based cryptography.

The Rank-Metric Codes working group began with a group discussion about the background literature and potential directions for research. The group then split into two main

subgroups; one focussing on computational classifications of MRD codes, and the other on algebraic methods for constructing and decoding families of codes.

The computational subgroup reviewed the known results, and discussed the feasibility of performing exhaustive searches for new parameters. Using a combination of pre-existing algorithms designed by some members, and incorporating some ideas from semifield theory, new nontrivial classifications were obtained. It was felt that there are many more parameters that are within reach with current computational power, particularly if semifield techniques can be adapted to the case of rectangular matrices. The group also considered searching for codes with prescribed automorphism groups.

The algebraic subgroup discussed the known results for decoding of Gabidulin codes, and discussed the possibility of extending these to the recently introduced family known as twisted Gabidulin codes. As the construction using linearized polynomials is similar for both families, initial investigations looked promising, and a coarse first estimate was obtained. However it was felt that there was much room for improvement by exploiting further properties of these codes, which will require a more in-depth analysis. The group also discussed potential constructions for new MRD codes similar to Gabidulin codes, as well as the need for codes with much less structure (for the purposes of avoiding weaknesses in the associated code-based cryptographic systems).

4.6 Age & Delay of Information

Jing Zhong (Rutgers University – New Brunswick, US)

The main focus of our working group was to introduce the age of information concept to participating coding theorists and explore potential age and delay problems in coding and storage. A brief introduction of status updating age and its connection to streaming source coding was presented and discussed during the first meeting. The disadvantage of using arithmetic coding instead of fixed-to-variable block coding due to large decoding delay was discussed, and an adaptive arithmetic coding scheme for fixed length message was proposed as a potential solution to avoid huge decoding delay. In the group meetings of the last two days, we moved to a recent topic noticed by the industry, which is the application of rateless code to multipath packet transmissions in real-time video streaming. A tutorial about spatially coupled LDPC codes was given, and the comparison between spatially coupled LDPC codes and traditional punctured LDPC codes in terms of delay and decoding probability was reviewed. At the end of seminar, we discovered several possible delay problems in file downloading from multiple servers.

Participants

- Iryna Andriyanova
  University of Cergy-Pontoise, FR
- Allison Beemer
  Univ. of Nebraska – Lincoln, US
- Jessalyn Bolkema
  Univ. of Nebraska – Lincoln, US
- Martin Bossert
  Universität Ulm, DE
- Michael Braun
  Hochschule Darmstadt, DE
- Eimear Byrne
  University College Dublin, IE
- Viveck Cadambe
  Pennsylvania State University –
  University Park, US
- Gerard Cohen
  Telecom Paris Tech, FR
- Alex Dimakis
  University of Texas – Austin, US
- Iwan M. Duursma
  University of Illinois –
  Urbana Champaign, US
- Michelle Effros
  CalTech – Pasadena, US
- Rafah El-Khatib
  EPFL – Lausanne, CH
- Tuvi Etzion
  Technion – Haifa, IL
- Christina Fragouli
  University of California at Los
  Angeles, US
- Heide Gluesing-Luerssen
  University of Kentucky, US

- Elisa Gorla
  Université de Neuchâtel, CH
- Marcus Greferath
  Aalto University, FI
- Venkatesan Guruswami
  Carnegie Mellon University –
  Pittsburgh, US
- Daniel Heinlein
  Universität Bayreuth, DE
- Camilla Hollanti
  Aalto University, FI
- Anna-Lena
  Horlemann-Trautmann
  EPFL – Lausanne, CH
- Christine A. Kelley
  Univ. of Nebraska – Lincoln, US
- P. Vijay Kumar
  Indian Institute of Science –
  Bangalore, IN
- Julia Lieb
  Universität Würzburg, DE
- Hans-Andrea Loeliger
  ETH – Zürich, CH
- Michael Luby
  Qualcomm Inc. – San Diego, US
- Felice Manganiello
  Clemson University, US
- Carolyn Mayer
  Univ. of Nebraska – Lincoln, US
- Sihem Mesnager
  University of Paris VIII, Telecom
  Panis Tech, FR
- Olgica Milenkovic
  University of Illinois –
  Urbana Champaign, US

- Elie Najm
  EPFL – Lausanne, CH
- Alessandro Neri
  Universität Zürich, CH
- Mario Osvin Pavcevic
  University of Zagreb, HR
- Sven Puchinger
  Universität Ulm, DE
- Tovohery Randrianarisoa
  Universität Zürich, CH
- Alberto Ravagnani
  Université de Neuchâtel, CH
- Johan Rosenkilde
  Technical Univ. of Denmark –
  Lyngby, DK
- Joachim Rosenthal
  Universität Zürich, CH
- John Sheekey
  University College Dublin, IE
- M. Amin Shokrollahi
  EPFL – Lausanne, CH
- Vitaly Skachek
  University of Tartu, EE
- Emina Soljanin
  Rutgers Univ. – Piscataway, US
- Pascal Vontobel
  The Chinese University of Hong
  Kong, HK
- Jing Zhong
  Rutgers University –
  New Brunswick, US