

# IEEE Information Theory Society Newsletter



Vol. 67, No. 4, December 2017

EDITOR: Michael Langberg, Ph.D.

ISSN 1059-2362

**Editorial committee:** Alexander Barg, Giuseppe Caire, Meir Feder, Joerg Kliewer, Frank Kschischang, Prakash Narayan, Parham Noorzad, Anand Sarwate, Andy Singer, and Sergio Verdú.

## President's Column

*Rüdiger Urbanke*

The wheels have been put in motion for the officer shift register to move once again. Michelle Effros, a BoG member for more than a decade, will shift out and take a well-deserved break. I would like to thank her for all her past service. During her presidency she oversaw a large number of Shannon Centennial events and her successful fund raising efforts made it possible for us to start the Shannon Documentary. Elza Erkip will be our President next year. Last time the Turks chased the Austrians we got croissants as a result. I expect great things to happen once again! Emina Soljanin shifts into the position of Vice President, and Helmut Boelsckei has been elected 2nd Vice President.



I would like to thank the outgoing BoG member Urbashi Mitra for her service. Stephen Hanly, Vince Poor, Aylin Yener, and Wei Yu were re-elected. And I welcome the incoming BoG members Daniella Tuninetti and Taschy Weissman.

In the past few years we have started a considerable number of outreach initiatives. Our most ambitious one, a Shannon documentary, is in full production and should be finalized in the first quarter of 2018. I am looking forward to the time where we can all enjoy the fruits of all these labors together.

Matthieu Bloch, together with Michelle Effros, Christina Fragouli, and Suhas Diggavi has been working diligently in the background to produce short movies that explain some of our Society's many technical contributions to a general audience. With the help of Brit Cruise, they have produced two such videos, one on Network Coding and one on Space-Time Codes. You can find these videos on Brit's Youtube channel "Art of the Problem" <https://www.youtube.com/watch?v=B0ZcAWEvjCA&t=11s> <https://www.youtube.com/watch?v=cbD4NsZQKYw&t=4s> This is just the beginning. At least two more videos are currently in production and we plan to continue this effort for the years to come.

But there is more. If the Shannon Documentary is targeted at our friends and parents, and the short educational movies at high school children, we did not forget the youngest ones. Anna Scaglione and Christina Fragouli have produced a children's book entitled "Information, in small bits." Get them while they are young!

In case you are looking for a romantic getaway for Valentine's Day—look no more. I suggest the beautiful Catamaran Resort, Pacific Beach, in sunny San Diego. Our most beloved ITA will take place there from February 11th till February 16th. ITA is no doubt one of the highest entropy conferences around, due to Alon's Herculean efforts. An innovative

format at a great venue, plenty of great talks, and constant surprises. And the biggest uncertainty of them all? When will we get the much anticipated invitation email this year? And did I mention already that you will get Anna's and Christina's new children's book as a present?

Our next ISIT is in breath-taking Vail from June 17 to 22. Don't forget to submit your best papers! This time we are trying something slightly different. The paper registration deadline is January 9th, 2018 11:59 Hawaii Time (UTC-10). No exceptions will be made. But the papers can be modified until January 12th, 2018 11:59 Hawaii Time (UTC-10). Acceptance notifications will be sent out by March 31, 2018.

For several years we have discussed how to add to our publications portfolio. Two proposals have emerged during all these discussions. The first is to convert our excellent newsletter (thanks to M. Langberg!) into a Magazine. This Magazine is mainly for our current members, bringing ideas from the outside into our community. The second is to add a Special Topics Journal exploring how information-theoretic ideas can be used in areas such as machine learning, signal processing, theoretical computer science, or biology.

*(continued on page 28)*

## From the Editor

Michael Langberg



Dear colleagues,

Our winter issue of 2017 opens with Rüdiger Urbanke's final column as President of the IT Society. Please join me in thanking Rudi for his dedication and inspiring leadership over the past year, and in warmly welcoming our incoming President Elza Erkip. We then proceed with an illuminating tutorial article by Viveck Cadambe and Pulkit Grover presenting the key ideas and challenges in applying coding techniques to the increasingly important task of distributed computing in the context of unreliable synchronization and computing components, "Codes for Distributed Computing: A Tutorial." The article is based on the material presented in a tutorial at the 2017 IEEE International Symposium on Information Theory (ISIT) in Aachen Germany over the summer. Many thanks to the contributors for their significant efforts!

The issue continues with an inspiring contribution by Aylin Yener reflecting on

the past decade of North American Schools of Information Theory (NASIT); an intriguing article by Haizi Yu and Lav R. Varshney on the combination of mathematics and music inspired by the work of the late Mary Elizabeth Moore "Betty" Shannon; Tony Ephremides's Historian's column; a report on "BalkanCom'17", a new conference on communications and networking, by Stark Draper; a call for nominations for several central awards in our community; minutes from the IEEE Information Theory Society Board of Governors meeting this summer at ISIT; and a list of recent publications from the IEEE Transactions on Information Theory, Foundations and Trends in Networking, Foundations and Trends in Theoretical Computer Science, and Problems of Information Transmission. Many thanks to all for their significant efforts in the preparation of their contributions!

With sadness, we conclude this issue with a tribute to Lotfi Zadeh, a pillar of the electrical engineering and computer science communities and the founder of fuzzy logic, who passed away on Sep. 6th at the age of 96. Thanks to Ahmad Taher Azar, from Benha University, Egypt for bringing the announcement forward.

Please help to make the newsletter as interesting and informative as possible by sharing with me any ideas, initiatives, or potential newsletter contributions you may have in mind. I am in the process of searching for contributions outside our community which may introduce our readers to new and exciting problems and, as such, broaden the influence of our society. Any ideas along this line will be very welcome.

(continued on page 15)

### IEEE Information Theory Society Newsletter

IEEE Information Theory Society Newsletter (USPS 360-350) is published quarterly by the Information Theory Society of the Institute of Electrical and Electronics Engineers, Inc.

Headquarters: 3 Park Avenue, 17th Floor, New York, NY 10016-5997.

Cost is \$1.00 per member per year (included in Society fee) for each member of the Information Theory Society. Printed in the U.S.A. Periodicals postage paid at New York, NY and at additional mailing offices.

**Postmaster:** Send address changes to IEEE Information Theory Society Newsletter, IEEE, 445 Hoes Lane, Piscataway, NJ 08854.

© 2016 IEEE. Information contained in this newsletter may be copied without permission provided that the copies are not made or distributed for direct commercial advantage, and the title of the publication and its date appear.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.



## Table of Contents

President's Column . . . . .	1
From the Editor . . . . .	2
Codes for Distributed Computing: A Tutorial . . . . .	3
A Walk Over the Memory Lane: One Decade of North American Schools of Information Theory (NASIT) . . . . .	16
On "Composing Music by a Stochastic Process": From Computers that are Human to Composers that are Not Human . . . . .	18
The Historian's Column . . . . .	20
BalkanCom'17: A New Conference on Communications and Networking . . . . .	21
Call for Nominations . . . . .	22
IEEE Information Theory Society Board of Governors Meeting . . . . .	23
Recent Publications . . . . .	29
Call for Papers . . . . .	35
Conference Calendar . . . . .	40

# Codes for Distributed Computing: A Tutorial

Viveck Cadambe and Pulkit Grover

This article is based on the material presented in a tutorial in the 2017 IEEE International Symposium on Information Theory (ISIT), Aachen Germany in June 2017. We write this article in two parts that mirror the structure of the tutorial presented in ISIT. First we focus on *shared memory emulation*—a classical problem in the area of distributed computing [60, 9]. We give a tutorial-like overview of shared memory emulation and its applications, and then describe relevant coding-theoretic formulations. The high level goal of shared memory emulation is to build algorithms that expose a distributed data storage system with certain desirable properties to external clients. Our description may therefore be viewed as a transition from the area of codes for distributed storage, where algorithmic aspects can often be ignored in coding-related studies, to distributed computing, where algorithmic aspects play a central role in problem formulation and solutions.

Second, we describe how codes can be used to perform reliable computations using unreliable processing elements. Just as sophisticated error-correction techniques that achieve the Shannon capacity have revolutionized communication on unreliable channels, use of error-correction techniques in computing has the potential to revolutionize next generation computing systems. In fact, techniques of replication and error-correction have been used in computing for decades now. In the last few years, information-theory community has contributed significantly to both fundamental limits and achievable strategies in this area. The purpose of this section is to put this recent information-theoretic work in the broader historical perspective of the area, and utilize that to suggest intellectually interesting and important research directions.

We admit that our goal is to not be comprehensive. Instead, this article is intended for graduate students and researchers interested in the area to learn about some of the key ideas, and is by its nature biased by our own interests and perspectives.

## 1. Part I: Shared Memory Emulation

### 1.1. Introduction

A read/write memory admits two operations:  $write(variablename, value)$  and  $read(variablename)$ . The goal of the *read/write shared memory emulation* problem is to implement a shared read/write memory over a distributed system of processing nodes (Fig. 1). For simplicity<sup>1</sup>, we focus here on a single variable and omit the *variablename* parameter. When the same read/write memory is concurrently accessed by multiple *client* nodes—possibly to jointly run some application or perform some task—it is usually referred to as a shared memory.

<sup>1</sup>A reader familiar with distributed systems theory will note that there can be a significant loss of generality in studying just a single variable, in particular, in systems that are not atomic. Specifically, consistency criteria weaker than atomicity are not “composable”, so the study of a single variable may not suffice from an algorithmic viewpoint, nonetheless, our description here is restricted to atomic variables and therefore focusing on single variable suffices.



Figure 1. The shared memory emulation problem.

Algorithms that solve the shared memory emulation problem provide theoretical underpinnings of cloud-based data-storage/database services that offer a property called *consistency*, which is explained in Sec. 1.2.1. There are numerous commercial and open source cloud based consistent data storage services such as Amazon Dynamo DB [24], Couch DB [7], Apache Cassandra DB [45], Berkeley DB [65], Redis DB [19]. Consistent data storage services are used in a wide variety of applications such as reservation systems, financial transactions, multi-player gaming [14], online social networks, and fog computing [79, 61]. Modern consistent data storage system design combines distributed computing theoretic approaches that design provably correct algorithms with systems engineering innovation to provide fast read and write operations. The study of shared memory emulation and design of consistent data storage services is an active area of research in distributed computing systems.

In consistent data storage services, the speed of access to data is critical, and therefore such data is commonly cached in relatively expensive high speed memory (RAM or solid state devices), and flushed to the magnetic disks, which is a less expensive, slower medium, only in the case of memory overflow. While data replication is often used in practice to provide fault tolerance, motivated by the need to use memory in the most efficient way possible, erasure coding based algorithms to provide consistent distributed storage services have been studied in the distributed computing theory and systems research communities [42, 27, 15, 16, 25, 53, 52, 71, 96, 20]. The use of erasure coding for such applications poses interesting challenges and research opportunities in information and coding theory, distributed algorithms, and optimization. We first describe some of the key concepts of shared memory emulation. We then describe an information theoretic framework to study the problem, and conclude with open directions of research.

### 1.2. The Shared Memory Emulation Problem

In this section, we first describe a distributed systems model, and then explain the concept of consistency. We then provide high level descriptions of the shared memory emulation problem and its solutions, where we explain some of the challenges of erasure coding based shared memory emulation.

#### 1.2.1. Distributed Systems Model

The shared memory emulation problem consists of client nodes and server nodes, where the clients nodes may be partitioned in to read and write clients (See Fig. 2). Write clients issue write




Figure 2. System model.




Figure 3. A typical execution of an ideal read-write memory.

operations, and read clients issue read operations<sup>2</sup>. We assume that there is a system with  $N$  server nodes and an arbitrary number of client nodes. The maximum number of server nodes that fail is  $f$ . The following assumptions distinguish the shared memory emulation model from commonly studied models of distributed storage systems in information theory research.

**A1—Arbitrary Asynchrony:** The nodes are all connected by (logical) point-to-point links<sup>3</sup> and the topology is assumed to be known to all nodes. The point-to-point links are themselves thought of as asynchronous links which are reliable, but their delay can be arbitrary and unbounded.

**A2—Nodes are computing devices:** The nodes are not merely storage devices but they are computing devices, and can therefore be used to execute fairly complex, interactive, protocols.

**A3—Decentralized Nature:** A node is unaware of the current state of any other node, and a node's knowledge of the state of another node is limited to what it can learn from the messages it receives on the incoming links.

In a system with  $N$  servers, we aim to design algorithms that operate correctly so long as the number of server failures no larger than some known threshold  $f$ . Note that because the system is asynchronous (assumption **A1**), a failed node cannot be distinguished from a very slow node; a node that does not obtain responses from another node cannot figure out if that node has failed or the messages from the node may simply be delayed. Furthermore, when a server node receives the updated data from a write client, it is not automatically aware of whether another node in the system has received this update yet. These aspects can make correct distributed algorithm design in the above model quite interesting and challenging. In practice, server nodes are typically in a cloud storage system, such as a data center, and client nodes are external devices or proxy/leader nodes in the storage system carrying out operations on behalf of an external clients.

<sup>2</sup>The split between read and write clients is a logical split—they can be physically be located on the same node.

<sup>3</sup>This system is often referred to as the *message passing* architecture.

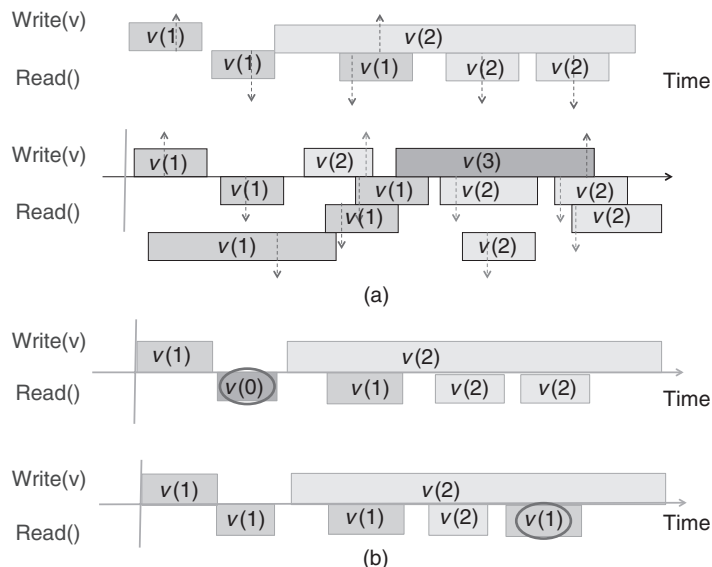


Figure 4. (a) Atomic executions: every operation “looks like” it took place instantaneously at some point in its interval, (b) Executions that are not atomic: we cannot place points in the intervals of operations and make the execution look like that of a correct instantaneous read-write memory.

**Concept of atomic consistency:** A typical execution of an *ideal* shared read-write memory is presented in Fig. 3. However, the execution of Fig. 3 is too idealistic to model in distributed systems, specifically, read and write operations are seldom instantaneous, but they take time. More specifically, the time scale of operation completion is often comparable to the inter-arrival time of operations; therefore operations may overlap<sup>4</sup>.

In systems where operations can overlap (e.g. Fig. 4), it is necessary to carefully define a set of rules that govern the possible outcomes of the write and read operations: such rules are known as *consistency* requirements. Consistency requirements are usually defined to ensure that (a) they are sufficiently relaxed so that they can be implemented in a realistic asynchronous distributed system, and (b) they are strict enough such that the overall execution is, in some sense, indistinguishable from one that could have taken place over an instantaneous read-write memory. For instance, an important requirement is that a read operation obtains the latest (in some sense) write operation. There are several formal ways of defining consistency, each of which is useful depending on the system and applications using the system. Here we give a high level explanation of an important consistency requirement known as *atomic consistency* or simply *atomicity* [54]<sup>5</sup>.

An execution is said to be atomic if every operation of the execution “looks like” it took place at some point in the interval of the operation. Examples of atomic executions are provided in Fig. 4a. Note that in each execution of Fig. 4a, we can place points in the interval of an operation such that, if we “shrink” the operation to that point, the overall execution will look like a correct execu-

<sup>4</sup>Operation overlap can be avoided through protocols that implement a *lock* on the system. However locks can make the system very slow [44], and a service that allows for concurrent overlapping operations is desirable.

<sup>5</sup>Atomicity is also referred to as *linearizability* [44] or strong consistency.

tion of an instantaneous read- write memory. Note that the executions of Fig. 4b do not satisfy this property, so they are not atomic. Implementations of atomic shared memory is important because it ensures modular application design. For instance, an application (such as a bank account or a reservation system) can be designed assuming the existence of instantaneous shared variables; using atomic variables in place of instantaneous variables continues to ensure that the overall application works correctly<sup>6</sup>.

### 1.2.2. Problem Statement and Solutions (Informal)

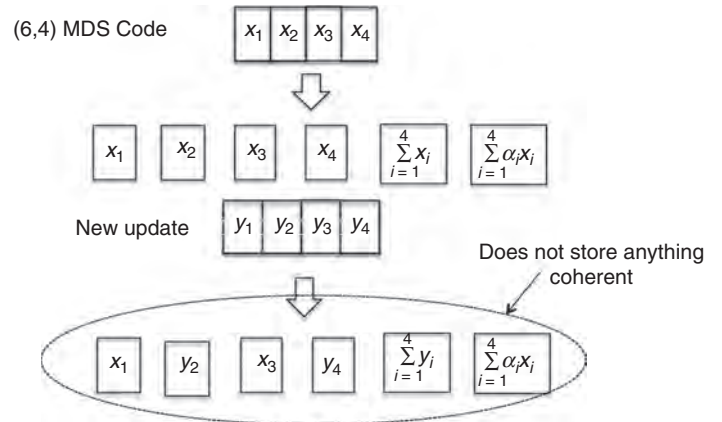
The goal of (atomic) shared memory emulation is to design protocols of write and read clients and servers so that the overall execution is always atomic, and every read and write operations eventually completes so long as the number of server failures is no bigger than  $f$ <sup>7</sup>. There are several solutions to the shared memory emulation problem as it is studied extensively in both theory and engineering research. A well-known solution comes from [8]; a reader interested in appreciating systems engineering aspects is also encouraged to read [24]. While [8, 24] use replication for fault tolerance, there is a rich literature in developing provably consistent erasure coding based shared memory emulation algorithms (a non-exhaustive list is ([1, 42, 27, 15, 16, 25, 52])), and also some system implementations [20]. For the purposes of the discussion here, we simply provide a high level understanding of the structure of typical shared memory emulation algorithms.

In shared memory emulation algorithms (e.g., [8, 35]), write operations send the new value to  $N$  server nodes and wait until getting acknowledgements from  $c_W$  servers before completing the operation. Similarly, a read operation sends a read request to all  $N$  servers and waits for responses from at least  $c_R$  nodes in the system before reading/decoding the value. To tolerate  $f$  server failures—that is, to ensure that write and read operations complete even if  $f$  servers fail—it is required that  $c_W, c_R \leq N - f$ . For every pair of complete write and read operations, there are at least  $c_R + c_W - N$  servers that received the value of the write operation, and responded to the read operation. Thus, for repetition based schemes [8, 24], where each server simply stores the latest version it receives in an uncoded form, choosing  $c_R, c_W$  so that  $c_R + c_W > N$ , suffices to ensure that every read operation that begins after a write operation completes will “see” the value of the write, or see a later value. Indeed, in such algorithms, if  $c_R + c_W > N$ , the replication-based protocols ensure atomic consistency. A natural extension of this principle indicates that when a maximum distance separable (MDS) code is used [42, 27, 16, 25], the dimension of the erasure code is chosen to be  $c_R + c_W - N$ . Thus, based on the structure of prior protocols, the requirement of ensuring consistency can be expressed as follows<sup>8</sup>:

<sup>6</sup>In practice, systems are sometimes designed with weaker requirements as compared with atomicity for the sake of better performance, i.e., faster read/write operations, see for e.g., [83, 11, 63],

<sup>7</sup>This property is known as the *non-blocking* or *wait-freedom*, property, since it also implies that overlapping operations must be able to proceed without blocking each other.

<sup>8</sup>The reader may note from Fig. 4a that ensuring consistency is more complex and involved than assumption **A4**, and requires careful protocol design; however, for the sake of the discussion here, we simply use **A4** as a proxy for the desired consistency criteria. Also, property **A4** assumes that there is some ordering of the writes/versions; protocols such as [8, 35] also take steps to ensure a non-unique ordering of the write operations.



**Figure 5. Use of erasure coding for shared memory emulation. The figure demonstrates why a server which receives a new updated codeword symbol cannot simply replace the stored codeword symbol.**

**A4—Consistency:** The latest among all the versions that have been propagated to at least  $c_W$  servers, or a later version, must be decodable from any set of  $c_R$  servers.

Erasure coding presents interesting algorithmic and coding challenges in asynchronous distributed systems where consistency is important. This is because, when erasure coding is used, no single server stores the data in its entirety; for instance, if a maximum distance separable (MDS) code of dimension  $k$  is used, each server only stores a fraction of  $1/k$  of the entire value. Therefore, for the read client (decoder) to decode some version of the data, at least  $k$  servers must send the codeword symbols corresponding to the same version. In particular, when a write operation updates the data, a server cannot delete the old version before ensuring that the new version has propagated to a sufficient number of servers. As a consequence, servers cannot simply store the latest version they receive; they have to store older versions at least until a sufficient number of codeword symbols corresponding to the newer version has been propagated (See Fig. 5). In fact, in situations where there are multiple concurrent writes, servers may have to store multiple codeword symbols, one corresponding to each concurrent operation. Indeed, the main algorithmic challenge of erasure coding based shared memory emulation algorithms developed previously, is to determine in a decentralized setting that a new version has propagated to a sufficient number of servers before deleting older versions [42, 27, 15, 16, 25, 53, 52, 15, 2, 1, 36]. Importantly for our purposes, the fact that servers have to store multiple versions can offset the storage cost gains of erasure coding, especially when the degree of concurrency is high. Next we summarize *multi-version codes* [87, 88, 89], which study the storage cost of shared memory emulation from an information-theoretic perspective.

### 1.3. A Coding Framework Related to Shared Memory Emulation

We here describe a coding framework called multi-version coding that simplifies the shared memory emulation problem and yet keeps its essential aspects that pertain to its storage cost. Specifically, in addition to fault tolerance, the multi-version

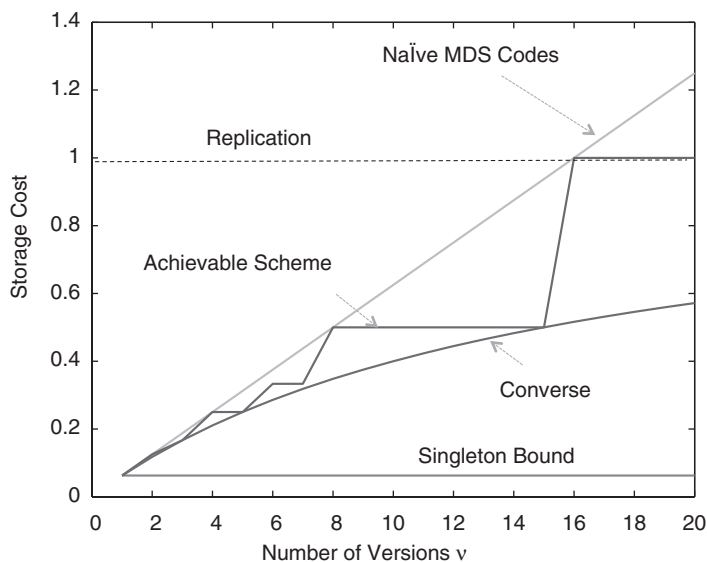


Figure 6. Summary of results of [89].

coding framework incorporates arbitrary asynchrony (A1), decentralized nature (A3) and consistency (A4) as discussed in Sections 1.2.1, 1.2.2. We will discuss modeling assumption (A2) in Section 1.4. The multi-version coding problem is parametrized by positive integers  $n, c_W, c_R, v$ . There are  $n$  servers, and there is a message/variable with  $v$  versions  $W_1, W_2, \dots, W_v$ . The versions are assumed to be totally ordered as  $W_1 < W_2 < \dots < W_v$ , with a higher ordered message version being interpreted as a more recent version as compared to a lower ordered version. We model asynchrony (A1) as follows: Every server receives an arbitrary subset of the versions (at one shot) and encodes the received versions. More formally, let  $\mathcal{S}_i \subseteq \{1, 2, \dots, v\}$  be the set of versions received by the  $i$ th server. Let  $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_n)$ ; note that  $\mathcal{S}$  is an element in the power set of  $\{1, 2, \dots, v\}^n$ . We refer to  $\mathcal{S}_i$  the *state* of server  $i$ , and  $\mathcal{S}$  as the *global state of the system*. Denoting  $\mathcal{S}_i = \{s_1, s_2, \dots, s_m\}$  the  $i$ -th server stores a codeword symbol generated by an encoding function  $\phi_{\mathcal{S}_i}^{(i)}$  that takes an input,  $(W_{s_1}, W_{s_2}, \dots, W_{s_m})$ , and outputs an element in  $\mathcal{Q}$ , where  $\mathcal{Q}$  is the alphabet of coding. Note that the decentralized nature (A3) of the system is implicit in the model since server  $i$ 's encoding function  $\phi_{\mathcal{S}_i}^{(i)}$  depends on  $\mathcal{S}_i$  alone, and does not depend on the the global state vector  $\mathcal{S}$ .

**Decoding Constraint:** For any state  $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_n)$  we refer to a message version that has propagated to at least  $c_W$  servers as a *complete* version. The server encoding functions  $\phi_{\mathcal{S}_i}^{(i)}, i = 1, 2, \dots, n$  have to be designed such that, the decoder must be able to recover from any  $c_R$  servers, the latest complete version as per the ordering  $\prec$ , or a later version. That is, a decoder must be able to decode, from any  $c_W$  servers,  $W_\ell$  where  $\ell \geq \max \{j : |\{i : j \in \mathcal{S}_i\}| \geq c_W\}$ . The decoding constraint has been defined based on consistency requirements (A4).

**Achievable schemes:** With replication, the storage cost per server is the size of one version. Now, suppose we aim to use an MDS code. For a decoder that connects to  $c_R$  servers, a complete version is present at least  $c_W + c_R - n$  of those servers. So using a dimension of  $c_W + c_R - n$  suffices for the decoder to decode the complete version. Furthermore, note that storing simply the latest version at a server

does not suffice, since from a server's perspective, it is not aware of which version is complete. Therefore, a server has to store a codeword symbol of a code of dimension  $c_W + c_R - n$  for each version it receives. In the worst case, the storage cost per server is  $(v/c_W + c_R - n)$ . In our preliminary work, we have improved upon the best of replication and erasure coding by varying the dimension of the code used for the  $v$  versions. Our results are depicted in Fig. 6. To appreciate the jagged nature of the achievable scheme, it can be verified that each server storing simply the latest version with dimension  $\lceil \frac{c_R + c_W - n}{v} \rceil$  suffices. The converse depicted in Fig. 6 carries an instructive message: *in distributed storage systems, in addition to fault tolerance, there is an inevitable price to be paid in terms of redundancy overhead to maintain consistency*. The converse derivation is combinatorial in nature, discovering the worst-case states that force a lower bound on the storage cost.

### 1.3.1. Where Do Classical Distributed Storage Erasure Codes Fit?

Classical erasure codes may be derived as a special case of multi-version codes for  $v = 1$ . From a modeling perspective, this translates to the system being synchronous, or completely centralized. For instance, if the system is synchronous, then all non-failed servers receive each version completely. Then storing simply the latest version with dimension  $c_R$  code suffices for the decoder. Even if the system is asynchronous, i.e., a server receives an arbitrary subset of the  $v$  versions in the system, a centralized system where each server knows the global system state leads to the classical erasure coding framework. To see this, observe that with global system state information, each server is aware of the the latest complete version. Therefore, a server with the latest complete version simply encodes it using a code of dimension  $c_W + c_R - N$  and stores the corresponding codeword symbol; a server that does not possess the latest common version does not store any information. These examples show that new coding ideas are required only when the model accounts for both the asynchronous nature and decentralized nature of the system.

## 1.4. Extensions and Open Problems

Classical erasure codes for distributed storage may be interpreted as an optimistic view of the system: it assumes that the system is synchronous, and that nodes have instantaneous and global system state information. In contrast, the multi-version coding framework takes a pessimistic/conservative view of storage systems (Fig. 7). Specifically the model assumes that the system is completely asynchronous, i.e., every version arrival pattern is possible; it also assumes that nodes do not have even stale or partial information of system state. However, in practice, we expect the system to operate somewhere in between these two extremes, specifically nodes may be able to opportunistically obtain stale or partial information of the system. This opens the door to many important open questions that enable code constructions.

- The multi-version coding is a worst-case formulation, where the decoding requirement as well as the storage cost characterization applies to every possible state. An open question is the potential storage cost reductions that may be obtained by studying the average storage cost, and/or allowing for a small probability of decoding error after invoking an appropriate probability measure on the state space.

- In several settings, either because of the topology or the internal gossip messages exchanged among the servers, servers obtain local and possibly stale side information (e.g., a server may know which server has some of the older versions.) In the extreme case where server has global and instantaneous side information, the system reverts back to the classical setting as explained in Section 1.3.1. The question of how to design codes that exploit stale or local side information about the states, and a quantitative understanding of the potential storage cost gains is a promising area of future work.
- The multi-version coding setting assumes that the different versions are independent. In some applications, subsequent writes and reads tend to have correlations that can be exploited to reduce the memory overhead in such applications. Our preliminary work [5] explores this direction and derives achievable schemes and characterizes their storage cost. However, questions related to optimality and practical low complexity code constructions remain open.

The second direction of open problems comes from understanding and developing connections between distributed computing systems and multi-version codes. We describe some of them briefly here.

- In our recent work [17], we have proved formal information theoretic lower bounds on the shared memory emulation problem inspired by converse results for the multi-version coding framework. Specifically, the result of [17] shows that the multi-version coding converse applies to the general shared memory emulation problem when  $v = 2$ , and for non-interactive protocols when  $v > 2$ . It is not known whether interactive protocols can improve upon the storage cost of multi-version coding—the reason for this gap in understanding is that multi-version coding does not capture assumption **A2** in Section 1.2.1 very well. While studying this question in the context of the general shared memory emulation problem might be challenging, a possible starting point is by making the toy model of [89] interactive.
- The shared memory emulation problem is a special case of the general *replicated state machine* [60, 9] problem in distributed computing, where the goal is to emulate an arbitrary state machine in a distributed asynchronous system in a fault tolerant manner. The replicated state machine problem is much more complicated than the shared memory emulation problem, and is an object of extensive study in distributed systems literature. Like shared memory emulation, solutions to the replicated state machine problem—for specific state machines—form the basis of several cloud based services (e.g., Google Spanner [21]). However, the question of how to encode state machines is not yet well studied by either the distributed systems or the coding theory communities. A good starting point is in reference [12], which generalizes basic coding theoretic concepts such as the Hamming distance to replicated state machines.

In addition to the above mentioned extension and open problems, a deeper understanding of the system from a networking and resource allocation viewpoint, which can guide an engineer on how




Figure 7. Open areas of research.

to choose coding parameters such as the length, dimension, and the extent of history to be stored, based on the total storage budget, frequency of updates, the degree of asynchrony, and other system level parameters is a broadly open and important area of research.

## 2. Part II: Coding for Distributed Computing Using Unreliable Components: Errors, Faults, Stragglers

The roots of computing using unreliable components go as far back as work of von Neumann in 1956 to a work [84]. He considers a problem where each gate is unreliable, and provides replication-type strategies with repeated majority decoding make the overall computation reliable. von Neumann’s work was inspired both by Shannon’s 1948 paper [72] and by his interest in understanding biological computation<sup>9</sup>, e.g., in our nervous system<sup>10</sup>. von Neumann’s work was followed up in a sequence of works, including those of Pippenger [66, 67], Taylor [77] (see also [78]), Hadjicostis [37], and Hajek and Weller [40] who provided achievable strategies for specific operations. Taylor’s work started the area of decoding using unreliable circuits, focusing on low-density parity-check codes of Gallager, and allowing for errors in his iterative decoding algorithm. This work has been followed up extensively recently, e.g. Varshney [80], Vasic et al. [81, 82], Dolecek et al. [94, 46], with an interesting recent work of Vasic showing that low-error rates in the decoding implementation can help *improve* the performance of decoding. Pippenger’s work [66] not only provided achievable strategies for linear computations, it also utilized work of Dobrushin and Ortyukov [26] to lay down fundamental limits by tightening the data-processing inequality in the context of scalar random variables. This direction on fundamental limits was extended and explored in the works of Evans and Schulman [33], and Erkip and Cover [32] on “quantified/strong” data-processing inequalities, that has led to recent resurgence in this direction, with broadening of problems to “information-dissipation” [68, 18, 6, 69].

In a deeply related but separate body of work, inspired by distributed and parallel computing systems that have been implemented in the last few decades, there has been significant work on addressing computing under “processor” unreliability, where a processor is assumed to be constituted by a large number of gates and attached memories.

<sup>9</sup>Shannon himself was examining the noisy computing problem at the same time, focusing on unreliable relays [62].

<sup>10</sup>It is now widely believed that the brain indeed uses error-correcting codes [73, 64] (his paper is titled “Probabilistic Logics and the Synthesis of Reliable *Organisms* from Unreliable Components”). Whether it does so to compute efficiently is far from established, although the thought behind Barlow’s famous and controversial “Efficient Coding Hypothesis” does suggest so [13].

Addressing processor unreliability is even today thought to be “one of top 10 challenges for exascale computing” [59] (i.e., larger-scale supercomputing, where processors can yield undetected errors). Coding-theoretic work in this direction was initiated by Huang and Abraham [47], who named it “Algorithm-Based Fault Tolerance” (ABFT), which is today a thriving research area with hundreds of papers (e.g. [43, 49, 37, 50]). These works provide techniques of error-correction for large-scale computing systems, including modern supercomputing systems. A closely related application area is distributed/cloud computing systems, where unreliability manifests itself in the form slow processing nodes (called *stragglers*), which significantly slow down the entire computation [23]. Use of coding techniques in cloud computing systems to address straggler bottlenecks was pioneered in the recent work of Lee et al. [56]. In [56], the authors demonstrated the power of ABFT-like coding techniques for linear computations via novel expected time-analysis based on exponential-tail models of processing times, in addition to experimental results. Interestingly, the coding techniques developed in the ABFT literature turn out to be suboptimal in the sense of error/erasure-tolerance for operations such as matrix-vector and matrix-matrix multiplication. Recent work in information theory has advanced on these constructions while obtaining, in some cases, fundamental limits as well. For instance, for the problem of matrix-matrix multiplication (Section 2.2), the work of Yu, Avestimehr, and Maddah-Ali [95] provides a coded computing construction that they call “Polynomial coding,” that provides scaling-sense improvements on ABFT by comparing ABFT’s performance with fundamental limits. This was followed by our own recent work [34] which provides scaling sense improvements on Polynomial codes. The information-theoretic approach has much to offer.

More recently, coded computing results have been obtained for convolutions [91, 31, 74, 95], solving linear inverse problems and PageRank [92], distributed gradient descent [30, 75, 76, 70, 41], linear regression and classification [56, 30], logistic regression [90], distributed iterative optimization [51, 10], and even separable non-linear functions [57], etc. For any computation that can be split into tasks, optimized approaches have been proposed by Joshi, Wornell, and Soljanin, and collaborators [86, 48, 85, 3, 4] for straggling processors, which focus on problems of task allocation using rigorous queuing-theoretic models. Coding techniques have also been used to reduce communication requirements even when processing nodes are completely reliable, and a notable work here is that on “Coded MapReduce” by Li, Avestimehr, and Maddah-Ali [58].

In the following, we will discuss the development of the field from the perspective of three key problems. In Section 2.1, we will discuss matrix-vector multiplication under processor level errors or stragglers [47, 56, 30, 75, 76, 70, 41] as well as gate-level errors [66, 93]. In Section 2.2, we will discuss advances in matrix-matrix multiplication, including recent works [47, 95, 34]. Finally, in Section 2.3, we will end with a recent coded-computing construction of Dutta et al. [29] that connects back to the inspiration of von Neumann’s seminal work: the brain. This construction provides reliability to large networks of the nonlinear McCulloch-Pitts neuron models that are revolutionizing computing and inference systems today.

## 2.1. Coded Matrix-Vector Multiplication

The problem of computing linear transforms of high-dimensional vectors is *the* critical step [22] in several machine learning and signal processing applications. Dimensionality reduction techniques

such as Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), taking random projections, etc., require the computation of short and fat linear transforms on high-dimensional data. Linear transforms are the building blocks of solutions to various machine learning problems, e.g., regression and classification etc., and are also used in acquiring and pre-processing the data through, e.g., filtering. Fast and reliable computation of linear transforms are thus a necessity for low-latency inference [22].

**Large processing nodes:** The goal is to compute the product  $\mathbf{A}\mathbf{y}$  of the matrix  $\mathbf{A}$  with the vector  $\mathbf{y}$ . We first discuss the case with “large” processing nodes, where e.g., the memory-size of each processing node can scale with the problem-size. The first strategy here was proposed in the seminal work of Huang and Abraham work on ABFT [47]. The idea of using coding to mitigate the *straggler effect* in distributed systems was pioneered by [56], which used what happens to be a special case of ABFT to speed up matrix-vector products. The authors in [56] show, both analytically (using models inspired from distributed systems) and experimentally (on Amazon’s EC2 cluster), that significant speed-ups in expected overall computation time can be obtained by using these techniques. The core idea of [56] is to code each column of the matrix  $\mathbf{A}$  by multiplying it with the same generator matrix of a linear code, resulting in a coded matrix  $\mathbf{A}_{\text{coded}}$ . Now, we distribute the rows of the matrix  $\mathbf{A}_{\text{coded}}$  among different processing nodes, and each node computes its own matrix-vector product. Because linear combinations of codewords of a linear code are also codewords, the resulting concatenation of outputs is also a codeword. Hence, the computation can proceed without waiting for a few straggling nodes by using erasure decoding to fill in for their outputs.

What if, due to communication or memory bottlenecks, each node cannot even compute a whole dot product of a row of of an  $M \times N$  matrix  $\mathbf{A}$  with the  $N \times 1$  vector  $\mathbf{y}$ ? Recently, our work [30] as well as Tandon et al. [75] (motivated by a different application) simultaneously arrived at strategies for coding matrix-vector multiplication by encoding the matrix into a sparse coded matrix. These codes, that we call “Short-Dot” codes because they distribute large matrix-vector products to processing nodes that compute short dot products, are illustrated below through an example.

Consider a  $2 \times 4$  matrix  $\mathbf{A} = [\mathbf{a}_1^T \ \mathbf{a}_2^T]^T = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \end{bmatrix}$  and a

$4 \times 1$  vector  $\mathbf{y} = (y_1, y_2, y_3, y_4)^T$ . Can we compute  $\mathbf{A}\mathbf{y}$  over  $n = 4$  nodes such that, (i) each node uses a shortened version of  $\mathbf{y}$ , i.e., at most three of four scalars  $y_1, y_2, y_3, y_4$ , and (ii) the overall computation is tolerant to one straggler, i.e., 2 of the three nodes suffice to recover  $\mathbf{A}\mathbf{y}$ ? Short-Dot codes use the following strategy: Node  $i$  computes  $(\mathbf{a}_1 + \mathbf{a}_2 + \mathbf{z}^i)\mathbf{y}$ ,  $i = 1, 2, 3, 4$  so that from any 3 of the 4 nodes, the polynomial  $p(x) = (\mathbf{a}_1\mathbf{y} + \mathbf{a}_2\mathbf{y}x + \mathbf{z}\mathbf{y}x^2)$  can be interpolated. Vector  $\mathbf{z} = (z_1, z_2, z_3, z_4)$  is chosen to satisfy  $\mathbf{a}_{1i} + \mathbf{a}_{2i} + \mathbf{z}_i^2 = 0$  for  $i = 1, 2, 3, 4$ , so that node  $i$  does not require  $y_i$ .

To understand the performance of Short-Dot from a quantitative perspective, consider a setting shown in Fig. 8 where each processing/worker node stores<sup>11</sup>  $MN/m$  linear combinations of the entries of  $\mathbf{A}$  and  $N/n$  non-zero entries of  $\mathbf{y}$ . Assume that  $P$  processing nodes perform the computation, we use “recovery threshold”  $K(m, n)$ —the minimum number of processing nodes required by a fusion node to recover the matrix-vector product  $\mathbf{A}\mathbf{y}$ —to measure

<sup>11</sup>Although we use the word “store” in our descriptions, note that the restrictions on the amount of  $\mathbf{A}$  and  $\mathbf{x}$  need not be only due to memory overhead; these restrictions may be motivated by communication costs as well.



its resilience. Note that the maximum number of straggling nodes that a scheme can tolerate is  $P - K(m, n)$ .  $K(1, 1)$  is clearly 1, and translates to a replication strategy. In the MDS code based

strategy of [56, 47],  $\mathbf{A}$  is split as  $\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_m \end{bmatrix}$  and node  $i \in \{1, 2, \dots, P\}$

computes  $(\sum_{j=1}^m g_{ij} \mathbf{A}_j) \mathbf{y}$ , where  $\mathbf{G} = (g_{ij})$  forms the generator matrix of a  $(P, m)$  MDS code<sup>12</sup>. The fusion node can recover  $\mathbf{A} \mathbf{y}$  from any  $m$  nodes, which gives  $K(m, 1) = m$ . Interestingly, going beyond  $n = 1$ —that is, using the full vector  $\mathbf{y}$  at all nodes—requires interesting coding strategies as demonstrated by the above example. Short-Dot obtains a recovery threshold of  $K(m, n) = (m/n) + P(1 - (1/n))$ ; via a converse result, this recovery threshold is approximately optimal [28, 30].

In fact, the number of non-zero entries in the stored, coded submatrix at each processing node is  $MN/mn$  since every row of the coded matrix only has  $N/n$  non-zero entries at *pre-defined* locations. Thus it requires even smaller memory than  $MN/m$ . The key difference between Short-Dot and MDS codes is that Short-Dot allows the vector  $\mathbf{y}$  to be shorter than its full-length  $N$ , thus allowing for cheaper communication of parts of  $\mathbf{y}$  to each processing node and computation of shorter dot-products at each processing node as compared to MDS code.

Concurrent work [75] also discovered *gradient coding*, which incorporates the essential coding idea of Short-Dot in the context of a single dot product in an important application: distributed gradient descent over data  $\mathbf{D} = (\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_P)$  where  $\mathbf{D}_i$  is the data stored in the  $i$ -th processing node, the gradient update step requires each processing node to compute a gradient  $g(\mathbf{D}_i)$ , which is used by a master node to compute  $\sum_{i=1}^P g(\mathbf{D}_i)$ . In [75], Tandon et al. view the operation of a master node as a single *dot-product* of the vector  $[1 \ 1 \ \dots \ 1]$  with the vector  $[g(\mathbf{D}_1) \ g(\mathbf{D}_2) \ \dots \ g(\mathbf{D}_P)]$ . They use a coding technique similar to Short-Dot to complete the iterations in presence of straggling worker nodes. They also introduce the notion of *partial stragglers*, [75] which opens up new research directions by modeling the fact that stragglers—unlike faulty nodes—eventually complete their operation and their outputs can be used in the eventual computation.

A direct application of matrix-vector multiplication is solving sparse linear inverse problems using iterative matrix-vector products, such as for finding eigenvectors of a matrix through, e.g., the PageRank algorithm. What is new in these problems is that the answer slowly converges to the true solution. The stragglers thus simply have fewer iterations compared to fast nodes, but their outputs are still useful. In [92], we use a decoding algorithm inspired by weighted least-squares to fully utilize the results from both fast nodes and stragglers by assigning weights to different nodes based on their proximity to convergence. Compared to erasure-coding based coded computing, this coding method achieves graceful degradation of remaining error with increasing number of stragglers.

<sup>12</sup>In fact, aggregating these coded matrices forms an MDS coded matrix,  $\mathbf{A}_{\text{coded}}$ . However, we use the term  $\mathbf{A}_{\text{coded}}$  more generally to denote a coded version of the matrix  $\mathbf{A}$  for an arbitrary linear code, not necessarily an MDS code.

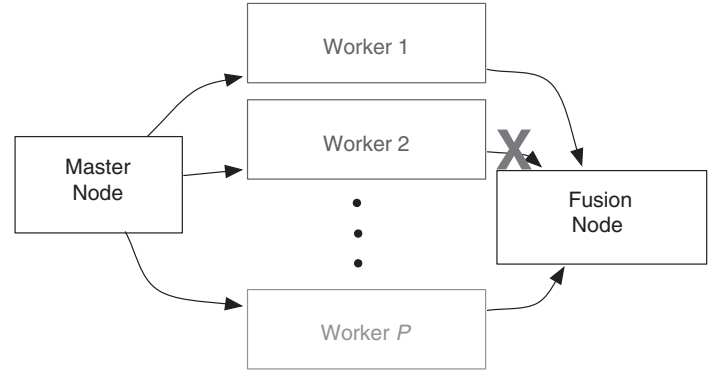


Figure 8. System model.

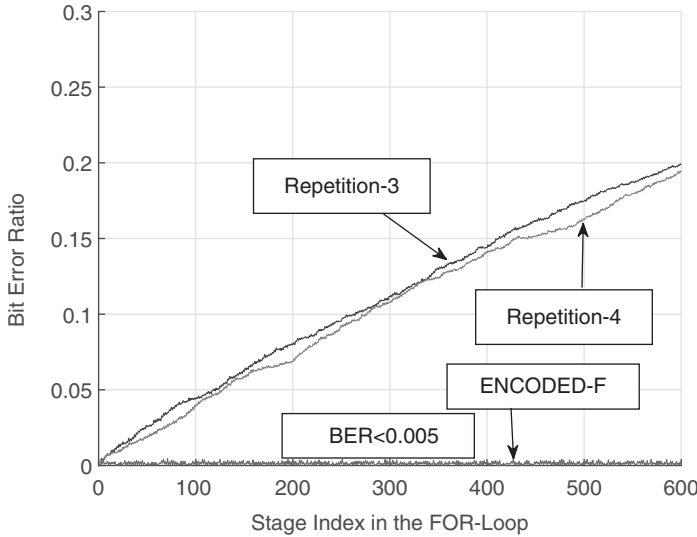
Both theoretical and experimental results show substantial advantages of the proposed algorithm over replication schemes<sup>13</sup>. Surprisingly, if the ordinary matrix-inverse algorithm is used for decoding, the remaining error shoots up because MDS coding matrices appear to be ill-conditioned, which is also recognized in the work of Haikin and Zamir [39, 38] on analog coding for erasure channels. By carefully combining the (partial) results of the straggling nodes (instead of ignoring them), we are able to circumvent the difficulty of ill-conditioned matrix inverse and achieve orders of magnitude reduction of remaining error in experiments.

**Small processing nodes (i.e., single gates):** When considering models with gate-level errors, the cost of error-detection can be significant, so it is important for the models to consider errors. Further, the implementation has to be fully decentralized: decoding itself can suffer from errors.

For this problem, the fundamental limits of “strong” data-processing inequality provide an important intuition: along any single “path” in the circuit, errors will accumulate, and information will dissipate. This seems to present a pessimistic picture. Is there any hope for reliable computation in the Shannon sense [72]? First, note that the overall computation error-probability cannot be lower than the error-probability of the last gate, so the best we can hope for is that the overall error probability is close to the last gate’s error probability. For the specific problems of binary matrix-vector multiplication, our recent work [93] shows that even with all gates noisy, this is achievable, and further, that sophisticated error-correction techniques that our community has developed offer scaling-sense advantages.

To describe our technique, it is useful to first discuss a fundamental limit, namely, Lemma 2, in the work of Evans and Schulman [33]. They derive an upper bound on mutual information between a binary input and the values carried by a set of wires by accumulating mutual information across those wires. This suggests an achievability: if the number of paths that a circuit can keep accumulating information from keeps increasing as the information dissipates along each individual path, it may

<sup>13</sup>It is interesting that the diversity gain achieved by the replication strategy here is smaller than that can be achieved in communications over random fading channels, because the results at stragglers are deterministic functions of those at fast nodes for the same inverse problem. Fundamental limit on obtained diversity gains are worthy of further investigation.



**Figure 9. A comparison of attained error-probability by our technique ENCODED with repetition/replication-type approaches. ENCODED is able to keep errors bounded even as the computation proceeds.**

be possible to offset the losses with the gains, and keep errors suppressed.

Our results [93] show that this is indeed attainable for computing a binary linear transform  $\mathbf{A}\mathbf{y}$ . To do so, we encode  $\mathbf{A}$ , obtaining  $\mathbf{A}_{\text{coded}}$  with coded columns, that is, each column is a codeword. The coding technique itself is a carefully chosen LDPC code. Now, instead of performing distributed dot-products of individual rows of  $\mathbf{A}_{\text{coded}}$  with  $\mathbf{y}$ , we perform scalar-vector multiplications of  $y_i$ , the  $i$ -th element of  $\mathbf{y}$ , with the  $i$ -th column of  $\mathbf{A}_{\text{coded}}$ . Note that the resulting scalar-vector product is a codeword itself. The key step is this next one: we add these codeword vectors across different indices  $i$  in a tree architecture, and at every intermediate node in the tree, we embed a noisy decoder (utilizing results from [94, 46]). These decoders enable repeated error-suppression as the computation advances (see Fig. 9), and the accumulation of information from various paths enables the computation to compensate for information dissipation. The resulting error probability is shown, through theoretical results and simulations, to remain bounded by a constant throughout the computation. Because the repeated error-suppression using embedded decoders is a critical aspect of our strategy, we call it “ENcoded Computation with Decoders Embedded,” or “ENCODED”.

## 2.2. Coded Matrix Multiplication

In this section, we focus on the problem of multiplying two  $N \times N$  matrices<sup>14</sup>  $\mathbf{A}$ ,  $\mathbf{B}$ . Consider the setting where each processing node stores  $N^2/m$  linear combinations of the entries of  $\mathbf{A}$  and  $N^2/n$  linear combinations of the entries of  $\mathbf{B}$ . Assume that  $P$  processing nodes perform the computation, we evaluate the straggler tolerance of a technique by its recovery threshold  $K(m, n)$  of the technique. As before,  $K(1, 1)$  is clearly 1, and translates to a repetition based strategy.

<sup>14</sup>We assume that both matrices are square for the sake of simplicity. The results of this section apply for multiplication of matrices of arbitrary dimension for mild assumptions on the matrix dimensions.

	Worker (:, 1)	Worker (:, 2)	Worker (:, 3)
Worker (1, :)	$\mathbf{A}_1$	$\mathbf{B}_1$	$\mathbf{B}_2$
Worker (2, :)	$\mathbf{A}_2$		
Worker (3, :)	$\mathbf{A}_1 + \mathbf{A}_2$		

**Figure 10. ABFT matrix multiplication [47] for  $P = 9$**

processing nodes with  $m = n = 2$ , where  $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$ ,  $\mathbf{B} = [\mathbf{B}_1 \ \mathbf{B}_2]$ .

The recovery threshold is 6.

A natural generalization of the MDS code based strategy, which we call *one-dimensional MDS coding*, where  $\mathbf{A}$  is encoded using and MDS code gives  $K(m, 1) = m$ ; similarly, encoding  $\mathbf{B}$  gives  $K(1, n) = n$ . We review here three strategies (i) ABFT matrix multiplication [47] (also called *product-coded matrices* in [55]), (ii) Polynomial codes [95] and (iii) MatDot codes [34] each with successively improving, i.e., smaller, recovery threshold. Rather than go into the technical details, we give three examples for the case where  $m = n = 2$ —i.e., each processing node stores half of  $\mathbf{A}$  and half of  $\mathbf{B}$ —that convey the ideas used. We begin by describing ABFT matrix multiplication.

**Example (ABFT Codes [47], Fig. 10, recovery-threshold = 6)** Consider two  $N \times N$  matrices

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \mathbf{B} = [\mathbf{B}_1 \ \mathbf{B}_2]$$

Can we compute  $\mathbf{A}\mathbf{B}$  over  $P$  nodes such that, (i) each node uses one linear combination of  $\mathbf{A}$  and one linear combination of  $\mathbf{B}$  and (ii) the overall computation is tolerant to  $P - 6$  stragglers, i.e., 6 nodes suffice to recover  $\mathbf{A}\mathbf{B}$ ? ABFT codes use the strategy as per Fig. 10, where 4 of the 9 worker nodes compute  $\mathbf{A}_i\mathbf{B}_j$ ,  $i, j \in \{1, 2\}$  and the remaining worker nodes compute  $\mathbf{A}_i(\mathbf{B}_1 + \mathbf{B}_2)$ ,  $(\mathbf{A}_1 + \mathbf{A}_2)\mathbf{B}_i$ ,  $(\mathbf{A}_1 + \mathbf{A}_2)(\mathbf{B}_1 + \mathbf{B}_2)$  for  $i = 1, 2$  respectively. The general principle of ABFT is to encode the rows of  $\mathbf{A}$  and the columns of  $\mathbf{B}$  using systematic MDS codes of dimension  $m, n$  respectively.

The question is whether the recovery threshold of 6 is optimal was addressed in [95], which gave the following, elegant, polynomial code construction.

**Example (Polynomial Codes [95], Fig. 11, recovery-threshold = 4)** Consider two  $N \times N$  matrices

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \mathbf{B} = [\mathbf{B}_1 \ \mathbf{B}_2]$$

can we compute  $\mathbf{A}\mathbf{B}$  over  $P$  nodes such that, (i) each node uses one linear combination of  $\mathbf{A}$  and one linear combination of  $\mathbf{B}$  and (ii) the overall computation is tolerant to  $P - 4$  straggler, i.e., 4 nodes suffice to recover  $\mathbf{A}\mathbf{B}$ ? Polynomial codes use the following strategy: Node  $i$  computes  $(\mathbf{A}_1 + \mathbf{A}_2i)(\mathbf{B}_1 + \mathbf{B}_2i^2)$ ,  $i = 1, 2, \dots, P$ , so that from any 4 of the  $P$  nodes, the polynomial  $p(x) = (\mathbf{A}_1\mathbf{B}_1 + \mathbf{A}_2\mathbf{B}_1x + \mathbf{A}_1\mathbf{B}_2x^2 + \mathbf{A}_2\mathbf{B}_2x^3)$  can be interpolated. Having interpolated the polynomial, the coefficient (matrices) can be used

to evaluate  $\mathbf{A}\mathbf{B}$  as  $\begin{bmatrix} \mathbf{A}_1\mathbf{B}_1 & \mathbf{A}_1\mathbf{B}_2 \\ \mathbf{A}_2\mathbf{B}_1 & \mathbf{A}_2\mathbf{B}_2 \end{bmatrix}$ .

In [34], we improve upon the recovery threshold of polynomial codes through a construction called *MatDot* codes. Unlike ABFT

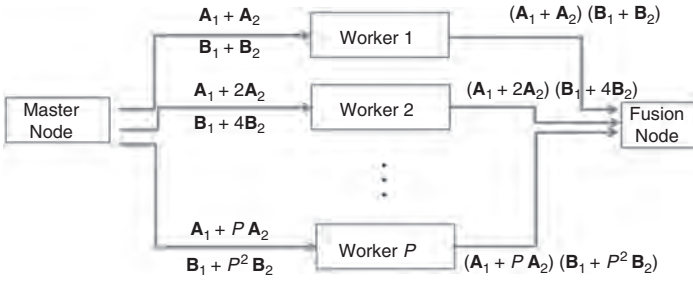


Figure 11. Polynomial Codes [95] with  $m = n = 2$  where

$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \mathbf{B} = [\mathbf{B}_1 \ \mathbf{B}_2]$ . The recovery threshold is 4.

and polynomial codes, MatDot splits the matrix  $\mathbf{A}$  column-wise and matrix  $\mathbf{B}$  row-wise.

**Example (MatDot Codes [34], Fig. 12, recovery-threshold = 3.)**  
Consider two  $N \times N$  matrices

$$\mathbf{A} = [\tilde{\mathbf{A}}_1 \ \tilde{\mathbf{A}}_2], \mathbf{B} = \begin{bmatrix} \tilde{\mathbf{B}}_1 \\ \tilde{\mathbf{B}}_2 \end{bmatrix}$$

can we compute  $\mathbf{AB}$  over  $P$  nodes such that, (i) each node uses one linear combination of  $\mathbf{A}$  and one linear combination of  $\mathbf{B}$  and (ii) the overall computation is tolerant to  $P - 3$  straggler, i.e., 3 nodes suffice to recover  $\mathbf{AB}$ ? MatDot codes use the following strategy: Node  $i$  computes  $(\tilde{\mathbf{A}}_1 + \tilde{\mathbf{A}}_2 i)(\tilde{\mathbf{B}}_1 i + \tilde{\mathbf{B}}_2)$ ,  $i = 1, 2, \dots, P$ , so that from any 4 of the  $P$  nodes, the polynomial  $p(x) = \tilde{\mathbf{A}}_1 \tilde{\mathbf{B}}_2 + (\tilde{\mathbf{A}}_1 \tilde{\mathbf{B}}_1 + \tilde{\mathbf{A}}_2 \tilde{\mathbf{B}}_2)x + \tilde{\mathbf{A}}_2 \tilde{\mathbf{B}}_1 x^2$  can be interpolated. Having interpolated the polynomial, the product  $\mathbf{AB}$  is simply the coefficient of  $x$ , that is  $\mathbf{AB} = \tilde{\mathbf{A}}_1 \tilde{\mathbf{B}}_1 + \tilde{\mathbf{A}}_2 \tilde{\mathbf{B}}_2$ .

In fact, if  $m = n = o(\sqrt{P})$  then the recovery threshold of one dimensional MDS, ABFT matrix multiplication, polynomial codes and MatDot codes is respectively  $\Theta(P), \Theta((m-1)\sqrt{P}), \Theta(m^2), \Theta(m)$ . There are differences between the schemes in terms of node communication costs and decoding costs; these and other related aspects are described in [34]. [34] also shows how to apply these ideas for multiplying more than two matrices, and generalizes and unifies Polynomial codes and MatDot codes to “PolyDot” codes that tradeoff between communication complexity and recovery threshold.

### 2.3. Coded Neural Networks

Deep Neural Networks (DNNs) are being extensively used in many inference applications. A DNN is a multilayer network of artificial neuron models (first developed by McCulloch and Pitts), where each layer performs a linear transform on its input vector from the previous layer. Training of DNNs is extremely time-intensive, and has two main steps: “feed-forward step” and “back-propagation.” In the feed-forward step, after computing a linear transform on the vector received from the layer to its left, each layer performs scalar nonlinear computations on each element of the vector that results from this linear transform, and then forwards it to the layer on its right. In the back-propagation step, the layer receives the back-propagated vector from a layer on its right, and computes a linear transform before forwarding it to the layer on its left. The vector received by a layer in the backpropagation step is also used to update the layers’ linear transform in preparation for the next iteration. An iteration consists of a feed-forward step and a back-propagation step—including the linear

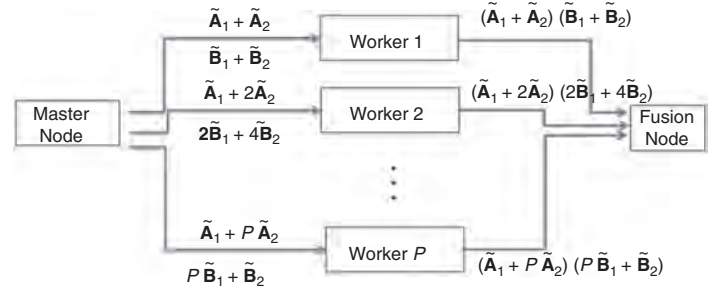


Figure 12. Matdot Codes [34] with  $m = n = 2$  where

$\mathbf{A} = [\tilde{\mathbf{A}}_1 \ \tilde{\mathbf{A}}_2], \mathbf{B} = \begin{bmatrix} \tilde{\mathbf{B}}_1 \\ \tilde{\mathbf{B}}_2 \end{bmatrix}$ . The recovery threshold is 3.

transformation update—executed by all layers, where the first layer receives a data point, and the last layer receives an evaluation of a loss function, as their inputs. In this description, for the sake of simplicity of exposition, we assumed that each iteration is for a single data point.

The most computationally intensive step in the training and testing of a DNN is a matrix-vector multiplication. However, a naive extension of ABFT techniques of coded matrix-vector multiplication (e.g. [47, 56]) would require us to encode matrices at every training iteration. This is because these matrices are updated at every iteration. This overhead of encoding can be enormous, and comparable to the cost of the computation itself, which is undesirable. However, our strategy—that we call *CodeNet*—is able to show that by carefully weaving coding into the computation of DNN, we are only required to code vectors (which is low complexity), and not matrices, at every iteration. This enables encoding and decoding on the fly. Suppose  $\mathbf{W}$  is the weight matrix at any layer of the DNN. Then the most computationally intensive operation in the feed-forward stage of the DNN is the matrix-vector product  $\mathbf{s} = \mathbf{W}\mathbf{x}$  where  $\mathbf{x}$  comes from the previous layer. Similarly in the back-propagation stage, the bottleneck operation is the computation of vector-matrix product  $\mathbf{c}^T = \boldsymbol{\delta}^T \mathbf{W}$  where  $\boldsymbol{\delta}$  comes from the next layer. Finally, in the update step, the  $\mathbf{W}$  is updated as  $\mathbf{W} - \mathbf{W} + \mu \boldsymbol{\delta} \mathbf{x}^T$ .

We divide the matrices and vectors into blocks as

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}_{0,0} & \mathbf{W}_{0,1} \\ \mathbf{W}_{1,0} & \mathbf{W}_{1,1} \end{bmatrix}, \mathbf{x} = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix} \text{ and } \boldsymbol{\delta}^T = [\boldsymbol{\delta}_0^T \ \boldsymbol{\delta}_1^T].$$

In CodeNet, at the beginning of training, the matrix  $\mathbf{W}$  is coded and stored distributedly (by paying a one-time cost) as follows:

$$\mathbf{W}_{\text{coded}} = \begin{bmatrix} \mathbf{W}_{0,0} & \mathbf{W}_{0,1} & \widehat{\mathbf{W}}_{0,2} & \widehat{\mathbf{W}}_{0,3} \\ \mathbf{W}_{1,0} & \mathbf{W}_{1,1} & \widehat{\mathbf{W}}_{1,2} & \widehat{\mathbf{W}}_{1,3} \\ \widehat{\mathbf{W}}_{2,0} & \widehat{\mathbf{W}}_{2,1} & \mathbf{x} & \mathbf{x} \\ \widehat{\mathbf{W}}_{3,0} & \widehat{\mathbf{W}}_{3,1} & \mathbf{x} & \mathbf{x} \end{bmatrix}$$

where each block is stored in a different processing node, and ‘x’ is simply denoting that no block exists at that location.

The straggler/error resilience for feed-forward and back-propagation are obtained through standard ABFT-type MDS coding, the difficulty is in maintaining the coding of the weight matrices even as they are updated because encoding the matrix at each iteration is computationally expensive. We now explain the strategy in a bit more detail, first discussing feed-forward and back-propagation steps, and then discussing how the coding is maintained in the update step:

For the feed-forward stage, one only uses the first two columns as follows:

$$\begin{bmatrix} s_0 \\ s_1 \\ \widehat{s}_2 \\ \widehat{s}_3 \end{bmatrix} = \begin{bmatrix} W_{0,0} & W_{0,1} \\ W_{1,0} & W_{1,1} \\ \widehat{W}_{2,0} & \widehat{W}_{2,1} \\ \widehat{W}_{3,0} & \widehat{W}_{3,1} \end{bmatrix} \mathbf{x} = \begin{bmatrix} W_{0,0}\mathbf{x}_0 + W_{0,1}\mathbf{x}_1 \\ W_{1,0}\mathbf{x}_0 + W_{1,1}\mathbf{x}_1 \\ \widehat{W}_{2,0}\mathbf{x}_0 + \widehat{W}_{2,1}\mathbf{x}_1 \\ \widehat{W}_{3,0}\mathbf{x}_0 + \widehat{W}_{3,1}\mathbf{x}_1 \end{bmatrix}$$

One can use decoding techniques similar to coded matrix-vector multiplication to get back  $s_0$  and  $s_1$  from the 4 coded vectors  $\{s_0, s_1, \widehat{s}_2, \widehat{s}_3\}$  under errors or stragglers. Similarly, for the back-propagation stage, one can only use the first 2 rows of the coded matrix as follows:

$$\begin{aligned} [c_0^T \quad c_1^T \quad \widehat{c}_2^T \quad \widehat{c}_3^T] &= \delta^T \begin{bmatrix} W_{0,0} & W_{0,1} & \widehat{W}_{0,2} & \widehat{W}_{0,3} \\ W_{1,0} & W_{1,1} & \widehat{W}_{1,2} & \widehat{W}_{1,3} \end{bmatrix} \\ &= \begin{bmatrix} \delta_0^T W_{0,0} & \delta_0^T W_{0,1} & \delta_0^T \widehat{W}_{0,2} & \delta_0^T \widehat{W}_{0,3} \\ +\delta_1^T W_{1,0} & +\delta_1^T W_{1,1} & +\delta_1^T \widehat{W}_{1,2} & +\delta_1^T \widehat{W}_{1,3} \end{bmatrix} \end{aligned}$$

Now, from the 4 coded vectors  $\{c_0^T, c_1^T, \widehat{c}_2^T, \widehat{c}_3^T\}$ , one can decode  $c_0^T$  and  $c_1^T$  under errors and stragglers. This ensures that both forward computation and back-propagation are resilient to errors. Importantly, by weaving coding in this manner, every sub-matrix is able to update itself at each iteration *without the need to encode matrices afresh*:

Suppose,  $\widehat{W}_{0,2} = W_{0,0} + W_{0,1}$ . For the update step, if the processing node only has two vectors:  $\delta_0^T$  and coded vector  $\hat{\mathbf{x}}_2 = \mathbf{x}_0 + \mathbf{x}_1$ , then it can update itself as

$$\widehat{W}_{0,2} + \mu \delta_0^T \hat{\mathbf{x}}_2^T = \underbrace{W_{0,0} + \mu \delta_0^T \mathbf{x}_0^T}_{\text{Updated } W_{0,0}} + \underbrace{W_{0,1} + \mu \delta_0^T \mathbf{x}_1^T}_{\text{Updated } W_{0,1}}$$

The recovery threshold of this strategy can be improved by utilizing more sophisticated codes, and will appear in [29].

## Acknowledgements

We would like to thank Kishori Konwar, Haewon Jeong, Sanghamitra Dutta, and Yaoqing Yang for giving insightful comments on this article. We would also like to thank Zhiying Wang, Nancy Lynch, Prakash Narayana Moorthy, Muriel Medard, A. Dimakis, Mohammad Maddah Ali, Salman Avestimehr, Ramy E. Ali, Farzin Haddadpour and Mohammad Fahim for their help during tutorial preparation. We are grateful to the National Science Foundation and the SRC SONIC center for funding our research, which led to this article.

## References

[1] Michael Abd-El-Malek, Gregory R Ganger, Garth R Goodson, Michael K Reiter, and Jay J Wylie. Fault- scalable byzantine fault-tolerant services. In *ACM SIGOPS Operating Systems Review*, volume 39, pages 59–74, 2005.

[2] Marcos Kawazoe Aguilera, Ramaprabhu Janakiraman, and Lihao Xu. Using erasure codes efficiently for storage in a distributed system. In *Proceedings of International Conference on Dependable Systems and Networks (DSN)*, pages 336–345, 2005.

[3] Mehmet Fatih Aktas, Pei Peng, and Emina Soljanin. Effective straggler mitigation: Which clones should attack and when? *arXiv preprint arXiv:1710.00748*, 2017.

[4] Mehmet Fatih Aktas, Pei Peng, and Emina Soljanin. Straggler mitigation by delayed relaunch of tasks. *arXiv preprint arXiv:1710.00414*, 2017.

[5] Ramy E Ali and Viveck R Cadambe. Consistent distributed storage of correlated data updates via multi-version coding. In *Information Theory Workshop (ITW), 2016 IEEE*, pages 176–180, 2016.

[6] Venkat Anantharam, Amin Gohari, Sudeep Kamath, and Chandra Nair. On hypercontractivity and a data processing inequality. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 3022–3026, 2014.

[7] J Chris Anderson, Jan Lehnardt, and Noah Slater. *CouchDB: the definitive guide*. O'Reilly Media, Inc., 2010.

[8] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. *J. ACM*, 42(1):124–142, January 1995.

[9] Hagit Attiya and Jennifer Welch. *Distributed computing: fundamentals, simulations, and advanced topics*, volume 19. John Wiley & Sons, 2004.

[10] N. Azian-Ruhi, S. Avestimehr, F. Lahouti, and B. Hassibi. Consensus-based distributed computing. In *Information Theory and Applications Workshop*, 2017.

[11] Peter Bailis, Shivaram Venkataraman, Michael J Franklin, Joseph M Hellerstein, and Ion Stoica. Probabilistically bounded staleness for practical partial quorums. *Proceedings of the VLDB Endowment*, 5(8):776–787, 2012.

[12] Bharath Balasubramanian and Vijay K. Garg. Fault tolerance in distributed systems using fused state machines. *Distributed Computing*, 27(4):287–311, 2014.

[13] Horace B Barlow. Possible principles underlying the transformations of sensory messages. *Sensory Communication*, pages 217–234, 1961.

[14] Ashwin R Bhambe, Jeffrey Pang, and Srinivasan Seshan. Colyseus: A distributed architecture for online multiplayer games. In *NSDI*, volume 6, pages 12–12, 2006.

[15] Christian Cachin and Stefano Tessaro. Optimal resilience for erasure-coded byzantine distributed storage. In *Dependable Systems and Networks, 2006. DSN 2006. International Conference on*, pages 115–124, 2006.

[16] Viveck R Cadambe, Nancy Lynch, Muriel Medard, and Peter Musial. A coded shared atomic memory algorithm for message passing architectures. In *2014 IEEE 13th International Symposium on Network Computing and Applications (NCA)*, pages 253–260, 2014.

[17] Viveck R Cadambe, Zhiying Wang, and Nancy Lynch. Information-theoretic lower bounds on the storage cost of shared memory emulation. In *Proceedings of the ninth annual ACM symposium on Principles of distributed computing*, PODC '16, pages 305–314. ACM, 2016.

[18] Flavio P Calmon, Yury Polyanskiy, and Yihong Wu. Strong data processing inequalities in power-constrained gaussian channels. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 2558–2562, 2015.

- [19] Josiah L Carlson. *Redis in Action*. Manning Publications Co., 2013.
- [20] Yu Lin Chen, Shuai Mu, Jinyang Li, Cheng Huang, Jin Li, Aaron Ogus, and Douglas Phillips. Giza: Erasure coding objects across global data centers. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, Santa Clara, CA, 2017. USENIX Association. <https://www.usenix.org/conference/atc17/program/presentation/chen-yu-lin>.
- [21] James C Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, Jeffrey John Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, et al. Spanner: Google’s globally distributed database. *ACM Transactions on Computer Systems (TOCS)*, 31(3):8, 2013.
- [22] William Dally. High-performance hardware for machine learning. *NIPS Tutorial*, 2015.
- [23] Jeffrey Dean and Luiz André Barroso. The tail at scale. *Communications of the ACM*, 56(2):74–80, 2013.
- [24] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Voshall, and Werner Vogels. Dynamo: Amazon’s highly available key-value store. In *SOSP*, volume 7, pages 205–220, 2007.
- [25] Dan Dobre, Ghassan Karame, Wenting Li, Matthias Majuntke, Neeraj Suri, and Marko Vukolić PoW-erStore: proofs of writing for efficient and robust storage. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 285–298. ACM, 2013.
- [26] Roland L’vovich Dobrushin and SI Ortyukov. Lower bound for the redundancy of self-correcting arrangements of unreliable functional elements. *Problemy Peredachi Informatsii*, 13(1):82–89, 1977.
- [27] Partha Dutta, Rachid Guerraoui, and Ron R Levy. Optimistic erasure-coded distributed storage. In *Distributed Computing*, pages 182–196. Springer, 2008.
- [28] S. Dutta, V. Cadambe, and P. Grover. “Short-Dot”: Computing Large Linear Transforms Distributedly Using Coded Short Dot Products. *ArXiv e-prints*, April 2017.
- [29] Sanghamitra Dutta, Ziqian Bai, Tze-Meng Low, and Pulkit Grover. Codenet: Large scale training of neural networks in presence of soft-errors and straggling processors. In *Preparation*, 2017.
- [30] Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. Short-dot: Computing large linear transforms distributedly using coded short dot products. In *Advances In Neural Information Processing Systems (NIPS)*, pages 2092–2100, 2016.
- [31] Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. Coded convolution for parallel and distributed computing within a deadline. In *IEEE International Symposium on Information Theory (ISIT)*, 2017.
- [32] Elza Erkip and Thomas M Cover. The efficiency of investment information. *IEEE Transactions on information theory*, 44(3):1026–1040, 1998.
- [33] William S Evans and Leonard J Schulman. Signal propagation and noisy circuits. *IEEE Transactions on Information Theory*, 45(7):2367–2373, 1999.
- [34] Mohammad Fahim, Haewon Jeong, Farzin Haddadpour, Sanghamitra Dutta, Viveck Cadambe, and Pulkit Grover. On the optimal recovery threshold of coded matrix multiplication. In *56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct 2017.
- [35] Rui Fan and Nancy A Lynch. Efficient replication of large data objects. In *DISC*, volume 3, pages 75–91. Springer, 2003.
- [36] Garth R Goodson, Jay J Wylie, Gregory R Ganger, and Michael K Reiter. Efficient byzantine-tolerant erasure-coded storage. In *2004 International Conference on Dependable Systems and Networks*, pages 135–144, 2004.
- [37] Christoforos N Hadjicostis and George C Verghese. Coding approaches to fault tolerance in linear dynamic systems. *IEEE Transactions on Information Theory*, 51(1):210–228, 2005.
- [38] Marina Haikin and Ram Zamir. Analog coding of a source with erasures. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2074–2078, 2016.
- [39] Marina Haikin, Ram Zamir, and Matan Gavish. Random subsets of structured deterministic frames have manova spectra. *Proceedings of the National Academy of Sciences*, 2017.
- [40] Bruce Hajek and Timothy Weller. On the maximum tolerable noise for reliable computation by formulas. *IEEE Transactions on Information theory*, 37(2):388–391, 1991.
- [41] Wael Halbawi, Navid Azizan-Ruhi, Fariborz Salehi, and Babak Hassibi. Improving distributed gradient descent using reed-solomon codes. *arXiv preprint arXiv:1706.05436*, 2017.
- [42] James Hendricks, Gregory R Ganger, and Michael K Reiter. Low-overhead byzantine fault-tolerant storage. *ACM SIGOPS Operating Systems Review*, 41(6):73–86, 2007.
- [43] Thomas Herault and Yves Robert. *Fault-Tolerance Techniques for High Performance Computing*. Springer, 2015.
- [44] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming, Revised Reprint*. Elsevier, 2012.
- [45] Eben Hewitt. *Cassandra: the definitive guide*. “O’Reilly Media, Inc.”, 2010.
- [46] Chu-Hsiang Huang, Yao Li, and Lara Dolecek. Gallager b ldpc decoder with transient and permanent errors. *IEEE Transactions on Communications*, 62(1):15–28, 2014.
- [47] Kuang-Hua Huang and Jacob A. Abraham. Algorithm-based fault tolerance for matrix operations. *IEEE transactions on computers*, 100(6):518–528, 1984.
- [48] Gauri Joshi, Yanpei Liu, and Emina Soljanin. On the delay-storage trade-off in content download from coded distributed storage systems. *IEEE Journal on Selected Areas in Communications*, 32(5):989–997, 2014.

- [49] J-Y Jou and Jacob A. Abraham. Fault-tolerant fft networks. *IEEE Transactions on Computers*, 37(5):548–561, 1988.
- [50] Jing-Yang Jou and Jacob A Abraham. Fault-tolerant matrix arithmetic and signal processing on highly concurrent computing structures. *Proceedings of the IEEE*, 74(5):732–741, 1986.
- [51] Can Karakus, Yifan Sun, and Suhas Diggavi. Encoded distributed optimization. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2890–2894, 2017.
- [52] Kishori M Konwar, N Prakash, Erez Kantor, Nancy Lynch, Muriel Médard, and Alexander A Schwarzmann. Storage-optimized data-atomic algorithms for handling erasures and errors in distributed storage systems. In *Parallel and Distributed Processing Symposium, 2016 IEEE International*, pages 720–729, 2016.
- [53] Kishori M. Konwar, N. Prakash, Nancy A. Lynch, and Muriel Médard. RADON: Repairable Atomic Data Object in Networks. In *20th International Conference on Principles of Distributed Systems (OPODIS 2016)*, volume 70 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:17. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [54] Leslie Lamport. On interprocess communication. *Distributed computing*, 1(2):86–101, 1986.
- [55] K. Lee, C. Suh, and K. Ramchandran. High-dimensional coded matrix multiplication. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 2418–2422, June 2017.
- [56] Kangwook Lee, Maximilian Lam, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran. Speeding Up Distributed Machine Learning Using Codes. *Machine Learning Systems Workshop, Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [57] Kangwook Lee, Ramtin Pedarsani, Dimitris Papailiopoulos, and Kannan Ramchandran. Coded computation for multicore setups. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pages 2413–2417, 2017.
- [58] Songze Li, Mohammad Ali Maddah-Ali, and A Salman Avestimehr. Coded mapreduce. In *Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 964–971, 2015.
- [59] Robert Lucas, J Ang, K Bergman, S Borkar, W Carlson, L Carrington, G Chiu, R Colwell, W Dally, J Dongarra, et al. Top ten exascale research challenges. *DOE ASCAC subcommittee report*, pages 1–86, 2014.
- [60] Nancy A. Lynch. *Distributed, Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.
- [61] Ruben Mayer, Harshit Gupta, Enrique Saurez, and Umakishore Ramachandran. Fogstore: Toward a distributed data store for fog computing. *CoRR*, abs/1709.07558, 2017.
- [62] Edward F Moore and Claude E Shannon. Reliable circuits using less reliable relays. *Journal of the Franklin Institute*, 262(3):191–208, 1956.
- [63] David Mosberger. Memory consistency models. *ACM SIGOPS Operating Systems Review*, 27(1):18–26, 1993.
- [64] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- [65] Michael A Olson, Keith Bostic, and Margo I Seltzer. Berkeley db. In *USENIX Annual Technical Conference, FREENIX Track*, pages 183–191, 1999.
- [66] N. Pippenger. On networks of noisy gates. In *Proc. Symp. Foundations Comput. Sci.*, pages 30–38, 1985.
- [67] N. Pippenger, G.D. Stamoulis, and J.N. Tsitsiklis. On a lower bound for the redundancy of reliable networks with noisy gates. *IEEE Trans. Inf. Theory*, 37(3):639–643, May 1991.
- [68] Yury Polyanskiy and Yihong Wu. Dissipation of information in channels with input constraints. *IEEE Transactions on Information Theory*, 62(1):35–55, 2016.
- [69] Maxim Raginsky. Strong data processing inequalities and log- $\phi$ -sobolev inequalities for discrete channels. *IEEE Transactions on Information Theory*, 62(6):3355–3389, 2016.
- [70] Netanel Raviv, Itzhak Tamo, Rashish Tandon, and Alexandros G Dimakis. Gradient coding from cyclic mds codes and expander graphs. *arXiv preprint arXiv:1707.03858*, 2017.
- [71] Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, and Susan Spence. Fab: building distributed enterprise disk arrays from commodity components. In *ACM SIGARCH Computer Architecture News*, volume 32, pages 48–58. ACM, 2004.
- [72] C.E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27(4):623–656, Oct 1948.
- [73] S Sreenivasan and I Fiete. Error correcting analog codes in the brain: beyond classical population coding for exponentially precise computation. *Nature Neuroscience*, 14:1330–1337, 2011.
- [74] Shreyas Sundaram and Christoforos N Hadjicostis. Fault-tolerant convolution via chinese remainder codes constructed from non-coprime moduli. *IEEE Transactions on Signal Processing*, 56(9):4244–4254, 2008.
- [75] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis. Gradient coding. In *Machine Learning Systems Workshop, Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [76] Rashish Tandon, Qi Lei, Alexandros G. Dimakis, and Nikos Karampatziakis. Gradient coding: Avoiding stragglers in distributed learning. In *International Conference on Machine Learning (ICML)*, pages 3368–3376, 2017.
- [77] M. G. Taylor. Reliable information storage in memories designed from unreliable components. *Bell Syst. Tech. J.*, 47(10):2299–2337, 1968.
- [78] Michael G Taylor. Reliable computation in computing systems designed from unreliable components. *Bell Labs Technical Journal*, 47(10):2339–2366, 1968.

- [79] Luis M Vaquero and Luis Rodero-Merino. Finding your way in the fog: Towards a comprehensive definition of fog computing. *ACM SIGCOMM Computer Communication Review*, 44(5):27–32, 2014.
- [80] Lav R Varshney. Performance of ldpc codes under faulty iterative decoding. *IEEE Transactions on Information Theory*, 57(7):4427–4444, 2011.
- [81] Bane Vasić, Predrag Ivanis, Srdan Brkic, and Vida Ravanmehr. Fault-resilient decoders and memories made of unreliable components. In *Information Theory and Applications Workshop (ITA), 2015*, pages 136–142, 2015.
- [82] Bane Vasić, Predrag Ivaniš, David Declercq, and Khoa Le-Trung. Approaching maximum likelihood performance of LDPC codes by stochastic resonance in noisy iterative decoders. In *Information Theory and Applications Workshop (ITA), 2016*, pages 1–9, 2016.
- [83] Werner Vogels. Eventually consistent. *Queue*, 6(6):14–19, 2008.
- [84] John Von Neumann. Probabilistic logics and the synthesis of reliable organisms from unreliable components. *Automata studies*, 34:43–98, 1956.
- [85] Da Wang, Gauri Joshi, and Gregory Wornell. Efficient Task Replication for Fast Response Times in Parallel Computation. In *ACM SIGMETRICS Performance Evaluation Review*, volume 42, pages 599–600, 2014.
- [86] Da Wang, Gauri Joshi, and Gregory Wornell. Using straggler replication to reduce latency in large-scale parallel computing. In *ACM SIGMETRICS Performance Evaluation Review*, volume 43, pages 7–11, 2015.
- [87] Zhiying Wang and Viveck Cadambe. Multi-version coding in distributed storage. In *2014 IEEE International Symposium on Information Theory (ISIT)*, pages 871–875, 2014.
- [88] Zhiying Wang and Viveck Cadambe. On multi-version coding for distributed storage. In *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 569–575, Sept 2014.
- [89] Zhiying Wang and Viveck R Cadambe. Multi-version coding—an information-theoretic perspective of consistent distributed storage. *IEEE Transactions on Information Theory*, 2017. Accepted to appear, doi: 10.1109/TIT.2017.2725273.
- [90] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Fault-tolerant distributed logistic regression using unreliable components. In *54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 940–947, 2016.
- [91] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Fault-tolerant parallel linear filtering using compressive sensing. In *IEEE International Symposium on Turbo Codes and Iterative Information Processing (ISTC)*, pages 201–205, 2016.
- [92] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Coding method for parallel iterative linear solver. *Advances In Neural Information Processing Systems (NIPS)*, to appear, arXiv:1706.00163, 2017.
- [93] Yaoqing Yang, Pulkit Grover, and Soumya Kar. Computing linear transformations with unreliable components. *IEEE Transactions on Information Theory*, 63(6), 2017.
- [94] SM Sadegh Tabatabaei Yazdi, Hyungmin Cho, and Lara Dolecek. Gallager b decoder on noisy hardware. *IEEE Transactions on Communications*, 61(5):1660–1673, 2013.
- [95] Qian Yu, Mohammad Ali Maddah-Ali, and Salman Avestimehr. Polynomial codes: an optimal design for high-dimensional coded matrix multiplication. *arXiv preprint arXiv:1705.10464*, 2017.
- [96] Heng Zhang, Mingkai Dong, and Haibo Chen. Efficient and available in-memory kv-store with hybrid erasure coding and replication. In *FAST*, pages 167–180, 2016.

---

## From the Editor (continued from page 2)

Announcements, news and events intended for both the printed newsletter and the website, such as award announcements, calls for nominations and upcoming conferences, can be submitted at the IT Society website <http://www.itsoc.org>. Articles and columns can be e-mailed to me at [mikel@buffalo.edu](mailto:mikel@buffalo.edu) with a subject line that includes the words “IT newsletter.”

The next few deadlines are:

Jan 10, 2017 for the issue of March 2018.

April 10, 2017 for the issue of June 2018.

December 2017

Please submit plain text, LaTeX or Word source files; do not worry about fonts or layout as this will be taken care of by IEEE layout specialists. Electronic photos and graphics should be in high resolution and sent as separate files.

I look forward to hearing your suggestions and contributions.

With best wishes,  
Michael Langberg.  
[mikel@buffalo.edu](mailto:mikel@buffalo.edu)

IEEE Information Theory Society Newsletter

## A Walk Over the Memory Lane: One Decade of North American Schools of Information Theory (NASIT)

Aylin Yener

I am writing this looking on the Pacific Ocean at an unusually warm summer day in San Francisco. As the local news outlets search their records to report “record-breaking” temperatures, I start browsing my own files, and emails, our society website with its abundant resources (including archived material for past schools and past issues of the newsletter), looking for “records” of significance for what I have been meaning to write about: the ten North American Schools of Information Theory completed to date, one in each year since 2008.

First a brief history of how these school have come about: A cold December day in 2006 in State College, Pennsylvania (I actually do not recall how cold it was, but likely near freezing if not below), Gerhard Kramer who at the time was with Bell Labs drove over to Penn State for a seminar, and a short visit. During our conversations on various things Information Theory, I found myself lamenting about the scarcity of information theorists on campus and lack of my students’ exposure to information theory other than the introductory course that I would teach. This of course was not a problem unique to Penn State either, and was common among many campuses around the country. The idea came up to organize a meeting that would be useful for students in particular, less formal than an ISIT, a workshop of sorts perhaps. Gerhard of course, was already very familiar with the long tradition of European school of Coding and Information Theory, and had abundance of experience of organizing workshops. An email trail in the next few days (from my records) reveals that it was him who suggested we try an “east coast workshop” at Penn State, which eventually morphed into the idea of the summer school.

The First Annual School of Information Theory (yes, optimistically named) was proposed at the Board of Governors Meeting at Allerton 2007. We proposed the school as a vehicle for students and postdocs to not only be exposed to excellent tutorials by the greats of information theory, but also for the junior members of



Students listening to the lecture at 2008 NASIT

our society to exchange ideas and present their latest results in a fully interactive environment. We would hold this at a university campus and pay for food and would not charge a registration fee. This we conveyed was a way for our society to pay it forward to the next generation of information theorists. The board was extremely supportive and provided us with the financial support to commit to the organization and get started. We assembled the rest from the generous support of local universities that chipped in as well as the federal government (DoD). The hope was that we



A group picture with enthusiastic attendees from 2008 NASIT at Penn State



would attract 30-40 attendees for the school and see if this was an idea worth continuing. What happened in the months followed was a response from the community that went well beyond our imagination, over a hundred students and postdocs signing up for the first school. So we tripled our efforts (and food orders), moved to a large classroom, and waited for the students to descend upon our small town in early June with enthusiasm matched to the colossal autumn football weekends (much more on my side as compared to the said weekends). Gerhard was infinitely helpful at every step, as was the remainder of our organizing committee, Nick Laneman, Ivana Maric, Lalitha Sankar, Brooke Shrader and Sennur Ulukus. The four days of the school were full of enthusiasm, by our lecturers (Toby Berger, Muriel Medard, Vince Poor, David Tse) and all of the attendees, students, postdocs, faculty that drove over from nearby states. Those four days will always remain in my memory as a most rewarding and pleasant experience. I am especially humbled by the amount of brilliance from the young minds we have had at the school, an occurrence that repeated itself every single year since. Upon inspection of the hundreds of pictures that our photographer took (all in my records, but not all of them made it to the web page), I counted over twenty students from that first school who are now exceedingly successful junior and mid-career faculty all over the world!

The second year at Northwestern was equally fantastic with about 40% growth in attendees and an amazing organizing team that included Randy Berry, Natasha Devroye, Dongning Guo and Daniela Tuninetti that left Gerhard and I much less to do! This was also the year Roberto Padovani provided his generous support to start the Padovani lectureship, a prestigious yearly award of our society since. Our first Padovani lecturer was Abbas El Gamal who awed us by teaching (literally) an entire book of Network Information Theory. Lectures by Dan Costello and Bruce Hajek and the keynote address by Bob Gallager were all superb. A memorable moment for me is captured in the picture on top of this page. It was really a sight to see how excited the students were to have the chance to speak with Bob, the circle that surrounded him lasted for quite a while!

By the end of the second NASIT, we felt confident that an annual school was sustainable and there was enough interest from the community, both on the side of the potential attendees, and volunteers to both organize and lecture at the schools. The third annual and the last one I was involved as an organizing co-chair was at USC. Gerhard Kramer was the chief organizer this time at his home campus and once again did an amazing job along with organizers that included Michelle Effros, Alex Dimakis and Tracey Ho. This school had over 200 attendees, the largest one to date. Despite the size, students found plenty of opportunities to present their findings (by the second school we switched to an all poster format for student presentations) and interact with their



**Students listening to Bob Gallager's wisdom filled advice at 2009 NASIT**

peers as well as the lecturers and other faculty in attendance. By the end of third NASIT, we knew this now was a tradition and that the annual schools would go on, and they did, with the generous support of the society, federal agencies such as NSF, centers and research institutes, the Padovani fund, and of course with the volunteers efforts.

Every year since 2008, NASIT ran with enthusiasm with volunteers putting tremendous effort to continue the "tradition" of excellent tutorials, enthusiastic presentations and research idea exchanges, good food and fun. Fast-forwarding to the last two years, the ninth school at Duke University in 2016 again had about 100 attendees, where I had the honor of lecturing. I have seen the same enthusiasm as in the first school then from a whole different generation. The topics of course evolved with time, information theory reaching out now more than ever to sister fields other than communications, but it was the same spirit of interactive discussion was there at every moment same as in the early schools. Same experience repeated the 10<sup>th</sup> Annual NASIT this year at Georgia Tech and was special in many ways. It brought a few of us who were in the first school together: Gerhard, Muriel and I were present to lecture, and most importantly, for the first time we had the organizing chair as an alumnus of NASIT! Yes, Matthieu Bloch was a newly minted



**Group Photo from 2016 NASIT at Duke University**

postdoc who was a student attendee in the first NASIT. So, we have come full circle in one decade, going from student to organizing the school. The paying forward indeed worked after all!

The enthusiasm of NASITs has been infectious and we now as a society support schools all over the world. In addition to an expanded annual European School of Information Theory (ESIT), Australia, India, East Asia hold annual schools, and South Africa had held one a couple of years ago as well. We now have a subcommittee (similar to the conference committee) that collects proposals and provides feedback to the organizers and eventually relays the

proposal to the Board of Governors for support. So, a well-working process is in place for proposing and organizing schools. Our sister societies also now hold summer schools with similar formats. Information theory society has been recognized by the IEEE, for its outstanding outreach programs, including the schools.

We have volunteers signing up to organize schools from the alumni of the school, and we now have a substantial number of school alumni who themselves have students they can encourage to participate in school. The future of the schools continues to inspire me. After all, it represents the future of information theory!

## On “Composing Music by a Stochastic Process”: From Computers that are Human to Composers that are Not Human

*Haizi Yu and Lav R. Varshney<sup>1</sup>*

With the passing of Mary Elizabeth Moore “Betty” Shannon on May 1, there has been renewed interest in her Bell Labs work at the intersection of music and mathematics (mutual loves she shared with her husband, Claude [1]). Here we discuss some context around her work described in the technical memorandum “Composing Music by a Stochastic Process” (1949), together with the intellectual arc of computational creativity for music composition that has carried forward to the present.

Betty Shannon studied mathematics at the New Jersey College for Women (now Douglass College at Rutgers University) on a full scholarship and graduated Phi Beta Kappa. The next day, she started work at Bell Laboratories in Manhattan as a computer in the mathematics department; she was later promoted to technical assistant [1], [2]. For a 200-year period in the history of science and technology, human workers called “computers” were employed to carry out scientific calculations by hand, whether in computing the 1758 return of Halley’s comet, mathematical tables during the Great Depression, or ballistics tables during WWII. Sometime during the war, Warren Weaver started calling computers “girls” and one member of his Applied Mathematics Panel defined the unit “kilogirl,” presumably a term for a thousand hours of computing labor [3]. In this era and beyond, e.g. at NACA, it was a significant and un-



sual achievement for a woman to get her name on a research report; human computers were even largely excluded from editorial meetings where reports were developed [4]. As such, it was exceptional for Bell Telephone Laboratories Technical Memorandum MM-49-150-29, “Composing Music by a Stochastic Process,” 15 November 1949 to list both John R. Pierce and Mary E. Shannon [5].

Though Pierce and Shannon were unaware at the time, their stochastic music composition algorithms followed in the tradition of W. Mozart, J. Haydn, M. Stadler, and K. P. E. Bach [6], but improved upon it considerably. In the memorandum, the authors introduced stochastic models to describe the generating process of chord progression in four-part harmonies, while carefully controlling the model behavior to resemble what a human composer would do in part writing. The rules for both chord construction (1-gram) and chord progression (2-gram) were largely borrowed from known music theory as well as personal specifications such as keeping common tones between adjacent chords. The set of rules defined the legal actions in the composition process as well as their chances to be selected. Therefore, the resulting chord progression was a realization of the stochastic process specified by the rule set.

Under modern taxonomy, the Pierce-Shannon composition model can be thought of as a probabilistic expert system for symbolic music (not audio waveforms). Both the rules and chances involved in this stochastic model are pre-specified rather than learned from data (a Markov random process rather than reinforcement learning),

<sup>1</sup> This work was supported in part by the IBM-Illinois Center for Cognitive Computing Systems Research (C3SR), a research collaboration as part of the IBM Cognitive Horizons Network.

and thus the algorithm runs mechanically as a sampling automaton once it starts; the algorithm is pre-fixed whereas its outputs are random realizations. Known weaknesses exist for this composition model, e.g. the composed music lacks a long-range plan due to being Markov and the quality of the results directly depends on the given rules whose design itself is an art. Indeed, as Pierce later described [6], “While the short-range structure of these compositions was very primitive, an effort was made to give them a plausible and reasonably memorable, longer-range structure.... the compositions were primitive rondos.”

In our eyes, the model was a significant achievement in computational creativity, especially for the 1940s when the algorithm was implemented by “throwing three specially made dice and by using a table of random numbers,” [6] and before J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon asserted creativity as a central goal in the formalization of artificial intelligence [7]. Further progress in Markov chain music was made by L. A. Hiller, Jr. and L. M. Isaacson of the University of Illinois at Urbana-Champaign, who formulated the rules of four-part, first-species counterpoint so that an electronic computer could choose notes at random and reject them if rules were violated. This music was curated and published as the *Illiic Suite for String Quartet* (1957) [8]. Even today, the Pierce-Shannon composition model is used as the core of many modern techniques in music composition and style replication. Perhaps the only change is that modern electronic computers enable extension of the model to higher-order  $n$ -grams that allow a long-range composition plan.

In reflecting on their work, John Pierce suggested [6] “the information-theoretic composer... will make up his composition of larger units [recognizable chords, scales, themes, or ornaments] which are already familiar in some degree to listeners through the training they have received in listening to other compositions. ... Perhaps the composer will surprise the listener a bit from time to time, but he won't try to do this continually.” But where do these concepts, rules, and notions of surprise come from?

Recent research in computer music, including our own [9]–[12], is concerned with automatically learning a hierarchy of human-interpretable composition rules—the laws of music theory directly from raw sheet music—and best methods to coherently break rules to personalize composition style. Interestingly, our approach to interpretable concept learning uses an iterative alternating optimization of information-theoretic functionals such as Bayesian surprise and mutual information, interpreted as cycling between a generative component (student) and a discriminative component (teacher) [9]–[11]. The algorithm not only reproduces much of standard music theory in textbooks, but also yields some novel music theory that our colleagues in the Music department find intriguing (see our project webpage at <https://mus-rover.csl.illinois.edu/>). Now, machines may not only execute composition rules automatically, but learn and break rules automatically too. They may even partner with human composers in creative conversations.

Some have described Claude and Betty Shannon as “one of the great mathematical marriages of our time: one that produced

path-breaking work” [1]; here we have tried to describe one further example of such path-breaking mathematical work in context. Computational creativity and computer music appear to be proceeding along a path of history that runs quite centrally through the work of Betty Shannon.

## Acknowledgment

We thank Charlie Bahr (Bell Labs) for sending us a copy of Technical Memorandum MM-49-150-29.

## References

- [1] J. Soni and R. Goodman, *A Mind At Play: How Claude Shannon Invented the Information Age*, New York: Simon & Schuster, 2017.
- [2] “Mary Elizabeth Moore ‘Betty’ Shannon,” *The Boston Globe*, 14 May 2017.
- [3] D. A. Grier, *When Computers Were Human*, Princeton: Princeton University Press, 2013.
- [4] M. L. Shetterly, *Hidden Figures: The American Dream and the Untold Story of the Black Women Mathematicians Who Helped Win the Space Race*, New York: HarperCollins, 2016.
- [5] J. R. Pierce and M. E. Shannon, “Composing Music by a Stochastic Process,” Bell Telephone Laboratories Technical Memorandum MM-49-150-29, 15 November 1949.
- [6] J. R. Pierce, *An Introduction to Information Theory: Symbols, Signals and Noise*, second revised edition, Dover Books, 1980.
- [7] J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” 31 August 1955.
- [8] L. A. Hiller, Jr., and L. M. Isaacson, *Illiic Suite for String Quarter*, Bryn Mawr, PA: Theodore Presser Company, 1957.
- [9] H. Yu, L. R. Varshney, G. E. Garnett, and R. Kumar, “Learning Interpretable Musical Compositional Rules and Traces,” in *ICML Workshop on Human Interpretability in Machine Learning*, June 2016.
- [10] H. Yu, L. R. Varshney, G. E. Garnett, and R. Kumar, “MUS-ROVER: A Self-Learning System for Musical Compositional Rules,” in *Proc. 4th Int. Workshop Musical Metacreation (MUME)*, June 2016.
- [11] H. Yu and L. R. Varshney, “Towards Deep Interpretability (MUS-ROVER II): Learning Hierarchical Representations of Tonal Music,” in *Proc. 5th Int. Conf. Learning Representations (ICLR)*, Apr. 2017.
- [12] H. Yu, T. Li, and L. R. Varshney, “Probabilistic Rule Realization and Selection,” to appear in *Proc. 30th Annu. Conf. Neural Information Processing Systems (NIPS)*, Dec. 2017.

## The Historian's Column

Most of us would probably agree that we live in turbulent times. The world has gone through very tough periods before and so it is tempting to distance ourselves from “current affairs” and maintain scientific “neutrality”. Storms come and go. True. However, they can leave a lot of destruction and damage in their wake. (Disclosure: I am writing these lines in the aftermath of the triple whammy delivered by Harvey, Irma, and Maria). Prudent individuals can minimize the harm if they plan and react properly. Therefore, I would argue that, even though we are scientists dedicated to Information Theory and its applications, it is important to not only observe but also to react and, why not, influence the path of the storm. After all, as nuclear physicists during and after World War II faced the heavy burden of responsibility and guilt for the potential misuse of their work, we, as contributors to the Information Revolution that has already changed the world, cannot stay silent about how our work may be used or abused. We belong to humankind and to the global society, in addition to our own local and national ones. Hence, it is possible (though not easy) to have a unified and common view regarding such matters.

So, what are the “burning” issues that prompted me to start this column in this solemn fashion? To begin with, we need to maintain a calm and patient attitude and employ our capacity for analytical thinking before we let emotional tendencies take the better of ourselves. There is turmoil in the world; and I do not mean only the political landscapes that have turned ugly almost everywhere. There is the threat and the actual conduct of wars that continue to ravage many parts of the globe. In addition, there is CHANGE in the air that is challenging traditions, attitudes, and cultures worldwide. The social fabric that has defined most of our attitudes and that was more or less prevailing everywhere before the Information revolution seems to be breaking apart. This change has been beneficial to many and yet it seems to be frightening to all. A person may need to change skills several times during their lifetime in the future as obsolescence has accelerated in everything we do. This is a truly frightening specter for many. Even the development of new fantastic technologies (like self-driving cars, extreme reliance on robots and automation, dazzling communication lines crisscrossing the world and our lives) can be scary. Religion, that generally provides a pillar of support to people as they navigate the turbulence in their lives, seems to be losing its grip. Taboos, like gender identity, family structures, etc. are also weakening. News, whether real or fake, and the social networking media spread disturbing rumors and threats. The energy reserves, the availability of healthy nutrition, and even the supply of clean water are all dwindling. All these raise intense feelings of insecurity that can transform sane people to fodder for demagogues and shake the foundations and stability of an increasingly overpopulated world.

Each one of us may share these thoughts to a higher or lesser degree but it is difficult to deny that, unlike disruptive developments of earlier eras, this time we are more aware of the fragility of our world. In our corner of the world of science, we have been seeing signs of this turmoil for some time now. We had the unabashed persecution of colleagues of Jewish origin or of political dissidents behind the Iron Curtain. The ISIT of 1968 that was scheduled to take place in Patras, Greece, had to be hastily rescheduled and moved to San Remo, Italy. Right after our symposia in Istanbul and Barcelona, we have seen serious unrest developing in those cities that might have affected our conferences there, had they

*Anthony Ephremides*



occurred just a few months earlier. The Signal Processing Society had to move its planned Spring 2018 major Symposium from South Korea to Canada, due to the potential threat of dangerous developments in the Korean peninsula. We experience new difficulties in the granting of visas and we see the freedom of movement around the world increasingly curtailed and in danger. So, what to do?

The explosion in the growth of the use of hand-held devices has been staggering. We all accept unquestioningly that this is a good thing; no, we all accept that this is a great thing and we are striving, through our work, to enhance this growth further. We ignore conservative voices that see doom and gloom down the road. Perhaps they are wrong. But, perhaps they are not. When I walk to class, I often see students on a roller board cutting in front of cars while their eyes are glued to the screen of their smartphones. Statistics about the amount of time spent collectively surfing on the social networks are reaching unbelievable levels. Is there time for anything else? Have people stopped thinking? or dreaming? or talking? Is humanity condemned to a life of idleness and nonsense exchanges over the ether? Perhaps President's Trump contribution to the world will be a growing disgust for tweeting.

Of course, there is the Altera pars. Through the social media, it is possible to organize against oppression. Through social media, it is possible to organize conferences in remote parts of the world within a few weeks. Through social media, it is possible to accomplish hundreds of worthwhile things.

The question is should we have a say, or a role, in deciding (or shaping) the way these advances in Information technology are to be used? A parallel question, asked perhaps too late, was whether the nuclear physicists should have had a role in deciding how nuclear power was to be used. Or, whether ballistics specialists should have a role in deciding how advanced weapons should be used. (Disclosure: again, these lines are written in the aftermath of the horrific shooting in Las Vegas). More specifically, shouldn't we revisit the mantra of Information Theory that discards the importance of the content (or the meaning, or the purpose) of a message over a communication link? Isn't it time (from a sociological perspective this time) to broaden the scope of our field? Already there have been investigations of the fundamental limits of Social Networks. Should we extend such investigations to the fundamental limits of the damage or the benefit that can result from the use of social networks?

These questions are disturbing and it is tempting to dismiss them as irrelevant or as daydreaming. However, the study of History is not intended only to record and interpret the past. It is mainly to help predict and shape the future. It is to learn from the past. Just as we do with handling the storms and other catastrophic events that mar our lives.

Ps: Time takes its toll; so, once again we need to update the picture that accompanies this column. The “age” of a picture should be much less than the age of the person it depicts 😊

# BalkanCom'17: A New Conference on Communications and Networking

Stark Draper

On Tuesday 30 May 2017 the first Balkan Conference on Communications (BalkanCom'17) was initiated in Tirana, Albania. The conference brought together lecturers, speakers, and students from a number of areas of communication amongst whom the Information Theory Society was well represented. The objectives of the conference are to establish a forum for high-quality research in the Balkan countries, to expose local students to state-of-the-art topics and to top-notch researchers, and to foster scientific interconnection and research in the region. This was the first time an international conference of communications had been held in Albania, and it set a good precedent for next year's meeting; to be held just up the coast in Montenegro.

The conference was bookended by tutorials. On the first day Prof. Ian Akyildiz of Georgia Tech delivered a tutorial on the ten key enabling technologies in the 5G roadmap. On the final day Prof. Huesyin Arslan of the University of South Florida looked past 5G in his tutorial on future methods of flexible radio access. In between attendees benefitted from three superb plenary talks. Prof. Georgios Giannakis of the University of Minnesota spoke on sketching techniques for big data. Prof. Gerhard Kramer of the Technical University of Munich presented information theoretic limits on fiber-optic communications. Prof. Alexander Proutiere from KTH Sweden discussed how online learning approaches



**Endrit Dosti (University of Oulu and UofT) presents on ultra-reliable communications at BalkanCom'17.**



**Amira Alloum (Nokia), Ido Tal (Technion), Özgür Gürbüz (Sabanci University) Stark Draper (UofT), Georg Böcherer (TUM), tour guide Dorarta, and Ingmar Land (Huawei) enjoy the Albanian seashore in Durrës at BalkanCom'17.**

can be used to solve resources allocation problems in wireless networking.

The technical program of the conference included both invited sessions and submitted talks and posters. Topics spanned networking, signal processing, information theory, and error-correction coding. Complementing the technical discussions was a panel discussion focusing on technological and educational development in the Balkans. The full program of BalkanCom'17 is available at <http://www.balkancom.info/index.html>.

Out of conference hours attendees enjoyed exploring Tirana, the bustling capital of Albania. Great food, friendly locals, and interesting sites and museums documenting the history of the country and region provided plenty of interesting discoveries. For the conference banquet, attendees were treated to a tour of the one of the largest Roman-era amphitheater in the Balkan peninsula located in the coastal Adriatic town of Durrës (Epidamnos in Greek, Durazzo in Italian, Drač in Serbian, or Dyrrachium in Latin). Durrës was the western terminus of the Roman Via Egnatia, which ran from Durrës to Thessalonica and on to Constantinople (present-day Istanbul). A visit to the Archeological Museum, and stroll along the seaside promenade led to a delicious seafood dinner, fueled by plenty of rakia.

Many thanks (Faleminderit!) are due the General Chairs Ian Akyildiz and Arjan Durresi, the TPC chairs Anthony Ephremides, Petar Popovski, and Mehmet Can Vuran, the rest of the organizing committee, student volunteers and participants. This was a great start to a new conference series in an under-represented region of Europe. We look forward to many more years, and hope that ITSoc members will consider submitted to the next edition. BalkanCom'18 will be held in Podgorica, Montenegro, on 4–8 June 2018.

# Call for Nominations

*(ordered by deadline date)*

## Thomas M. Cover Dissertation Award

The IEEE Information Theory Society Thomas M. Cover Dissertation Award, established in 2013, is awarded annually to the author of an outstanding doctoral dissertation.

**NOMINATION PROCEDURE:** Nominations and letters of endorsement must be submitted by **January 15, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/cover-award> for details.

## IEEE Joint ComSoc/ITSoc Paper Award

The Communications Society/Information Theory Society Joint Paper Award recognizes outstanding papers that lie at the intersection of communications and information theory. Any paper appearing in a ComSoc or ITSoc publication during the preceding three calendar years is eligible for the award.

**NOMINATION PROCEDURE:** Nominations and letters of endorsement must be submitted by **February 15, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/honors/comsoc-information-theory-joint-paper-award/comsoc-itsoc-paper-award-nomination-form> for details. Please include a statement outlining the paper's contributions.

## IEEE Information Theory Society Claude E. Shannon Award

The IEEE Information Theory Society Claude E. Shannon Award is given annually to honor consistent and profound contributions to the field of information theory.

**NOMINATION PROCEDURE:** Nominations and letters of endorsement must be submitted by **March 1, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/shannon-award> for details.

## IEEE Information Theory Society Aaron D. Wyner Distinguished Service Award

The IT Society Aaron D. Wyner Service Award honors individuals who have shown outstanding leadership in, and provided long standing exceptional service to, the Information Theory community.

**NOMINATION PROCEDURE:** Nominations and letters of endorsement must be submitted by **March 1, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/wyner-award> for details.

## IEEE Fellow Program

Do you have a colleague who is a senior member of IEEE and is deserving of election to IEEE Fellow status? If so, please submit a nomination on his or her behalf to the IEEE Fellow Committee. The deadline for nominations is **March 1, 2018**.

IEEE Fellow status is granted to a person with an extraordinary record of accomplishments. The honor is conferred by the IEEE Board of Directors, and the total number of Fellow recommendations in any one year is limited to 0.1% of the IEEE voting membership. For further details on the nomination process please consult: <http://www.ieee.org/web/membership/fellows/index.html>

## IEEE Information Theory Society Paper Award

The Information Theory Society Paper Award is given annually for an outstanding publication in the fields of interest to the Society appearing anywhere during the preceding two calendar years. The purpose of this Award is to recognize exceptional publications in the field and to stimulate interest in and encourage contributions to fields of interest of the Society.

**NOMINATION PROCEDURE:** Nominations and letters of endorsement must be submitted by **March 15, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/honors/information-theory-paper-award/itsoc-paper-award-nomination-form> for details. Please include a statement outlining the paper's contributions.

## IEEE Information Theory Society James L. Massey Research & Teaching Award for Young Scholars

The purpose of this award is to recognize outstanding achievement in research and teaching by young scholars in the Information Theory community. The award winner must be 40 years old or younger and a member of the IEEE Information Theory Society on January 1st of the year nominated.

**NOMINATION PROCEDURE:** Nominations and supporting materials must be submitted by **April 30, 2018**. All nominations should be submitted using the online nomination forms. Please see <http://www.itsoc.org/honors/massey-award/nomination-form> for details.

### IEEE Awards

The IEEE Awards program pays tribute to technical professionals whose exceptional achievements and outstanding contributions have made a lasting impact on technology, society and the engineering profession. For information on the Awards program, and for nomination procedures, please refer to <http://www.ieee.org/portal/pages/about/awards/index.html>

# IEEE Information Theory Society Board of Governors Meeting

**Location:** Aachen, Germany

**Date:** 25 June 2017

**Time:** The meeting convened at 12:50 pm (GMT+2); the meeting adjourned at 5:11 pm.

**Meeting Chair:** Rüdiger Urbanke

**Minutes taken by:** Stark Draper

**Meeting Attendees:** Jeff Andrews\*, Alexander Barg, Andrew Barron, Matthieu Bloch, Helmut Bölcskei, Suhas Diggavi, Alex Dimakis, Stark Draper, Michelle Effros, Elza Erkip, Christina Fragouli, Stephen Hanly, Tara Javidi, Michael Langberg, Matt LaFleur#, Ubli Mitra, Pierre Moulin, Prakash Narayan, Krishna Narayanan, Alon Orlitsky, Vincent Poor, Anand Sarwate#, Emina Soljanin, Daniela Tuninetti, Rüdiger Urbanke, Emanuele Viterbo, Michelle Wigger, Aaron Wagner#, Greg Wornell, Aylin Yener, Wei Yu (Remote attendees denoted by \*, non-voting attendees by #).

**Business conducted between meetings:** Between the Feb. 2017 and Jun. 2017 Information Theory Society (ITSoc) Board of Governors (BoG) meetings, a number of items of business were conducted and voted upon by email. These included an election and the passing of a number of motions. These items and results are summarized below:

1) Gerhard Kramer was elected to the Nominations and Appointments Committee.

2) **Motion:** "The BoG is requested to consider a motion to approve the following appointment to the Editorial Board of the IEEE Transactions on Information Theory: Max Costa, Shannon Theory." The motion was put forward by Transactions Editor-in-Chief Prakash Narayan and was passed by the BoG.

3) **Motion:** "Nominate Ayer Ozgur to serve as co-chair of Student Outreach Subcommittee from Jan. 2017–Dec. 2018." The motion was approved.

4) **Motion:** "To approve a new simplified procedure to handle approvals of technical co-sponsorships (TCS) for conference." In summary: For conferences that have been approved by the BoG for two consecutive years, the Conference Committee can independently authorize IT Society Technical Co-Sponsorship (TCS), provided that the Committee reaches consensus and that the conference has not undergone major changes. In all cases, TCS decisions revert to the BoG for review every 6 years. The motion passed.

5) **Motion:** "To approve an allocation of \$45k USD to make three additional educational videos." This would be a continuation of an initiative to develop videos to explain network coding and space-time codes, with those first two videos nearing completion. The motion passed.

6) **Motion:** "Approve the necessary funds to buy around 800 copies of the new Shannon biography *A mind at play: How*

*Claude Shannon invented the Information Age* (this is equal to the number of ISIT participants) to be distributed as gifts at ISIT in Aachen." The motion passed.

- 7) **Motion:** "Approve the ITSoc BoG minutes from the Feb. 2017 meeting." The motion passed.
- 8) **Motion:** "To support the 2017 Indian School with the amount requested (\$12k USD)." The motion passed.

At 12:50pm local time, ITSoc president Rüdiger Urbanke called the meeting to order. He started by reviewing the agenda.

**Motion:** A motion was made to approve the agenda. The motion was passed unanimously.

- 1) **President's Report:** Rüdiger presented the President's report. Rüdiger started by proposing a new approach to running ITSoc BoG meetings going forward, one that is being prototyped in this meeting. In order to free up time for discussion of forward-looking topics, which at this meeting would center on the proposals for a new journal and a new magazine, many issues were dealt with by email prior to the meeting. Thus, Rüdiger will try to move through the standard reports quickly, with an aim of finishing all within two hours. He requested the BoG to provide feedback on the success of the approach and for further suggestions. Rüdiger then discussed the new approach to distributing the T-IT table of contents via email with one objective being to increase click rate. He updated the BoG on the Shannon movie. While some graphics and historic shooting at Princeton remain to be done, the target is to have the movie completed by the end of 2017.
- 2) **Treasurer's Report:** Treasurer Daniela Tuninetti presented the treasurer's report. She focused on the 2018 budget, shortly due to the IEEE. The IEEE Board of Directors approved a target goal of 2.5% operational surplus (excluding new initiatives) for IEEE overall. As noted in previous meetings, ITSoc's challenge is that there has been a continuing decline in revenue from Transactions in the past few years (which the IEEE predicts will continue) and, in the same period, the BoG has aimed to reduce income from conferences. Therefore unless we can, e.g., increase the click-rate to increase the income coming from the Transactions, there will be a need to increase revenue from conferences. She presented two proposals to balance the budget in the near term: (i) reduce subsidies to schools and (ii) increase the conference surplus targets. In 2018 year we can further increase revenue by eliminating the subsidy provided to those that receive the printed version of the Transaction. It is not clear whether this subsidy can be eliminated even more quickly, but Daniela's proposal is to get there as soon as possible. A question was raised about the level of editing of the Transactions, which a couple years back was reduced to "moderate". There is no lower-level of editing, so no further savings are possible there.

**Motion:** "Change the pricing of the paper copy of the Transition to be cost neutral as soon as possible." Daniela made the motion. Aylin Yener seconded the motion. The motion passed.

Daniela next discussed new initiatives. She suggested ITSoc continue into 2018 some of the new initiatives that started in 2017 as a continuation of the outreach activities for the 2016 Shannon Centennial in 2016. These could continue to be supported under the "3% rule" is needed. These might include (i) further educational videos, (ii) continued development of the children's book, and (iii) an additional subsidy to purchase copies of the Shannon biography that attendees will receive at ISIT. There was a discussion by the BoG on how schools might be reconfigured or augmented to be classified as new initiatives. For example, schools could include undergraduates and/or K-12 outreach, new aspects that would also provide the graduate students attending the opportunity to contribute to outreach.

**Motion:** "Move for an additional \$45k USD to be spent on video clips, \$15k USD more for the speaker series, \$20k USD for the children's book on information theory, \$20k USD for K-12 and undergraduate outreach, all under the for 3% rule, pending IEEE approval." The motion passed unanimously.

- 3) **Online Committee:** Online Committee Chair Anand Sarwate quickly reviewed the current outreach activities by the Committee, received feedback from the BoG on how the various activities (the Shannon documentary, YouTube tutorials, the children's book) are promoted, and discussed how the activities will be archived. He summarized current outreach via email lists, including distributing the table-of-contents of each month's Transactions.
- 4) **Pilot Video Project:** On behalf of Matthieu Bloch, Michelle Effros gave an update on the pilot video project. The two pilot videos are almost complete. One is on network coding, the other on space-time coding. Both videos were developed in partnership with Brit Cruise, an experienced developer of online educational videos. The target audience was high school students. There is a call out to the ITSoc community for proposals to develop additional introductory videos. There is budgeting for three more videos. Proposals should be sent to Matthieu Bloch.
- 5) **Fellows Committee:** Fellows Committee Chair Helmut Bölcskei thanked the members of the Committee. There were 9 nominations this year. The outcome will be announced in November.
- 6) **Massey Award Committee:** Massey Committee Chair Helmut Bölcskei thanked the members of the Committee. The Committee considered all 8 nominations that were received since the inception of the award in 2015. Aaron Wagner was selected. The field was quite strong. Helmut noted that there was some confusion regarding the nomination deadline, the age cutoff, etc. Necessary clarifications will be made and posted.
- 7) **Membership Committee:** Membership Committee Chair Emina Soljanin told the BoG that this year's Padovani Lecturer, Prof. Amin Shokrollahi, was not able to get a visa to visit the US to deliver his lecture in person at the North American School of Information Theory (NASIT). Instead he had to deliver his lecture remotely. Regarding chapter



activities, there are two new chapters, one in Italy that has been formed, and a second in Switzerland that has been approved. There is a need for more requests for Distinguished Lecturer (DL) tours. Such requests must come from chapters and there were only two requests this year. In an effort to increase the number of tours for DLs, the DLs have been invited to the Chapter's Lunch at ISIT 2017. Emina reviewed activities of the Outreach Subcommittee at ITA2017 (including the panel discussions "You and your research") and (then upcoming) ISIT2017 activities including roundtables on a dozen important topics. The Outreach Subcommittee has also been looking into the effectiveness of the Mentoring Program, with a report forthcoming. The activities of the Student Subcommittee were also reviewed, including outreach efforts and meet-the-Shannon-Lecturer lunch. Finally, Emina discussed the reorganization and consolidation of the Membership Committee that was initiated at the October 2016 BoG meeting.

- 8) **Nominations and Appointments Committees:** Nominations and Appointments (N&A) Committee Chair Michelle Effros first discussed the protocol the Committee follows that leads to the creation of the list of BoG nominees. At least 12 candidates from the N&A Committee (for which no voting is needed), individual BoG members can also nominate (these require the support of at least five BoG members), and non-BoG ITSoc members may nominate. Each year six new BoG members are elected in the fall along with one from each under-represented region. Michelle presented the list of 13 BoG candidates for 2018–2020. Aspects the Committee considers in the final determination of the list include: BoG members completing their first term, geographical representation, seasoned wisdom vs. new energy, demonstrated Society leadership, and diversity (broadly interpreted). She asked if there were any nominations at the meeting, of which there were none.
- 9) **Constitution and Bylaws Committee:** Michelle Effros chairs this Committee as part of her duties as Second Past President. Michelle described to the BoG how there are a number of confusing sections in the Constitution and Bylaws, and a number of internal inconsistencies. This year Michelle is reviewing the documents to clean up these aspects. The BoG will receive updated text for the Bylaws prior to the October 2016 meeting.
- 10) **Nominations to Presidential chain:** Rüdiger next requested nominations to the presidential chain of offices. He first asked for nominations for the second vice president: (i) Emanuele Viterbo was nominated by Elza Erkip and seconded by Prakash Narayan; (ii) Helmut Bölcskei was nominated by Emina Soljanin and seconded by Pierre Moulin. Rüdiger next asked for nominations for the first vice president: Emina Soljanin was nominated by Alon Orlitsky and seconded by Christina Fragouli. Rüdiger then requested for nominations for president: Michelle Effros was nominated by Elza Erkip and seconded by Suhas Diggavi. The election will be conducted by email.
- 11) **Awards Committee:** Awards Committee Chair Elza Erkip discussed the Best Paper award. A report was distributed to the BoG a couple weeks prior to the meeting. She talked

through the nominations for the best paper award and the process followed by the Committee. There was a discussion amongst the BoG of how better to identify candidate papers and to foster nominations.

**Motion:** "To accept the report of the Awards Committee." This was moved by Rüdiger and seconded by Alexander Barg. The motion passed unanimously.

**Motion:** "To award the paper per the recommendation of the Awards Committee." This was moved by Vincent Poor and seconded by Ubli Mitra. The motion passed unanimously.

- 12) **Schools Subcommittee:** Schools Subcommittee Chair Aylin Yener presented the proposal for the North American School of Information Theory (NASIT'19) to be held at Texas A&M University.

**Motion:** "To support 2019 NASIT at Texas A&M in the amount of \$20000 USD." The motion was approved.

- 13) **Publications Committee:** Publications Committee Chair and Transaction Editor-in-Chief (EiC) Prakash Narayan began by reminding the BoG of the new system of management of the Transactions. This includes an EiC and an Executive Editor (EE) both serving contemporaneous 18-month terms. Alexander Barg is the current EE. Prakash next turned to the PRAC (Periodicals Review and Advisory Committee) report. This is an IEEE Technical Activities Board (TAB) review that every IEEE publication goes through each five years. The objective is to ensure timeliness and quality, compliance with IEEE policies and procedures, and financial help as well as to provide suggestions for improvements and thoughts on best practices. Main issues discussed in the report include submission-to-publication (sub-to-pub) times, reduction in click rates, and a technical concern about fast-reject. Prakash review for the BoG some statistics about the Transactions: roughly 1100 papers have been submitted each of the past few years, the average sub-to-pub time has been reduced from over 20 months in 2013 to 15.3 now, the average annual page count is about 7500 pages across roughly 500 articles. Prakash was asked about the impact factor of the Transactions and the effect of papers of particular significance. Prakash told the BoG that the impact factor can vary quite unpredictably, e.g., the impact factor was 1.7 the week prior to the ISIT then jumped to 2.5 the week of the BoG meeting. His strategy will be to stick to the basics and then to diversity and improvise judiciously.

Prakash told the BoG he had formed an ad-hoc group for self examination consisting of the current and past two EiCs, the EE, and the current Executive Editorial Board. The goals of the Ad-Hoc Group were to figure out how to make the Transactions appeal to a broader community, reduce sub-to-pub times, and how to improve impact factor and click rates on IEEEExplore. The Group backed the following recent steps: (i) increase the number of associate editors (AEs) to 60, and (ii) try to accelerate inevitable decisions. Regarding the latter, Prakash discussed the process for fast rejection of out-of-scope or technically deficient submissions. For out-of-scope, fast rejection is handled by the EE. To reject based

on technical deficiency the assigned AE discusses the article with the EiC and the EE; all three must concur on a decision to reject.

Prakash then discussed some new initiatives backed by the Ad-Hoc Group intended to enhance the appeal of the Transactions, and also to increase the impact factor and clicks on IEEEExplore. These can be categorized into “Quickies” that are already in motion and “Slowies” which will be initiated beginning in 2018. The quickies include distribution of each month’s Transaction’s table of contents via email, efforts by the authors and AEs to ensure that, in their list of references, all articles published in the Transactions cite published versions of background material rather than ArXiv versions. The slowies include publishing invited (and reviewed) cross-cutting articles in the Transactions that draw on ideas from information theory and emerging developments in complementary fields (including, e.g., bioinformatics, machine learning cyber physical systems, theoretical computer science, physics). For these contributions the EiC will serve as AE. The intent is to have 3–4 articles per year where the co-authors have at least one “collective foot” planted firmly in ITSoc. To date three groups of authors have been invited.

Prakash then pointed the BoG (and the general ITSoc community) to information for reviewers posted by the former EiC Frank Kschischang (<http://www.comm.utoronto.ca/trans-it/reviewer-info.shtml>). To summarize one point, regarding whether to accept a request to review an article, Frank quotes the mathematician Paul Halmos: “Papers to be refereed don’t improve by sitting for a few months at the bottom of the pile on your desk, and you don’t save time and energy by postponing the day when you refuse the job.” Thus, Frank suggests that when “deciding whether or not to accept a review request, a potential reviewer should take on the review only if it is something they have time for immediately.” He again quotes Halmos: “Zero or infinity. Do the job now or do it never, and, in either case, say which, now.”

In conclusion, Prakash reviewed AE retirements and presented the slate for new editors.

**Motion:** “The BoG is requested to consider a motion to approve the EiC’s list of appointments to the Editorial Board of the IEEE Transactions on Information Theory.” Seconded by Vince Poor. The slate was approved unanimously.

**14) New publications (background):** Starting at the 2016 ISIT BoG meeting, the possibility of starting new ITSoc publications has been discussed. The two possibilities considered were a magazine and a journal of special topics. At the October 2016 BoG meeting in Chicago a consensus formed to move forward by organizing an ad-hoc committee on the topic. The Committee is being co-chaired by Elza Erkip and Jeff Andrews.

**15) Proposal for a magazine:** Ad-hoc Committee Co-Chair Elza Erkip first reviewed the rationale for a Magazine. She started by putting the ITSoc Newsletter in some context. The Newsletter is only read by ITSoc members, it is not covered by search engines, articles that appear therein are not cited (and

therefore do not lead to clicks on IEEEExplore). In contrast, an ITSoc Magazine would provide a venue for tutorial-style papers (covering both core information theory and new directions), would increase the visibility of the field of information theory, would increase clicks on IEEEExplore, and generally would increase the visibility of ITSoc. She reviewed the impact factors of various IEEE publications. Elza noted that the numbers were as of the morning of the BoG meeting, recognizing that, as already noted by the EiC, impact factors can fluctuate “quickly.” Regarding traditional journals the impact factor of the Transactions is 2.7, that of the Transactions on Communications is 4, that of the Transactions on Wireless is 2.9, and that of the Transactions on Signal Processing is 4.3. Regarding the two IEEE special issue journals of related societies (Communications and Signal Processing), the Journal on Special Areas of Communications has an impact factor of 8.1 and the Journal of Special Topics in Signal Processing has an impact factor of 5.3. Turning to magazines, the impact factors are uniformly higher: the Communications Magazine has an impact factor of 10.4, the Wireless Magazine has an impact factor of 8.9, and the Signal Processing Magazine has an impact factor of 9.6. The deduction Elza made is that, judged solely on impact factor, a special-issue journal and/or a magazine could make sense.

Elza next focused in on the idea of a magazine, reviewing possible content. Possible content could include technical articles (aimed at a wide audience, discussing “hot” topics, providing historical perspective, and could be a mix of invited and contributed articles—reviewed in all cases), columns (such as those in the newsletter, including from the ITSoc President, from the EiC, the student column, puzzles, and also new columns, e.g., educational tips, book reviews, industrial perspective). Turning to the mechanics of technical articles, the current thinking is that certain issue would be focused on specific topics, though not necessarily all. (The Signal Processing Magazine follows such a mixed format.) For focus issues a team of guest editors would take lead, teams that are not necessarily core-ITSoc would be welcome (or invited), and there would be a combination of invited teams and those selected from an open call. In addition to the focus topics, the Shannon lecturer and ISIT/ITW plenary speakers would continue to be incited to contribute articles (as in the Newsletter) and there would be space for other invited or contributed articles. Regarding the logistics of publication, the magazine would be published quarterly (four per year), there would be (soft) page limits, and likely there would be a print as well as an electronic version. It appears that the general IEEE readership is more keen on having physical copies of magazines than Transactions or Journals. The editorial board would consist of an EiC and an Executive Editorial Board (EEB) of 2–3 members. There would be a team of senior editors appointed by the EiC in consultation with the EEB that would, e.g., select/invite guest editors and oversee the guest editor review process as well as the non-focus issues. Column editors would be appointed in the same manner.

A number of questions were raised by the Ad-Hoc Committee. The first of these is accessibility. Would the magazine be available on IEEEExplore, would the table-of-contents be pushed by email to ITSoc members (as has recently been started for the Transactions), would we link to the magazine from the ITSoc Website, could we make the current issue available

to all on the ITSoc website? (Regarding the last point, the Signal Processing Society does this.) A second question raised would be the level of effort involved. A magazine is a larger effort than the Newsletter and would require a more active editorial board. For it to be successful a magazine would require significant participation from the ITSoc community and from the IEEE. Third is the question of financial impact. There are significant costs associated with production (and design, printing, delivery), but also revenue (from subscriptions, advertising, and via IEEEExplore). Projections related to magazine proposals from sister societies indicate (the expectation of) profitability in about 3 years.

Finally, Elza reviewed the process for establishing a new IEEE journal or magazine (a letter of intent, a phase I proposal, and a phase II proposal), and the recommendation of the Ad-Hoc Committee. The Committee's recommendation is that the BoG proceed with a letter of intent. Elza indicated to the BoG that she had prepared a pair of motions to support the submission of a letter of intent, and to support costs incurred by those involved in the proposal (e.g., travel costs to IEEE). In terms of timeline, if approved and no significant delays, the magazine could be launched in late 2018 or early 2019.

**Motion:** "ITSoc should submit a letter of intent for a magazine to IEEE."

**Motion:** "Authorize reimbursement of reasonable travel expenses for in-person presentations at the necessary TAB meetings."

The two motions were moved by Andrew Barron and seconded by Ubli Mitra. There was discussion and voting was delayed until after the discussion of the proposal on a journal on special topics. When the vote was conducted the motions passed: 21 votes were in favor and 1 was opposed.

Following Elza's presentation, there was an extended discussion amongst the BoG. Topics discussed included advertising, naming (to increase the scope of the magazine), and how to leverage the existing efforts and infrastructure of the Newsletter to provide a firm foundation for a new magazine. One worry raised was the potential of siphoning off of AE and community resources from the Transactions (in terms of both editorial effort and technical contributions). One response to this concern was that many ITSoc community members actively contributed to other societies' magazines and that this type of publication attracts a different sort of technical contribution when compared to publications in Transactions. Further, a magazine would have much broader visibility than does the Newsletter. There was a sense that articles need to be attracted to the magazine from outside the core information theory community if the magazine is to be successful.

16) **Proposal for a new journal on selected topics:** On behalf of Ad-hoc Committee Co-Chair Jeff Andrews (who was in attendance via telecom and created the presentation) Rüdiger Urbanke presented the Ad-Hoc Committee's report on a new special topics journal. Rüdiger started by noting that the information theory community has been very successful, so

successful in fact that communications has become a commodity. He observed that NIPS (the Conference on Neural Information Processing Systems) was started by information theorist Ed Posner so ITSoc could have evolved to captured more of the work being done in important areas such as machine learning, as well as other related areas such as quantum information theory and cryptography. Thus, since communications is now somewhat commoditized, this is an important moment to look for opportunities to bring some of the developments in these related areas back to the information theory community. As one example, ITSoc members are already contributing heavily to the Journal of Special Areas in Communications (JSAC) and to Foundations and Trends (FnT) journals. Contributions to those publications don't bring direct benefits to the ITSoc community. The overriding hope is that a new special topics journal would bring some of those contributions back to our community, would help establish more and stronger connections to the machine learning and other related communities, and would raise the profile of the Transactions as a suitable venue for papers from those communities.

Rüdiger reviewed the goals of the special topics journal. On the technical side objectives are to broaden the use and appreciation of information theory in related fields, to expose the information theory community to new ideas and research directions, and to create venues for focused discussion of special topics. On the logistical front, a special topics journal would be a broad and welcoming venue for topics related to (but not core to) information theory, the journal would have a low (and low variance) 9–12 month sub-to-pub time, and the new journal could relieve some pressure on the Transactions. Financially, such a journal could provide a new and significant long term stream of revenue.

Some logistics were also reviewed. First, Rüdiger reviewed the approval process, which is the same as for the magazine discussed above. In contrast to the magazine, the special topics journal is envisioned as an electronic-only journal. It would publish about 4 to 6 special issues a year, and there would most likely be page limits. Rüdiger reviewed preliminary financial projections as of year four. He discussed the proposed governance structure: an EiC and steering committee plus a set of 8-10 senior editors who would supervise the special issues. A few ideas on possible initial topics were presented.

**Motion:** "ITSoc should submit a letter of intent for a special topics journal to IEEE," and

**Motion:** "Authorize reimbursement of reasonable travel expenses for in-person presentations at the necessary TAB meetings."

The two motions were moved by Vincent Poor and seconded by Ubli Mitra. Following discussion, the motions passed with 17 votes in favor, 7 abstentions, and 1 opposed.

Following Rüdiger's presentation an extended discussion took place. A initial question raised concerned technical papers submitted to the new journal that in current circumstances would rather have been submitted to the Transactions, and the impact on the latter. The Committee's hope is that, while

certainly some papers would appear in the new journal that would have been submitted to the Transactions, many more papers would eventually be attracted to the new journal that would never have been submitted to the Transactions. The special topics issues would provide the kernel for additional follow-on submissions to the Transactions from outside the information theory community, and for initial involvement of non-ITSoc members with the ITSoc community, through the mechanism of guest editorship of the new journal.

Building on this last point followed a discussion of creating a new journal that could serve as a home to, e.g., NIPS or ICML (the International Conference on Machine Learning) papers that currently don't have a great venue for journal publication. In particular, our society might provide a good venue for a subset of learning theory papers that match our society's culture, perspective and tools. Further, it was felt that, in addition to attracting contribution in such research areas, this initiative will only be really successful if contributions are two-way; if we members of ITSoc can contribute aspects of our tools and culture to the problems and culture of those areas. A question was then raised at how dire the situation really is, whether these proposals on new publications are too conservative, since they look a lot like what other societies are doing. In the past a proposal for a journal on "Information" had been floated, perhaps something like that would be a bold step.

Other points of discussion raised in the BoG discussion included the following. First was a concern whether the consistently high quality of the Transactions could be maintained in a new journal. Second was discussion of the formal relationship between the Transactions (especially special issues of the Transactions) and the new journal, and how article submitted to one that do not quite fit or can't be accommodated could be routed to the other. The fact that the editorial board will be separate might pose a bit of a challenge here. A third point was an opportunity provided by special topics journals, namely the opportunity to write review articles (of the focus area of the issue), a type of article not common in the Transactions. A final point of discussion

was whether staffing the editorial board of a new journal would increase the already existing challenge in attracting AEs to the Transactions.

Rüdiger summarized the discussion by stating that while the BoG appears generally to be in support of both measures, perhaps we don't know yet if there is sufficient person power to support the simultaneous establishment of two new journals. The tactical question thereby raised is whether the BoG should move forward with two letters of intent at this point or only one. In order to determine which new publication had more support, the thought was posited that an editorial board (or boards) should first be organized. While that is logical, on the other hand, it will likely be easier to identify people willing to commit to editorial roles once the BoG makes a firm commitment. At this point the BoG returned to the two sets of motions presented above, passing both.

17) **ArXiv, new ee.IT listing:** Madhu Sudan presented a discussion of how the information theory section of ArXiv can be searched. Currently information theory postings are listed on both "cs.IT" and "math.IT". That said, a new "ee" (electrical engineering) section is shortly to be added to the ArXiv, making an "ee.IT" listing possible. Madhu first asked the BoG whether the BoG would like information theory posting to be listed in ee.IT. Next, assuming the response to an ee.IT listing was positive, how did the BoG want information theory posting to be cross-listed ("overlays" in ArXiv terms) if it won't be allowed to be listed in all three categories. Madhu told the BoG that currently it is rare to have even two overlays (as IT currently has), so there may be pushback from the committee that oversees the ArXiv to add a third. The BoG was polled and all members wanted ee.IT to be included. The BoG further expressed a generally strong preference (if having three overlays is not allowed) to keep cs.IT over math.IT if only one overlay is allowed.

18) **Adjournment:** The meeting adjourned at 5:11 local time.

---

## From the President *(continued from page 1)*

At the BoG meeting in Aachen we decided to send Letters of Intent regarding both projects to IEEE. This has been done. The next step is to submit so-called Phase I proposals. Both of these projects are ambitious and getting them realized will take several years. But they are also potentially transformative.

Elza Erkip has championed the Magazine proposal and Jeff Andrews has been driving the Special Topics Journal. I would like to thank them both.

This is my last newsletter column. One year can seem long but it is short when it comes to planning and executing large projects. Our managing structure of five officers who shift yearly

has many advantages. It spreads out the workload and you always have people to discuss with and to ask for advice. But it also has its downsides. Since from our side essentially only the president participates at the Technical Advisory Board meetings of IEEE, it leaves no time to forge connections and to take advantage of IEEE's many internal projects and resources. Many Societies have two year terms. This might be a good topic for future discussions.

It was fun to be king, it is better to have Elza! Thanks for letting me write all these columns.

*Ruedi Out*

# Recent Publications

## IEEE Transactions on Information Theory

### Table of content for volumes 63(9), 63(10), 63(11)

Vol. 63(9): Sep. 2017.

SIGNAL PROCESSING, SPARSE RECOVERY, LEARNING, ESTIMATION		
<i>D. Stotz, E. Riegler, E. Agustsson, and H. Bölcskei</i>	Almost Lossless Analog Signal Separation and Probabilistic Uncertainty Relations	5445
<i>A. Kolleck and J. Vybíral</i>	Non-Asymptotic Analysis of $\ell_1$ -Norm Support Vector Machines	5461
<i>L. Jacques</i>	Small Width, Low Distortions: Quantized Random Embeddings of Low-Complexity Sets	5477
<i>U. J. Mönich and H. Boche</i>	A Two Channel System Approximation for Bandlimited Functions	5496
<i>M. Behr and A. Munk</i>	Identifiability for Blind Source Separation of Multiple Finite Alphabet Linear Mixtures	5506
<i>S. Tang</i>	Universal Spatiotemporal Sampling Sets for Discrete Spatially Invariant Evolution Processes	5518
<i>P. Pad, K. Alishahi, and M. Unser</i>	Optimized Wavelet Denoising for Self-Similar $\alpha$ -Stable Processes	5529
<i>S. Ubaru, A. Mazumdar, and Y. Saad</i>	Low Rank Approximation and Decomposition of Large Matrices Using Error Correcting Codes	5544
<i>S. Park, N. Y. Yu, and H.-N. Lee</i>	An Information-Theoretic Study for Joint Sparsity Pattern Recovery With Different Sensing Matrices	5559
<i>G. Gül and A. M. Zoubir</i>	Minimax Robust Hypothesis Testing	5572
<i>V. Konev and S. Vorobeychikov</i>	Quickest Detection of Parameter Changes in Stochastic Regression: Nonparametric CUSUM	5588
QUANTUM INFORMATION THEORY		
<i>M. Dalai and A. Winter</i>	Constant Compositions in the Sphere Packing Bound for Classical-Quantum Channels	5603
<i>K. Feng, L. Jin, C. Xing, and C. Yuan</i>	Multipartite Entangled States, Symmetric Matrices, and Error-Correcting Codes	5618
<i>J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu</i>	Sample-Optimal Tomography of Quantum States	5628
<i>H.-A. Loeliger and P. O. Vontobel</i>	Factor Graphs for Quantum Probabilities	5642
<i>K. Horodecki</i>	On Distinguishing of Non-Signaling Boxes via Completely Locality Preserving Operations	5666
CODING THEORY AND TECHNIQUES		
<i>V. Guruswami and M. Wooters</i>	Repairing Reed-Solomon Codes	5684
<i>S. Li, C. Li, C. Ding, and H. Liu</i>	Two Families of LCD BCH Codes	5699
<i>Y.-C. Huang and K. R. Narayanan</i>	Construction $\pi_A$ and $\pi_D$ Lattices: Construction, Goodness, and Decoding Algorithms	5718
<i>S. Kudekar, T. Richardson, and A. R. Iyengar</i>	Analysis of Saturated Belief Propagation Decoding of Low-Density Parity-Check Codes	5734
<i>Y.-Y. Jian, H. D. Pfister, and K. R. Narayanan</i>	Approaching Capacity at High Rates With Iterative Hard-Decision Decoding	5752
<i>Y. Yehezkeally and M. Schwartz</i>	Limited-Magnitude Error-Correcting Gray Codes for Rank Modulation	5774
<i>E. Nieminen</i>	On Quadratic Permutation Polynomials, Turbo Codes, and Butterfly Networks	5793
<i>K. V. Rashmi, N. B. Shah, and K. Ramchandran</i>	A Piggybacking Design Framework for Read-and Download-Efficient Distributed Storage Codes	5802
<i>Q. Yan, M. Cheng, X. Tang, and Q. Chen</i>	On the Placement Delivery Array Design for Centralized Coded Caching Scheme	5821
<i>L. Wang, S. Hu, and O. Shayevitz</i>	Quickest Sequence Phase Detection	5834
SHANNON THEORY		
<i>W. Mao and B. Hassibi</i>	Capacity Analysis of Discrete Energy Harvesting Channels	5850
<i>W. Huleihel and Y. Steinberg</i>	Channels With Cooperation Links That May Be Absent	5886
<i>S. Rini and S. Shamai</i>	On the Capacity of the Carbon Copy onto Dirty Paper Channel	5907
<i>O. Johnson</i>	Strong Converse for Group Testing From Finite Blocklength Results	5923
<i>Y. Li, A. Kavčić, and G. Han</i>	Capacity of Multilevel NAND Flash Memory Channels	5934
<i>N. Soltani and A. J. Goldsmith</i>	Directed Information Between Connected Leaky Integrate-and-Fire Neurons	5954

	<b>SOURCE CODING</b>	
<i>R. G. Wood, T. Linder, and S. Yüksel</i>	Optimal Zero Delay Coding of Markov Sources: Stationary and Finite Memory Codes	5968
	<b>SECURE COMMUNICATION</b>	
<i>Y. Chen, O. O. Koyluoglu, and A. Sezgin</i>	Individual Secrecy for the Broadcast Channel	5981
	<b>GAUSSIAN CHANNELS</b>	
<i>J. Gutiérrez-Gutiérrez, P. M. Crespo, M. Zárrega-Rodríguez, and B. O. Hogstad</i>	Asymptotically Equivalent Sequences of Matrices and Capacity of a Discrete-Time Gaussian MIMO Channel With Memory	6000
<i>Y. Wang and M. K. Varanasi</i>	Degrees of Freedom of the Two-User MIMO Broadcast Channel With Private and Common Messages Under Hybrid CSIT Models	6004
<i>A. Gholami Davoodi and S. A. Jafar</i>	Transmitter Cooperation Under Finite Precision CSIT: A GDoF Perspective	6020
<i>A. Khina, Y. Kochman, and U. Erez</i>	The Dirty MIMO Multiple-Access Channel	6031
	<b>COMMUNICATION NETWORKS</b>	
<i>A. Sharma, K. Jagannathan, and L. R. Varshney</i>	Queuing Approaches to Principal-Agent Communication Under Information Overload	6041
<i>Y. Yang, S. Kar, and P. Grover</i>	Graph Codes for Distributed Instant Message Collection in an Arbitrary Noisy Broadcast Network	6059
<i>N. Xie and S. Weber</i>	On the Aloha Throughput-Fairness Tradeoff	6085
	<b>SEQUENCES</b>	
<i>D. Tang, C. Carlet, X. Tang, and Z. Zhou</i>	Construction of Highly Nonlinear 1-Resilient Boolean Functions With Optimal Algebraic Immunity and Provably High Fast Algebraic Immunity	6113

## Vol. 63(10): Oct. 2017.

	<b>SEQUENCES</b>	
<i>S. Jain, F. Farnoud, and J. Bruck</i>	Capacity and Expressiveness of Genomic Tandem Duplication	6129
<i>S. Mesnager, C. Tang, and Y. Qi</i>	Generalized Plateaued Functions and Admissible (Plateaued) Functions	6139
<i>C. Tang, Z. Zhou, Y. Qi, X. Zhang, C. Fan, and T. Helleseeth</i>	Generic Construction of Bent Functions and Bent Idempotents With Any Possible Algebraic Degrees	6149
<i>K. T. R. Boothby and D. J. Katz</i>	Low Correlation Sequences From Linear Combinations of Characters	6158
<i>Z. Heng, C. Ding, and Q. Yue</i>	New Constructions of Asymptotically Optimal Codebooks With Multiplicative Characters	6179
<i>Z. Sun, X. Zeng, C. Li, and T. Helleseeth</i>	Investigations on Periodic Sequences With Maximum Nonlinear Complexity	6188
	<b>CODING THEORY AND TECHNIQUES</b>	
<i>S. Moloudi, M. Lentmaier, and A. Graell i Amat</i>	Spatially Coupled Turbo-Like Codes	6199
<i>J. Gómez-Torrecillas, F. J. Lobillo, and G. Navarro</i>	A Sugiyama-Like Decoding Algorithm for Convolutional Codes	6216
<i>M. Schwartz and P. O. Vontobel</i>	Improved Lower Bounds on the Size of Balls Over Permutations With the Infinity Metric	6227
<i>M. Shi, Y. Guan, and P. Solé</i>	Two New Families of Two-Weight Codes	6240
<i>H. Liu and Y. Maouche</i>	Some Repeated-Root Constacyclic Codes Over Galois Rings	6247
<i>M. Braverman, R. Gelles, J. Mao, and R. Ostrovsky</i>	Coding for Interactive Communication Correcting Insertions and Deletions	6256
<i>C.-J. Wu, C.-H. Wang, and C.-c. Chao</i>	UEP Constructions of Quasi-Cyclic Low-Density Parity-Check Codes via Masking	6271
<i>S. Brkic, P. Ivaniš, and B. Vasić</i>	Majority Logic Decoding Under Data-Dependent Logic Gate Failures	6295
<i>M. Ye and A. Barg</i>	Explicit Constructions of Optimal-Access MDS Codes With Nearly Optimal Sub-Packetization	6307
<i>S. Goparaju, A. Fazeli, and A. Vardy</i>	Minimum Storage Regenerating Codes for All Parameters	6318
	<b>SHANNON THEORY</b>	
<i>C. T. Li and A. El Gamal</i>	Distributed Simulation of Continuous Random Variables	6329
<i>E. Fachini and J. Körner</i>	Hamilton Paths With Lasting Separation	6344
<i>E. Ahmed and A. B. Wagner</i>	Coding for the Large-Alphabet Adversarial Channel	6347
<i>N. Weinberger and N. Merhav</i>	Channel Detection in Coded Communication	6364
<i>A. Padakandla and S. S. Pradhan</i>	An Achievable Rate Region Based on Coset Codes for Multiple Access Channel With States	6393
<i>J. Nam</i>	Capacity Bounds for the $K$ -User Gaussian Interference Channel	6416
<i>H. Ghoflan and G. Kramer</i>	Models and Information Rates for Multiuser Optical Fiber Channels With Nonlinearity and Dispersion	6440

<i>Ain-ul-Aisha, L. Lai, Y. Liang, and S. Shamai</i>	On the Sum-Rate Capacity of Poisson MISO Multiple Access Channels	6457
<i>S. L. Fong and V. Y. F. Tan</i>	A Tight Upper Bound on the Second-Order Coding Rate of the Parallel Gaussian Channel With Feedback	6474
<i>Q. Wang, S. Jaggi, M. Médard, V. R. Cadambe, and M. Schwartz</i>	File Updates Under Random/Arbitrary Insertions and Deletions	6487
<b>SOURCE CODING</b>		
<i>A. Painsky, S. Rosset, and M. Feder</i>	Large Alphabet Source Coding Using Independent Component Analysis	6514
<i>S. Bilotta</i>	Variable-Length Non-Overlapping Codes	6530
<i>L. Dębowski</i>	Regular Hilberg Processes: An Example of Processes With a Vanishing Entropy Rate	6538
<i>Y. Xu, J. Chen, and Q. Wang</i>	The Sum Rate of Vector Gaussian Multiple Description Coding With Tree-Structured Covariance Distortion Constraints	6547
<b>GAUSSIAN NETWORKS</b>		
<i>A. Gholami Davoodi and S. A. Jafar</i>	Generalized Degrees of Freedom of the Symmetric $K$ User Interference Channel Under Finite Precision CSIT	6561
<i>N. Naderializadeh, A. El Gamal, and A. S. Avestimehr</i>	Fundamental Limits of Non-Coherent Interference Alignment via Matroid Theory	6573
<i>C. Hao, B. Rassouli, and B. Clerckx</i>	Achievable DoF Regions of MIMO Networks With Imperfect CSIT	6587
<i>B. Yuan, H. Sun, and S. A. Jafar</i>	Replication-Based Outer Bounds: On the Optimality of “Half the Cake” for Rank-Deficient MIMO Interference Networks	6607
<i>M. Locatelli and Z.-Q. Luo</i>	On the Complexity of Optimal Power Allocation in a Multi-Tone Multiuser Communication System	6622
<b>COMMUNICATION NETWORKS</b>		
<i>J. K. Sundararajan, D. Shah, M. Médard, and P. Sadeghi</i>	Feedback-Based Online Network Coding	6628
<i>A. Sengupta, R. Tandon, and O. Simeone</i>	Fog-Aided Wireless Networks for Content Delivery: Fundamental Latency Tradeoffs	6650
<i>G. Fanti, P. Kairouz, S. Oh, K. Ramchandran, and P. Viswanath</i>	Hiding the Rumor Source	6679
<i>J. Zhao</i>	Probabilistic Key Predistribution in Mobile Networks Resilient to Node-Capture Attacks	6714
<b>SPARSE RECOVERY, SIGNAL PROCESSING, LEARNING, ESTIMATION</b>		
<i>F. E. Rezagah, S. Jalali, E. Erkip, and H. V. Poor</i>	Compression-Based Compressed Sensing	6735
<i>M. Yuan and C.-H. Zhang</i>	Incoherent Tensor Norms and Their Applications in Higher Order Tensor Completion	6753
<i>F. Hedayati and P. L. Bartlett</i>	Exchangeability Characterizes Optimality of Sequential Normalized Maximum Likelihood and Bayesian Prediction	6767
<i>J. Jiao, K. Venkat, Y. Han, and T. Weissman</i>	Maximum Likelihood Estimation of Functionals of Discrete Distributions	6774
<b>COMPLEXITY AND CRYPTOGRAPHY</b>		
<i>C. L. Canonne, V. Guruswami, R. Meka, and M. Sudan</i>	Communication With Imperfectly Shared Randomness	6799
<i>R. Dowsley and A. C. A. Nascimento</i>	On the Oblivious Transfer Capacity of Generalized Erasure Channels Against Malicious Adversaries: The Case of Low Erasure Probability	6819
<b>COMMENTS AND CORRECTIONS</b>		
<i>N. Merhav</i>	Correction to “The Generalized Stochastic Likelihood Decoder: Random Coding and Expurgated Bounds”	6827

---

Vol. 63(11): Nov. 2017.

<b>SIGNAL PROCESSING, SPARSE RECOVERY, LEARNING, ESTIMATION</b>		
<i>E. Mohammadi, A. Fallah, and F. Marvasti</i>	Sampling and Distortion Tradeoffs for Indirect Source Retrieval	6833
<i>S. Waldron</i>	A Sharpening of the Welch Bounds and the Existence of Real and Complex Spherical $t$ -Designs	6849
<i>S. Haghghatshoar and E. Abbe</i>	Polarization of the Rényi Information Dimension With Applications to Compressed Sensing	6858
<i>N. Nguyen, D. Needell, and T. Woolf</i>	Linear Convergence of Stochastic Iterative Greedy Algorithms With Sparse Constraints	6869
<i>L. Zheng, A. Maleki, H. Weng, X. Wang, and T. Long</i>	Does $\ell_p$ -Minimization Outperform $\ell_1$ -Minimization?	6896

	<b>SHANNON THEORY</b>	
<i>H. W. Chung, B. M. Sadler, and A. O. Hero</i>	Bounds on Variance for Unimodal Distributions	6936
<i>P. Delgosha and A. Gohari</i>	Information Theoretic Cutting of a Cake	6950
<i>H. Tyagi, S. B. Venkatakrishnan, P. Viswanath, and S. Watanabe</i>	Information Complexity Density and Simulation of Protocols	6979
<i>S. Kuzuoka and S. Watanabe</i>	On Distributed Computing for Functions With Certain Structures	7003
<i>S. H. Lim, C.-Y. Wang, and M. Gastpar</i>	Information-Theoretic Caching: The Multi-User Case	7018
<i>F. Ye, K. W. Shum, and R. W. Yeung</i>	The Rate Region for Secure Distributed Storage Systems	7038
<i>A. Thangaraj</i>	Dual Capacity Upper Bounds for Noisy Runlength Constrained Channels	7052
<i>S. T. Jose and A. A. Kulkarni</i>	Linear Programming-Based Converse for Finite Blocklength Lossy Joint Source-Channel Coding	7066
<i>P. A. Stavrou, C. D. Charalambous, and C. K. Kourtellaris</i>	Sequential Necessary and Sufficient Conditions for Capacity Achieving Distributions of Channels With Memory and Feedback	7095
<i>L. Ong</i>	Optimal Finite-Length and Asymptotic Index Codes for Five or Fewer Receivers	7116
	<b>CODING THEORY AND TECHNIQUES</b>	
<i>A. Barg and A. Mazumdar</i>	Group Testing Schemes From Codes and Designs	7131
<i>M. Aldridge</i>	The Capacity of Bernoulli Nonadaptive Group Testing	7142
<i>A. Sharma and C. R. Murthy</i>	Computationally Tractable Algorithms for Finding a Subset of Non-Defective Items From a Large Population	7149
<i>Z. Chang, J. Chrisnata, M. F. Ezerman, and H. M. Kiah</i>	Rates of DNA Sequence Profiles for Practical Values of Read Lengths	7166
<i>J. Jeong, E. J. Kim, and S. Oum</i>	The “Art of Trellis Decoding” Is Fixed-Parameter Tractable	7178
<i>Y. Xia, N. Li, X. Zeng, and T. Helleseth</i>	On the Correlation Distribution for a Niho Decimation	7206
<i>S. Li, C. Ding, M. Xiong, and G. Ge</i>	Narrow-Sense BCH Codes Over $GF(q)$ With Length $n = \frac{q^m - 1}{q - 1}$	7219
<i>C. Hu</i>	Explicit Construction of AG Codes From a Curve in the Tower of Bassa-Beelen-Garcia-Stichtenoth	7237
<i>C. Shangguan, X. Wang, G. Ge, and Y. Miao</i>	New Bounds for Frameproof Codes	7247
<i>S. Shao, T. Liu, C. Tian, and C. Shen</i>	On the Tradeoff Region of Secure Exact-Repair Regenerating Codes	7253
<i>H. Lee and J. Lee</i>	An Outer Bound on the Storage-Bandwidth Tradeoff of Exact-Repair Cooperative Regenerating Codes	7267
<i>C. Li, S. Weber, and J. M. Walsh</i>	On Multi-Source Networks: Enumeration, Rate Region Computation, and Hierarchy	7283
<i>A. De Bonis and U. Vaccaro</i>	$\epsilon$ -Almost Selectors and Their Applications to Multiple-Access Communication	7304
<i>M. Kovavcević</i>	A Note on Parallel Asynchronous Channels With Arbitrary Skews	7320
	<b>COMMUNICATIONS</b>	
<i>F. Etezadi, A. Khisti, and J. Chen</i>	A Truncated Prediction Framework for Streaming Over Erasure Channels	7322
<i>D. Grant and M. K. Varanasi</i>	Feasibility of Single-Beam Interference Alignment in Multi-Carrier Interference Channels	7352
	<b>SECURE COMMUNICATION</b>	
<i>J. Liu, P. Cuff, and S. Verdú</i>	Secret Key Generation With Limited Interaction	7358
<i>M. M. Mojahedian, M. R. Aref, and A. Gohari</i>	Perfectly Secure Index Coding	7382
<i>A. J. Budkuley, B. K. Dey, and V. M. Prabhakaran</i>	Communication in the Presence of a State-Aware Adversary	7396
<i>M. Nafea and A. Yener</i>	Secure Degrees of Freedom for the MIMO Wire-Tap Channel With a Multi-Antenna Cooperative Jammer	7420
<i>L. Yu, H. Li, and W. Li</i>	Joint Source-Channel Secrecy Using Uncoded Schemes: Towards Secure Source Broadcast	7442
	<b>COMMUNICATION NETWORKS</b>	
<i>F. Xu, M. Tao, and K. Liu</i>	Fundamental Tradeoff Between Storage and Latency in Cache-Aided Wireless Interference Networks	7464
<i>Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff</i>	Update or Wait: How to Keep Your Data Fresh	7492
<i>T. Wang, Y. Shen, A. Conti, and M. Z. Win</i>	Network Navigation With Scheduling: Error Evolution	7509
<i>M. Johnston, I. Kessler, and E. Modiano</i>	Channel Probing in Opportunistic Communication Systems	7535
	<b>QUANTUM INFORMATION THEORY</b>	
<i>L. Lami, C. Hirche, G. Adesso, and A. Winter</i>	From Log-Determinant Inequalities to Gaussian Entanglement via Recoverability Theory	7553



	<b>COMPLEXITY AND CRYPTOGRAPHY</b>	
<i>N. Alon, K. Efremenko, and B. Sudakov</i>	Testing Equality in Communication Graphs	7569
<i>A. Canteaut, S. Duval, and L. Perrin</i>	A Generalisation of Dillon's APN Permutation With the Best Known Differential and Nonlinear Properties for All Fields of Size $2^{4k+2}$	7575
<i>H. Wojewódka, F. G. S. L. Brandão, A. Grudka, K. Horodecki, M. Horodecki, P. Horodecki, M. Pawłowski, R. Ramanathan, and M. Stankiewicz</i>	Amplifying the Randomness of Weak Sources Correlated With Devices	7592
	<b>COMMENTS AND CORRECTIONS</b>	
<i>F. Nan and V. Saligrama</i>	Comments on the Proof of Adaptive Stochastic Set Cover Based on Adaptive Submodularity and Its Implications for the Group Identification Problem in "Group-Based Active Query Selection for Rapid Diagnosis in Time-Critical Situations"	7612
<i>L. Bazzi</i>	Corrections to "Weight Distribution of Cosets of Small Codes With Good Dual Properties"	7615
<i>A.-L. Horlemann-Trautmann, F. Manganiello, M. Braun, and J. Rosenthal</i>	Correction to "Cyclic Orbit Codes"	7616

## Foundations and Trends® in Networking

Volume 12, Issue 1–2

Network and Protocol Architectures for Future Satellite Systems

Tomaso de Cola, Alberto Ginesi, Giovanni Giambene, George C. Polyzos, Vasilios A. Siris, Nikos Fotiou, and Yiannis Thomas

## Foundations and Trends® in Theoretical Computer Science

Volume 13, Issue 1–2

Coding for Interactive Communication: A Survey

Ran Gelles

## Problems of Information Transmission

Volume 53, Issue 3, July 2017

Two comparison theorems for distributions of Gaussian quadratic forms

M. V. Burnashev

Pages 203–214

On coupling of probability distributions and estimating the divergence through variation

V. V. Prelov

Pages 215–221

A note on random coding bounds for classical-quantum channels

M. Dalai

Pages 222–228

A special class of quasi-cyclic low-density parity-check codes based on repetition codes and permutation matrices

F. I. Ivanov

Pages 229–241

Nonlinear q-ary codes with large code distance

S. A. Stepanov

Pages 242–250

Propelinear codes related to some classes of optimal codes

I. Yu. Mogilnykh, F. I. Solov'eva

Pages 251–259

Analytical approach to multiband filter synthesis and comparison to other approaches

A. B. Bogatyrev, S. A. Goreinov, S. Yu. Lyamaev

Pages 260–273

Adaptive search for one defective in the additive group testing model

V. S. Lebedev  
Pages 274–278

On projectively invariant points of an oval with a distinguished exterior line

A. M. Balitskii, A. V. Savchik, R. F. Gafarov  
Pages 279–283

On the real complexity of a complex DFT

I. S. Sergeev  
Pages 284–293

Information-Theoretic method for classification of texts

B. Ya. Ryabko, A. E. Gus'kov, I. V. Selivanova  
Pages 294–304

Erratum to: “Remark on balanced incomplete block designs, near-resolvable block designs, and  $q$ -ary constant-weight codes”

L. A. Bassalygo, V. A. Zinoviev  
Page 305

Erratum to: “MDS codes in Doob graphs”

E. A. Beshpalov, D. S. Krotov  
Page 306

# Call for Papers

## 2018 International Zurich Seminar on Information and Communication

February 21 – 23, 2018



The 2018 International Zurich Seminar on Information and Communication will be held at the Hotel Zürichberg in Zurich, Switzerland, from Wednesday, February 21, through Friday, February 23, 2018. High-quality original contributions of both applied and theoretical nature are solicited in the areas of:

Wireless Communication

Information Theory

Coding Theory and its Applications

Detection and Estimation

Data Storage

Optical Communication

Fundamental Hardware Issues

Network Algorithms and Protocols

Network Information Theory and Coding

Cryptography and Data Security

Invited speakers will account for roughly half the talks. In order to afford the opportunity to learn from and communicate with leading experts in areas beyond one's own specialty, no parallel sessions are anticipated. All papers should be presented with a wide audience in mind.

Papers will be reviewed on the basis of a manuscript (A4, not exceeding 5 pages) of sufficient detail to permit reasonable evaluation. Authors of accepted papers will be asked to produce a manuscript not exceeding 5 pages in A4 double column format that will be published in the Proceedings. Authors will be allowed twenty minutes for presentation.

The deadline for submission is **September 17, 2017**.

Additional information will be posted at

<http://www.izs.ethz.ch/>

We look forward to seeing you at IZS.

Amos Lapidoth and Stefan M. Moser, Co-Chairs.





# Call for Papers

## 52nd Annual Conference on Information Sciences and Systems

### March 21, 22, & 23, 2018

**Princeton University - Department of Electrical Engineering  
and Technical Co-sponsorship with the  
IEEE Information Theory Society**

Authors are invited to submit previously unpublished papers describing theoretical advances, applications, and ideas in the fields of information sciences and systems including:

- Information Theory
- Coding Theory
- Image Processing
- Communications
- Signal Processing
- Machine Learning
- Statistical Inference
- Security and Privacy
- Energy Systems
- Networking
- Systems and Control
- Biological Systems

Electronic submissions of up to 6 pages (in Adobe PDF format) including 3-4 keywords must be submitted by **December 11, 2017**. Submissions should be of sufficient detail and length to permit careful reviewing. Authors will be notified of acceptance no later than **January 17, 2018**. Final manuscripts of accepted papers are to be submitted in PDF format no later than **January 31, 2018**. These are firm deadlines that will permit the distribution of Electronic Proceedings at the Conference. Accepted Papers will be allotted 20 minutes for presentation, and will be reproduced in full (up to 6 pages) in the conference proceedings. IEEE reserves the right to exclude a paper from post-conference distribution (e.g., removal from IEEE Xplore) if the paper is not presented by the author at the conference.

For more information visit us at: <http://ee-ciss.princeton.edu/>

<b>CONFERENCE COORDINATOR</b>	<b>PROGRAM DIRECTORS</b>	<b>IMPORTANT DATES</b>
<p><b><u>Lisa Lewis</u></b> Dept. of Electrical Engineering Princeton University Princeton, NJ 08544 Phone: (609) 258-6227 <b>Email: <a href="mailto:ciss@princeton.edu">ciss@princeton.edu</a></b></p>	<p><b><u>Prof. H. Vincent Poor</u></b> <b><u>Prof. Sergio Verdú</u></b> Dept. of Electrical Engineering Princeton University Princeton, NJ 08544</p>	<p><b>Paper submission deadline:</b> <b><u>December 11, 2017</u></b></p> <p><b>Notification of acceptance:</b> <b><u>January 17, 2018</u></b></p> <p><b>Final accepted manuscript due:</b> <b><u>January 31, 2018</u></b></p>



## Call for Papers

### The First Workshop on the Age of Information (AoI Workshop)

Honolulu, HI, USA April 16, 2018

<https://www.eng.auburn.edu/AoIWorkshop/>

The Age of Information is a new concept that can serve as a performance metric for characterizing the freshness of information, as a means for fundamental research, and as a tool in numerous applications. Recent research advances on the Age of Information suggest that many well-known design principles (e.g., for providing high throughput and low delay) that lead to the success of traditional data networks are inappropriate and need to be re-examined for enhancing information freshness in the rapidly emerging real-time applications.

The **1st Workshop on the Age of Information (AoI Workshop)** will be held in Honolulu on April 16, 2018, jointly organized with the [IEEE INFOCOM Conference 2018](#). Topics of interest include, but are not limited to:

- Age of Information Analysis and Optimization
- Age-based Source and Channel Coding
- Age of Information and Information Theory
- Real-time Signal Tracking and Estimation
- Age of Channel State Information
- Age of Information in Robotics and Control Systems
- Age of Information and Security
- Age of Information and Networking
- Age of Information and Game Theory
- Data Freshness in Caches and Databases
- Fresh Big Data
- Fresh Data for Online Learning
- Applications of Age of Information

The AoI Workshop calls for original and unpublished papers no longer than **6 pages**. The reviews will be single blind. The manuscripts should be formatted in standard IEEE camera-ready format (double-column, 10-pt font) and be submitted as PDF files (formatted for 8.5x11 inch paper). Manuscripts should be submitted as PDF files through the EDAS website.

#### Important Dates

Submission Deadline: Saturday, 30 December 2017 (11:59 pm EDT)

Notification of Acceptance: Wednesday, 7 February 2018

Camera-ready Deadline: Wednesday, 28 February 2018

#### Workshop Organizers

Yin Sun, Auburn University

Anthony Ephremides, University of Maryland

#### Technical Program Co-chairs

Yin Sun, Auburn University

Anthony Ephremides, University of Maryland

Bo Bai, Huawei



## Symposium Committee

### General Co-Chairs

Masayoshi Ohashi Fukuoka Univ  
Mehul Motani National Univ of Singapore

### General Secretaries

Yutaka Jitsumatsu Kyushu Univ  
Yasuyuki Nogami Okayama Univ  
Kazushi Mimura Hiroshima City Univ

### Finance

Tomoharu Shibuya Sophia Univ  
Kui Cai Singapore Univ of Tech & Design

### Publicity

Brian M. Kurkoski JAIST  
Takayuki Nozaki Yamaguchi Univ

### Publications

Masanori Hirotomo Saga Univ  
Yu Morishima National Inst of Tech, Suzuka College

### Registration

Shigeaki Kuzuoka Wakayama Univ

### Local Arrangement

Han Mao Kiah National Univ of Singapore  
Shun Watanabe Tokyo Univ of Agriculture and Tech

## Technical Program Committee

### TPC Co-Chairs

Yasutada Oohama Univ of Electro-Communications  
Pascal Vontobel The Chinese Univ of Hong Kong

### Secretary

Tomohiro Ogawa Univ of Electro-Communications

### Poster Session

Vincent Y. F. Tan National Univ of Singapore

## ISITA2018

October 28–31, 2018  
Singapore

The International Symposium on Information Theory and Its Applications (ISITA) is a leading conference on information theory. Since its inception in 1990, ISITA has been an exciting forum for interdisciplinary interaction, gathering leading researchers to discuss topics of common interest in the field of information theory. In 2018, the biennial ISITA will be held October 28–31 in Singapore at the Grand Copthorne Waterfront.

### Call for Papers

Interested authors are invited to submit papers describing novel and previously unpublished results on topics in information theory and its applications, including, but not limited to:

- Error Control Coding
- Coded Modulation
- Communication Systems
- Detection and Estimation
- Signal Processing
- Rate-Distortion Theory
- Stochastic Processes
- Network Coding
- Shannon Theory
- Coding Theory and Practice
- Data Compression and Source Coding
- Data Storage
- Mobile Communications
- Pattern Recognition and Learning
- Multi-Terminal Information Theory
- Cryptography and Data Security
- Applications of Information Theory
- Quantum Information Theory

### Paper Submission

Authors should submit papers according to the guidelines which will later appear on the conference website:

<http://www.isita2018.org/>

This link points to the permanent site <http://www.isita.ieice.org/2018/>. Accepted papers will appear in the symposium proceedings. To be published in *IEEE Xplore*, an author of an accepted paper must register and present the paper. IEEE does not guarantee inclusion in *IEEE Xplore*.

### Schedule

**Paper submission deadline April 6, 2018**

**Acceptance notification June 30, 2018**

Further information on the technical program, plenary talks, social events and registration will be posted on the symposium web site as it becomes available.

### Financial Support

The Telecommunications  
Advancement Foundation



### Sponsor

Research Society of Information Theory and Its Applications,  
Engineering Sciences Society, IEICE



### Technical Co-Sponsor

IEEE Information Theory Society



Photo: Flickr/Koiaeh



Welcome to ITW 2018 in Guangzhou! ITW 2018 solicits and welcomes original contributions on the frontiers of information theory, coding theory and their applications, as well as the frontiers with other fields such as data science, biology and signal processing. The conference structure consists of a daily plenary seminar followed by two parallel sessions throughout the day. Guangzhou is the third largest city in mainland China with a history of over 2,000 years. The conference will take place at the Sun Yat-sen Kaifeng Hotel, located within the university campus where the attendees can explore many historic architectures and artifacts, including the famous Swacey Hall, Xing Pavilion, Scholar Archway and many others. The conference also provides ample social events for better interactions among the participants. With appreciation and anticipation, we look forward to welcoming you in Guangzhou.

#### Scope of Submission

Original papers on Information and Coding Theory are encouraged for submission. The scope of submissions includes, but is not limited to

- Information Theory and its Applications
- Frontiers of Coding Theory and Practice
- Boundaries between Information Theory and Data Science, Biology and Signal Processing
- Network Information Theory
- Network Coding and Distributed Storage
- Information Theoretic Security



#### Important Dates

Paper submission : May 18, 2018  
 Acceptance notification : August 13, 2018  
 Final paper submission : September 13, 2018  
 Tutorial proposal submission: March 1, 2018  
 Tutorial acceptance notification: March 15, 2018

#### General Chairs

Pingzhi Fan - Southwest Jiaotong University, China  
 Aria Nosratinia - University of Texas at Dallas, USA  
 Li Chen - Sun Yat-sen University, China

#### Technical Program Chairs

Krishna Narayanan - Texas A&M University, USA  
 Dongning Guo - Northwestern University, USA  
 Pascal Vontobel - The Chinese University of Hong Kong, Hong Kong SAR, China  
 Daniela Tuninetti - University of Illinois at Chicago, USA

#### Publication Chairs

Shenghao Yang - The Chinese University of Hong Kong (Shenzhen), China  
 Qin Huang - Beihang University, China

#### Publicity Chairs

Xiaojun Yuan - University of Electronic Science and Technology of China, China  
 David Mitchell - New Mexico State University, USA

#### Financial Chair

Baodian Wei - Sun Yat-sen University, China

#### Chinese IT Society Liaison

Yunjiang Wang - Xidian University, China



Call For Papers

## Conference Calendar

DATE	CONFERENCE	LOCATION	WEB PAGE	DUE DATE
December 4–8, 2017	IEEE GLOBECOM	Singapore	<a href="http://globecom2017.ieee-globecom.org/">http://globecom2017.ieee-globecom.org/</a>	Passed
January 7–10, 2018	SIAM: ACM-SIAM Symposium on Discrete Algorithms (SODA)	New Orleans, Louisiana, USA	<a href="https://www.siam.org/meetings/da18/">https://www.siam.org/meetings/da18/</a>	Passed
February 21–23, 2018	2018 International Zurich Seminar on Information and Communication	Zurich, Switzerland	<a href="http://www.izs.ethz.ch/">http://www.izs.ethz.ch/</a>	September 17, 2017
March 17–22, 2018	IEEE Wireless Communications and Networking Conference (WCNC)	Barcelona, Spain	<a href="http://wcnc2018.ieee-wcnc.org/">http://wcnc2018.ieee-wcnc.org/</a>	Passed
March 21–23, 2018	52nd Annual Conference on Information Sciences and Systems (CISS)	Princeton University, USA	<a href="http://ee-ciss.princeton.edu/">http://ee-ciss.princeton.edu/</a>	December 11, 2018
April 16, 2018	The First Workshop on the Age of Information (AoI Workshop)	Honolulu, HI, USA	<a href="https://www.eng.auburn.edu/AoIWorkshop/">https://www.eng.auburn.edu/AoIWorkshop/</a>	December 30, 2017
May 7–11, 2018	16th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)	Shanghai, China	<a href="http://www.wi-opt.org/">http://www.wi-opt.org/</a>	December 10, 2017
May 7–11, 2018	European School of Information Theory (ESIT)	Bertinoro, Italy	—	—
June 17–22, 2018	IEEE International Symposium on Information Theory (ISIT)	Vail, Colorado, USA	<a href="http://www.isit2018.org">http://www.isit2018.org</a>	January 9, 2018
October 28–31, 2018	International Symposium on Information Theory and Its Applications (ISITA)	Singapore	<a href="http://www.isita2018.org">http://www.isita2018.org</a>	April 6, 2018
June 25–29, 2018	50th Annual ACM Symposium on the Theory of Computing (STOC)	Los Angeles, CA, USA	<a href="http://acm-stoc.org/stoc2018/">http://acm-stoc.org/stoc2018/</a>	Passed

Major COMSOC conferences: <http://www.comsoc.org/conf/index.html>