Adversarial Queuing Theory

ALLAN BORODIN

University of Toronto, Toronto, Ontario, Canada

JON KLEINBERG

Cornell University, Ithaca, New York

PRABHAKAR RAGHAVAN

IBM Almaden Research Center, San Jose, California

MADHU SUDAN

Massachusetts Institute of Technology, Cambridge, Massachusetts

AND

DAVID P. WILLIAMSON

IBM T. J. Watson Research Center, Yorktown Heights, New York

Abstract. We consider packet routing when packets are injected continuously into a network. We develop an adversarial theory of queuing aimed at addressing some of the restrictions inherent in probabilistic analysis and queuing theory based on time-invariant stochastic generation. We examine the stability of queuing networks and policies when the arrival process is adversarial, and provide some preliminary results in this direction. Our approach sheds light on various queuing policies in simple networks, and paves the way for a systematic study of queuing with few or no probabilistic assumptions.

Categories and Subject Descriptors: C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*store and forward networks*; G.3 [**Mathematics of Computing**]: Probability and Statistics—*queuing theory*

---

1. Introduction

Motivated by the study of dynamic continuous (or dynamic) packet routing, we introduce a new approach to the analysis of queuing systems. In the context of packet routing, our objective is to study stability and bounds on routing delay for various networks and scheduling policies. Our approach is based on an adversarial generation of packets (i.e., jobs) so that positive results (e.g., stability and upper bounds on queue sizes) are more robust in that they do not depend on particular probabilistic assumptions about the input sequences. Negative results (e.g., instability) can also be used to suggest corresponding results for more traditional stochastic queuing model assumptions. Moreover, the adversarial model may be a better (or at least safer) model of arrival processes in applications such as heterogeneous ATM networks.

New applications in communications networks and complex manufacturing systems have recently led to significant advances in queuing theory. In particular, the issue of stability for "open multiclass queuing networks" is now much better understood. In a preliminary version of our work (see Borodin et al. [1996]), we did not fully take into account the dramatic progress in this area. For example, using fluid model limits (see Dai [1995], Dai and Meyn [1995], and Rybko and Stolyar [1992]), Bramson [1996] has recently shown that for independent and time-invariant input distributions (say, for example, Poisson arrivals), FIFO scheduling is stable for *any* class-independent service time distribution (including constant service time, the standard assumption in packet routing) as long as the necessary load conditions (i.e., total expected arrival rate at any server is less than the expected service rate) are satisfied. Hence, queuing theory clearly provides a general methodology for studying continuous packet routing. However, our initial results coupled with the subsequent compelling results of Andrews et al. [2001], and other more recent results (for example, see Aiello et al. [1998], Andrews [2000], and Andrews and Zhang [2000]) demonstrate that the adversarial approach provides another useful perspective to an important and fast evolving field. In fact, as one might expect in such a well-studied field, our "new approach" is not entirely new. The "leaky bucket" model studied by Cruz [1991a; 1991b] should also be considered as a (restricted) adversarial model.

We introduce our adversarial model and for definiteness present this model in the context of packet routing. More generally this adversarial approach can be applied to any queuing network; see Tsaparas [1997]. Surprisingly, in spite of the power of the adversary, certain networks and scheduling disciplines can be proven to be "universally stable." That is, certain networks are stable under any greedy scheduling rule and certain scheduling rules are stable for any network. In contrast, we will also see that some natural scheduling disciplines can be quite ill-behaved.

2. Related Work

Queuing theory and analysis is, of course, a well-developed and highly utilized field of study. We cannot endeavor to survey all the applicable literature. We will

only attempt to briefly review the most relevant results concerning multiclass[1] queuing networks. We will also briefly review some relevant results from the field of packet routing. For basic definitions in queuing theory, we refer the reader to standard texts such as Kelly [1979], Kleinrock [1975], and Walrand [1988]. The reader may wish to skip this section and proceed directly to the new definitions in Section 3.

2.1. MULTICLASS QUEUING NETWORKS. We can view dynamic packet routing as one restrictive (but still quite important and nontrivial) type of multiclass queuing network. In particular, while oblivious routing necessitates different classes[2] (i.e., one for each path), the service time distributions are all identical and independent of the class as well as independent of the number of jobs (i.e., packets) waiting for service at any edge.

In a series of papers beginning with the work of Lu and Kumar [1991] and Rybko and Stolyar [1992], it was shown that the load conditions are not in general sufficient to guarantee stability. These initial "counterexamples" (to what seems a very natural conjecture) were first derived for scheduling rules based on the priority of a class. Seidman [1994] (for deterministic arrivals and service times) and then Bramson [1994a] for Poisson arrivals and exponential service times, showed that even FIFO can be made to be unstable. Indeed, Bramson [1994b] showed that FIFO can be unstable at an arbitrarily small ratio of arrival rate to service rate. All these examples require that the service time distribution at a given server is class dependent.

Rybko and Stolyar [1992] were able to analyze the stability of a particular network by first arguing about stability in a "fluid model abstraction" of this network. This approach was then formalized as a general approach in the work of Dai [1995] and Dai and Meyn [1995]. They show a "meta theorem" whereby stability in the fluid model abstraction of a queuing network implies stability in the given queuing network. We note that the fluid model abstraction only considers the mean rate of service disregarding the nature of the service distribution. (Since the fluid model must necessarily depend on the scheduling rule and there is neither a formal definition of a scheduling rule nor a procedure showing how to transform the scheduling rule into the corresponding equations, we have chosen to call this a meta theorem. But the meaning and applicability of this result is clear.) A Kelly-type network is a possibly multiclass network but one in which the service time distribution at a server is class independent. Using the fluid model approach, Bramson [1996] has shown that FIFO scheduling in a Kelly-type network using FIFO scheduling is stable. In particular then, under mild assumptions on the nature of the (mutually independent) input distributions of the classes, continuous packet-routing networks (having constant time service) using FIFO scheduling are stable networks. The Bramson FIFO stability result for the fluid model stands in direct contrast to the result of Andrews et al. [2001], which shows that FIFO can be unstable in the adversarial model.

---

[1] In a multiclass queuing network, jobs are partitioned into different classes with the understanding that all jobs in a class are indistinguishable in the sense that they have the same service requirements.
[2] Note that from a queuing theory perspective, the Bernoulli (or Markovian) network assumption discussed below avoids the need for multiple classes.

2.2. THE CRUZ PERMANENT SESSION MODEL. Cruz [1991a] introduces an input model designed to capture the burstiness of inputs in communication networks. Cruz defines a (σ_i, ρ_i) regulator as an input process that controls the rate of a particular input session (i.e., a path in the network) so that during any time interval $[t_1, t_2]$, the input traffic (for this path) is bounded by $\sigma_i + \rho_i(t_2 - t_1)$ units of traffic. The rate ρ_i corresponds to the long-term average rate of flow while the constant σ_i bounds the burstiness of the input. Such an input process can be viewed as the output of a "leaky bucket model" of flow control where packets are dropped whenever the above constraints cannot be satisfied. The total induced traffic rate $\rho(e)$ on any edge e is defined as $\Sigma_{i:e \in P_i} \rho_i$. Andrews [2000] refers to this input model as a (σ, ρ) regulated session model. Following Andrews and Zhang [2000], we will refer to this input model as the permanent session model (in contrast to the adversarial model in which sessions are viewed as temporary). Cruz [1991a; 1991b] develops some basic properties of such an input model including the study of stability for acyclic networks relative to a class of greedy scheduling rules.

Tassiulas and Georgiadis [1996] adopt the permanent session model for the analysis of routing packets[3] in a (unidirectional) ring G. They show that for any greedy scheduling rule $\mathcal{S}$ that the system $(G, \mathcal{A}, \mathcal{S})$ is stable for any adversary $\mathcal{A}$ that corresponds to the (σ, ρ) permanent session model under the load constraint that $\rho(e) < 1$ for every edge e. Although the concept of an adversary is implicit, these results by Cruz, and Tassiulas and Georgiadis, can be viewed as initial results in the development of an adversarial queuing theory. Following the adversarial FIFO instability result of Andrews et al. [2001], a recent result of Andrews [2000] proves that FIFO can also be unstable in the Cruz permanent session model.

For $\rho < 1$, we can model the permanent session model as a (w, ρ') adversary for any $\rho' > \rho$ and a sufficiently big window w. Note that this corresponds to a (restricted) adversary which is controlling individual input streams (i.e., particular paths in the network) rather than an adversary that is globally controlling the entire input process. The more general (w, ρ) adversary thus intuitively seems to provide a better model for ATM networks having heterogeneous and frequently changing rates of traffic. Recently Andrews and Zhang [2000] have shown that there are known scheduling rules (namely, Generalized Processor Sharing and Weighted Fair Queuing), which are unstable (for the "baseball graph" introduced by Andrews et al. [2001]) in the full adversarial model but which (by results of Parekh and Gallager [1993; 1994]) are universally stable in the permanent session model. Our DAG result and the unidirectional ring result of Andrews et al. [2001] therefore provide (respectively) a significant generalization of the Cruz and Tassiulas and Georgiadis results.

2.3. STATIC ROUTING. There has been recent work showing how stability results and delay bounds for continuous packet routing may be derived from analogous results in the *static* setting. In a static routing problem, we are given a finite set of packets, each with an assigned path, and we wish to schedule the

---

[3] To be more precise, the scheduling policies in Tassiulas and Georgiadis [1996] also allow processor-sharing rules and the transmission of packets viewed as streams of traffic that can be routed in a cut-through mode as well as discrete packets that are transmitted in a store-and-forward mode. This added generality is not important for our purposes.

motion of each packet along its respective path so as to minimize the maximum arrival time. A seminal result in this regard is the well-known bound of Leighton [1994] that every static packet routing problem in which the maximum path length is D and the maximum number of packets using an edge is C has a schedule in which each packet reaches its destination within $O(C + D)$ units of time. A centralized, polynomial-time algorithm to compute such a schedule is given by Leighton et al. [1999]. Rabani and Tardos [1996] and Ostrovsky and Rabani [1997] developed distributed algorithms for the static problem of Leighton, Maggs, Rao with bounds that approach the $O(C + D)$ achieved by the centralized methods of Leighton et al. [1994; 1999].

The work of Rabani and Tardos [1996] establishes a connection with the adversarial model developed in this paper, in the following sense. Consider any algorithm for the n-packet static problem with parameters C and D defined above that produces a schedule with completion time $(1 + \delta)C + g(n)D + f(n)$, where $\delta > 0$ is a constant and $f(n) \geq \log n$. Then it can be converted into an algorithm for continuous packet routing that is stable against an arbitrary $(w, 1 - \varepsilon)$ adversary, where ε is a function of δ and each packet has an inverse polynomial probability of being "dropped" before reaching its destination. This latter condition, that an algorithm may drop packets with small probability, is not present in our model.

Broder et al. [1997] also present a general method for transforming static routing algorithms into continuous ones. They focus on a model in which the switching nodes in the network have bounded buffers. Here a distinction is drawn between *input nodes* in the network, where packets are generated, and all other *intermediate nodes* of the network. The queue at an input node may grow arbitrarily large, but the queue at an intermediate node v can grow no larger than an absolute bound B, after which packets cannot enter v. Broder et al. [1997] show that any static packet routing algorithm in this model that satisfies some technical conditions can be transformed into a stable algorithm for continuous packet routing, against an arbitrary adversary of sufficiently small constant rate.

The connections between the static and continuous settings appears implicitly in the development of other protocols in adversarial queuing theory. The analysis showing a polynomial bound on queue size for the randomized universally stable protocol of Andrews et al. [2001], for example, builds on techniques used in the proof of the static result of Leighton et al. [1994]; the protocol, however, itself has a simple description that is independent of the analysis.

2.4. CONTINUOUS PACKET ROUTING. Without any explicit use of queuing theory results, Leighton [1990] analyzes one-bend routing on $n \times n$ arrays; the paths in one-bend routing are acyclic and in fact one-bend routing on arrays turns out to be easier to analyze than routing on cyclic networks such as rings. Leighton considers the case where each injected packet has a random destination[4] and packets are injected at each node according to a Bernoulli distribution with rate $\alpha < 4/n$.

---

[4] We note that continuous packet routing results to date are restricted to the case in which packets have random destinations. One of our main goals is to extend such results to the case where a packet is given a specified destination and path at the time of its injection.

These assumptions (on the routing scheme, on the injection rate at each node, and on the fact that packets have random destinations) imply that the induced traffic rate on any edge is less than one (the service rate of each edge). Leighton [1990] is able to provide a detailed analysis that shows that for the "farthest-to-go" scheduling discipline at each edge queue, the network is stable and "with high probability" queue sizes are bounded by a small constant. Moreover, the expected delay of a packet is bounded by a constant with high probability. These results are strengthened in Kahale and Leighton [1995] where the same results are proven for *any* scheduling discipline. Somewhat weaker results are derived when the underlying graph is a ring. For the ring, stability and bounded delay results are shown for FTG scheduling. The precise nature of these packet delay results seems beyond what we can hope to derive from applying general queuing theory results. On the other hand, the Leighton and Kahale–Leighton results are specific to certain networks and utilize the assumption of random destinations.

Stamoulis and Tsitsiklis [1991] consider the case of layered networks under the assumption of Bernoulli routing. In queuing networks with Bernoulli routing, jobs are indistinguishable (i.e., a single class network) and the next server taken is a Markov process (i.e., a probabilistic function of the last server and independent of previous history). In the context of packet routing, the Bernoulli assumption means that the next edge to be traversed is a random function of the last edge traversed and is independent of the packet identity. Stamoulis and Tsitsiklis consider such networks under the assumption of Poisson arrivals. They observe that for layered Bernoulli routing networks, the distribution on the network states (i.e., the queue size at each edge) that results from a network with constant time edge traversal and FIFO scheduling is statistically dominated by the state distribution obtained using a processor (i.e., edge) sharing scheduling discipline. They apply this observation to random destination routing in layered networks (e.g., the butterfly) and in networks that can be layered (e.g., dimension-by-dimension routing to random destinations in the hypercube).[5] Thus bounds on expected queue sizes for constant-time edge traversal can be inferred from results about the analogous network which assumes processor sharing for edge traversal. Since the latter assumption results in a product form network, standard queuing theory analysis can yield constant bounds on expected queue sizes and from this follow (by Little's Theorem) bounds on the expected time in the network.

Following similar experiments by Mitra and Cieslak [1987] for the Omega network, Harchol-Balter and Black [1994] simulate packet routing on array networks and conjecture that the queue sizes that obtain under the assumption of exponentially-distributed edge-traversal times are an upper bound for the queue sizes obtained with constant-time edge traversal. Mitzenmacher [1994] proves this conjecture for one bend routing on arrays. It is tempting to believe that the experiments and conjecture of Harchol-Balter and Black [1994] apply to a much wider context. Indeed, independent of the results in packet routing, there are many queuing-theoretic results pertaining to generalized Jackson networks. In particular, under the appropriate load conditions and some very general assump-

---

[5] Stamoulis and Tsitsiklis [1991] allow a more general, non-uniform, selection of random destinations than in Leighton [1990] and Kahale and Leighton [1995]; e.g., nearby destinations can be more probable.

tions on the arrival and service time distributions, Meyn and Down [1994] establish stability for such networks. These results apply to the more general case that each server has its own service time distribution. However, Harchol-Balter and Wolfe [1995] give evidence that it will not be a simple task to apply the approach in Stamoulis and Tsitsiklis [1991] to the general study of dynamic packet routing. First, they show that the "layered" assumption in Stamoulis and Tsitsiklis [1991] is not necessary as they are able to derive the same results for any Bernoulli routing network. But they also show that without the Bernoulli network assumption, it is no longer necessarily true that the set of delays for FIFO with constant-time servers is stochastically dominated by processor-sharing servers.

3. *The Adversarial Model*

We begin with an informal discussion of packet routing, adversaries and scheduling policies. A *routing network* is a directed graph. Time proceeds in discrete *steps*. A *packet* is an atomic entity that resides at a node at the end of any step. A packet must travel along a path in the network from its *source* to its *destination*, both of which are nodes in the network. When the packet reaches its destination, we say that it is *absorbed*. During each step, a packet may be sent from its current node along one of the outgoing edges from that node. At most one packet may travel along any edge of the network in a step. Any packets that wish to travel along an edge e at a particular time step but are not sent wait in a *queue* for edge e. The *delay* of a packet is the number of steps which the packet spends waiting in queues.

At each step, an *adversary* generates a set of requests. In this paper, a *request* is a path specifying the route followed by a packet. We say that the adversary *injects* a set of packets when it generates a set of requested paths. We restrict ourselves to the case in which the path traversed by each packet is fixed at the time of injection (i.e., nonadaptive routing), so as to be able to focus on the queuing rather than routing aspects of the problem. (Recently, Aiello et al. [1998] have successfully extended the adversarial model to adaptive routing.)

Clearly an unrestricted adversary can flood the network with packets, demanding more bandwidth than available. Therefore, we need to introduce a restriction analogous to the load condition imposed in queuing theory. Let w be an arbitrary positive integer, e any edge in the network and τ any sequence of w consecutive time steps. We define $N(\tau, e)$ to be the number of paths injected by the adversary during time interval τ that traverse edge e. It is common in packet routing to assume that all paths are simple paths. However, in more general queuing applications, one may not want to assume simple paths and hence in the definition of $N(\tau, e)$ we need not assume simple paths; thus, we would count the multiplicity of the number of times a particular path traverses e. For any $\rho > 0$, we define a (w, ρ) *adversary* that injects paths subject to the following *load condition*: for every sequence τ of w consecutive time steps and for every edge e, $N(\tau, e)/w \leq \rho$. We say that such a (w, ρ) adversary injects packets at *rate* ρ with *window size* w. A rate ρ adversary is a (w, ρ) adversary for some w.[6]

---

[6] We follow the definition of a (w, ρ) adversary as it appears in Andrews et al. [2001]. In our original paper [Borodin et al. 1996], we only defined and proved results for path-packing adversaries (to be

Clearly, if the rate ρ were greater than one, an adversary (or any input generation process) would congest some edge (since the service rate of every edge is assumed to be 1) and the network would be unstable. Hence, we will hereafter assume that $\rho \leq 1$. We note that in its full generality, there are no computational requirements on how the adversary chooses its requests in any given step and (as will be formalized below). The adversary's choice of input packets is simply a function of the history of the packet routing that has taken place thus far. One can of course place further restrictions on the nature of the adversary. For example, a (w, ρ) *path-packing adversary* is further restricted so that the paths requested at any step must be edge disjoint. A more stringent restriction (a *single path adversary*) requires that at most one packet be injected at each step. In general, we have a *request collection* of permissible request sets; at each step, the adversary must pick one request set (= set of paths) from this request collection.

Additionally, we consider stochastic adversaries. A (w, ρ) *stochastic adversary* is also specified by a request collection, a rate and a window size; now, however, at each step the adversary has a probability distribution (possibly different for each step) over its request collection, and at each step draws a set of requests from the specified distribution. Now, the quantity $N(\tau, e)$ is a random variable induced by the distributions used by the adversary during the time steps τ. If τ denotes the w time steps $t + 1, \ldots, t + w$, we let H_t denote the entire history of packet arrivals up to (and including) step t. The appropriate load condition[7] is that for every sequence τ of w consecutive time steps and for every edge e, $E[N(\tau, e)|H_t]/w \leq \rho$. In order to achieve the greatest degree of generality, we would like to impose the fewest conditions on a stochastic adversary and still be able to derive stability results. Once again, it is clear that $\rho \leq 1$ is necessary to avoid flooding some edge. Moreover, it is not hard to see that even for the simplest network consisting of one edge, a rate 1 stochastic adversary can be unstable. Simply consider an adversary which injects zero packets with probability 1/2 and injects two packets with probability 1/2; after t steps the expected queue size (which is acting like a random walk on the line $[0, \infty)$) is approximately $\sqrt{t}$. Hence, we will be concerned with stochastic adversaries having rate $\rho < 1$. However, this bound on the expectation (i.e., the first moment) is not sufficient for our purposes and we will also have to impose a bound on the pth moment (see Lemma 2) for $p > 2$. We will then say that a (w, ρ) stochastic adversary is *properly bounded* if there exist constants $p > 2$, $\epsilon > 0$ and V such that for all sequences τ of w consecutive time steps $t + 1, \ldots, t + w$ and all edges e, $E[N(\tau, e)|H_t]/w \leq 1 - \epsilon$ and $E[N(\tau, e)^p|H_t] \leq V$. Note that this model is quite general as it allows the adversary to adaptively modify the distribution at each time step. In particular, it includes the special case (as is often assumed in queuing theory) of a fixed, time-invariant input distribution (e.g., say a Poisson or constant rate arrival process) for each possible request. It also subsumes the case of oblivious packet routing for packets that are generated

defined) although we implicitly suggested a more general adversary. In subsequent discussions, we independently adopted the general (w, ρ) adversarial framework. Stability proofs for the general adversary that are presented in this paper are extended versions of our earlier proofs given for path-packing adversaries.

[7] In our conference paper [Borodin et al., 1996], we incorrectly stated the load condition as the unconditional expectation $E[N(\tau, e)]/w \leq \rho$.

at each node with randomly chosen destinations (e.g., as studied in Kahale and Leighton [1995], Leighton [1990], and Stamoulis and Tsitsiklis [1991]). To differentiate such stochastic adversaries from the nonstochastic adversaries defined above, we will refer to the nonstochastic adversaries as *deterministic* adversaries. Clearly, deterministic adversaries are a special case of stochastic adversaries.

To the best of our knowledge, these adversarial models provide the first framework for studying queuing models where we do not essentially assume independent input streams. Moreover, the analysis of these models provide not only results about stability but also quantitative bounds on queue lengths and delays.

A *scheduling policy* specifies, for each edge *e* and each time step, which packet (amongst those waiting) is to be moved along edge *e*. A *greedy* scheduling policy (called a *work-conserving policy* in the terminology of queuing theory) always specifies some packet to move along edge *e* if there are packets waiting to use edge *e*. In this paper, we restrict ourselves to deterministic greedy scheduling policies. Examples of natural greedy scheduling policies include the following:

—FIFO (*First-In-First-Out*). This is also called FCFS (First-Come-First-Served).
—LIS (*Longest-In-System*). A packet originates at a specified time, and priority is given to the packet that has been in the network for the longest amount of time.
—NIS (*Newest-In-System*). Similar to LIS but now priority is given to the packet that has spent the least amount of time in the network.
—FTG (*Furthest-To-Go*). Each packet has a prescribed path, and priority is given to the packet which has the largest number of edges still to be traversed.
—NTG (*Nearest-To-Go*). Similar to FTG but now priority is given to the packet that has the smallest number of edges still to be traversed.

All of these scheduling policies require some tie-breaking rule in order to be unambiguously defined. For our purposes, we can assume that whenever we are proving a positive (e.g., stability) result, the adversary can break the tie. For a negative result, we can assume any well-determined tie-breaking rule. It should also be clear that all of these scheduling policies can be extended to more general queuing networks. For example, FTG would be extended to MTTG (Most-service-Time-To-Go) where priority is given to that job which has the most service time remaining (summing the required service over all servers still scheduled to be visited).

In order to make the above discussion more precise, we need some additional definitions. A *packet P* is a triple $(ID, p, \tilde{t})$ where ID is some unique identifier, p is the path (from the origin to the destination of the packet) that the packet must traverse, and $\tilde{t}$ is the time that the packet was injected into the network (at its origin). The $state_t(P)$ of a packet P at time t is the vector $(P, i, t_1, \ldots, t_i)$ where i is the number of time steps in which P has traversed an edge in its path and t_j is the time of the jth such traversal. We define the *configuration* $config_t(G)$ of a network G at time t as $config_t(G) = \{state_t(P_k) | P_k$ is a packet that is present in the network at time $t\}$. Note that the configuration implicitly specifies the states of packets in each (edge) queue of the network. In particular, we will assume that time begins at step $t = 0$ and thus $config_0(G)$ denotes the initial configuration of the network. We can then formally define a *deterministic adversary* $\mathcal{A}$ as a function $\mathcal{A} : \bigcup_{j=0}^{\infty} [j \times \Pi_{t=0}^{j} config_t(G)] \to \{P_k | P_k$ is a new

packet}; that is, $\mathcal{A}(\tilde{t}, \Pi_{t=0}^{\tilde{t}} \, config_t(G))$ specifies the new packets that the adversary $\mathcal{A}$ injects into the network at time $\tilde{t}$ and it does so based on the entire history of the routing up to this point of time. A *stochastic adversary* $\mathcal{A}_S$ is similarly defined as a function $\mathcal{A}_S \colon \bigcup_{j=0}^{\infty} [j \times \Pi_{t=0}^{j} \, config_t(G)] \rightarrow \{D_k | D_k$ is a distribution on new packets}. That is, a stochastic adversary chooses a distribution for the next set of requests and this distribution is a function of the history of the network configurations. At any time t, once the distribution has been chosen, the inputs are then randomly generated according to this distribution. A *deterministic scheduling policy* $\mathcal{S}$ is defined as having the same domain as a deterministic adversary with its range being the set of packets that exist in the system at time $\tilde{t}$; that is, in its full generality we can consider global scheduling rules based on the entire history to date. However, in practice and for the purpose of this paper, we will mainly be concerned with distributed scheduling policies where the packet selected to traverse any given edge e is determined by the states of the packets that are presently in the queue associated with edge e. We note that all the known deterministic scheduling rules fit into this framework. We could also easily define randomized scheduling policies but will not do so since we do not consider such policies in this paper. However, we note that randomized scheduling policies are often utilized in packet routing; see, for example, Rabani and Tardos [1996].

Our main goal will be to prove the *stability* of arbitrary and particular greedy scheduling policies $\mathcal{S}$ on various networks G and against various classes Υ of adversaries.[8] Similar to Andrews et al. [2001], we define a *network system* as the tuple $(G, \Upsilon, \mathcal{S})$. We will say that a network system $(G, \Upsilon, \mathcal{S})$ is *stable*, if for every initial configuration $C_0(G)$ there is a constant M (which may depend on the size of the network G, the initial configuration and the rate and window parameters of the adversary class Υ) such that for every adversary $\mathcal{A}$ in the class of adversaries Υ, when the network system is executed with initial configuration $C_0(G)$ against adversary $\mathcal{A}$ (i.e., is executed using the deterministic adversary $\mathcal{A}$ to generate packets and the scheduling policy $\mathcal{S}$ to route packets), the maximum number of packets in any queue is bounded by M.[9] For stochastic adversaries, we say that the network system $(G, \Upsilon, \mathcal{S})$ is stable if for every initial configuration $C_0(G)$ there is a constant M as above such that for every stochastic adversary $\mathcal{A}_S$ in the class Υ, when the network system is executed with initial configuration $C_0(G)$ against adversary $\mathcal{A}_S$ then at all times the expected number of packets in any queue is bounded by M. We will say that a network G is *universally stable* with respect to a class Υ of adversaries if $(G, \Upsilon, \mathcal{S})$ is stable for all greedy scheduling policies, and similarly a scheduling policy $\mathcal{S}$ is universally stable if $(G, \Upsilon, \mathcal{S})$ is stable for all networks G.[10] Andrews et al. [2001] show that for

[8] For example, certain networks G are stable against the class of stochastic adversaries of rate $\rho < 1$ for any greedy scheduling rule.

[9] Equivalently, we can say that the maximum number of packets in the system is bounded by some constant M. In queuing theory one sometimes defines stability in terms of the existence of a limiting stationary distribution for the state of the network. Our definition is consistent with other standard uses of the term and furthermore one cannot hope for a limiting distribution in an adversarial model.

[10] For the study of the universal stability of a scheduling rule, we can assume that the initial configuration is empty. Informally, we argue that any network can be modified by appending "input trees" to each node and using these input trees to adversarially construct any desired initial configuration.

undirected networks (i.e., in the sense that each undirected edge represents two directed edges) it is decidable if a packet routing network is universally stable. Based on Andrews et al. [2001], Goel [1997] gives a simple characterization for the universal stability of directed and undirected networks. Bertsimas et al. [1996] use linear programming to decide universal stability for all 2-station fluid model networks and conjecture that such a test exists for all fluid model queuing networks.

Rather than define stability in terms of maximum queue size, one could also define stability in terms of the maximum delay incurred by any packet. Clearly, if no packet is delayed by more than M steps then the maximum queue size is also bounded by M. Conversely, for any (w, ρ) adversary with rate $\rho < 1$, if the maximum queue size is bounded by M, then a packet can be delayed in any queue by at most cw steps where c is large enough that $cw > M/(1 - \rho)$. (In cw steps, at most $cw\rho$ packets enter the queue and cw packets leave the queue while it remains nonempty.) Clearly, when $\rho = 1$, it is possible to have a given packet delayed forever even though the system is stable under the given definition.

4. *Results*

It is perhaps tempting to conjecture that for every network G, every greedy queuing discipline is stable for a deterministic adversary with injection rate $\rho < 1$ (or even $\rho = 1$). Our first result shows that every directed acyclic network is universally stable for rate 1 deterministic adversaries.

(1) *If G is a DAG, and Υ is the class of deterministic rate 1 adversaries, then $(G, \Upsilon, \mathcal{S})$ is stable for every greedy queuing discipline $\mathcal{S}$. That is, every DAG is universally stable against the class of rate 1 deterministic adversaries.*

Unfortunately, universal stability against rate 1 adversaries does not extend to the case of graphs with directed cycles, as our next result shows.

(2) *Let G be a unidirectional ring network on $n \geq 3$ nodes and let Υ be the class of deterministic rate 1 single path adversaries. Then the network system (G, Υ, LIS) is not stable. Indeed, for every initial configuration, there is a rate 1 adversary $\mathcal{A}$, injecting a single path on every step, which will force the network to have an unbounded number of packets. Similarly, FIFO is unstable for rate 1 single path adversaries.*

This instability result is in contrast with the next result.

(3) *If G is a unidirectional ring and Υ is the class of deterministic rate 1 adversaries, then (G, Υ, FTG) is stable.*

All our stability results also hold for stochastic adversaries of bounded variance for injection rates bounded away from 1.

In our conference paper [Borodin et al. 1996], we asked if our stability results for the ring and DAGs could be extended to all networks and all greedy scheduling rules. Our results were soon significantly extended by Andrews et al. [2001]. They show that the ring is universally stable with respect to the class of rate $\rho < 1$ deterministic adversaries. They also show that certain scheduling policies $\mathcal{S}$ (such as Newest-in-System (NIS), LIS and FTG) are *universally stable* against the class of rate $\rho < 1$ deterministic adversaries. However, they demon-

strate that certain common scheduling policies (such as FIFO and Nearest To Go (NTG)) are *not* universally stable. Specifically, for FIFO (respectively, NTG) scheduling, they show that there is some network and initial configuration for which some rate $\rho > 0.798$ (respectively, rate $\rho > 0.62$) path-packing adversary causes instability.

In hindsight, the instability of NTG (packet routing) is well motivated by the queuing instability examples[11] of Lu and Kumar [1991] and Rybko and Stolyar [1992]. Note that our instability result for FIFO on the ring with rate 1 and the FIFO instability result of Andrews et al. [2001] stands in contrast to Bramson's [1996] fluid stability result for class-independent service time networks with FIFO scheduling. We strengthen the Andrews et al. [2001] NTG instability result by showing the following.

(4) *For* every *rate $\rho > 0$, there exists a network and initial configuration for which some rate ρ adversary causes NTG to be unstable. Indeed, although motivated by initial adversarial results, this turns out not to be an adversarial result rather, the "adversary" will only set the initial configuration and then will be injecting all paths at a constant rate*!

Thus, in addition to providing a framework for obtaining very general stability results, the adversarial model is able to distinguish between different scheduling policies and to suggest instability results in more classical settings. This sensitivity to the scheduling policy is not unnatural, yet prior queuing theoretic results do not seem to highlight this.

The remainder of the paper is structured as follows. Section 5 discusses stability bounds for acyclic networks G. Section 6 discusses results for the ring. Our NTG instability result is outlined in Section 7. We conclude in Section 8 with a number of open problems.

5. *Dags and Meshes*

Perkins and Kumar [1989] show that for constant rate arrivals there is a class of scheduling rules (called "clear a fraction") with respect to which all acyclic networks are stable. Using a fluid model, Down and Meyn [1995] consider a specific acyclic network and show that it is universally stable. Again using a fluid model, Dai [1995] shows that every acyclic network is universally stable for rate less than 1. These stability results for fluid models then can be applied to obtain stability results for the more standard "stochastic queuing networks" (i.e., with time-invariant input distributions). We extend these results by showing that for packet routing, acyclic networks are universally stable in the (deterministic and stochastic) adversarial setting. Recently, Gamarnik [1998] showed how to use fluid models to show that if a packet routing network is universally stable in the fluid model, then it is universally stable in the deterministic adversarial model. However, the fluid model results do not seem to provide quantitative bounds on queue lengths.

---

[11] These queuing network examples do not show that NTG is unstable for packet routing (where we assume identical service time distributions) but can be easily modified to show packet routing instability for some priority based scheduling policy.

THEOREM 1. *Let G denote an arbitrary directed acyclic graph, $\mathcal{S}$ an arbitrary greedy protocol, and Υ the class of deterministic rate 1 adversaries. Then $(G, \Upsilon, \mathcal{S})$ is stable.*

PROOF. For e an edge of G, let $Q_t(e)$ denote the queue at edge e at time t, and let $A_t(e)$ denote the number of packets (not already absorbed) that have arrived by time t and are eventually destined to cross edge e. For any adversary $\mathcal{A}$ that injects at rate 1, there exists some window size w, such that for any window of time $(t - w, t]$ and for any edge e the adversary can inject at most w packets during this window that are destined to cross the edge e.

We define a function $\psi(\cdot)$ inductively on the edges as follows: For an edge e, suppose $f_1, \ldots, f_k$ are edges entering the tail of e. (Notice there may be no such elements and then we just take k to be zero.)

$$\psi(e) = \max\{2w, Q_0(e)\} + \sum_{i=1}^{k} \psi(f_i).$$

We claim that for all $t = l \cdot w \geq 0$ and all $e \in G$, we have

$$A_t(e) \leq \psi(e) \tag{1}$$

Note that for any time t' such that $t - w < t' < t$, $A_{t'}(e) \leq A_{t-w}(e) + w$. The theorem will then follow, since $\Sigma_e \, \psi(e)$ gives an absolute upper bound on the number of packets in the system, in terms of the initial configuration.

The proof of this claim proceeds by induction on l. The claim clearly holds when $l = 0$. Now let $t = lw$ for $l \geq 1$. Suppose e is an edge whose tail is entered by edges $f_1, \ldots, f_k$. We consider two cases.

Case 1. $A_{t-w}(e) \leq w + \Sigma_{i=1}^{k} \psi(f_i)$. In this case, we use the fact that in w time steps the number of new packets inserted that wish to cross the edge e_j is at most w. Thus, in this case

$$A_t(e) \leq A_{t-w}(e) + w \leq 2w + \sum_{i=1}^{k} \psi(f_i) \leq \psi(e).$$

Case 2. $A_{t-w}(e) > w + \Sigma_{i=1}^{k} \psi(f_i)$. By the inductive assertion we have that $A_{t-w}(f_i) \leq \psi(f_i)$. But notice that $A_{t-w}(e)$ is at most $Q_{t-w}(e) + \Sigma_{i=1}^{k} A_{t-w}(f_i)$. Thus, we have

$$Q_{t-w}(e) \geq A_{t-w}(e) - \sum_{i=1}^{k} A_{t-w}(f_i)$$

$$> w + \sum_{i=1}^{k} \psi(f_i) - \sum_{i=1}^{k} A_{t-w}(f_i)$$

$$\geq w.$$

In other words, the number of packets queued at edge e must be at least w at the beginning of time step $t - w$. Thus, in the next w time steps, at least one packet crosses the edge e in every step. But the number of packets inserted which wish to cross this edge may go up by at most w in w time steps. Thus, in this case also, we have $A_t(e) \le A_{t-w}(e) \le \psi(e)$.

The upper bound on the number of packets in the system that follows from this proof is exponential in the number of edges (more precisely, the depth) of G. Indeed, Andrews et al. [2001] show that for every m, there is an $O(m^2)$-node DAG G and a $(1 - 1/(m + 2))$ rate path-packing adversary for which the scheduling rule NIS (or FTG) can be forced to have queue size 2^{m-1}. Another example of exponentially long (in the depth but not the size) queues for DAGs can be found in Cruz [1991b]. For the special case of rooted tree networks (with edges directed away from the root), the proof of Theorem 1 gives a much better bound; in this case, all the sets $Q(e)$ have linear (in the depth) size.

It is now reasonably easy to obtain essentially the same stability result for stochastic adversaries (satisfying some minimal conditions as previously indicated). We need the following Martingale type lemma due to Pemantle and Rosenthal [1999]:

LEMMA 2. *Let* $X_1, X_2, \ldots$ *be a sequence of nonnegative random variables satisfying the following properties*:

(1) *There exists positive constants* α *and* β *such that for all* $x_1, \ldots, x_n$ *with* $x_n > \beta$,

$$E[X_{n+1} - X_n | X_1 = x_1, \ldots, X_n = x_n] \le -\alpha.$$

(2) *There exists a positive constant* θ *and a* $p > 2$ *such that for all* $x_1, \ldots, x_n$

$$E[|X_{n+1} - X_n|^p | X_1 = x_1, \ldots, X_n = x_n] \le \theta.$$

Then there exists M (M *is a function of* X_0, α, β *and* θ) *and* t_0 *such that for all* $t \ge t_0$, $E[X_t] \le M$.

An example in Pemantle and Rosenthal [1999] shows that condition (ii) above cannot be replaced by a bounded second moment. We also note that by bounding higher moments we immediately obtain improved results on the tail probabilities (i.e., the probability that X_n will exceed $c \cdot M$). Indeed, Hajek [1982] had previously shown that a bound on exponential moments (replacing condition (ii) above) yields exponentially decreasing bounds on the tail probabilities of the X_n and assuming an absolute bound on $X_{n+1} - X_n$ allows us to use basic results concerning super-Martingales (see, for example, the text by Durrett [1995]) to determine bounds on $E[X_n]$ and the tail probabilities.

THEOREM 3. *Let* G *denote an arbitrary directed acyclic graph,* $\mathcal{S}$ *an arbitrary greedy protocol, and* Υ *the class of properly bounded stochastic adversaries with rate* $1 - \epsilon$ *for some* $\epsilon > 0$. *Then* $(G, \Upsilon, \mathcal{S})$ *is stable.*

PROOF. Let $\mathcal{A}_S$ be a rate $(1 - \epsilon)$ adversary and let w be an appropriate window size; that is, in any window of w consecutive time steps and for any history of packet injections preceding this window, for every edge e, the expected number of packets injected by $\mathcal{A}_S$ that need to cross e is bounded by $(1 - \epsilon)w$. For every edge e in G, define $Q_t(e), A_t(e)$ as in the proof of Theorem 1. We

also define $\psi(e)$ in a manner similar to the proof of Theorem 1. Let $f_1, \ldots, f_k$ be edges entering e. Then $\psi(e) = \max\{2w, Q_0(e)\} + \Sigma_{l=1}^{k} (\psi(f_l) + \epsilon w)$. (Notice that for $\epsilon = 0$ this is exactly what we had in the previous proof.)

Unlike the deterministic adversary case, we will not be able to claim that $A_t(e) \leq \psi(e)$. Instead we use a potential function and show that if the potential function is larger than a specified quantity, the potential is expected to decrease in the next w time steps.

To define the potential, let $e_1, \ldots, e_m$ be a numbering of the edges in topological order (i.e., if $i < j$, then no directed path in G contains the edge e_j followed by the edge e_i). The potential function associated with e_i is defined as:

$$\phi_t(e_i) = \max\{A_t(e_i), \psi(e_i)\}.$$

The potential associated with the whole network is:

$$\Phi_t = \sum_i \lambda^{m-i} \phi_t(e_i),$$

where λ is a positive real number greater than 1 whose exact value will be chosen later.

Following the notation used in the definition of a stochastic adversary, let H_{t-w} denote the entire history of packet injections up to and including step $t - w$. In order to prove that the expectation of this potential function must decrease if it is too large, we prove the following bounds.

CLAIM 4

(1) $A_t(e_i) \geq A_{t-w}(e_i) - w$
(2) $E[A_t(e_i)|H_{t-w}] \leq A_{t-w}(e_i) + (1 - \epsilon)w.$
(3) *Let i be the smallest index (if it exists) such that for every $j < i$, $A_{t-w}(e_j) \leq \psi(e_j)$ and $A_{t-w}(e_i) > \psi(e_i)$. (If such an i does not exist, then say $i = m + 1$.) Then for every $j < i$, $E[A_t(e_j)] \leq \phi_{t-w}(e_j)$ and (if $i \leq m$) $E[A_t(e_i)|H_{t-w}] \leq A_{t-w}(e_i) - \epsilon w$.*

PROOF. Part (1) follows from the fact that at most one packet can cross any edge in a given time step. Part (2) follows from the fact that the expected number of packets that are injected in time $(t - w, t]$ and wish to cross the edge e_i is at most $(1 - \epsilon)w$. To show part (3) we use an induction argument similar to that in the proof of Theorem 1. Suppose for some $j \leq i$ that e_j is an edge whose tail is entered by edges $f_1, \ldots, f_k$. We consider two cases.

Case 1. $A_{t-w}(e_j) \leq w + \Sigma_{l=1}^{k} (\psi(f_l) + \epsilon w)$. Notice first that in this case that $A_{t-w}(e_j) \leq \psi(e_j) - w$ and (since $j < i$) we only need to show that $E[A_t(e_j)] \leq \phi_{t-w}(e_j)$. We use the fact that in w time steps the expected number of newly inserted packets that wish to cross the edge e is at most $(1 - \epsilon)w \leq w$. Thus, in this case

$$E[A_t(e_j)|H_{t-w}] < A_{t-w}(e_j) + w \leq \psi(e_j) - w + w = \psi(e_j) \leq \phi_{t-w}(e_j).$$

Case 2. $A_{t-w}(e_j) > w + \Sigma_{l=1}^{k} (\psi(f_l) + \epsilon w)$. We will show in this case that $E[A_t(e_j)|H_{t-w}] \leq A_{t-w}(e_j) - \epsilon w$. Notice that this is sufficient to prove the assertion for both cases $j < i$ and $j = i$.

Since $f_l = e_{j'}$ for some $j' < i$, by the definition of i, we have that $A_{t-w}(f_l) \le \psi(f_l) < \psi(f_l) + \epsilon w$. But $A_{t-w}(e_j)$ is at most $Q_{t-w}(e_j) + \Sigma_{l=1}^k A_{t-w}(f_l)$. Thus we have

$$Q_{t-w}(e_j) \ge A_{t-w}(e_j) - \sum_{l=1}^k A_{t-w}(f_l)$$

$$> w + \sum_{l=1}^k (\psi(f_l) + \epsilon w) - \sum_{l=1}^k A_{t-w}(f_l)$$

$$\ge w.$$

In other words, the number of packets queued at edge e must be at least w at the beginning of time step $t - w$. Thus in the next w time steps one packet crosses the edge e in every step. But the expected number of packets inserted which wish to cross this edge may go up by at most $(1 - \epsilon)w$ in w time steps. Thus, in this case, we have $E[A_t(e_j)|H_{t-w}] \le A_{t-w}(e_j) - \epsilon w$. $\square$

We now conclude by observing that if $\Phi_{t-w} > \beta = \Sigma_{i=1}^m \lambda^{m-i}(\psi(e_i) + \epsilon w)$, then there must exist a smallest i such that $\phi_{t-w}(e_i) > \psi(e_i) + \epsilon w$. Using the claim above, we conclude that

$$\begin{aligned} E[\Phi_t|H_{t-w} \wedge \Phi_{t-w} > \beta] - \Phi_{t-w} &= \sum_{j<i} \lambda^{m-j}(E[\phi_t(e_j)|H_{t-w}] - \phi_{t-w}(e_j)) \\ &\quad + \lambda^{m-i}(E[\phi_t(e_i)|H_{t-w}] - \phi_{t-w}(e_i)) \\ &\quad + \sum_{j>i} \lambda^{m-j}(E[\phi_t(e_j)|H_{t-w}] - \phi_{t-w}(e_j)) \\ &\le 0 - \epsilon w \lambda^{m-i} + \sum_{j=i+1}^m (1 - \epsilon)w\lambda^{m-j} \\ &\le -\epsilon w \lambda^{m-i} + (1 - \epsilon)\frac{w\lambda^{m-i}}{\lambda - 1} \\ &= \frac{w\lambda^{m-i}}{\lambda - 1}(-\epsilon\lambda + \epsilon + 1 - \epsilon) \\ &\le \frac{-w\epsilon}{2} \text{ (provided } \lambda \ge 2/\epsilon) \end{aligned}$$

The final inequality above determines our choice of λ, which we set to $2/\epsilon$. Thus we conclude that if the potential Φ_{t-w} is high enough, then after w time steps the potential Φ_t is expected to decrease by at least $\alpha = \epsilon w/2$. It should also be clear that by assuming $\mathcal{A}_S$ is properly bounded, we also know that $E[|\Phi_t - \Phi_{t-w}|^p]$ is bounded by some constant θ for $p > 2$ since Φ_t is linear in the $A_t(e)$. Letting $X_i = \Phi_{i \cdot w}$, the theorem follows from Lemma 2 since the potential Φ_t is an upper bound on the number of packets in the system at time t.

COROLLARY 5. *Consider any* (*say two dimensional*) *mesh as a routing network and consider the case of one bend routing. That is, packets are first sent along their originating row to the destination column and then traverse along this column until reaching the destination. Then, for any scheduling rule* $\mathcal{G}$, *and for the class* Υ *of rate* $1 - \epsilon$ *stochastic* (*and hence deterministic*) *one bend adversaries, the network* ($G, \Upsilon, \mathcal{G}$) *is stable.*

PROOF. We sketch the proof for a two dimensional N by N mesh and deterministic adversaries. The proof is similar to the argument used by Kahale and Leighton [1995] when they consider one bend routing in the context of Bernoulli distributed inputs destined to random destinations. Let $\epsilon > 0$ and consider a ($w_1, 1 - \epsilon$) adversary for any window w_1. We can consider packets traversing a row as if they are traversing a one dimensional line (i.e., in each direction a very restricted DAG) with inputs generated by the adversary. By Theorem 1, all packets reach their column destination within $c(w_1 + N)$ steps for some constant c. Now consider the packets that enter a particular column (say column j) during any interval [$t, t + w_2$] of w_2 time steps. Any such packet has either arrived during this window of time or was generated at a step $t' \in [t - c(w_1 + N), t)$. There are therefore at most $c(w_1 + N)(1 - \epsilon) + w_2(1 - \epsilon)$ packets in the network that are destined to traverse any particular edge in column j during the interval [$t, t + w_2$]. Now for any $\epsilon' < \epsilon$, $c(w_1 + N)(1 - \epsilon) + w_2(1 - \epsilon) < w_2(1 - \epsilon')$ for w_2 sufficiently large. With regard to column j, we can think of the ($w_1, 1 - \epsilon$) adversary as a ($w_2, 1 - \epsilon'$) adversary. Applying Theorem 1 again (for column j thought of as a line), every packet that enters column j at time t will reach its destination by time $t + c(w_2 + N)$. □

This proof easily generalizes to any dimensional array and also to stochastic adversaries. The same idea can be applied to toroidal networks using the Andrews et al. [2001] universal stability results for unidirectional rings. However without the one bend routing assumption, the instability results of Andrews et al. [2001] show that (two-dimensional) meshes are not universally stable networks.

6. *The Ring*

There are now a number of independent results (for different queuing model assumptions) showing that the (unidirectional) ring is universally stable for rate < 1 (i.e., the total rate of service required at any server is less than 1). For class independent service rates, Dai and Weiss [1996] prove universal stability for fluid models (and hence for time-invariant stochastic queuing networks under the same assumption of class independent service rates). Tassiulas and Georgiadis [1996] establish the analogous result for packet routing using the Cruz [1991a] leaky bucket model. Andrews et al. [2001] show that the ring is universally stable with respect to deterministic adversaries. We shall now show that the scheduling rule FTG is stable at rate 1 for the ring. In contrast, neither LIS (which is a universally stable scheduling rule at any rate less than 1) nor FIFO are stable at rate 1 for the ring. We then extend the FTG proof to stochastic adversaries with rate $1 - \epsilon$.

Throughout this section, our underlying graph G will be the n-node unidirectional cycle, with vertices numbered $0, \ldots, n - 1$. For the purpose of packet routing the unidirectional assumption is not usually a restriction as we most often

assume simple one directional paths and clearly paths in the clockwise direction will not interfere with paths in the counterclockwise direction. For definiteness let us assume packets are being routed in a clockwise direction. Even this case is quite non-trivial, both from the point of view of classical queuing theory and within our adversarial setting. Indeed as demonstrated in Andrews et al. [2001], a "slight" extension of the ring network (allowing two edges between each pair of adjacent nodes) is enough to show that common scheduling policies like FIFO and NTG (nearest-to-go) can be unstable.

6.1. INSTABILITY OF LIS AND FIFO AT RATE 1.   We show in Section 6.2 that the FTG protocol is stable at injection rate 1 on the ring. On the other hand, we now exhibit simple adversaries with injection rate 1 that cause instability on the ring for the Longest-in-System protocol (priority to the packet that was injected longest ago) and the FIFO protocol (queues are maintained in First-Come-First-Served fashion).

THEOREM 6.   *There is a deterministic adversary $\mathscr{A}$* (*respectively, an adversary $\mathscr{A}'$*) *that injects single paths onto the ring at rate* 1, *such that $\mathscr{A}$* (*respectively, $\mathscr{A}'$*) *will force the scheduling rule LIS* (*respectively, the scheduling rule FIFO*) *to have unbounded size queues*.

PROOF.   We first describe the adversary $\mathscr{A}$ for LIS. For simplicity of presentation, assume that each path requested by $\mathscr{A}$ will be a "self-loop"—a path that traverses all the edges of the ring in sequence. It is not difficult to refine this argument so that the adversary injects shorter paths.

$\mathscr{A}$ works as follows:

—For $k = 1, 2, 3, \ldots$
—Inject kn self-loops in sequence at node 1.
—Inject kn self-loops in sequence at node 0.

It is easy to verify by induction on $k \geq 1$ that at the end of iteration k of this process, there will be one packet at node 1 (destined for node 0) and $kn - 1$ packets queued at node 0 (also destined for node 0). Thus the number of packets becomes unbounded.

The adversary for the FIFO case is similar but a bit more complicated. To simplify the presentation, we prove this case by contradiction. Say there is an absolute bound M such that the number of packets in the system is bounded by M for every adversary. Let us first show how to construct an adversary that contradicts this bound M. The adversary, $\mathscr{A}'$ first does an "initial" phase 0 and then works in phases indexed by k as follows:

—For $k = 1, 2, 3, \ldots$
—Inject $(M + 1)n$ self-loops in sequence at node $k (\bmod n)$.

The invariant that will be established is that at the end of phase k all packets in the network are destined for node k with one packet at each queue other than the queue on the edge from $k \to k + 1$, which has Q_k packets in its queue where the Q_k's form a monotone increasing sequence in k. Thus, by the end of M phases we derive a contradiction.

We assume an empty initial configuration. For phase 0, the adversary simply injects n self loops at node 0. Thus the invariant is initially satisfied for $k = 0$

with $Q_0 = 1$. Assume that the invariant is true at the end of phase $k - 1$. We make some observations on the transient behavior in phase k. We start by observing that once the invariant is established for some time step during phase k, it continues to hold in all subsequent time steps in phase k. Next notice that throughout this phase there is at most one packet queued in every queue other than queues k and $k - 1$. The number of packets queued at $k - 1$ is monotone nonincreasing and the number of packets queued at k is monotone non-decreasing. Furthermore, since the queues at nodes k and $k - 1$ are of length at most M, no packet waits at any queue for more than M time steps. Thus, after at most Mn time steps in phase k, every packet injected in phase 0 to $k - 1$ has reached its destination. In at most n more time steps, we reach the invariant that all the queues other than $k - 1$ and k have exactly one packet in their queue. To conclude, we need to show that the queue at node $k - 1$ does eventually drain down to having 0 and then 1 packet; and that the queue size at k at the end of phase k is strictly larger than the queue size at $k - 1$ at the beginning of this phase.

To verify the first part, notice that the queue size at node $k - 1$ (which consists of both phase $k - 1$ packets destined for node $k - 1$ as well as phase k packets destined for node k) is upper bounded by the number of packets in the system with destination $k - 1$. This observation is easily verified by induction on time. Moreover, since the queue size goes down every time a packet with destination $k - 1$ reaches its destination, the queue becomes empty (for one time step). Finally, we need to verify that Q_k, the number of packets at queue k at the end of phase k, is greater than Q_{k-1}. Assume otherwise. Then, this implies that the total work remaining in the system has not increased. But this can not be the case, since there is at least one queue, namely at node $k - 1$, that was idle for one time step (immediately after the last packet destined for $k - 1$ reached its destination) during phase k. Since n units of work are added at each time step, the only way the workload does not go up is if every queue remains non-idle in every time step in phase k. This concludes the proof for the bound M.

In order to show that there is (one) adversary that defeats every bound M, we simply keep changing the goal of the adversary so that after it defeats a given M, it resets all queues to be empty (by not doing any injections) and then proceeds to defeat $M + 1$, etc. $\square$

6.2. THE FURTHEST-TO-GO PROTOCOL. In this section, we prove that the FTG protocol is stable for the ring, first (in contrast to the LIS and FIFO instability result at rate 1) for a deterministic adversary at injection rate 1, and then for a stochastic adversary at injection rate $1 - \epsilon$. We are assuming that all packets are traveling a simple path (say in a clockwise direction).

We first define a quantity $A(i, j, t)$ for $0 \leq i, j \leq n - 1$. (Throughout this section, all arithmetic on node names is mod n.) The quantity $A(i, j, t)$ denotes the number of packets at time t in the queues in nodes $i, i + 1, \ldots, j$ (inclusive) which need to cross the edge from node $i - 1$ to node i. We assume that this quantity is measured at the end of time step t (i.e., at time step t, packets are inserted, then moved, after which the value of A is determined). Let $\mathcal{A}$ be a deterministic (w, ρ) adversary with $\rho = 1$. We next define an appropriate potential function $\phi(i, t) = \max\{\phi_1(i, t), w + n - 1\}$, where $\phi_1(i, t) =$

$\max\{A(i, k, t) + (i + n - 1 - k): i \leq k \leq i + n - 1\}$. The crux of the argument is the following lemma:

LEMMA 7. *Let $\mathcal{A}$ be a (w, ρ) adversary with $\rho = 1$. Then for all $t \geq 0$, $\phi(i, t + w) \leq \phi(i, t)$.*

PROOF. Clearly the only way $A(i, j, t)$ can increase is due to the insertion of packets and can only go up by one per insertion of a packet that needs to cross edge $(i - 1, i)$. (Note that by assuming simple paths any packet entering the queue at node i from the queue at $i - 1$ will not need to traverse the edge $(i - 1, i)$ again, and thus is not counted in $A(i, j, t)$.) Thus, after the w time steps $t + 1, t + 2, \ldots, t + w$, no $A(i, j, t)$ can increase by more than w and hence $\phi(i, t)$ will increase by at most w due to all the packet insertions during these steps.

If $\phi_1(i, t + \ell) \leq n - 1$ at any step $t + \ell$ with $0 \leq \ell \leq w - 1$, then $\phi_1(i, t + w)$ can be at most $w + n - 1$ and hence $\phi(i, t + w) = w + n - 1 \leq \phi(i, t)$. If $\phi_1(i, t + \ell) > n - 1$ throughout these w time steps, then we argue that each routing step causes ϕ_1 to decrease by one (offsetting any increase due to the injection step). To see this, note that there must be at least one packet in the system at every time step by the definition of ϕ_1. Let k be any index that maximizes the quantity $B = \max\{A(i, k, t) + (i + n - 1 - k): i \leq k \leq i + n - 1\}$. If $k > i$, the queue for edge $(k, k + 1)$ must be nonempty and contain a packet destined to cross edge $(i - 1, i)$ or else the index $k - 1$ would yield a larger value for B. If $k = i$, the queue for edge $(i, i + 1)$ must be nonempty since $\phi_1(i, t + \ell) > n - 1$. By the FTG protocol some such packet must traverse the edge $(k, k + 1)$ during this routing step and hence the quantity B must decrease. $\square$

THEOREM 8. *Let G denote the n-node cycle and Υ the class of rate 1 deterministic adversaries (say with window size w). Define $Q_0(G)$ to be the total number of packets initially in the network (i.e., the sum of all edge queue sizes) and let $Q_0 = \max\{Q_0(G), w\}$. Then (G, Υ, FTG) is stable and furthermore there are never more than $n(Q_0 - 1) + w$ packets in the system.*

PROOF. For all k, $A(i, k, 0) \leq Q_0$, and so $\phi_1(i, 0) = \max\{A(i, k, 0) + (i + n - 1 - k): i \leq k \leq i + n - 1\}$ is at most $Q_0 + i + n - 1 - i = Q_0 + n - 1$, since the $(i + n - 1 - k)$ term is maximized with $k = i$. Also, $w + n - 1 \leq Q_0 + n - 1$, by definition of Q_0. Thus, $\phi(i, 0) = \max(\phi_1(i, 0), w + n - 1) \leq Q_0 + n - 1$.

Consider any time $t = mw$ for m a non-negative integer. By Lemma 7, we have $\phi(i, mw) \leq \phi(i, 0) \leq Q_0 + n - 1$. Hence, $A(i, i - 1, mw) + (i + n - 1 - (i - 1)) = A(i, i - 1, mw) + n \leq \phi(i, mw) \leq Q_0 + n - 1$. Subtracting n from both sides, we have $A(i, i - 1, mw) \leq Q_0 - 1$.

The total number of packets in the system at time $t = mw$ can then be upper bounded by $\Sigma_{i=0}^{n-1} A(i, i - 1, mw)$, which will be no greater than $n(Q_0 - 1)$. Finally, for any time t with $mw < t \leq (m + 1)w$, we have $\Sigma_{i=0}^{n-1} A(i, i - 1, t) \leq w + \Sigma_{i=0}^{n-1} A(i, i - 1, mw)$. $\square$

We now extend the above theorem to the case of a stochastic adversary with injection rate $1 - \epsilon$.

THEOREM 9. *Let G denote the n-node cycle and* Υ *the class of properly bounded stochastic adversaries with rate* $1 - \epsilon$ *for some* $\epsilon > 0$. *Then* (G, Υ, FTG) *is stable*.

The proof here is analogous to the proof for deterministic adversaries. We need an analogue of Lemma 7 so as to apply Lemma 2.

LEMMA 10. *For all i, and for all* $t \geq 0$, $E[\phi(i, t + w)|\phi(i, t) \geq n + 2w] - \phi(i, t) \leq -w\epsilon$.

PROOF. We argue as in the deterministic case, that as packets are inserted during steps $t + 1, \ldots, t + w$, the potential increases by at most the number of inserted packets during these time steps which wish to cross the edge $i - 1$ to i. By the definition of a $(w, 1 - \epsilon)$ stochastic adversary, this implies that the expected increase in ϕ due to insertions during time steps $t + 1, \ldots, t + w$ is bounded by $w(1 - \epsilon)$. It remains to observe (as in the deterministic case) that subject to the condition that $\phi(i, t) \geq n + 2w$, the decrease in ϕ during these steps due to packet routing is exactly w.

We combine these facts (about the expected increase due to insertions and the decrease due to routing steps) to obtain $E[\phi(i, t + w)|H_t] \leq \phi(i, t) + w(1 - \epsilon) - w$ provided $\phi(i, t) \geq n + 2w$. □

Given Lemma 10 we can now apply Lemma 2 so as to complete the proof of the Theorem.

7. NTG

Two natural scheduling rules for packet routing are the Nearest-To-Go (NTG) and Furthest-To-Go (FTG) policies. Andrews et al. [2001] prove that FTG is stable for any packet routing network at rate $\rho < 1$. Tsaparas [1997] generalizes the Andrews et al. [2001] stability result for FTG to more general queuing networks by proving the universal stability of Most-Time-To-Go (MTTG) scheduling. (When all service times are identical as we assume in packet routing networks, then MTTG becomes FTG.)

In contrast to the stability of FTG, Andrews et al. [2001] show that there is a simple 6-node network such that NTG scheduling is unstable for a rate $\rho = 0.62$ path-packing adversary.

We extend this result in two ways. First, we show for every $\rho > 0$ there is a queuing network for which NTG is unstable at rate ρ. Moreover, this instability will occur (for some initial configuration) even if packets are generated with constant rate ρ. Thus, this instability result is not an adversarial result (except for the setting of the initial configuration) but was motivated by the adversarial approach. (Some experimental evidence suggests that this instability would still occur for Poisson arrivals and an empty initial configuration.)

We shall now describe the network G and only sketch the intuitive reason for instability. (The details of the instability proof can be found in Tsaparas.) Let $\rho > 0$ be given and without loss of generality assume that $1/\rho$ is an integer. We first construct a "toroidal queuing network" $\tilde{G}$ (where the servers are the nodes rather than the edges) and describe the instability result in terms of $\tilde{G}$. It is then easy to convert $\tilde{G}$ to a packet routing network G for which the same instability phenomenon occurs.

FIG. 1.   Bad network for NTG scheduling.

Let $\bar{G}$ be an $n \times n$ torus with $n\rho > 2$ and n even. There will be n classes of jobs. A class i job initiates at node $(i, i + 1)$ and follows the following path of (node) servers $(i, i + 1), (i, i + 2), \ldots, (i, i - 1), (i + 1, i), (i + 2, i), \ldots,$ $(i - 1, i)$. (See Figure 1. All node addresses are computed mod n.)

For every i, a new class i job will be generated every $1/\rho$ steps. For every node, exactly two job classes will pass through any given node so that the total induced load per step on any server is exactly 2ρ.

For the NTG scheduling rule, jobs moving along a column have priority over jobs moving along a row. Thus, class i jobs when moving along column i will intersect and have priority over all other job classes. The intuitive idea to achieve instability is to initially establish two "walls" of packets, say in column i and column $j = i + n/2$. By a wall of packets on column i, we mean that every node (k, i) with $k \neq i$ is occupied and in addition there may be a large number of class i jobs queued in node $(i, i - 1)$,[12] which will continue to keep column i occupied. Such a wall will prevent other jobs from making progress. Eventually, the walls at columns i and j will drain out but at the same time a pair of walls will have been formed at columns $i + 1$ and $i + 1 + n/2$. Moreover, we claim that the total number of jobs in the system has increased during this "phase" where the walls are moving over one column. To see this, suppose that there are M class i jobs in the wall at the start of the phase (and the same for class $j = i + n/2$). Now assume that there is a total of T jobs in the system at the start of the phase. The phase lasts M steps during which time exactly $nM\rho$ new jobs have entered the system and $2M$ jobs have left the system. Thus, at the end of the phase (and the start of the next phase), there are $T' = T + nM\rho - 2M > T$ jobs in the system since $n\rho > 2$.

Tsaparas [1997] provides a careful proof of this intuitive idea. In particular, his proof requires an initial configuration which is very "symmetric" and then shows that this symmetry is preserved throughout the process so that new walls are formed without any breaks occurring in the walls.

---

[12] To be more precise, we consider this queue of jobs at node $(i, i - 1)$ to include other class i packets which will reach $(i, i - 1)$ during the present "phase" when a wall exists at column i.

Finally, we need to indicate how to convert $\tilde{G}$ to a packet routing network G. For every node, $\kappa = (i, j)$ in $\tilde{G}$, we have two nodes u_κ and v_κ in G. The edge set of G consists of a "server" edge $e_\kappa = (u_\kappa, v_\kappa)$ for every node κ in $\tilde{G}$ and a "connecting edge" $f_{(\kappa, \kappa')} = (v_\kappa, u_{\kappa'})$ for every pair of nodes κ, κ' in $\tilde{G}$ for which some job traverses from κ to κ' in the queuing network; $\kappa = (i, j)$ and $\kappa' = (i', j')$ and either $i' = i + 1$ and $j' = j$ or $i' = i$ and $j' = j + 1$. In the obvious way, each job class in the queuing network will determine a packet class (and its path) in the packet routing network. We claim that for $n > 3\rho$ that the packet routing network system $(G, NTG, \mathcal{A})$ is unstable where $\mathcal{A}$ is an "adversary" which is generating packets in each class at constant rate ρ. Again the proof of this claim can be found in Tsaparas [1997].

8. *Conclusion and Open Problems*

We have introduced a new approach for the study of queuing networks. Although our motivation comes from continuous packet routing, we note that the adversarial approach can be applied more generally to the wider field of queuing theory (as is done in Tsaparas [1997]). Clearly more general queuing networks (in which jobs can revisit the same server many times and/or where service time distributions are class dependent) offer additional challenges. It remains an interesting open question to find some set of general conditions (depending, for example, on the rate, queuing discipline, and underlying network) that are sufficient to guarantee stability. Further work is also needed to better bound queue sizes and packet delays, and to understand specific networks such as arrays and hypercubes. We mention a few of the many open problems:

—For packet routing, can FIFO be made unstable for arbitrarily small positive rates of injection in the adversarial model? More generally, does stability at some rate $\rho_1 > 0$ imply stability for all $\rho < 1$ (for a "natural" scheduling policy)?

—Considering the ring as a queuing network allowing arbitrary (i.e., reversing) paths, is every scheduling policy stable?

—Is NTG unstable for a Poisson input model for packet routing and for arbitrarily small rates? In particular, does the instability result of Section 7 extend from constant rate arrivals to Poisson arrivals? More generally, when does adversarial instability from some given initial configuration imply instability with Poisson arrivals (and say an empty initial configuration). Note that while FIFO is unstable in an adversarial setting (a special case of class independent service times), it is stable for Poisson arrivals and class independent service times.

—Is there a generic transformation of stability results from deterministic to stochastic adversaries?

—Our queue size bound for DAGs is exponential in the length d of the longest path in the network. Indeed Andrews et al. [2001] show that NIS and FTG can be made to suffer such exponential size queues on DAGs. What is the best bound on queue size in DAGs that holds for the FIFO scheduling rule? Recently Adler and Rosén (personal communication) show that LIS has polynomial size queues (as a function of d) on any DAG. Andrews and Zhang [2000] show that using LIS, a packet ρ can have exponential delay as a function

of l, the path length of the packet. In their construction, l is a logarithm in the network depth d. Andrews et al. [1996] construct a randomized ("LIS based") scheduling rule which has (with high probability) polynomially-sized queues and note that this randomized policy can be converted to a centralized deterministic scheduling policy with polynomially sized queues. The obvious question is whether there exists a (natural) decentralized, deterministic scheduling rule that has polynomially-sized (as a function of either d or the size of the network) queues for all networks.

—What can be said about stability of nonadaptive routing in a network where each edge e is capable of simultaneously transmitting some number n_e packets in one step? Specifically, if $(G, \mathcal{A}, \mathcal{S})$ is universally stable for any (w, ρ) adversary $\mathcal{A}$, does it follow that $(G', \mathcal{A}', \mathcal{S})$ is also universally stable for any (w, ρ) adversary $\mathcal{A}'$ where G is a graph where all edges have bandwidth 1 and G' is the same graph with arbitrary edge bandwidths $\{n_e\}$. Here the load condition is modified in the obvious way so that in any window τ of w time steps and for any edge e, the adversary injects at most $N(\tau, e) \leq \rho \cdot w \cdot n_e$ packets which cross edge e. (This can also be viewed as a very special case of adaptive routing by viewing a "high bandwidth" edge as a set of edges.) More generally, what "network transformations" preserve stability? Recent results along these lines have been derived by Borodin et al. [2001].

—Which (natural or efficient) scheduling rules are universally stable for deterministic rate 1 adversaries? Subsequent to asking this question, Gamarnik [1999] has proven that the "nearest to origin" NTO scheduling rule (which gives priority to packets that have traversed the fewest edges thus far) is universally stable for deterministic rate 1 adversaries. It is easy to adapt Gamarnik's proof to show that FTG is also universally stable for deterministic rate 1 adversaries. These proofs yield an exponential bound on the total number of packets in the network. Gamarnik also observes that using either NTO or FTG, a rate 1 adversary can prevent some packets from ever reaching their destination. He asks if there is a scheduling rule (perhaps a modification of NTO or FTG) that can guarantee bounded or at least finite delivery time with respect to deterministic rate 1 adversaries.

REFERENCES

AIELLO, W., KUSHILEVITZ, E., OSTROVSKY, R., AND ROSÉN, A. 1988. Adaptive packet routing for bursty adversarial traffic. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC)* (Dallas, Tex., May 23–26). ACM, New York, pp. 359–368.

ANDREWS, M. 2000. Instability of FIFO in session-oriented networks. In *Proceedings of the 11th Annual ACM–SIAM Symposium on Discrete Algorithms* (San Francisco, Calif., Jan. 9–11). ACM, New York, pp. 440–447.

ANDREWS, M., AWERBUCH, B., FERNÁNDEZ, A., LEIGHTON, F. T., LIU, Z., AND KLEINBERG, J. 2001. Universal stability results for greedy contention-resolution protocols. *J. ACM 48*, 1 (Jan.), 39–69.

ANDREWS, M., AND ZHANG, L. 2000. The effects of temporary sessions on network performance. In *Proceedings of the 11th Annual ACM–SIAM Symposium on Discrete Algorithms* (San Francisco, Calif., Jan. 9–11). ACM, New York, pp. 448–457.

AWERBUCH, B., AND LEIGHTON, F. T. 1994. Improved approximations for the multicommodity flow problem and local competitive routing in networks. In *Proceedings of the 26th Annual ACM*

Symposium on Theory of Computing (Montreal, Que., Canada, May 23–25). ACM, New York, pp. 487–496.

BERTSIMAS, D., GAMARNIK, D., AND TSITSIKLIS, J. N. 1996. Stability conditions for multiclass fluid queuing networks. *IEEE Trans. Automat. Cont. 41* (Nov.), 1618–1631.

BORODIN, A., KLEINBERG, J., RAGHAVAN, P., SUDAN, M., AND WILLIAMSON, D. 1996. Adversarial queuing theory. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (Philadelphia, Pa., May 22–24). ACM, New York, pp. 376–385.

BORODIN, A., OSTROVSKY, R., AND RABANI, Y. 2001. Routing networks with speeds and capacities. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Theory of Computing*. ACM, New York.

BRAMSON, M. 1994a. Instability of FIFO queuing networks. *Ann. Appl. Prob. 4*, 414–431.

BRAMSON, M. 1994b. Instability of FIFO queuing networks with quick service times. *Ann. Appl. Prob. 4*, 693–718.

BRAMSON, M. 1996. Convergence to equilibria for fluid models of FIFO queuing networks. *Que. Syst. 22*, 5–45.

BRODER, A., FRIEZE, A., AND UPFAL, E. 1997. A general approach to dynamic packet routing with bounded buffers. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 390–399.

CRUZ, R. L. 1991a. A calculus for network delay, Part I. *IEEE Trans. Inf. Theory 37* (Jan.), 114–131.

CRUZ, R. L. 1991b. A calculus for network delay, Part II. *IEEE Trans. Inf. Theory 37* (Jan.), 132–141.

DAI, J. G. 1995. On positive Harris recurrence of multiclass queuing networks: A unified approach via fluid limit models. *Ann. Appl. Prob. 5*, 49–77.

DAI, J. G., AND MEYN, S. P. 1995. Stability and convergence of moments for multiclass queuing networks via fluid limit models. *IEEE Trans. Automat. Cont. 40* (Nov.), 1889–1904.

DAI, J. G., AND WEISS, G. 1996. Stability and instability of fluid models for reentrant lines. *Math. Oper. Res. 21* (Feb.).

DOWN, D. D., AND MEYN, S. P. 1995. Stability of acyclic multiclass queuing networks. *IEEE Trans. Automat. Cont. 40*, 5, 916–920.

DURRETT, R. 1995. *Probability: Theory and Examples*, 2nd ed. Duxbury Press, Belmont, Calif.

GAMARNIK, D. 1998. Stability of adversarial queues via fluid models. In *Proceedings of the 39th Symposium on the Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 60–70.

GAMARNIK, D. 1999. Stability of adaptive and non-adaptive packet routing policies in adversarial queuing networks. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing* (Atlanta, Ga., May 1–4). ACM, New York, pp. 206–214.

GOEL, A. 1997. Stability of networks and protocols in the adversarial queuing model for packet routing. Stanford Univ. Tech. Note STAN-CS-TN-97-59. Stanford Univ., Stanford, Calif.

HAJEK, B. 1982. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Prob. 14*, 502–525.

HARCHOL-BALTER, M., AND BLACK, P. E. 1994. Queuing analysis of oblivious packet-routing networks. In *Proceedings of the 5th Annual ACM–SIAM Symposium on Discrete Algorithms* (Arlington, Va., Jan. 23–25). ACM, New York, pp. 583–592.

HARCHOL-BALTER, M., AND WOLFE, D. 1995. Bounding delays in packet-routing networks. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing* (Las Vegas, Nev., May 29–June 1). ACM, New York, pp. 248–257.

JACKSON, J. R. 1963. Jobshop-like queuing systems. *Manage. Sci. 10*, 131–142.

KAHALE, N., AND LEIGHTON, F. T. 1995. Greedy dynamic routing on arrays. In *Proceedings of the 6th Annual ACM–SIAM Symposium on Discrete Algorithms* (San Francisco, Calif., Jan. 22–24). ACM, New York, pp. 558–566.

KELLY, F. P. 1979. *Reversibility and stochastic networks*. Wiley, New York.

KLEINROCK, L. 1975. *Queuing Systems*. Wiley, New York.

LEIGHTON, F. T. 1990. Average case analysis of greedy routing algorithms on arrays. In *Proceedings of the 2nd Annual ACM Symposium on Parallel Algorithms and Architectures* (Island of Crete, Greece, July 2–6). ACM, New York, pp. 2–10.

LEIGHTON, F. T., MAGGS, B. M., AND RAO, S. B. 1994. Packet routing and job-shop scheduling in O(congestion + dilation) steps. *Combinatorica 14*, 167–180.

LEIGHTON, F. T., MAGGS, B. M., AND RICHA, A. W. 1999. Fast algorithms for finding O(congestion + dilation) packet routing schedules. *Combinatorica 19*, 3, 375–401.

LOEVE, M. 1978. *Probability Theory*, vol. 2. Springer-Verlag, New York.

LU, S. H., AND KUMAR, P. R. 1991. Distributed scheduling based on due dates and buffer priorities. *IEEE Trans. Automat. Cont. 36*, 12, 1406–1416.

MANI CHANDY, K., HOWARD, J., AND TOWSLEY, D. 1977. Product form and local balance in queuing networks. *J. ACM* (Apr.). 250–263.

MANSOUR, Y., AND PATT-SHAMIR, B. 1991. Greedy packet scheduling on shortest paths. In *Proceedings of the 10th Annual ACM Symposium on Principles of Distributed Computing* (Montreal, Que., Canada, Aug. 19–21). ACM, New York, pp. 165–176.

MEYN, S. P., AND DOWN, D. 1994. Stability of generalized Jackson networks. *Ann. Appl. Prob. 4*, 124–148.

MITRA, D., AND CIESLAK, R. 1987. Randomized parallel communications on an extension of the omega network. *J. ACM 34*, 4 (Oct.), 802–824.

MITZENMACHER, M. 1994. Bounds on the greedy routing algorithm for array networks. In *Proceedings of the 6th Annual ACM Symposium on Parallel Algorithms and Architectures* (Cape May, N.J., June 27–29). ACM, New York, pp. 346–353.

OSTROVSKY, R., AND RABANI, Y. 1997. Universal O(congestion + dilation + $\log^{1+\epsilon} N$) local control packet switching algorithms. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (El Paso, Tex., May 4–6). ACM, New York, pp. 644–653.

PAREKH, A. K., AND GALLAGER, R. G. 1993. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Trans. Netw. 1*, 3, 344–357.

PAREKH, A. K., AND GALLAGER, R. G. 1994. A generalized processor sharing approach to flow control in integrated services networks: The multiple-node case. *IEEE/ACM Trans. Netw. 2*, 2, 137–150.

PEMANTLE, R., AND ROSENTHAL, J. S. 1998. Moment conditions for a sequence with negative drift to be uniformly bounded in L^r. Tech. Rep. 9814. Dept. *Statistics*, Univ. Toronto, Toronto, Ont., Canada.

PEMANTLE, R., AND ROSENTHAL, J. S. 1999. Moment conditions for a sequence with negative drift to be uniformly bounded in L^r. *Stoch. Proc. Their Appl. 82*, 143–155.

PERKINS, J. R., AND KUMAR, P. R. 1989. Stable distributed real-time scheduling of flexible manufacturing/assembly/disassembly systems. *IEEE Trans. Automat. Cont. 34*, 139–148.

RABANI, Y., AND TARDOS, E. 1996. Distributed switching in arbitrary networks. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (Philadelphia, Pa., May 22–24). ACM, New York, pp. 366–375.

RYBKO, A. N., AND STOLYAR, A. L. 1992. Ergodicity of stochastic processes describing the operation of open queuing networks. *Prob. Inf. Trans. 28*, 199–220.

SEIDMAN, T. I. 1994. 'First come, first serve' can be unstable. *IEEE Trans. Automat. Cont. 39*, 2166–2171.

STAMOULIS, G. D., AND TSITSIKLIS, J. N. 1991. The efficiency of greedy routing in hypercubes and butterflies. In *Proceedings of the 3rd Annual ACM Symposium on Parallel Algorithms and Architectures* (Hilton Head, S.C., July 21–24). ACM, New York, pp. 248–260.

TASSIULAS, L., AND GEORGIADIS, L. 1996. Any work-conserving policy stabilizes the ring with spatial re-use. *IEEE/ACM Trans. Netw. 4*, 2, 205–208.

TSAPARAS, P. 1997. Stability in adversarial queuing theory. M.Sc dissertation. Dept. of Computer Science, Univ. of Toronto, Toronto, Ont. Canada.

WALRAND, J. 1988. *An Introduction to Queuing Networks*. Prentice-Hall, Englewood Cliffs, N.J.