Information Theory and Communication Networks: an Unconsummated Union

Anthony Ephremides and Bruce Hajek

Abstract— **Information theory has not yet had a direct impact on networking, although there are similarities in concepts and methodologies that have consistently attracted the attention of researchers from both fields. In this paper, we review several topics that are related to communication networks and that have an information theoretic flavor, including multiaccess protocols, timing channels, effective bandwidth of bursty data sources, deterministic constraints on datastreams, queueing theory, and switching networks.**

Keywords— **Communication networks, multiaccess, effective bandwidth, switching**

I. INTRODUCTION

Information theory is the conscience of the theory of communication; it has defined the "playing field" within which communication systems can be studied and understood. It has provided the spawning grounds for the fields of coding, compression, encryption, detection, and modulation and it has enabled the design and evaluation of systems whose performance is pushing the limits of what can be achieved. Thus, it constitutes a scientific success story of almost unparalleled proportions to which we pay tribute during this golden anniversary year of its birth.

However,information theory has not yet made a comparable mark in the field of communication networks, the sister field and natural extension of communication theory, that is today, and is likely to remain for many years, the center of activity and attention in most information technology areas. The principal reason for this failure is twofold. First, by focusing on the classical point-to-point, source-channel-destination model of communication, information theory has ignored the bursty nature of real sources. Early on there seemed to be no point in considering the idle periods of source silence or inactivity. However, in networks, source burstiness is the central phenomenon that underlies the process of resource sharing for communication. Secondly, by focusing on the asymptotic limits of the trade-off between accuracy and rate of communication, information theory ignored the role of delay as a parameter that may affect this trade-off. In networking, delay is a fundamental quantity, not only as a performance measure, but, also, as a parameter that may control and affect the fundamental limits of the rate-accuracy trade-off.

In fact, part of the reason why information theory did not go far enough in providing a solid theoretical foundation for networking is the urgency for rapid resolution of practical network design problems that has contributed to the creation of an anti-intellectual bias in parts of the networking community. At the same time, information theory has not done much to dispel that bias.

During its early development, information theory did consider multiuser systems [1], [2] and much of the subsequent work on such systems tried to capture (and did) many of the fundamental differences between the classical, stand-alone, single-channel case and that of the shared channel in multiuser systems. For example, it was realized that although feedback from the receiver to the source did not have an effect on channel capacity in single-user memoryless systems, it did have an effect in the case of multiuser systems [3]. But, still, the study of these systems has continued to be conducted in the restricted framework of nonbursty and delay-insensitive sources.

In this paper we will not address multiuser information theory, which is reviewed elsewhere in this issue [4]. Yet, there is a major need for a better synthesis between multiuser information theory and the networking topics discussed in the sequel.

In the past few years, the impact of the development of wireless systems, such as cellular networks, on information theory has been to steer the attention of its powerful principles and techniques toward the deeper significance of feedback information, in the form of channel measurement, and its effect on the choice of adjustable parameters such as transmission power and rate. But even in this case, the main thrust of the work continues to ignore the intrinsic role of delay and burstiness. Nonetheless, it has spawned the rapid development of the field of multiuser detection (see [5] in this issue and, for a more thorough account, [6]). In a sense, both multiuser information theory and multiuser detection theory represent major forays of information theory toward the field of networks that, so far, have revealed insights but have not yet produced the deep breakthrough that will have the same definitive impact on networking as it did on single point-to-point communication. We will not be addressing multiuser detection theory here either. It represents a distinct, self-sufficient field that, nonetheless, has intrinsic connections to both multiuser information theory and networking.

Just as communication systems were designed and built during the pre-Shannon years based mostly on heuristics, empirical knowledge, and partial dependence on theories of related fields (such as propagation, filtering, etc.), it is fair to say that today communication networks are designed and built based on similarly inadequate principles and techniques. And, yet, there is increasing evidence that the catalytic (almost messianic) effect of Shannon's work

Anthony Ephremides is with the Department of Electrical Engineering and the Institute for Systems Research at the University of Maryland, College Park, MD. E-mail: tony@eng.umd.edu

Bruce Hajek is with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. E-mail: b-hajek@uiuc.edu

on point-to-point communication may be brought about on the field of networks as well, either by the elaboration and enhancement of the same fundamental ideas of information theory that caused the revolution that started in 1948, or by some novel and revealing breakthroughs of a different kind that are, however, just as likely to come from information theorists or people with information theoretic training and background.

This assertion is based on the fact that some of the most influential and far-reaching advances in the field of networking, as well as some of the most intriguing observations about network behavior, originated from information theory scientists. It should be remembered that for the first twenty years, or so, of information theory, very little was actually accomplished in bridging the gap between theory and practice in point-to-point communication.

It is the intention of this article to document this assertion and to describe in more detail the relationship between information theoretic ideas and networking. Thus, although the final chapters of the impact of information theory on networking have not been written yet, we intend to review what has been achieved so far, inadequate and incomplete thought it may be, but to also speculate about the powerful potential of information theory to shape the future of communication networks.

The article is organized as follows. The next section reviews early work. Section III discusses timing channels and related topics, that include the protocol information needed to properly identify packets, and the existence of covert channels associated with the timing of signals or packets. Section IV reviews two methods for quantifying the effective data rate of a bursty source, that are similar to the use of entropy and rate-distortion functions as measures of effective data rates. One method is based on the theory of large deviations in queues, and the other on a calculus of deterministic constraints. Section V discusses the problem of random multiaccess communication, in which information plays a key role in a distributed setting. Two works combining aspects of multiaccess and information theory are discussed. Section VI briefly discusses queueing theory and its relation to information theory. Section VII discusses switching networks. The theory of basic switching network design is intertwined with information theory, and switching networks form the heart of the nodes within large communication networks. Section VII concludes the paper with a look to the future.

II. EARLY WORK

A. Network Layers

The principal communication network that existed during the formative years of information theory was the circuit-switched telephone network that was, by and large, conceived of, and operated as, a conglomeration of individual point-to-point links. The origins of the ideas of message and packet switching that have transformed the way communication networks are thought of, can be found in the emergence of computer communication and the interconnection of the, then, so-called, interface message processors. Among the first who formulated the backbone elements of packet-switched networking was Kleinrock who, first, in his original work [7] that was based on his Ph.D. dissertation and, subsequently, in his two-volume book [8] on queueing systems, popularized many of the innovative intricacies and challenges of communication networks.[1] A substantial volume of other work in the late sixties and early seventies [9]-[15], mostly by computer scientists and engineers, and the global interest on the still embryonic, but rapidly growing, field led to the formulation of the seven-layer, Open System Interconnection (OSI) framework and to the useful, at the time, separation between the physical, the link, and the higher layers.

There is a strong revisionist feeling today with respect to the notion of layering. It is increasingly realized that the original convenience and structure provided by the layering concept is superseded by the inherent coupling between the layers in almost every aspect of network operation. The artificiality of layer definition is apparent in some cases and concealed in some others. A case in point is that the original seven-layer OSI model left no natural place for multiaccess mechanisms, which parallel certain physical, link, and network layer mechanisms in the OSI model.

Nonetheless, the layered framework of network study has helped considerably in isolating individual networking problems that have been successfully attacked. Interestingly, in a way, the layering idea is first found in Shannon's work. Shannon clearly conveyed that the discrete, digital channel that he studied is a layer above the underlying physical, analog channel and the process of channel coding is a layer below the process of data compression. It is a pity that, in his work, Shannon did not expand on this concept. It might have saved a lot of time for networking researchers who, in a sense, reinvented the concept and first implemented it in an awkward framework.

B. Protocol Overhead

The first to recognize the significance of networking to information theory, both in terms of the challenges as well as the opportunities it presented, was Gallager who in 1973 [16] offered a clear vision of the natural connection between the two areas[2]. And he was the first to point out the fundamental significance of source burstiness and its relationship to information rate. In his landmark paper [22] on the subject, he considered a simple multiplexer of a finite number of sources, each of which was transmitting symbols from a ternary alphabet $(0, 1, i)$, where i indicated idleness and, therefore, did not carry message information. The critical

[1] The serendipitous presence of Claude Shannon in Kleinrock's Ph.D. defense committee may have been the forecaster of the bond between the two fields.

[2] We do not consider Shannon's work on the two-way channel [17] to be a genuine grasp of networking; this may be a debatable point, however. If we do accept that it does, then, again, Shannon must be credited with the prophesy of almost all aspects of the field of communication. It is also interesting to note that in [18] (as well as in [19] and [20]), a first look at max-flow min-cut relationships is provided; thus, the notion of flow approximations, widely used in network studies, was again, first noted by Shannon (among others). This notion was pursued further by Elias in [21].

observation was that, nonetheless, i did carry information. It carried the "message-start" or "message-end" information and was, therefore, an important participant in the information output of the source. By assuming geometrically distributed message-lengths and idle-period-lengths and independence among the sources, he computed the entropy of the sources, and hence the channel capacity needed to transmit the generated information with constant delay. The surplus (over the mean rate of data bits) was due to the "start" and "end" information inherently generated by the transitions between the "on" state and the "off" state. Gallager called this surplus the "protocol" information, since it represented the overhead price that had to be paid to accommodate the multiplexing of the bursty sources. The remarkable result is that this information can dominate the total transmitted amount of information. This work is discussed in more detail in Section III.

The key contribution of that early work was to show that even in the simplest of networks the need of overhead protocol information can expand significantly the amount of needed resources. This observation provided crucial conceptual and quantitative explanation to the alarming experience of early network engineers who found that their designs sometimes required that each packet carry a substantial amount of overhead bits, and so appeared to be very inefficient.

However, there has been almost no impact of this work on the design of practical systems. The reason is that the actual overhead inefficiency of most currently used network protocols is so large that the portion of overhead that handles burstiness is relatively limited and thus tolerable. This is not to say, however, that the timing overhead idea will not find application at some point in the future.

C. ALOHA and Multiaccess Protocols

At about the same time another early contributor to information theory, Abramson, proposed a simple idea that, perhaps because of its simplicity and its potency, had a major impact on the entire field of multiaccess communication [23]. Confronted with the practical difficulty of ensuring access to the mainframe computers of the University of Hawaii by terminals located in the outer islands of the state, Abramson proposed the simplest of ideas – pure random access. Eventually known as the ALOHA protocol, the simple scheme of attempting transmission randomly, independently, distributedly, and based on simple quantized feedback from the receiver, fertilized (if not created) the field of local area networks (whether radio-based or cable-based) and triggered an avalanche of work on what came to be known as the multiaccess problem. Subsequently, sophisticated schemes were proposed that combined ideas of fixed allocation (such as TDMA or FDMA), with reservations and contention to create the familiar protocols of Carrier Sense Multiple Access (CSMA) and CSMA-CD (CSMA with collision detection). To this day, the basic problem of access, that was so brilliantly illuminated by Abramson's ALOHA ideas, remains generally unresolved although quite thoroughly understood. The beauty of the ALOHA protocol was enhanced by Abramson's method of analysis, described in Section V.

It should be noted that there is no aspect of information theory that is directly involved in the entire story of ALOHA and random access. Still there is a flavor that is unmistakably information-theoretic in the formulation, exposition, and interpretation of this simple protocol. The ability of the model to capture what is essential in the contention process brings to mind the familiar models of the Binary Symmetric channel (BSC) or the additive white Gaussian channel (AWGN), or, later, the multiaccess adder channel, all of which share with ALOHA the same simplicity and predictive power.

The explosive spread of interest in the collision channel model and the problem of multiaccess communication that Abramson's work generated, led naturally to sophisticated and detailed analyses of modifications that would guarantee stability, and to an eventual refinement and redefinition of the problem that identified its connection to the problem of group testing or collision resolution. The area is briefly reviewed in Section V of this article.

D. Routing

Another early landmark in the history of contributions to the field of networks by information theory (or information theorists) is the resolution of the question of minimum-delay routing in packet-switched, store-and-forward networks. In a network of fixed topology and given source-destination node pairs with associated input traffic levels, the (very practical and important) question was to determine the optimal routing paths that yield minimum weighted total average delay. For clear implementation reasons (to reduce state-information latency and overhead and to ensure improved survivability and robustness) dynamic and distributed solutions were preferable. In [24], Gallager presented a concise and direct formulation of the problem accompanied by an elegant solution that permitted each node, based on simple periodic information exchanges with its neighbors, to determine the best next step in the path of each "commodity" (i.e., source-destination pair). The proposed algorithm yields convergence to the optimum and is even able to successfully "chase" a shifting optimum provided that conditions in the network (such as input traffic and topology) change at a rate less than the convergence rate of the algorithm.

It was realized soon after the publication of [24] that Gallager's algorithm is an independently derived solution to a special case of convex optimization problems eminently studied and analyzed by Bertsekas [25]. This realization led to the collaboration between these two authors that produced the classic text on networking [26] that summarizes the field in the most complete and scientifically sound fashion. The algorithm originally proposed in [24] and modified accordingly in [25] is compatible with the class of distributed Bellman-Ford-type algorithms [27] and charts a journey in the connection between networking and distributed algorithms, graph theory, and optimization. It is ironic that the relationship between communication networks and control system methodology (another, not fully explored and exploited relationship) that has been identified in [28] as well as in several subsequent publications and forums [29], was actually first pointed out, through [24], by information theorists.

Information theorists played a part in the origins of the field of distributed network protocols. Early in the implementation of packet switching networks it became clear that protocols are needed to coordinate a network. For example, upon start up, each node in a network might first learn the identity of its neighbors. Then through message passing, a spanning tree might be discovered by the nodes in order to serve as a backbone for the passing of control information. One basic question, quite natural for an information theorist, is what sorts of things are possible. For example, if nodes can enter and exit a network, and if routing tables are to be maintained, is it necessary to use sequence numbers? (The answer is no [30]). Another basic question, quite natural for an information theorist to ask, is "How many messages must be passed to accomplish a task." This is known as the communication complexity. For example, the paper by Gallager et al. [31] gives an efficient distributed algorithm for finding a minimum weight spanning tree given weights on the edges connecting nodes.

There has been much more in the brief history of networking that can be attributed to information theoretic thinking. Much of it is reviewed in this article. What is even more interesting is what has not yet been done that can be done by information-theoretic methods. We attempt to provide some glimpses to some of these opportunities as well.

III. TIMING CHANNELS

There are interesting connections between information theory and the timing of packets in a communication network. We first mention a source coding problem and then some channel coding problems that arise in connection with timing. Early in the development of computer communication, asynchronous communication emerged in which data is sent in packets. A packet is a finite sequence of bits. Typically packets generated by a source in a communication network are to be reproduced at a destination. This necessitates the use of some mechanism such as start or stop flags, or headers indicating packet length, or synchronization and fixed packet lengths. In a pioneering paper, Gallager [22] quantified the amount of protocol information per packet that is needed for reconstruction of the packets at the destination with a specified mean delay. Gallager took the interesting stance that, "to an information theorist, a protocol is a source code for representing control information." For example, if the delay per packet is to be identically constant, then the protocol must convey not only the values of the bits within the packets, but it must also convey the generation time of the packets and the packet lengths. As pointed out in [22], if the packet lengths are small compared to the random interarrival times, then the protocol information required per packet can far exceed the mean number of data bits in a packet.

For example, a datastream generated by a bursty source could consist of data bits (0's and 1's), interspersed with strings of i's representing idle time slots. The information rate required to reconstruct the source exactly is not just the mean arrival rate of data bits, but is equal to the entropy rate of the source viewed as one with the ternary alphabet $\{0, 1, i\}$.

Gallager realized that constant delay reproduction of packets is too strong of a requirement, so he explored the protocol information required to reconstruct a sequence of packets within a certain mean delay. This gives rise to the formulation of a rate-distortion problem, where the distortion measure is mean delay, and the rate is the protocol information per packet. The source generates packets according to a Poisson point process of specified rate, and the protocol must convey sufficient information per packet to enable reconstruction of the packets with specified mean delay. It is also assumed that packets are presented at the destination in order, and that each packet is presented at the destination only after it is generated by the source. For example, the time axis could be divided into intervals of length $2D$, and all packets arriving during each such interval could be delivered at the destination at the end of the interval. Then the required protocol information per packet related to arrival times would simply be the entropy of the number of arrivals per period divided by the mean number of arrivals per period. Note that the output would not determine the exact arrival times. A different rate-distortion function for Poisson processes was defined and identified by Verdú [32].

Even if enough protocol information is provided to identify the packets at the destination within a specified delay, such delay may be unobtainable due to the possible queueing delay experienced by bursty datastreams transmitted by constant rate transmitters. This issue is addressed in Sections IV and VI.

The flip side of the coin is that timing can be used to convey information. For example, if a source can make use of the three symbols $\{0, 1, i\}$, then through coding it could send information at rate $\log_2 3 \approx 1.585$ bits per channel use. The actual capacity of this channel is thus higher than the naive thought that at most one bit per channel use can be conveyed. In some situations, the information carrying capacity hidden in packet timing can be undesirable. For example, suppose that an agent is only authorized to send (or only pays for sending) particular types or amounts of information. The packets sent by the agent might be monitored. However, the agent could transmit additional information covertly by encoding it into the timing of packets.

Another example of a timing channel is the phone ringing channel. One party can convey information to a second party at no charge. The second party never answers the phone but only observes the times that it rings, that are controlled by the first party so as to convey a message (see [33] for a mathematical formulation and capacity result). A so-called two-ring, four-ring answering machine conveys information in the reverse direction as follows. It answers after four rings if it contains any unplayed messages, and after two rings otherwise. When picking up messages remotely, the owner hangs up after three rings, knowing there are no unplayed messages.

One countermeasure for covert communication is to introduce "timing noise" into the communication channel. A device that randomly delays packets could be inserted on all output lines in an effort to mask timing information. One possible device is a simple single-server queue with random service times. In fact, the Shannon capacity of the single server queue with service times that are independent, identically distributed random variables with some fixed mean μ was recently identified by Anantharam and Verdú [33]. They found that the capacity of such a queue, when over the long run packets transit the queue at rate λ, is given by the surprisingly simple formula $C(\lambda) = \lambda \log(\mu/\lambda)$ for $0 \leq \lambda \leq \mu$. This capacity tends to zero as either λ tends to zero (since then there are few packets to convey information) or λ tends to μ, since then the queue is nearly always full of packets so that the time between outputs is often just that of the service time distribution. Remarkably, the capacity of the $\cdot/M/1$ queue does not increase with feedback information. In addition, if the exponential service time distribution is replaced by another with the same mean, then the capacity cannot decrease [33]. The Shannon capacity of a discrete-time queue is addressed in [34] and [35].

As an aside, we briefly note an application of [33] to the source coding problem of [22]. Given a rate λ Poisson process of packet arrivals and a mean delay constraint D, the rate-distortion problem of [22] involves randomly delaying the points by at most D on average, in such a way as to minimize the mutual information per packet between the input and output streams. One could simply try taking the single server exponential queue as the delaying mechanism. The service rate μ should be selected to be $\mu = \lambda + 1/D$, so that the mean delay induced by the queue is D. The mutual information between input and output, divided by the input rate, is thus $R_\lambda(D) = \log(1 + 1/\lambda D)$. This is an upper bound on the rate-distortion (where distortion is delay) function of [22]. For small λD this bound asymptotically coincides with the lower bound on the distortion rate-distortion function given by Gallager, therefore eliminating a small gap left in his paper. However, for large λD a bound in [22] is smaller, indicating that the single-server exponential server queue is not a mutual-information minimizing delay mechanism for a Poisson input source.

There are many less obvious examples of covert communication channels within distributed computing systems and computer networks. For example, a multiple level security system is to offer services to clients with different levels of security. There may be two clients, one low and one high, and the system should restrict, and ideally completely prevent, the flow of information from high to low. One scenario is known as the computer processing unit (CPU) scheduling channel, and dates back to [36] and [37]. (See [38] for more background and citations.) Both clients submit tasks to their respective queues, one low queue and one high queue. The tasks are served by a single processor, that divides its service among the two queues in a round-robin fashion. Each client observes the completion times of the jobs that it submits to the queues. The question is, can the high client send information to the low client? The answer is clearly yes. The high client, depending on what message it wants to send to the server, carefully controls the times that it places jobs in its own queue. For its part, the low client carefully submits jobs to the queue and observes the sequence of response times. From the response times, the low client can learn the message that the high client intended to covertly send.

In another scenario, the low client sends a datastream to the high client. (The data sent need not be covert–the point of the multiple level security is to prevent information exchange in the reverse direction.) On the high side of the system there is a finite buffer into which the data is first placed, and later it is taken up by the high client. Suppose there is some protocol that gives feedback from the high side to the low side in order to acknowledge receipt of the data, or to warn the low side client to slow down because the buffer is nearly full, or to notify the low side client that the buffer did overflow and drop packets that must be retransmitted. Again, the question is, can the high client send information to the low client? Yes, the high client can carefully choose when to read data from the buffer, which influences the feedback messages from the high side buffer to the low side client. In this way the high side client can convey messages to the low side client.

In either of the above scenarios, the existence of other clients that are not participating in the covert communication could be considered to cause noise on the covert channel, giving rise to subtle and complex multiuser communication channels [39].

To summarize this section we note that there is much to the theory and practice of timing information and timing channels which remains to be understood, especially in network scenarios. Additionally, information theoretic ideas can play an important role in providing such understanding.

IV. TRAFFIC MODELING

There has been an extensive effort since the inception of packet-switched communication networks to characterize the traffic carried by networks. The work aims to account for the bursty nature of many data sources. In this section, two concepts arising in this work that strike us to be particularly close to the ideas and principles of information theory are reviewed: the effective bandwidth of datastreams, and deterministic traffic constraints.

Recently, there has been a keen interest in accounting for the observations of many studies of traffic in real networks, that indicate that datastreams exhibit self-similar behavior. That is, the random fluctuations in the arrival rate of packets appears to be nearly statistically the same on different time scales, ranging over several orders of magnitude. We touch on this development briefly in the context of the effective bandwidth of a self-similar Gaussian source. An extensive annotated bibliography on the subject is given in

[40].

A. Effective Bandwidth of a Datastream

One of the primary goals of information theory is to identify the effective information rate of a data source. The entropy or the rate-distortion function of a data source may be thought of as such. The theory of effective bandwidth, described in this section, has a similar goal. The word "bandwidth" is in this context an entrenched misnomer for data rate. Another connection between the theory of effective bandwidth of datastreams and information theory is that much of the theory of effective bandwidth is based on large deviations theory, which intersects Shannon's theory of information. Moreover, more direct connections between the theory of effective bandwidth and Shannon's theory of information are possible. For example, perhaps an "effective-bandwidth versus distortion" function can be computed for some nontrivial sources.

A major way that the theory of effective bandwidth differs from the Shannon theory is that it treats the flow of data bits as it would the flow of a fluid. The values of the bits are not especially relevant. The idea is that individual connections or datastreams carried by a network may be variable in nature. The data rate of each source may be constant in time, but a priori unknown, in which case we suppose the rate of such a source to be random. Or the sources can have time-varying rates. Suppose many variable datastreams are multiplexed together onto a line with a fixed capacity (measured in bits per second). Because of statistical multiplexing, the multiplexer has less work to do than if all the datastreams were sending data at the peak rate all the time. Therefore a given datastream has an effective bandwidth (that depends on the context) somewhere between the mean and peak rate of the stream.

To illustrate the ideas in the simplest setting first, we begin by considering a bufferless communication link, following Hui [41], [42]. The total offered load (measured in bits per second, for example) is given by

$$X = \sum_{j=1}^{J} \sum_{i=1}^{n_j} X_{ji}$$

where J is the number of connection types, n_j is the number of connections of type j, and X_{ji} is the data rate required by the ith connection of type j. Assume that the variables X_{ji} are independent, with the distribution of each depending only on the index j. If the link capacity is C then the probability of overload, $P[X > C]$, can be bounded by Chernoff's inequality:

$$\log P[X \geq C] \leq \log E[e^{s(X-C)}] = s(\sum_{j=1}^{J} n_j \alpha_j(s) - C) \quad (1)$$

where $\alpha_j(s)$ is given by

$$\alpha_j(s) = \frac{1}{s} \log E[e^{sX_{ji}}]. \quad (2)$$

(The bound (1) is trivial in case $\alpha_j(s) = +\infty$ for some j). Thus, for a given value of γ, the quality of service constraint $\log P[X > C] \leq -\gamma$ is satisfied if the vector $n = (n_1, \ldots, n_J)$ lies in the region

$$\mathcal{A} = \{n \in R_+^J : \min_{s>0} [s(\sum_{j=1}^{J} n_j \alpha_j(s) - C)] \leq -\gamma\} = \cup_s \mathcal{A}(s) \quad (3)$$

where

$$\mathcal{A}(s) = \{n \in R_+^J : \sum_{j=1}^{J} n_j \alpha_j(s) \leq C - \frac{\gamma}{s}\}. \quad (4)$$

The *complement* of $\mathcal{A}$ relative to R_+^J is convex. Let n^* be on the boundary of $\mathcal{A}$ (think of n^* as a "nominal" value of the vector n). A polyhedral subset of $\mathcal{A}$, delineated by a hyperplane tangent to the boundary of $\mathcal{A}$ at n^*, is given by $\mathcal{A}(s^*)$, where s^* achieves the minimum in (3). Thus, any vector $n \in Z_+^J$ satisfying

$$\sum_{j=1}^{J} n_j \alpha_j(s^*) \leq C - \frac{\gamma}{s^*} \quad (5)$$

satisfies the quality of service constraint. Once C, γ, and s^* are fixed, the sufficient condition (5) is rather simple. The number $\alpha_j(s^*)$ is the effective bandwidth of a type j connection, and $C - \gamma/s^*$ is the effective capacity. Condition (5) is analogous to the condition in classical information theory that ensures that a particular channel is capable of conveying several independent data sources within specified average distortions, namely that the sum of the rate distortion functions evaluated at the targeted distortions should be less than or equal to the channel capacity.

A caveat regarding the use of (5) is in order: For large values of s^* the value of $\alpha_j(s^*)$ can be very sensitive to variations in the upper tail of the distribution of X_{ji}.

As long as the random variables X_{ji} are not constant, the function α_j is strictly increasing, and ranging from the mean, $E[X_{ji}]$ as $s \to 0$, to the peak (actually the essential supremum, $\sup\{c : P[X_{ji} > c] > 0\}$) of X_{ji}. Note that the effective bandwidth used depends on the variable s^*. Such dependence is natural, for there is a tradeoff between the degree of statistical multiplexing and the probability of overload, and the choice of the parameter s^* corresponds to selecting a point along that tradeoff curve. As the constraint on the overflow probability becomes more severe, a larger value of s^* is appropriate. For example, if γ is very large, then the sets $\mathcal{A}(s)$ are nonempty only for large s, so that the choice of s^* is also large, meaning that the effective bandwidths will be near the peak values.

The set $\mathcal{A} \cap Z_+^J$, where $\mathcal{A}$ is defined in (3), is only a subset of the true acceptance region $\mathcal{A}_o$, defined by

$$\mathcal{A}_o = \{n \in Z_+^J \ : \ \log P[X > C] \leq -\gamma\}.$$

However, the sets $\mathcal{A}$ and $\mathcal{A}_o$ are asymptotically equivalent in the following sense. Let $\mathcal{A}/C$ (resp. $\mathcal{A}_o/C$) denote the set $\mathcal{A}$ (resp. $\mathcal{A}_o$) scaled down by a factor C. Note that $\mathcal{A}/C$ depends on C and γ only through the ratio C/γ. Then the Hausdorff distance between the sets $\mathcal{A}/C$ and $\mathcal{A}_o/C$ tends to zero as C and γ tend to infinity with C/γ fixed [43]. This follows from Cramér's theorem (see [44]), to the effect that Chernoff's bound gives the correct exponent.

So far, only a bufferless link confronted with demand that is constant over all time has been considered. The notion of effective bandwidth can be extended to cover sources of data that vary in time, but that are statistically stationary and mutually independent [45], [46], [47]. Let $X_{ji}[a,b]$ denote the amount of data generated by the ith connection of type j during an interval $[a,b]$. We assume that the process X is stationary in time. Set

$$\alpha_j(s,t) = \frac{1}{st}\log E[e^{sX_{ij}[0,t]}]. \tag{6}$$

For t fixed, the function α_j is the same as the one-parameter version of α_j considered above, applied to the amount of work generated in an interval of length t. Beginning with the well-known representation of Loynes for the stationary queue length, $Q(0) = \sup_{t\geq 0} X[-t,0] - tC$, we write

$$\begin{aligned}
&\log P[Q(0) > B] \\
&\quad = \log P[\sup_{t\geq 0}\{X[-t,0] - tC\} > B] && (7)\\
&\quad \sim \sup_{t\geq 0} \log P[X[-t,0] - tC > B] && (8)\\
&\quad \sim \sup_{t\geq 0}\min_{s\geq 0}\left[st\sum_{j=1}^{J} n_j\alpha_j(s,t) - s(B+tC)\right] && (9)
\end{aligned}$$

The symbol "$\sim$" used in (8) and (9) denotes that the ratio between the quantities on either side of it tend to one. This asymptotic equivalence is justified by limit theorems in at least two distinct regimes: (1) the buffer size B tends to infinity with n and C fixed, and (2) the elements of the vector n, the capacity C, and the buffer space B all tend to infinity with the ratios among them fixed. Under either limiting regime, the line (8) is justified by the fact that the probability of the union of many rare events (with probabilities tending to zero at various exponential rates) is dominated by the probability of the most probable of those events. The line (9), which represents the use of the Chernoff bound as in (1), relies on the asymptotic exactness of the Chernoff bound (Cramér's theorem or more general large deviations principles such as the Gärtner-Ellis theorem [44]).

Equations (7)-(9) suggest that the effective bandwidth to be associated with a connection of type j is $\alpha_j(s^*,t^*)$, where t^* achieves the supremum in (9), and s^* achieves the minimum in (9) for a nominal value n^* of n. The approximate condition for meeting the quality of service requirement $\log P[Q(0) > B] \leq -\gamma$ for n near n^* is then

$$\sum_{j=1}^{J} n_j\alpha_j(s^*,t^*) \leq C + \frac{B}{t^*} - \frac{\gamma}{s^*t^*}$$

This region scales linearly in γ if n^*, B and C scale linearly in γ, and asymptotically becomes a tight constraint as $\gamma \to \infty$. The value t^* is the amount of time that the system behaves in an unusual way to build up the queue length just before the queue length exceeds B. The quantity $C + B/t^* - \gamma/s^*t^*$ is the effective capacity of the link. Following [48], we call t^* the critical time scale. In the first limiting regime, described above, t^* tends to infinity, so the effective bandwidth becomes $\alpha_j(\infty, s^*)$. Use of the Gärtner-Ellis theorem of large deviations theory allows the limit theorems in the first limiting regime to be carried out for a wide class of datastreams with memory.

The above approximation simplifies considerably in the case that the datastream rate is Gaussian. In particular, suppose also that there is only one class of customers (so we drop the index j and let n denote the number of connections) and that for each i, $X_i(0,t]$ is a Gaussian random variable with mean λt and variance $V(t)$. The corresponding effective bandwidth function is $\alpha(s,t) = \lambda + sV(t)/2t$. Inserting this into (9) and then performing the minimization over s yields that

$$\log P[Q(0) > B] \sim -n\inf_t \frac{((c-\lambda)t + b)^2}{2V(t)} \tag{10}$$

where b is the buffer space per connection (so $B = nb$) and c is the capacity per connection ($C = nc$).

Suppose $V(t)/t^{2H}$ converges to a finite constant σ^2 as t tends to infinity, where H, known as the Hurst parameter, typically satisfies $1/2 \leq H < 1$. If $H = 1/2$, we see the process does not exhibit long range dependence. In particular, if X has independent increments (therefore the increments of a Brownian motion with drift λ and diffusion parameter σ^2), then $V(t) = \sigma^2 t$ and moreover (10) holds with exact equality.

If $H > 1/2$ (but still $H < 1$) then the critical time scale t^* is still finite. That is, even in the presence of long range dependence, the critical time scale is still finite in the limiting regime of C, B and n tending to infinity with fixed ratios among them [48]. The value of $V(t)$ for t larger than t^* therefore does not influence the approximation.

See [49] and [43] for extensive surveys on effective bandwidth, and [40] for a very extensive bibliographic guide to self-similar datastream models and their use. The paper [50] presents significant bounds and analysis related to notions of equivalent bandwidth with a different terminology. Finally, the paper [51] connects the theory of effective bandwidths to thermodynamics and statistical mechanics.

B. Network Engineering Through Traffic Constraints

An alternative to treating datastreams with statistical methods is to impose deterministic constraints on the data admitted into the network. The responsibility for respecting the constraints might lie with the end user, or it could be policed at the network entry points. The selection of which constraints would be imposed on a particular datastream would be done at the time a connection is requested, possibly in conjunction with a pricing mechanism. In return, the network should be able to provide a guaranteed quality of service (such as specified maximum transit time) for a particular connection. The type of constraints used should satisfy the following requirements:

Flexibility The constraints should allow for a controlled degree of burstiness on the part of data sources.

Easy to Enforce or Monitor It should be easy to police a datastream (through dropping or delaying part of the stream) to produce an output stream satisfying the constraints. Also, it should be easy to determine whether a datastream is meeting a particular declared set of constraints.

Operational Significance to the Network It should be possible for the network to exploit the constraints on admitted datastreams in order to deliver performance guarantees. This may entail, for example, providing end-to-end delay guarantees by bounding the delay for each device or link transitted.

A popular datastream constraint, introduced by Cruz in [52], [53], is the (σ, ρ) constraint, defined as follows. Consider a datastream described by a function $(R(t) : t \geq 0)$, where $R(t)$ denotes the amount of data generated up to time t. Assume that $R(0) = 0$, and that R is right continuous. Clearly R is nondecreasing. Let $\sigma \geq 0$ and $\rho > 0$. The stream R is said to satisfy the (σ, ρ) constraint if $R(t) - R(s) \leq \sigma + \rho(t - s) \quad \text{whenever} \quad s < t$. We discuss briefly why this particular constraint satisfies the requirements above.

First, regarding flexibility, the constraint allows a stream to contain an occasional burst of size σ, as long as in between the bursts the data rate falls below ρ enough. Secondly, in order to enforce a (σ, ρ) constraint, or to monitor a datastream to see whether it is in compliance with the constraint, a so-called "leaky bucket" regulator can be used. A leaky bucket regulator operates as follows. Imagine a bucket that holds tokens, such that tokens arrive at rate ρ. Tokens that arrive to find the bucket full are lost (this represents the leaking from the bucket). Data packets that arrive at the input of the regulator instantaneously take a token from the bucket with them and then pass through the network. However, if no tokens are available in the bucket for a given data packet, then the packet may be queued until a token becomes available, or the packet may be simply dropped. In practice, the scheme can be implemented by using a single counter, that is incremented at rate ρ and is decremented whenever a packet passes through (as long as the counter is not already at zero).

Finally, the (σ, ρ) constraint has operational significance for a network. For example, if a (σ, ρ) stream passes through a buffered link with a constant service rate C, then the delay at the buffer will never exceed $D = \sigma/(C - \rho)$, and the output stream satisfies the (σ', ρ) constraint for $\sigma' = \sigma + \rho D$. The basic approach taken by Cruz [52], [53] allowed arbitrary or first-come, first-served order-of-service when multiple datastreams arrive at a link. Bounds on network transit delay were derived. Parekh and Gallager [54], [55] showed how tighter bounds on network transit delay can be obtained in a network through the use of datastream constraints and generalized processor sharing (weighted round robin) scheduling disciplines at network nodes. The paper [54] also introduced the important concept of a service curve, that summarizes the performance of a server using the generalized processor sharing discipline. The notion of service curves has been refined, beginning with [56], in order to provide a calculus characterizing both sources and servers in a unified framework with an appealing algebraic structure. Additionally, [57] indicates how to provide transit delay guarantees through the use of deadlines at intermediate nodes and earliest-deadline-first scheduling. Many concepts can be formulated in both a stochastic framework and in a deterministically constrained framework. For example, delay bounds in a switch under deterministic constraints at the input and output ports are given in [58], and a notion of equivalent bandwidth for datastreams satisfying (σ, ρ) (and peak) constraints is given in [59].

V. MULTIACCESS COMMUNICATION

The problem of multiaccess communication arises in the consideration of the simplest possible, nontrivial multiuser system. A common receiver is accessed by N sources through a common channel. The principal motivating practical applications are (i) "cable" local area networks and (ii) "radio" local area networks. In either case, the main ingredient of the problem is the contention among the sources and the need to share the channel resource.

The approach taken by multiuser information theory is to consider the N sources as abstract digital emitters that produce bits at constant rates $R_1, R_2, \ldots, R_N$, and to aim at characterizing the region of values of the R_i's that (with appropriate encoding) permit error-free communication to the receiver. This approach is amply explored elsewhere in this issue [4].

An intermediate approach, taken rather recently by researchers who are motivated by the cellular communication paradigm, continues to consider nonbursty, continuously transmitting sources, but it gives up the asymptotic approach of multiuser information theory. It focuses on finite performance criteria and goals. This approach has become known as the multiuser detection theory approach to multiaccess communication. It is also explored elsewhere in this issue [5]. The key notion is that, in principle, it is possible to improve upon the performance of the traditional matched-filter-based receivers that are optimal in single-user, AWGN channel environments. The details of implementation become especially interesting when Code Division Multiple Access (CDMA) signals are used, when adaptive antenna arrays are used to provide diversity transmission or reception, and when fading channels are encountered.

Coding, detection, good channel modeling, source burstiness and delay are all important issues. The canonical multiaccess network model described below focuses on the later two, whereas multiuser information theory and multiuser detection theory focus on the first three. In the terminology of layers, the topics of this paper are more at a multiaccess (MAC) layer or network layer, and the other topics are more at the physical layer. Research on the canonical multiaccess network model, in which data packets are dealt with as "black boxes" whose internal structure is irrelevant, helped to crystallize some basic concepts of multiaccess communication, especially regarding bursty sources and delay. However, the physical and MAC cannot be cleanly separated (see discussion in Section II.A), so that the areas of multiuser information theory, multiuser detection, and multiaccess networking issues are best understood or developed in concert.

The canonical networking model of multiaccess, considers the so-called collision channel as its basic resource model. This channel is time-slotted (the non-time-slotted version introduces nonessential variations that are nowhere as significant as the differences between synchronism and lack thereof in single-user channels or in channels in which the "bit-structure" of the packets is not ignored). The signals transmitted by the sources are modeled by fixed-length packets (the bit-content of which are irrelevant), each of which fits snugly within one channel time slot. If two or more users transmit their packets in the same slot, none of the packets are correctly received (i.e., a collision is said to occur). The users are informed about the outcome of events in each slot by a variety of feedback structures. The simplest assumes instantaneous ternary feedback (denoted by 0,1, or e) that indicates to all sources whether the slot was unutilized or idle (denoted by "0"), was utilized successfully through a single, and hence successful, transmission (denoted by "1"), or was wasted through a collision (denoted by "e"). There have been many variations of this structure (binary feedback, M-ary feedback, in which the number of colliding packets is known, delayed feedback, etc., etc.). They are adequately reviewed in [26], [60]. The fundamental behavior exhibited in this model is impervious to these perturbations. So, we focus here on the simple, ternary, instantaneous feedback, even though this particular model is (at least almost) never encountered in practice.

The basic question was to determine allowable transmission strategies of the N sources that can achieve high aggregate "throughput" with small access delay. If a magic genie could coordinate the transmissions, then the channel would act like a multiplexer with throughput one packet/slot, and the resulting delay would be caused only by congestion due to possibly bursty arrival streams, rather than by the access problem per se. Without such a genie, throughput near one can still be obtained by the use of a sophisticated distributed algorithm such as an adaptive version of time-division multiplexing. However, it is believed that to achieve throughput near one for very large N, the mean delay must also be large.

The first consideration of this model by Abramson [23] made the additional natural simplification that the number of sources N is infinite. Such an assumption, unnatural though it may appear at first, is a clever and useful one in that, first of all, it lower bounds the performance of a finite-user system (since it amounts to a pessimistic assumption that each user's packets may compete against each other). In particular, if for a given throughput rate the mean average packet delay is finite for the infinite N model, then bounded delay can be achieved uniformly over all large finite N. (Researchers believe that the converse is true as well, but we know of no proof of such a converse.) More importantly, the infinite N assumption permits the decoupling of the analysis from the non-essential details of each source's storage of incoming packets. With an infinite number of users and a finite combined offered data rate of λ packets/slot, each source will only generate a single packet in its lifetime and thus there is no need to track queueing delays at each terminal. Thus, the multiaccess channel model was coupled from the outset with the assumption of aggregate input data that was generated by a Poisson process of rate λ.

A. The ALOHA Multiaccess Protocol

The next question was, of course, to determine the protocol for packet transmission and retransmission. As mentioned earlier, Abramson proposed the original, simple, random access in which a terminal attempts transmission as soon as its packet is generated and, if unsuccessful, continues to attempt transmission after a random waiting period. This is the ALOHA protocol. By assuming (incorrectly) that the aggregate data process (that includes new and retransmitted data) is also Poisson of rate G and by assuming (incorrectly) that this protocol yields a steady-state equilibrium, it is a trivial exercise to determine that $\lambda = Ge^{-G}$. This equation captures the essence of ALOHA. It implies that the maximum achievable throughput is equal to $e^{-1} \sim 0.36$ and occurs at $G = 1$. It further implies that there is a bistable behavior (i.e., for the same value of λ there are two possible corresponding values, G_1 and G_2, of the total data rate). By turning the situation around and abandoning the stability assumption, one can still use this equation to see that the actual ALOHA behavior (as confirmed by experiments) will produce a deteriorating throughput ($\lambda \to 0$) and an increasing total transmission intensity ($G \to \infty$) as more and more terminals get "blocked" and thus slide into the retransmission mode.

The bottom line of the ALOHA analysis is that uncontrolled random access, in both theory and practice, is a poor performer (no surprise). Left alone under pure ALOHA, the system disintegrates. With appropriate controls that steer G around its optimal values of 1, only 36% of the "capacity"[3] of the collision channel is utilized. Clearly there should be better ways of legislating transmission and retransmission rights to improve performance. Indeed, for almost two decades after the introduction of the ALOHA concept, massive research (much of which is accounted and summarized in [60]) ensued, with the goal of determining the ultimate capabilities of random access; that is, determining the maximum stable throughput over the collision channel. And, yet, the simple ideas of the ALOHA protocol galvanized everyone's think-

[3] The term "capacity" is used here in the sense of maximum achievable throughput and has nothing to do with the concept of Shannon channel capacity.

ing about channel access in general. And, eventually, practical and well performing protocols were developed, that actually mix the random access element with ingredients of reservation and the concept of fixed access (like the standard Carrier-Sensing-Multiple-Access with Collision-Detection (CSMA-CD)). Such protocols might not have been invented without the catalytic effects of ALOHA, even though many of the assumptions in the ALOHA model are far from being practical.

Naturally, the first subsequent attempts centered around the modification and stabilization of ALOHA. Metcalfe [61] and Lam and Kleinrock [62] were the first to suggest control mechanisms that reduce the retransmission rates of individual sources when the transmission intensity increases. Based on the observed ternary feedback, it is possible to adjust the packet retransmission probability so as to keep G close to 1. The papers [63] and [64] independently gave the first proofs that finite mean delay can be achieved for the canonical model (with Poisson arrivals, corresponding to infinite N). Several other stabilization algorithms were given, including an interesting one of Rivest [65] based on Bayesian estimation.

The famous exponential "back-off" algorithm, that basically reduces the retransmission probability of a packet by a factor of 2 every time the packet experiences a collision, preoccupied the minds of many researchers for a while. It was initially conjectured that this algorithm would stabilize ALOHA's behavior but it was eventually shown that (somewhat surprisingly) it didn't for the case of $N = \infty$. For the values of N encounted in practice, the exponential backoff protocol and many other protocols are adequate, even though they would lead to instability for $N = \infty$, or to bistability or large mean delay for very large finite N.

In parallel, practically oriented engineers started incorporating elements of the real environment in the ALOHA protocol. For example, the ability to "listen" to the channel, and determine whether it is in use or not, should be used to avoid unnecessary collisions. Thus, CSMA was born, and its variants, based on the values of its various parameters (like persistence in transmission, or propagation delay) and on whether a packet is divisible or not (i.e., whether a detected collision can be aborted before the full length of a slot is wasted), were painstakingly analyzed [66] and were shown to yield throughput performance that did approach the limit of 1 packet/slot.

B. Conflict Resolution

To information theorists, however, this thinking was unsatisfying. Before understanding and exhausting the possibilities of what is achievable with the basic model, the rush to explore its modifications (useful in practice, though the latter may be) suggested lack of intellectual tenaciousness. So, it was not surprising that, as a segment of the community pursued the development of practical protocols that depended on ALOHA to variable extents, information theorists relentlessly continued to pursue the basic collision channel model.

The major thrust began when Capetanakis [67] and Tsybakov and Mikhailov [68] adopted a radically different approach to the problem of retransmission, that was also suggested by Hayes [69] in a somewhat different context. Capetanakis and Tsybakov and Mikhailov explored the simple idea that every collision should be resolved before additional transmissions could be permitted. What better way to resolve a collision than subdivide the sources of the collided packets into groups and permit those groups to transmit one at a time in a TDMA fashion? Thus, the connection of conflict resolution to group testing was identified.

It may be argued that, in so doing, one mixes pure random access with fixed sharing and/or reservations, depending on how one views the allocation of the slots to the subgroups of the collided users. This is true; however, this is done in response to the channel feedback, without violating the basic assumptions of "indivisible" packets, and without introducing additional features of the environment into the model. Thus, it penetrates the essence of the conflict resolution process.

Capetanakis started by considering a finite number of users (2^n), with known binary identities of length n. Each user could be thought of as a leaf of a binary tree of depth n. After a collision, one half of the users were allowed to attempt retransmission in the next slot (say, the half that corresponded to the upper half of the tree); if a success or an idle occurred, the users in the bottom half of the tree were enabled next. If a collision occurred, the subgroup was subdivided again into two subgroups and the process was repeated. The end of such a search through a given subset of the tree could be detected by the occurrence of two successive slots with successful transmissions and thus, one by one, all subgroups would be explored (with the size of each subgroup being as large as the feedback information would permit).

Such a search was, indeed, similar to that of statistical group testing methods that were introduced in the first half of this century. Soon, the partitioning method that was based on user-ID was replaced by an equivalent random experiment with binary outcomes, performed independently by each user involved in the collision. In this way, the method of Capetanakis could be performed on the canonical, infinite-user ALOHA model. The first results were not spectacular. The basic tree-algorithm (as it came to be known) was achieving a maximum throughput that was slightly higher than that of ALOHA (it was, in fact, 0.43). The big difference, however, was that the protocol was stable. So long as the input data rate was less than 0.43, the successful throughput rate was equal to the input rate.

One difficulty with the tree algorithm (as well as with all subsequent variations) was that it didn't offer itself to an elegant analysis. To track (and prove) the stability and to calculate the length of the conflict resolution period (which is a measure of packet delay and another quantity of fundamental interest in the networking view of multiaccess communication), one had to resort to rather abstruse and lengthy derivations, the likes of which have been referred to at times as "brute-force" methods, or as "19th century mathematics".

The ideas of Capetanakis and Tsybakov and Mikhailov excited the community (more so its information-theoretically inclined members). Several people on both sides of the Iron Curtain started thinking seriously about this new view of conflict resolution. Among others, Gallager, Massey, Berger, Humblet, Mikhailov, Moseley, Tsybakov, all contributed insights and suggestions that led to a series of improvements to the basic tree algorithm that gradually yielded higher values of maximum stable throughput. There is little value in recounting them here; in detail they were reviewed in [60], and most of them were building blocks that helped clarify the essence of the splitting process.

Eventually, the most natural formulation that emerged parsed the packets of the different users on the basis of time of arrival [70], [71]. So, in slot t (just after having resolved all collisions that were caused by packets that were generated prior to an earlier time slot t'), two parameters needed to be chosen: (i) the length Δ of the next interval to be resolved, i.e., the interval from t' to $t'+\Delta$, so that all packets that arrived at instants within that interval would form the next group that would be "searched", and (ii) the fraction α of that interval that would be searched next if a collision occurred when all arrivals in interval $(t', t'+\Delta)$ were enabled. The search was to proceed pretty much as in the basic tree algorithm of Capetanakis. That is, if the channel feedback was 0 or 1, this marked the end of the (in this case, very brief) current conflict resolution period. If the feedback was e (collision), the users in the first fraction α of the original interval would be enabled next. In case of collision that fraction would be subdivided anew (by the same fraction α); in case of success, the "enabled" interval would shift starting from $t'+\alpha\Delta$ and extending to $t'+\Delta$; and in case of an idle slot, the enabled interval would start from $t'+\alpha\Delta$ but would only extend to the fraction α of the interval $(t'+\alpha\Delta, t'+\Delta)$. The reason for the last choice resulted from the crucial observation, that was first made on Capetanakis's algorithm, that if a collision is followed by an idle, another collision is certain to occur if the entire balance of the originally enabled subgroup, that produced the collision in the first place, is enabled again. Thus, it is important to anticipate this occurrence and explore only a subset of that balance.

Another important observation is that if a collision follows upon the heels of another, there is no information about the contents of the unexplored portion of the first interval that yielded the first collision. Thus, instead of, when its time comes, visiting the unexplored portion alone, (of length $(1-\alpha)\Delta$), it is preferable to enable a full-length interval from $t'+\alpha\Delta$ to $t'+\alpha\Delta+\Delta$ (or to the current slot t, whichever is less).

These intricacies of the algorithm are clearly explained in [26]. The analysis of it, however, has been similarly plagued by the need for inelegant, computationally intensive methods that have aimed at establishing the same two performance indices of interest, i.e., the maximum stable throughput and the average packet latency. Clearly, though, by mapping the entire process of splitting into the time axis, based on time of arrival, one can see that both quantities (i.e., stability and delay) are captured by the "lag" between the current time t and the time of completed resolutions t'. Thus, the difference $t-t'$ is closely related to the duration of the conflict resolution period (and hence the packet delay) as well as to the "drift" of the resolution process. Unless $t-t'$ approaches a limiting distribution, the process is unstable.

The precise calculation of the maximum stable throughput of the FCFS (first-come, first-served) splitting algorithm (as it was eventually known) was accomplished in [72] through the policy iteration method of dynamic programming (where the problem was posed as one of optimization, i.e., maximization of the stable throughput, with respect to the choices of Δ and α). The precise calculation relies on extensive computations and thus the numerical accuracy of the results has been a question of some dispute. If α is decided to be chosen as $1/2$ and the optimization is carried out only with respect to Δ, it was determined that $\Delta \cong 2.6$ slots and the corresponding maximum stable throughput should be 0.4871 packets/slot (a significant improvement over ALOHA and the basic tree algorithm). However, the optimal value of α is not $1/2$, but, rather, very slightly less than $1/2$. There is no easy explanation for this but it does yield slightly higher throughput (0.487117 as claimed by Moseley and Humblet in [72]). Even more puzzling is an observation by Vvedenskaya and Pinsker [73] that the throughput can gain another small increment by somewhat modifying the lengths of the intervals after a large number of collisions. A possible resolution of these somewhat variable numerical values is offered by Verdú in an overlooked technical note [74], where the maximization problem is formulated in an elegant, iterative fashion that bypasses the need for complicated dynamic-programming-based reasoning. The result is that the FCFS algorithm with $\alpha = \frac{1}{2}$ yields a throughput of 0.487117, the Tsybakov-Mikhailov version with $\alpha = 0.485$ yields 0.487694, and the precise calculation by Verdú yields 0.487760.

It should be mentioned that the value of 0.4871 obtained by Gallager follows from a very elegant and simple argument based on the drift of the quantity $t-t'$ and also bypasses the obscuring mathematical details.

At the same time, a great deal of effort had been focused on looking at the problem from the other end. That is, by assuming that additional information is available and by determining the corresponding maximal throughput, one can obtain upper bounds on the throughput in the original problem. The first to obtain such a bound was Pippenger [75] who showed that the maximum stable throughput cannot exceed $\sim$0.73. A series of similar efforts followed and the currently known least upper bound [76] is 0.587. Some researchers conjectured that the optimal value might be 0.5, but this claim was quickly abandoned as baseless.

In this brief (about five-year), but intense, saga about zeroing-in on the maximum stable throughput of random access over the collision channel, there was a modest degree of similarity to the quest for establishing the true capacity of a channel in the usual Shannon-theoretic sense. The problem here had nothing to do with Shannon capacity and it was mostly an academic exercise of limited practical value.

The study of the collision resolution problem did not stop after the derivation of the results quoted above. A myriad of possible extensions and modifications were possible and many of them were pursued in considerable depth. For example, the issue of feedback delay or feedback errors, the issue of new users coming into the system (or old users dropping out) in the middle of a resolution period, the issue of multiple levels of feedback, and many other variations have been looked at over the years and continue to be looked at today, albeit with somewhat diminished interest. Again, many of these variations are reviewed in [26].

C. Finite-user ALOHA

Useful though the infinite-user model is, it is also worthwhile to examine systems with a finite number of users, in which, or course, a user's own packets don't collide with each other on the channel. In this case, each user generates a finite percentage of the total input data and, thus, it is necessary to queue up the arriving packets. Even if the other features of the collision channel model remain the same, the problem is now transformed in a significant way. It becomes a problem of non-standard queueing theory (i.e., one in which successive service times in each queue are not independent and/or in which service time durations are not independent of the arrival processes), known also as a problem of interacting, or coupled, queues.

Consider, N users, each with an infinite buffer and receiving packets independently at a rate λ_i. Thus, the total input rate is $\sum_{i=1}^{N} \lambda_i$. Each user attempts to transmit the packet at the head-of-the-line position in the queue in each slot with probability p_i (irrespectively of whether this is the first attempted transmission or a retransmission). The feedback from the collision channel is as before (0,1, or e). This model encapsulates the ALOHA protocol in a queueing environment. A central question is to determine the values of the rates λ_i, $i = 1, \ldots, N$, for which the average delay in all of the queues is finite.

With Bernoulli arrivals (or any other independent, identically distributed arrivals) this problem can be accurately modeled in a straight-forward way as an N-dimensional random walk. From the early work by Fayolle and Iasnogorodski [77] to more recent works by Szpankowski [78], Sidi and Segall [79], Rao and Ephremides [80], Anantharam [81], and others, it has become well-known that such chains cannot be easily solved. Thus, much of the work has concentrated on obtaining outer and inner bounds to the region of stability. A key idea that has yielded some of these bounds relies on partially decoupling the queues by considering as "bounding" systems those in which some of the queues stochastically dominate their counterparts in the original one and, hence, their stability implies the stability of the original system. Note, also, that in the N-user model, it is possible that some of the queues may be stable and others unstable. Recently, an index of "potential instability" for each queue was obtained [82], given by $\frac{\lambda_i(1-p_i)}{p_i}$. The meaning of this index is that if the queues are ranked on the basis of this index, that is, if queue i is stable, all queues j ($j < i$) are also stable and if queue i is unstable, all queues j ($j > i$) are also unstable.

If the transmission probability vector p can be adjusted as a function of the arrival rates (but not as a function of backlogs and feedback), we are led to consider a capacity region that is the union of arrival rate regions over all vectors p. Before examining such capacity region for the queueing model, we shall discuss the capacity region defined by Abramson [83] (and summarized in [8, vol. II]). The definition corresponds to the throughput vectors achieved by a saturated ALOHA system in which all users always have packets to send, so considerations of queueing and delay are avoided. Suppose that user i transmits in each slot with probability p_i, independently from slot to slot, and independently of other users. The success probability for user i is then $\lambda_i = p_i \prod_{j:j \neq i}(1-p_j)$. Abramson's capacity region, that we write as $\mathcal{C}_A$, is given as the set of all vectors $\lambda = (\lambda_1, \ldots, \lambda_N)$ obtained in this way, as the vector $p = (p_1, \ldots, p_N)$ varies. Abramson showed that the upper boundary of $\mathcal{C}_A$ is the set of those vectors λ obtained when p is a probability vector, meaning that it is desirable for the mean number of transmissions per slot to be one. In the special case of two users, Abramson showed the region reduces to

$$\mathcal{C}_A = \{(\lambda_1, \lambda_2) : \sqrt{\lambda_1} + \sqrt{\lambda_2} \leq 1\}. \tag{11}$$

Now we return to the queueing model. From the perspective of a given user i, the assumption that all the other users are busy is a pessimistic one. Therefore, the buffered ALOHA network is ergodic if the vector of arrival rates falls within $\mathcal{C}_A^o$, where $\mathcal{C}_A^o$ is the set of all $\lambda \in R_+^N$ such that for some p (depending on λ), $\lambda_i < p_i \prod_{j:j \neq i}(1-p_j)$. Equivalently, $\mathcal{C}_A^o$ is $\mathcal{C}_A$ with all points on the upper boundary deleted, and for the case of $N = 2$ users:

$$\mathcal{C}_A^o = \{(\lambda_1, \lambda_2) : \sqrt{\lambda_1} + \sqrt{\lambda_2} < 1\}. \tag{12}$$

Tsybakov and Mikhailov [84] first published this result, and moreover they showed for the case $N = 2$ that the region $\mathcal{C}_A^o$ is the *complete* stability region for the ALOHA network, rather than a proper subset of it. Specifically, with $N = 2$ and independent, identically distributed arrivals at each of the two users with means λ_i (and finite variance) per slot, the buffered ALOHA network is ergodic if and only if $\lambda \in \mathcal{C}_A^o$. Anantharam [81] showed that $\mathcal{C}_A^o$ is also the entire stability region for any N, but only for particular (unrealistic) arrival sequences that for different users are weakly statistically dependent. Anantharam's result suggests that $\mathcal{C}_A^o$ may well be equal to the stability region for any N and independent arrivals, but it also shows that the issue depends on subtle details about the interactions of the queues that are probably unimportant in applications.

D. Models with Elements of Multiuser Information Theory

Two noteworthy models that involve elements of multiuser information theory, and either queueing or collision access (but not both at once) are discussed in this section. One is the model of a collision channel without feedback, introduced by Massey and Mathys [85]. It is assumed that there is no feedback, and moreover there is not even a way for the users to synchronize their transmissions. Forward error correction, rather than a retransmission protocol, is thus needed to achieve reliability. In this sense, the model is similar to the models used in the multiuser information theory literature, initiated by Shannon. On the other hand, the model differs from the usual models of multiuser information theory in that, to quote [85], "information is transmitted only in the contents of packets and not also in the timing of access attempts."

Massey and Mathys identify the capacity region, and show that it doesn't depend on whether the system is slot-synchronized or whether zero-error (rather than arbitrarily small error) probability is required. The capacity (=zero-error capacity) region they obtained is precisely the region $\mathcal{C}_A$ obtained by Abramson [83]. Massey and Mathys noted in [85] that it "seems somewhat surprising" that precisely the same set of rates can be achieved error-free without feedback as can be achieved under the slotted ALOHA system with feedback. See [80] for further elaboration. Is it a meaningless coincidence? Perhaps, but not likely. If not, then, what is the significance of it and in what way do the two very different notions relate to each other? To this day, there is no answer.

The other model we mention, proposed by Teletar and Gallager [86], combines elements of queueing theory and information theory for multiaccess communication. As in the buffered ALOHA model discussed above, a finite number of users accumulate randomly arriving packets to be transmitted. The packets are sent using forward error correction, and a small amount of feedback is available from the receiver to the users. Teletar and Gallager investigate the use of optimal codes (known to exist by random coding arguments, and hence the connection to information theory) for the forward error correction. The time needed to send a packet is variable in length, since the number of active users fluctuates. The feedback allows the receiver to notify the transmitter as soon as the receiver is able to decode a packet, for otherwise the transmitter would not know when to cease transmitting (redundant) bits pertaining to the packet. The resulting dynamics of the queuing process are much like processor sharing, in which the service rate experienced by a user is roughly inversely proportional to the number of active users.

The scheme of Teletar and Gallager is similar to IS-99 [87], the link level data protocol recently designed for use with the IS-95 CDMA cellular standard. Under the IS-99 link protocol, a standard Transport Control Protocol packet is divided into 32 frames, that are each transmitted using Code Division Multiple Access. Negative selective acknowledgments are sent by the receiver to a user to compensate for frame errors, that typically occur in one or two percent of the frames. A user with data sends a frame each 20 milliseconds with a probability p, where p is varied dynamically based on feedback from the base station, that is monitoring the signal-to-interference level. Thus, the transfer speed per user tends to diminish with the number of users.

E. Interaction between physical and higher network layers

Another direction in which network multiaccess communication has become intriguing is that of spatial diversity. With the increasing importance of sectorized and directional antennas, let alone adaptive antenna arrays, the possibility of space-division-multiple-access (SDMA) has becomes a reality. In the field of multiuser detection theory, which represents an intermediate stage between multiaccess information theory and network multiaccess, there has been considerable activity centered on detailed signal modeling, power-control, antenna patterns, channel interference models, and receiver structures that yield "throughput" results indirectly as functions of the required quality-of-service. That is, bit-error-rates are calculated as functions of the transmission rates, the transmission powers, the channel bandwidth, and the other design parameters.

At the networking level, it is of interest again (at least as a first approach) to suppress the system details into rigid "black boxes" and to attempt to capture the effect of directionality as a means of aiding in the resource allocation . It should be mentioned, as an additional example of the different viewpoints of multiuser detection theory and networking, that the phenomenon of "capture" (meaning that one of many competing simultaneously transmitted signals may be correctly received by a single receiver) can be modeled in very different ways. At one end, by taking into account the synchronization preambles of the signals (especially in the case of CDMA), the received powers, and the exact times of arrival, it is possible to get a detailed and accurate micro-model of how capture occurs and to then analyze its effects. At the other end, the networking view of capture has simply assumed that if multiple packets are simultaneously received, then either one packet can be successfully received with some probability p, or the packet with the highest received power is correctly received. Based on such simple modeling, one can derive the effect of capture on the throughput of the otherwise classical collision channel with the associated random access protocol. This was, in fact, done since the very early days of the history of multiaccess communication [88].

Recently [89], some attempts have been made to combine (up to a certain extent) the "black-box" mentality on capture with the detection-theoretic needs for more detailed modeling. The motivation for some of this work has been to study the role of energy conservation in wireless networks as a means of network control. Thus, the model in [82] assumes that the length of the packet is not constant anymore. Rather, for a fixed number of symbols per packet, it is possible to adjust its length (i.e., the rate of transmission) and keep the detectability criterion of the signal-to-interference ratio unaffected, provided that the transmission power is adjusted simultaneously. The effect on throughput performance is clear. If the packets shrink in length, and if the packet input rate stays constant, the overlaps that cause collisions become less likely. At the same time increasing the transmission power depletes battery energy faster, unless the energy savings, from having less wasteful transmissions due to reduced packet overlaps, prevail. At the same time, keeping several distinct power levels among the users facilitates capture (which enhances throughput). A first study of this elaborate trade-off shows that the throughput, (as well as the normalized throughput per energy-unit) is maximized if all users transmit at peak-power [82].

The problem becomes more intriguing if the coupling to the physical layer is permitted to strengthen. For example, if the modulation choice is not fixed, then the value of the threshold for the signal-to-interference ratio to ensure detectability at the desired bit error rate changes nonlinearly and the effect of the transmission power (vis-a-vis the transmission rate and the associate packet length) is unclear; furthermore, the vertical interlayer coupling can also be strengthened if a multiuser detector is assumed at the receiver. These ideas are still premature, and they represent only initial thoughts of current concern in the field of multiaccess communication.

But let us return to the issue of spatial diversity. The networking view is to simply consider the packets of the different users as having not only a "time-of-arrival" coordinate but also a "space-location" coordinate (limited for the moment to the single dimension of planar angle-of-arrival). Thus, the collision channel can now be studied as before with the simple additional feature that the receiver is able to focus an ideal beam onto the location that it chooses, along with a chosen value of beamwidth angle. This is simply equivalent to enabling the transmissions of subgroups of users not only by sorting out their time-of-arrival, but also their location. At first glance, it might appear that this increased capability can produce an increase in the value of the maximum stable throughput. After some thinking, however, it is not surprising to see that, as shown in [90], this additional degree of freedom cannot increase the throughput.

Of course, the interesting case is the one that involves more than a single beam; in that case it is clear that throughput gains can be indeed realized. This is another case for which work is just beginning and it is premature to report any definitive ideas or progress.

F. Wireless Networks

The rapid growth of wireless networking today is causing continued interest in a variety of multiaccess communication problems. The most prominent type of a wireless network today is the one based on the cellular model. In that model, the base nodes are accessed by mobile nodes, but are interconnected among themselves via a wired infrastructure that is part of the telephone switched network. The principal issues that need to be addressed in such networks include spectral efficiency (in terms of spatial frequency reuse), power control (for combating the near-far problem of CDMA signals), hand-offs among base-stations nodes, and mobile tracking as nodes move from cell-to-cell. Additional issues that deal with specialized applications (like mobile computing, multicasting, or information distribution) have to do with database structures, and signal compression. All these together transcend the confines of the subject of multiaccess communication. They have captivated the interest of the networking community (and the dollars of the community at large), and it is not clear what role information theory can play in it. However, multiuser detection theory and some recent work by Tse [91], Hanly [92], Knopp and Humblet [93], and Gallager and Medard [94], that study a variety of subjects associated with cellular models (such as effective wireless bandwidth, power-rate control, etc.), have a strong information-theoretic flavor. In addition, the theory of compression and multiple descriptions will undoubtedly play a key role in those wireless applications that deal with information distribution and database access.

But there is another form of wireless network that has emerged recently as a subject of great interest (especially in military applications) called all-mobile networks. They are described by a variety of other names like peer-to-peer, flat, multihop, ad hoc, and others. All-mobile networks have a large number of nodes with no hierarchy (no base station nodes) and no fixed infrastructure. All nodes may move and constantly change their neighbor sets. They all share the same frequency band and must communicate with each other in a flexible fashion that permits all kinds of services (data, voice, video, etc.). Clearly, they involve all the problems encountered in cellular networks plus many more. Early work on such networks [95] identified the need for, and methods to achieve, distributed, self-reconfiguration and has established some principles (or more accurately, problem areas) that govern their design and operation. An increasing segment of the networking community is zeroing-in on them and it is too soon to tell how information theoretic ideas may contribute to their study. Multiaccess communication, however, is a central issue for these networks and we would be remiss if we did not identify these networks in this section.

In conclusion, the burgeoning field of channel access, from its early modest phases to its current complex and multifaceted profile, has been one of the principal areas in which information theory has played, and will likely continue to play, a major role. The complexity and multitude of multiaccess-related issues that arise today (especially in the area of wireless networks) has led much of the networking community to a state of mild confusion. We believe that the simplicity and sharpness of information-theoretic ideas may yet penetrate the field further and illuminate those issues that are basic and fundamental.

VI. Queueing Theory

Queueing theory has provided the most useful analytical tools in the study of communication and computer networks. It has offered a natural foundation for delay analysis and has also been the source of sophistication for the description of complex interplay among network parameters.

Despite its central role in the theoretical side of networking, queueing theory remains, for the most part, uncoupled to information theory. It is, of course, closely connected to the theory of stochastic processes and, to the extent that the latter is related to information theory, one may claim that there is a certain connection between the two fields. But beyond the limited similarity in terms of asymptotics and stochastic analysis, there is no fundamental bond between the two disciplines. In fact, if there was such a bond, the missing link between delay and information theory would have been uncovered by now as well.

Actually, queueing theory has displayed much more affinity to control theory. Stochastic control of simple queueing models [96], dynamic adjustment of retransmission probabilities in random access systems [63], optimal routing [26], flow control, and many other networking problems have been fruitfully cast in the framework of control and optimization theory. A thorough survey of that connection can be found in [28], where it is shown how the methodology of system theory applies naturally to networking. Furthermore, the theory of discrete-event systems [29] has also found applicability to problem of network design and operation [97].

Yet, there have been some hopeful, albeit feeble, signs that the right way of combining information theory and queueing theory may, indeed, be taking shape. In Section III of this paper we mentioned the pioneering work of Anantharam and Verdú [33] on the Shannon capacity of a queue. In addition to determining the capacity of the simple queue channel, this work sets a landmark in the study of the two fields. Viewing a service system as a channel may prove to be nothing more than a whimsical, cute exercise; yet, it may prove to have a catalytical role in creating a common platform for the joint study of information-theoretic and queueing-theoretic systems. It may represent a pivotal moment in the history of the two fields. The interesting (common) part of that history has yet to be written, however. The enthusiasm of those who saw in this work an opportunity to advance the coupling of the two disciplines was quickly tempered by the difficulty of extending the approach to even the slightest perturbation of the plain $M/M/1$ system. For example, simply adding one more server (i.e., considering an $M/M/2$ system) complicates the analysis considerably. And yet, the two-server model would be invaluable in shedding more light on the interplay between information and waiting, since it captures the notion of increased bandwidth and parallel service.

The fact, discussed in Section V.D, that the queuing theoretic capacity region for multiaccess communication coincides (at least for N=2 stations) with the capacity region of the collision channel without feedback of Massey and Mathys [85] **may** (just may) be something fundamental tying queueing theory to information theory. The identity of the regions may be an instance of a yet undiscovered broader principle. The work of Teletar and Gallager [86], also discussed in Section V.D, illustrates some of the significant interactions between queueing and physical layer considerations, many of which involve elements of information theory. The notion of effective bandwidth, described in Section IV.A, is grounded in queueing theory, and as we mentioned it has some natural compatibility with information theory.

There have been other approaches recently that also attempt a joint study of information-theoretic and queueing-theoretic system aspects. For example, the use of variable rate source coding in conjunction with congestion control combines rate distortion theory with buffer management. It doesn't reach into any level of profundity, but it does permit (at least) a phenomenological coupling. In [98], Tse considered a version of this problem that can be thought of, not only as a study of the trade-off between information fidelity and congestion, but, also, as a means of coupling among the OSI networking layers, at least as far as quality of service is concerned.

VII. Switching Networks

There is a natural interplay between Shannon information theory and the theory of switching, routing and sorting in interconnection networks. The classical example is a circuit switch with n inputs and n outputs, interconnected by wires and relays (or crosspoints). Each relay has two states: open or closed, so that the number of internal states of the network is 2^R, where R is the number of relays. Suppose the switch is to be capable of connecting the n inputs to the n outputs according to any permutation. Because different permutations require different network states, the network must have at least $n!$ network states. This requires that $2^R \geq n!$ or that $R \geq \log_2 n \sim n \log n$. The earliest published account of this idea is that of Shannon [99]. Similarly, if the network is constructed of component switches of fixed in-degree and out-degree s and links between them, with each component switch capable of handing any of the $s!$ possible permutations of input to output connections, at least a constant times $n \log n$ such switches (for fixed s) are needed to connect any input to any output.

A simple, elegant construction of switching networks with the minimum number of two-by-two switches (within a factor of two) is the Beněs network, attributed by Beněs [100] to Slepian, Duguid and Le Core. A simple algorithm, now known as the 'looping algorithm", was given for the determination of routes. The Beněs network is not well suited to dynamic operation, in that if a set of routes are in progress and a new route between an idle input and idle output is requested, then rerouting of existing connections is sometimes required.

Thus, in addition to being able to route any permutation, it is also desirable that a switch be able to emulate a full crossbar switch in a *dynamic* fashion. The strongest form of this property, termed *strict sense nonblocking*, is the following: whenever a set of compatible routes are already carried by the network, and an idle input and an idle output are identified, it is possible to assign a route to the new input-output pair that is compatible with the routes already given. There is no information-theoretic argument that rules out the existence of strict sense nonblocking switches with complexity $O(n \log n)$, and indeed they were shown to exist by Pinsker and Bassalygo [101]. Pinsker and Bassalygo first showed the existence of bipartite graphs with certain expansion properties. Several stages of switches were then interconnected using such graphs at each step, so that from any idle input, or any idle output, strictly more than half the idle center state lines can be reached, so that there exist and end-to-end connection between the idle input and idle output. The construction of Pinsker and Bassalygo was nonexplicit, because the existence of the expanders was only shown by a random construction. That is, it was shown that with nonzero probability (in fact with probability tending to one as the size tends to infinity) a randomly constructed regular bipartite graph has the desired expansion property.

Just as algebraic coding theory seeks to find explicit and structured solutions to replace the nonexplicit constructions in Shannon's coding theorems, so too have researchers worked to find explicit and structured solutions for the construction of strict sense nonblocking networks. A breakthrough came in the paper of Margulis [102], who proposed a construction of related graphs with an expansion property, and used deep theorems from the theory of group representations to prove the expansion property. Gabbar and Galil [103], using relatively elementary methods of harmonic analysis, provided explicit constructions of expanders with explicit (though large) bounds on the required size. See [104] for a more detailed account of the chronology given here, including an exposition of the construction and proof of [103]. Of many notable improvements in explicit constructions of expanders that followed, we mention the work of [105].

In addition, with the growth of data over networks in packetized form, circuit switched connections have evolved to packet oriented connections such as virtual circuit connections or pure one-at-a-time datagram packet routing. Packet routing is closely connected to the theory of sorting networks. For example if a batch of input packets are addressed to the outputs in a one-to-one fashion, then routing the packets may be done exactly by a sorting network. The story regarding existence and explicit constructions for sorting networks somewhat parallels that for circuit switching networks. The explicit constructions of sorting networks with the minimum required order of complexity $O(n \log n)$, starting with [106], are much too large for currently practical implementation, whereas the sorting network of Batcher, with complexity $n \log n^2$, is quite effective for small networks.

The above story of probabilistic constructions followed later by explicit constructions parallels the development of channel codes. Recently a more concrete connection between the topics was made by Sipser and Spielman [107], who used expander graphs to construct a new family of asymptotically good, linear error-correcting codes with linear time sequential decoding algorithms.

The search for asymptotically optimal complexity strict sense expanders and sorting networks has so far been primarily one of theoretical consequence. In practical networks, switches are engineered only to have a small probability of internal blocking. This is akin to using codes, such as turbo codes, that have small minimum distance but still have a small error probability.

Information theoretic ideas are applied in [108] in the context of switching networks using deflection routing of packets. Deflection routing implies that all packets entering a node in one time slot exit the node in the next time slot. While the transit delay in a node is thus minimized, the drawback is that sometimes a packet exits a node on a link that does not help the packet progress towards its destination, in which case we say the packet is deflected. A lower bound on the mean number of hops a packet needs to travel is given in [108], assuming there are two outgoing links per node and that a packet is independently deflected with probability q in each slot. The lower bound is roughly the entropy of the probability distribution of the packet destination divided by the Shannon capacity of a binary symmetric channel with crossover probability q. The idea is that by observing the progress of a packet, an observer learns the destination of the packet, and such information is conveyed in spite of the the deflections, that are essentially noise. If there is only a single source node, the lower bound can be asymptotically achieved, through the use of a graph based on good channel codes for the binary symmetric channel [108]. For the more natural case in which any node can be a source or destination node, there is a gap between the lower bound and the mean number of hops needed for packets in the graph constructed in [108].

VIII. Future Work

Several problem areas from networking may hold considerable potential for information-theoretic analysis, and the opportunities to impact actual system implementation abound. The development of communication networks to support heterogeneous datastreams in heterogeneous networks promises to continue at a torrid pace for the next decade and beyond. This trend is fueled by the demand for higher speed, lower delay communication, anywhere, anytime. The distinction between computing and communication will increasingly blur, as network resource allocation involves interactions among fluctuations in datastreams (due to bursty sources), fluctuations in link capacities (due to fading channels and node mobility), and fluctuations of demand for computer cycles in multiprocessor environments. As the feature sizes of Very Large Scale Integrated (VLSI) chips decrease, the performance of the interconnects suffers relatively more than the performance of the devices [109]. Thus VLSI designers will have to confront increasingly slow and unreliable data links within a chip. Massive network communication problems emerge, and information theory should have a role to play in it.

The use of sophisticated antenna arrays for communication in fading channel environments is not well understood, though it seems that feedback provided by protocols can play an important role. Much more development in multiuser detection theory, including better channel modeling, will be needed, and much of that may be difficult to cleanly separate from network issues. Information theory could play a significant role in the mix.

While information theorists have made important contributions to the theory of automatic repeat request protocols [110], for the most part information theorists have invested much more effort in forward error control. Still, the use of feedback and automatic repeat request is sometimes clearly preferable to forward error correction. Consider for example a synchronous binary erasure channel in which each transmitted bit reaches the receiver with probability p, and is replaced by a null symbol otherwise, and the outcomes of different transmission are mutually independent. The Shannon capacity of this channel is p bits/second, and if immediate error-free feedback is available, then simply repeating each bit until it is successfully received achieves the capacity and at the same time minimizes the delay. The scheme is essentially unaffected if p is unknown or even time varying in an arbitrary way. In contrast, a forward error correcting scheme is almost unworkable in this circumstance. On the other hand, forward error control is typically better when feedback is not available or comes with long delay, and when the channel is well modeled. Better error control mechanisms, integrating both forward error correction and automatic repeat protocols, are needed in the context of networks. Feedback and delay considerations, as well as bit error probabilities, are important.

The interaction of source coding with network-induced delay cuts across the classical network layers and has to be better understood. The interplay between the distortion of the source output and the delay distortion induced on the queue that this source output feeds into may hold the secret of a deeper connection between information theory. Again, feedback and delay considerations are important.

But information theory has not paid full attention to the subtleties of feedback. Even though it was Shannon himself who chose to speak on the notion of feedback [111] in his speech that inaugurated the Shannon Lecture series, the important result [112] that feedback does not improve single-user memoryless channel capacity stymied somewhat the growth of interest on the issue of feedback. Even when it was shown that [3] feedback may increase multiuser capacity, the main action continued to be based on constant-rate transmission. Thus, feedback was incorporated only in the form of (to use a uniquely information-theoretic term for feedback) side-information at the transmitter. It was not considered in the networking context of timing, delay, and sporadic transmission.

A recent trend to replace circuits on long-haul backbone networks with packet switched data, and the evolution of the Internet, raises numerous challenges regarding fault-tolerance and security. We touched on the problem of covert channels, but many other security issues would seem amenable to information theoretic style approaches.

Our review of the topics in which information theory and networking seem to make contact suggest that the union between the two fields remains unconsummated. Yet, information theorists have maintained active attention to several problems of communication networks and several of their contributions strongly suggest a deeper relationship. In addition, many of the theoretical techniques used in network modeling and analysis are similar to those used in information theory. Finally, there are several aspects of networking, some old (like error control), some new (like uses of antenna diversity), and some spawned by recent observations (like timing in channels and queues) that hold promise for the eventual establishment of a firm and clear relationship between the two fields.

Acknowledgments

The authors would like to acknowledge very helpful comments from an anonymous reviewer and from Dr. Richard Blahut.

References

[1] R. Ahlswede, "Multiway communication channels," *Proceedings 2nd International Symposium on Information Theory*, Tsahkadsor, Armenia, 1971.
[2] A. El Gamal and T.M. Cover, "Multiple user information theory," *Proceedings of the IEEE*, Vol. 68, pp. 1466-1483, 1980.
[3] N.T. Gaarder and J.K. Wolf, "The capacity region of a multiple access discrete memoryless channel can increase with feedback," *IEEE Transactions on Information theory*, Vol. 21, pp. 100-102, 1975.
[4] R. Calderbank and S. Verdú, "
," *this issue.*
[5] T. Kailath and H.V. Poor, "Modern statistical processing: a retrospective," *this issue.*
[6] S. Verdú, *Multiuser Detection*, Cambridge University Press, New York, 1998.
[7] L. Kleinrock, *Communication Nets*, Dover Publications, New York, 1964.
[8] L. Kleinrock, *Queueing Systems, Vol. 1 &2*, John Wiley & Sons, 1975.
[9] V.G. Cerf and R.E. Kahn, "A protocol for packet network communication," *IEEE Transactions on Communications*, Vol. 22, pp. 637-648, 1974.
[10] H. Frank and W. Chou, "Topological optimization of computer networks," *Proceedings of the IEEE*, Vol. 60, pp. 1385-1397, 1972.
[11] I. Rubin, "Communication networks: message path delays," *IEEE Transactions on Information Theory*, Vol. 20, pp. 738-745, 1974.
[12] B. Meister, H. Muller, and H. Rudin, "New optimization criteria for message switching networks," *IEEE Transactions on Communications*, Vol. 19, pp. 256-260, 1971.
[13] F.W. Heart, R.E. Kahn, S.M. Ornstein, W.R. Crowther, and D.C. Walden, "The interface message processor for the ARPA computer network," *AFIPS Conference Proceedings*, Vol. 36, pp. 551-567, 1970.
[14] H. Frank and I.T. Frisch, *Communication, Transmission, and Transportation Networks*, Addison-Wesley, 1971.
[15] S.D. Crocker, J.F. Heafner, R.M. Metcalfe, and J.R. Postel, "Function-oriented protocols for the ARPA computer network," *AFIPS Conference Proceedings*, Vol. 40, pp. 271-279, 1972.
[16] R.G. Gallager, "Information theory - the first 25 years", Keynote Address, *IEEE International Symposium on Information Theory*, Ashkelon, Israel, June 1973.
[17] C.E. Shannon, "Two-way communication channels," *Proceedings Fourth Berkeley Symposium on Probability and Statistics*, pp. 611-644, 1960 (reprinted in D. Slepian, ed., *Key Papers in the Development of Information Theory*, IEEE Press, NY 1974 and in N. Sloane, A. Wyner, Eds., *Collected Papers of C.E. Shannon*, IEEE Press, NJ, 1993).
[18] P. Elias, A. Feinstein, and C.E. Shannon, "A note on the maximal flow through a network," *IEEE Transactions on Information Theory*, December 1956.
[19] G.B. Dantzig and D.R. Fulkerson, "On the max-flow min-cut theorem of networks," in *Linear Inequalities.* Princeton, N.J. *Ann. of Math Studies* 38, 1956.
[20] L.R. Ford and D.R. Fulkerson, "Maximal flow through a network," *Canad. J. Math.*, vol. 8, pp. 399-404, 1956.

[21] P. Elias, "Networks of Gaussian channels with applications to feedback systems," *IEEE Transactions on Information Theory*, pp. 493-501, July 1967.

[22] R.G. Gallager, "Basic limits on protocol information in data communication networks," *IEEE Transactions on Information Theory*, Vol. 22, pp. 385-399, 1976.

[23] N. Abramson, "The ALOHA system - another alternative for computer communications," *AFIPS Conference Proceedings*, Vol. 37, pp. 281-285, 1970.

[24] R.G. Gallager, "A minimum delay routing algorithm using distributed computation," *IEEE Transactions on Communications*, Vol. 23, pp. 73-85, 1977.

[25] D.P. Bertsekas, "Dynamic models of shortest path routing algorithms for communication networks with multiple destinations," *Proceedings of the IEEE Conference on Decision and Control*, Ft. Lauderdale, FL, pp. 127-133, 1979.

[26] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, First edition 1987, second edition 1992.

[27] L. Ford, Jr. and D. Fulkerson, *Flows in Networks*, Princeton University Press, 1962.

[28] A. Ephremides and S. Verdú, "Control and optimization methods in communication network problems," *IEEE Transactions on Automatic Control*, Vol. 34, pp. 930-942, 1989.

[29] C. Cassandras, *Discrete-Event Systems: Modeling and Performance Analysis*, R.D. Irwin, Inc. and Aksen Associates, Inc., 1993.

[30] J.M. Spinelli and R.G. Gallager, "Event driven topology broadcast without sequence numbers," *IEEE Trans. on Communications*, vol. 37, pp. 468-474, 1989.

[31] R.G. Gallager, P.A. Humblet, and P.M. Spira, "A distributed algorithm for minimum-weight spanning trees," *ACM Trans. Programming Language Syst.*, vol. 5, pp. 66-77, 1983.

[32] S. Verdú, "The exponential distribution in Information Theory," *Problemy Peredachi Informatsii*, vol. 32, no. 1, pp. 100-111, Jan.-Mar. 1996. English version in *Problems of Information Transmission*, vol. 32, no. 1, pp. 86-95, Jan.-Mar. 1996.

[33] V. Anantharam and S. Verdú, "Bits through queues," *IEEE Trans. Information Theory*, vol. 42, pp. 4-18, January 1996.

[34] A.S. Bedekar and M. Azizoglu, "The information theoretic capacity of discrete-time queues," *IEEE Trans. Information Theory*, vol. 44, pp. 446-461, March 1998.

[35] J.A. Thomas, "On the Shannon capacity of discrete time queues," *IEEE International Symposium on Information Theory*, Ulm, Germany, June, 1997.

[36] S. B. Lipner "A comment on the confinement problem," *Proc. of the 5th Symposium on Operating System Principles*, University of Texas at Austin, pp. 192-196, Nov. 1975.

[37] M. Schaefer, B. Gold, R. Linde, and J. Scheid, "Program confinement in KVM/370," In *Proc. of the Annual Conference*, ACM, pp. 404-410, Seattle, Washington, October 1977.

[38] I.S. Moskowitz, S.J. Greenwald, and M.H. Kang, "An analysis of the timed Z-channel," In *Proc. IEEE Symp. on Security and Privacy*, pp. 2-11, Oakland, California, May 6-8, 1996.

[39] M.H. Kang, I.S. Moskowitz, and D.C. Lee, "A network pump," *IEEE Transactions on Software Engineering*, vol. 22, no. 5, pp. 329-337, May 1996.

[40] W. Willinger, M.S. Taqqu, and A. Erramilli, "A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks," In *Stochastic Networks Theory and Applications*, F.P. Kelly, S. Zachary and I. Ziedins, Eds., Oxford Science, pp. 339-366, 1996.

[41] J.Y. Hui, "Resource allocation for broadband networks," *IEEE J. Selected Areas of Communications*, vol. 13, no. 6, 1598-1608.

[42] J.Y. Hui *Switching an Traffic Theory for Integrated Broadband Networks*, Kluwer, Boston, 1990.

[43] F.P. Kelly, "Notes on effective bandwidth," in *Stochastic Networks Theory and Applications*, F.P. Kelly, S. Zachary and I. Ziedins, Eds., Oxford Science, pp. 141-168, 1996.

[44] A. Shwartz and A. Weiss *Large Deviations for Performance Analysis, Queues, Communication and Computing*, Chapman & Hall, London, 1995.

[45] D.D. Botvich and N. Duffield, "Large deviations, the shape of the loss curve, and economies of scale in large multiplexers," *Queueing Systems*, vol. 20, 293-320, 1995.

[46] A. Simonian and J. Guibert, "Large deviations approximations for fluid queues fed by a large number of sources," *IEEE J. Selected Areas of Communications*, vol. 13, no. 6, 1081-1090.

[47] C. Courcoubetis and R. Weber, "Buffer overflow asymptotics for a switch handling many traffic sources," *J. Applied Probability*, vol. 33, no. 3, 886-903, 1996.

[48] B.K. Ryu and A. Elwalid, "The importance of long-range dependence of VBR traffic engineering: myths and realities," *Proc. ACM SIGCOMM*, 3-14, August 1996.

[49] G. de Veciana, G. Kesidis, and J. Walrand, "Resource management in wide-area ATM networks using effective bandwidths," *IEEE J. Selected Areas of Communications*, vol. 13, no. 6, 1081-1090.

[50] C.S. Chang, "Stability, queue length, and delay of deterministic and stochastic queueing networks," *IEEE Trans. on Automatic Control*, vol. 39, no. 5, pp. 913-931, May 1994.

[51] J.Y. Hui and E. Karasan, "A thermodynamic theory of broadband networks with application to dynamic routing," *IEEE J. Selected Areas of Communications*, vol. 13, no. 6, 991-1003, 1995.

[52] R.L. Cruz, "A calculus for network delay, part I: network elements in isolation," *IEEE Transactions on Information Theory*, vol. 37, pp. 114–131, Jan. 1991.

[53] R.L. Cruz, "A calculus for network delay, part II: network analysis," *IEEE Transactions on Information Theory*, vol. 37, pp. 132–141, Jan. 1991.

[54] A.K. Parekh and R.G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single node case," *IEEE/ACM Trans. on Networking*, vol. 1, no. 3, pp. 344-357, 1993.

[55] A.K. Parekh and R.G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the multiple-node case," *IEEE INFOCOM 93*, San Francisco, 521-530, March 1993.

[56] R.L. Cruz, "Service burstiness and dynamic burstiness measures: a framework," *Journal of High Speed Networks*, vol. 1, no. 2, pp. 105-127, 1992.

[57] G.L. Georgiadis, R. Guérin, V. Peris, and K.N. Sivrajan, "Efficient network QOS provisioning based on per node traffic shaping," *IEEE/ACM Transactions on Networking*, vol. 4, no. 4, pp. 482-501, August 1996.

[58] T. Weller and B. Hajek, "Scheduling nonuniform traffic in a packet switching system with small propagation delay," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 813-823, December 1997.

[59] A. Elwalid, D. Mitra, and R.H. Wentworth, "A new approach for allocating buffers and bandwidth to heterogeneous regulated traffic in an ATM node," *IEEE J. Selected Areas of Communications*, vol. 13, no. 6, pp. 991-1003, 1995.

[60] Special Issue on Random-Access Communications, *IEEE Transactions on Information Theory*, Vol. 31, March 1985.

[61] R. Metcalfe, "Steady state analysis of a slotted and controlled Aloha system with blocking," *Proceedings 6th Hawaii Conference on System Science*, Honolulu, Hawaii, 1973.

[62] S. Lam and L. Kleinrock, "Packet switching in a multiaccess broadcast channel: dynamic control procedures," *IEEE Transactions on Communications*, Vol. 23, pp. 891-904, 1975.

[63] B. Hajek and T. Van Loon, "Decentralized dynamic control of a multiaccess broadcast channel", *IEEE Transactions on Automatic Control*, Vol. 27, pp. 559-569, 1982.

[64] V. A. Mikhailov, "Geometrical Analysis of the stability of Markov chains in R^n and its application to throughput evaluation of the adaptive random multiple access algorithm," *Problems of Information Transmission*, 24, pp. 47-56, January-March 1988.

[65] R. Rivest, "Network control by Bayesian broadcast," *IEEE Transactions on Information Theory*, pp. 323-328, May 1987.

[66] L. Kleinrock and F.A. Tobagi, "Packet switching in radio channels: part 1: CSMA mode and their throughput-delay characteristics," *IEEE Transactions on Communications*, Vol. 23, pp. 1400-1416, 1975.

[67] J.I. Capetanakis, *The multiple access broadcast channel: protocol and capacity considerations.* Ph. D. Diss., MIT, Dept. of Electrical Engineering and Computer Science, 1977. Also *IEEE Transactions on Information Theory*, Vol. 25, pp. 505-515, Sept. 1979.

[68] B.S. Tsybakov and V.A. Mikhailov, "Free synchronous packet access in a broadcast channel with feedback," *Prb. Peredachi Infor.* (USSR). vol. 14, no. 4, pp 32-59, 1978.

[69] J. Hayes, "An adaptive technique for local distribution," *IEEE Transactions on Communications*, Vol. 26, pp. 1178-1186, 1978. (Also Bell Telephone Laboratory Memo TM-76-3116-1, 1976).

[70] R.G. Gallager, "Conflict resolution in random access broadcast

networks," *Proceedings AFOSR Workshop on Communication Theory and Applications*, Provincetown, MA, pp. 74-76, 1978.

[71] B.S. Tsybakov and V.A. Mikhailov, "Random multiple access of packets: part and try algorithms," *Problemy Peredachi Informatsii* Vol. 16, pp. 65-79, 1980.

[72] J. Moseley and P. Humblet, "A class of efficient contention resolution algorithms for multiple access channels," *IEEE Transactions on Communications*, Vol. 33, pp. 145-151, 1985.

[73] N. Vvedenskaya and M.S. Pinsker, "Non-optimality of the part-and-try algorithm," *Abstracts of the International Workshop of Convolutional Codes & Multiuser Communication*, Sochi, Russia, pp. 141-148, 1983.

[74] S. Verdú, "Computation of the efficiency of the Moseley Humblet contention resolution algorithm: a simple method," *Proceedings of the IEEE*, Vol. 74, pp. 613-614, 1986.

[75] N. Pippenger, "Bounds on the performance of protocols for a multiple access broadcast channel," *IEEE Transactions on Information Theory*, Vol. 27, pp. 145-151, 1981.

[76] V.A. Mikhailov and B.S. Tsybakov, "Upper bound for the capacity of a random multiple access system," *Problemy Peredachi Informatsii*, Vol. 17, pp. 90-95, 1981.

[77] G. Fayolle and R. Iasnogorodski, "Two coupled processors: the reduction to a Riemann-Hilbert problem," *Wahrscheinlichkeitstheorie*, pp. 1-27, 1979.

[78] W. Szpankowski, "Stability conditions for some multiqueue distributed systems: buffered random access systems," *Advances in Applied Probability*, Vol. 26, pp. 498-515, 1994.

[79] M. Sidi and A. Segall, "Two interacting queues in packet-radio networks," *IEEE Transactions on Communications*, Vol. 31, pp. 123-129, 1983.

[80] R. Rao and A. Ephremides, "On the stability of interacting queues in a multiple-access system," *IEEE Transactions on Information Theory*, Vol. 34, pp. 918-930, 1988.

[81] V. Anantharam, "The stability region of the finite-user slotted ALOHA protocol," *IEEE Transactions on Information Theory*, Vol. 37, pp. 535-540, 1991.

[82] W. Luo and A. Ephremides, "Effect of packet lengths and power levels on random access systems," *Proceedings 35th Allerton Conference on Communication, Control, and Computation*, Allerton, IL, 1997.

[83] N. Abramson, 'Packet switching with satellites, " *AFIPS Conference Proceedings*, Vol. 42, pp. 695-702, 1973.

[84] B.S. Tsybakov and V.A. Mikhailov, "Ergodicity of a slotted ALOHA system," *Problemy Peredachi Informatsii*, vol. 15, no. 4, pp. 73-87, October-December 1979. English translation in *Problems in Information Transmission*, pp. 301-312, vol. 15, no. 4, pp. 301-312, October-December 1979.

[85] J.L. Massey and P. Mathys, "The collision channel without feedback," *IEEE Transactions on Information Theory*, Vol. 31, pp. 192-204, 1985.

[86] E. Teletar and R.G. Gallager, "Combining queueing theory and information theory for multi-access," *IEEE J. Selected Areas in Communications*, vol. 13, no. 7, pp. 963-969, August 1995.

[87] Electronic Industries Alliance Standard TIA/EIA/IS-99: *Data Service Option Standard for Wideband Spread Spectrum Digital Cellular System (July, 1995)*, 1995.

[88] L. Roberts, "Aloha packet system with and without slots and capture," *ASS Note 8, ARPA*, Stanford Research Institute, June 1972.

[89] R.O. LaMaire, A. Krishna, and M. Zorzi, "Optimization of capture in multiple access radio systems with Rayleigh fading and random power levels," in *Multiaccess, Mobility, and Teletraffic for Personal Communications* (Eds. B. Jabbari, P. Godlewski, X-Lagrange), Kluwer, Boston, MA, pp. 331-336, 1996.

[90] K. Sayrafianpour and A. Ephremides, "Can spatial separation increase throughput of conflict resolution algorithms?" *IEEE International Symposium on Information Theory*, Ulm, Germany, June 1997.

[91] D.N. Tse and S.V. Hanly, "Multiuser demodulation: effective interference, effective bandwidth, and capacity," *Proc. 35th Allerton Conference on Communication, Control, and Computing*, Allerton, IL, 1997.

[92] S.V. Hanly, "Congestion measures and capacity constraints in spread spectrum networks," in *Multiaccess, Mobility and Teletraffic for Personal Communications*, (Eds. B. Jabbari, P. Godlewski, X. Lagrange), Kluer, Boston, MA, pp. 321-336, 1996.

[93] R. Knopp and P. Humblet, "Channel control and multiple access," in *Multiaccess, Mobility, and Teletraffic for Personal Communications*, (Eds. B. Jabbbari, P. Godlewski, X. Lagrange), Kluwer, Boston, MA, pp. 321-336, 1996.

[94] R.G. Gallager and M. Medard "Bandwidth scaling for fading channels," *IEEE International Symposium on Information Theory*, Ulm, Germany, June, 1997.

[95] A. Ephremides, J. Wieselthier, and D. Baker, "A design concept for reliable mobile radio networks with frequency hopping signaling," *Proceedings of the IEEE*, Vol. 75, pp. 56-73, 1987.

[96] A. Ephremides, P. Varaiya, and J. Walrand, "A simple dynamic routing problem," *IEEE Transactions on Automatic Control*, Vol. 25, pp. 690-693, 1980.

[97] J. Wieselthier, C. Barnhart, and A. Ephremides, "Ordinal optimization of admission control in wireless multihop integrated networks via standard clock simulation," *NRL Tech. Report 5521-95-9781*, August 1995.

[98] D.N. Tse, *Variable-rate lossy compression and its effects on communication networks*, Ph.D. Dissertation, LIDS-TH:2269, MIT, Cambridge, MA, September 1994.

[99] C.E. Shannon, "Memory requirements in a telephone exchange," *Bell System Technical Journal*, vol. 29, pp. 343-349, 1950.

[100] V.E. Beněs, "On rearrangeable three-stage connecting networks," *Bell System Technical Journal*, vol. 41, pp. 1481-1492, 1962.

[101] L.A. Bassalygo and M.S. Pinsker, "Complexity of an optimum nonblocking switching network without reconnections," *Probl. Peredachi Inform.* (English translation in *Problems of Information Transmission*) vol. 9, no. 1, pp. 84-87, 1973.

[102] G.A. Margulis, "Explicit constructions of expanders," *Probl. Peredachi Inform.* (English translation in *Problems of Information Transmission*) vol. 9, no. 4, pp. 71-80, 1973.

[103] O. Gabber and Z. Galil, "Explicit constructions of linear-sized superconcentrators," *Journal of Computer and System Sciences*, vol. 22, pp. 407-420, 1981.

[104] N. Pippenger, "Telephone switching networks," in *The Mathematics of Networks*, Stefan A. Burr, Ed., Proceedings of the Symposium on Applied Mathematics, vol. 25, American Mathematical Society, pp. 101-134, 1982.

[105] A. Lubotzky, R. Phillips, and P. Sarnak, "Ramanujan graphs," *Combinatroica*, vol. 8, no. 3, pp. 261-277, 1988.

[106] M. Ajtai, J. Komlos and E. Szemeredi, "An $O(N \log N)$ sorting network," *Proceedings of the 15th Annual ACM Symposium on the Theory of Computing*,pp. 1-9, ACM, 1983.

[107] M. Sipser and D.A. Spielman, "Expander codes," *IEEE Trans. on Information Theory*, vol. 42, pp. 1710-1722.

[108] B. Hajek and R.L. Cruz, "On the average delay for routing subject to independent deflections," *IEEE Trans. on Information Theory*, vol. 39, pp. 84-91, January 1993.

[109] *National Technology Roadmap for Semiconductors*, Semiconductor Industry Association, (URL www.sematech.org), 1997.

[110] S. Lin and D.J. Costello, *Error Control Coding: Fundamentals and Applications*, Prentice-Hall, 1983.

[111] C.E. Shannon, "The wonderful world of feedback," First Shannon Lecture, *IEEE International Symposium on Information Theory*, Ashkelon, Israel, 1973.

[112] R.G. Gallager, *Information Theory and Reliable Communication*, Wiley, 1968.

Anthony Ephremides received his B.S. degree from the National Technical University of Athens (1967), and M.S. (1969) and Ph.D. (1971) degrees from Princeton University, all in Electrical Engineering. He has been at the University of Maryland since 1971, and currently holds a joint appointment as Professor in the Electrical Engineering Department and the Institute of Systems Research (ISR). He is co-founder of the NASA Center for Commercial Development of Space on Hybrid and Satellite Communications Networks established in 1991 at Maryland as an off-shoot of the ISR.

He was a Visiting Professor in 1978 at the National Technical University in Athens, Greece, and in 1979 at the EECS Department of the University of California, Berkeley, and at INRIA, France. During 1985- 1986 he was on leave at MIT and ETH in Zurich, Switzerland.

He was the General Chairman of the 1986 IEEE Conference on Decision and Control in Athens, Greece and will be the Technical Program Chairman of the 1999 IEEE INFOCOM and the 2000 IEEE International Symposium on Information Theory. He has served as Associate Editor for communication Networks of the IEEE Transactions on Information Theory (1988-1991). He has also been the Director of the Fairchild Scholars and Doctoral Fellows Program, an academic and research partnership program in Satellite Communications between Fairchild Industries and the University of Maryland. He won the IEEE Donald E. Fink Prize Paper Award (1992). He has been the President of the Information Theory Society of the IEEE (1987), has served on the Board of Governors of that Society for many years and has been a member of the Board of the IEEE (1989 and 1990).

Dr. Ephremides' interests are in the areas of communication theory, communication systems and networks, queueing systems, signal processing, and satellite communication.

His email address is: tony@eng.umd.edu.

Bruce Hajek (M'79-SM'84-F'89) received a B.S. in Mathematics and an M.S. in Electrical Engineering from the University of Illinois in 1976 and 1977, and a Ph. D. in Electrical Engineering from the University of California at Berkeley in 1979. He is a Professor in the Department of Electrical and Computer Engineering and in the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign, where he has been since 1979. He served as Associate Editor for Communication Networks and Computer Networks for the IEEE Transactions on Information Theory (1985-1988), as Editor-in-Chief of the same Transactions (1989-1992), and as President of the IEEE Information Theory Society (1995). His research interests include wireless and highspeed communication networks, stochastic systems, combinatorial and nonlinear optimization and information theory.