Guest Editorial Communication Methodologies for the Next-Generation Storage Systems

Lara Dolecek, Mario Blaum, Jehoshua Bruck, Anxiao (Andrew) Jiang, Kannan Ramchandran, and Bane Vasic

INFORMATION is being generated at a rapid pace. To sustain this information revolution, next-generation data storage must be more practical, economical, and reliable than ever before.

There are presently more than half a million data centers world-wide and this number is expected to rise even further. Companies like Google, Facebook, and Microsoft have already invested billions of dollars into their data center infrastructure. To make such storage systems economically viable, new solutions based on efficient and reliable data storage must now be invented. A key enabling technology at both device- and network-level of reliable storage is fundamentally rooted in mathematical models, algorithms, and techniques for data communication. The goal of this special issue is to offer a publication venue to advance the state of the art in communication methods for emerging storage systems.

At the device physics level, a wide range of new technologies are being actively investigated. Solid state drives (NAND flash based and DRAM based) are rapidly gaining importance due to their energy efficiency at high access speeds. Such non-volatile memories (NVMs) have already found use in consumer electronics and in enterprise storage, but still lag hard disk drives in terms of the cost. While very promising, these technologies present new challenges in lifetime and reliability, which will require non-conventional solutions.

At the system level, massive distributed storage networks, data centers, and cloud storage systems desperately need new methods to improve storage efficiency and data transfer speeds. Computer architectures, file systems, data classification and search, security, and interaction with data networks have to be revisited in the light of the new technologies, and will greatly influence the way new storage systems are built.

The last JSAC issue on related topics was *Data Communication Techniques for Storage Channels and Networks* and was published in February 2010. As demonstrated by the scope of the papers included in the present issue, there have been several exciting developments since the last related issue, ranging from device-level characterization and coding to NVM-inspired algebraic methods for non-conventional channel models to novel repair techniques for large-scale distributed storage systems.

This issue consists of 22 high-caliber papers with contributions from both academia and industry. The papers are organized into the following six sections: (i) Channel Modeling and Signal Processing Algorithms for Emerging Memory Technologies, (ii) Error Control Coding Techniques for Flash Memories, (iii) Algebraic Methods with Applications to Non-Volatile Memories, (iv) Polar Codes with Application to Storage, (v) Performance Limits of Storage Systems, and (vi) Codes for Distributed Network Storage.

It is known that NVMs have special properties that are not well-captured by the conventional communication models. The first group of papers covers several timely topics related to channel modeling, signal detection, and interference, and threshold evaluation in non-volatile memories. Results of this type are particularly useful in emerging practical systems employing dense memories.

Motivated by physical characteristics of flash, phase change memories (PCM), and resistive RAM, the paper "Rewritable Storage Channels with Hidden State," by R. Venkataramanan, S. Tatikonda, L. A. Lastras-Montaño, and M. Franceschini, studies a novel storage model that captures the changes in reliability as a function of the write input, past number of writes, and random noise, as well as the variability across cells (that is often caused by fabrication variability). The paper offers a lower bound on the capacity of a general rewritable channel with a hidden state, with a constraint on the average number of writes per cell. Additionally, the paper presents an asymptotically optimal coding scheme for uniformly distributed write noise and hidden cell states.

Multi-level flash devices offer increased densities over conventional flash memories that store a single bit per cell. However, inter-cell interference (ICI) has emerged as a serious issue in multi-level flash since the narrowing of the voltage ranges associated with different levels results in reduced

L. Dolecek is with the Electrical Engineering Department, University of California, Los Angeles (UCLA), Los Angeles, CA 90095, USA (e-mail: dolecek@ee.ucla.edu).

M. Blaum is with IBM Almaden Research Center, San Jose, CA 95125, USA (e-mail:mblaum@us.ibm.com).

J. Bruck is with the Electrical Engineering Department, California Institute of Technology, Pasadena, CA 91125, USA (e-mail:bruck@caltech.edu).

A. Jiang is with the Computer Science and Engineering Department and the Electrical and Computer Engineering Department of Texas A&M University, College Station, TX 77843, USA (e-mail:ajiang@cse.tamu.edu).

K. Ramchandran is with the Electrical Engineering and Computer Science Department at the University of California at Berkeley, Berkeley, CA 94720, USA (e-mail:kannanr@eecs.berkeley.edu).

B. Vasic is with the Department of Electrical and Computer Engineering, the University of Arizona, Tucson, AZ 85721, USA (e-mail:vasic@ece.arizona.edu).

S. Low is the J-SAC Board Representative for this issue of IEEE Journal on Selected Areas in Communications.

Digital Object Identifier 10.1109/JSAC.2014.140501

0733-8716/14/$31.00 © 2014 IEEE

robustness to errors associated with programming adjacent cells. The paper "Optimal Detector for Multilevel NAND Flash Memory Channels with Intercell Interference," by M. Asadi, X. Huang, A. Kavcic, and N. P. Santhanam, investigates a timely topic of designing an optimal detector for multi-level flash in the presence of ICI. This paper is the first to investigate a soft-output detector without channel output quantization for multi-level flash. The paper considers both 1-D and 2-D channel models and provides optimal detectors. Additionally, the paper offers low-complexity, practical methods for combatting small levels of ICI.

Given the importance of addressing inter-cell interference (ICI) in multi-level flash channels, development of codes that mitigate the effects of ICI has emerged as a key research topic. Constrained coding is a promising technique for preventing harmful ICI patterns. The paper "Constrained Codes that Mitigate Inter-Cell Interference in Read/Write Cycles for Flash Memories," by M. Qin, E. Yaakobi, and P. H. Siegel, further explores this promising approach. The paper studies balanced constrained codes that forbid particularly harmful ICI patterns, and establishes the asymptotic rate of such codes. It also studies the integration of constrained codes with write-once memory (WOM) codes, an emergent technique for rewriting data in flash memories under limited program and erase (P/E) cycles. The characterization of the sum-capacity for the resultant combined coding scheme is also presented.

Another important problem that arises in flash memories, and in multi-level memories in particular, is that of correctly yet efficiently identifying threshold levels. The paper "Adaptive Threshold Read Algorithms in Multi-Level Non-Volatile Memories," by E. Hemo and Y. Cassuto, studies the problem of using thresholds to identify different (integer) levels in a number of cells, while minimizing the number of times the thresholds are used. Among other results, the authors present one-dimensional and two-dimensional search algorithms, show that the two-dimensional algorithm is superior, and provide performance bounds.

Flash memories are being broadly considered for storage applications that demand extremely high reliability levels. This now necessitates the development of high-performance channel codes that can meet stringent reliability requirements. The second group of papers explores new coding methods that are well-suited for emerging flash memories, including LDPC codes, Raptor codes and Reed-Solomon codes.

In the paper "Concatenated Raptor Codes in NAND Flash Memory," by G. Yu and J. Moon, the authors propose two concatenated coding schemes based on fixed-rate Raptor codes for NAND flash memory error correction systems. One method is designed for off-line recovery of uncorrectable pages and the other method is for a normal read-mode page-error correction. The authors describe efficient encoders and decoders necessary to accommodate large length of the Raptor code, and they show significant performance gains. Both off-line and on-line schemes use hard-decisions from the channel; the inner decoder (of a BCH or product code) converts them to packet-erasures handled then by the Raptor code.

LDPC codes are known to perform exceptionally well in a number of data communication and storage applications. They are actively being pursued for use in NVMs. A discernible feature of dense flash memories is that the error rates are not equally distributed across different bits stored in a cell nor across different cell levels. The paper "Structured Bit-Interleaved LDPC Codes for MLC Flash," by K. Haymaker and C. A. Kelley, explores the use of LDPC codes to address the difference in raw error rates across the two bit types (i.e., the most significant bit and the least significant bit in a two-bit cell) in multiple level cell (MLC) flash that store two bits per cell. The paper presents code design techniques for both binary and non-binary LDPC codes that take advantage of this error asymmetry. The code designs employ bit-interleaved coded modulation and varying node degrees for improved decoding thresholds.

The paper "Enhanced Precision Through Multiple Reads for LDPC Decoding in Flash Memories," by J. Wang, K. Vakilinia, T-Y. Chen, T. Courtade, G. Dong, T. Zhang, H. Shankar and R. Wesel, presents a method for optimizing the word-line voltages for multiple reads in a flash channel based on maximizing mutual information of the underlying quantized channel. The authors provide examples where maximizing mutual information minimizes the frame error rate, and examples where maximizing mutual information does not necessarily minimize the frame error rate. The authors also show an important property: a channel code (e.g., an LDPC code) that is the best for a given precision is not necessarily the best code for a different precision choice.

The topic of extracting soft information is also explored in the next paper. It is known that LDPC codes achieve good decoding results when soft information is provided. However, obtaining soft information requires multiple flash reads resulting in degraded performance. An interesting solution to the problem is proposed in the paper "Classification Codes for Soft Information Generation from Hard Flash Reads," by M. N. Kaynak, P. R. Khayat and S. Parthasarathy. The idea is to augment the redundancy by using tensor product codes with error detection capability. The indication of errors provided by the detection mechanism is then used to create soft information from hard reads for use by the LDPC decoder. Simulation results demonstrate that with relatively low overhead the new scheme improves the quality of the results of the LDPC decoding.

The paper "Coded Modulation Using Lattices and Reed-Solomon Codes, with Applications to Flash Memories," by B. M. Kurkoski, describes a coded modulation scheme where low-dimensional lattices are used to define a constellation, and Reed-Solomon (RS) codes are used for error correction. The approach is based on using low-dimensional lattices rather than standard one- and two-dimensional constellations. Such lattices have better distance properties while permitting efficient soft-input decoding. The proposed scheme combines a lattice code with an RS code so that the RS code corrects errors produced by the decoder of the lattice code. The author compares the lattice-RS scheme with a BCH-coded scheme and demonstrates performance gains between 1-1.8 dB on the additive white Gaussian noise channel.

Due to their special physical properties, emerging NVM technologies have also recently inspired several new mathematical formulations that have in turn helped expand the repertoire of available algebraic and combinatorial coding-

theoretic techniques. The third group of papers discusses such novel, NVM-driven algebraic and combinatorial methods, including new data representation approaches and the use of non-standard distance metrics and channel models.

Rank modulation is a recently proposed technique for data representation wherein the information in a block of cells is represented by the relative rankings of charge levels in the cells rather than the absolute values of the cells themselves. The advantage of rank modulation is that it naturally circumvents several problematic issues arising in flash memories (such as write asymmetry and costly erases). Recent work has explored a use of constant-weight codes (i.e., codes in which all the codewords over bipolar symbols have the same weight) for local rank modulation. In the paper "Constant Weight Codes: An Approach Based on Knuth's Balancing Method," by V. Skachek and K. A. S. Immink, the authors take a further step in the study of constant weight codes. They extend the previous result by Knuth for the case of the balanced constant weight codes (corresponding to the weight zero) to a more general scenario, present new code constructions, and discuss code redundancies in this generalized setting.

In the rank modulation set-up, the errors due to inadvertent changes in charge are represented as certain transformations from one permutation to another permutation, where the permutation denotes the relative cell ranking in a block. While Kendall-tau distance previously studied in the NVM literature captures the behavior of errors with small magnitudes well, it may not be sufficient in the case where charge drops exceed the difference between adjacent levels. In this scenario, it is worthwhile to investigate the length of the largest common subsequence of two permutations. Ulam distance conveniently captures this quantity. The paper "Multipermutation Codes in the Ulam Metric for Nonvolatile Memories," by Farzad Farnoud (Hassanzadeh) and Olgica Milenkovic, considers multipermutation codes as a generalization of permutation codes where each message is encoded as a permutation of the elements of a multiset. The paper considers multipermutation codes both in the Ulam and Hamming metrics, provides bounds on code cardinalities, code constructions, and offers accompanying decoders.

A method for enhancing the lifetime of a flash device is to maximize the number of writes before a costly erase becomes necessary. This task is well captured by the write-once memory (WOM) set up wherein one seeks to maximize the amount of stored information over a prescribed number of writes such that the newly stored information does not decrease the values previously stored in the memory. Coding for WOM has received substantial attention in recent literature due to its applicability in flash systems although most of the existing literature has focused on the binary setting. With the advent of the dense flash technologies storing multiple bits per cell, it is an interesting problem to construct non-binary WOM codes. The paper "Lattice-Based WOM Codes for Multilevel Flash Memories," by A. Bhatia, M. Qin, A. R. Iyengar, B. M. Kurkoski, and P. H. Siegel, studies an important geometric problem of how to partition an n-dimensional space into disjoint regions such that these regions offer upper bounds on the codebook sizes for multiple-write non-binary WOM codes. For the case of two cells, the paper also provides a scheme to effectively design codes that fit the optimal shaping regions.

Recent years have witnessed an increased interest in theory and application of polar codes due to the fact that they achieve capacity of symmetric memoryless channels while also having explicit construction. Their large stopping distance and low complexity implementation make them attractive candidates for data storage application. However, this low-complexity decoding comes with a penalty of a large decoding delay due to sequential nature of the successive cancellation algorithm. Also, the performance for short to moderate lengths polar codes is inferior to that of other codes, such as LDPC codes. Recently, a number of methods have been proposed to address these drawbacks. The next two papers in this Special Issue introduce fast decoding algorithms of polar codes.

The paper "Fast Polar Decoders: Algorithm and Implementation," by G. Sarkis, P. Giard, A. Vardy, C. Thibeault, and W. J. Gross, describes an algorithm, hardware architecture, and an FPGA implementation of a very fast (gigabit-per-second) polar decoder. The proposed architecture circumvents the problem of the inherently sequential decoding process which is the main bottleneck affecting the data throughput. The authors use the properties of high-rate codes to accelerate speed. The proposed architecture outperforms the existing implementations of polar code decoders.

The state-of-the-art soft-output decoder for polar codes is based on belief propagation and performs well at the cost of high processing and storage requirements. The paper "Low-Complexity Soft-Output Decoding of Polar Codes," by U. U. Fayyaz and J. R. Barry, proposes a low-complexity version of soft-output decoding of polar codes which combines successive cancellation and belief propagation. The proposed decoder outperforms the existing ones while significantly reducing processing and storage requirements. For moderate length polar codes on the dicode channel the authors show an example where the proposed decoder complexity is only 4% of the total complexity of the belief propagation decoder, and an example of a long code which requires less than half the memory required by a belief propagation decoder.

The remaining collection of papers in the Special Issue deal with the system-level reliability. Distributed storage systems provide reliable access to large amounts of data by spreading the redundancy in form of coding over several individually unreliable storage nodes. Due to its many applications, distributed storage is currently a thriving research area. The fifth group of papers discusses various limits and trade-offs in distributed storage.

It had been an open question whether there exists a non-vanishing gap between the optimal bandwidth-storage trade-off of the functional-repair regenerating codes and that of the exact-repair regenerating codes. The paper "Characterizing the Rate Region of the (4,3,3) Exact-Repair Regenerating Codes," by C. Tian, proves the existence of the gap by showing a complete characterization of the rate region of the (4,3,3) exact-repair regenerating codes. The paper explores a new computer-aided approach for proving the information-theoretic rate region, and uses problem-specific properties to reduce the complexity of such an approach. An explicit binary code for the region of interest is also constructed.

The paper "Update Efficiency and Local Repairability Limits for Capacity Approaching Codes," by A. Mazumdar, V. Chandar, and G. W. Wornell, deals with a problem that has drawn attention in recent coding literature: given a distributed system with many nodes, design codes that can retrieve a number of failed nodes from the surviving ones. Maximum distance separable (MDS) codes are a popular approach to address this problem. Interestingly, in practice, the most common failure consists of only one node. In that case, it is desirable to recover the failed node from a minimal number of surviving ones (the authors call this the local property) and, at the same time, the updates should affect as few nodes as possible. The trade-offs between the correcting capability of the code and the local property are thoroughly studied in this paper.

A related problem is explored in the next paper. While codes are typically used for managing the reliability of distributed storage systems, the added information redundancy can be employed for improving read performance. In particular, the paper "On the Delay-Storage Trade-Off in Content Download from Coded Distributed Storage Systems," by G. Joshi, Y. Liu and E. Soljanin, considers a distributed storage system where content is encoded using an (n, k) MDS code. In a download operation, the information is read in parallel from the n disks and is reconstructed using the first k information packets that are available. The paper provides rigorous analysis results of the expected download time using a fork-join queuing model and demonstrates potential performance benefits of the proposed download scheme.

Continuing on with the theme of distributed storage, the last group of papers is focused on various coding techniques for networked storage and on their practical properties including cost of repair and locality.

The paper "A Repair Framework for Scalar MDS Codes," by K. Shanmugam, D. S. Papailiopoulos, A. G. Dimakis, and G. Caire, presents efficient repair strategies for distributed storage when the codes are designed over some higher extension field. These codes are thus scalar in the extension field but are vector codes in the base field. Exploiting this idea, repair strategies for the case of $(5, 3)$- and $(6, 4)$-RS codes are presented. Strategies for more general parameters that outperform naïve schemes of repair are also discussed.

A new class of RAID-6 type of codes called MDR codes are introduced in the paper "MDR Codes: A New Class of RAID-6 Codes with Optimal Rebuilding and Encoding," by Y. Wang, X. Yin and X. Wang. The proposed MDR codes have MDS property. Additionally, they have the property that in order to reconstruct one single failure, they need access to half of the nodes, achieved by judiciously using the Q parity. In traditional RAID systems, all or nearly all of the surviving nodes are used by emplying only the P parity. In recent literature, it was shown that the recovery time in some known codes, like EVENODD or RDP codes, can be improved to 75% of the time required with traditional decoding using the P parity. The theoretical estimates indicate that the improvement can be on the order of 50%. For example, recently proposed so-called zig-zag codes meet this theoretical estimate. The codes presented here achieve the same goal, but they are binary codes, as opposed to the zig-zag codes, which require a field of size at least 3. The trade-off is that the optimal updating property of the zig-zag codes is sacrificed in order to use a binary field.

The paper "Optimal Locally Repairable Linear Codes," by W. Song, S. H. Dau, C. Yuen, and T. J. Li, also deals with the problem of locality. Locality is defined as (r, δ) locality, which means that given an $[n, k, d]$ linear code, any symbol i may be reconstructed from a set S_i of size $|S_i|$ such that $|S_i| - \delta + 1 \leq r$. The authors distinguish between information locality (in which only the information symbols have locality) and all-symbol locality (in which the local reconstruction does not distinguish between information and parity symbols). The authors present a bound for locally repairable linear codes, and they call such codes optimal. Then, they proceed to study properties of these optimal codes.

Locality property is also exploited in the paper "Repairable Fountain Codes," by M. Asteris and A. G. Dimakis. The paper studies fountain codes with nice locality properties that allow for local decoding. The approach undertaken in this work is to investigate systematic, random codes of fixed degree per coding symbol. Additionally, the paper offers a new mathematical technique for analyzing the rank of sparse random matrices with specific structure over finite fields that may be of broader interest.

Lastly, coding methods for cloud storage systems that minimize the cost of repairing failed nodes is studied in the paper "Irregular Fractional Repetition Code Optimization for Heterogeneous Cloud Storage," by Q. Yu, C. W. Sung, and T. H. Chan. Contrary to homogeneous systems, the authors assume different costs of repair per node and per edge. The idea is to provide a coding scheme that minimizes the storage cost of all nodes as well as the cost of all possible node failures. At the same time, there should be a certain number of subsets of the nodes that allow for total reconstruction in case the remaining nodes have been lost. This problem is a very complex one, and the authors propose some constructions towards its solution. They also provide some simulations showing the advantages of their approach, that is based on what they call MDS-IFR codes, which are concatenations of MDS codes with Irregular Fractional Repetition (IFR) codes.

We wish to thank all the reviewers for providing detailed and expert reviews that have significantly aided in paper selection process and helped improve the final content of the accepted papers. We also hope that these expert reviews were useful to authors whose papers could not be accepted. Additionally, we wish to thank the authors for submitting top quality papers that have collectively made this special issue a timely and a comprehensive reference. Finally, we would like to express our gratitude to the JSAC team: Professor Steven Low who has served as our mentor, Dr. Martha Steenstrup, the current Editor-in-Chief, Ms. Laurel Greenidge the current Executive Editor, and Ms. Sue Lange, the current Production Manager, for their time and efforts in bringing together this special issue. We hope that the present issue will serve as a lasting reference to researchers working in this timely area.

Lara Dolecek (S'05–M'10–SM'13) is an Assistant Professor with the Electrical Engineering Department at the University of California, Los Angeles (UCLA). She holds a B.S. (with honors), M.S. and Ph.D. degrees in Electrical Engineering and Computer Sciences, as well as an M.A. degree in Statistics, all from the University of California, Berkeley. She received the 2007 David J. Sakrison Memorial Prize for the most outstanding doctoral research in the Department of Electrical Engineering and Computer Sciences at UC Berkeley. Prior to joining UCLA, she was a postdoctoral researcher with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. She received Northrop Grumman Excellence in Teaching Award, Intel Early Career Faculty Award, University of California Faculty Development Award, and Okawa Research Grant all in 2013, NSF CAREER Award in 2012, and Hellman Fellowship Award in 2011. She is an Associate Editor for Coding Theory for IEEE Transactions on Communications and for IEEE Communication Letters and is the lead guest editor for IEEE JSAC special issue on emerging data storage. Her research interests span coding and information theory, graphical models, statistical algorithms, and computational methods, with applications to emerging systems for data storage, processing, and communication. She is a Senior Member of IEEE since 2013.

Mario Blaum (F'00) was born in Buenos Aires, Argentina, in 1951. He received the degree of Licenciado from the University of Buenos Aires in 1977, the M. Sc. degree from the Israel Institute of Technology in 1981 and the Ph. D. degree from the California Institute of Technology in 1984, all these degrees in Mathematics. From January to June, 1985, he was a Research Fellow at the Department of Electrical Engineering at Caltech. In August, 1985, he joined the IBM Research Division at the Almaden Research Center. In January, 2003, the division was transferred to Hitachi Global Storage Technologies, where he was a Research Staff Member until February 2009. In August 2009 he rejoined the IBM Almaden Research Center. From September 1990 to September 1991 he was a Consulting Professor at Stanford University, where he taught the course in Error-Correcting Codes, which he did again in 1997. In 2008 and 2009 he was a Lecturer at Santa Clara University, where he taught Error-Correcting Codes. Since 2005 he is an Academic Adviser at the Universidad Complutense of Madrid. He was Associate Editor of IEEE Transactions on Information Theory from April 2009 to April 2012. Dr. Blaum is author and co-author of numerous papers and patents related to error-correction and synchronization for storage technologies, comprising hard disks, tape and flash memories. He is a Fellow of the IEEE since 2000.

Jehoshua Bruck (S'86-M'89-SM'93-F'01) received the B.Sc. and M.Sc. degrees in electrical engineering from the Technion-Israel Institute of Technology, Haifa, Israel, in 1982 and 1985, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1989.

He is the Gordon and Betty Moore Professor of computation and neural systems and electrical engineering at the California Institute of Technology, Pasadena, CA. His current research interests include information theory and systems and the theory of computation in biological networks. His industrial experience includes working with the IBM Almaden Research Center, as well as cofounding and serving as Chairman of Rainfinity, acquired by EMC in 2005; and XtremIO, acquired by EMC in 2012.

Dr. Bruck is a recipient of the Feynman Prize for Excellence in Teaching, the Sloan Research Fellowship, the National Science Foundation Young Investigator Award, the IBM Outstanding Innovation Award, and the IBM Outstanding Technical Achievement Award. His papers were recognized in journals and conferences, including winning the 2010 IEEE Communications Society Best Student Paper Award in Signal Processing and Coding for Data Storage for his paper on codes for limited-magnitude errors in flash memories, the 2009 IEEE Communications Society Best Paper Award in Signal Processing and Coding for Data Storage for his paper on rank modulation for flash memories, the 2005 A. Schelkunoff Transactions Prize Paper Award from the IEEE Antennas and Propagation Society for his paper on signal propagation in wireless networks, and the 2003 Best Paper Award in the 2003 Design Automation Conference for his paper on cyclic combinational circuits.

Anxiao (Andrew) Jiang (S'00-M'05-SM'12) received the B.Sc. degree in electronic engineering from Tsinghua University, Beijing, China in 1999, and the M.Sc. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, California in 2000 and 2004, respectively.

He is currently an Associate Professor in the Computer Science and Engineering Department and the Electrical and Computer Engineering Department at Texas A&M University in College Station, Texas. He has been a visiting professor at California Institute of Technology, University of California in San Diego, Ecole Polytechnique Federale de Lausanne (EPFL) and Singapore Data Storage Institute (DSI), and a consulting researcher at HP Labs, EMC and Microsoft Research. His research interests include information theory, data storage, networks and algorithm design.

Prof. Jiang is a recipient of the NSF CAREER Award in 2008 for his research on information theory for flash memories and a recipient of the 2009 IEEE Communications Society Best Paper Award in Signal Processing and Coding for Data Storage.

Kannan Ramchandran (F'05) received Ph.D from Columbia University in 1993. He is a Professor of Electrical Engineering and Computer Science at UC Berkeley, where he has been since 1999. He was on the faculty at UIUC from 1993 to 1999, and with AT&T Bell Labs from 1984 to 1990. Prof. Ramchandran is a Fellow of the IEEE. He has published extensively in his field, holds over a dozen patents, and has received several awards for his research and teaching including an IEEE Information Theory Society and Communication Society Joint Best Paper award for 2012, an IEEE Communication Society Data Storage Best Paper award in 2010, two Best Paper awards from the IEEE Signal Processing Society in 1993 and 1999, an Okawa Foundation Prize for outstanding research at Berkeley in 2001, and an Outstanding Teaching Award at Berkeley in 2009, and a Hank Magnuski Scholar award at Illinois in 1998. His research interests are broadly in the area of distributed systems theory and algorithms intersecting the fields of signal processing, communications, coding and information theory, and networking. His current systems focus is on large-scale distributed storage, large-scale collaborative video content delivery, and biological systems, with research challenges including latency, privacy and security, remote synchronization, sparse sampling, and shotgun genome sequencing.

Bane Vasic (S'90-SM'02-F'12) is a Professor of Electrical and Computer Engineering and Mathematics at the University of Arizona and a Director of the Error Correction Laboratory. Dr. Vasic is an inventor of the soft error-event decoding algorithm, and the key architect of a detector/decoder for Bell Labs magnetic recording read channel chips which were regarded as the best in industry. Different variants of this algorithm were implemented in virtually all magnetic hard drives. His pioneering work on structured low-density parity check (LDPC) error correcting codes and invention of codes has enabled low-complexity iterative decoder implementations. Structured LDPC codes are today adopted in a number of communications standards, and are the considered for adoption in extremely high-density magnetic recording, nonvolatile memories and optical communications systems. He was also involved in a Digital Versatile Disc (DVD) Standardization Group and is involved in the IEEE Working Group on Error Correction Coding for Non-Volatile Memories. His work on characterization of theoretically achievable recording densities for multi-track and multi level recording systems lead to practical constrained codes for DVD. Dr. Vasic is known for his theoretical work in error correction coding theory and codes on graphs which has led to characterization of the hard decision iterative decoders of LDPC codes, and design of codes and decoders for binary symmetric channel with best error-floor performance known today. He is a co-founder of Codelucida, a startup company developing advanced error correction solutions for communications and data storage. He is an IEEE Fellow and da Vinci Fellow.