IEEE Information Theory Society Newsletter

Vol. 63, No. 2, June 2013 Editor: Tara Javidi ISSN 1059-2362

President's Column

Gerhard Kramer

The IEEE mission statement is "to foster technological innovation and excellence for the benefit of humanity". The IEEE has several boards including a member and geographic activities board, a standards board, and a technical activities board (TAB). The TAB is responsible for "directing the advancement of the theory and practice of electrical, electronics, communications and computer engineering … and their application for the benefit of IEEE members worldwide and for the general public." The TAB consists of seven committee chairs, ten division directors, the society/technical council presidents, the TAB secretary, and staff. Our society is in division IX that includes the signal processing, aerospace and electronic systems, geoscience and remote sensing, intelligent transportation systems, oceanic engineering, and vehicular technology societies.

The IEEE held its first TAB meeting of 2013 in Austin, Texas, in February. I had the pleasure of representing our society at its five-year review during this meeting. The review was supposed to have taken place in New Brunswick, New Jersey, in November 2012, but Hurricane Sandy intervened. The purpose of the review is to ensure that each society "maintains its vitality and technical leadership in its field of interest and is interacting appropriately with other entities."

The review went well, I think. The IEEE society review committee (SRC) notes that "our society is particularly proud about maintaining the high quality of its transactions, high participation at conferences, the level of collegiality and volunteerism that pervades its culture and the substantial involvement of its members in the technical activities of many societies." I hope that you agree! The SRC views our society as having a "strong inward focus" that has served us well and made us successful. The word "inward" is sometimes applied critically, but (as noted) many of our members are involved in the activities of other societies. The word reminds me of Shannon's admonition that "we must keep our own house in first class order … only by maintaining a thoroughly scientific attitude can we achieve real progress in communication theory and consolidate our present position."

The SRC appreciates the many changes we made in response to the recommendations they gave to Bixio Rimoldi (President 2007) at the last review. For example, the SRC praised our strong focus on students through our student committee, information theory schools, student paper award, conference student events, and mentoring. They highlighted our focus on under-represented membership through the women in the information theory society (WithITS) program. The SRC further highlighted the success of the online committee that provides us with a visible web presence, as well as the implementation of a distinguished lecturer program.

The main suggestions for improvement are to explore ways to expand our non-academic membership (we are a rather academic-oriented society) and to consider developing formal strategic planning that includes milestones, metrics, responsible persons, and resources. In the recent past, or at least during my time on the society's board of governors, strategic planning was done by ad hoc committees or by individuals who put in the time and effort to make their ideas successful. But perhaps this is a good moment to think about formal strategic planning. This will be one of the items to consider amongst the board members.

The remaining space gives me the opportunity to bring to your attention the outstanding contributions of two groups of individuals. First, I would like to appreciate our past and present online committee chairs Nick Laneman and Matthieu

continued on page 3

From the Editor

Tara Javidi

Dear IT Society members,

The second issue of 2014, in addition to our popular and regular contribution by our historian Tony Ephremides and our puzzle master Solomon Golomb, contains Abbas El Gamal's Shannon Lecture summary which was delivered last year in Boston. In addition, Øyvind Ytrehus has kindly prepared an illustrated report from the International Workshop on Coding and Cryptography which was held in April in Norway. Last but not least, on behalf of all ISIT 2013 participants, I would like to thank Elza Erkip who has prepared a "travel" note, pointing out her favorite spots in Istanbul.

As a reminder, 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 jointly at the IT Society website http:// www.itsoc.org/, using the quick links "Share News" and "Announce an Event." Articles and columns also can be e-mailed to me at ITsocietynewsletter@ece.ucsd.edu with a subject line that includes the words "IT newsletter." The next few deadlines are:

Issue	Deadline
September 2013	July 10, 2013
December 2013	October 10, 2013
March 2014	January 10, 2014

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 hear your suggestions (especially regarding the new column) and contributions.

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.

Table of Contents

The Historian's Column

Anthony Ephremides

We recently considered the indignities we all suffer at the dining table during workshops and conferences. The reactions from many readers confirmed and amplified the outrage that the "hospitality" business causes through their blatant disregard of quality constraints and their assumption of our capacity for unlimited abuse. Encouraged by these reactions, I hereby follow up on my promise to propose counter-measures of culinary resistance.

To begin with, we should realize that we have the capacity to strike back. There are several ways. The first method, which could be considered as the ultimate passive resistance, is to find the will power to forego the banquet offerings. What do I mean by that? Not to simply go elsewhere for a decent meal but, rather, to sit at the banquet table and, smilingly, leave everything untouched. We could even take out our own sandwiches and let the waiters take the plastic wraps off the table along with the intact plate contents. Imagine the logistical problem this would create. One thousand bowls of soup, returned to the kitchen, would have to be disposed of. One thousand plates of potatoes, green peas, gravy, and stringy fish would have to be stacked up in containers and returned to the trashcan of gastronomy. Indeed, despite the wide-spread guilt that this would cause, when we know that there is hunger in many parts of the world, the effectiveness of this reaction would be indisputable. The architects of "trash-gourmet" who develop their menus in rooms filled with artificial ingredients would feel shell-shocked as they would have to scramble to get rid of the returned fare. Hopefully, it would make them think twice next time they design a menu.

But clearly, this is not enough. We need to be proactive and imaginative. So, I suggest that every conference establish alongside the technical program committee a special "culinary affairs" committee. This committee would issue a "call for recipes", that would be distributed and disseminated along with the "call for papers". The proposed recipes would have to be no longer than two single-spaced pages and should provide evidence (even through simulation) that they would be enjoyed by the participants. They should include estimates of costs, list of ingredients, references, and explanations of how they would fit in the theme of the conference. For example, multi-grain bread would be appropriate for multi-packet reception capability and raw vegetables would match sessions of unprocessed observations. Double-or-triple-deck sandwiches would be a good fit for Big Data. Fish dishes would be good companions to "fishy" assertions. Well-done steak (unfortunately) would have to be accepted as a reward for well- done papers. The list can go on, but, I am sure, you get my drift. Those recipes that would be accepted would constitute the basis for the offered meals. Observance of cost bounds would be a plus but it would not be a sufficient condition for acceptance. Originality would be highly prized and eventual acceptance would lead to a possible alternative career path towards excellence in cooking. There could even be a "best recipe" award and a "best student recipe" award. Eventually, the establishment of a parallel culinary track in our conferences would lead to an ultimate "get-even" achievement or an ultimate revenge. Those who attain recognition for their suggestions could one day be called upon to do the cooking for a conference of the "hospitality industry". Imagine the expression on our faces as we would prepare opossum soup with grits or margarine-laced wonder bread for the head table of the annual gathering of dining service professionals!

But, then again, this is not enough. We need to send an unequivocal message that substandard fare cannot be tolerated in our meetings. Thus, I recommend establishing a regular column in all publications we control, in which there will be an "ombudsman"-like questions-and-answer structure where horror stories of substandard meals would be presented (in a "dear Tony", let's say, format, emulating the "dear Abby" columns), where thoughtful responses would follow the presented complaints and constructive advice would be offered as to how to avoid future recurrences.

This column does not have the ambition to assume this role and function. However, I would not be adverse to entertaining reports and complaints followed by therapeutic counter-measures to undo the damage. Like, right after a banquet in the bowels of Hilton International (somewhere), where horror stories about the fare would be told, I could suggest a number of restaurants that would "restore" faith in sound nutrition by providing appropriate antidotes. Like veal cheeks after overdone chewy steak, or succulent rigatoni after lumpy pasta, or nuessli-salad after decaying lettuce leaves, or steamed smoked ham after nitrate-laden cold cuts, or freshly roasted and brewed aromatic Arabica after brown water coffee, or, or, or ….

Errata: In my previous column I stated that the role of the prison warden in the movie with Kirk Douglas was Peter Fonda. Many alert readers noted the error. It was of course Henry Fonda and the movie was "There was a wicked man".

President's Column *continued from page 1*

Bloch. Both played the key roles in planning and implementing the online resource that we benefit from so much today. I have had the treat of working with Nick on a special issue and with Matthieu on the schools of information theory; their responsiveness and helpfulness are character traits that I admire. Second, our annual symposium in Istanbul is rapidly approaching and I'd like to acknowledge the organizers' hard work. The general chairs are Elza Erkip and Erdal Arikan, and the technical program chairs are Amos Lapidoth, Igal Sason, Jossy Sayir, and Emre Telatar (Amos and Emre are serving for a 2nd time).

I look forward to seeing you at ISIT in July.

Networks

Shannon Lecture, ISIT 2012,Boston, Massachusetts

Abbas El Gamal,
Department of Electrical Engineering, Stanford University,
350 Serra Mall, Stanford, CA 94305, USA,
Email: abbas@ee.stanford.edu, May 2, 2013

I. Introduction

The past four decades have witnessed revolutionary advances in technological networks for computation and communication. These advances have been driven by several breakthroughs:

> mathematical theories, most notably Shannon's theory of information and Turing's theory of computation;
>
> architectures (von Neumann computer, cellular and packet-switched networks, and layered protocols) and algorithms (signal processing, coding, optimization, and control) that have been motivated directly or indirectly by these theories;
>
> rapid advances in integrated circuit (VLSI) technologies (Moore's law, radio frequency circuits, computer aided design), that have made the implementation of these architectures and algorithms possible.

Growing up as an electrical engineer during this era, I have been especially intrigued by problems in networks. I spent most of my career studying theoretical and applied problems in various types of computation and communication networks. The focus of my theoretical work has been on studying limits on the performance of networks and how to achieve these limits, which is the unifying theme of my lecture.

Following in the footsteps of many previous Shannon award winners, my lecture has a strong autobiographical component mixed with historical perspective. More importantly, it is a tribute to the exceptional people I learned from, was inspired by, and collaborated with in this area.

Upon settling on the unifying theme of my lecture, I pondered whether to go deep into one topic or to skim the surface of several topics. I decided on the latter approach, to quote Berlekamp's Shannon lecture, because "it maximizes the chance that each of you will find at least one topic interesting." My lecture does not contain any new results. The problems I present are informally stated and the results are given with only proof sketches or with no proofs at all. This is because the focus is not only on the results but also on the stories behind these results.

My lecture is organized roughly chronologically in four parts. The first part is on early work on network information theory (NIT) that has received attention only in recent years. The second and third parts are on problems in computation networks that I worked on in the 1980s and that fit well with the theme of my lecture. The fourth part is on teaching NIT in a simple and unified manner, which I have dedicated much of my time to in the past ten years. I conclude with some remarks on future research in this area.

II. Network Information Theory

Previous Shannon lecturers spoke of the golden age of information theory, which should be referred to as the first golden age since there have been other periods of great progress in our field. This first golden age took place in the 1950s and early 1960s mostly at the MIT Laboratory for Information and Decision Systems (LIDS), which produced ten Shannon award winners, starting with Shannon himself, Fano, Elias, Gallager, Root, Massey, Berlekamp, Forney, Ziv, and Kailath. It is a great honor to have received the Shannon award at MIT.

I started my graduate work during the second golden age of information theory, which took place in the 1970s and early 1980s with significant contributions by members of the Information Systems Laboratory (ISL) at Stanford, which produced three previous Shannon award winners Cover, Kailath, and Gray. I was fortunate to do my PhD in the right place and at the right time. I was also fortunate to do my PhD with the right advisor, Tom Cover. I worked on several basic NIT problems—broadcast channels, relay channels, and multiple access channels with correlated sources.

Soon after graduating from Stanford, I started teaching a course on NIT and worked with exceptional researchers some of whom were graduate students at the time. We worked on a number of other basic NIT problems—multiple description coding, relay networks, channel with state, and the interference channel. However, many of the results in my thesis and this subsequent work received attention only recently. Some of these results, to borrow Gallager's metaphor, seemed like "leaves on the knowledge tree" with no potential subsequent developments, but turned out much later to have been budding shoots which developed into healthy branches with many leaves of their own. In the following two subsections, I discuss two of these results.

A. Compress–Forward for the Relay Channel

The relay channel (RC) is a canonical example of a multihop network in which a sender X_1 wishes to communicate a message $M \in [1:2^{nR}]$ to a receiver Y_3 with the help of a relay; see Figure 1. In network information theory, we seek to find the capacity of this channel, which is the highest achievable rate R in bits/transmission, and the optimal coding scheme that achieves it. This relay channel problem was first studied by van der Meulen in his 1971

PhD thesis on multi-way networks. The capacity of this channel is not known in general and it is considered as one of the key open problems in NIT.

I first learned about the relay channel problem while visiting the University of Hawaii (UH) in the spring of 1976. Tom Cover was invited by Norm Abramson (who was his doctoral advisor at Stanford) to spend a month at UH and he asked me to join him. At that time, Norm Abramson was leading the celebrated ALOHA packet radio network project, which produced many of the key architectural concepts and protocols of modern data communication networks. David Slepian, one of the early giants of information theory and a Shannon award winner, mentioned the relay channel problem to me and suggested that I work on it for my thesis.

My work with Tom on the relay channel is described in the paper "Capacity Theorems for the Relay Channel" [1]. In this paper we established upper and lower bounds on the capacity of the relay channel and showed that they are tight in some special cases. We established what is now known as the cutset upper bound on the capacity.

Theorem 1 (Cutset upper bound):

$$C \le \max_{p(x_1,x_2)} \min\{I(X_1, X_2; Y_3), I(X_1; Y_2, Y_3|X_2)\}.$$

This upper bound was directly motivated by the celebrated max-flow min-cut theorem by Ford and Fulkerson [2]. The bound has a very intuitive interpretation: Given any sequence of codes with diminishing probability of error, the rate of this sequence of codes cannot be higher than the rate of communication from (X_1, X_2) together to Y_3 and also the rate from X_1 to (Y_2, Y_3) together. The cutset bound turned out to be tight for almost all relay channels with known capacities. I later extended this bound to networks [3] and this extension has been shown to be tight for most networks with known capacity regions.

In the 1979 paper, we also proposed several block Markov coding schemes in which $b-1$ messages $M_j \in [1:2^{nR}]$, $j \in [1:b-1]$, are sent over b blocks of n-transmissions and the codewords transmitted in each block can depend on the message sent in the previous block. The first block Markov scheme we introduced is what is now known as decode–forward (DF). In this scheme, the relay recovers the message M_j sent in block j and coherently cooperates with the sender to communicate the previous message M_{j-1} to the relay. We showed that this scheme achieves the following lower bound on the capacity.

Theorem 2 (Decode–forward lower bound):

$$C \ge \max_{p(x_1,x_2)} \min\{I(X_1, X_2; Y_3), I(X_1; Y_2|X_2)\}.$$

In Theorem 1 of [1], we showed that this bound coincides with the cutset bound if the relay channel is physically degraded. At first, the assumption of physical degradedness did not seem well motivated at all. We were surprised, however, to find a very well motivated setup in which the relay channel is physically degraded, which is when there is feedback from the receiver to the relay as depicted in Figure 2. What makes this result even more surprising is that it is a rare example of a channel for which the capacity is not known, but the capacity with feedback is known.

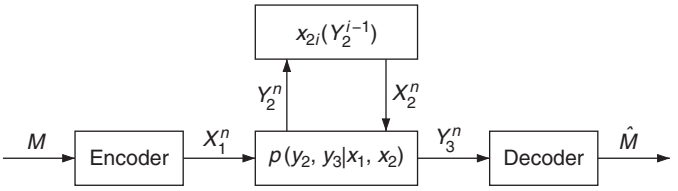

Fig. 1 Relay channel.

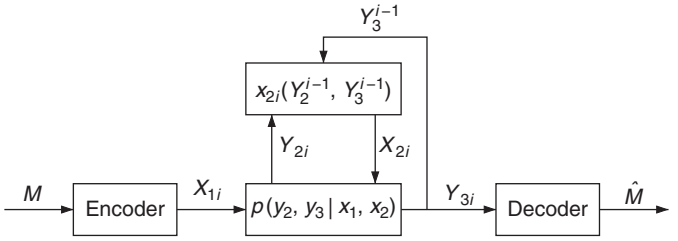

Fig. 2 Relay channel with feedback.

We then observed that when the channel from the sender to the relay is not stronger than that to the receiver, DF does not perform well because it requires the relay to recover the entire message. This led us to develop two other block Markov schemes. The first of these two schemes is partial decode–forward (PDF), which is a natural extension of DF. In this scheme, the relay recovers only part of the message M_j' sent in block j and coherently cooperates with the sender to communicate M_{j-1}' to the receiver. The rest of the message M_j'' is sent to the receiver directly. This schemes yields the lower bound,

$$C \ge \max_{p(u,x_1,x_2)} \min\{I(X_1, X_2; Y_3), I(U; Y_2|X_2) + I(X_1; Y_3|X_2, U)\},$$

which is a special case of Theorem 7 in [1]. In [4], Aref and I showed that this scheme is optimal for some nontrivial classes of relay channels and networks for which the capacity also coincides with the cutset bound. So early on, we established optimality results for both DF and PDF.

The last scheme we introduced in the 1979 paper is what is now known as compress–forward (CF). In this scheme, the relay doesn't recover any part of the message, but rather sends a description $\hat{Y}_2^n$ of its received sequence to the receiver Y_3 in a manner similar to the Wyner–Ziv scheme for lossy compression with side information. This scheme yields the following bound.

Theorem 3 (Compress–forward lower bound):

$$C \ge \max_{\substack{p(x_1)p(x_2)p(\hat{y}_2|y_2,x_2):\\ I(X_2;Y_3)\ge I(Y_2;\hat{Y}_2|X_2,Y_3)}} I(X_1; \hat{Y}_2, Y_3|X_2).$$

Unlike the DF and PDF schemes, however, we did not establish any optimality results for CF and it seemed like a dead-end. In fact the landmark book by Cover and Thomas [5] does not mention this scheme at all!

Over 20 years later, CF turned out to be as or more fundamental than DF. In 2007, Cover and Kim [6] showed that CF is optimal for a deterministic RC. In 2009, Aleksic, Razaghi, and Yu [7] found another example where CF is optimal. More interestingly, their example shows that the cutset bound is not tight in general. Around the same time, Mohseni, Zahedi and I [8] stumbled upon the alternative characterization of the CF bound

$$C \geq \max_{p(x_1)p(x_2)p(\hat{y}_2|y_2,x_2)} \min\{I(X_1, X_2; Y_3) - I(Y_2; \hat{Y}_2 | X_1, X_2, Y_3), \\ I(X_1; \hat{Y}_2, Y_3 | X_2)\},$$

which has a similar form to the cutset bound.

In what at first seemed like a separate line of investigation from the relay channel, in 2000, Ahlswede, Cai, Li, and Yeung [9] developed the celebrated network coding scheme for multicast graphical networks. Their scheme was generalized to various classes of deterministic networks by Ratnakar and Kramer [10] and Avestimehr, Diggavi, and Tse [11] and to erasure networks by Dana, Gowaikar, Palanki, Hassibi, and Effros [12]. Recently, Kim and I [13] developed the noisy network coding scheme, which naturally extends the aforementioned alternative characterization of CF to networks and includes network coding and its extensions to deterministic and erasure networks as special cases. Thus, the CF scheme, which at first did not seem very promising, turned out to be the start of a general coding scheme for noisy networks.

B. Capacity of Deterministic Interference Channel

The second result that initially seemed like a dead end, but later helped lead to some very interesting work is the capacity region of a class of deterministic interference channels [14]. The interference channel (IC) is a canonical example of a single-hop network in which sender-receiver pairs wish to communicate to each other over a shared medium, such as a wireless channel. For the case of a two sender-receiver pair depicted in Figure 3, sender 1 wishes to communicate a message $M_1 \in [1 : 2^{nR_1}]$ to receiver 1 and sender 2 wishes to communicate a message $M_2 \in [1 : 2^{nR_2}]$ to receiver 2. We seek to find the capacity region, which is the set of simultaneously achievable rate pairs (R_1, R_2), and the coding scheme that achieves it. This interference channel model was first studied by Ahlswede [15], a Shannon award winner. The capacity region

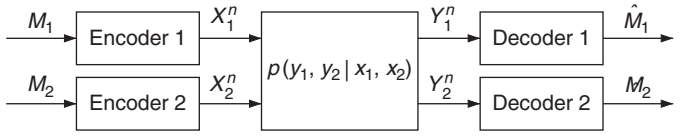

Fig. 3 Interference channel.

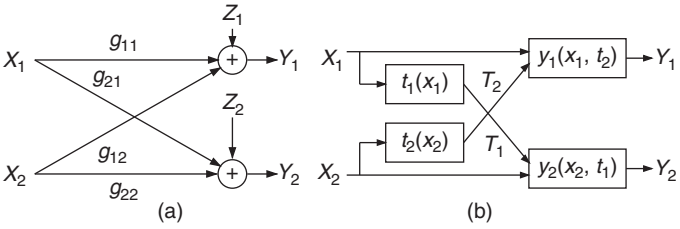

Fig. 4 (a) Gaussian interference channel. (b) Injective deterministic interference channel.

of the interference channel is also not known in general and it is considered as one of the key open problem in NIT.

In 1981, Han (a Shannon award winner) and Kobayashi [16] established the tightest known inner bound on the capacity region. Max Costa, who was a student at ISL in the early 1980s, was interested in investigating the optimality of the Han–Kobayashi scheme for the Gaussian interference channel depicted in Figure 4-(a). This somehow led us to introduce the deterministic channel depicted in Figure 4-(b) in which the functions y_1 and y_2 are injective in the interfering signals t_1 and t_2, respectively. This condition holds for example when the functions y_1 and y_2 are addition as in the Gaussian model.

Costa and I showed that the Han–Kobayashi scheme is optimal for this class. This result remains as the only interference channel example for which the Han–Kobayashi scheme in its full generality is optimal. It is also the first converse proof to use the idea of a genie (although this term was introduced much later), which has been used in several converses since then, e.g., [17], [18]. At the time, however, we did not see a precise connection between this deterministic model and the more practically motivated Gaussian IC.

Twenty five years later, the dots were connected between our deterministic model and the Gaussian IC. Etkin, Tse, and Wang [18] used our result in their proof of the 1-bit theorem for the Gaussian IC, which shows that the Han–Kobayashi inner bound is within 1-bit per dimension of the capacity region of the Gaussian interference channel. The connection between our deterministic model and the Gaussian IC was formalized further by Telatar and Tse [19]. Avestimehr, Diggavi, and Tse [11] later showed that a special class of our deterministic model can be used to approximate Gaussian interference channels in high SNR. Their work has spawned a new area of investigation in network information theory. Here again a result that at first seemed like a dead end turned out to be the beginning of a promising new direction in the field.

III. VLSI Theory

My first faculty job was at USC, which produced five Shannon award winners, Reed, Golomb, Viterbi, Welch, and Gray. It was indeed a privilege to be a colleague of some of the communication and coding giants. The greatest influence on my career during this time, however, came from Carver Mead, a renowned Caltech device physicist and a National Medal of Technology winner, who is considered the father of modern VLSI design. At the time, integrated circuits (chips) were designed by hand using a trial-and-error approach and the design process was highly compartmentalized—system and algorithm designers knew nothing about chip design and chip designers knew nothing about algorithms or system design. Carver was among the first to recognize that this trial-and-error approach will not scale with Moore's law. He introduced the so-called silicon compiler approach to design and evangelized it through a very influential textbook with Lynn Conway, *Introduction to VLSI Systems* [20]. His approach has enabled the design of today's complex chips not only by dramatically improving design productivity, but also by making it possible for system and algorithm developers to become directly involved in high level chip design.

Meeting Carver while at USC changed the course of my career. I ended up spending a significant amount of my time working on various applied and theoretical problems in VLSI design and design automation. I will describe two results in this area that fit well with the theme of my lecture.

A. VLSI Complexity

Around the same time I met Carver Mead, Thompson [21] introduced the notion of VLSI area-time complexity in his doctoral dissertation at CMU. I will briefly describe his basic result. Consider a rectangular chip for computing a function g of n binary variables; see Figure 5. Assume that the computation network (circuit) for g is layed out on a grid in which each grid point can have an input, an output, a logic gate, a memory cell, or a constant number of wire crossings; and every grid line can carry a constant number of wires. Further assume that each wire can carry only a constant number of bits per clock cycle of operation. These assumptions are not as restrictive as they may seem and the results can be extended to other physical computing devices. The performance metrics that Thompson considered are chip area A in square grid units and compute time T in clock cycles. Thompson then used the following cutset argument to establish a lower bound on AT^2. Bisect the chip along its width such that each side of the bisected chip has $n/2$ inputs. Now, let I be the minimum number of bits exchanged during computation between the two sides of the bisection in any chip that computes g. From the above assumptions, $I \le cW \times T \le c\sqrt{A}\,T$. Squaring both sides, we obtain $AT^2 \ge cI^2$, or in order notation $AT^2 = \Omega(I^2)$.

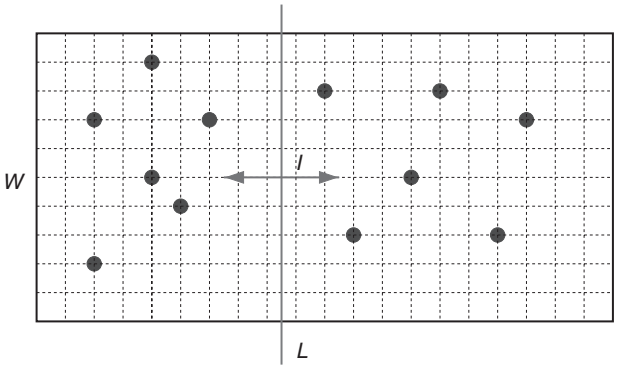

Fig. 5 Thompson's chip model and cutset argument.

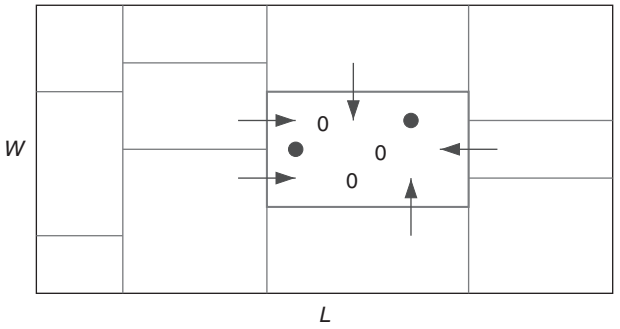

Fig. 6 Partitioning argument used in the proof of the lower bound on the AT^2 complexity of decoding.

Hence, to evaluate this lower bound, one needs to find the minimum amount of information flow to compute a function with distributed inputs. Thompson and others showed that this bound is order optimal or close to optimal for computing many popular functions such as sorting, DFT, and matrix multiplication.

Because of my interest in information theory and coding, we decided to study the VLSI complexity of error correction coding. Consider a chip that implements encoding or decoding for an (n,R,t) error correction code (ECC), where n is the block length, R is the rate, and t is the number of errors that the code can correct. Earlier arguments by Savage [22] on the circuit complexity of coding imply that any such chip must satisfy the lower bound $AT^2 \ge n\log(t+1)$. This lower bound, however, turned out to be very loose. Using Thompson's approach, Greene, Pang, and I [23], showed that

$$AT^2 = \Omega(nR^2t),$$

which is a much tighter bound than the bound by Savage, especially for large t, for example, when the code has nonzero rate. However, we could not use Thompson's bisection argument to establish this bound because in some cases one can use two separate t-error correcting codes each with half the block length, and place their encoders or decoders side by side on the chip, in which case no communication needs to take place between them.

To overcome this difficulty, we partition the chip into $2n/t$ rectangular blocks such that each block has exactly $Rt/2$ outputs; see Figure 6. It is easy to see that on average each block has $t/2$ inputs and perimeter $\Theta(\sqrt{At/n})$. Now using a pigeon hole argument, there exists a block with at most t inputs and perimeter less than or equal to twice the average. Now, suppose that all the chip inputs in this block are set to zero by errors, then since we are using a t-error correcting code, at least $Rt/2$ bits must flow into the block to determine the correct outputs from this block. But the number of bits flowing into the block $I \le cT \times$ (block perimeter). Combining these bounds, we obtain $AT^2 = \Omega(nR^2t)$.

In [23], we extended this result to obtain a lower bound on any chip for encoding or decoding an error correction code with block length n, rate R, and error probability P_e.

This work received absolutely no attention in the 30 years since we presented it at the MIT conference on VLSI. In fact I had completely forgotten about it until a few months ago, Goldsmith and Grover told me that they stumbled upon our paper while working on the limits on power consumption of error correction chips [24].

B. Configuring VLSI Arrays Around Defects

The second problem in VLSI theory that I will discuss concerns limits on configuring VLSI arrays around defects. VLSI chips are subject to manufacturing defects. Discarding every chip that has some defect can make the yield (fraction of "good" chips) unacceptably low. This has motivated much work on techniques for fault tolerance. One way to achieve fault tolerance is to treat the defects as temporal or intermittent noise and to use coding. However, since defects do not change with time, coding redundancy (in terms of time and space) can be reduced either by finding the defects after manufacture and using this knowledge as side information at the encoder as in the work of Kuznetsov and Tsybakov on memory with defects [25], which led to the Gelfand–Pinsker [26] and the Costa [27] writing on dirty paper schemes;

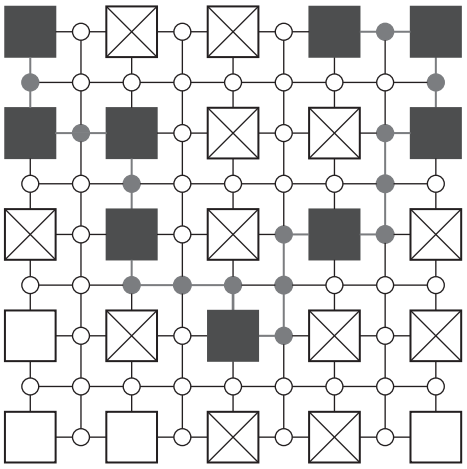

Fig. 7 Array of $\sqrt{n} \times \sqrt{n}$ processors configured into a chain of k good processors.

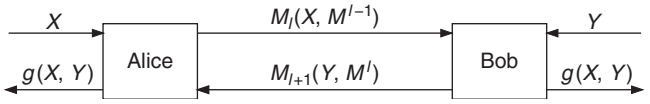

Fig. 8 Communication complexity setup.

or by using the information about the defects to reconfigure the system around them as is routinely performed in computer memory chips.

Greene and I investigated the latter approach for processor arrays. I will give an example of this work. Suppose we wish to build a chip with a chain of k processors. Assume that each processor is defective with probability p independent of all other processors. One way to build such a chain is to build a chip with a larger number of processors n and configurable connections and then configure k good processors into a chain. By the law of large numbers, this can be done with high probability (whp) if $n > k/(1-p)$. Since the purpose of configuration is to reduce the chip area overhead needed for fault tolerance, we would like n to be as close as possible to this limit. However, if $n = \Theta(k)$, it is easy to show that the longest interprocessor connection grows as $\Theta(\log n)$ whp, which can result in unacceptably high interprocessor delay, a key metric in chip design.

It turns out that this large delay problem can be solved by using a 2-D array of $\sqrt{n} \times \sqrt{n}$ processors with configurable connections and configuring k good processors into a chain as shown in Figure 7. Greene and I [28] showed that this can now be done with high probability with $n = \Theta(k)$ and with a *constant* interprocessor connection length! Our proof of this result used percolation theory.

Although this work did not lead to practical implementations, it motivated several key inventions in field programmable gate arrays (FPGAs) [29], which are chips that can be electrically configured to implement different digital systems. These types of chips are now widely used in electronic systems and for prototyping complex chips.

IV. Communication Complexity

Network information theory, which is the topic of the first part of my lecture, deals mainly with limits on information flow for communication. We assumed an information theoretic model with large transmission blocks and diminishing probability of error and aimed to find a single letter characterization of the achievable rates or at least bounds on them.

VLSI complexity, which is the topic of the second part of my lecture, deals with the question of how much information needs to be exchanged between various parts of a chip to be able to compute a function. We assumed a deterministic, single-instance model and zero error, and aimed to find the minimum amount of bits exchanged or bounds on it.

Around the same time Thompson introduced VLSI complexity, Andrew Yao, a Turing award winner, independently introduced the notion of communication complexity [30], which studies the limits on information flow for computing, for example, as encountered in Thompson's work. Yao considered two communication nodes Alice and Bob. Alice observes X and Bob observes Y, where X and Y are drawn from finite sets; see Figure 8. Alice and Bob wish to compute the same function $g(X, Y)$. To achieve this goal, they communicate over a noiseless two-way link in rounds such that the message sent by Alice in odd round l is a function of X and all previously transmitted messages, and similarly for Bob. The questions Yao posed are: what is the communication complexity $C(g)$, which is the minimum number of bits that need to be exchanged between Alice and Bob in order for both of them to compute g, and what is the protocol that achieves $C(g)$.

Pang, Orlitsky, and I resolved several conjectures in Yao's original work. I will give an example of this work which was suggested to us by Tom Cover.

Example (Communication complexity of cyclic shift) [31]: Let X be a binary n-sequence and Y be an arbitrary cyclic shift of X. The function that both Alice and Bob wish to compute is the shift amount, which ranges between 0 and $n-1$ (so in effect they each wish to know the other's sequence). It is easy to see that $C(g) \le n + \log n$, and in the paper with Orlitsky, we established a general lower bound which, when specialized to this example, gives $2\log n(1 - 2^{-n/2})$. As can be seen, there is a very large gap between these two bounds.

We showed that the shift can be determined with only $2\log n$, which is very close to the lower bound. The scheme is quite simple. View X, Y, and their shifts as binary numbers and let Z be the unique largest such number among all shifts, which both Alice and Bob can separately find. Alice sends the shift amount from Z to X to Bob, and Bob sends the shift amount from Z to Y to Alice. Clearly both can now find the number of shifts between X and Y with at most $2\log n$ bits of communication.

A. Noisy Broadcast Network

Yao's communication complexity setup assumes zero error communication. Real-world computing networks, however, suffer from noise. Hence, a natural question to ask is: What is the communication complexity of computing a function in the presence of noise. This question is intimately related to work on reliable computing using unreliable components studied by von Neumann, Moore, Shannon, Elias, Dobrushin, and Winograd, among many others, under different setups and with different conclusions; see for example [32].

Fig. 9 Noisy broadcast network setup.

Motivated by this work and by the question of communication complexity under noise, in a 1984 workshop organized by Cover and Gopinath [33] I proposed a toy problem of reliable distributed computing that later became known as the noisy broadcast network problem. Consider a broadcast network of n nodes in which each node has an arbitrary bit $s_j \in \{0,1\}$; see Figure 9. Node 1 (or all the nodes) wishes to compute the parity of all the bits in the network. To compute the parity, the nodes communicate in rounds over the network. Without loss of generality, assume that only one bit is transmitted in each round. Each transmitted bit from each node j can depend on its source bit s_j and past received bits. The way I modeled the noise is to assume that each bit is received via an independent binary symmetric channel (BSC) with parameter p, that is, each received bit is in error with probability p independent of all other received bits. The questions I posed are: what is the communication complexity C_ϵ, which is the minimum number of bits that need to be exchanged so that node 1 can compute the parity of all the bits in the network with probability of error $P_e < \epsilon$, $\epsilon < 1/2$, and what is the protocol that achieves C_ϵ.

Since each node (except node 1) must transmit its bit at least once, one can obtain a trivial lower bound $C_\epsilon = \Omega(n)$. We can also easily establish the upper bound $C_\epsilon = O(n \log n)$—each node broadcasts its bit $\log n$ times so that node 1 can estimate each bit whp and then compute the parity with the desired error probability. Can we do better in terms of the number of transmissions needed? Note that every time a bit is transmitted there are $n-1$ independent observations of it in the network. Hence, there is more than enough information to recover the bit whp. The question is how to combine these observations without using too much communication.

Right after I presented this problem at the Cover–Gopinath workshop, Gallager [34] devised a very clever scheme that requires only $n \log \log n$ transmissions.

In Gallager's scheme, each node first broadcasts its bit $\Theta(\log \log n)$ times. The nodes are pre-partitioned into groups of $\Theta(\log n)$ nodes. Each node in the network then estimates the parity of its group from the bits it has received and broadcasts its estimate. Finally, node 1 makes a reliable estimate of each group's parity and adds them up modulo 2 to estimate the overall parity with the desired P_e.

Gallager showed that computing all the bits reliably can be done also with only $O(n \log \log n)$ transmissions. His scheme is similar to the one for computing the parity except that now each node broadcasts an estimate of the parity of a different group of nodes.

This problem received no attention from the information theory community after Gallager's work. In 1997, Yao [35] became quite excited about this problem and encouraged the theoretical CS community to study it as a simple example of reliable distributed computing. Since then there have been tens of papers on this problem and many variations of it. In 2008, Goyal, Kindler, and Saks [36] showed that Gallager's scheme for recovering all the bits is order optimal. More surprisingly, they showed that for computing the parity only $\Theta(n)$ bits is needed!! Their scheme estimates the Hamming weight of the bits from which the parity can then be recovered with the desired error probability. Each node broadcast its bit a constant number of times, and then the Hamming weight is estimated from answers to $O(n)$ binary questions about the nodes' estimates of the weight.

V. Teaching Network Information Theory

The third golden age of information theory started in the mid 1990s and has been fueled by the advent of the Internet and wireless cellular communication with contributions by many researchers around the world. I was drawn back to network information theory mainly by students interested in the applications of NIT to communications and multimedia. This lead to collaborations with some great students and researchers. We worked on some of the old problems as well as on some new ones such as at the intersection of information theory and networking. However, the most significant project I have been involved in since then is how to teach NIT in a unified and accessible manner.

I started teaching network information theory again in 2002. This class had several of our rising stars, including Young-Han Kim. We developed a fairly comprehensive set of notes on the subject and converted them into a textbook [13]. The book presents the models, results, and techniques of NIT in a simple and unified manner that makes the subject accessible to a wide audience. I will share my views on how to teach NIT and what to teach in a course on NIT. There are several viable ways to teach a first course on information theory and for details about these different ways, I refer the reader to Sergio Verdú's wonderful Shannon lecture. However, there is currently only one viable way to teach NIT in a unified manner, which is using random coding, typicality, and the weak converse.

There are many definitions of typicality in the literature. We chose the notion of robust typicality by Orlitsky and Roche [37] because it can be used to prove achievability for all discrete memoryless systems using a few simple lemmas. Also, under this definition of typicality, lossless source coding becomes an immediate corollary of lossy source coding, which not only unifies these two source coding problems, but more importantly because in some cases it is easier to first prove results for lossy source coding and then specialize them to the lossless case than to prove them for lossless source coding directly, for example, see Sections 11.2 and 21.1 in [13].

What about achievability proofs for the very popular Gaussian models? The way we decided to prove achievability for Gaussian sources and channels is to first extend the achievability proofs for their discrete memoryless counterparts to the discrete memoryless models with cost, then to discretize the signals in the original model and apply appropriate limits [38]. A key point I would like to emphasize here is that even if one is interested only in Gaussian models, it is far better to prove the result (whenever possible)

first for the DM counterparts. Just imagine where information theory would be today if Shannon had considered only Gaussian channels.

If we follow this approach for teaching NIT, the students need to know only basic probability and some facts about MSE estimation and convexity. The course syllabus itself can be customized to the audience depending on whether it is a first or a second course in information theory, whether the audience is interested in theory or applications, for example to wireless communication, or whether they are interested in channel or source coding.

As an example, the NIT course I teach at Stanford is aimed at EE graduate students in information theory, communications, networking, and multimedia. The course focuses on models, coding schemes, and techniques rather than pure lemma–theorem–proof. For achievability, I first introduce some useful and general lemmas, and then go through the proofs for some of the basic network building blocks and give only rough sketches of the proofs for more complex network models. For the converse, I go over a few proofs in detail just to illustrate the key approaches and techniques used. The course includes final projects—some proposed by us and others proposed by the students. Many of these projects were developed further into conference and journal publications.

VI. Looking Ahead

My lecture presented examples of work on the performance limits for computation and computing networks. Some of the results were unexpected like good jokes to quote Cover. Some were like buds on a tree—it is difficult predict how they will develop. Others led to unintended consequences (e.g., configuring VLSI arrays led to the development of FPGAs). There were several recurring themes in the setups of the different problems and in the proof techniques. Finally I argued that NIT is now ready to be taught to a wider audiences. In particular, the communication community should seriously consider teaching NIT as part of a standard graduate curriculum on communications and networking.

Looking ahead, there is still much to be done on on performance limits of networks. In particular, there are many open problems in NIT, for example, the capacity of the broadcast, interference, and relay channels. These problems appear to be quite difficult and as a result many of the talks on NIT open with: Here is the information flow problem. We don't know the capacity even for simple building blocks, so let's do something else—approximate, find capacity scaling, study wired networks, …. These are all interesting and fruitful research directions, but pursuing them should not discourage us from continuing to make progress on the basic open problems in the field. After all, it is these basic problems that have had the most impact on both theory and practice. In working on these basic problems we should be patient and not despair because we cannot find the answer. We should keep in mind what Shannon said in his 1956 Bandwagon paper, "Seldom do more than a few of nature's secrets give way at one time."

Beyond continuing to work on the basic open problems, we should also think broadly. The founder of our field worked in many areas [39]. In particular there are many exciting opportunities in new types of networks, such as smart power grids, nano and quantum computing, biological networks, social networks, and economic networks. In deciding on problems to work on, we should take guidance from Shannon, "I am more interested in the elegance of a problem. Is it a good problem, an interesting problem?"

VII. Acknowledgments and Dedication

In addition to the people I mentioned in my lecture, I would like to also thank Young-Han Kim, Alon Orlitsky, Bernd Bandemer, John Gill, and Pulkit Grover for help with the preparation of this lecture. Most of the work in this lecture was supported in part by DARPA and NSF grants.

Finally, I dedicate this lecture to my mentor, colleague, and friend Tom Cover who passed away less than four months before I delivered this lecture.

References

[1] T. M. Cover and A. El Gamal, "Capacity theorems for the relay channel," *IEEE Trans. Inf. Theory,* vol. 25, no. 5, pp. 572–584, Sep. 1979.

[2] L. R. Ford, Jr. and D. R. Fulkerson, "Maximal flow through a network," *Canad. J. Math.,* vol. 8, no. 3, pp. 399–404, 1956.

[3] A. El Gamal, "On information flow in relay networks," in *Proc. IEEE National Telecomm. Conf.*, New Orleans, LA, Nov. 1981, vol. 2, pp. D4.1.1–D4.1.4.

[4] A. El Gamal and M. R. Aref, "The capacity of the semideterministic relay channel," *IEEE Trans. Inf. Theory,* vol. 28, no. 3, p. 536, May 1982.

[5] T. M. Cover and J. A. Thomas, *Elements of Information Theory,* 2nd ed. New York: Wiley, 2006.

[6] T. M. Cover and Y.-H. Kim, "Capacity of a class of deterministic relay channels," in *Proc. IEEE Int. Symp. Inf. Theory, Nice,* France, June 2007, pp. 591–595.

[7] M. Aleksic, P. Razaghi, and W. Yu, "Capacity of a class of modulo-sum relay channels," *IEEE Trans. Inf. Theory,* vol. 55, no. 3, pp. 921–930, 2009. [Online]. Available: http://dx.doi.org/10.1109/TIT.2008.2011518

[8] A. El Gamal, M. Mohseni, and S. Zahedi, "Bounds on capacity and minimum energy-per-bit for AWGN relay channels," *IEEE Trans. Inf. Theory,* vol. 52, no. 4, pp. 1545–1561, 2006.

[9] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory,* vol. 46, no. 4, pp. 1204–1216, 2000.

[10] N. Ratnakar and G. Kramer, "The multicast capacity of deterministic relay networks with no interference," *IEEE Trans. Inf. Theory,* vol. 52, no. 6, pp. 2425–2432, 2006.

[11] A. S. Avestimehr, S. N. Diggavi, and D. N. C. Tse, "Wireless network information flow: A deterministic approach," *IEEE Trans. Inf. Theory,* vol. 57, no. 4, pp. 1872–1905, Apr. 2011.

[12] A. F. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of wireless erasure networks," *IEEE Trans. Inf. Theory,* vol. 52, no. 3, pp. 789–804, 2006.

[13] A. El Gamal and Y.-H. Kim, *Network Information Theory.* Cambridge: Cambridge, 2011.

[14] A. El Gamal and M. H. M. Costa, "The capacity region of a class of deterministic interference channels," *IEEE Trans. Inf. Theory,* vol. 28, no. 2, pp. 343–346, 1982.

[15] R. Ahlswede, "The capacity region of a channel with two senders and two receivers," *Ann. Probability,* vol. 2, no. 5, pp. 805–814, 1974.

[16] T. S. Han and K. Kobayashi, "A new achievable rate region for the interference channel," *IEEE Trans. Inf. Theory,* vol. 27, no. 1, pp. 49–60, 1981.

[17] G. Kramer, "Outer bounds on the capacity of Gaussian interference channels," *IEEE Trans. Inf. Theory,* vol. 50, no. 3, pp. 581–586, 2004.

[18] R. Etkin, D. N. C. Tse, and H. Wang, "Gaussian interference channel capacity to within one bit," *IEEE Trans. Inf. Theory,* vol. 54, no. 12, pp. 5534–5562, Dec. 2008.

[19] İ. E. Telatar and D. N. C. Tse, "Bounds on the capacity region of a class of interference channels," in *Proc. IEEE Int. Symp. Inf. Theory, Nice,* France, June 2007, pp. 2871–2874.

[20] C. Mead and L. Conway, *Introduction to VLSI Systems,* Philippines, 1980.

[21] C. Thompson, "A complexity theory for vlsi," Ph.D. Thesis, Carnegie–Mellon University, Pittsburgh, PA, Aug. 1980.

[22] J. Savage, "Complexity of decoders: II – computational work and decoding time," *IEEE Trans. Inf. Theory,* vol. 17, no. 1, pp. 77–85, 1971.

[23] A. El Gamal, J. Greene, and K. Pang, "VLSI complexity of coding," in *Proc. of the MIT Conference on Advanced Research in VLSI,* Cambridge, MA, Apr. 1984, pp. 150–158.

[24] P. Grover, A. Goldsmith, and A. Sahai, "Fundamental limits on complexity and power consumption in coded communication," in *Proc. IEEE Int. Symp. Inf. Theory,* MIT, MA, Jul. 2012.

[25] A. V. Kuznetsov and B. S. Tsybakov, "Coding in a memory with defective cells," *Probl. Inf. Transm.,* vol. 10, no. 2, pp. 52–60, 1974.

[26] S. I. Gelfand and M. S. Pinsker, "Coding for channel with random parameters," *Probl. Control Inf. Theory,* vol. 9, no. 1, pp. 19–31, 1980.

[27] M. H. M. Costa, "Writing on dirty paper," *IEEE Trans. Inf. Theory,* vol. 29, no. 3, pp. 439–441, 1983.

[28] J. Greene and A. El Gamal, "Configuration of VLSI arrays in the presence of defects," *Journal of the Association for Computing Machinery,* vol. 31, no. 4, pp. 694–717, 1984.

[29] J. Rose, A. El Gamal, and A. Sangiovanni-Vincentelli, "Architecture of field-programmable gate arrays," *Proceedings of the IEEE,* vol. 81, no. 7, pp. 1013–1029, 1993.

[30] A. C.-C. Yao, "Some complexity questions related to distributive computing," in *Proc. 11th Ann. ACM Symp. Theory Comput.,* Atlanta, Georgia, 1979, pp. 209–213.

[31] A. El Gamal and A. Orlitsky, "Interactive data compression," in *Proc. 25th Ann. Symp. Found. Comput. Sci.*, Washington DC, Oct. 1984, pp. 100–108.

[32] N. Pippenger, "Development in the synthesis of reliable organisms from unreliable components," in *Proc. of Symposia in Pure Mathematics,* vol. 50, 1990, pp. 311–324.

[33] T. Cover and B. Gopinath, *Open Problems in Communication and Computation.* Springer-Verlag, 1987.

[34] R. Gallager, "Finding parity in a simple broadcast network," *IEEE Trans. Inf. Theory,* vol. 34, no. 2, pp. 176–180, 1988.

[35] A. Yao, "On the complexity of communication under noise," in *5th ITCS,* 1997.

[36] N. Goyal, G. Kindler, and M. Saks, "Lower bounds for the noisy broadcast problem," *SIAM J. Comput.,* pp. 1806–1841, 2008.

[37] A. Orlitsky and J. R. Roche, "Coding for computing," *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 903–917, 2001.

[38] R. J. McEliece, *The Theory of Information and Coding.* Reading, MA: Addison-Wesley, 1977.

[39] N. J. Sloan and A. D. Wyner, *Claude Shannon: Collected papers.* New Jersey: IEEE Press, 1993.

Workshop report (Øyvind Ytrehus):

The International Workshop on Coding and Cryptography,

April 15–19, 2013, Bergen, Norway

WCC is a series of biennial workshops on coding theory, cryptography, and discrete mathematics, alternating between Paris and Bergen. This year's edition took place in Bergen, with 117 participants from approximately 36 countries. A COST workshop on network coding was integrated in the last two days of the WCC program. Pre-proceedings from the workshop, as well as other related information, can be found in the WCC web page at http://www.selmer.uib.no/WCC2013/ .

The WCC organizers would like to thank all participants for making the WCC an exciting, enjoyable, and pleasant event.

An Afternoon in Istanbul

Elza Erkip,
Polytechnique Institute of NYU

Istanbul is a world-class city; rich with history, endowed with great geography and a medley of people and cultures. One cannot do justice with a one-page article to Istanbul's great sights. So instead of giving a rundown of Istanbul's classic tourist sights, which you can find in any respectable guidebook, I will tell you about my favorite spot in Istanbul. It is a stretch of the Bosphorus, on the European side, from Arnavutkoy to Bebek. You can walk by the sea, watch people fish, and occasionally swim, see big tankers and small boats navigate this magical waterway.

Arnavutkoy, which means "Albanian village" in Turkish, has great fish restaurants near the water and more affordable eateries on the street that extends inland from the police station. A favorite Turkish pastime is eating ice cream or sunflower seeds as you stroll by the water. You can walk to Bebek in less than an hour even with a leisurely stroll. Bebek means "baby", which is a shortened version of "eye baby" which in turn means "pupil" as in "pupil of Bosphorus," a reference to its spectacular location. This is the place to sit by the water and have a glass of Turkish tea in the park by the mosque or one of the cafes that are lined by the Bosphorus.

If you have time, a pleasant option to extend your trip on the Bosphorus is to take a ferry from Arnavutkoy to Kanlica, a small port on the Anatolian side of the Bosphorus known for its yogurt (not the frozen variety, but the real thing). Or you can stay on the European side and go to Ortakoy, a cluster of restaurants, cafes and shops near the "first bridge" as locals call the southern suspension bridge that connects the two continents.

If you can spend a full day on the Bosphorus, check out the cruise operated by the city ferries: http://www.sehirhatlari.com.tr/en/timetable/full-bosphorus-cruise-362.html. For 25 YTL (~$15) you can go up all the way to Anadolu Kavagi, right where the Bosphorus meets the Black Sea, climb up the old Yoros castle for breathtaking views, have a seafood lunch by the sea and ride the ferry back. You can get off at Kanlica on the way back for yogurt or at Ortakoy for some shopping. It is easier to catch the ferry from Besiktas if you are leaving from Taksim area where the conference center is located.

GOLOMB'S PUZZLE COLUMN™

Word Puzzles

Solomon W. Golomb

For a change of pace, this column will be more verbal than mathematical.

1) Hidden "bits".

List all the English words you can think of that contain the consecutive letters b-i-t. (They can be at the beginning, the middle, or the end of a word.) List only one example for each "root meaning." That is, don't add prefixes like de-, re-, un-, or suffixes like –s, -ing, -(t)ion unless they change the basic root meaning.

2) Hidden "numbers".

a) I like to hide the names of numbers in longer words that preceed and follow the number name by 2 letters at each end. For example, TEN can be concealed in ANTENNA. Find such concealments for the names of the numbers 1, 2, 8, 9, and 10. (See if you can find three different concealing words for each of the numbers.)

b) For the numbers 3, 4, 5, 6, 7, 30, 40, 80, and 100, find longer words that contain their names, with the letters of the number name in the correct sequential order, but with other letters which may preceed, follow, and/or be interspersed. For example, FIVE can be found in FestIVE. More than one concealing word for each of these numbers is encouraged.

c) I can find ZERO in a South American coin, hidden in the pattern _ _ _ ZE_RO. Also, the Russian word for "Lake" is OZERO (almost a "double zero"). Can you find an English word, of at least 5 letters, in which the four letters ZERO occur consecutively, without any other letters being inserted?

d) What English words contain the consecutive letters _ _ _ ONETWO_ _ _? What Comic Book females are named _ _ TWO_ _ _? The common misspelling "RELEVENT" contains "ELEVEN". Can you coin a reasonable word to form _ELEVEN _ _ _ _?

GOLOMB'S PUZZLE COLUMN™

Using Roots of Unity Solutions

Solomon W. Golomb

1) We put the n vertices of a regular n-gon on the n*th* roots of unity in the complex plane, and rotate it to have the point P on the positive real axis. Then P is at the point $z = d + i\cdot 0 = d$, and the vertices are at $e^{2\pi ik/n}$ for $k = 0, 1, 2, \ldots, n-1$.

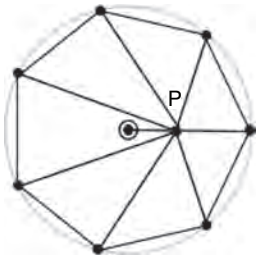

Thus the product of the distances from P to each of the n vertices is given by $\prod_{k=0}^{n-1}|z - e^{2\pi ik/n}| = \left|\prod_{k=0}^{n-1}(z - e^{2\pi ik/n})\right| = |z^n - 1| = 1 - d^n$, since $0 \le d \le 1$.

2) We place the $a \times b \times c$ box in the first octant of xyz-space, and use the usual coordinates (x, y, z), with $1 \le x \le a, 1 \le y \le b, 1 \le z \le c$, for its cells. We set $\eta = e^{2\pi i/n}$, a complex root of $z^n - 1 = 0$, and assign the value of η^{x+y+z} to the cell at (x, y, z). Over any $1 \times 1 \times n$ brick filling the box, the sum of the values of the cells in the brick will be 0, since two of x, y, z will be constant, while the third will take each of the values from 0 to $n-1$, modulo n, in some order, so that over the brick the values will sum to $\eta^0 + \eta^1 + \eta^2 + \cdots + \eta^{n-1} = 0$. If the box is fully packed with bricks, the sum of the values over all the cells in the box will be the sum over all the bricks, and therefore also 0. However this sum is $\sum_{x=1}^{a}\sum_{y=1}^{b}\sum_{z=1}^{c}\eta^{x+y+z} = \sum_{x=1}^{a}\eta^x\sum_{y=1}^{b}\eta^y\sum_{z=1}^{c}\eta^z = ABC = 0$, where A, B, and C are complex numbers, and in order for their product to be 0, one of the factors must be 0. However, $\eta^1 + \eta^2 + \cdots + \eta^t = 0$ if and only if t is a multiple of n; so for A, B, or C to be equal to 0, (at least) one of a, b, or c must be a multiple of n.

3) Corresponding to each positive integer n, we use z^n, and for the set of all positive integers we have $z + z^2 + \cdots + z^n + \cdots = z/(1-z)$. For each arithmetic progression $P_i = \{a_ik + b_i\}$, we use the smallest b_i that makes $a_ik + b_i$ positive at $k = 1$, and associate it with $\sum_{k=1}^{\infty} z^{a_ik+b_i} = z^{b_i}\sum_{k=1}^{\infty} z^{a_ik} = \frac{z^{a_i+b_i}}{1-z^{a_i}}$. If the r progressions exactly cover the positive integers, we will have $\frac{z}{1-z} = \sum_{i=1}^{r}\frac{z^{a_i+b_i}}{1-z^{a_i}}$, valid for $|z| < 1$. The only singularity of $z/(1-z)$ occurs as $z \to 1^-$. However, $(z^{a_i+b_i}/(1-z^{a_i}))$ approaches infinity as z goes to each complex a_i*th* root of unity. If $a_0 = \max_{1\le i\le r} a_i$, and if there is only one arithmetic progression of the form $a_0k + b_0$ (any b_0) among the progressions P_i, then as z approaches *any* complex a_0*th* root of unity, the right side of the above equation will increase in magnitude without limit, while $|z/(1-z)|$ will remain bounded, a contradiction.

4) For prime p, we wish to represent the p^2 numbers $0, 1, 2, \ldots, p^2 - 1$ as sums $s + t$, with $s \in S$ and $t \in T$, where S and T each contain p non-negative integers. This requires $1 + z + z^2 + \cdots + z^{p^2-1} = S(z)T(z)$, where $S(z) = \sum_{s\in S} z^s$ and $T(z) = \sum_{t\in T} z^t$. We know that $z^m - 1 = \prod_{d|m}\Phi d(z)$, where $\Phi_d(z)$ is the "cyclotomic polynomial" whose roots are the $\phi(d)$ primitive d*th* roots of unity, and the product is taken over all positive divisors d of m. With $m = p^2$, the only positive divisors are $1, p$, and p^2; so $1 - z^{p^2} \equiv (1-z)(1 + z + z^2 + \cdots + z^{p^2-1}) = \Phi_1(z)\Phi_p(z)\Phi_{p^2}(z)$, where $\Phi_1(z) = 1 - z, \Phi_p(z) = 1 + z + z^2 + \cdots + z^{p^2-1}$ and $\Phi_{p^2}(z) = 1 + z^p + z^{2p} + \cdots + z^{(p-1)p}$. Dividing by $1 - z = \Phi_1(z)$, we see that $S(z)T(z) = \Phi_p(z)\Phi_{p^2}(z)$, and since all cyclotomic polynomials are irreducible, we must have, as S and T, the sets of exponents of $\Phi_p(z)$ and $\Phi_{p^2}(z)$, namely $\{0, 1, 2, \ldots, p-1\}$ and $\{0, p, 2p, \ldots, (p-1)p\}$, in either order.

Notes. With $m^2 = 16$ (for non-prime $m = 4$), from $1 + z + z^2 + \cdots + z^{15} = (1 - z^{16})/(1-z) = \Phi_2(z)\Phi_4(z)\Phi_8(z)\Phi_{16}(z) = (1+z)(1+z^2)(1+z^4)(1+z^8)$, there are three ways to partition these four irreducible binomials into two products of two factors each, to get three different solutions for S and T. These are: $\{0, 1, 2, 3\} \oplus \{0, 4, 8, 12\}$; $\{0, 1, 4, 5\} \oplus \{0, 2, 8, 10\}$; and $\{0, 1, 8, 9\} \oplus \{0, 2, 4, 6\}$. See what you can find for $m^2 = 100$, where 100 has 9 positive divisors. (One of the "non-standard" representations $S + T$ has $S = \{0, 1, 2, 3, 4, 50, 51, 52, 53, 54\}$ and $T = \{0, 5, 10, 15, 20, 25, 30, 35, 40, 45\}$.)

In the Blogosphere...

In consultation with various IT society "heavy weights," I decided to add a column in which I include pointers to some interesting blog items around. The items, for now, are essentially indication of my personal taste and limited time but I hope folks will send in their own pointers and their suggested blog posts to add diversity.

Information Theory Blog

An Extremal Conjecture: Experimenting with Online Collaboration

Posted on March 5, 2013 by Thomas Courtade

I have an extremal conjecture that I have been working on intermittently with some colleagues (including Jiantao Jiao, Tsachy Weissman, Chandra Nair, and Kartik Venkat). Despite our efforts, we have not been able to prove it. Hence, I thought I would experiment with online collaboration by offering it to the broader IT community.

In order to make things interesting, we are offering a $1000 prize for the first correct proof of the conjecture, or a $250 award for the first counter example! Feel free to post your thoughts in the public comments. You can also email me if you have questions or want to bounce some ideas off me.

Although I have no way of enforcing them, please abide by the following ground rules:

1) If you decide to work on this conjecture, please send me an email to let me know that you are doing so. As part of this experiment with online collaboration, I want to gauge how many people become involved at various degrees.

2) If you solve the conjecture or make significant progress, please keep me informed.

3) If you repost this conjecture, or publish any results, please cite this blog post appropriately.

One final disclaimer: this post is meant to be a brief introduction to the conjecture, with a few partial results to get the conversation started; it is not an exhaustive account of the approaches we have tried.

1. The Conjecture

Conjecture 1. *Suppose X, Y are jointly Gaussian, each with unit variance and correlation ρ. Then, for any U, V satisfying $U - X - Y - V$, the following inequality holds:*

$$2^{-2I(Y;U)}\,2^{-2I(X;V|U)} \geq (1-\rho^2) + \rho^2 2^{-2I(X;U)}2^{-2I}(Y;V|U). \qquad (1)$$

An Ergodic Walk

C.R. Rao and Information Geometry

Posted on April 13, 2013 by Anand Sarwate

On Lalitha's recommendation I read Frank Nielsen's paper "Cramer-Rao Lower Bound and Information Geometry," which is a survey how C.R. Rao's work has impacted information geometry. I remember spending some time in grad school trying to learn information geometry (mostly for fun), but since it ended up not being particularly useful in my research, I'm afraid a lot of it has leaked out of my ears. This paper has a short introduction to the Cramer-Rao lower bound and an introduction to information geometry which might be a nice read for some of the readers of this blog. It's certainly faster than trying to read Amari's monograph! In particular, it goes over the "highlights" of geodesics and other geometric features on the manifold of probability distributions.

The paper mentions the sub-family of f-divergences known as α-divergences, which are given by

$$D_\alpha(p\|q) = \frac{4}{1-\alpha^2}\left(1 - \int p(x)^{(1-\alpha)/2} q(x)^{(1+\alpha)/2} dx\right)$$

The KL divergence is $D_{-1}(p\|q)$—you have to take the limit as $\alpha \to -1$. Within this family of divergences we have the relation $D_\alpha(p\|q) = D_{-\alpha}(q\|p)$. Consider a pair of random variables (X, Y) with joint distribution P_{XY} and marginal distributions P_X and P_Y. If we take $q = P_X P_Y$ and $p = P_{XY}$ then the mutual information is $D_{-1}(p\|q)$. But we can also take

$$D_{-1}(P_X P_Y \,\|\, P_{XY}) = D_1(P_{XY} \,\|\, P_X P_Y)$$

Thus it turns out that the "lautum information" defined by Palomar and Verdú is a special case of this: it's the 1- divergence between the the joint distribution and the product of the marginals. While their paper mentions the lautum information is an f-divergence, it doesn't discuss this connection to this family of divergences. Nielsen's paper calls this the "reverse Kullback-Leibler divergence," but some googling doesn't seem to indicate that this is a common term, or indeed if it has some use in information geometry. Palomar and Verdú give several operational interpretations of the lautum information.

ANNOUNCING

IEEE ITSoC Student Committee Video Contest

Step 1: Make an < 10 min video on any topic related to Information Theory

Step 2: Post on youtube and contest page itsoc-competition.ece.utexas.edu

Step 3: Join in for the video-award ceremony at ISIT

Over $1000 in prizes + fame & pride!

Competition ends July 1st 2013
Contact sriram@ece.utexas.edu for details

Many will enter, few will win

Information Theory Workshop

http://itw2013.tsc.uc3m.es

ITW 2013 will be held on the campus of the University of Seville. Technical contributions are solicited in all areas of Information Theory with special emphasis on innovative and interdisciplinary research related to:

- Bioinformatics
- Coding
- Communication
- Compression
- Machine Learning
- Networks
- Security
- Signal Processing
- Spectrum Sharing
- Statistics

Dates to remember

Notification:

July 12th, 2013

Early Registration:

July 26th, 2013

Plenary Speakers

Ernest Fraenkel

László Gyorfi

Michael Kearns

Sergio Verdú

Martin Wainwright

Venkat Anantharam

Gabor Lugosi

Shuki Bruck

Amin Shokrollahi

Suhas Diggavi

Rudiger Urbanke

Amos Lapidoth

Naftali Tishby

Technical Chairs

Albert Guillén i Fàbregas

Michael Honig

Alon Orlitsky

Organizing Committee

Pedro Crespo

Alfonso Martínez

Juan J. Murillo-Fuentes

Javier Payán

Fernando Pérez-Cruz

Matilde Sánchez-Fernández

FIFTY-FIRST ANNUAL ALLERTON CONFERENCE

ON COMMUNICATION, CONTROL, AND COMPUTING

October 2–4, 2013
Call for Papers

The Fifty-First Annual Allerton Conference on Communication, Control, and Computing will be held from Wednesday, October 2 through Friday, October 4, 2013, at Allerton House, the conference center of the University of Illinois. Allerton House is located twenty-six miles southwest of the Urbana-Champaign campus of the University in a wooded area on the Sangamon River. It is part of the fifteen-hundred acre Robert Allerton Park, a complex of natural and man-made beauty designated as a National natural landmark. Allerton Park has twenty miles of well-maintained trails and a living gallery of formal gardens, studded with sculptures collected from around the world.

Papers presenting original research are solicited in the areas of communication systems, communication and computer networks, detection and estimation theory, information theory, error control coding, source coding and data compression, network algorithms, control systems, robust and nonlinear control, adaptive control, optimization, dynamic games, multi-agent systems, large-scale systems, robotics and automation, manufacturing systems, discrete event systems, multivariable control, computer vision-based control, learning theory, cyber-physical systems, security and resilience in networks, VLSI architectures for communications and signal processing, and intelligent transportation systems.

Information for authors: Regular papers suitable for presentation in twenty minutes are solicited. Regular papers will be published in full (subject to a maximum length of eight 8.5″ x 11″ pages, in two column format) in the Conference Proceedings. Only papers that are actually presented at the conference can be included in the proceedings, which will be available after the conference on IEEE Xplore.

For reviewing purposes of papers, a title and a five to ten page extended abstract, including references and sufficient detail to permit careful reviewing, are required.

Manuscripts must be submitted by **Tuesday, July 9, 2013**, following the instructions at the Conference website: http://www.csl.uiuc.edu/allerton/**.**

Authors will be notified of acceptance via e-mail by August 7, 2013, at which time they will also be sent detailed instructions for the preparation of their papers for the Proceedings.

Final versions of papers to be presented at the conference will need to be submitted electronically by October 6, 2013.

Conference Co-Chairs: Tamer Ba ar and Olgica Milenkovic
Email: *allerton-conf@illinois.edu* URL: www.csl.illinois.edu/allerton/

COORDINATED SCIENCE LABORATORY AND THE DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

University of Illinois at Urbana-Champaign

Call for Papers

2014 International Zurich Seminar on Communications

February 26 - 28, 2014

The 2014 International Zurich Seminar on Communications will be held at the Hotel Zürichberg in Zurich, Switzerland, from Wednesday, February 26, through Friday, February 28, 2014.

High-quality original contributions of both applied and theoretical nature are solicited in the areas of:

- Wireless Communications
- Information Theory
- Coding Theory and its Applications
- Detection and Estimation
- MIMO Communications
- Optical Communications
- 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 4 pages) of sufficient detail to permit reasonable evaluation. Authors of accepted papers will be asked to produce a manuscript not exceeding 4 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 22, 2013**.

Additional information will be posted at

`http://www.izs.ethz.ch/`

We look forward to seeing you at IZS.

Amos Lapidoth and Helmut Bölcskei, Co-Chairs.

C A L L f o r P A P E R S

BlackSeaCom 2013

Technically Sponsored by IEEE Communications Society

3-5 July 2013




First International Black Sea Conference on Communications and Networking
Sheraton Hotel, Batumi, Georgia (http://www.blackseacom.net)

As the name suggests, the BlackSeaCom series of conferences will be held in the countries surrounding the Black Sea. The goal of BlackSeaCom is to bring together visionaries in academia, research labs and industry from all over the world to the shores of the Black Sea. Here they will address many of the outstanding grand challenges that exist in the areas of communication and networking while having an opportunity to explore this exciting and dynamic region that has a rich history.

The first edition of the conference will take place at the Sheraton Hotel on July 3, 4 and 5, 2013 in Batumi, Georgia--a beautiful major port city on the Eurasian Crossroad. The theme of the conference, **Communication and Networking in the Land of the Golden Fleece**, is fitting for the chosen venue (Georgia) to set the goal of discovery and value in our field paralleling the ancient myth of the "Golden Fleece".

We seek original completed and unpublished work not currently under review by any other journal/magazine/conference. Topics of interest include, but are not limited to:

- Information theoretic analysis of wireless networks
- Communication protocols (transport, routing, link, and physical layers)
- Cross-layer design and optimization
- Theoretical analysis frameworks
- Resource management
- Energy efficiency, resiliency, reliability, and robustness
- Deployment scenarios and experiences
- Infrastructures, Platforms, test beds, and software
- Wireless networks, Ad hoc and mesh networks, Wireless sensor networks
- Cognitive radio networks
- Communication in challenging environments (underwater, underground, tunnels/mines, space, disasters)
- Internet of things
- Nano-scale networks
- Integration of heterogeneous networks

Submission Guidelines: Papers should describe original work and should be no more than 8 pages in the IEEE double column proceedings format including tables, figures and references. Note that accepted papers up to 6 pages will be published in the proceedings as well as in **IEEE Xplore** with no additional charge. Exceeding pages will be charged an additional fee.

To submit a paper, please use the instructions at: http://www.blackseacom.net

IEEE reserves the right to exclude a paper from distribution after the conference (e.g. removal from **IEEE Xplore**) if the paper is not presented at the conference.

Important dates:

Paper submission deadline:	**15 March 2013**
Notification of acceptance:	**1 May 2013**
Camera-ready papers due:	**1 June 2013**

Honorary Co-Chairs:
Giuli Alasania, Vice Rector, International Black Sea University, Tiblisi, Georgia
Tayfun Acarer, Chairman, Information and Communication Technologies Authority, Turkey

General Co-Chairs:
Ian F. Akyildiz, Georgia Institute of Technology, USA
Mehmet Ulema, Manhattan College, USA

Technical Program Co-Chairs:
Anthony Ephremides, University of Maryland at College Park, USA
Eylem Ekici, Ohio State University, USA

Publicity Co-Chairs:
Rao uf Boutaba, University of Waterloo, Canada
Dario Pompili, Rutgers University, USA

Publication Chair:
Tommaso Melodia, State University of New York- Buffalo, USA

Local Arrangement Chair:
Lasha Ephremidze, University of Tbilisi, Georgia

Finance Chair:
M. Can Vuran, University of Nebraska-Lincoln, USA

Web Chair:
Josep M. Jornet, Georgia Institute of Technology, USA

Regional Advisors:
Cabir Erguven, International Black Sea University, Georgia
Yevgeni Koucheryavy, St. Petersburg State University of Telecomm., Russia
Tuna Tugcu, Bogazici University, Turkey
Yuriy Prokopenko, National Technical University of Ukraine,Ukraine

Steering Committee:
Ian F. Akyildiz, Georgia Institute of Technology, USA
Anthony Ephremides, University of Maryland at College Park, USA
Mehmet Ulema, Manhattan College, USA

Sponsored by International Black Sea University, Georgia

ICITS 2013

7th International Conference on Information Theoretic Security

Singapore, November 28–30, 2013
`http://www.spms.ntu.edu.sg/mas/conference/icits2013/`

Call for Papers

This is the seventh in a series of conferences that aims to bring together the leading researchers in the areas of information theory, quantum information theory, and cryptology. Papers on all technical aspects of information-theoretic security and quantum information-theoretic security are solicited for submission to ICITS 2013. Areas of interest include, but are not restricted to:

Unconditional security	Lattices and cryptography	Quantum information theory
Quantum cryptography	Secret sharing	Network coding security
Authentication codes	Multiparty Computation	Physical models & assumptions
Wiretap channels	Bounded storage model	Physical layer security
Randomness extraction	Oblivious transfer	
Codes and cryptography	Nonlocality and nonsignaling	

Two types of contributed presentations will take place in ICITS 2013. The *Conference Track* will act as a traditional conference, consisting of original papers with published proceedings in the Lecture Notes in Computer Science series. The *Workshop Track* will operate more like an informal workhsop, with papers that have appeared elsewhere or that consist of work in progress.

Important Dates

Conference Track submissions deadline	**Friday, July 5, 2013, 13.00 GMT**
Conference Track notification	Friday, August 30, 2013
Proceedings version	Friday, September 20, 2013
Workshop Track submissions Deadline	**Friday, August 2, 2013, 13.00 GMT**
Workshop Track notification	Thursday, September 19, 2013

Conference Organization

General Chairs	Frédérique Oggier (*NTU, Singapore*) and Miklos Santha (*CQT, Singapore*)
Program Chair	Carles Padró (*NTU, Singapore*)

The 2013 School of Information Theory is organized by **Center for Science of Information (http://soihub.org),** a National Science Foundation science and technology center, and is sponsored by the IEEE Information Theory Society. Hosted at Purdue University from Tuesday, June 4 to Friday, June 7, 2013, the school provides a venue where doctoral and postdoctoral students can meet to learn from distinguished professors in information theory, and form friendships and collaborations. This year the school will introduce several interdisciplinary topics in the emerging field of science of information. Students will present their own research via a poster during the school. Although the focus is on information theory, interdisciplinary topics are welcome, e.g., topics related to mathematics, physics, biology, control, networking, etc.

Important Dates:
Applications: April 1, 2013
Acceptance Decisions: April 15, 2013
Registration: May 1, 2013

Program Overview:

Mornings: Lectures by invited speakers, TBA
Afternoons: Presentations and posters by students
Evening: Special events/activities

Organizing Committee:

- General Chair: Wojciech Szpankowski (Purdue University)
- Andrea Goldsmith (Stanford University)
- Sergio Verdu (Princeton University)
- Deepak Kumar (Bryn Mawr College)
- Olgica Milenkovic (University of Illinois)
- Todd P. Coleman (UC San Diego)
- Mark D. Ward (Purdue University)
- Brent Ladd (Purdue University)
- Barbara Gibson (Purdue University)
- Bob Brown (Purdue University)

Advisor:

- Gerhard Kramer (Technical University of Munich)

For updates, application, and further details: http://www.itsoc.org/north-american-school-2013/

Conference Calendar

DATE	CONFERENCE	LOCATION	WEB PAGE	DUE DATE
April 14–19, 2013	**32nd IEEE International Conference on Computer Communications (INFOCOM 2013)**	Turin, Italy	http://infocom.di.unimi.it/	Passed
April 15–19, 2013	**International Workshop on Coding and Cryptography (WCC 2013)**	Bergen, Norway	http://www.selmer.uib.no/WCC2013/	Passed
April 22–26, 2013	**2013 IEEE European School on Information Theory (ESIT 2013)**	Ohrid, Republic of Macedonia	http://www.itsoc.org/european-school-2013	Passed
May 8–9, 2013	**2013 Iran Workshop on Communication and Information Theory (IWCIT)**	Tehran, Iran	www.IWCIT.org	Passed
May 13–17, 2013	**WiOpt 2013**	Tsukuba Science City, Japan	http://www.wi-opt.org/	Passed
June 2–5, 2013	**2013 77th Vehicular Technology Conference (VTC2013-Spring)**	Dresden, Germany	http://www.ieeevtc.org/vtc2013spring/	Passed
June 4–7, 2013	**2013 IEEE North American School of Information Theory**	West Lafayette, Indiana, USA	http://www.itsoc.org/north-american-school-2013/	Passed
June 9–13, 2013	**IEEE International Conference on Communications (ICC 2013)**	Budapest, Hungary	http://www.ieee-icc.org/	Passed
June 23–26, 2013	**2013 IEEE Communication Theory Workshop (CTW 2013)**	Phuket, Thailand	http://www.ieee-ctw.org/	Passed
July 3–5, 2013	**1st International Black Sea Conference on Communications and Networking (BlackSeaCom 2013)**	Batumi, Georgia	http://www.blackseacom.net/	Passed
July 7–12, 2013	**2013 IEEE International Symposium on Information Theory (ISIT 2013)**	Istanbul, Turkey	http://www.isit2013.org/	Passed
September 9–13, 2013	**2013 IEEE Information Theory Workshop (ITW 2013)**	Seville, Spain	http://itw2013.tsc.uc3m.es/	Passed
October 2–4, 2013	**51st Annual Allerton Conference on Communication, Control, and Computing**	Monticello, Illinois, USA	http://www.csl.illinois.edu/allerton/	July 9, 2013
November 3–6, 2013	**Asilomar Conference on Signals, Systems, and Computers (ASILOMAR 2013)**	Pacific Grove, CA, USA	http://www.asilomarssc.org/	May 1, 2013
December 9–13, 2013	**2013 IEEE Global Communications Conference (GLOBECOM 2013)**	Atlanta, GA, USA	http://www.ieee-globecom.org/	Passed

Major COMSOC conferences: http://www.comsoc.org/confs/index.html