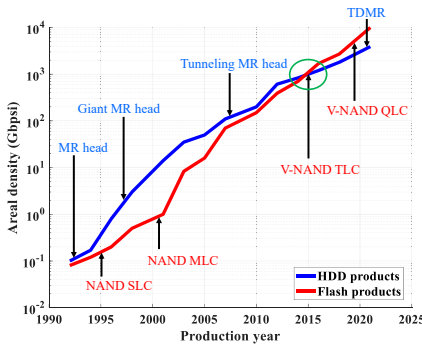


Ahmed Hareedy

Research Statement



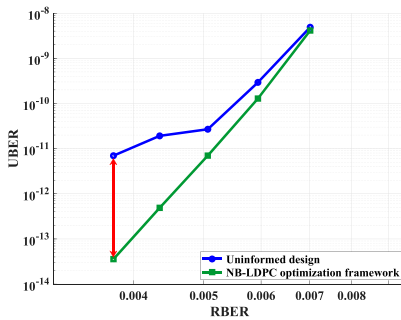
Data storage density trends.

I develop new methods of data processing for modern solid-state and magnetic data storage devices that exploit advances in the underlying physics. I devise coding techniques that improve performance of cloud storage as well as distributed computing and machine learning. I am intrigued by new problems in coding and information theory that are illuminated either by unparalleled access to data and computing or by new technologies, e.g., blockchain, DNA data storage, and quantum computing.

With the rapidly growing amounts of data being accessed, processed, and stored by modern applications, e.g., internet of things (IoT) applications, data storage devices are required to be increasingly denser. This continuous and rapid density increase results in new characteristics associated with sources of errors. I propose effective coding techniques that take full advantage of these new characteristics to increase the reliability of modern data storage. From devices to clouds, my techniques improve performance and reduce latency of storage systems.

Graph-Based Codes for Data Storage

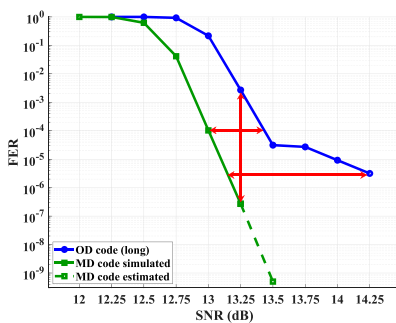
Low-density parity-check (LDPC) codes are specified by graphs, and they are the error-correction technique of choice in data storage because of their capacity-approaching performance and the degrees of freedom they offer. These codes also find application in 5G wireless communication systems. A key feature of graph-based codes is that decoding errors can be mapped to particular detrimental structures in the graph defining the code.



Error rates of NB-LDPC codes in Flash.

Non-binary LDPC (NB-LDPC) codes and spatially-coupled (SC) codes are among the exciting recent developments in graph-based codes. I significantly improved the performance of these codes in various data storage systems by first identifying the graphical structures implicated in system errors, and then eliminating as many of these structures as possible. I devised a general framework for optimizing NB-LDPC codes, which effectively works for Flash and magnetic recording (MR) systems. This framework leads to performance gains of up to 2.5 orders of magnitude compared with uninformed designs. I also introduced an approach to design high performance SC codes for AWGN, Flash, and MR channels. This channel-aware approach takes advantage of all the available degrees of freedom to

construct SC codes that outperform block codes of the same length and rate, which answered an open question about SC codes. My designs have the potential to be employed in other applications as well. I am also interested in the statistical prediction of performance for graph-based codes in modern data storage systems.



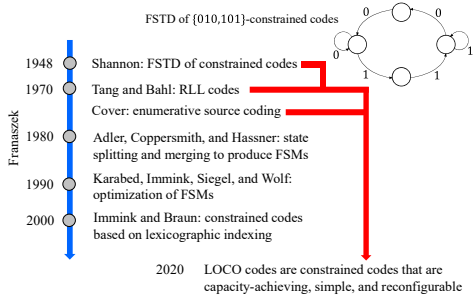
Error rates of MD/OD codes in MR.

Multi-dimensional (MD) graph-based codes couple different copies of a one-dimensional (OD) code to mitigate different types of interference and channel non-uniformity in modern storage systems. I used MD coupling to optimize the elimination of detrimental objects from the graph defining an MD code. On realistic models of Flash and MR systems, I demonstrated significant gains in program/erase (P/E) cycles and signal-to-noise ratio (SNR), respectively, compared with OD codes with similar length and rate. These performance gains are in the code threshold, waterfall, and error floor regions, and they translate to significant increases in the lifetime and density of the storage device.

I won the *Best Paper Award* (GLOBECOM 2015), the *Memorable Paper Award* (NVMW 2018), the *Distinguished Ph.D. Dissertation Award* (UCLA 2019), and the *Best Student Paper Award* (DSTC 2020) for my innovations in graph-based codes for storage.

Constrained Codes for Storage and Transmission


Interference in MR devices results from the read-head sensitivity, and it results from parasitic capacitances in Flash devices. Constrained codes make it possible to mitigate interference, and thus improve performance. They find application in data storage (MR, Flash, optical) devices and in computer standards for data transmission. The data storage industry adopts coding techniques that combine constrained codes, designed to avoid interference-causing patterns, and LDPC codes, designed to correct the errors that remain.



History and my LOCO codes.

In 1948, Shannon represented a constraint via a finite-state transition diagram (FSTD), and introduced the notion of capacity. In the 1980s and the 1990s, connections with symbolic dynamics led to methods of encoding and decoding, based on finite-state machines (FSMs), at or very close to capacity. For binary alphabets, it is not straightforward to apply such methods to construct near capacity-achieving codes. This goal becomes even more difficult for non-binary alphabets, which are suitable for non-binary physical gates in Flash devices with more than two levels per cell. If also the problematic patterns change as the storage device ages, it becomes very challenging to reconfigure the code.

By returning to the roots of constrained coding, I devised a general method for designing constrained codes based on lexicographic indexing. The method starts from the finite set of patterns to forbid, and derives encoding and decoding algorithms by going through systematic steps. The new codes are collectively named lexicographically-ordered constrained (LOCO) codes. LOCO codes are capacity-achieving codes offering low complexity, where reconfiguration is as simple as reprogramming an adder.




LDPC+LOCO integration in MR.

1) I developed binary symmetric LOCO (S-LOCO) codes for one-dimensional MR systems. On an industry-recommended model of a modern MR system, I achieved a density gain of about 20% by using an S-LOCO code to protect only the parity bits of an SC (LDPC) code. 2) The need to increase storage density in Flash is driving evolution from SLC (2-level) to MLC (4-level), TLC (8-level), QLC (16-level), and PLC (32-level) devices. I designed q -ary asymmetric LOCO (QA-LOCO) codes for Flash memories with q levels per cell, $q \geq 2$. QA-LOCO codes can contribute to achieving significant lifetime gains with normalized rates more than 0.95. 3) Two-dimensional magnetic recording (TDMR) is an emerging technology, where inter-track interference and TD flying height variations

result in TD non-uniformity. I developed LOCO codes that forbid TD patterns derived from the discrete approximate TD channel impulse response, and demonstrated notable performance gains on a realistic TDMR model.

Centralized and Decentralized Cloud Storage



Centralized cloud storage network.

Because of the strict constraints on latency imposed by IoT and machine-learning applications, distributed systems are becoming increasingly attractive. Moreover, the blockchain technology is moving various systems towards decentralization. In cloud storage, symbols of a codeword to be stored in a cloud are distributed over its local storage nodes (local servers in a centralized system), which may be subject to erasures. I contributed to hierarchical cloud storage solutions, which are based on algebraic coding, to increase reliability. These solutions are for centralized and blockchain-based decentralized cloud storage, and they provide local and global, i.e., multi-level, erasure-correction capability. These solutions support: 1) Heterogeneity: data lengths in various clouds are allowed to differ. 2) Scalability: new clouds can be added with minimum changes to the existing network, which reduces cost. 3) Flexibility: a cloud that has its data unexpectedly becoming hot, i.e., of higher demand, can split into smaller clouds to reduce

latency. Our solution for decentralized cloud storage works for any network topology, i.e., it is topology-adaptive.

I am currently working on devising hierarchical codes and decoders to address errors rather than erasures. This work finds application in solid-state storage solutions supporting fast multi-task systems, e.g., computational storage solutions, which have potential to be used in self-driving cars.

Future Research Directions

Modern graph-based codes and applications: My current results on MD codes for Flash and MR systems are quite promising. I will improve my systematic SC and MD code design frameworks to increase performance gains. I will develop an asymptotic analysis that further explains experimental results obtained for data storage, and is able to predict performance in other areas such as wireless communications. I am also interested in reducing the complexity and latency of decoders for both SC and MD codes, and exploring implications for distributed systems.

Three-dimensional cross-point (3D XPoint) and resistive random-access memory (ReRAM) operate at higher speeds than Flash memory. I want to develop new classes of graph-based codes and decoders that can meet the latency requirements of such revolutionary non-volatile memory technologies.

LOCO codes in Flash and TDMR systems: I demonstrated significant density gains in MR systems by using a LOCO code to protect only the parity bits of an LDPC code, which limits the rate loss. I will apply this idea of unequal data protection to Flash systems, and I will combine my MD graph-based codes with LOCO codes for TDMR and other modern MD storage systems. I will develop a theoretical analysis that explains performance gains in both solid-state and magnetic storage systems. I expect this analysis to be of interest in other areas such as communications.

Applying machine learning in data storage: It is very difficult to develop accurate mathematical models for channels underlying next-generation storage technologies such as Q/P-LC Flash and TDMR. Machine learning can help here. I am interested in using errors collected before the graph-based decoder to learn patterns that need to be forbidden via the constrained code. As the device ages and forbidden patterns change, I propose responding by reconfiguring the LOCO code. I am interested in using errors the graph-based decoder cannot correct to learn detrimental graphical structures to eliminate. I also want to use machine learning to predict performance and improve iterative decoding in modern storage and communication systems.

Coding for low latency reliable distributed computing: I am intrigued by novel straggler-resilient coding techniques to speed up distributed computing and distributed machine learning. I want to investigate classes of graph-based codes in this area. The performance of such codes and the sparse nature of their parity-check matrices are promising. However, new metrics and low latency decoding schemes need to be devised.

The end of Moore's law was about a decade ago. Computer architects have been searching for new solutions to speed up processing, e.g., involving GPUs. One idea is to bring data storage units closer to distributed computing units (multi-core processing units) to significantly reduce the data transfer overhead. This framework is named computational storage. I am excited about developing coding solutions that enable low latency computational storage without compromising reliability. I am also interested in data processing methods for in-memory computing systems. These systems have the potential to overcome all speed limitations that are due to memory access.

Coding and sequencing for DNA data storage: The DNA data storage technology promises increasing the density and the lifetime of storage devices by orders of magnitude. I am very interested in collaborating with professors from other departments to understand how DNA molecules operate, and therefore identify the coding and sequencing techniques that are suitable for this emerging technology.

Error correction in the quantum world: Quantum computers promise to be more capable of solving certain problems than any classical computer, and these devices are moving out of the lab and are becoming generally programmable. Error-correction techniques are required in the quantum world to ensure that computations are being performed and results are being stored reliably. I am very interested in collaborating with professors from other departments to understand channels underlying a quantum computing system. Based on this understanding, I want to develop effective, low latency quantum error-correction techniques that abide by system restrictions. As demonstrated by my research interests and achievements, I am highly interested in interdisciplinary research.