#111 (12)

Spatial Reuse in Multihop Packet Radio Networks

LEONARD KLEINROCK, FELLOW, IEEE, AND JOHN SILVESTER, SENIOR MEMBER, IEEE

Invited Paper

Multihop packet radio networks present many challenging problems to the network analyst and designer. The communication channel, which must be shared by all of the network users, is the critical system resource. In order to make efficient use of this shared resource, a variety of channel access protocols to promote organized sharing have been investigated. Sharing can occur in three domains: frequency, time, and space. This paper is mostly concerned with sharing and channel reuse in the spatial domain. A survey of results on approaches to topological design and associated channel access protocols that attempt to optimize system performance by spatial reuse of the communication channel is presented.

I. Introduction

The complete design space for packet radio networks is indeed quite complex, as can be seen by scanning this and other papers in this issue. Although there are many results for fully connected packet radio networks (in which all nodes can directly communicate with one another), the analysis of mutihop networks (in which some nodes act as store-and-forward repeaters since not all nodes can directly communicate), is much more complex. Rarely do we have closed-form solutions, and thus optimization of these networks is not an easy task.

In tuning the performance of these networks, we are trying to optimize the use of the scarce system resource—the channel. The key to an effective resource-sharing scheme is the ability to share or reuse the resource as efficiently as possible. There are three main domains in which we can achieve channel sharing: frequency, time, and space.

Frequency reuse and its management can be as straightforward as FDMA or can use more modern approaches such as CDMA and spread spectrum [20], [33]. Unless otherwise stated, we assume that a single frequency is assigned for use by nodes in the network and no spread-spectrum technique is used, and so we only have to be concerned with time and space allocation.

Time reuse and its management is usually referred to as the channel access protocol [12]; see [31] for a good survey or [32] for some comparative studies (by simulation) for radio networks. We can classify channel access protocols in terms of how much knowledge about the system state they use in making decisions. The simplest protocols (e.g., ALOHA) require no system state information to operate; more complex protocols (e.g., CSMA), which tend to have better performance, typically require additional state knowledge (such as the state of the neighbors, obtained by listening to the channel for CSMA). We can expect that additional (say, global) information would allow even better performance to be obtained if, of course, there were: i) some way of obtaining this information and ii) fast algorithms available to make use of the information and decide who should transmit in real time. Later in this paper we discuss some algorithms that select sets of nodes to transmit based on global information. Since this information is typically not possible to collect in a real network, we should think of these algorithms as providing bounds on network performance.

This leaves spatial reuse–the main focus of this paper. The key point in spatial reuse is that when a node is transmitting in some part of the network, it is possible to reuse the same frequency and time in another part of the network with no (significant) interference. This effect is due to propagation loss or the fact that one part of the network is shielded from other areas (perhaps by natural obstructions such as hills or the nature of the radio wave propagation). A simple model to represent a packet radio network is a (directed) graph in which the nodes correspond to the transmitters and a link is present from A to B if node B is able to successfully receive messages from A by direct communication (we will typically assume that for a link to be operational, there must be a link from B to A as well as a link from A to B, i.e., the link must be bidirectional). Nodes that are too distant to communicate directly have no link joining them. From this graph we can easily determine if

Manuscript received June 17, 1986; revised September 5, 1986. The work of L. Kleinrock was supported by the Defense Advanced Research Projects Agency of the U.S. Department of Defense under Contract MDA903-82-C-0064. The work of J. Silvester was supported by the Army Research Office (ARO) under Contract DAAG29-84-K-0084.

L. Kleinrock is with the Computer Science Department, University of California at Los Angeles, Los Angeles, CA 90024, USA.

J. Silvester is with the Communications Sciences Institute, Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089-0272, USA.

0018-9219/87/0100-0156$01.00 © 1987 IEEE

two links can be used simultaneously (on the same frequency and in the same time slot) without interference. Of course, this graph gives a very simplistic view of interference; we discuss more precise models later in the paper.

Several topological problems immediately suggest themselves to the network analyst: i) Given a specific topology, how should the set of links to be used (in each time slot) be selected? ii) What is a desirable topology? This can be split into two questions: iia) If the option of locating nodes is available, where should the nodes be located? and iib) If the node locations are given, how should transmission power, directionality of the transmission, routing, etc., be controlled? In the rest of this paper we survey work that has attempted to answer these questions.

II. Models and Assumptions

Perhaps the most difficult aspect of comparing work on packet radio networks is identifying a common set of assumptions—not only those assumptions made for the tractability of the mathematics, but more importantly those of system operation. The reason that so many different sets of assumptions (and indeed opinions) abound is that an exact description of network operation for a generic system is not possible (nor necessarily meaningful), and so each author tends to focus on those aspects of the problem that he considers important to the particular research he is pursuing. In this section, we have tried to collect together the main assumptions that are used in papers on spatial reuse.

A. Protocols

Many different channel-access protocols are possible. Most of the work that we describe assumes that the network operates in a slotted mode, and uses (Slotted) ALOHA or TDMA as the access protocol. This is not because these protocols are necessarily considered the best, but rather because an attempt is being made to reduce the design space and focus on simple access protocols with the hope that simpler system level models to investigate the more global topological and spatial reuse issues will result.

B. Channel Model

The first set of assumptions is in regard to the operation of the communication channel itself. The main issues for determining network performance and hence investigating spatial reuse are: i) can node A communicate to node B (and at what error rate)? and ii) if node A is sending a message to node B, does C's transmission interfere with the AB transmission (and how is this impact measured in terms of the error rate)? The answer to these questions depends on the assumptions made with regard to the following:

1) *Radio Propagation Effects:* Much of the work on packet radio networks assumes a very simple model for radio propagation. Given a certain transmission power, it is assumed that if two nodes, A and B are closer than a certain distance (transmission range), they are able to communicate (without error). In addition, a transmission by A, say, will interfere with the reception of a third party's transmission at node B (causing an error with probability one). Thus A is able to communicate with B without error if there are no other nodes within a distance r of B that transmit during the transmission from A to B, where r is the distance limit of a transmitter (farther than which, the transmission is not successful due to background noise). If, however, one of B's neighbors does transmit, then the transmission from A to B is destroyed (with probability one). It is thus assumed that in order to communicate with a more distant node, either a relay should be used or the transmission power should be increased so that the intended destination is now within range. It is also typically assumed that all nodes transmit with the same power (range). The main criticism of this model is that it does not take relative distances into account, for example, a nearby interfering node will cause more damage than a distant one. Recently, Sousa and Silvester [28] have taken a different approach. They consider that all links in the network potentially exist, but that the *error rate* on a link is determined by the length of the link and the current activity and location of nodes in the vicinity of the receiver (actually, the activity of all nodes in the network is included, but the contribution of distant nodes is minimal). All links are thus probabilistic in this model.

2) *Directionality:* It is also typically assumed that the antennas are omnidirectional, even though there is much to be gained by using directional antennae [4].

3) *Interference:* As noted above, it is typically assumed that interference is binary in nature, i.e., the packet is destroyed if any other transmission within a critical region occurs and is successful if no other transmission coincides with the transmission of interest. Since all nodes are using the same frequency, it is assumed that nodes cannot transmit and receive at the same time.

4) *Capture Effects:* Capture (in the narrow-band context) is the ability of a receiver to correctly receive the strongest of several interfering signals. Most of the papers on packet radio network performance assume no capture—there are a few notable exceptions, however. The original paper on Slotted ALOHA, [21], also contained analysis for radios with capture. Fratta and Sant, [5], also studied this problem for both centralized and distributed fully connected networks. Later in this paper we present some work which looks at spatial reuse when capture occurs [18], [30].

This issue becomes even more complex, (ie., more choices) when we look at the use of spread-spectrum signaling for packet radio [9]. See [20] in this issue for additional discussion of the spread-spectrum issues.

1) *Multiple Simultaneous Receptions:* One of the major benefits of spread spectrum is the fact that the desired signal can be decoded in the presence of other (hostile or friendly) interfering signals. From a multiple-access viewpoint, this means that we may have several transmissions overlapping in time, space, and frequency that are all successful. In papers addressing network performance, this is typically modeled by assuming that some threshold number of transmissions can simultaneously occur before the noise level becomes too high for successful decoding of the packets. The threshold model does not tell the whole story, however; many researchers have worked on more detailed models of multiuser interference in spread-spectrum networks, references to which can be found in the paper by Pursley in this issue.

2) *Capture Effects (revisited):* Capture in a spread-spectrum system has a different connotation than for FM. It is the ability of the *first* signal to arrive to a receiver to be

"locked-on to," so that later arriving signals can be rejected (and appear as noise).

3) *Simultaneous Transmission/Reception:* For certain types of spread-spectrum systems it may be possible for nodes to transmit and receive at the same time, which is different from the narrow-band case.

Although most of the results on packet radio have been derived for narrow-band systems, many can (and indeed have been) generalized to the spread-spectrum case. The specific assumptions that are made depend on the details of the spread-spectrum system in use (modulation, coding, protocol, etc.), and are beyond the scope of this paper.

C. Topological Structure

Even if we have a detailed understanding of the channel (model), in order to determine the *network topology*, i.e., which nodes can communicate with which others, the locations of the nodes must be specified. There are three prevalent assumptions that are used here:

1) *Regular Structure:* Nodes are considered to be located in some regular pattern on the (n-dimensional hyper-) plane, or on the vertices of a regular polyhedron. The advantage of this approach is that we can make assumptions that all nodes are statistically equivalent, which greatly simplifies the (analytical) problem.

2) *Continuum of Nodes:* In this model, nodes are considered to be continuously present throughout the space of interest (typically the infinite plane). This continuum of nodes is then considered to generate traffic at some rate per unit area. The advantage here is that we can assume the existence of a node at any convenient location.

3) *Random Locations:* Nodes are considered to be randomly distributed in the space of interest (again typically the infinite plane or some finite subset). This complicates matters in that assumptions about homogeneity are a little harder to swallow, but the Poisson assumption often allows closed-form solutions. For this case we can define λ to be the density of nodes per unit area. Then the probability of finding k nodes in a region of area A is

$$\Pr\,[k \text{ in } A] = \frac{(\lambda A)^k}{k!}\, e^{-\lambda A}.$$

It is often more convenient to work with the expected number of nodes in a transmission range of size r. We note that this corresponds to the expected number of neighbors or average degree, and we represent it by d

$$d = \lambda \pi r^2$$

for a network embedded in the plane.

These topological models are chosen mostly for their analytical convenience, and we are led to ask how well they represent "real networks." It is quite difficult to find studies of network structure for real networks in the open literature. Few measurements have been made since most of the networks with capabilities similar to those that we are discussing in this paper are in the experimental or early development phase. Fortunately, a good model for predicting network connectivity exists, called the Longley–Rice model [16], which incorporates effects of irregular terrain—an item which is completely ignored in all analytical network performance studies. This model has been built into a tool for studying network connectivities in different locations. Networks generated using this model, [22], show a close structural resemblance to networks generated using a Poisson distribution of nodes and fixed transmission radius, as described above.

D. Traffic Models and Analytical Approaches

Whenever discussing network performance, we must be careful to specify the traffic model being used. Although this is covered in more detail in other papers in this issue, we present the key aspects that pertain to models used for spatial reuse.

1) *Traffic Models:* When we turn to look at the traffic models that have been used we again find a wide variety of assumptions, which correspond to the analytical approach being used. The important parameters are as follows:

a) *Arrival Process:* Arrivals are typically assumed to be Poisson for continuous time protocols and Bernoulli for slotted systems. Many of the analyses are concerned with throughput only and assume the "heavy traffic model" as discussed below. For this scenario there is no arrival process, since it is assumed that nodes are always busy. In this case, a transmission control scheme is usually employed so that the arrivals to the channel are Poisson or Bernoulli.

b) *Nodal Buffering:* Performance analysis of protocols for fully connected networks have typically assumed that nodes are unbuffered (i.e., have a single buffer, which holds the packet that is currently being transmitted). This allows the state description to be just the number of backlogged nodes (having something to transmit). While this is a satisfactory model for single-hop communication, it is inadequate for multihop networks, so infinite buffering is usually assumed. If we are only concerned with throughput, infinite buffering ensures that we have a conservative system, i.e., there is no loss due to overflows, and that no deadlocks exist. When we turn to delay, the fact that we must model the buffers leads us into great difficulties, since the state description becomes multidimensional and the associated Markov chains have no closed-form solutions.

c) *Packet Length Distributions:* For slotted systems, packets are usually assumed to be of fixed length equal to the slot size. For continuous time systems, packet links are usually assumed to be exponentially distributed (see [33]).

d) *Traffic Matrix:* In discussing the performance of *multihop* networks, the traffic matrix will have significant impact. For most studies, a uniform traffic matrix is assumed, i.e., for a total network traffic level of γ, a fraction $\gamma/n(n-1)$ is the traffic from any particular node to any other. Some (throughput) models are concerned only with the *local throughput*, which corresponds to the rate at which packets succeed over one hop, see [26] for example. (To obtain network throughput, this must be divided by some measure of the average path length.) Since worrying about specific topologies and traffic matrices is too complex, a common assumption is that the load on all nodes is homogeneous. Local throughput is then computed and divided by the network average path length to obtain an estimate of the global network average throughput. For real networks, the center of the network will see more traffic than the edge (since the shortest paths are more likely to traverse the network center) and thus the homogeneity assumption tends to overestimate performance. This fact was noted for regular networks in [27] and an estimate of the impact of the homogeneity assumption for random networks is made by Hajek, [6], where he finds that the center of the network sees 2.2 times as much traffic as the average loading.

2) Traditional Analytical Approaches: In this section, we give an overview of the common approaches to analyze the performance of multihop packet radio networks. The approaches can be classified as follows:

a) *Exact solution:* As noted elsewhere, the exact solution is not available and we are constantly forced to use approximations, such as homogeneity, independence, etc. Exact solutions can occasionally be used for very simple network topologies (see [23] for example).

b) *Abramson:* The usual modeling approach is based on that used by Abramson in his original ALOHA paper [1] and Roberts [21] in the Slotted ALOHA version. Each node schedules transmission according to a memoryless process and acts independently from other nodes. Of course, the success or failure of any particular transmission is dependent on the state of other nodes in the network. For fully connected (unbuffered) networks, the system state can be represented by the number of busy nodes. In order to extend this to multihop networks and model store-and-forward traffic, approximations must be used. A typical assumption is that all nodes are statistically identical.

c) *Boorstyn's approach* [3] (discussed in detail elsewhere in this issue [33]): While this is a powerful numerical technique for evaluating the throughput of some complex multihop systems, the analysis depends upon an independence assumption which can lead to results far from the correct results when the true dependence is strong (later, we describe some cases of synchronized systems in which such is the case). Since the technique is numerical, we do not have any analytical expressions that can be used for (analytical) experimentation with topology, routing, etc. Furthermore, it does not give any direct results for delay.

Typical studies of spatial reuse follow the approach of Abramson, i.e., they consider a slotted network operating under heavy traffic and assume that node i transmits in each slot with equal probability p_i. Homogeneity assumptions regarding the network topology are also made (whether it is random, regular, or a continuum).

3) Novel Analytical Approaches: As noted above, most of the tractable analytical approaches require the use of strong assumptions and approximations. Consequently, we would welcome more powerful analytic tools. Unfortunately, this is a problem of coupled queues, whose general solution is known to be out of reach at the present time. Part of the difficulty comes from the fact that the approach one usually takes is to attempt a solution for the detailed behavior of each node in the network, when one really wishes only global results (e.g., the total system throughput).

Fortunately, Yemini [36] has suggested an approach for evaluating the system behavior of interacting queues which takes a macroscopic view. Specifically, he recommends the use of methods from statistical mechanics. Using this approach, he defines what is known as the "interaction potential" to solve for the equilibrium behavior directly. The model he uses is one in which the interference graph of a multi-access system (i.e., two nodes—terminals—are connected by an edge in the graph if the transmission by one can potentially cause interference with the other); the interference may be a collision (as in a multihop packet radio system) or blocking (as in a finite-storage Jackson network). An idle (or busy) node remains so for an exponentially distributed amount of time with rate λ (or μ). The equilibrium probability $p(A)$ that a set A of nodes is busy is given by

$$p(A) = \begin{cases} \rho^{|A|/z}, & \text{for } A \in C \\ 0, & \text{otherwise} \end{cases}$$

where C is the set of cliques (defined more precisely in Section III-A below), $\rho = \lambda/\mu$, and z is given by

$$\begin{aligned} z &= \sum_{A \in C} \rho^{|A|} \\ &= \sum_i \alpha_i \rho^i \end{aligned}$$

and α_i is the number of cliques with i nodes. $z = z(\rho)$ is known as the "partition" function of this interference and has an analogy with the partition function found in statistical mechanics models of mechanical systems. The partition function provides a complete description of allowable transmissions among the cliques, and this may be used to derive the equilibrium behavior of the system.

Yemini then goes on to draw the analogy to statistical mechanics and finds the equivalence in the interacting queues model to the thermodynamic quantities of energy, pressure, volume, temperature, and entropy. For example, the energy of a "microstate" corresponds to the size of the clique set. The global energy of the system corresponds to the system throughput. The pressure corresponds to a measure of the average rate of blocked transmissions. He applies these ideas to a number of diverse examples, yielding the system performance directly.

Whereas this approach has its own problems, and one can challenge some of the assumptions necessary for the model, it appears to be a fresh new direction to obtaining the macroscopic behavior of complex systems of interfering queues directly.

III. Surveying the Analytical Results

In multihop packet radio networks, we are faced with a number of very difficult analytic and design problems. Among these outrageous networking problems we include the selection and optimization of the channel-access method, the determination of appropriate transmission ranges and topological structure, and the design of the routing procedure. The studies presented in this section focus on two main research areas: i) attempts to optimize the access protocol or find throughput bounds over all access protocols; and ii) attempts to optimize network performance by modification of the network structure.

A. Optimizing the Access Protocol

All queueing analysts pray each night for the world to be *deterministic*. This is because queues are formed in systems where the arrival and service times are random.

Determination (or even definition!) of an optimal channel access algorithm is not so simple for multihop radio networks with arbitrary topology, as discussed by Silvester,

[25], with regard to optimal frequency or time slot assignment, and also in the paper by Nelson and Kleinrock [17], discussed in more detail below. As noted above, though, we expect the optimal algorithms to schedule transmissions in a deterministic manner. This is the topic of this section.

By using deterministic scheduling, we expect to eliminate most of the queues that form. However, in addition to randomness, there is another source of queueing which comes about due to overloads! The overloads we have in mind do not persist for long intervals, but rather, they occur for short bursts (the system is designed to be able to keep up, on the average).

We begin with the results reported by Nelson and Kleinrock [17], in which they studied a packet radio network where the nodes are randomly distributed over the infinite plane according to a two-dimensional Poisson point process. The average number of nodes within range of each node is d, and the network is assumed to be connected. Using a deterministic scheduling method, they found the maximum number of simultaneous transmissions that could be supported with no collisions, and from this they found the probability of a successful transmission. This probability forms an upper bound on all possible access protocols which use neither the knowledge of the directionality nor of the location of any nodes. The basic approach to this analysis is to recognize that a set of nodes may simultaneously transmit with no collisions if they are all mutually at least three hops away from each other. For a given Poisson density, one can construct an ideal layout for those nodes which would be allowed to transmit simultaneously by placing them on the vertices of equilateral triangles; these triangles tesselate the infinite plane, as shown in Fig. 1. The sides of the triangles have a length equal to

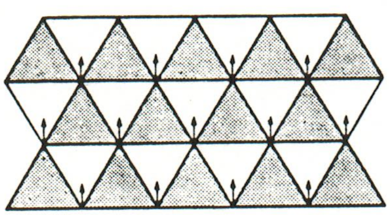

Fig. 1. Mapping of vertices to triangles.

$\overline{X}$, the average distance between nodes which are three hops apart in the original random network. In this idealized topology, each vertex of the tesselation corresponds to a transmitting node. It is shown that $\overline{X} = 2$ (under the assumption that the transmission radius of each packet radio node is $r = 1$), an intriguingly simple result! Denoting the expected fraction of successful transmissions per node in the packet radio network by $f(d)$, it is then shown that

$$f(d) = 0.9069/d$$

where we recall that d is the average number of nodes within the transmission range of any node. This is the success rate for our ideal protocol which maximizes the number of transmissions that can take place simultaneously with no collisions. This upper bound on throughput was compared in [17] to the original random network by generating a large number of such random networks of different densities d, and finding the maximum number of nodes which were mutually a distance of at least three hops from each other; the comparison was excellent and the upper bound was found to be rather tight.

Since $\overline{X}$, the average distance between nodes three hops apart in the random graph, is $\overline{X} = 2$, we see that this gives us an ideal topology in which each node is surrounded by a unit circle tangent to the corresponding circles of its six neighbors, as shown in Fig. 2. This corresponds very nicely

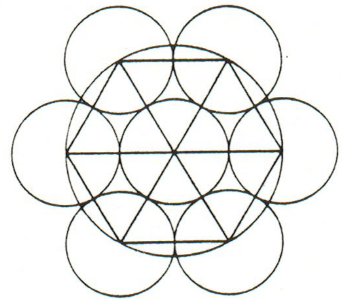

Fig. 2. Transmission areas.

with the known result that such a topology of tangent circles optimally covers the plane with fixed radius circles in which no region may be doubly covered. Each unit circle corresponds to a node at its center which successfully transmits a packet.

In [17], the authors then consider any possible improvement in success rate that comes about if the circles are allowed to overlap slightly. This does give a small improvement and results in the new bound

$$f'(d) = 0.9278/d.$$

The results for $f(d)$ and $f'(d)$ show that the optimal (per hop) success rate for a random packet network is

$$f(d) \cong 1/d$$

this last being a simple upper bound. We note that this tells us that the average traffic which should be generated in a "region" (i.e., within a unit circle which is the range of a packet radio unit at the center of the circle) is approximately one packet per slot; this result has been seen before in a number of contexts, for example, in [35] and indeed in [1].

An interesting comparison of this success rate to that of Slotted ALOHA and to CSMA is made in the paper. It is shown that Slotted ALOHA is only 39.6 percent as efficient as our optimal protocol, and the CSMA is only 48.5 percent as efficient. We note that CSMA's performance in this multihop environment is far inferior to its performance in a single-hop environment (where it achieves efficiencies in the vicinity of 90 percent). This is perhaps not so surprising, since a multihop environment automatically implies that some radios are "hidden" from others (in that they cannot hear some others), and CSMA is known to degrade badly in the presence of even a small fraction of hidden terminals.

In a related paper, Nelson and Kleinrock [19] evaluate the performance of a deterministic, collision-free, scheduling algorithm for a packet radio network in which node locations are fixed and known. The protocol is referred to as Spatial TDMA. This access protocol is suggested as an alternative to wire-line systems in which line capacities are not at all easily changed as traffic demands change; indeed, the ability to dynamically alter the effective line capacity between two nodes simply by changing the assignment of slots within the TDMA frame is offered as an interesting advantage over wire-line nets.

The Spatial TDMA protocol operates as follows. A (simplex) *channel* is a directed pair (of nodes), for which communication can take place from the tail node to the head node. If bidirectional radio connectivity exists then two channels, one in each direction, are used. A *clique* is then defined as a set of channels, such that communication on any channel of the clique does not interfere with any other channel of the clique; that is, a clique permits all of its members to simultaneously transmit successfully. A maximal clique is a clique to which no additional pairs can be added. A clique cover is a set of maximal cliques which collectively contain all directed radio pairs in the network. We assume that the channel is synchronized into slots, of length equal to a packet transmission time. Spatial TDMA defines a repeating frame which contains a fixed number of slots, with each slot being assigned to a unique clique in the clique cover. A slot assigned to a clique may be used for transmission by any (and all) pairs in that clique, and only by those in that clique. (A pair is said to use the slot if a transmission is sent from a node at the tail of the direction arrow to the node at the arrow head.) Spatial TDMA for multihop packet radio networks is a generalization of TDMA for single-hop networks. Thus the capacity assigned to any one of the directed radio pairs is equal to the sum of all slots assigned to all cliques of which that pair is a member.

The authors develop an approximate analysis for the mean response time of a newly generated packet to make its way one hop into the network. For this purpose, they use the fluid approximation [12] to evaluate the backlog at the channel connecting each directed radio pair. From the TDMA structure, it can be seen that a given channel has three modes:

i) It will be in the *input* mode whenever the TDMA assignment allows packets to arrive from neighboring nodes to the tail node of the channel which must be transmitted over that channel (at some later time).

ii) It will be in the *service* mode whenever any clique of which it is a member is assigned a slot by the TDMA protocol.

iii) It will be in the *idle* mode when neither of the above cases is true.

When in the receiving mode, it is capable of receiving traffic from certain of its neighboring radios which are within range. During the transmitting mode, it is capable of transmitting. During the idle period, it is capable of neither action. However, during the entire frame (i.e., at all times), it is capable of receiving (external) traffic from its locally attached Host.

The authors assume that both the internal and external traffic are Bernoulli. They use a fluid approximation [12] to evaluate the mean response time for this channel. That is, the backlog grows at a rate (packets per slot) equal to the internal plus external Bernoulli rates during the input mode. During the service mode, the backlog grows at a rate equal to the external Bernoulli rate minus one, and during the idle mode, the backlog grows at the external Bernoulli rate. An example is shown below, in Fig. 3(a) and the corresponding backlog is shown in Fig. 3(b).

From these considerations, one easily calculates the average backlog which, when divided by the average number

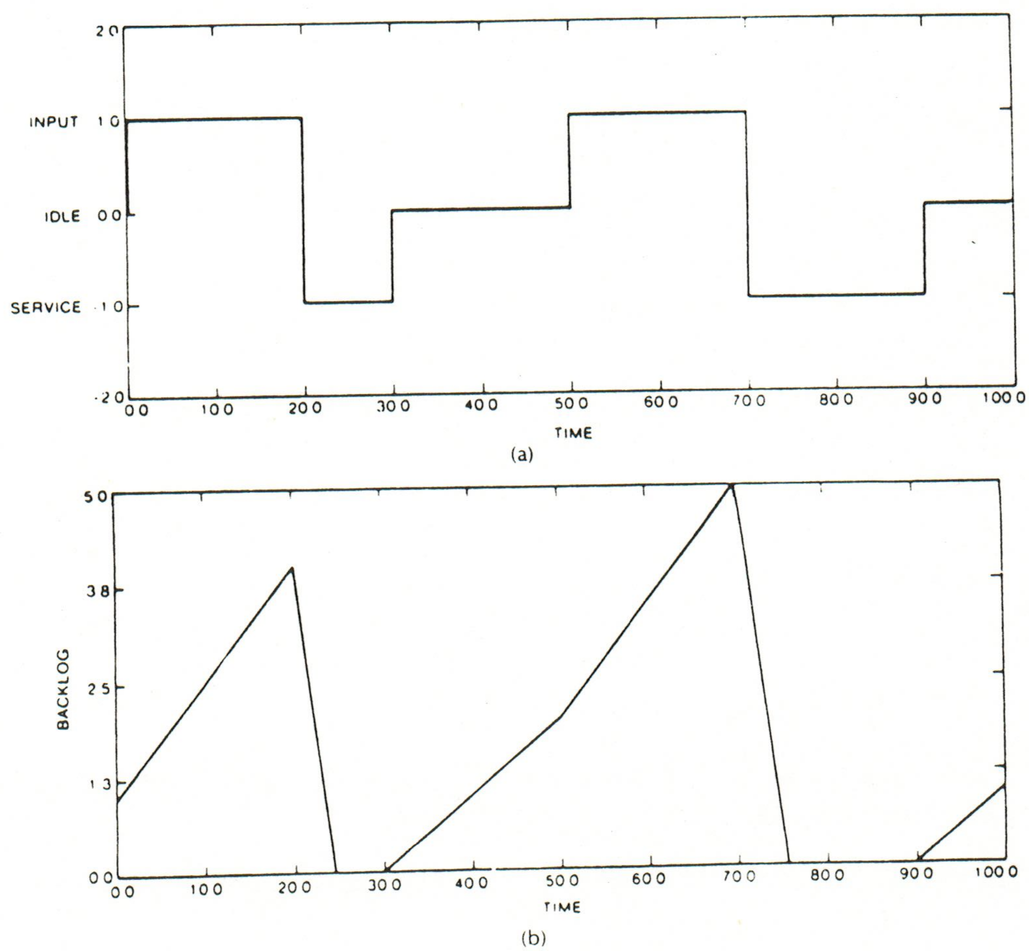

Fig. 3. (a) An example time frame. (b) The backlog for the time frame shown.

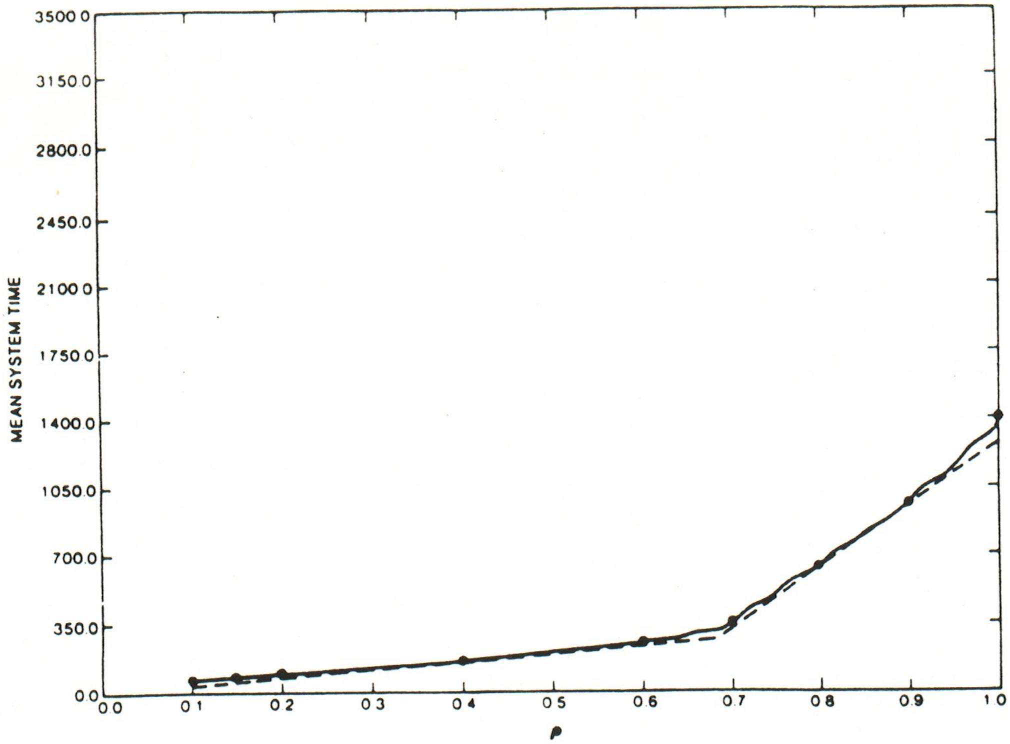

Fig. 4. Mean time in system.

of arrivals to that channel in a frame, yields the mean response time (Little's Result [10]). It can be seen that the order in which the input, service, and idle modes are distributed in a frame will greatly influence the mean response time. The optimum ordering is to chop the input, service, and idle periods into infinitesimal intervals which alternate with each other in that order; this will minimize the mean response time.

An example plot of the response time versus the utilization of the channel is shown in Fig. 4. Here we see that the response time exhibits a piecewise-linear behavior which is explained in terms of the distribution of the various modes in a frame. Observe that the fluid approximation is predicting delays since the channel goes through periods of overload during the frame cycle. The analytic performance (solid lines) were compared to simulation results (dotted lines) which were obtained from a single Spatial TDMA node whose frames were randomly generated with the same parameters as those used in the approximation; the fit is amazingly good.

The authors also provide an algorithm for approximating the optimum capacity assignment to the many channels in the network. The problem of determining an optimum clique cover is related to generalized graph coloring, see [25], and has yet to be solved.

Additional results are available in the literature which support the observation that determinism reduces queues and waiting times. For example, Hajek [7] has shown that if one wishes to send a fraction p of arrivals from an arbitrary renewal process (i.e., interarrival times which are independent and identically distributed) into a queueing system with an exponential server, then the optimum way (i.e., the way to minimize the mean response time) is to split this stream of arrivals "deterministically." That is, one should choose the most regular sequence of customers in the arriving stream to feed into the queueing system. For example, if $p = \frac{1}{2}$, then every other customer should be allowed in. More specifically, the nth customer should be allowed in if the following expression equals unity:

$$\lfloor (n+1)p \rfloor - \lfloor np \rfloor$$

where $\lfloor x \rfloor$ is the greatest integer less than or equal to x. This result is not unlike that which was found for Spatial TDMA since, in that system, one attempts to spread the receiving mode out as uniformly (i.e., as regularly) as possible in a cycle.

Deterministic behavior corresponds, in a real sense, to synchronized behavior. The more synchronized a system is, the more regular, or deterministic, it is. A result exhibiting the advantages of synchronization may be found in [34]. In this work, the system studied was a tandem network of N nodes, as shown in Fig. 5. The interesting thing about

Fig. 5. A tandem network.

this tandem is that it is a model of a simple multihop system. In the figure, we see that all the traffic enters at node N and is destined for node 1. The model assumed that time was slotted, with slot size equal to the transmission time of the fixed-size packets. Heavy traffic was assumed, in that all nodes were assumed to always have a packet to send. The access (or transmission scheme) was taken to be Slotted ALOHA; node i transmits a packet in a slot with probability p_i, independently of all other nodes. Propagation time was assumed to be zero, and acknowledgments were available immediately and at no cost. Node i hears only its neighbors: nodes $i-1$ and $i+1$ (except that nodes 1 and N each hear only one neighbor). All traffic moves down the chain from node N to node 1. For this model, the nodal probabilities p_i were found which maximized the system throughput; for most practical tandems (of more than four nodes), the maximal throughput is very close to $\frac{4}{27}$. However, if one changes the policy so that each node is rude (i.e., transmits with probability $p_i = 1$ whenever that node has a packet to send), then the throughput goes to $\frac{1}{3}$ (that is, exactly 2.25 times as high). What happens in this rude system is that the tandem gets synchronized and every third node will contain a packet; these packets will march down the tandem in lockstep, arriving at the destination at a rate of one packet every three slots. In [15, ch. 6], an attempt was made to remove the poor assumption of independent transmissions in the first access scheme. The exact analysis is extremely difficult since it leads to a system of N coupled queues (a problem whose analytic complexity was also recognized in [34]. Therefore, an approximation to the behavior of the throughput in the region $p \cong 1$ was developed. This analysis showed that an increase in system throughput occurs even when $p < 1$. Moreover, the independence assumption works well when $p \ll 1$. Once again, we see that synchronous (and, indeed, nearly synchronous) systems exhibit improved behavior; what is especially nice about the tandem is that it synchronizes itself! (Other tandem-like systems with this property are discussed in [34].)

Another interesting system which exhibits "self-synchronization" is described in [15, ch. 7]. In this multihop system, nodes are uniformly spaced along a bus and they share the bus capacity using CSMA in a slotted mode; that is, they are "polite," in that a node will not transmit if it "senses" another node's transmission. The ith node has a probability p_i of sensing the channel in an attempt to transmit. The bus is a "fast" bus, in the sense that the propagation delay between two adjacent nodes is greater than or equal to the packet transmission time (i.e., the ratio of these two, a, is extremely large). For such a fast bus, it is shown that the total system throughput increases when any of the p_i's increase (i.e., the channel is self-stabilizing), and the total throughput can be made to approach one packet per slot! This is a fascinating result, since conventional wisdom argues that the efficiency of CSMA degrades badly as a increases. The reason for the improvement is the synchronization that takes place; when a node succeeds in transmitting a packet, it tends to take over the bus by shutting everyone else out. Thus a node tends to capture the bus and, under heavy traffic, will send a possibly long sequence of uninterrupted, uncollided packets. Once again, synchronism (determinism) provides an improvement in system behavior. Note that this simply increases one node's throughput, but does not scale up the throughput of all nodes.

B. Optimizing the Topology and Routing

In this section we look at papers that are attempting to optimize the performance by modifying the network topological structure: regular location of nodes versus random location; specification of the average number of neighbors that each node has; directional or omnidirectional antennas; route selection.

1) Giant Stepping: Perhaps the first contribution to this area was in an unpublished note [11], which posed the question: "Is it better to take many short hops, or a few long ones?" for a network with a continuum of nodes and the ability to arbitrarily adjust communication range (power). If a small range is used, many short hops are needed but there is little contention for the channel in each hop (only a few other nodes within transmission range of the receiver). If a long range is used, only a few long hops are necessary, but the transmission for each hop must contend with much more traffic.

Nodes are assumed to be distributed in a continuum over the infinite plane. Traffic is generated with a Poisson density λ packets per second per unit area. If r is the transmission range of a node, and D is the (mean) distance requirement from source to destination, then a packet must take

$$\bar{h} = \frac{D}{r}$$

hops, on average, assuming that there is always a node available for reception in exactly the right location (at distance r along the vector from source to destination). The end-to-end delay T is therefore given by

$$T = \bar{h}t(\lambda\pi r^2)$$

where $t(x)$ is the one-hop delay when there is a total traffic level of x contending for the attention of the receiver. Kleinrock shows that the optimal value of r must satisfy

$$\frac{dt}{dr} = \frac{t}{r}.$$

Considering a graph of $t(r)$ against r, this relation is satisfied by taking a ray from the origin and increasing the slope until the ray touches the delay curve $t(r)$. This is an important result. It says that for a given traffic load there is an *optimal range* that should be used for packet transmission to minimize delay.

Taking a simple delay model (M/M/1), the optimum radius to minimize delay r^* is given by

$$r^* = \sqrt{\frac{\mu C}{3\lambda\pi}}$$

where $1/\mu C$ is the mean message transmission time. We note that as the load is increased, i.e., $\lambda \to \infty$, the radius should be decreased toward zero. The implication is that to obtain high throughput, very small transmission ranges should be used.

The general findings of this paper are quite interesting and it has lead to much of the work described below. The main problems with this approach is the assumption that repeaters are available where needed (exactly on the edge of a transmission radius). This is valid for large r, but if we want to maximize throughput, the results indicate that the network will have low average degree. Section III-B3 describes more recent work that has tried to overcome this problem. Another criticism of this paper might be that interference between nodes is not accounted for. This, however, is not really a problem since, as noted in the paper, more complex delay models fit the framework of the model.

2) Regular Structures: In [2] and [27] we find investigations of the performance of multihop packet radio networks havings a regular structure. Akavia and Kleinrock look at planar networks and study the relative performance of the regular tesselations of the plane. Silvester and Kleinrock look at these and other regular topologies, such as loop networks and planar networks of higher degree.

The channel-access protocol used is Slotted ALOHA and the homogeneity assumption is made. Consider a regular topology in some space (linear or planar), with each node having exactly d neighbors. Assuming that a node transmits with probability p in any slot (heavy traffic model), we have that the probability of success (local throughput) for an arbitrary node is

$$s = p(1 - p)^{d-1}$$

(ignoring edge effects). The optimal p value to maximize local throughput is $p = 1/d$, a not unfamiliar result. For a network with n nodes, the total number of successful transmissions or network local throughput, s_{net}, is

$$s_{net} = \frac{n}{d}\left(1 - \frac{1}{d}\right)^{d-1}.$$

This formulation works for any network with regular structure using Slotted ALOHA. In order to find real network throughput and evaluate the effect of the number of neighbors, we divide the network local throughput by the network average path length $\bar{h}$, to account for the multiple transmissions that a message must take en route from source to its final destination. For a one-dimensional (or ring) network, the average path length is linearly proportional to the number of nodes divided by the degree

$$\bar{h} = c\frac{n}{d}.$$

The optimal topology is selected by evaluating the constant of proportionality for each case. Network throughput is independent of the number of nodes in the network and, as shown in [27], is fairly insensitive to the actual average degree used. For a loop network, optimal network throughput is $2/e$ (for large n).

For two-dimensional networks, however, the average path length is proportional to the square root of the number of nodes divided by the average degree

$$\bar{h} = c\sqrt{\frac{n}{d}}$$

with the specific constant of proportionality c depending on the topology. As noted in Section III-B1, low connectivities are the best selection to maximize throughput. The three obvious choices are the 3 regular tesselations of the plane: triangular ($d = 6$), grid ($d = 4$), and hexagonal ($d = 3$). In [2], the authors show that the best selection is the hexagonal tesselation.

In [27], the authors also note that the center of the network is more heavily loaded than the edges and they evaluate the link flows for all links in a (Manhattan) grid network. They then compute the network throughput, assuming that the central node will saturate first and hence determine the network throughput. The total network throughput γ is found to be

$$\gamma = 0.08\sqrt{n}$$

i.e., proportional to the square root of the number of nodes in the network.

3) Random Structures: Having seen that the best topology to maximize throughput in *regular* networks uses the lowest possible connectivity, we are lead to consider *random* networks. There have been a series of papers dealing with the selection of optimal transmission ranges (average degree) for random networks (again using Slotted ALOHA). The first paper to address this topic was [14]. Later papers, [4], [6], [8], [18], [30] represent progressive refinements of the model, consideration of different strategies, and modifications of the assumptions as to how disconnected nodes should be treated.

The approach is to compute the expected progress toward the destination for an arbitrary transmission, assuming a homogeneous traffic environment and Poisson node location. In Fig. 6, node S transmits a packet to node R on its way to the final destination D. The progress Z toward the destination is defined to be

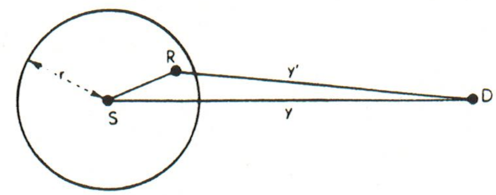

Fig. 6. Progress in an random network.

$$Z = (Y - Y') \cdot \Pr[\text{success } S \rightarrow R].$$

The basic model to find Pr [success] proceeds in a similar fashion to the regular network case, except that we must condition the success probability on the number of nodes in range, since this is now a random variable. Computation of the progress term, $Y - Y'$, varies depending on the assumption made concerning what to do with nodes that are either disconnected or have no neighbors in the right direction that can forward the packet toward the final destination. Following the approach of [30], we have the following expression for the success probability:

$$s = p(1 - p)e^{-pd}(1 - e^{-d})$$

where the first term, p, is the probability that the source transmits, the $1 - p$ term is the probability that the destination is not transmitting, the e^{-pd} is probability of no interference from other nodes, and the term $1 - e^{-d}$ is the probability that some node is in range. The optimal transmission probability is about $1/(d + 1)$, ($d + 1$ to account for the source node—note that this again corresponds to a total environment traffic of unity).

The progress term depends on the specific routing algorithm used. The simplest is called "Most Forward with Fixed Range" (MFR) in which the packet is sent to that node within the fixed transmission range that maximizes the forward progress. For this case, the expected forward progress is evaluated using standard geometrical techniques. Differences between the papers are mostly concerned with how to handle the case where the node has no neighbor in the forward direction.

Kleinrock and Silvester [14] determined that the optimal transmission range should be such that the average degree is 5.89 (about 6) and also noted that the performance is fairly insensitive to using a larger value, but very sensitive to using lower values. This fact was emphasized in simulation studies that were performed [24] in which it was found to be difficult to produce a connected network for small values of the average degree. Later Takagi and Kleinrock [30] improved this model and estimated the optimal average degree to be about 8. Hou and Li, [8], providing a more precise analysis find that 6 is a better estimate of the average degree to optimize throughput for the fixed transmision radius case. Their model does not allow transmissions to result in negative progress toward the destination (in the case where no neighbor exists in the right direction) and they provide a more careful investigation of the regions in which interfering nodes must be located. (Of course, not allowing traffic to go backwards does not resolve what should actually be done in such cases in a real network. One approach is to allow nodes to adjust their transmission range so that connectivity is achieved. This idea is discussed more fully later.)

Hou and Li [8] also introduce two other routing strategies: "Nearest with Forward Progress" (NFP) and "Most Forward with Variable Radius" (MVR). NFP adjusts the transmission radius so that a node is found that allows progress in the desired direction (the first in the desired direction). MVR is similar to MFR, except that once the repeater has been identified, the range is reduced to exactly that needed. NFP is shown to give better performance.

Hajek [6], uses a different measure, efficiency, to study the network performance. The efficiency is the expected progress divided by the area covered by the transmission. Using this measure, he finds that the optimal value for the average degree is about 3. The difference here is that the range is increased until a terminal in the right direction is located. This approach results in an overall increase in efficiency of 85 percent. Chang and Chang, [4] have also studied the use of directional antennas and showed that performance is substantially increased since interference is largely eliminated. The problem with adjustable transmission radii and directional antennas is that the locations of other nodes must be known.

In [30], the authors also consider range optimization for CSMA and find that the improvement over Slotted ALOHA is small (16 percent) due to the hidden terminal problem. The approach is similar to that used for Slotted ALOHA as described above.

In [18], the authors investigate the effect of capture under a slightly different routing algorithm. They model the case where a packet is forwarded to any neighbor in the right direction with equal probability. Takagi and Kleinrock [30] also investigated the effect of capture on their optimal range study described above. They find that throughput is improved by 36 percent for a similar optimal average degree.

Finally, we turn to the paper of Sousa and Silvester, [28], which is a divergence from previous approaches in that they use a propagation model to determine the interference contribution at the receiver due to every other transmitting node in the network. Knowing the interference and the signal strength from the desired transmission (determined using the same propagation model), allows the probability that the signal-to-noise ratio is above the threshold needed for successful reception to be determined, and hence the probability of a successful transmission. The paper describes the model for direct-sequence spread-spectrum network assumptions, but can also be applied to narrowband systems. The optimization is no longer, "at what range should I transmit," rather it is "to whom should I address the transmission" so that the expected progress (distance times probability of success) is maximized. For the spread-spectrum system and assuming that a repeater can be located exactly where needed, the authors find that the optimal strategy is to address the packet such that there are $1.3\sqrt{K}$ terminals between the transmitter and addressed repeater, where K is the multi-user capability of the spread-spectrum system (typically about 100 for a system with a processing gain of 1000). Power adjustment does not affect the performance of the system modeled here, since background noise has been ignored (considered insignificant compared to multi-user interference) and the power of all nodes scale together without affecting the signal-to-noise ratio. If nodes are allowed to *individually* adjust their power, however, network performance can be improved. Sousa and Silvester have investigated some of these capabilities using a graphics modeling package that they have developed [29]. Although this is an interesting abstract model, most spread-spectrum systems use frequency hopping or a combination of direct sequence and frequency hopping, in which case the power of interfering signals may be less important than the number of them.

IV. Concluding Remarks and Directions for Future Work

Let us summarize the most important findings. As we might have expected, determinism helps. By having a deterministic structure (known number of interferers, known locations) the throughput is significantly higher than a random topology (without even considering the troublesome problem of what to do about the disconnected nodes that occur in random networks). Similarly, the use of deterministic scheduling protocols also allows better performance. The problem of course with deterministic structure and schedules is how can they be achieved in an operational network?

By optimizing the network structure, be it random or regular, a *local* throughput which is *linear* with respect to the number of nodes in the network is obtained. When the number of hops that a message must take is taken into account, the *network* throughput is found to be proportional to the *square root* of the number of nodes in the network. In addition, the analysis indicates that small degrees should be used in order to maximize throughput. For *regular* networks, this should be taken to the extreme of letting each node have only three neighbors (i.e., covering the plane with a hexagonal tesselation). For *random* networks a higher degree is necessary to overcome the uncertainty of being able to find a node to act as a repeater that is situated in the right direction. The optimal number of neighbors is in the range of 6–8, depending on how the (difficult) cases of disconnected nodes and no forward progress are handled.

We can identify (but not solve!) several problems for continued study. An obvious one is the study of optimal network structure to minimize delay. Another is the development of operational protocols where each node is allowed to make local decisions as to when to transmit, what power to use, and which node to select as the next node on the path to the final destination as a function of the current traffic loading. Optimal selection of (re-)transmission control parameters in conjunction with range control is an additional problem that has not been solved.

Algorithms to *design* radio networks are lacking. The typical design problem might be: Given a current network (locations of nodes) to satisfy a communication requirement, how much power should each node use (local power control)?; where should additional nodes be deployed to improve network performance?; which nodes should be moved to improve network performance? In summary, although some general design guidelines have been found, optimization of radio networks still presents many challenging problems.

References

[1] N. Abramson, "The ALOHA system—Another alternative for computer communication," in *AFIPS Conf. Proc.* (Fall Joint Computer Conf., 1970), pp. 281–285.

[2] G. Y. Akavia and L. Kleinrock, "Performance tradeoffs and hierarchical designs of distributed packet switching communication networks," Computer Systems Modeling and Analysis Group, Department of Computer Science, School of Engineering, UCLA, Rep. UCLA-ENG-7952, Sept. 1979.

[3] R. Boorstyn and A. Kershenbaum, "Throughput analysis of multihop packet radio," in *Proc. ICC*, pp. 13.6.1–13.6.6, June 1980.

[4] C. J. Chang and J. F. Chang, "Optimal design parameters in a multihop packet radio network," in *Proc. Globecom*, pp. 15.5.1–15.5.5, Nov. 1984.

[5] L. Fratta and D. Sant, "Some models of packet radio networks with capture," in *Proc. Int. Conf. on Computer Communications*, pp. 155–161, Oct. 1980.

[6] B. Hajek, "Adaptive transmission strategies and routing in mobile radio networks," in *Proc. Conf. Inform. Sci. Syst.*, pp. 373–378, Mar. 1983.

[7] —–, "Extremal splittings of point processes," presented at the Statistical Research Conf. dedicated to the memory of Jack Kiefer and Jacob Wolfowitz, Cornell University, Ithaca, NY, July 1983.

[8] T-C. Hou and V. O. K. Li, Transmission range control in multihop packet radio networks," *IEEE Trans. Commun.*, vol. COM-34, no. 1, pp. 38–44, Jan. 1986.

[9] R. Kahn, S. Gronemeyer, J. Burchfiel, and R. Kunzeman, "Advances in packet radio technology," *Proc. IEEE*, vol. 66, no. 11, pp. 1468–1496, Nov. 1978.

[10] L. Kleinrock, *Queueing Systems, Volume I: Theory*. New York, NY: Wiley-Interscience, 1975.

[11] —–, "On giant stepping in packet radio networks," UCLA Packet Radio Temporary Note #5, PRT 136, Mar. 1975.

[12] —–, *Queueing Systems, Volume II: Computer Applications*. New York, NY: Wiley-Interscience, 1976.

[13] L. Kleinrock and G. Y. Akavia, "On a self-adjusting capability of random access networks," *IEEE Trans. Commun.*, vol. COM-32, no. 1, pp. 40–47, Jan. 1984.

[14] L. Kleinrock and J. A. Silvester, "Optimum transmission radii in packet radio networks or why six is a magic number," in *Proc. Nat. Telecommunications Conf.* (Birmingham, AL, Dec. 1978), pp. 04.3.1–04.3.5.

[15] H. Levy, "Non-uniform structures and synchronization patterns in shared channel communication networks," Ph.D. dissertation, University of California, Los Angeles, 1984. UCLA, Computer Science Dep. Rep. 840049, Aug. 1984.

[16] A. Longley and P. Rice, "Prediction of tropospheric radio transmission loss over irregular terrain—A computer method," ESSA, U. S. Dep. of Commerce, NTIS 676-874, ERL-79-ITS67, 1968.

[17] R. Nelson and L. Kleinrock, "Maximum probability of successful transmission in a random planar packet radio network," in *IEEE Infocom 83 Proc.* (Apr. 18–21, 1983, San Diego, CA), pp. 365–370.

[18] —–, "The spatial capacity of a Slotted ALOHA multihop packet radio network with capture," *IEEE Trans. Commun.*, vol. COM-32, no. 6, pp. 684–694, June 1984.

[19] —–, "Spatial TDMA: A collision-free multihop channel access protocol," *IEEE Trans. Commun.*, vol. COM-33, no. 9, pp. 934–944, Sept. 1985.

[20] M. B. Pursley, "The role of spread spectrum in packet radio networks," this issue, pp. 116–134.

[21] L. G. Roberts, "ALOHA packet system with and without slots and capture," *Comput. Commun. Rev.*, vol. 5, pp. 28–42, Apr. 1975 (originally appeared as a Packet Radio Note in 1972).

[22] P. F. Sass, "Connectivity analysis of UHF ground radio networks," Research and Development Tech. Rep. CECOM-TR-84-2, U. S. Army Communications and Electronics Command, Fort Monmouth, NJ.

[23] M. Sidi and A. Segall, "A three-node packet radio network," *IEEE Trans. Commun.*, vol. COM-32, no. 12, pp. 1336–1339, Dec. 1984.

[24] J. A. Silvester, "On the spatial capacity of packet radio networks," Computer Science Department, Computer Systems Modeling and Analysis Group, Tech. Rep. UCLA-ENG-8021, UCLA, Los Angeles.

[25] —–, "Perfect scheduling in multi-hop broadcast networks," in *Proc. 6th Int. Conf. on Computer Communication* (London, England, Sept. 1982), pp. 449–454.

[26] J. A. Silvester and L. Kleinrock, "On the capacity of single-hop Slotted ALOHA networks for various traffic matrices and transmission strategies," *IEEE Trans. Commun.*, vol. COM-31, no. 8, pp. 983–991, Aug. 1983.

[27] —–, "On the capacity of multihop Slotted Aloha networks with regular structure," *IEEE Trans. Commun.*, vol. COM-31, no. 8, pp. 974–982, Aug. 1983.

[28] E. S. Sousa and J. A. Silvester, "Determination of optimum transmission ranges in a multi-hop spread spectrum network," in *Proc. MILCOM*, Oct. 1985.

[29] —–, "A graphical package for experimentation with spread spectrum packet radio networks," in *Proc. MILCOM*, Oct. 1986.

[30] H. Takagi and L. Kleinrock "Optimal transmission ranges for randomly distributed packet radio networks," *IEEE Trans. Commun.*, vol. COM-32, no. 3, pp. 246–257, Mar. 1984.

[31] F. A. Tobagi, "Multi-access protocols in packet communication systems," *IEEE Trans. Commun.*, vol. COM-28, pp. 468–488, Apr. 1980.

[32] F. A. Tobagi and D. H. Shur, "Performance evaluation of channel access schemes in multihop packet radio networks with regular structure by simulation," CSL Tech. Rep. 85-278, Stanford Univ., June 1985.

[33] F. A. Tobagi, "Modeling and performance analysis of multihop packet radio networks," this issue, pp. 135–155.

[34] Y. Yemini, "On channel sharing in discrete-time, multi-access broadcast communication," Ph.D. dissertation, University of California, Los Angeles, 1980. Rep. UCLA-ENG-8061, UCLA School of Engineering and Applied Science.

[35] Y. Yemini and L. Kleinrock, "On a general rule for access control or silence is golden . . .," in *Proc. Int. Symp. on Flow Control in Computer Networks* (Versailles, France, Feb. 1979).

[36] Y. Yemini, "A statistical mechanics of distributed resource sharing mechanisms," in *IEEE Infocom '83 Proc.* (Apr. 18–21, 1983, San Diego, CA), pp. 531–539.

Leonard Kleinrock (Fellow, IEEE) received the B.S. degree in electrical engineering from the City College of New York in 1957 (evening session) and the M.S.E.E. and Ph.D.E.E. degrees from the Massachusetts Institute of Technology, Cambridge, in 1959 and 1963, respectively.

While at MIT, he worked at the Research Laboratory for Electronics, as well as with the computer research group of Lincoln Laboratory in advanced technology. He joined the faculty of the University of California at Los Angeles in 1963. Presently, he is a Professor of Computer Science there, while his research interests focus on local area networks, computer networks, performance evaluation, and distributed systems. He has had over 150 papers published and is the author of five books, the most recent being *Solutions Manual for Queueing Systems, Volume II*. He is co-director of the UCLA Computer Science Department Center for Experimental Computer Science and is a well-known lecturer in the computer industry. He has received numerous best paper and teaching awards, including the Communications Society 1975 Leonard G. Abraham Prize Paper Award. In 1982, as well as having been selected to receive the CCNY Townsend Harris Medal, he was co-winner of the L. M. Ericsson Prize presented by His Majesty King Carl Gustaf of Sweden, for his outstanding contribution in packet switching technology. In July of 1986, he received the 12th Marconi International Fellowship Award presented by His Royal Highness Prince Albert, brother of King Baudoin of Belgium, for his pioneering work in the field of computer networks. He is also founder and President of Technology Transfer Institute, a computer/communications seminar and consulting organization located in Santa Monica, CA.

Dr. Kleinrock is a member of the National Academy of Engineering, a Guggenheim Fellow, a member of IBM Science Advisory Committee, and in 1986 he became a member of the Computer Science and Technology Board of the National Research Council.

John Silvester (Senior Member, IEEE) was born in Kent, England, in 1950. He received the B.A. and M.A. degrees in mathematics and operations research from the University of Cambridge, Cambridge, England, in 1971 and 1975, respectively, the M.S. degree in statistics and computer science from West Virginia University, Morgantown, in 1973, and the Ph.D. degree in computer science from the University of California, Los Angeles, in 1980.

Since 1979 he has been at the University of Southern California, Los Angeles, where he is an Associate Professor in the Department of Electrical Engineering. He teaches courses in Computer Architecture, Computer Networks, Queueing Theory, and Computer Systems Performance Evaluation. His current research interests are performance evaluation and modeling of computer communication networks; packet radio and spread-spectrum radio networks; local area networks; and multiple-access protocols. He is the author of over 30 technical papers and has lectured both in the United States and abroad. He consults for many of the leading networking companies as well as the U. S. Army and is currently Director of Consulting for Technology Transfer Institute. He has served on the program committees of major conferences and is currently on the INFOCOM conference board.

Dr. Silvester is a member of the ACM and is currently Chairman of the Technical Committee for Computer Communications of the Communications Society.