Panel on "New Perspectives on Information Theory" IEEE Information Theory Workshop, Paraty, October 20, 2011

Venkat Anantharam (Berkeley)
Giuseppe Caire (USC)
Max Costa (Campinas)
Gerhard Kramer (Technical University Munich)
Raymond Yeung (Chinese University, Hong Kong)
Sergio Verdú, Moderator (Princeton)

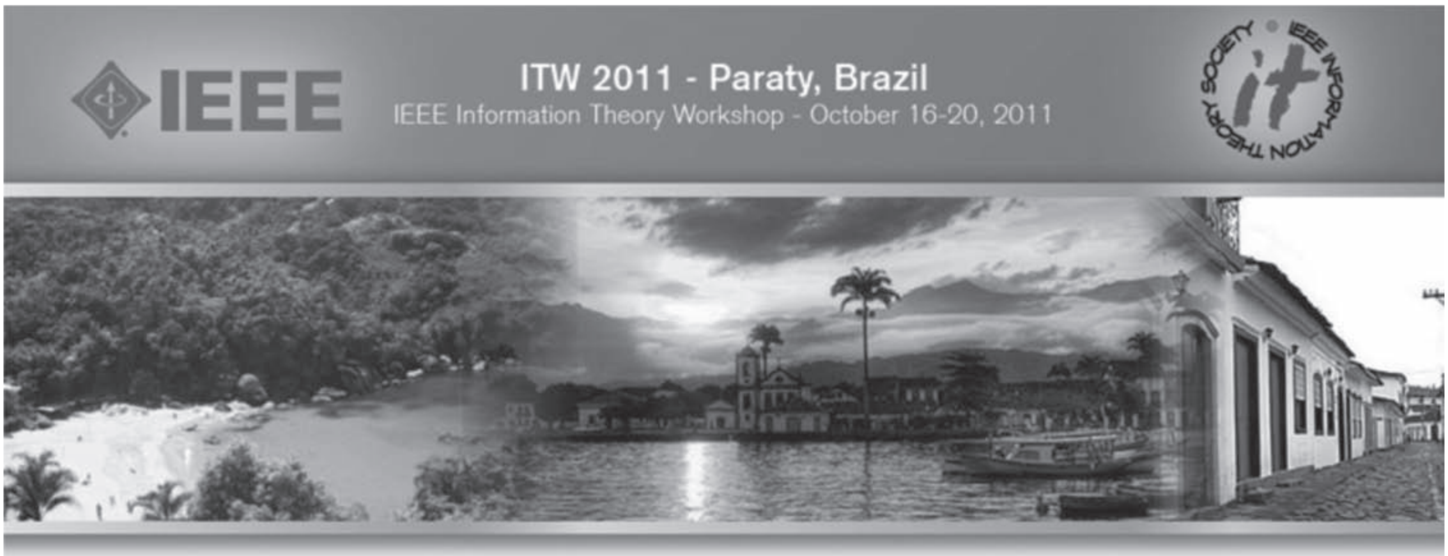

A video of the panel discussion (including questions from the audience) can be found in http://media.itsoc.org/itw2011/ What follows is an edited transcript.

Verdú *Seventeen years ago at the ITW that was held in Moscow, I organized a similar panel on the future of Information Theory with the participation of Dick Blahut, Imre Csiszár, Dave Forney, Prakash Narayan and Mark Pinsker. In preparation for this panel I have asked our panelists to read the transcript of that panel (published in the December 1994 issue of this newsletter) and discuss the ways in which that panel's predictions were and were not accurate.*

Costa Well, it´s been said that it is difficult to make predictions, specially about the future. The 1994 panel predictions were good in many aspects, but they could not guess those areas that appeared from nowhere and brought completely new tools and perspectives to the field. There was another situation in which this happened. Estill Green, a VP of Bell Labs, also made some bold and courageous predictions on telecommunications, looking from 1961 into that technology in the year 2012. Bob Lucky commented on those forecasts in 1999, and pointed out the areas in which the forecasts came close to what was happening, and others that were far out. Forecasts usually estimate the increase or reduction of some variables based on the anticipated development of certain known technologies. They are less precise when the actual changes are produced by a complete switch of paradigm, or a totally brand new technology that takes over the field. It was like that that Green's predictions had problems by excluding the changes produced by optical fiber communications and by the advent of the internet. Yet, Green´s predictions were not completely off because he was counting on a very fast growth of video telephony. Bell Labs was investing heavily in video telephony at that time and of course we all know that that project did not go as expected. The idea that an image is worth a thousand words also goes for the rates, and Green was expecting a much greater traffic demand than actually happened from video, but an even greater traffic actually happened for other reasons, like the internet. So it is very difficult to make predictions on a very long stretch of time. Even attempts to update the course of Green's predictions for 2012 only 12 years ago estimated 1 Tbps as the capacity limit of an optical fiber, and we are already seeing lab transmissions of 100 Tbps in a single fiber and 26 Tbps from a single laser source.

Anantharam: There were some really striking omissions. For one thing, there was no mention at all of LDPC codes or turbo codes, which I think has really been one of the defining features of research in our subject for the past decade or so.

Verdú: Actually, the turbo codes had been just presented at ICC, just about a year before, and in fact very few people believed that that paper was correct. And the LDPCs in Gallager's thesis had not been rediscovered yet.

Anantharam: Yes, so that speaks to the point about the danger of predicting the future; another thing to note is that network coding had not yet been invented at the time. From the point of view of someone who is more interested in theoretical problems, one of the striking things was that some of the problems that were posed as being the major open problems in the field, such as the capacity regions of broadcast channels, relay channels, interference channels–they are all basically still open. Of course, we know a lot more about them and understand them in approximate senses. Also there was a dynamic having to do with what one really means by Information Theory. For instance, Csiszár was one of the panelists and he presented himself as viewing Information Theory more as a mathematical science than a science that is focused entirely on applications. Now, that's not exactly where I stand. I think of Information Theory as a science of understanding how to extract what is relevant about interactions between distributed entities in order for them to coordinate in some way. Communication theory to me is just one aspect of that and in fact much of the discussion in our community is focused on the notion of information in communication theory, but since this panel is meant to look forward, I would venture to say that we are likely to move in directions where the notion of information is coming from fields that are not so close to communication theory, for instance biology, as a prominent example, or neuroscience. So if one is venturing a prediction over a time scale as large as that between 1994 and 2011, I would say that if you are looking at this field 15 to 20 years from now it might not be as communication theory centric as it is now.

Yeung In the 1994 transcript we read that Pinsker and Csiszár were trying to explore the line whether Information Theory is just applied mathematics but also some kind of pure mathematics. My own point of view is that Information Theory is both applied mathematics and pure mathematics, because there is some pure mathematics content in it. In the 1994 transcript, Pinsker brought up the application of Information Theory to ergodic theory and that was a very major impact. Csiszár talked about the application of Information Theory to statistics. This is more like a specific topic as far as I am concerned. But nevertheless, it shows that Information Theory has some pure mathematics elements in it, which is always my belief. In fact I would say that it is this particular reason that has been keeping me very interested in Information Theory. And along this line, during the last 15 years or so, the work by a few others and also by myself on the study of entropy functions has established a link with a few other branches of mathematics. Now we know that there is a common structure between entropy functions, group theory, network coding, and also Kolmogorov complexity. The results along this line also have implications in conditional independence in probability theory. Also, most recently we are interested in how we can use differential entropy functions to study inequalities regarding positive definite matrices. So definitely, I believe that there are some pure mathematics elements in Information Theory. Another thing is that Information Theory is a very special kind of applied mathematics if you want to think of it this way. In many branches of applied mathematics, researchers in the field are users of the results from pure mathematics. You can think of Information Theory being a branch of applied probability, but it is also very different from probability theory, because if you talk to a probability theorist, most likely he or she would know very little about Information Theory. That's why very often in Information Theory we have to develop the tools from scratch. That's something rather unique about Information Theory. Another thing which I find very unique about Information Theory is that for other branches of applied mathematics, we very often use mathematics to model a physical system, for example, control theory, signal analysis, mechanics, and other things. But in Information Theory we are using mathematics to model information, so there is one layer of abstraction, but information by itself is also something rather abstract, instead of being a physical system. So in Information Theory there are two layers of abstraction, which makes the subject very subtle and interesting as far as I am concerned. And therefore after so many years, we still from time to time find that there are some basic things which are not well understood and even not well defined, and that's has been keeping me very interested in the subject.

A few years ago I met a student of Kolmogorov (Albert Shiryaev) who was visiting our university in Hong Kong. When he knew that I was working on Information Theory, he was very interested in talking to me. In fact, Alon Orlitsky was also visiting us at that moment. I told him that I am not a mathematician by training, and so for the mathematics I use all the time I know them quite well, but there are a good number of branches of mathematics that I know nothing about. He said, "That's okay. Shannon was an engineer, too." But he actually was not entirely correct because Shannon got his PhD in mathematics. Then he went on to talk about how Kolmogorov thought about Shannon. I knew that Kolmogorov had worked on Information Theory, but I did not know that Kolmogorov had such a high regard of Shannon. In fact, when Shannon's work was translated to Russian, somehow the section on entropy power inequality was missing. Subsequently, Kolmogorov found out that was the case. He was very upset and said, "How come they could have missed this very important result by Shannon!"

A few years ago (I think it was two years ago) I visited Qualcomm. There I met a Russian guy and we talked about Shannon and Kolmogorov, and then he told me, "You guys in the West compiled a collection of papers by Shannon around 2000. In Russia we had it already in 1963!" Then he rushed to his own office, brought a copy of that collection, and showed it to me. It was all Russian. There I found a word which I thought meant "Shannon." I said, "Is this Shannon's name in Russian?" He said, "Yes." Then he showed me the preface which was written by Kolmogorov. He said, "Basically Kolmogorov said that Shannon had tremendous engineering insight, although he did not really prove anything." That's Kolmogorov's point of view.

Caire What is striking to me while reading the 1994 panel transcripts is the absence of the two keywords that have dominated the last 20 years of communication engineering: internet and wireless. The discussion in 1994 was mainly focused on issues like mathematics vs. engineering, or Shannon theory vs. coding theory. In contrast, the magic word "network" was almost entirely missing. As a matter of fact, network Information Theory has proven to be a formidable area of research both in terms of its richness (the abundance of open problems is almost a life insurance for information theorists . . . . we will never be out of jobs) and in terms of its impact on technology, especially in wireless (3G came in 1996, 4G is happening now, 802.11n with space-time coding, spatial multiplexing happened a few years ago, and the next big thing is to handle interference and multihop relaying). In the wired domain, of course, network coding and its variant application to distributed storage systems (essential for cloud computing) is gaining a lot of traction too.

Kramer I enjoyed the back-and-forth between the engineers (Forney, Blahut), mathematicians (Csiszár, Pinsker), and those perhaps in between (Verdú, Narayan). My apologies if I am categorizing people too much, but our panel does not seem to have the same span of folks as then.

Concerning succeeding and erring on future trends, a few things caught my eye: In contrast to what Giuseppe said, networks were mentioned by several people as an important future topic. Blahut correctly predicted the importance of energy (p. 9) in communications and computing. This is now, 17 years later, a hot topic! Verdú was right in predicting that the probabilistic school is gaining ground on the combinatorial school, in the sense that turbo and LDPC codes dominated coding for many years (with a few exceptions from the then-unknown area of network coding). He also correctly predicted the importance of the interference channel, even if this took a little longer to happen. I think Verdú was too cautious with his statement that "Maybe the day will come when a software package will enable the engineer to closely approach capacity on almost any channel with the technology of the day. Admittedly, I am afraid it is us who will be dead when that day arrives." But about 7 years later, you could download degree distributions for LDPC codes that approach capacity on almost any (practical) channel.

Verdú I am going to stand by the statement I made in 1994. The key is "almost any channel". In fact, even the capacity of many practical single-user channels is still an open problem (e.g. channels with frequency selective and time selective channels, deletion channels). We are nowhere near the point of having an algorithm that given a black box will find not only its capacity but near-optimal codes. Think of the counterpart in data compression, where we do not need to know anything about the source in order to approach its fundamental limits, provided some general technical conditions are satisfied.

In 1994, I did mention the renaissance of physical layer research, but reality ended up exceeding anyone's expectations on the future relevance of Information Theory to the practical world, and in particular the wireless world. The wireless revolution would prove to be a godsend for Information Theory.

Yeung Toby Berger said in the mid-1990's that the rise of wireless is to the advantage of EE (and hence also to IT) because it is something not easily eaten up by CS. He was absolutely right. We have already witnessed the failures of the Google Phone and then Windows Phone.

Verdú I find it interesting that one of the questions I put to the panel in 1994 was "Is Information Theory dead and if not, what evidence do we have to the contrary?" I wouldn't ask that question today: We have come a long way since then and we have now a much higher collective self-esteem and optimism about the future of our field.Why is that?

Kramer One major change since 1992 is the appearance of turbo codes and more generally iterative processing. I like to say that turbo codes made mutual information an engineering quantity. I recall that while I was doing my Master's degree work in 1991, most people believed that that the cut-off rate was the practical limit of reliable communication.

But getting back to the 1994 panel, you had started the discussion with the question "Is Information Theory dead?" and so the focus of the first few pages was of course on this question. I think that the best response was that as long as Information Theory continues to attract the best and brightest young people, Information Theory will continue on as it always has. This response was given by Csiszár (twice! on p. 4, right column, and on p. 10, top left), and by Blahut (p. 11). One can ask what will attract such people, and that will be having difficult and relevant problems (see Forney, p. 9). "Relevant" can mean many things, of course. One thing I especially appreciated from the article was Blahut's statement on p. 5 that it is not necessary to be defensive and negative concerning Information Theory. He wrote that he was living in the "decade of information". And now that decade has stretched on for at least 17 years. I think it will stretch on for many years more. (See Sergio's recent book review on "The Information" in the Sep. issue of the IT Transactions where he laments the focus on 1948.) Blahut was refreshingly positive in his closing statement.

Caire In my opinion, the periodically emerging question about whether IT is dead, or is dying, is essentially due to a complex of inferiority that, at various degrees of intensity, is permeating our community. In fact, I would revert the question and notice that there are other fields, such as, "communications", "signal processing" and "networks", which in the past have been regarded as "eating up" Information Theory, are now getting swallowed by Information Theory.

Today, it is the general understanding of a very broad research and industrial development community that the "Shannon approach" is the winner: extracting the essence of a problem, characterizing its fundamental limits, and designing systems tightly inspired by the optimal or near-optimal strategies stemming from the information theoretic investigation has proven to yield superior results in several areas, and especially in wireless networks. As an example, it is not an accident that systems designed on the basis of engineering common-sense (an infamous example being the 3G CDMA-based system) led to disappointing performance and years of delay in deployment, and that the present generation of networks (LTE, WLANs) is based on OFDM, which is (oversimplifying) what Shannon teaches us to do over a Gaussian frequency selective channel.

If you take a look at 3GPP or IEEE 802.11 standardization forums, it is really striking to notice that what we do percolates almost immediately from theorems to system proposals. In this sense, at least in the realm of wireless communications, Information Theory is the clear winner, "our" approach has become "the approach", and the distance between new discoveries, even the most exotic and probably difficult to implement, such as interference alignment, and the industrial R&D, has become nonexistent. This is also due to Moore's Law: today we can implement on an iPhone 4 algorithms that 20 years ago would have required a powerful mainframe. Such abundant computation power and memory, even in small devices, allows the adoption of Information Theory-driven approaches that in the past were unthinkable.

If a problem exists, right now, is not whether Information Theory is dead or alive, but the lack of recognition that, as a scientific community, we seem to suffer. Our ideas are grabbed very quickly and used by larger and wealthier groups of researchers and system designers, without always giving the correct credit to the originators. It is sufficient to look at conferences such as SIGCOMM, but sometimes also to ICC and Globecomm, to understand what I am talking about.

Costa Before I address the point, let me make a comment on what Raymond just mentioned. It´s true that some of the proofs of what Shannon could see were not there, like the EPI, for example, the entropy power inequality, that was later proved by Stam and Blachman. But he could have the tremendous insight to see it, and maybe feel that it was essentially the isoperimetric inequality. This amazing insight that Shannon had was also what led him to the random coding argument. I remember a class that Tom Cover gave in a course on Information Theory in which he said: "If I had that idea of the random coding argument, I think I would just go home and sober up." Also to point to the recognition of the work of Shannon, he arrived unexpectedly at the 1985 ISIT in Brighton. Nobody knew that he was going to attend the meeting and there was a big commotion. Bob McEliece is reported to have said that it was just like if in a conference of physicists someone had announced that Newton was present.

Now, we see obituaries of Information Theory come up from time to time. Right now is actually a time that we have a better perspective. Coding theory has also gone through that, and obituaries were announced for coding theory a number of times. One of the early announcements was closely followed by the invention of trellis coded modulation and all the burst of activity that it generated. A few years later, another such obituary was challenged by the creation of turbo codes in 1993, and by the rediscovery of LDPC codes. So it is very risky to make categorical statements regarding the end of an area. Many new things are always coming up to second guess the less optimistic forecaster.

It seems odd to imagine that Information Theory may be perishing when we have just witnessed the dawn of the Information Age. To mention changes that are occurring in a number of schools, the traditional engineering denominations are being replaced by names like information engineering, energy engineering, environmental engineering, and so on. These changes point to the importance of paying attention to the resources, and being resourceful is definitively one of the highlights of Information Theory.

Yeung I just want to add something to what Max just said about Shannon's engineering insight. Well, it has been very amazing to me that while Shannon can be sloppy, there is not a single incidence that he is found wrong. Can someone correct me? It's really amazing! I mean in most mathematical subjects if you do your proofs sloppily, you are bound to get some wrong results, but not in the case of Shannon. I really don't understand.

Verdú I don't think "sloppy" is the right word. The Bell Labs Technical Journal in 1948 is not the IT Transactions in 2011. Shannon wrote a very readable paper aiming to reach a very wide audience. He knew exactly what he was doing. In the paper itself he spelled out a few abstract mathematical arguments, although admittedly he did not include the epsilons and deltas in each and every proof.

Kramer I hope you disagree with me but perhaps the field has become more uniform? For example, at ISITs I think there are fewer pure mathematicians and there are fewer people "directly" connected to industry than before.

Verdú The entropy of topics in ISITs has definitely decreased. Perhaps, today fads play a bigger role in people's choice of fields? I definitely miss the mathematicians, people like Paul Shields, Rudi Ahlswede, Janos Korner, Kati Marton and the whole Russian school. On the positive side, many of the engineers at places like Qualcomm, Bell Labs, IBM, HP Labs are PhDs well-versed in Information Theory.

Caire As a matter of fact, there have been years where an enormous number of sessions was focused on just one topic, and referenced only a handful of papers (e.g., at the peak of Trellis Coded Modulation, the topic distribution was essentially a delta function at Ungerboeck). Right now, we have topics such as biology, compressed sensing and matrix reconstruction, wireless networks, network coding, machine learning, and many more.

Verdú Going back to answering my question on the renaissance of Information Theory, I think that much of this renaissance comes from: 1) the great influx of young talent in the last two decades; 2) the successes of information-theory-driven technologies such as: sparse-graph codes, universal data compression, voiceband modems, discrete multitone modulation (DSL), multiuser detection and MIMO (Shannon tells us don't treat digitally modulated signals as thermal noise: exploit their structure), space-time codes, opportunistic signaling, network coding, etc. I believe that there is increasing realization in industry that Information Theory provides the only reliable guidance for sound efficient design. It used to be that technology was way behind theory, and the big challenge was "how to do it?". Often now the bottleneck is not implementation but lack of complete theoretical understanding. The challenge is "what to do?"

In this respect, it is useful to contrast IT with the Complexity Theory community within Theoretical Computer Science: a relatively small community of abstract thinkers occupied with fundamental limits of efficient computation. The holy grail $NP \neq P$ seems more elusive than ever. Lately they are devoting a lot of efforts to studying the role of randomness in computation, the foundations of cryptography, interactive proofs, quantum computational complexity. A leading member of that community, Andy Yao, describes the situation as "a bunch of monkeys climbing trees in order to reach the moon." The big difference is that the real world of technology keeps us honest and relevant. That gives us a lot of credibility for the outside world. For us a 10% improvement grabs our attention, for the complexity theory people there is little difference between linear and n^{1000}.

Anantharam I might like to add a couple of notes of caution while we are busy congratulating ourselves on how successful we have been. I think we have also been great beneficiaries of a lot of serendipity. My own view of research is that there is a certain kind of randomness involved in the generation of new ideas and many of the ideas that have driven the field are not really of our own doing. There are great individual results that have come out perhaps with atypical frequency over the last couple of decades which one could not really have predicted. Quite apart from that, there is also the general evolution of technology. Moore's law, as Giuseppe mentioned, the driver that hardware provides to create problems where we are relevant is something that we shouldn't forget. In fact, one can mention the possible applications of some of our techniques at levels that were not even conceived of at the time of the previous panel, for instance coding at the level of communication over VLSI buses. So my main point in the context of this discussion is that we should be a little bit humble about why we are so successful—it's not entirely of our own making.

Verdú *Just like in 1994 I think it is futile to try to predict what the disruptive problems that will revolutionize the field in the future. Nevertheless, it is useful to discuss those current topics with the highest chance to have an impact on the future development of Information Theory.*

Kramer I suppose we all have our favorite current topics that we feel are important now and that we can predict will be important in the future. One of my favorites is Information Theory applied to optical channels, including optical fiber (MIMO is hot), free space, non-coherent vs. coherent, quantum, and so forth. A second favorite topic of mine is whether and how one can transfer the substantial progress on understanding relay and interference channel capacity into wireless systems. A third favorite is the same question concerning network coding and its application to distributed storage.

Caire At the risk of being disproved and laughed about by those who will read the transcript of this panel in a few years from now, I am going to try a forecast. What I'd like to see in the next 5–6 years from now is the development of a "communication theory for networks''. We gained an enormous insight about network Information Theory, in understanding interference and relaying. Nevertheless, we are very far from a ''plug-and-play'' set of techniques around which novel physical layer architectures can be actually designed. To make an analogy, in point-to-point communications we had Ungerboeck TCM, and now Turbo and LDPC codes followed by some form of bit-interleaving and mapping onto modulation alphabets. These techniques have been widely studied at the point that they have become standard tools around which systems based on point-to-point links can be safely designed. Still, in order to handle interference we rely on orthogonal (or quasi-orthogonal) access, and treating interference as noise, or as ''collisions''. We are still very far from the point where a new set of codes in the signal space (lattice codes? polar codes applied to multiuser problems?) can be used as basic building blocks for a robust system design. Of course, the risk of not filling this gap between Information Theory and communication theory (and therefore, practice of system design) is that these areas will remain confined in the purely theoretical domain and they may eventually fade away.

Anantharam As I said in the beginning, much of the success we have had in this field is centered around problems that are in the communication theory arena, but I think there are vast realms out there that are waiting to be conquered by what you might call "information-theoretic thinking". Shannon basically brought information-theoretic thinking to bear on communication theory. But there are aspects of nature, for instance, which have to have been designed with the concept of the optimization of some kind of information content in view. When you have a lot of interacting entities, either entities in nature or entities that you want to design, for instance when you want to design a biological system, which is eventually going to happen, there has to be a thinking, both in the engineered design and in nature as it came up with the designs that we are aware of today, which involves a notion of some information aspect that was optimized in enabling the coordination between the interacting entities. I am not sure on what time scale this will happen, but for instance biology is advancing at an enormous rate, so it could very well be twenty years, maybe longer, I think what we will see is a success of information-theoretic thinking in a lot of fields that have nothing to do with communication theory. Of course we are going to see all these great things happening in communication theory, which is very close to our hearts, but if you were to come to a meeting with the title "Information Theory workshop" twenty years from now, maybe we will have, say 20% of the papers where people are discussing for instance why a certain organ in the body works the way it does because the cells have enabled the coordination between themselves by optimizing some kind of information measure. We are really waiting for the Shannon for all of these different fields. That Shannon hasn't arrived and it is not clear if that Information Theory will look the way Shannon's Information Theory does, but I think we can rest assured that the Shannon will eventually come.

Verdú We are waiting for the second coming of the Messiah.

Costa I mentioned something about the increase in capacity that we have already seen, with fibers transmitting at 100 Tbps. But just to mention something that is anticipated on the demand side, the traffic in the internet is supposed to multiply by four in the next four years. We are supposed to have 15 billion network connected devices by 2015, and the expected overall traffic is supposed to be one zetabyte (10^{21}) per year in 2015. So when we get together for ISIT in Hong Kong in 2015, we will be able to check on these predictions.

In the long run, not focusing strictly in communications, but in the broader aspects of Information Theory, I think there will be a number of changes that will come up, particularly because of some of the connections that were already mentioned, with economics and biology. In fact we have already seen some BCH code structures in some proteins, and I think that the secrets of biological codes will be revealed little by little, protein by protein. Of course that will have a tremendous impact on what we will be able to do. I think polar codes will also have a great impact, and they will be extended to multiple user channels. We have already seen some of this happening with multiple access channels. Modern coding theory is now basically prevalent and to some extent has replaced the drive in algebraic coding theory, but I think that algebraic coding theory will make a strong come back in network coding applications and in the design of cyclic and quasi-cyclic LDPC codes with more predictable performances and substantially decreased error floors. Also quantum codes will come to be something that people will relate to in a more pragmatic way. And their anticipated impact, I believe, is enormous.

Another thing that I notice is that I remember Tom Cover, many years ago, talking about multiple user Information Theory as the foundation to a more general network understanding. Of course the setting at that time was completely different. (Incidentally, again Shannon was the first one to write something about multiple user channels with his two-way channel paper in 1961.) Network coding has brought techniques to surpass the classical max-flow-min-cut limits, and even the famous butterfly network has seen improvements in some cases, when the intermediate router does not need to access all the inputs, but needs just a function of these inputs. Even though we have had these tremendous advances in network coding theory we have not yet seen a marriage of that theory with the basic tiles or bricks that form multiple user structures, like broadcast channels, multiple access channels, interference channels and relay channels. I think there will be more development in these areas, and it may be a stretch, it may take a long time, but eventually we will see some conciliation and integration between the approaches of network coding and multiple user Information Theory.

Also source coding is still far from achieving the limits. I must say that I didn't expect to be alive on the day that channel coding limits would be approached as they are, within a small fraction of a dB. I really thought that this would be happening after my time. Now we can ponder that source coding is still not at that point, and hopefully we will still be around when those limits are approached within a fraction of a dB. More effective ways to combine source and channel coding will also became prominent. Improved and new inequalities will continue to extend the power and the scope of Information Theory tools.

I believe there are many fronts in which Information Theory will continue to bring significant contributions, both in practical technologies and in pure scientific and mathematical issues. So rather than thinking that Information Theory will eventually come to some type of blockade, I am more inclined to think that Information Theory will never die.

I would like to quote Karl Popper on something similar to the idea of monkeys trying to reach the moon. My son Bruno is a philosopher and we have some interesting discussions about this sort of thing. This is something that he told me. Karl Popper used to say that we may be very different in the ways we do things and in our abilities and knowledge about things, but in our infinite ignorance we are all alike.

Yeung A few years ago I had a chat with Prakash Narayan who was a panelist 17 years ago. I was pointing out the fact in the control theory community people had been using optimization tools for decades. Prakash made a very interesting point. He said that once the structure of a problem is exposed, what remain are algorithms and optimization. So, I think at least in the context of communications, in the Information Theory community we are going to see more and more of that.

Since we are still on the broad topic of "New Perspectives for Information Theory," I want to pick up a point Sergio mentioned a little while ago, regarding the lack of entropy of topics at ISIT. Personally, this actually bothers me quite a bit. I remember visiting Jim Massey in 2000 at Copenhagen. I was mentioning to Jim that these days research has become so competitive in many areas that if you don't publish something immediately, then very likely 3 months later somebody else will publish the same result. And Jim said, "In that case you shouldn't publish the result." I am not sure whether this is the best way to survive in today's research environment. Ideally we should all be working on problems that we think are important instead of following what the trend is doing all the time. But in the United States in particular, research is pretty much driven by funding. Whatever they call for you have to work on it, although you can do things in disguise. You have to follow the game.

Verdú Looking at my crystal ball, I see going forward: breakthroughs in multiuser Information Theory; intersections of Information Theory with machine learning, with signal processing, with compressed sensing, with theoretical computer science. Maybe one day we will think of the beginning of the XXI century as the era of bad quality: dropped cellphone calls, bad skype connections, lousy youtube video quality. This should put pressure in narrowing the gap between lossy compression theory and practice, and to that end one of the requisites is to learn how we can fool the eye and the ear more effectively than today. New approaches drawn from other fields such as random matrices and statistical physics methods are gaining prominence. And finally, non-asymptotic Information Theory: many practical applications are characterized by short messages or strict delay constraints. In the non-asymptotic regime we do not have the luxury of the closed-form formula, but we can still get very tight bounds as a function of delay.

Anantharam Other modern mathematical tools are also being brought to bear on Information Theory problems, e.g. from additive number theory in problems of interference alignment, and new kinds of concentration inequalities from our improved understanding of concentration of measure. Extracting structure from randomness is central to many branches of mathematics.

Verdú *It is time for me to thank all the attendees, my fellow panelists, and Valdemar da Rocha and Sueli Costa for organizing this workshop and providing the impetus for the organization of this panel.*

Addendum *Ioannis Kontoyiannis of the Athens University of Economics and Business was scheduled to participate in the panel but had to cancel his appearance. Here are some of his thoughts regarding the panel discussion.*

Kontoyiannis Reading the transcript of the panel discussion that was held at ITW in 1994, one notices that our community has made absolutely no progress in answering "foundational" questions like, "Is Information Theory part of applied mathematics or is it an engineering discipline?" I consider this a great success. About 15 years ago, the speaker in a philosophy of science seminar I was attending remarked that, when a field enters existential, esoteric discussions of this kind, it is usually a sign of intellectual decline. My (admittedly self-serving) view is that, as a field, we have been so successful that we can afford to avoid entertaining these questions seriously. When things are looking up, one rarely worries about the meaning of life. This success, as far as I can judge, has been facilitated to a significant extent by the combination of two distinct qualities. The field of Information Theory is defined by basic problems we wish to solve; and the community is fearless in bringing in the right tools to attack these problems. We are not a "methods looking for problems" discipline, and we have been open to the use of whatever new mathematical tool works – from the traditional analytical and probabilistic machinery of applied mathematics to the use of elliptic curves, random matrices, additive combinatorics, and so on.

On the other hand, the field is mature enough that it is now seen by researchers in numerous other areas as a collection of useful tools for their problems. The near simultaneous appearance of special issues on "Information Theory in neuroscience" in our Transactions and in the Journal of Computational Neuroscience is strong evidence of this trend. The editorial in the JCN special issue concludes: "Information Theory is thriving in the neuroscience community, and the long efforts are bearing fruit, as diverse research questions are being approached with more elaborate and refined tools. […] Information Theory is firmly integrated in the fabric of neuroscience research, and a progressively wider range of biological research in general, and will continue to play an important role in these disciplines."

Shannon's bandwagon warning notwithstanding, it is probably a safe prediction that this trend – information theoretic-ideas and tools being systematically applied in biology and perhaps in the other sciences – will continue and it will grow. In the reverse direction, another recent –though somewhat less noticeable– trend has been the growing use of information-theoretic concepts in core mathematics research. Although this was advocated by Kolmogorov almost 30 years ago ("Information Theory must precede probability theory and not be based on it. [...] The concepts of Information Theory […] can acquire a certain value in the investigation of the algorithmic side of mathematics as a whole"), progress has perhaps been slower and less flashy than the corresponding successes in, e.g., biology. But there are numerous examples – including Perelman's proof of the Poincaré conjecture and the celebrated Green-Tao theorem on the existence of arithmetic progressions in the primes – where Shannon entropy and the associated "technology" have served as important intellectual guidelines for major mathematical breakthroughs. This is another direction that I believe will continue strong and will gain momentum.

Finally, one of the essential components of our trade has to do with building foundations. Given a new communications scenario – be it a new technology with different physical characteristics, a new biological setting describing the communication between two distinct parts of an organism, or a new type of network model like those we have been studying in recent years arising in social media interactions – we abstract its fundamental characteristics and provide a rigorous "language" for its study. Keeping an open mind – and open doors – towards such new problems virtually guarantees a healthy outlook and a wealth of opportunities. A recent success story in this direction is the area of "compressed sensing." This could well have become a sub-field of statistics or harmonic analysis. The fact that it was embraced by the Information Theory community is a testament to both our open-mindedness and our strength.

I cannot resist one last comment. We really need to figure out how to do lossy compression effectively in practice!

Report on the Princeton CCI Workshop on Counting, Inference, and Optimization on Graphs

*Jin-Yi Cai, Michael Chertkov, G. David Forney, Jr.,
Pascal O. Vontobel, and Martin J. Wainwright (co-organizers)*

Over 100 participants attended an interdisciplinary workshop on "Counting, Inference, and Optimization on Graphs" at Princeton University, NJ, November 2–5, 2011. The workshop was organized by the authors under the auspices of the Center for Computational Intractability (CCI) at Princeton.

The workshop was originally inspired by the recognition of connections between certain duality results in the theory of codes on graphs and recent work on "holographic" algorithms in theoretical computer science. Ultimately, topics included holographic algorithms, complexity dichotomy theorems, capacity of constrained codes, graphical models and iterative decoding algorithms, and exact and approximate calculation of partition functions of graphical models. The participants had a wide range of backgrounds, including theoretical computer science, information and coding theory, statistical physics, and statistical inference.

The program is listed below. Copies of slides, references to related papers, and videos of some of the talks are available on the conference website at <http://intractability.princeton.edu/blog/2011/05/workshop-counting-inference-and-optimization-on-graphs>.

The participants were enthusiastic about the quality of the talks, the stimulation of various cross-disciplinary dialogues, and the excellent arrangements provided by the Center for Computational Intractability.

Program:

Leslie Valiant, "Holographic algorithms"

Jin-Yi Cai, "Complexity dichotomy theorems for counting problems"

Mark Jerrum, "Approximating the partition function of the ferromagnetic Ising model"

Leslie Ann Goldberg, "Approximating the Tutte polynomial (and the Potts partition function)"

Martin Loebl, "Complexity of graph polynomials"

Predrag Cvitanović, "Dynamical zeta functions: What, why, and what are they good for?"

Michael Chertkov, "Gauge transformations and loop calculus: General theory and applications to permanents"

Moshe Schwartz, "Networks of relations in the service of constrained coding"

Vladimir Chernyak, "Planar and surface graphical models which are easy"