Adversarial Models for Priority-Based Networks*

C. Àlvarez[1], M. Blesa[1], J. Díaz[1], A. Fernández[2], and M. Serna[1]

[1] Dept. de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya
Jordi Girona 1–3, 08034 Barcelona
{alvarez,mjblesa,diaz,mjserna}@lsi.upc.es
[2] Dept. Ciencias Experimentales e Ingeniería, Universidad Rey Juan Carlos
Tulipán s/n, Campus de Móstoles, 28933 Madrid
anto@gsyc.escet.urjc.es

Abstract. We propose several variations of the adversarial queueing model to cope with packets that can have different priorities, the *priority* and *variable priority* models, and link failures, the *failure* and *reliable* models. We address stability issues in the proposed adversarial models. We show that the set of *universally stable* networks in the adversarial model remains the same in the four introduced models. From the point of view of queueing policies we show that several queueing policies that are universally stable in the adversarial model remain so in the priority, failure and reliable models. However, we show that LIS, a universally stable queueing policy in the adversarial model, is not universally stable in any of the other models, and that no greedy queueing policy is universally stable in the variable priority model. Finally we analyze the problem of deciding stability of a given network under a fixed protocol. We provide a characterization of the networks that are stable under FIFO and LIS in the failure model. This characterization allows us to show that deciding network stability under FIFO and LIS in the proposed models can be solved in polynomial time.

1 Introduction

The model of Adversarial Queueing Theory (AQT) proposed by Borodin et al. [10] considers the time evolution of a packet-routing network as a game between an adversary and a queueing policy. At each time step the adversary may inject a set of packets to some of the nodes. For each packet the adversary specifies the sequence of edges that it must traverse, after which the packet will be absorbed. If more than one packet try to cross an edge e at the same time step, then the queueing policy chooses one of these packets to be sent across e. The remaining

* Work partially supported by the FET Programme of the EU under contract number IST-2001-33116 (FLAGS), and by the Spanish CICYT project TIC-2001-4917-E. The second author was also supported by the Catalan government with the predoctoral grant 2001FI-00659.

B. Rovan and P. Vojtáš (Eds.): MFCS 2003, LNCS 2747, pp. 142–151, 2003.

packets wait in the queue. This game then advances to the next time step. The main goal of the model is to study *stability issues* of the network, under different *greedy* queueing policies. Stability is the property of deciding whether at any time the maximum number of packets present in the system is bounded by a constant that may depend on system parameters. Recall that a protocol is *greedy* if whenever there is at least one packet waiting to use an edge, the protocol advances a packet through the edge.

In the *adversarial* model the adversary is restricted by a pair (r, b) where $b \geq 0$ is *the burstiness*, and $0 < r < 1$ is the *injection rate*. The adversary must obey the following rule $N_e(I) \leq r|I| + b$, where $N_e(I)$ denotes the number of packets injected by the adversary during a time interval I that have paths containing the edge e[1] [10,5].

In this paper we consider a generalization of the adversarial model which takes into account the possibility that packets may have different priorities and we explore some dynamic network settings.

Priority Models. Considering priorities is a natural approach to model nowadays' networks. Today's networked applications, such as data mining, e-commerce and multimedia, are bandwidth hungry and time sensitive. These applications need networks that accommodate to these requirements and guarantee some *Quality of Service* (QoS). Classifying and prioritizing network traffic flow are basic tasks.

We are interested in analyzing the power of an adversary that can prioritize the packets. We will consider two settings, in the first one each packet have a fixed priority and in the second one the adversary is allowed to modify the priority of a packet at any time step. Consequently, we define two new models for adversarial queueing theory, the *priority* and the *variable priority*. When packets have priorities, each edge has a queue associated to every possible priority. If at a certain time, more than one packet tries to cross the same edge e, the queueing policy chooses one of these packets to send across e, from the non-empty queue with highest priority. The limitations on the adversary are the same as in the adversarial model.

Models for Dynamic Networks. Inspired by the priority models and by the growing importance of wireless mobile networks where some connections between nodes may fail or change quickly and unpredictably, we also consider some variations of the adversarial model for *dynamic networks*, in which edges can appear and disappear arbitrarily. Note that in the priority model, we can *simulate* the failure of an edge e by injecting a packet of length 1 in e with a priority higher than any other packet in the queue of e. Once this packet has been injected none of the remaining packets of the queue can be served until the next time step. It seems natural to introduce models for dynamic networks in which the adversary controls not only the packet arrivals, but also the edge failures. The constraints of an adversary are defined obeying the rule: the number of packets introduced by the adversary during interval I, which have paths containing e, can not be

[1] Recall that in [10] the model is defined over windows of fixed size w and the equation $N_e(I) \leq r|I|$, for $|I| = w$. It is known that both models are equivalent [16].

greater than a fraction of the number of times that e is alive in this interval. Furthermore, as packets must follow a pre-specified path, the adversary should not be able to make an edge fail perpetually. We assume that a packet cannot cross a failed link, and that during an edge failure the packets that arrive wait queued at the head of the edge.

Let $F_e(I)$ be the number of steps during time interval I for which the edge e is down. In order to guarantee that we keep a bound on the maximum number of failures of an edge e in any time interval, we propose the *failure* model, in which the adversary is controlled by a common bound on both the injection and the edge failures, according to the restriction

$$N_e(I) + F_e(I) \leq r|I| + b. \tag{1}$$

Observe than in this case, during a given interval, the injection rate limits the maximum number of failures and the maximum number of packet injections per edge.

With the aim of allowing a higher amount of edge failures, we define a new dynamic model. To do so, we introduce an additional parameter α. The adversary is characterized by (r, b, α) , where r, b are defined as before and $r \leq \alpha \leq 1$. For any edge e and any interval I, the adversary must obey the constraint

$$N_e(I) + \alpha F_e(I) \leq r|I| + b. \tag{2}$$

In the model defined by equation (2), we can consider two extreme cases: For $\alpha = 1$ we obtain the constraint (1) defining the failure model. For $\alpha \leq r$ we get a model in which an edge can be down all the time. Notice that in the case $r < \alpha \leq 1$, if the adversary produces a failure of an edge, then it is forced to recover the edge after $b/(\alpha - r)$ steps, otherwise it will violate inequality (2). We are interested in the latter property and will use the term *reliable model* to denote an adversary with parameters (r, b, α) where $0 < r < 1$, $b > 1$, and $r < \alpha \leq 1$ and constrained by inequality (2).

Greedy Protocols. Through this paper the term *network* will refer to digraphs which may have multiple edges but no loops. As in [10,5] we will only consider greedy protocols which apply their policies to the queues at the edges according to some local or global criteria. The systems acts synchronously. The main queueing protocols we consider are: FIFO (First In First Out), LIFO (Last In First Out), SIS (Shortest In System), LIS (Longest In System), NTG (Nearest To Go), FTG (Furthest To Go), NFS (Nearest From Source), FFS (Farthest From Source) and NTG-LIS, which works as NTG, but resolving ties using the LIS protocol.

It is known that FTG, NTS, SIS and LIS are universally stable in the adversarial model while FIFO, LIFO, NTG and FFS are not [5].

Related Work. Two adversarial models for dynamic networks have been proposed in [7] and [6]. In both models the injected packets are defined by specifying only source and destination, and thus they are not forced to follow a pre-specified path. The dynamic models proposed in this paper consider the case in which the injected packets are defined specifying the sequence of edges that they must traverse. Our models and the dynamic models proposed in [7] and in [6] have the common characteristic that, for every interval I, the adversary can not inject to any edge e (or to any set S of nodes for the model in [6]), more packets than the number of packets that e can absorb (or the edges with only one extreme in S).

An interpretation of system failure as a *slowdown* in the transmission or *link capacity*, instead of link failure was studied in [9], in both models packets are injected with a pre-specified path.

In the *dynamic slowdown model*, a packet p suffers slowdown $s_e(t)$ while crossing edge e at time t, that means that it p starts to traverse the link at time t and arrives to the tail of e at time $t + s_e(t)$. During this transfer time the packets that want to cross e wait in the queue of e. In the *static* case every link e has a fixed slowdown s_e. This situation has some similarities with the failure model as the slowdown s incurred by a packet traversing a link e can be interpreted as "link e *fails* during $s - 1$ steps". However, there is a difference, in the slowdown model p is delayed after leaving e's queue, while in the failure model p waits in e's queue. This means that when e is recovered, the next packet to be served might not be p.

In the *capacity* model every edge e in a network has capacity $c_e(t)$ at time step t. They also consider a static case in which the capacity does not depend on time. In step t a link is able to transmit simultaneously up to $c_e(t)$ packets.

The main results in the paper are that every universally stable network remains universally stable in the slowdown and the capacity models, even in the dynamic case. That SIS, NTS and FTG remain universally stable in all the models. The situation is different for LIS since it is universally stable in the static slowdown model but it is not in the dynamic slowdown and capacity models. Even though we can interpret that a link fails at any time step with $c_e(t) = 0$ the proof that LIS is not universally in the dynamic capacity model uses two non zero capacities (see Theorem 3.1 [9]).

Our Contributions. We address stability issues in the proposed adversarial models. Let us recall that a network is *stable under a protocol and an adversary* if the number of packets in the system at any time step remains bounded. Our first results concern with *universal stability* of networks. First we show that the property that a network is stable under any adversary and queueing policy remains the same in the adversarial, priority, variable priority, failure and reliable models.

From the point of view of queueing policies, we show that NFS, SIS and FTG, that are universally stable in the adversarial model [5], remain so in the failure, reliable, and priority models. However, we show that LIS, a universally stable queueing policy in the adversarial model [5], is not universally stable in the failure, reliable and priority models. Moreover, we show that no greedy protocol is universally stable in the variable priority model.

Finally we analyze the problem of deciding stability of a given network under a fixed protocol. We provide a characterization of the networks that are stable under FIFO and LIS in the failure model. This characterization is the same as the one given in [3] in the adversarial model for universal stability and for stability under NTG-LIS. Thus, our results show that for FIFO and LIS the stability problem in the failure model can be solved in polynomial time. Let us observe that the characterization of FIFO stability in the adversarial model remains an open problem [4].

Due to the lack of space we will omit all the proofs they can be found in the full version [1].

2 Universal Stability of Networks

A communication system is formed by three main components: A network G, a scheduling protocol $\mathcal{P}$ and a traffic pattern $\mathcal{A}$ which is represented by an adversary. The concept of *universal stability* applies either to networks or protocols.

Let $\mathcal{M}$ denote a model in the set {*adversarial*, *reliable*, *failure*, *priority*, *variable priority* } as defined in the previous section. Given a network G, a queuing policy $\mathcal{P}$ and a model $\mathcal{M}$, we say that for a given adversary $\mathcal{A}$ following the restrictions of $\mathcal{M}$, the system $S = (G, \mathcal{A}, \mathcal{P})$ is *stable in the model* $\mathcal{M}$ if, at any time step, the maximum number of packets in the system is bounded by a fixed value that may depend on system parameters. The pair $(G, \mathcal{P})$ is *stable in the model* $\mathcal{M}$ if, for any adversary $\mathcal{A}$ following the restrictions of $\mathcal{M}$, the system $S = (G, \mathcal{A}, \mathcal{P})$ is stable in $\mathcal{M}$. A network G is *universally stable in the model* $\mathcal{M}$ if, for any greedy queuing policy $\mathcal{P}$, the pair $(G, \mathcal{P})$ is stable in $\mathcal{M}$. A greedy protocol $\mathcal{P}$ is *universally stable in the model* $\mathcal{M}$ if, for any digraph G, the pair $(G, \mathcal{P})$ is stable in $\mathcal{M}$.

In order to compare the power of adversaries in different adversarial models, we introduce the concept of simulation. We say that an adversary $\mathcal{A}$ in model $\mathcal{M}$ *simulates* an adversary $\mathcal{A}'$ in model $\mathcal{M}'$ when for any network G and any protocol $\mathcal{P}$ if $(G, \mathcal{A}, \mathcal{P})$ is stable in $\mathcal{M}$, then $(G, \mathcal{A}', \mathcal{P})$ is stable in $\mathcal{M}'$.

Lemma 1. *Any (r, b)-adversary in the adversarial model can be simulated by an (r, b)-adversary in the failure model. Any (r, b)-adversary in the failure model is an $(r, b, 1)$-adversary in the reliable model. Any (r, b)-adversary in the failure model can be simulated by an (r, b)-adversary in the priority model using two priorities. Any (r, b, α)-adversary in the reliable model can be simulated by an $(r + 1 - \alpha, b)$-adversary in the failure model.*

Observe that the failure and reliable models are equivalent. So any stability or instability result for one model applies to the other as well. We will consider in the following only the failure and priority models.

Now we can state our main result in this section.

Theorem 1. *Given a digraph G, the following properties are equivalent*

1. *G is universally stable in the adversarial model,*
2. *G is universally stable in the failure model,*
3. *G is universally stable in the priority model, and*
4. *G is universally stable in the variable priority model.*

3 Universal Stability of Protocols

In this section we address the universal stability property in the failure and priority models, from the point of view of the queuing policy. We will consider the six basic protocols presented in the introduction. Recall that FTG, NTS, SIS and LIS are universally stable in the adversarial model while FIFO, LIFO, NTG, and FFS are not [5]. Since any adversary in the adversarial model can be seen as an adversary in the other models, FIFO, LIFO, NTG, and FFS are not universally stable in the failure and priority models.

First we show how the behavior of any (r, b)-adversary for network G in the priority model can be simulated by an (r', b')-adversary for a network G' in the adversarial model, under the same protocol $\mathcal{P}$ in the case that $\mathcal{P} \in \{\text{FTG}, \text{NTG}, \text{NFS}, \text{FFS}\}$.

Let G be a directed graph and A_π an adversary in the priority model, assigning at most π priorities and with injection rate (r, b), where $b \geq 0$ and $0 < r < 1$. Every injected packet p has a priority π_p in the ordered interval $[1, \ldots, \pi]$, being one the lowest priority.

Lemma 2. *For any system $S = (G, \mathcal{A}_\pi, \mathcal{P})$ in the priority model, for $\mathcal{P} \in \{\text{FTG}, \text{NTG}\}$, there is a system $S' = (G', \mathcal{A}', \mathcal{P})$ in the adversarial model such that, G is a subgraph of G', $\mathcal{A}'$ has injection rate (r, b), if a packet p is injected in S at time t with route r, a packet p' is injected in S' at time t with a route r' obtained by concatenating r and a path of edges not in G, and if p crosses edge e at time t' in S, p' crosses e at time t' in S'.*

To get a similar result for NFS and FFS we have to relate two different networks.

Lemma 3. *For any system $S = (G, \mathcal{A}_\pi, \mathcal{P})$ in the priority model, for $\mathcal{P} \in \{\text{NFS}, \text{FFS}\}$, there is a system $S' = (G', \mathcal{A}', \mathcal{P})$ in the adversarial model such that, G is a subgraph of G', $\mathcal{A}'$ has injection rate (r, b'), where $b' = r(\pi - 1)d + b$, if a packet p is injected in S with route r, a packet p' is injected in S' with a route r' obtained by concatenating and a path of edges not in G and r, and if p crosses edge e at time t in S, p' crosses e at time t in S'.*

As a consequence of the previous lemmas, and the results in [5] we get,

Theorem 2. FTG *and* NFS *are universally stable in the priority, failure and reliable models.*

We can show the universal stability of SIS in the priority model by following similar arguments to that of Lemma 2.2 in [5] for showing the universal stability of SIS in the adversarial model and induction on the number of priorities.

Theorem 3. SIS *is universally stable in the priority, failure and reliable models.*

The next result states the non universal stability of LIS in the failure model, and therefore in all the other models. We show that the graph U_1 of Figure 1 is not stable under LIS in the failure model. Putting this result together with Lemma 1 we get,

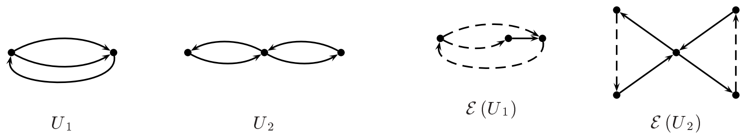

Fig. 1. The two subgraphs characterizing universal stability in the adversarial model (see [3]), and their extensions replacing an edge by a path

Theorem 4. LIS *is not universally stable in the failure, reliable and priority models.*

Finally, we consider universal stability of protocols in the variable priority model. We show that the graph U_1 of Figure 1 is not stable under any greedy protocol in the variable priority model.

Theorem 5. *There is no greedy protocol universally stable in the variable priority model*

4 Stability under a Protocol

In this section we analyze the complexity of the problem of deciding whether a given network G is stable under a fixed protocol. Few results are known for this problem. Before formally stating our results, we need to introduce some graph theoretical definitions. We will consider the following subdivision operations over digraphs:

- The *subdivision of an arc* (u, v) in a digraph G consists in the addition of a new vertex w and the replacement of (u, v) by the two arcs (u, w) and (w, v).
- The *subdivision of a 2-cycle* (u, v), (v, u) in a digraph G consists in the addition of a new vertex w and the replacement of (u, v), (v, u) by the arcs (u, w), (w, u), (v, w) and (w, v).

Given a digraph G, $\mathcal{E}(G)$ denotes the family of digraphs formed by G and all the digraphs obtained from G by successive arc or 2-cycle subdivisions. Given a family of digraphs $\mathcal{F}$, $\mathcal{S}(\mathcal{F})$ denotes the family of digraphs that contain a graph in $\mathcal{F}$ as a subgraph.

Figure 1 provides the two basic graphs needed to characterize universal stability, and the shape of the extensions of those graphs. This basic family provides the characterization of the stability properties. It is known that a digraph is universally stable in the adversarial model if and only if $G \notin \mathcal{S}(\mathcal{E}(U_1) \cup \mathcal{E}(U_2))$ [3] . The same property characterizes network stability under NTG-LIS [3] and FFS [2]. It is also known that, for a given digraph G, checking whether $G \notin \mathcal{S}(\mathcal{E}(U_1) \cup \mathcal{E}(U_2))$ can be done in polynomial time [3]. Further results for undirected graphs and other variations can be found in [5] and [4]. Nothing is known about the complexity of deciding stability in the adversarial model for any other queueing policy. In the following we will provide a similar characterization of stability in the failure model under FIFO and LIS.

4.1 FIFO Stability under the Failure Model

Much effort have been devoted to study stability and instability properties of FIFO recently. The FIFO protocol was shown not to be universally stable [5]. A network-dependent absolute constant is provided in [12] such that FIFO is stable against any adversary with smaller injection rate. A lower bound of 0.749 for the instability is calculated in [13]. This bound was decreased to 0.5 [15]. In [11] it is shown that FIFO is stable if the injection rate is smaller than $1/(d-1)$. Recently, it has been proved that FIFO can become unstable at arbitrarily low injections rates [8,14].

We show that the two basic graphs given in Figure 1, as well as their extensions, are not stable under FIFO in the failure model. As a more general result, for network U_2 we can show instability in the adversarial model.

Lemma 4. *Any graph in $\mathcal{S}(\mathcal{E}(U_1) \cup \mathcal{E}(U_2))$ is not stable under* FIFO *in the failure model.*

As we have pointed out before, all networks $G \notin \mathcal{S}(\mathcal{E}(U_1) \cup \mathcal{E}(U_2))$ are universally stable in the adversarial model. Taking into account that if a network has an unstable subnetwork it is also unstable we get the following result.

Theorem 6. *Let G be a digraph, the pair (G, FIFO) is stable in the failure model if and only if G is universally stable in the adversarial model.*

A corollary of this result is the equivalence between FIFO stability in the failure model and universal stability in the adversarial model. Furthermore, as instability in the failure model implies instability in the priority and reliable models, the characterization of FIFO stability remains the same in the priority and reliable models. Observe also that stability under FIFO can be checked in polynomial time for the failure, priority and reliable models.

4.2 LIS Stability under the Failure Model

The LIS protocol gives priority to the packet that was longer in the system, i.e., that joined the network earlier. In [5], the LIS protocol was shown to be universally stable in the adversarial model, with $O(b/(1-r)^d)$ queue size per edge and delay of the packets in the order $O(b/(1-r)^d)$. However, as we have shown the protocol is not universally stable in the failure model.

We proceed as in the case of FIFO by showing, respectively, the instability of the basic graphs given in Figure 1 and their extensions.

Lemma 5. *Any graph in $\mathcal{S}(\mathcal{E}(U_1) \cup \mathcal{E}(U_2))$ is not stable under* LIS *in the failure model.*

Therefore, as in the case of FIFO, we have

Theorem 7. *A digraph G is stable under* LIS *in the failure model if and only if G is universally stable in the adversarial model.*

4.3 The Variable Priority Model

For the variable priority model the situation is simpler. It is easy to adapt the LIS instability proofs for U_1, U_2 and their extensions. This together with Theorem 1 gives the following result.

Theorem 8. *Let $\mathcal{P}$ be any greedy protocol. A digraph G is stable under $\mathcal{P}$ in the failure model if and only if G is universally stable in the adversarial model.*

5 Conclusions and Open Problems

We have proposed several variations on the adversarial model to cope with packet priorities and link failures. We have studied universal stability from the point of view of both, the network and the queueing policy. We have also addressed the complexity of deciding stability under a fixed protocol.

We have shown that in the adversarial, failure, reliable, priority and variable priority models, the set of networks that are universally stable remains the same. The models present a different behavior with respect to the universal stability of protocols, since LIS is universally stable in the adversarial model, but it is not universally stable in the other models. In contrast, we have shown that there are no universally stable protocols for the variable priority model.

We have proposed a new and natural way to model the behavior of queueing systems in dynamic networks. Our results compared to the slowdown models introduced in [9] show that the power of an adversary in the failure and in the dynamic slowdown model is quite similar. In both cases the LIS protocol is not universally stable. However the static slowdown model is less powerful than the failure model as LIS remains universally stable [9].

The argument used in the proof of Theorem 4.1 in [9] can be used to show how to construct an adversary in the variable priority model that simulates an adversary in the dynamic slowdown model. It would be of interest to find constructions, similar to those given in Lemmas 2 and 3 to relate the power of the slowdown and failure models without changing the protocol.

Regarding the dynamic capacity model, the authors frequently use the trick of injecting $c - c_e(t)$ dummy packets which only need to traverse link e. This can be done without violating the load condition for a network with static capacity c provided that $c_e(t) > 0$ (see Theorems 3.3 and 3.4 [9]). It will be of interest to analyze the case with zero capacities.

It remains as an open problem to show the existence of a protocol that is universally stable in the failure model but it is not in the priority model.

All the already known characterizations of stability under a protocol are equivalent to universal stability in the adversarial model, even in the variable priority model. It is an interesting open question to know whether there is any protocol $\mathcal{P}$, not universally stable, for which there are networks that are not universally stable but that are stable under $\mathcal{P}$. Finally let us point that deciding stability under FIFO in the adversarial model is open, and that also nothing is known about characterizations of stability under LIFO.

References

1. C. Àlvarez, M. Blesa, J. Díaz, A. Fernández and M. Serna. Adversarial models for priority-based networks. Technical Report LSI-03-25-R, Software department, Universitat Politècnica de Catalunya, 2003.
2. C. Àlvarez, M. Blesa, J. Díaz, A. Fernández and M. Serna. The complexity of deciding stability under FFS in the adversarial model. Technical Report LSI-03-16-R, Software department, Universitat Politècnica de Catalunya, 2003.
3. C. Àlvarez, M. Blesa, and M. Serna. A characterization of universal stability in the adversarial queueing model. Technical Report LSI-03-27-R, Software department, Universitat Politècnica de Catalunya, 2003.
4. C. Àlvarez, M. Blesa, and M. Serna. Universal stability of undirected graphs in the adversarial queueing model. In *14th ACM Symposium on Parallel Algorithms and Architectures* (SPAA'02), pages 183–197, Winnipeg, Canada, August 2002. ACM Press New York, USA.
5. M. Andrews, B. Awerbuch, A. Fernández, J. Kleinberg, T. Leighton, and Z. Liu. Universal stability results for greedy contention–resolution protocols. *Journal of the ACM*, 48(1):39–69, 2001.
6. E. Anshelevich, D. Kempe, and J. Kleinberg. Stability of load balancing algorithms in dynamic adversarial systems. In *34th. ACM Symposium on Theory of Computing* (STOC'02), pages 399–406, 2002.
7. B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler. Simple routing strategies for adversarial systems. In *42th. IEEE Symposium on Foundations of Computer Science* (FOCS'01), pages 158–167, 2001.
8. R. Bhattacharjee and A. Goel. Instability of FIFO at arbitrarily low rates in the adversarial queueing model. Technical Report 02-776, Department of Computer Science, University of Southern California, Los Angeles, USA, 2002.
9. A. Borodin, R. Ostrovsky, Y. Rabani. Stability Preserving Transformations: Packet Routing Networks with Edge Capacities and Speeds. In *ACM-SIAM Symposium on Discrete Algorithms* (SODA'01), pages 601–610, 2001.
10. A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, and D. Williamson. Adversarial queueing theory. *Journal of the ACM*, 48(1):13–38, 2001.
11. Charny, A. and Le Boudec, J.-Y. Delay Bounds in a Network With Aggregate Scheduling In *Proc. First International Workshop on Quality of future Internet Services.* Berlin, Germany, 2000.
12. J. Diaz, D. Koukopoulos, S. Nikoletseas, M. Serna, P. Spirakis, and D. Thilikós. Stability and non-Stability of the FIFO Protocol. In *13th annual ACM Symposium on Parallel Algorithms and Architectures* (SPAA'01), pages 48–52, 2001.
13. D. Koukopoulos, M. Mavronicolas, S. Nikoletseas, and P. Spirakis. On the stability of compositions of universally stable, greedy contention-resolution protocols. In D. Malkhi, editor, *Distributed Computing, 16th International Conference*, volume 2508 of *Lecture Notes in Computer Science*, pages 88–102, Springer. 2002.
14. D. Koukopoulos, M. Mavronicolas and P. Spirakis. FIFO is unstable at arbitrarily low rates. ECCC, TR-03-16, 2003.
15. Z. Lotker, B. Patt-Shamir, and A. Rosén. New stability results for adversarial queuing. In *14th ACM Symposium on Parallel Algorithms and Architectures* (SPAA'02), pages 175–182, Winnipeg, Canada, 2002.
16. A. Rosén. A note on models for non-probabilistic analysis of packet switching networks. *Information Processing Letters*, 84:237–240, 2002.