Receptivity: A Measure of Computer Networks' Ability to Accommodate Concurrent Communication

John K. Antonio
School of Electrical Engineering
Purdue University
West Lafayette, IN 47907
Tel: (317) 494-6416
E-mail: jantonio@ecn.purdue.edu

ABSTRACT

In this paper we introduce a network performance measure, called receptivity, which quantifies computer networks' ability to accommodate concurrent communication. Because of limited resources and/or poor resource management, networks may have no choice but to reject a communication request for at least one origin-destination pair. Receptivity is basically defined as the probability that such a rejection does not occur, i.e., it is the probability that the network will be "receptive" to all expected concurrent communication requests. An insightful graph-theoretic result is derived which upper bounds the maximum number of concurrently communicating origin-destination pairs for which a given computer network can accommodate. The bound is dependent on several important network parameters including the number of network links, the maximum link capacity, and a measure of the network's average minimum hop distance. Next, the upper bound result is used as a basis for estimating receptivity. It is then shown that the estimate for receptivity compares favorably with simulated values of receptivity for several example networks. The utility of this derived estimate for receptivity is twofold. First, it can be computed quickly (polynomial time). Second, its simple analytic form provides the network architect with clear insight into some of the inherent limitations and tradeoffs associated with network design parameters in terms of how they affect the receptivity of the network.

I. INTRODUCTION

A. Basics

Designing optimal computer network topologies has been a subject of considerable interest in recent years, see for example [1–12]. Naturally, the most fundamental issue associated with this area of research is the question: What constitutes an "optimal" computer network? Intuitively, the notion of optimality depends on the expected application for the network. For instance, in reference [3] the authors present a methodology for optimally interconnecting large geographically distributed computer networks. The overall objective of the design procedure in [3] is, roughly speaking, to interconnect a geographically distributed set of computers so as to minimize the expected average delay in sending packets of data through the network. Practical constraints related to capital investment (the number, physical length, and the bandwidth of the interconnecting data links) are incorporated as design constraints in the formulation of this optimization problem. In other applications there is the requirement to interconnect hundreds or even thousands of individual processing elements (PEs) into a parallel network computer [10,11,26–28]. Here, the network design problem is generally influenced by whether the network computer is to operate as a "special purpose" or "general purpose" parallel network computer.

In this paper our first aim is to define a new performance measure, *receptivity*, which gives an indication of how receptive a particular computer network[1] is to accommodating concurrent communication. The measure of a network's receptivity can be meaningfully applied to virtually any type of a computer network, including: special purpose, general purpose, distributed, and/or parallel. Our premise for defining receptivity stems from the fact that whenever origin-destination pairs must contend for network resources, the result is low overall throughput. Our claim is that by improving the receptivity of a network, the overall throughput is necessarily improved as well. As would be expected, the employment of techniques such as optimal routing [1,13,17,20,23,25–29], scheduling [15,19,30], and flow control [14,21,22] can improve the receptivity of a computer network, however, these techniques are not our main concern in this paper.

Our primary focus shall be to investigate how a network's topological structure affects its measure of receptivity. To this end, a graph-theoretic result is derived which provides an upper bound for the number of concurrently communicating origin-destination pairs which can be accommodated by a given computer network. The bound is dependent on topological attributes such as minimum hop distances, number of edges, and network diameter.[2] The bound is shown to be asymptotically tight for several topologically distinct example networks. The bound is then used as a cornerstone for developing a simple formula to estimate the receptivity of general computer networks. This theoretically based estimate for the receptivity of a network is useful on two counts. First, it can be computed quickly (relative to the amount of computation time generally required to determine the exact value of receptivity). Second, its simple form naturally exposes important tradeoffs and limitations associated with certain network design choices—in terms of how the values of these design parameters impact the network's measure of receptivity. For instance, under certain conditions it turns out that the estimate for receptivity is approximately inversely proportional to the diameter of the network. Thus, in such cases the network designer should certainly strive to minimize the diameter of the network graph. Of course in reality the minimization of network diameter may be tempered by other practical design constraints such as a hard limit on the number of interconnecting links that can be afforded, the maximum allowable node degree, reliability requirements, and fault tolerance.

For more in-depth tutorials on various types of computer

[1] Throughout the remainder of this paper, the term *computer network* is meant to include the so-called *network computers*, a term often associated with certain classes of general purpose parallel computers.

[2] The diameter of a network is defined as the largest minimum hop distance, taken over all possible origin-destination pairs.

CH2979-3/91/0000-0358 $1.00

networks (ranging from large distributed data networks to massively parallel computer networks), the reader is referred to [1,5,10,11,26–28].

B. Organization of the Paper

In Section II we introduce the network model and the associated notations which are used throughout the remainder of the paper. The assumed network model is coarse in that it is does not explicitly model many of the fine-grain implementation details, but, rather, accounts for them using certain macro-level assumptions. The utility of this type of network model is that it exploits important common network features for a wide class of computer networks. Section III contains the formal definition of network receptivity along with some simple illustrative examples for which receptivity can be easily determined (by inspection). Section IV addresses the issue of computing the exact value of receptivity for general computer networks. In Section V, a graph-theoretic result is derived which provides for an "easy-to-compute" estimate of receptivity. Also, this result is useful in that its simple form provides a clear indication of some of the inherent limitations and trade-offs associated with maximizing receptivity under practical design constraints. Finally, it is shown that the easy-to-compute estimate of receptivity compares favorably with simulated values of receptivity for several example networks.

II. The Network Model

A. Preliminaries

We shall model an n-node computer network with a directed and connected graph $G = (\mathcal{N}, \mathcal{L})$, where $\mathcal{N} = \{1, \ldots, n\}$ denotes the set of nodes and $\mathcal{L}$ is the associated set of interconnecting links. There is a direct link connecting node i to node j if and only if $(i,j) \in \mathcal{L}$. Associated with each link $(i,j) \in \mathcal{L}$ is a link capacity, as defined below.

Definition: The *capacity* of link (i,j), denoted as C_{ij}, is defined as the maximum possible bit-rate of incoming data that node i can receive—possibly simultaneously from any or all neighboring nodes—and successfully transmit to node j in finite time, via link (i,j).

Note that the above definition of a link's capacity hides the details of how the transmissions are accomplished, i.e., whether or not queues are used at each outgoing link; what—if any—type of coding scheme is employed, etc.

Also associated with each link is a flow rate, defined below.

Definition: The *flow* on link $(i,j) \in \mathcal{L}$ at time t, denoted as $f_{ij}(t)$, is defined as the total rate at which data is received at node i, for link (i,j), at time t.

Definition: We shall let Ω denote the set of all possible origin-destination pairs. Thus, using standard set notation we have

$$\Omega = \{(i,j) : i,j \in \mathcal{N} \;\;\&\;\; i \neq j\}. \qquad (1)$$

So the total number of possible origin-destination pairs is given by $|\Omega| = n(n-1)$. To avoid notational confusion between links in $\mathcal{L}$ and origin-destination pairs in Ω, we shall often refer to generic elements in Ω as ω. Also, for each $\omega \in \Omega$, we let O_ω and D_ω denote the origin and destination node, respectively, associated with the origin-destination pair ω. Associated with each origin-destination pair is a demand function, as defined below.

Definition: For each $\omega \in \Omega$ there is a non-negative *demand function*, denoted as $r_\omega(t)$, which indicates the rate at which data is requested to be sent from node O_ω to node D_ω, at time t.

Definition: For a given network graph $G = (\mathcal{N}, \mathcal{L})$, let $\mathcal{P}_\omega$ denote the set of all logical paths from O_ω to D_ω, defined for each $\omega \in \Omega$.

Definition: For each $\omega \in \Omega$ and each path $p \in \mathcal{P}_\omega$, let $x_p(t)$ denote the flow on path p at time t (i.e., the rate at which data is being sent along path p).

Based on the above notations and definitions, the following flow conservation and constraint equations are needed in order to define the concept of a valid routing.

$$x_p(t) \geq 0, \qquad \text{for all } p \in \mathcal{P}_\omega, \omega \in \Omega. \qquad (2)$$

$$\sum_{p \in \mathcal{P}_\omega} x_p(t) = r_w(t), \qquad \text{for all } \omega \in \Omega. \qquad (3)$$

$$f_{ij}(t) = \sum_{\substack{\text{all paths } p \\ \text{containing link } (i,j)}} x_p(t), \qquad \text{for all } p \in \mathcal{P}_\omega, \omega \in \Omega. \qquad (4)$$

$$f_{ij}(t) \leq C_{ij}, \qquad \text{for all } (i,j) \in \mathcal{L}. \qquad (5)$$

Equation (2) is an obvious nonnegativity constraint for the flow rates on all logical paths. Equation (3) ensures that the total flow rate on all paths from O_ω to D_ω is equal to the requested demand, $r_\omega(t)$. Equation (4) states that the total flow on link (i,j), denoted $f_{ij}(t)$, equals the sum total of all path flows $x_p(t)$ for which link (i,j) belongs. Finally, equation (5) states that the flow rates on each link must not exceed that link's capacity. Next, we formally define a valid routing.

Definition: For a given network graph $G = (\mathcal{N}, \mathcal{L})$ with link capacities C_{ij}, $(i,j) \in \mathcal{L}$, and a given set of origin-destination demand functions, $r_\omega(t)$, $\omega \in \Omega$, a *valid routing* is defined as a set of path flows, $\{x_p(t) : p \in \mathcal{P}_\omega, \omega \in \Omega\}$, which satisfy equations (2) through (5).

We note here that for a given network graph and a given set of demand functions, exactly one of the following statements is true:

- There exists at least one valid routing.
- No valid routing exists.

If one or more valid routing exists, then optimal routing could be used to further determine which of these valid routings is the "best." For example, in large data networks an optimal routing is often defined as a routing which minimizes the overall average queuing delay in the network. In this paper, however, we shall not concern ourselves directly with determining optimal routing solutions. Instead, we shall focus our attention on a more fundamental issue: the existence of valid routings. In Section V we shall address the issue of estimating the number of nonzero demand functions for which a given network can accommodate, under the assumption that the nonzero demands are uniformly bounded away from zero. As might be expected, the underlying topology of the given network is of prime importance.

B. Routing Models

For the purposes of this paper we shall assume one of two types of routing: (1) packet switched or (2) circuit switched. In packet switched routing, a message originating at node O_ω and destined for node D_ω is divided into small packets of data which are subsequently routed to the destination along (possibly) several distinct paths interconnecting O_ω to D_ω. The destination node must then reconstruct the original message by piecing together the incoming packets (which do not, generally speaking, arrive in the proper order). In circuit switched routing, a single virtual circuit is established between the origin and destination node. Once a virtual circuit is established, the origin node sends data to the destination node as if the two nodes were directly connected.

The main thrust of this paper is, essentially, to quantify an upper bound for the number of concurrently communicating origin-destination pairs which a given computer network can accommodate. Therefore, we shall not concern ourselves directly with the issue of the time complexity associated with solving optimal routing problems. We are more concerned with the amount of concurrent communication that is achievable; given that some type of (intelligent) routing scheme is used. So, with respect to circuit switched routing, we are not concerned with the latency associated with setting up a virtual circuit. Instead, we are interested in whether or not a circuit is available. In terms of packet switched routing, we are not concerned with the computational time required by each node to determine the optimal next neighbor to forward a packet to. Instead, we are interested in answering the more fundamental question of whether or not there is an available neighbor to which a packet may be forwarded.

Recall the constraint equation (4) of the previous subsection. From this equation note that a nonzero flow on path p at time t instantaneously adds $x_p(t) > 0$ amount of flow to each link flow $f_{ij}(t)$, where link (i,j) is on path p. Under the assumption of circuit switching, the constraint equation (4) is an accurate model since once a virtual circuit is established, the interconnecting links along a path act as a single interconnecting link, and as such, are used concurrently. On the other hand, under the assumption of packet switching this constraint equation represents a worst-case estimate for the amount of concurrent utilization of the links along path p. For instance, consider a particular situation in which $x_{p_o}(t) = 1$ packet/sec, for some $p_o \in \mathcal{P}_{\omega_o}$, and for some $\omega_o \in \Omega$. So, starting at time t a single packet is forwarded from the origin node O_{ω_o} to some neighbor node j, and therefore link (O_{ω_o}, j) is utilized (for at least one second). However, the remaining links along the path from the node j to the destination node D_{ω_o} are not used during the time interval $[t, t+1]$. With a little thought, one sees that if the originating message consists of a relatively large number of packets (with respect to, for example, the diameter of the network) then the links along the path will be "filled" much like a pipeline. Therefore, in terms of the packet switching assumption, the results obtained in this paper—which employ equation (4) for modeling packet switched routing—are generally worst case estimates of the amount of concurrent communication an actual packet switching network is able to accommodate. Simply stated, during the time that the links along a (packet switched) path are being filled or emptied, all of the links are not, in reality, used concurrently; nevertheless, in this paper we shall assume for simplicity that they are.

C. Modeling the Demand Functions

Our goal here is to provide a realistic model for the demand functions,[3] $r_\omega(t)$. We assume that time is discretized, that is $t \in I^+$, where $I^+ = \{1, 2, 3, \ldots\}$. At each time $t \in I^+$, each demand function $r_\omega(t)$ assumes one of two possible states—either the *active state* or the *inactive state*. The active state refers to the situation in which the demand function assumes a positive (nonzero) value, while the inactive state is defined when $r_\omega(t) = 0$, at time t. Whenever a demand function is in the active state, we further require that the actual value it assumes is taken from a finite set containing N_r distinct real values, denoted by $\Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_{N_r}\}$, where $\gamma_i > 0$, for all $i \in \{1, 2, \ldots, N_r\}$. So, if $r_\omega(t)$ is in the active state, then $r_\omega(t) \in \Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_{N_r}\}$. We define the minimum and maximum values contained in Γ as $\gamma_{\min} = \min_{\gamma \in \Gamma}\{\gamma\}$ and $\gamma_{\max} = \max_{\gamma \in \Gamma}\{\gamma\}$.

The description for the demand functions provided thus far allows for a rich class of stochastic variants. We now characterize the class of demand functions we shall consider throughout the remainder of this paper. First, for each $\omega \in \Omega$ and each $t \in I^+$ we assign an associated Bernoulli probability, denoted as $0 \leq q_{\omega,t} \leq 1$, which represents the probability that the demand function $r_\omega(t)$ is in the inactive state. Thus, the probability that $r_\omega(t)$ is in the active state is given by $1 - q_{\omega,t}$. Now, whenever $r_\omega(t)$ is in the active state, then the actual nonzero real value assumed by the function $r_\omega(t)$ is governed by a probability density function, denoted as $P_{\omega,t}(\gamma)$, which is defined for each $\gamma \in \Gamma$. So, for each $\omega \in \Omega$ and each $t \in I^+$, the associated density function $P_{\omega,t}(\gamma)$ can be chosen as any nonnegative function (defined over the finite set of real values $\gamma \in \Gamma$) which satisfies the following condition:

$$\sum_{\gamma \in \Gamma} P_{\omega,t}(\gamma) = 1. \tag{8}$$

We note here that each demand function is completely characterized by the following 3-tuple: $\{q_{\omega,t}, P_{\omega,t}(\gamma), \Gamma\}$.

III. Receptivity

A. Defining Receptivity

For a given network graph (including associated link capacity values) and a given characterization of the demand functions, receptivity is defined to be the probability that a valid routing exists. Later in the sub-section we provide a formal definition of receptivity. In order to do so, however, we must introduce a few preliminary definitions.

Definition: The set of *active origin-destination pairs at time t*, denoted as $W(t)$, is defined by

$$W(t) = \{\omega : r_\omega(t) > 0, \omega \in \Omega\}. \tag{9}$$

It is important to note that for a given characterization of all of the demand functions (i.e., given the 3-tuple $\{q_{\omega,t}, P_{\omega,t}(\gamma), \Gamma\}$, for each $\omega \in \Omega$, and each $t \in I^+$) the likelihood that a particular origin-destination pair, say $\omega_o \in \Omega$, will belong to $W(t)$ depends only on the value of the associated Bernoulli probability, $q_{\omega_o,t}$. For instance, if $q_{\omega_o,t} = 0.4$ (for all t), then there is a 60% chance that ω_o will belong to $W(t)$. Of course the actual nonzero value of $r_{\omega_o}(t)$, assuming $\omega_o \in W(t)$, depends on the associated density function $P_{\omega_o,t}(\gamma)$ defined over the finite set Γ.

Next, we recall that all possible sets of active origin destination pairs can be classified in one of two ways. Namely, either: (1) There exists an associated valid routing or (2) An associated valid routing does not exist. If there exists a valid routing for a particular set of active origin-destination pairs, we call the set a concurrently communicating origin-destination pair set, as defined below.

Definition: Consider a network $G = (\mathcal{N}, \mathcal{L})$, with associated link capacities C_{ij}, $(i,j) \in \mathcal{L}$ and a particular set of active origin-destination pairs $W(t)$. We say that $W(t)$ is a *concurrently communicating origin-destination pair set*, if there exists a valid routing for the demand functions $r_\omega(t)$, $\omega \in W(t)$. Also, we adopt the notation $W(t) \sim W_{cc}$ to signify that $W(t)$ is a concurrently communicating origin-destination pair set.

[3] Strictly speaking, the demand functions are actually realizations of a particular stochastic process, however, we shall nevertheless call them functions.

We now state the formal definition of receptivity.

Definition: Given a network graph $G = (\mathcal{N}, \mathcal{L})$ with link capacities $C_{ij}, (i,j) \in \mathcal{L}$ and a characterization of the demand functions $\{q_{\omega,t}, P_{\omega,t}(\gamma), \Gamma\}$, for each $\omega \in \Omega$ and each $t \in I^+$. The *receptivity* of the network graph G at time t, denoted as $\rho_G(t)$, is defined by

$$\rho_G(t) = P\left(W(t) \sim W_{cc}\right). \qquad (10)$$

In words, the receptivity of G at time t is the probability that the set of active origin-destination pairs at time t is a concurrently communicating origin-destination pair set. Obviously, the value of $\rho_G(t)$ is always between zero and one. So, for example, if $\rho_G(t) = 0.95$, then the interpretation is that the network G will be able accommodate 95% of all expected concurrent communication requests.

B. Simple Illustrative Examples

In order to appreciate the interplay between demand pattern statistics, network topology, and the associated measure of receptivity, it is instructive to look at some simple examples.

Example 1: Let $G = (\mathcal{N}, \mathcal{L})$ be an arbitrary network graph with associated link capacities $C_{ij}, (i,j) \in \mathcal{L}$, and let $C_{\min} = \min_{(i,j)\in\mathcal{L}}\{C_{ij}\}$. Assume the demand functions are characterized as follows: $q_{\omega,t} \geq 0$, for all $\omega = (i,j) \in \mathcal{L}$ and all t, and $q_{\omega,t} = 1$ for all $\omega \notin \mathcal{L}$. Second, assume the value of $\gamma_{\max}$ is no greater than $C_{\min}$, however, the value of $\gamma_{\min}$ is arbitrary, as are the density functions, $P_{\omega,t}(\gamma)$, for all ω and all t.

Now, since $q_{\omega,t} = 1$, for all $\omega \notin \mathcal{L}$, we note that there is a zero probability of having origin-destination pairs which are not nearest neighbors request communication (i.e., be in the active state). Also, since the value of $\gamma_{\max}$ is assumed to be less than or equal to the minimum link capacity, $C_{\min}$, we conclude that the network is always receptive. That is, $\rho_G(t) = 1$.

Example 2: As in Example 1, let $G = (\mathcal{N}, \mathcal{L})$ be an arbitrary network graph with associated link capacities C_{ij}. Unlike Example 1, however, here we assume that all link capacities are equal, that is, we let $C_{ij} = C$, for all $(i,j) \in \mathcal{L}$. Also, the demand function characterizations (described below) are slightly different from those assumed for Example 1. Assume $q_{\omega,t} = 0$, for all $\omega = (i,j) \in \mathcal{L}$, $q_{\omega_o,t} = q_o \neq 1$, for some $\omega_o \notin \mathcal{L}$, and $q_{\omega,t} = 1$ for all $\omega \notin \{\mathcal{L} \cup \{\omega_o\}\}$. The value of $\gamma_{\max}$ is assumed to be exactly equal to C, and the density functions are as follows (for all ω and all t):

$$P_{\omega,t}(\gamma) = \begin{cases} 0 & \text{if } \gamma \neq C \\ 1 & \text{if } \gamma = C \end{cases}.$$

Putting the description of this example in words; there is a zero probability that non-nearest neighbors will be in the active state, with the exception of the single origin-destination pair ω_o, which is in the active state with a probability of $1 - q_o$, where $q_o \neq 1$. Also, all active origin-destination pairs request (with probability one) the entire capacity of a single link (which are all assumed to have equal capacities). With some thought, it is easy to verify that the receptivity for this example is given by $\rho_G(t) = q_o$. This example shows how the activity of a single origin-destination pair can limit the overall receptivity of of a network. It should be clear that if we were to assume that $0 < q_{\omega,t} \leq 1$, for all $\omega = (i,j) \in \mathcal{L}$ (instead of $q_{\omega,t} = 0$), then we have $q_o < \rho_G(t) \leq 1$. In conclusion, we note that in order to guarantee that this network is always receptive (i.e., $\rho_G(t) = 1$) we must have $q_o = 1$.

IV. Exact Calculation of Receptivity

It turns out that calculating the exact value of a network's receptivity is not (generally speaking) a computationally trivial task. Recall that the total number of distinct origin-destination node pairs for an n-node network is given by $|\Omega| = n(n-1)$. So, one can enumerate all possible combinations of potentially active origin-destination pair sets. For instance, there are $n(n-1)$ possible origin-destination pair sets containing only one element each. Likewise, there are $\binom{n(n-1)}{2}$ possible origin-destination pairs sets containing two elements each, and so on down the line. Based on this type of enumeration, we have that the total number of potentially active origin-destination pair sets, denoted as N_W, is given by

$$N_W = \sum_{k=0}^{n(n-1)} \binom{n(n-1)}{k} = 2^{n(n-1)}. \qquad (11)$$

For convenience, we label the i^{th} possible active origin-destination pair set as W_i, where $i \in \{1, \ldots, N_W\}$. Now, for a given characterization of the demand functions, one can compute the probability that the actual active origin destination pair set at time t, denoted as $W(t)$, is equal to the i^{th} possible origin-destination pair set, W_i. By a straightforward calculation we have

$$P\left(W(t) = W_i\right) = \prod_{\omega \in W_i} (1 - q_{\omega,t}) \prod_{\omega \in \Omega - W_i} q_{\omega,t}, \;\; i \in \{1, \ldots, N_W\}. \qquad (12)$$

The above formula gives the probability that the actual set of active origin destination pairs at time t is equal to each one of the N_W number of possible active origin-destination pair sets. Next, we shall enumerate (and find the associated probabilities for) all possible permutations of values which the demand functions associated with each possible active origin-destination pair set may assume. Recall that there are only a finite number of distinct values for which any single demand function (in the active state) may assume, i.e., $r_\omega(t) \in \Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_{N_r}\}$ and $N_r < \infty$. Therefore, the number of possible distinct sequences of values for the demand functions associated with the set W_i, is given by $(N_r)^{|W_i|}$ (where $|W_i|$ denotes the number of elements in W_i). We let $\Gamma_{W_i}(t)$ denote the actual sequence of nonzero values associated with the active demand functions in $W(t)$, assuming $W(t) = W_i$. (In other words, given that $W(t) = W_i$, $\Gamma_{W_i}(t)$ is the actual sequence of values associated with each demand function in $W(t)$.) Next, we let $\Gamma^j_{W_i}$, $j \in \{1, 2, \ldots, (N_r)^{|W_i|}\}$, denote the j^{th} possible distinct sequence of values for the active demand functions in W_i. So, we compute the probability that $\Gamma_{W_i}(t) = \Gamma^j_{W_i}$ as follows.

$$P\left(\Gamma_{W_i}(t) = \Gamma^j_{W_i} \mid W(t) = W_i\right) = \prod_{\gamma \in \Gamma^j_{W_i}} P_{\omega,t}(\gamma), \qquad (13)$$

where the ω appearing in $P_{\omega,t}(\gamma)$ is understood to be to the origin-destination pair which corresponds to the associated $\gamma \in \Gamma^j_{W_i}$.

For a given graph $G = (\mathcal{N}, \mathcal{L})$ with link capacities C_{ij}, $(i,j) \in \mathcal{L}$, we know that either there is, or there is not a valid routing associated with each possible sequence $\Gamma^j_{W_i}$. We define an indicator function, denoted as $\mathcal{I}_G\left(\Gamma^j_{W_i}\right)$, as follows:

$$\mathcal{I}_G\left(\Gamma^j_{W_i}\right) = \begin{cases} 1 & \text{if } \exists \text{ a valid routing on } G \text{ assoc. with } \Gamma^j_{W_i} \\ 0 & \text{otherwise} \end{cases}$$

So, the receptivity is computed as

$$\rho_G(t) = \sum_{i=1}^{N_W} \sum_{j=1}^{(N_r)^{|W_i|}} \mathcal{I}_G(\Gamma^j_{W_i}) P\left(\Gamma_{W_i}(t) = \Gamma^j_{W_i}\right), \qquad (14)$$

where, from equations (12) and (13), we note that

$$P\left(\Gamma_{W_i}(t) = \Gamma^j_{W_i}\right) = \prod_{\gamma \in \Gamma^j_{W_i}} P_{\omega,t}(\gamma) \prod_{\omega \in W_i} (1 - q_{\omega,t}) \prod_{\omega \in \Omega - W_i} q_{\omega,t}. \qquad (15)$$

Simply stated, equation (14) is the summation of the probabilities of all possible outcomes, for which there exists a valid routing. So, for example, if the network G is such that there exists a valid routing for every possible outcome, except for one particular $\Gamma^{j_o}_{W_{i_o}}$, for which $P\left(\Gamma_{W_i}(t) = \Gamma^{j_o}_{W_{i_o}}\right) = 0.05$, then its associated measure of receptivity is given by $\rho_G(t) = 0.95$.

It should be clear that computing the exact value of receptivity as per equation (14) is not practical, even for fairly small networks.

V. Estimating Receptivity

A. An Upper Bound for Concurrent Communication

In this sub-section we derive a theorem which provides an upper bound for the number origin-destination pairs for which a given network can provide concurrent communication. The bound is used later in this section as a basis for estimating a given network's measure of receptivity. It is also noted that the bound is useful in that it provides the network designer some immediate insights into the types of network topologies which tend to have good measures of receptivity.

Let h_ω denote the shortest hop distance associated with the origin-destination pair $\omega \in \Omega$, and let ℓ_p denote the number of hops along path p, defined for all paths $p \in \mathcal{P}_\omega$, $\omega \in \Omega$. Based on the above notation, note that for each path $p \in \mathcal{P}_\omega$, the following condition is obviously satisfied: $\ell_p \geq h_\omega$, for all $\omega \in \Omega$. We also define $x_{\min}$ as the minimum nonzero path flow, taken over all nonzero path flows:

$$x_{\min} = \min_{\substack{\omega \in \Omega \\ p \in \{p: x_p(t) > 0, p \in \mathcal{P}_\omega\}}} \{x_p(t)\}. \qquad (16)$$

Theorem: For a given network graph $G = (\mathcal{N}, \mathcal{L})$ with link capacities $C_{ij}, (i,j) \in \mathcal{L}$, the size of all possible concurrently communicating origin-destination pair sets, denoted by $|W_{cc}|$, is bounded above by

$$|W_{cc}| \leq \frac{C_{\max}}{x_{\min}} \frac{|\mathcal{L}|}{h_{\text{avg}}}, \qquad (17)$$

where $|\mathcal{L}|$ is the number of network links, $C_{\max}$ is the maximum link capacity, defined by $C_{\max} = \max_{(i,j) \in \mathcal{L}}\{C_{ij}\}$, $x_{\min}$ is the minimum nonzero path flow (defined above), h_{avg} is the average minimum hop distance between origin-destination pairs in W_{cc}, defined by $h_{\text{avg}} = \frac{1}{|W_{cc}|} \sum_{\omega \in W_{cc}} h_\omega$.

Proof: Assume $W(t) = W_{cc}$ is a concurrently communicating origin-destination pair set. The objective is the show that all possible choices for W_{cc} must satisfy the above inequality.

Since the origin-destination pairs $\omega \in W_{cc}$ can communicate concurrently (by assumption), then there exists (by definition) at least one associated valid routing. Now, for any such valid routing we have that the flow on any link (i,j), denoted $f_{ij}(t)$, is less than the associated capacity, C_{ij}, for all $(i,j) \in \mathcal{L}$. Furthermore, since the value of $f_{ij}(t)$ is the sum total of all (nonzero) path flows which contain link (i,j), we conclude that the maximum number of active paths sharing link (i,j) is necessarily less than or equal to $\frac{C_{ij}}{x_{\min}}$, which is bounded by $\frac{C_{\max}}{x_{\min}}$. So, we have that

$$\textit{Number of Active Paths Sharing Link } (i,j) \leq \frac{C_{\max}}{x_{\min}}, \quad \text{for all } (i,j) \in \mathcal{L}. \qquad (18)$$

Next, note that the total number of hops taken by all active paths is given by

$$\textit{Total Number of Hops Taken by All Active Paths} = \sum_{\substack{\omega \in W_{cc} \\ p \in \{p: x_p(t) > 0, p \in \mathcal{P}_\omega\}}} \ell_p, \qquad (19)$$

where ℓ_p is the number of hops along path p. Now, by equation (18) we have that each link is used by no more than $\frac{C_{\max}}{x_{\min}}$ distinct paths, so the total number of hops taken by all active paths must be bounded by $\frac{C_{\max}}{x_{\min}}|\mathcal{L}|$. Combining this with equation (19), we have

$$\sum_{\substack{\omega \in W_{cc} \\ p \in \{p: x_p(t) > 0, p \in \mathcal{P}_\omega\}}} \ell_p \leq \frac{C_{\max}}{x_{\min}}|\mathcal{L}|. \qquad (20)$$

Now, since $|\{p : x_p(t) > 0, p \in \mathcal{P}_\omega\}| \geq 1$, for each $\omega \in W_{cc}$ and since $\ell_p \geq h_\omega$, for all $p \in \mathcal{P}_\omega$, and for each ω, we have that

$$\sum_{\omega \in W_{cc}} h_\omega \leq \sum_{\substack{\omega \in W_{cc} \\ p \in \{p: x_p(t) > 0, p \in \mathcal{P}_\omega\}}} \ell_p. \qquad (21)$$

Combining equations (20) and (21) then yields

$$\sum_{\omega \in W_{cc}} h_\omega \leq \frac{C_{\max}}{x_{\min}}|\mathcal{L}|. \qquad (22)$$

Finally, by defining $h_{\text{avg}} = \frac{1}{|W_{cc}|} \sum_{\omega \in W_{cc}} h_\omega$, the result is proven.
□

The result of the above theorem coincides with intuition. For a fixed network topology one would expect that if the value of $C_{\max}$ is increased, then the potential for accommodating more concurrently communicating origin-destination pairs should increase as well. Similar statements can be made regarding the parameters $x_{\min}$ and $|\mathcal{L}|$, i.e., increasing $|\mathcal{L}|$ and/or decreasing $x_{\min}$ should increase the potential for more concurrent communication. Finally, the fact that the bound is inversely proportional to the average minimum hop distance of the origin-destination pairs in W_{cc} also makes intuitive sense. That is, if the set of origin-destination pairs are such that each origin-destination pair is "close" in terms of minimum hop distances, then fewer links will (generically) be utilized when all active origin-destination pairs are concurrently communicating. On the other hand, if the origin-destination pairs tend to be "far away" from each other, then the number of links utilized for communication by just a single origin-destination pair may be on the order of the diameter of the network.

In general, the value of h_{avg} is bounded above and below by

$$1 \le h_{\text{avg}} \le d, \tag{23}$$

where d denotes the diameter of the network. Substituting these bounds into equation (17), we have

$$|W_{cc}| \le |\overline{W_{cc}}| \stackrel{\text{def}}{=} \frac{C_{\max}}{x_{\min}} \frac{|\mathcal{L}|}{1} \tag{24}$$

and

$$|W_{cc}| \le |\underline{W_{cc}}| \stackrel{\text{def}}{=} \frac{C_{\max}}{x_{\min}} \frac{|\mathcal{L}|}{d}, \tag{25}$$

where $|\overline{W_{cc}}|$ and $|\underline{W_{cc}}|$ denote maximum and minimum upper bounds for $|W_{cc}|$, respectively. Note that the value of $|\overline{W_{cc}}|$ provides a optimistic (best case) bound, and is obtained from equation (17) by assuming all active origin-destination pairs are separated by only one hop. On the other hand, the value of $|\underline{W_{cc}}|$ is a conservative (worst case) bound, since it is based on the assumption that all active origin-destination pairs are d hops apart. From this later bound, we note (generically) that the size of concurrently communicating origin-destination pairs sets is inversely proportional to the diameter of the network.

Tightness of the Upper Bound

Simulation studies presented here indicate that "on the average," the conservative upper bound $|\underline{W_{cc}}|$, of equation (25), is asymptotically tight. The simulations were set-up as follows. First, a simple topological structure is chosen: linear, star, or square mesh network.[4] The value of $C_{\max}/x_{\min}$ is fixed, and the number of nodes in the network, say n, is fixed. The set of active origin-destination pairs W is initialized to the empty set. Then, an active origin-destination pair and associated demand is chosen at random and added to the set of active origin-destination pairs. The optimal routing problem is solved[5] (using the gradient projection based algorithm of [29]). If a valid routing exists (i.e., $f_{ij} < C_{ij}$, for all $(i,j) \in \mathcal{L}$), then another origin-destination pair is chosen at random and added to the set of active origin-destination pairs. This procedure is continued until the solution of the optimal routing solution is no longer valid. At this point, the number of active origin-destination pairs are stored and the procedure is done again (with W initialized to the empty set). The average of the number of origin-destination pairs accommodated for each run is then computed; denote this quantity as $|W|_{\text{avg}}$.

Let $K = C_{\max}/x_{\min}$ in the following description. For the n-node linear network depicted in Fig. 1(a), note that the number of links is given by $|\mathcal{L}| = n-1$ and the the diameter is obviously $d = n-1$. Thus, according to equation (25) the least upper bound for $|W_{cc}|$ is given by

$$|\underline{W_{cc}}| = K\left(\frac{n-1}{n-1}\right) = K. \tag{26}$$

So, assuming that K is independent of n (which is reasonable), we see that the least upper bound for $|W_{cc}|$ is a constant. Fig. 1(b) shows the simulation results for the linear network, which indicate that the simulated value of $|W|_{\text{avg}}$ is independent of n as well.

Similar numerical studies were done for the n-node star network shown in Fig. 2(a). For this network note that the number of links is $|\mathcal{L}| = n-1$ and the diameter is $d = 2$. Therefore, according to the the least upper bound, we have

$$|\underline{W_{cc}}| = K\left(\frac{n-1}{2}\right). \tag{27}$$

The results of the numerical studies for the star network are summarized by the graph of Fig. 2(b). Note that the value of $|W|_{\text{avg}}$ increases linearly with n, as expected.

To further justify the theory, simulations were done for the n-node mesh network, see Fig. 3(a). In an n-node mesh network, we have $|\mathcal{L}| = K_1 n$ and $d = K_2\sqrt{n}$, thus the least upper bound becomes

$$|\underline{W_{cc}}| = K_3\sqrt{n}, \tag{28}$$

where K_1, K_2, and K_3 are constants independent of n. The results of the simulations for the mesh network given in Fig. 3(b), agree with the theoretically derived least upper bound. Namely, $|W|_{\text{avg}} = O(\sqrt{n})$.

Note that the linear, star, and mesh networks all have $|\mathcal{L}| = O(n)$. However, even though the star network can accommodate the most origin-destination pairs (on the average), there is a severe tradeoff in terms of reliability. That is, if the the center node of a star network fails, the entire network becomes completely disconnected. While the linear network is also 1-connected, the overall integrity of the linear network is better than that of the star network. Finally, the mesh network strikes a reasonable compromise in that the it is 3-connected, while providing $|\underline{W_{cc}}| = O(\sqrt{n})$.

It is noted further that the linear, star, and mesh topologies were used here primarily because their simple and regular structures simplified the programming task (as the number of nodes was increased). Future simulations are planned for more irregular and complicated structures. Of particular interest is the balanced hierarchically clustered topology, which is described in more detail in the concluding section.

B. A Simple Estimate of Receptivity

The upper bound result of the previous sub-section is used here as a basis for estimating receptivity. We offer the following line of reasoning in deriving our estimate. First, note that the actual value of $x_{\min}$ can be known only after a valid routing has been determined. However, we would rather not go to the trouble of determining a valid (or optimal) routing for every potential W_{cc}. Therefore, we propose making the assumption that $x_{\min} \approx \gamma_{\min}$. We believe this to be a reasonable approximation, since our practical experience with various optimal routing algorithms—see [17] and [18], for example—indicate that the values of the nonzero path flows are usually on the same order of magnitude as associated demand $r_\omega(t)$, which in our case is bounded below by $\gamma_{\min}$.

Second, we propose estimating h_{avg} as follows:

$$h_{\text{avg}} \approx \tilde{h}_t \stackrel{\text{def}}{=} \frac{\sum_{\omega\in\Omega}(1-q_{\omega,t})h_\omega}{\sum_{\omega\in\Omega}(1-q_{\omega,t})}. \tag{29}$$

The reasoning for the above approximation is that it measures, in some sense, the typical or expected value of h_{avg} taken over all W_{cc}'s, without having to explicitly compute the value of h_{avg} for every possible W_{cc}. Note that if the $q_{\omega,t}$'s in equation (29) are equal for all ω (i.e., $q_{\omega,t} = q_t$, for all ω) then the approximation for h_{avg} given in equation (29) reduces to the original definition of h_{avg}, with $W_{cc} = \Omega$.

Based on the approximations made thus far, we have the following new bound for concurrent communication:

$$|W_{cc}| \le \frac{C_{\max}}{\gamma_{\min}} \frac{|\mathcal{L}|}{\tilde{h}_t}, \tag{30}$$

[4] These simple topologies were used because their diameters are well defined as a function of the number of nodes. Also, their high degree of regularity simplified the programming task as the number of nodes was increased.

[5] Actually, the optimal routing problem is solved only for the mesh network, since for both the linear and star networks, there is only one path between each pair of nodes.

where $\tilde{h}_t$ is defined as

$$\tilde{h}_t = \frac{\sum_{\omega\in\Omega}(1-q_{\omega,t})h_\omega}{\sum_{\omega\in\Omega}(1-q_{\omega,t})}. \quad (31)$$

Now, if the value of $|W(t)|$ (based on some characterization of the demand functions) is, with a high probability, close to the value of the bound given for $|W_{cc}|$ in equation (30), then one would expect the network to have a relatively poor measure of receptivity. On the other hand, if the value of $|W(t)|$ is, with a high probability, substantially less than the bound of equation (30), then one would expect that the network have a relatively good measure of receptivity. It is exactly this line of reasoning that we propose the following estimate for receptivity.

$$\rho_G(t) \approx P\left(|W(t)| \leq \alpha \frac{C_{\max}}{\gamma_{\min}} \frac{|\mathcal{L}|}{\tilde{h}_t}\right), \quad (32)$$

where α is a scaling constant, $0 < \alpha \leq 1$. By choosing α sufficiently close to one, the approximation is optimistic in that it provides a high estimate of the receptivity; on the other hand, by choosing α close to zero, the estimate is generically pessimistic in that it under estimates the actual receptivity. We show in the next sub-section that choosing $\alpha = \frac{1}{2}$ often provides a reasonably accurate estimate of receptivity.

Unfortunately, computing the exact value of the above probability in equation (32) is still not trivial—it generically requires the summing of $O(N_W)$ of the terms given in equation (12). In particular, the value of $P(|W(t)| \leq k)$ is given by

$$P(|W(t)| \leq k) = \sum_{\substack{i\in\{1,2,\ldots,N_W\} \\ |W_i|\leq k}} \prod_{\omega\in W_i}(1-q_{\omega,t}) \prod_{\omega\in\Omega-W_i} q_{\omega,t}. \quad (33)$$

Our next step, then, is to provide an approximation to the above probability. Consider first the simple case in which the Bernoulli probabilities associated with each ω are equal. That is, assume $q_{\omega,t} = q_t$, for all $\omega \in \Omega$ and all t. Under this assumption, it is easy to verify that the probability density function, for the size of the active origin-destination pair set, is the binomial distribution. So, we have

$$P(|W(t)| = k) = \binom{n(n-1)}{k}(1-q_t)^k q_t^{n(n-1)-k},$$

$$\text{for all } k = 0, 1, \ldots, n(n-1). \quad (34)$$

Now, by the De Moivre-Laplace limit theorem, [31] (special case of the central limit theorem), we have that if the random variable X has a binomial distribution with parameters N and $p = 1-q$; then for fixed $a < b$

$$P(a \leq X \leq b) \to \Phi\left(\frac{b-Np}{\sqrt{Npq}}\right) - \Phi\left(\frac{a-Np}{\sqrt{Npq}}\right) \quad \text{as } N \to \infty, \quad (35)$$

where $\Phi(\cdot)$ is the standard cumulative normal distribution, defined by

$$\Phi(z) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{z} e^{-t^2/2}dt. \quad (36)$$

So, our approximation, under the assumption that $q_{\omega,t} = q_t$ for all ω and all t, becomes

$$\rho_G(t) \approx \Phi\left(\frac{\left(\alpha\frac{C_{\max}}{\gamma_{\min}}\frac{|\mathcal{L}|}{\tilde{h}_t}\right) - n(n-1)(1-q_t)}{\sqrt{n(n-1)(1-q_t)q_t}}\right). \quad (37)$$

Now, to handle the general situation[6] we propose approximating the actual density with a binomial distribution having parameters $n(n-1)$ and $\tilde{p}_t = 1-\tilde{q}_t$, where $\tilde{q}_t$ is defined by

$$1-\tilde{q}_t = \frac{1}{n(n-1)}\sum_{\omega\in\Omega}(1-q_{\omega,t}). \quad (38)$$

We note that by choosing the parameter $\tilde{q}_t$ of the approximating binomial distribution as per equation (38), the expected value for $|W(t)|$ (based on the approximating binomial distribution) is the same as that of the actual distribution. That is, for the actual distribution we have $E[|W(t)|] = \sum_{\omega\in\Omega}(1-q_{\omega,t})$, and for the approximating binomial distribution we have $E[|W(t)|] = n(n-1)(1-\tilde{q}_t)$. Based on all of the above approximations, we next state our most general formula for estimating receptivity.

The Simple Estimate for Receptivity:

$$\rho_G(t) \approx \Phi\left(\frac{\left(\alpha\frac{C_{\max}}{\gamma_{\min}}\frac{|\mathcal{L}|}{\tilde{h}_t}\right) - n(n-1)(1-\tilde{q}_t)}{\sqrt{(n-1)n(1-\tilde{q}_t)\tilde{q}_t}}\right), \quad (39)$$

where

$$\tilde{h}_t = \frac{\sum_{\omega\in\Omega}(1-q_{\omega,t})h_\omega}{\sum_{\omega\in\Omega}(1-q_{\omega,t})}. \quad (40)$$

and

$$\tilde{q}_t = 1 - \frac{1}{n(n-1)}\sum_{\omega\in\Omega}(1-q_{\omega,t}). \quad (41)$$

We note that both $\tilde{h}$ and $\tilde{q}_t$ can be computed (serially) in $O(n^2)$ time, since $|\Omega| = n(n-1) = O(n^2)$. Also, the time required to evaluate the integral associated with $\Phi(\cdot)$ (either numerically or with a standard table) is independent of the size of the network.

C. Simulation Studies

We compared the estimate of receptivity, given in equation (39), with simulated values of receptivity for the three network topologies shown in Figs. 1(a), 2(a), and 3(a). In all cases, we assume that $q_{\omega,t} = q$, for all ω and t. So, at each point in time the probability that any origin-destination pair is in the active state is $1-q$. Figs. 4, 5 and 6 are comparisons between estimated and simulated values of receptivity for the linear, star, and mesh topologies, respectively. The simulated values were obtained by computing (for each value of $1-q$) the percentage of trials for which there exists a valid routing. For each trial, a set of active origin-destination pairs was randomly generated, based on the associated value of $1-q$. Three-hundred trials were done for each of the 200 data points between $1-q=0$ and $1-q=0.2$ (stepsize = 0.01).[7] As would be expected, as the value of $1-q$ is increased, the value of receptivity decreases. Also, for a fixed value of $1-q$ and a fixed topology, the receptivity decreases as the size of the network increases. This is due to the fact that for all topologies considered, the total number of data links grows only as $O(n)$, while the number of possible origin-destination pairs increases with $O(n^2)$. In all cases, the parameter α in the estimate of receptivity was chosen as $\frac{1}{2}$.

[6] By general situation, we mean one in which all of the values of $q_{\omega,t}$ are not equal and thus the underlying density (for the size of the active origin-destination pair set) is not a true binomial distribution.

[7] Simulated values of receptivity were computed (instead of exact values) since the number of active origin-destination pairs sets required to be tested for exact calculation, i.e., N_W of equation (11), is too large, even with $n = 10$.

All computation was done on a Sun 3/60 workstation. Typical cpu times required to compute simulated values of receptivity for the 20 node networks (over the range $1-q=0$ to $1-q=0.2$) were on the order of several hours. The cpu times required to compute the corresponding estimated values, as per equation (39), were on the order of a couple of seconds. Integration of the standard cumulative normal distribution was accomplished with an IMSL numerical integration routine [32].

VI. Conclusions and Future Work

The primary contribution of this paper is a new performance measure for computer networks called receptivity. The receptivity of a network is defined as the probability that the network is able to accommodate expected concurrent communication requests. It is pointed out that routing decisions can certainly have an impact on the overall receptivity of a given network; however, in this paper our focus is on the more fundamental issue of determining the relationship between the topology of a network and the corresponding measure of receptivity. A theorem is derived for bounding the number of origin-destination pairs which a given network topology can accommodate concurrently. The bound is then used as a basis for developing a simple estimate for receptivity. It is shown that this easy-to-compute estimate for receptivity does predict simulated values of receptivity reasonably well. Furthermore, for the small networks used in the simulation studies we found that the estimate for receptivity can be computed several orders of magnitude faster than simulated values. (Exact calculation of receptivity requires an exponential-type computation time, an thus was not attempted in this paper.)

The result of the theorem (for bounding the number of concurrently communicating origin-destination pairs) shows, generically, that the receptivity of a network is inversely proportional to the average hop distance between communicating origin destination pairs. So, from a topology design point of view; in order to maximize receptivity, one should strive to locate those origin-destination pairs—which are in the active state with a high probability—close together (in the sense of minimum hop distance). In future work we plan to use this premise as a basis for solving the topology design problem in large data networks. Two classes of network topologies which seem to be particularly well suited are the balanced hierarchically clustered topologies, defined in [16], and the sparse networks, developed in [27].

References

[1] D. P. Bertsekas and R. G. Gallager, *Data Networks.* Englewood Cliffs, NJ: Prentice-Hall, 1987.

[2] J. K. Antonio, "Characterizing Optimal Topological Structures for a Class of Large Distributed Data Networks," *Proc. 1990 IEEE Int. Symp. on Circuits and Systems,* New Orleans, LA, May 1–3, 1990, pp. 2396–2399.

[3] M. Gerla and L. Kleinrock, "On the Topological Design of Distributed Computer Networks," *IEEE Trans. Commun.,* vol COM-25, pp. 48–60, Jan. 1977.

[4] R. Hinden, J. Haverty, and A. Sheltzer, "The DARPA Internet: Interconnecting Heterogeneous Computer Networks with Gateways," *IEEE Computer Magazine*, vol. 16, pp. 38–48, Sept. 1983.

[5] D. P. Bertsekas, J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods.* Englewood Cliffs, NJ: Prentice-Hall, 1989.

[6] D. A. Reed and H. D. Schwetman, "Cost-Performance Bounds for Multimicrocomputer Networks," *IEEE Trans. Computers*, vol. C-32, pp. 83–95, Jan. 1983.

[7] L. D. Wittie, "Communication Structures for Large Networks of Microcomputers," *IEEE Trans. Computers*, vol. C-30, pp. 264–273, April 1981.

[8] B. Gavish, "A General Model for the Topological Design of Computer Networks," *Proc. GLOBCOM '86,* 1986, pp. 1584–1588.

[9] B. Gavish and K. Altinkemer, "A Parallel Savings Heuristic for the Topological Design of Local Access Tree Networks," *Proc. IEEE INFOCOM '86*, Miami, FL, 1986, pp. 130–147.

[10] Y. Saad, M. H. Schultz, "Topological Properties of Hypercubes," *IEEE Trans. Computers,* vol. 37, pp. 867–872, July 1988.

[11] H. J. Siegel, *Interconnection Networks For Large-Scale Parallel Processing,* Lexington, MA: Lexington Books, 1985.

[12] D. B. West, "Construction of Communication Networks Satisfying Bounded Fan-in of Service Ports," *IEEE Trans. Comp.* vol. 37, pp. 1148–1151, 1988.

[13] M. Schwartz and T. E. Stern, "Routing Techniques Used in Computer Communication Networks," *IEEE Trans. on Commun.*, vol. COM-28, pp. 539–552, April 1980.

[14] D. P. Bertsekas, "Optimal Routing and Flow Control Methods for Communication Networks," in *Analysis and Optimization of Systems*, A. Bensoussan and J.L. Lions, Eds. New York: Springer-Verlag, pp. 615–643, 1982.

[15] P. P. Bhattacharya and A. Ephremides, "Optimal Scheduling with Strict Deadlines," *IEEE Trans. Automat. Contr.*, vol. 34, pp. 721–728, July 1989.

[16] J. K. Antonio, G. M. Huang, and W. K. Tsai, "A Fast Distributed Shortest Path Algorithm for a Class of Hierarchically Clustered Data Networks," submitted to the *IEEE Trans. on Computers*, 1988.

[17] W. K. Tsai, G. M. Huang, J. K. Antonio, and W. T. Tsai, "Distributed Iterative Aggregation Algorithms for Box-Constrained Minimization Problems and Optimal Routing in Data Networks," *IEEE Trans. Automatic Control*, vol. AC-34, pp. 34–46, Jan. 1989.

[18] J. K. Antonio, G. M. Huang, and W. K. Tsai, "Deriving Time Complexities for a Class of Distributed Gradient Projection-Based Optimal Routing Algorithms," to appear in the *Proc. 1990 Conf. on Dec. and Control*, Dec. 1990.

[19] T. Hirayama, "Optimal Service Assignment in a Finite-Source Queue," *IEEE Trans. Automatic Control*, vol. 34, pp. 67–75, Jan. 1989.

[20] J. N. Tsitsiklis and D. P. Bertsekas, "Distributed Asynchronous Routing in Data Networks," *IEEE Trans. Automat. Contr.*, vol. 31, pp. 325–332, Apr. 1986.

[21] Z. Rosberg, I. S. Gopal, "Optimal Hop-by-Hop Flow Control in Computer Networks," *IEEE Trans. Automat. Contr.*, vol. 31, pp. 813–822, Sep. 1986.

[22] K. H. Muralidhar and M. K. Sundareshan, "Combined Routing and Flow Control in Computer Communication Networks: A Two-Level Adaptive Scheme," *IEEE Trans. Automat. Contr.*, vol. 32, pp. 15–25, Jan. 1987.

[23] A. Gersht and A. Shulman, "Optimal Routing in Circuit Switched Communication Networks," *IEEE Trans. on Commun.*, vol. COM-37, pp. 1203–1211, Nov. 1989.

[24] B. Gavish and I. Neuman, "A System for Routing and Capacity Assignment in Computer Communication Networks," *IEEE Trans. on Commun.*, vol. COM-37, pp. 360–366, April 1989.

[25] A. Orda and R. Rom, "Routing With Duplication and Elimination in Computer Networks," *IEEE Trans. on Commun.*, vol. COM-36, pp. 860–866, July 1988.

[26] E. E. Swartzlander, Jr. and D. J. Heath, "A Routing Algorithm for Signal Processing Networks," *IEEE Trans. Comp.* vol. 28, pp. 567–572, 1979.

[27] E. E. Swartzlander, Jr. *VLSI Signal Processing Systems,* Hingham, MA: Kluwer Academic Publishers, 1986.

[28] M.-S. Chen, K. G. Shin, and D. D. Kandular, "Addressing, Routing, and Broadcasting in Hexagonal Mesh Multiprocessors," *IEEE Trans. Computers*, vol. C-39, pp. 10–18, Jan. 1990.

[29] D. P. Bertsekas, B. Gendron and and W. K. Tsai, "Implementation of an Optimal Multicommodity Network Flow Algorithm Based on Gradient Projection and a Path Flow Formulation," Mass. Inst. Technol., Cambridge, MA, LIDS Rep., p. 1364, 1984.

[30] T. L. Casavant, J. G. Kuhl, "A Taxonomy of Scheduling in General-Purpose Distributed Computing Systems," *IEEE Trans. on Software Engineering,* vol. SE-14, no. 2, pp. 141–154, Feb. 1988.

[31] A. M. Mood, F. A. Graybill, and D. C. Boes, *Introduction to the Theory of Statistics,* New York, NY: McGraw-Hill Book Co., 1974.

[32] IMSL Library, IMSL Inc., Houston, Texas.

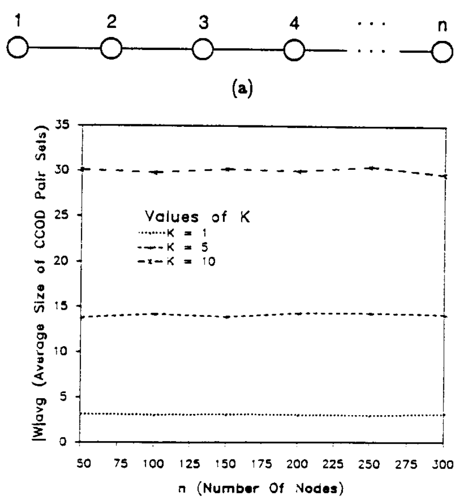

Fig. 1 (a) The n-node linear network. (b) The results of simulation studies for the linear network. Note that the value of $|W|_{avg}$ is independent of n.

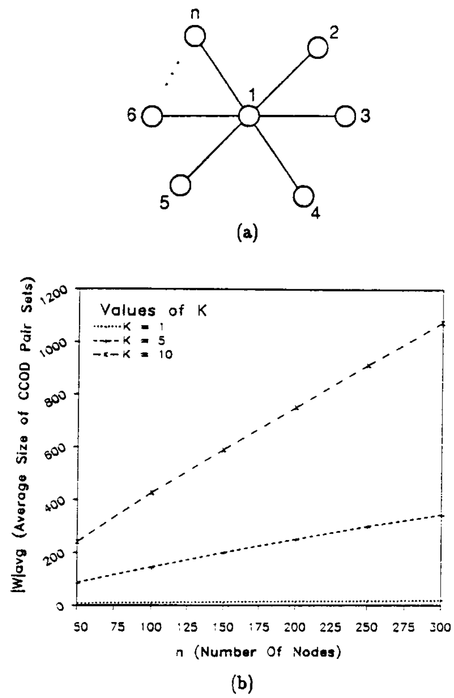

Fig. 2 (a) The n-node star network. (b) The results of simulation studies for the star network. Note that the value of $|W|_{avg}$ increases linearly with n.

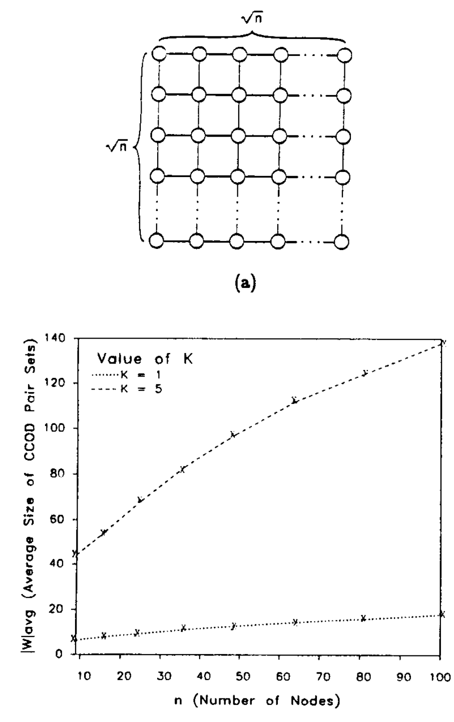

Fig. 3 (a) The n-node mesh network. (b) The results of simulation studies for the mesh network. Note that the value of $|W|_{avg}$ increases with the square root of n.

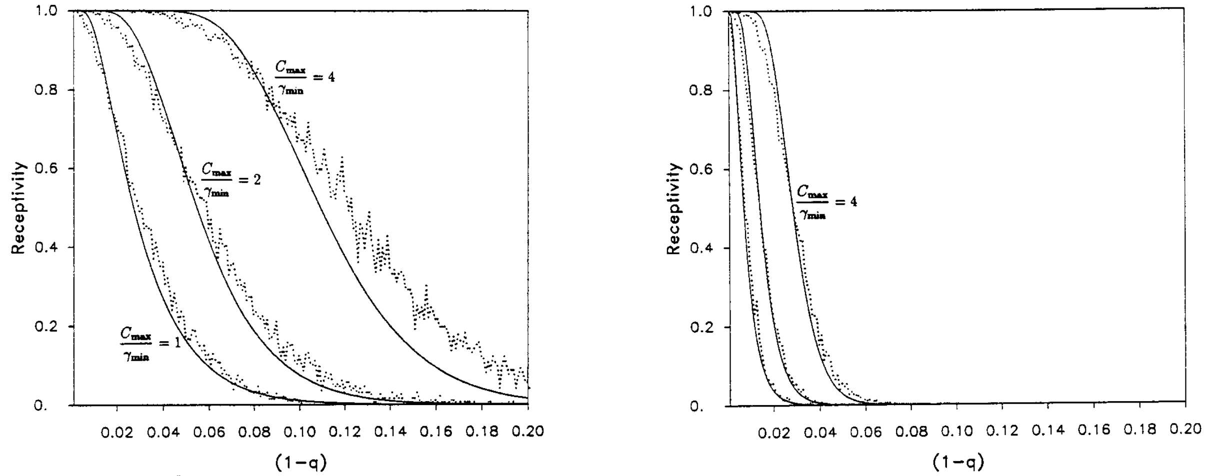

Fig. 4 Comparison of simulated and estimated values of receptivity for the linear network. Dotted lines are simulated values, solid lines are estimated values. (Left Graph: 10-nodes; Right Graph: 20-nodes).

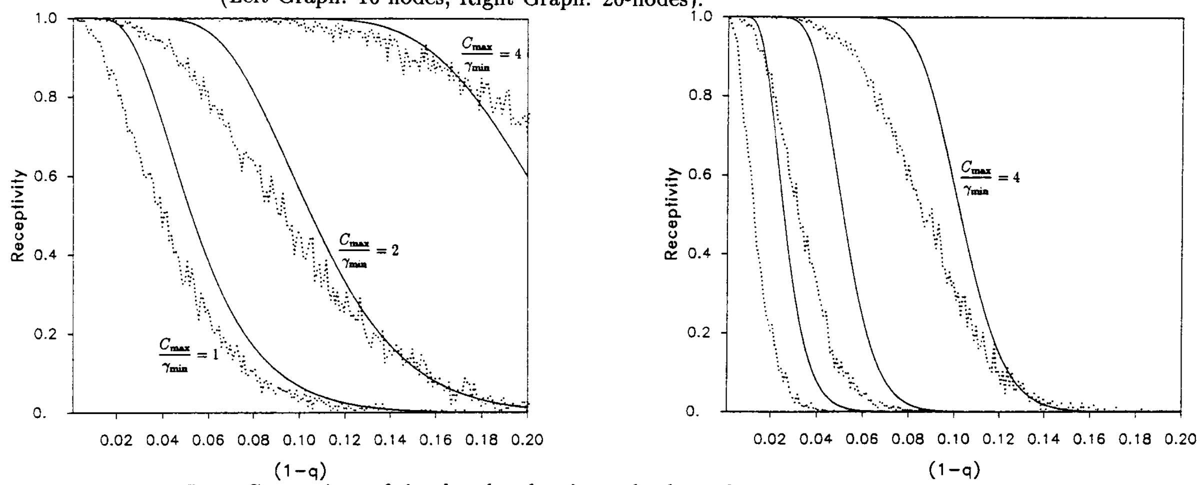

Fig. 5 Comparison of simulated and estimated values of receptivity for the star network. Dotted lines are simulated values, solid lines are estimated values. (Left Graph: 10-nodes; Right Graph: 20-nodes).

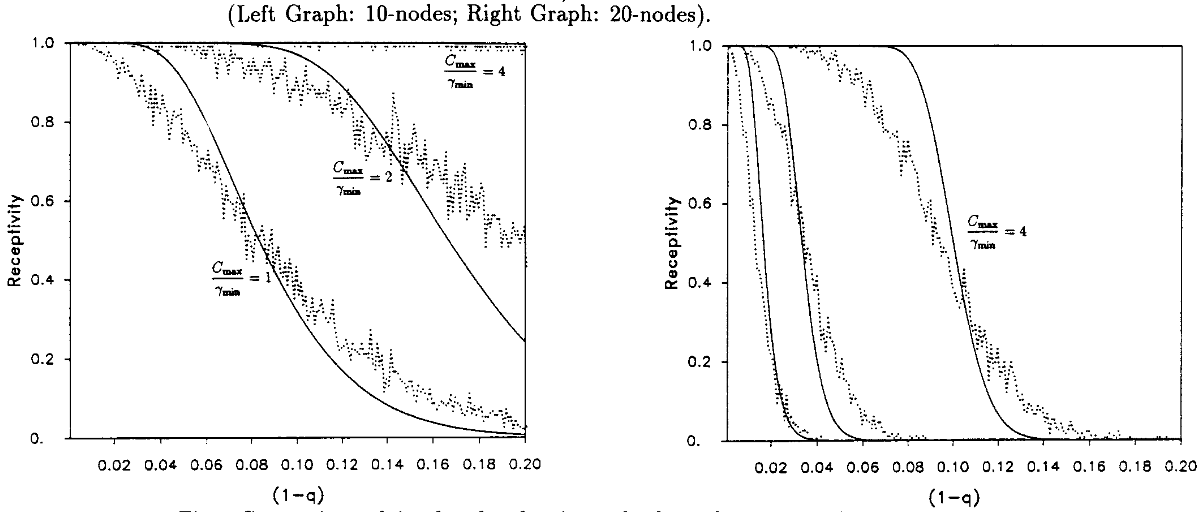

Fig. 6 Comparison of simulated and estimated values of receptivity for the mesh network. Dotted lines are simulated values, solid lines are estimated values. (Left Graph: 9-nodes; Right Graph: 25-nodes).