Guest Editorial
Channel Modeling, Coding and Signal Processing for Novel Physical Memory Devices and Systems

Shayan Srinivasa Garani, *Senior Member, IEEE*, Tong Zhang, *Senior Member, IEEE*, Ravi Hiranand Motwani, *Senior Member, IEEE*, Haralampos Pozidis, *Senior Member, IEEE*, and Bane Vasić, *Fellow, IEEE*

THE digital universe is doubling every two years and expected to reach an unwieldy 44 zettabytes into the next decade. To cope with the ever increasing need for storing, transmitting and retrieving huge amounts of data, cloud storage, data centers and other massively distributed storage networks have emerged. These rely on efficient memory technologies at the physical level for speed, reliability and energy efficiency.

Advancements in the theoretical foundations from seventy plus years of signal, information and communication theories have been used with phenomenal success in data storage systems starting from punch cards to the state-of-the-art magnetic, optical and flash devices. In all these technologies, coded information bits are densely packed within the physical medium as a change in magnetic flux, optical interference effects or electron charge and subjected to various read, write and media conditions.

At the recording technology level, recent advances in data storage technologies, such as, 3D non-volatile memories (NVM) and two-dimensional magnetic recording (TDMR) are bound to transform the storage industry in the next few years. These technologies are inherently multi-dimensional. Further, some technologies such as flash and optical storage admit multi-level recording unlike magnetic recording which is binary based. These advancements require a fundamental understanding into the channel characterization aspects under practical physical constraints. By developing reasonable model abstractions, we can assess data storage density limits along with practical ways for approaching them using sophisticated signal processing and coding algorithms.

At the system level, we need new channel architectures catering to high throughput and better energy efficiencies with lower cost for realizing and integrating the memory hierarchy into a working piece of Silicon. Often, this translates to finding the right balance between algorithmic sophistication purely driven by the signal-to-noise ratio (SNR) of the medium and an amenable hardware solution that satisfies other system level constraints such as power, area and throughput.

The previous issue of JSAC on data storage was Communication Methodologies for the Next-Generation Storage Systems and was published in May 2014. We have several new developments since then dedicated to physical data storage. This issue has 19 high quality papers covering various physical storage technologies with contributions from both academia and industry. The papers are organized into the following categories (a) device-level characterization, modeling and capacity analysis, system level optimization and coding aspects for flash memories (b) statistical analysis of discrete grain models, signal processing and related channel engineering aspects for TDMR (c) new coding methods for resistive memories, fast decoders for futuristic memories and spatially-coupled Euclidean geometry codes for storage and (d) novel signal processing aspects to mitigate gain mismatch and signal detectors for tackling auto-regressive noise applicable to existing storage channels.

Flash memories have special constraints during the reading, writing and erasure processes. Read-write asymmetricity and degradation of memory cells due to aging, program erase cycles, and inter-cell interference require accurate channel characterization and modeling along with novel programming and channel coding towards system optimization. The first group of papers covers several timely topics related to flash storage including channel characterization, modeling and capacity analysis, intercell interference (ICI) mitigation, channel coding and cell programming for faster multi-level flash memories. These results are particularly useful in advancing the channel and systems engineering aspects of current flash memories.

Accurate modeling of the static voltage distribution can be used to determine errors occuring within a flash cell for a read reference voltage so that the flash controller can adjust the read reference voltage for minimizing the error rate, thereby increasing the life time of the device. Also, knowing how the dynamic distribution changes over time facilitates flash controller mechanisms for dynamically adjusting the system

S. S. Garani is with the Department of Electronic Systems Engineering, Indian Institute of Science, Bengaluru 560012, India (e-mail: shayan.gs@dese.iisc.ernet.in).

T. Zhang is with the Department of Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: tong.zhang@ieee.org).

R. H. Motwani is with Intel Corportation, Santa Clara, CA 95052 USA (e-mail: ravi.h.motwani@intel.com).

H. Pozidis is with IBM Research Zürich CH-8803, Switzerland (e-mail: hap@zurich.ibm.com).

B. Vasić is with the Department of Electrical and Computer Engineering, The University of Arizona, Tucson, AZ 85719 USA (e-mail: vasic@ece.arizona.edu).

Digital Object Identifier 10.1109/JSAC.2016.2603713

parameters such as the error correcting code (ECC) strength, read reference and passthrough voltages for better reliability.

The paper "Enabling accurate and practical online flash channel modeling for modern MLC NAND flash memory" by Luo et al. proposes a new, low-complexity flash memory model based on experimental characterization of the threshold voltage distribution from 1X-nm multi-level cell (MLC) flash chips. Their model based on a modified version of the t-distribution and power law captures the threshold voltage distribution and predicts the future distribution with wear leveling with less than 3% modeling error. This work is useful to enhance the life time of the device beyond manufacturer specifications.

Channel models for flash memories are important to estimate the capacity of the underlying flash memory channel. In the paper "On the capacity of the beta-binomial channel model for multi-level cell flash memories" by Taranalli et al., the authors focus on flash channel modeling from an information-theoretic view point. They propose a truncated-support beta-binomial model (TS-BBM) and derive the capacity. Using empirical error statistics from 1X-nm and 2Y-nm MLC flash chips, the channel capacity is evaluated as a function of the program/erase cycles giving an upper bound on the coding rates applicable for these channels.

Flash memories inherently admit incremental step pulse programming (ISPP) to program data concurrently over word lines. The paper "Flash memories: ISPP renewal theory and flash design trade-offs," by Asadi et al. models incremental step pulse programming (ISPP), the method used for programming levels in flash memory, using renewal process theory. This is probably the first model of flash threshold voltage that attempts to take into account the nature of the underlying ISPP programming algorithm, bringing an interesting twist to the existing NAND flash memory modeling literature. The paper proceeds to analyze intercell interference, and the interplay between step size and write latency optimization.

The paper "Performance of multilevel flash memories with different binary labelings: A multi-user perspective," by Huang et al. studies the performance of different decoding schemes for multilevel flash memories where each page in every block is encoded independently. The authors consider both low-complexity approaches treating interference as noise decoding and relatively high-complexity approaches using successive cancelation decoding. The achievable rate regions and sum rates of both decoding schemes are obtained for different flash channel models. The effect of different binary labelings of the cell levels is also studied, and the optimal labeling for each decoding scheme and channel model is identified.

In the paper "Capacity of the MLC NAND flash channel" by Parnell et al., the authors provide a framework for assessing the symmetric capacity of NAND pages using statistics collected from 1Y-nm flash memories. The symmetric capacity is calculated at various instances during the P/E cycling and data retention phases for both hard and soft-output channels. The authors show that it is possible to significantly increase the endurance by nearly 85% by configuring variable rate ECC architectures to match the distribution of capacity within that block.

Constrained codes along with ECC is the most common channel coding configuration in almost all storage devices. In the paper "Mitigation of inter-cell interference in flash memory with capacity-approaching variable-length constrained sequence codes" Cao and Fair propose variable-length constrained codes to mitigate ICI in an all-bitline flash memory with multi-page programming for single-level cell (SLC), MLC, and triple-level cell (TLC) configurations. They provide a construction exploiting inherent error control properties within such constrained codes to limit error propagation and show improved reliability when coupled to another ECC.

Write-once memory (WOM) codes are useful for several rewrites on flash cells before erasures, but prone to increase in errors due to ICI. In the paper "d-imbalance WOM codes for reduced intercell interference in multi-level NVMs" by Hemo and Cassuto, the authors provide simple constructions for the design of WOM codes to reduce ICI. They derive an upper bound on the number of guaranteed writes of a d-imbalance WOM code and prove the optimality of their construction. This work is useful for flash memories to reduce the bit error rate.

For low-density parity-check (LDPC) codes widely used in NAND flash memories, system performance is closely tied to the number of bits per message used by the message-passing decoder. The paper "LDPC decoding mappings that maximize mutual information," by Romero and Kurkoski proposes message-passing decoding mappings that maximize mutual information to decode regular LDPC codes. In this paper, finite-length results for various regular LDPC code rates show that using 4 bits per message is sufficient to perform close to theoretical limits, achieving full belief propagation decoding performance. This is attractive for data storage applications where high-rate codes are used and a few bits per message for the soft-decision decoding are needed to obtain an acceptable performance without increase the memory latency for the reading process.

Low-Density Parity-Check (LDPC) codes are now actively being considered for modern dense storage devices, such as multi-level flash and hard disk drives. As many emerging memory devices exhibit an increased level of asymmetry (e.g., 3D Flash), LDPC codes optimized for symmetric, AWGN-like channels, are not optimum for these applications. The paper "A general non-binary LDPC code optimization framework suitable for dense flash memory and magnetic storage," by Hareedy et al. proposes a general combinatorial framework for the analysis and design of non-binary LDPC (NB-LDPC) codes for asymmetric channels. The authors develop a general code optimization framework, and demonstrate its effectiveness on the realistic highly-asymmetric normal-Laplace mixture flash channel. Moreover, the authors show that their framework can be customized to optimize NB-LDPC codes for channels with memory and other asymmetric channels.

The paper "Minimal maximum-level programming - combined cell mapping and coding for faster multi-level-cell memory," by Berman and Birk describes a novel method, Minimal Maximum-Level Programming (MMLP) for storing data in multi-level memory cells. The authors propose to write data sectors to the cells of a word line in an incremental

manner so that only small increments in the threshold voltage of the underlying cells are required during programming. The MMLP concept presents an interesting approach towards reducing write latency in Flash memory. The technique is also applicable to phase change memories (PCMs).

TDMR is an emerging magnetic storage technology that aims to achieve the highest areal density possible by storing a bit over a magnetic grain. In TDMR, the sectors are inherently two-dimensional with reduced track pitch and bit widths, leading to severe 2D inter-symbol interference (ISI). The read channels engineering requires accurate media models and powerful two-dimensional signal processing and coding algorithms for overcoming 2D ISI, timing artifacts and jitter noise due to irregular media grain positions as well as electronics noise during reading. The second group of papers cover topics related to statistical analysis of the discrete grain model and equalization and signal detection techniques for shingled and multi-track detection applicable to this emerging area of magnetic storage.

In the paper "MCMC methods for drawing random samples from the discrete-grains model of a magnetic medium," Das and Kashyap revisit the discrete grain model comprising of four different rectangular shapes for modeling the distribution of grains on a magnetic medium and analyze the statistical properties. The authors use the Metropolis-Hastings algorithm to design a Monte-Carlo Markov chain method for sampling from a probability distribution arbitrarily close, in total variation distance, to a target distribution of all tilings of these basic shapes. The authors also present some bounds and conjectures on the mixing times of the underlying Markov chains. These results are useful for estimating the time taken by the MCMC methods for generating a random tiling sampled from the target probability distribution.

Using turbo-equalization with 2-D row/column detector and an irregular repeat-accumulate (IRA) code, the paper "Turbo-equalization for two dimensional magnetic recording using Voronoi model averaged statistics" by Mehrnoush et al. report a 6.5% increase in density over a Voronoi media grain model. The authors also consider a non-linear function to map extrinsic output loglikelihood ratios from the detector to the decoder to overcome bit overwrites.

Shingled magnetic recording based hard disk drives are in commercial use today. In this technology, data is written over narrow overlapping tracks. In the paper "Bidirectional decision feedback modified Viterbi detection (BD-DFMV) for shingled bit-patterned magnetic recording (BPMR) with 2D sectors and alternating track widths," by Wang and Kumar, the authors achieve a 2 dB gain in SNR by using multiple readers and bi-direction decision feedback equalization with alternatingly varying track widths over the uniform track width case under identical channel conditions.

In the paper "The rotating-target algorithm for jointly detecting asynchronous tracks," Sadeghian and Barry discuss joint detection of multiple tracks that are written asynchronously in two-dimensional magnetic recording. The proposed ROTAR algorithm, uses the joint Viterbi algorithm and a time-varying target that results from embedding the asynchrony of the tracks into the underlying target. Simulation results in the case of K=2 tracks with N=2 readback waveforms indicate that the proposed algorithm is significantly better than a previously reported detector that separately detects two tracks and closely matched a genie-aided detector.

The third group of papers covers coding methods for new memories, fast decoders for emerging memories and Euclidean geometry based spatially-coupled LDPC codes for storage applications. Resistive random-access memories (RRAMs) and phase change memories (PCMs) are considered potential replacements to the flash technology due to better speed performance. In RRAMs the resistance across a normally insulating dielectric material can be made to conduct with the application of a large voltage. Unlike flash memories, data in RRAMs allow rewriting of data without erasures. Improved write endurance and write power consumption are key challenges to realizing the potential of RRAMs. In the paper "Locally rewritable codes for resistive memories", Kim et al. propose locally rewritable codes (LWC) for correcting multiple stuck-at defects in resistive memories. These codes are motivated by locally repairable codes (LRCs) well known in distributed storage and useful to reduce the problems of write endurance and write power consumption. The authors quantify the writing cost based on experimental data.

Storage class memories need fast decoding error correcting codes to be able to correct errors in just a few nanoseconds. In the paper "Fast decoding ECC for future memories," by Amato et al., the authors consider ultra fast double and triple correcting BCH codes within a novel coded systems architecture to avoid bottlenecks in the decoding process. They provide a gate level design of these architectures with an analysis of hardware-oriented implementations of the underlying finite field operations of the code.

In the paper "Euclidean geometry based spatially-coupled LDPC codes for storage," Xie et al. present a method for constructing binary spatially-coupled (SC) low-density parity-check codes based on Euclidean geometry (EG) LDPC codes for next-generation storage applications. The authors propose a two dimensional edge-spreading process to construct the base matrix of EG-SC LDPC codes, which consists of matrix unwrapping, periodically time-varying of a protograph. The error rate performance of the constructed EG-SC LDPC codes are evaluated using a weighted bit-flipping decoding algorithm which is suitable for mobile storage applications. Numerical results show that the bit error rate performance of the constructed EG-SC LDPC codes is superior to that of their EG LDPC block codes counterpart in a low raw error rate regime.

The last group of papers deal with novel signal processing aspects for memory applications.

In the paper "Minimum Pearson distance detection using mass-centered codewords in the presence of unknown varying offset" by Immink and Skachek, the authors consider a new approach to the transmission and storage of binary coded data over a channel, where a Pearson distance- based detector is used instead of the conventional Euclidean based detector. This formulation is useful in mitigating unknown channel gain,

and time-varying varying offset applicable to optical storage and flash memories. The authors also study the properties of a code satisfying the center-of-mass constraints used in conjunction with the Pearson distance detector and derive the redundancy for asymptotic codeword lengths along with an analysis on word error rate.

Noise predictive maximum likelihood Markov based Viterbi detectors have been popularly used in magnetic and optical storage channels. In the paper "Near optimal Viterbi algorithm for storage channels with linear regressive noise," Soltanpur and Cruz investigate noise predictive sequence detection algorithms for intersymbol interference channels with linear regressive (non-Markov) noise. The multivariate moving average noise model leads to an extended noise predictive maximum likelihood (NPML) detector. The simulation results indicate that the advantage of the proposed vector noise predictive (VNP) Viterbi algorithm may be significant the noise statistics deviate from the Gauss-Markov assumption.

Channel engineering is an important component along with media and recording research to meet the continued demands for emerging high density storage technologies with low read/write access times and low power constraints. Each storage technology presents unique technological challenges that must be overcome via coding, signal processing and system level optimization to encode, write and read data reliably. Magnetic and optical storage channels have evolved from a 1D paradigm to 2D. Signal processing in these channels extends all the way from the analog side to the digital back end, where powerful capacity achieving ECCs are used to overcome the channel limitations. On the other hand, flash memories are inherently digital and pose read/write asymmetry issues and constraints on block programming of cells that requires ECC architectures best suited for these applications. Similarly, RRAMs are based on stuck-at-defect models requiring short length algebraic codes most suited for these memories. The right combination of signal processing and coding cognizant of the channel limitations is important to realize a viable memory system. In this journal special issue, we have many interesting papers that cover various channel engineering aspects of different state-of-the-art physical storage technologies. It is expected that future data storage technologies will be multi-level and multi-dimensional. This will require a fundamentally new approach to address these technological challenges.

We wish to thank all the reviewers for providing timely, constructive and detailed reviews that helped us in the paper selection process and significantly improved the final content of the accepted papers. We also hope that these expert reviews were useful to the authors whose papers could not be accepted to this issue. Additionally, we wish to thank the authors for submitting top quality papers that have collectively made this special issue a timely, comprehensive and a long lasting reference. Finally, we would like to express our sincere gratitude to the JSAC team: Prof. Douglas Leith who has served as our mentor, Prof. Muriel Medard, the current Editor-in-Chief, Ms. Laurel Greenidge, the current Executive Editor, Ms. Sharon T. Nutter, the current Journals Coordinator, and Ms. Lauren Briede, the current staff Editorial Assistant, for their time, efforts and encouragement in bringing together this special issue. We hope that the present issue will serve as a lasting reference to researchers working in this niche area.

Shayan Srinivasa Garani (SM'11) received the Ph.D. degree in electrical and computer engineering with a minor in mathematics from the Georgia Institute of Technology, Atlanta GA, USA. He is currently a Faculty Member with the Department of Electronic Systems Engineering, Indian Institute of Science, where he directs the Physical Nanomemories, Signal, and Information Processing Laboratory. Prior to joining academics, he managed advanced read channels research at Western Digital (USA) and directed several external university research programs. He was the key inventor and an architect for advanced coding and signal processing solutions that went into hard disks and flash memories. He holds 12 US patents and several others pending in key disruptive technology areas within data storage. He is a Senior Member of the Optical Society of America and Chair of the Photonic Detection Group.

Tong Zhang (S'97–M'02–SM'08) received the Ph.D. degree in electrical engineering from the University of Minnesota in 2002. He is currently a Professor with the Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute, NY, USA. He has co-authored over 150 refereed papers in memory circuits and systems, VLSI signal processing, and computer architecture. He has graduated 15 Ph.D. students. He has served as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II, and the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He served as the Technical Program Co-Chair of the 2012 ACM Great Lakes Symposium on VLSI, and the 2012 IEEE Workshop on Signal Processing Systems, and the General Co-Chair of 2013 ACM Great Lakes Symposium on VLSI.

Ravi Hiranand Motwani (SM'00) received the Ph.D. degree in electrical engineering from the Indian Institute of Science, Bangalore, in 1998. He is currently the Program Manager for the ECC activities for the Non-Volatile Memory Solutions Group with Intel Corportation, Santa Clara, CA, USA. He is also the ECC Lead Manager of flash memory and 3-D cross point solutions. He has previously held research positions with Philips Research Labs, The Netherlands, Broadcom Corporation, USA, and faculty positions with IIT Kanpur, and the National University of Singapore. He has over 18 years of experience in ECC for storage and communications. He has over 40 US patents and 40 papers in peer-reviewed conferences, and journals.

Haralampos Pozidis (SM'10) received the Ph.D. degree in electrical engineering from Drexel University, Philadelphia, PA, USA, in 1998. He was with Philips Research, Eindhoven, The Netherlands. He is currently managing the Non-volatile Memory Systems Group, IBM Research Labs, Zürich, Switzerland. He has been involved in read channel design for DVD and Blu-ray disc optical recording formats with technology transfers to optical drive products with Philips. He has played a key role in developing the world's first scanning probe-based data storage system, and the Millipede. His current focus is on developing flash memory controllers for all-flash arrays as well as phase change memory technology and system solutions. He holds over 70 US and European patents. He has co-authored over 100 publications in the areas of solid-state memory technology, probe-based data storage, control systems technology, and optical data storage. He has received the 2009 Control Systems Technology Society Award and the 2009 IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY Best Paper Award. He is also a Principal Research Scientist and an IBM Master Inventor.

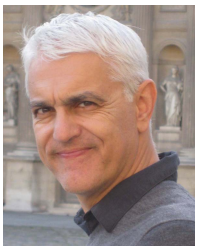

Bane Vasić (F'11) received the Ph.D. degree in electrical engineering from the University of Nis, Serbia, in 1994. He is currently a Professor of electrical and computer engineering and mathematics with The University of Arizona and the Director of the Error Correction Laboratory. He is also an inventor of the soft error-event decoding algorithm, and the key architect of a detector/decoder for Bell Labs data storage chips which were regarded as the best in industry. He is a Chair of the IEEE Data Storage Technical Committee, an Editor of two previous IEEE JSAC Special Issues on Data Storage Channels and a member of the Editorial Board of the IEEE TRANSACTIONS ON MAGNETICS. He is a da Vinci Circle Fellow, Fulbright Scholar, and the Founder and Chief Science Officer of Codelucida Inc.