

An Adversarial Construction of Instability Bounds in LIS Networks

Dimitrios Koukopoulos

Abstract—In this work, we study the impact of dynamically changing *link slowdowns* on the stability properties of *packet-switched* networks under the *Adversarial Queueing Theory* framework. Especially, we consider the *Adversarial, Quasi-Static Slowdown Queueing Theory* model, where each link slowdown may take on values in the two-valued set of integers $\{1, D\}$ with $D > 1$ which remain fixed for a long time, under a (w, ρ) -adversary. In this framework, we present an innovative systematic construction for the estimation of adversarial injection rate lower bounds, which, if exceeded, cause instability in networks that use the *LIS* (Longest-in-System) protocol for contention-resolution. In addition, we show that a network that uses the *LIS* protocol for contention-resolution may result in dropping its instability bound at injection rates $\rho > 0$ when the network size and the high slowdown D take large values. This is the best ever known instability lower bound for *LIS* networks.

Keywords—Network stability, quality of service, adversarial queueing theory, greedy scheduling protocols.

I. INTRODUCTION

ONE of the most important features of today large-scale communication networks, such as the *Internet*, is their *robustness*. Robustness is the ability of communication despite network link failures. As the Internet evolves into a ubiquitous communication infrastructure and supports increasingly important services, its dependability in the presence of various failures becomes critical. These failures can degrade system performance and lead to service disruption. Thus, the study of performance and correctness properties of real networks which suffer from link failures becomes a necessity. This study could help on detecting, understanding and overcoming the conditions leading to these mentioned negative effects, as well as helping to their prevention.

A. Motivation-Framework

We are interested in the behavior of *packet-switched networks* in which packets arrive dynamically at the *nodes* and they are routed in discrete time steps across the *links*. Recent years have witnessed a vast amount of work on analyzing

packet-switched networks under *non-probabilistic* assumptions (rather than stochastic ones); we work within a model of *worst-case* continuous packet arrivals, originally proposed by Borodin *et al.* [7] and termed *Adversarial Queueing Theory* to reflect the assumption of an *adversarial* way of packet generation and path determination. A major issue that arises in such a setting is that of *stability*-- will the number of packets in the network remain bounded at all times? The answer to this question may depend on the *rate* of injecting packets into the network, the *slowdown* of the links, which is the time delay which is suffered by outgoing packets in order to be forwarded on a link, and the *protocol* that is used to resolve the conflict when more than one packet wants to cross a given link in a single time step. The underlying goal of our study is to establish the stability properties of networks when packets are injected by an adversary (rather than by an oblivious randomized process) and the link slowdowns are chosen by the same adversary in a dynamic way.

Most studies of packet-switched networks assume that one packet can cross a network link (an edge) in a single time step. This assumption is well motivated when we assume that all network links are identical. However, a packet-switched network can contain different types of links, which is common especially in large-scale networks like Internet. Also, a real network can suffer from link failures due to natural disasters (like hurricanes), human action (like hacker attacks) or by unintentional software failures. Then, it is well motivated to assign a slowdown to each link. Furthermore, if each link slowdown takes on values in the two-valued set of integers $\{1, D\}$ for $D > 1$, D takes on large values and each value remains fixed for a long time, then we can consider approximately as a link failure the assigning of slowdown D to a link, while the assigning of unit slowdown to a link can be considered as the proper service rate. Therefore, the study of the stability behavior of networks under our model of quasi-static slowdowns can be considered as an approximation of the fault-tolerance of a network where links can temporarily fail (infinite slowdown). The goal of this study is to provide an insight towards detecting, understanding, and overcoming the conditions leading to performance degradation and service disruption of today's communication networks during network attacks or failures.

In this work, we consider the impact on the stability behavior of networks if the adversary besides the packet injections in paths which it determines, it also can set the slowdowns of network edges in each time step. This subfield

Manuscript received July 8, 2005.

Dimitrios Koukopoulos is with the Research and Academic Computer Technology Institute, 61 Riga Feraiou Str., 26110 Patras, Greece (corresponding author to provide phone: +30-26410-34473; fax: +30-26410-58075; e-mail: koukopou@ceid.upatras.gr).

of study was initiated by Borodin *et al.* in [8]. Note that we continue to assume uniform packet sizes.

Roughly speaking, a protocol P is *stable* [7] on a network G against an adversary A of rate ρ if there is a constant B (which may depend on G and A) such that the number of packets in the system is bounded at all times by B . On the other hand, a protocol P is *universally stable* [7] if it is stable against every adversary of rate less than 1 and on every network. We also say that a network G is *universally stable* [7] if every greedy protocol is stable against every adversary of rate less than 1 on G . We consider a *greedy* contention-resolution protocol-- ones that always advance a packet across a queue (but one packet at each discrete time step) whenever there resides at least one packet in the queue. This protocol is *LIS* (Longest-in-System) that gives priority to the least recently injected packet into the network.

B. Contribution

We define here the *weakest* possible adversary of dynamically changing network *link slowdowns* in the context of Adversarial Queueing Theory where the adversary may set link slowdowns to any of two integer values 1 and D ($D > 1$ is a parameter called high slowdown). In the classical Adversarial Queueing Theory only one slowdown value is available to the adversary. Moreover, once a link slowdown takes on a value, the value stays fixed for a continuous time period proportional to the number of packets in the system at the time of setting the slowdown to the value. We call this the *Adversarial, Quasi-Static Slowdown Queueing Theory* model. In this framework, we obtain the following results:

- We present for the first time to the best of our knowledge a systematic construction for the estimation of injection rate lower bounds, which, if exceeded, cause instability in networks where the *LIS* protocol is running on their nodes for contention-resolution.
- We study the behavior of the presented systematic construction when it approaches its limits. More specifically, we present a *size-parameterized* network which uses the *LIS* protocol for contention-resolution that is unstable at arbitrarily low injection rates when the network size takes large values and the link slowdowns can be changed dynamically between unit and a large enough value D . The drop of the instability bound is proportional to the increase of the high slowdown D and the network size. This result is the first one that shows instability at arbitrarily low injection rates for a protocol that has been proved universally stable in the classical Adversarial Queueing Model. Till now instability bounds of $\sqrt{2}-1$ or more for the *LIS* protocol have been proved only on adversarial models where the network link capacities can be changed dynamically [8], [14].

The combinatorial constructions of networks and adversaries that we have employed for showing that the *LIS* protocol can be unstable for arbitrarily low injection rates when link slowdowns can change dynamically significantly extend ones that appeared before in [7], [12]-[14]. In more

detail, some of the tools we devise in order to obtain constructions of networks and adversaries that imply improved bounds are the following:

- We employ combinatorial constructions of networks with multiple successively pairs of parallel queues; we judiciously use such paths for the simultaneous injection of various non-overlapping sets of packets. Also, this construction allows the adversary to inject a set of packets at a time period over a path with unit slowdown edges, while the previously injected sets of packets are delayed in another queue due to its high slowdown D .
- We use the technical notion of *investing flow*; this is some special case of packet flow that we use in our adversarial constructions that consist of inductive *phases*. Roughly speaking, an investing flow injects packets in a phase some of which will remain in the system till the beginning of the next phase, in order to guarantee the inductive hypothesis for the next phase.

C. Related Work

Adversarial Queueing Theory was developed by Borodin *et al.* [7] as a more realistic model that replaces traditional stochastic assumptions in Queueing Theory [9] by more robust, worst-case ones. It received a lot of interest and attention in the study of stability and instability issues (see, e.g., [2], [4], [10]-[13]). The universal stability of various natural greedy protocols (*SIS*, *LIS*, *NTS* and *FTG*) was established by Andrews *et al.* [4]. Also, several greedy protocols such as *NTG* have been proved unstable at arbitrarily low rates of injection in [15].

Borodin *et al.* in [8] studied for the first time the impact on stability when the edges of a network can have capacities or slowdowns. They proved that the universal stability of networks is preserved under this varying context. Also, it was shown that many well-known universally stable protocols (*SIS*, *NTS*, *FTG*) do maintain their universal stability when the link capacity or slowdown is changing dynamically, whereas the universal stability of *LIS* is not preserved. More specifically Borodin *et al.* in [8] presented for the first time an instability bound of $\rho > D / (D - 1) > 0.5$ for the *LIS* protocol. This work was further extended by Koukopoulos *et al.* in [14] proving lower bounds of $\sqrt{2}-1$ on the injection rates that guarantee instability for the *LIS* protocol under an adversary of dynamically changing *link capacities*.

Alvarez *et al.* in [1] presented some variations of the adversarial queueing model for dynamic networks, the failure and reliable models where the adversary controls the edge failures. They proved that the universal stability of networks is preserved under this varying context. Also, it was shown that many well-known universally stable protocols (*SIS*, *NTS*, *FTG*) do maintain their universal stability, whereas the universal stability of *LIS* is not preserved. Furthermore, Alvarez *et al.* in [3] proposed three different ways of failure management and studied how they influence on the stability of faulty communication networks.

Two other proposals for dynamic networks have been

initiated in [5] and [6]. In both cases the injected packets are defined by specifying only source and destination, and thus are not forced to follow a pre-specified path. The adversary is restricted to guarantee that a static multi-commodity problem has a solution. Stability results are obtained using a load balancing algorithm, for the case that the adversary injection rate is one, and the packets have a unique common destination. The main difference in both models is that in [6] the adversary has to provide a solution to the associated multi-commodity problem, while in [5] the injection pattern must obey a condition that guarantees the existence of the solution.

D. Road Map

The rest of this paper is organized as follows. Section II presents model definitions. Section III demonstrates instability bounds for the LIS protocol. We conclude, in Section IV, with a discussion of our results and some open problems.

II. PRELIMINARIES

The model definitions are patterned after those in [7], adjusted to reflect the fact that the edge slowdowns may vary arbitrarily as in [8], but we address the weakest possible model of changing slowdowns. We consider that a routing network is modelled by a directed graph $G = (V, E)$. Each node u of the set V represents a communication switch, and each edge e of the set E represents a link between two switches. In each node, there is a buffer (queue) associated with each outgoing link. Time proceeds in discrete time steps. Buffers store packets that are injected into the network with a route, which is a simple directed path in G . A *packet* is an atomic entity that resides at a buffer at the end of any step. It must travel along paths in the network from its *source* to its destination, both of which are nodes in the network. When a packet is injected, it is placed in the buffer of the first link on its route. When a 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. Edges can have different integer slowdowns, which may or may not vary over time. Denote $D_e(t)$ the *slowdown* of the edge e at time step t . That is, we assume that if a packet p is scheduled to traverse the edge e at time t , then packet p completes the traversal of e at time $t + D_e(t)$ and during this time interval, no other packet can be scheduled on e .

Let $D > 1$ be an integer parameter. We demand that for all e and for all t , $D_e(t) \in \{1, D\}$ (i.e. each edge slowdown can get only two values, high and low). We also demand for each edge e that $D_e(t)$ stays at some value for a continuous period of time at least equal to $f(\rho, D)s$ time steps, where s is the number of packets in the system at the time of setting the link slowdown to the value and $f(\rho, D)$ is a function of the injection rate ρ of the adversary in the network and the high link slowdown D . We call this the *Adversarial, Quasi-Static Slowdown Queueing Theory Model*. This model is the weakest possible of the models that are implied by [8].

Any packets that wish to travel along an edge e at a particular time step, but they are not sent, they wait in a queue

for the edge e . At each step, an *adversary* generates a set of requests. A *request* is a *path* specifying the route that will be followed by a packet. In this work, it is assumed, as it is common in packet routing, that all paths are simple paths with no overlapping edges. We say that the adversary generates a set of packets when it generates a set of requested paths. Also, we say that a packet p *requires* an edge e at time t if the edge e lies on the path from its position to its destination at time t . We restrict our study to the case of *non-adaptive routing*, where the path that is traversed by each packet is fixed at the time of injection, so that we are able to focus on queueing rather than routing aspects of the problem. There are no computational restrictions on how the adversary chooses its requests at any given time step.

Fix any arbitrary positive integer $w \geq 1$. For any edge e of the network and any sequence of w consecutive time steps, define $N(w, e)$ to be the number of paths that are injected by the adversary during the time interval of w consecutive time steps requiring to traverse the edge e . For any constant ρ , $0 < \rho \leq 1$, a (w, ρ) -adversary is an adversary that injects packets subject to the following *load condition*: For every edge e and for every sequence τ of w consecutive time steps,

$$N(\tau, e) \leq \rho \sum_{t \in \tau} \frac{1}{D_e(t)} \quad (1)$$

We say that a (w, ρ) -adversary injects packets at rate ρ with *window size* w . The assumption that ρ is less than 1 ensures that it is not necessary a priori that some edge of the network is congested (that happens when $\rho > 1$).

In order to formalize the behavior of a network under the adversarial, quasi-static slowdown queueing theory model, we use the notion of *system*. A triple of the form $\langle G, A, P \rangle$ where G is a network, A is an adversary and P is the used protocol on the network queues is called a *system*.

A *contention-resolution* protocol specifies, for each pair of an edge e and a time step, which packet among those waiting at the tail of the edge e will be moved along the edge e . A *greedy* contention-resolution protocol always specifies some packet to move along the edge e if there are packets waiting to use the edge e . In this work, we restrict attention to the LIS (Longest-in-System) protocol that gives priority to the least recently injected packet into the network.

In the adversarial constructions we study here for proving instability, we split time into *phases*. In each phase, we study the evolution of the *system configuration* by considering corresponding *time rounds*. For each phase, we inductively prove that the number of packets of a specific subset of queues in the system increases in order to guarantee instability. This inductive argument can be applied repeatedly, thus showing instability. Furthermore, we assume that there is a sufficiently large number of packets s_0 in the initial system configuration. This will imply instability results for networks with an *empty* initial configuration, as it was established by Andrews *et al.* [4]. For simplicity, and in a way similar to that in [4] and in works following it, we omit floors and ceilings from our analysis, and we, sometimes, count time steps and

packets only roughly. This may only result to loosing small additive constants, while it implies a gain in clarity.

III. INSTABILITY BOUNDS FOR LIS

In this section, we prove that the LIS protocol can become unstable for arbitrarily low injection rates. In our proof, we denote by X_i the set of packets that are injected into the system in the i^{th} round of a phase. These packet sets are characterized as *investing flows* because only packets from these sets will remain in the system at the beginning of the next phase contributing in packet accumulation.

A. A Parameterized Network Family

We provide here a parameterized family of networks N_l . The motivation that led us to such a parameterization in the network topology is *two-fold*:

- The existence of many pairs of parallel queues in the network allows the adversary to inject an investing flow at a time round over a path with unit slowdown edges, while the previously injected investing flows are delayed in another queue due to its high slowdown D .
- Such a parameterized network topology construction enables a parameterized analysis of the system configuration evolution into distinguished rounds whose number depends on the parameterized network topology.

B. A Systematic Adversarial Construction for Instability

Lower Bounds

The main ideas of the adversarial construction we present here are:

- The accurate tuning of the duration of each round of every phase j (as a function of the high slowdown D , the injection rate ρ and the number of packets in the system at the beginning of phase j , s_j) to maximize the growth of the packet population in the system.
- The careful setting of the slowdowns of some edges to D for specified time intervals in order to accumulate packets.
- The careful injections of packets in order to guarantee that the load condition is satisfied.

We consider an instance of the parameterized network family (network N_l , see Fig. 1). We show:

Theorem 1. Let $\rho'=0.0057$. For the network N_l where $l > 1000$ is a parameter linear to the number of network queues, there is an adversary A of rate ρ that can change the link slowdowns of N_l between the two integer values 1 and $D > 1000$ such that the system $\langle N_l, A, LIS \rangle$ is unstable for every $\rho > \rho'$.

Proof: The construction of the adversary A is broken into phases.

Inductive Hypothesis: At the beginning of phase j (suppose j is even), there are s_j packets that are queued in the queues f_{4l-9}', f_{4l-6}' (in total) requiring to traverse the edges e_0, f_1 .

Induction Step: At the beginning of phase $j+1$, there will

be $s_{j+1} > s_j$ packets that will be queued in the queues f_{4l-9}, f_{4l-6} (in total) requiring to traverse the edges e_1, f_1' .

We will construct an adversary A such that the induction step will hold. Proving that the induction step holds, we ensure that the inductive hypothesis will hold at the beginning of phase $j+1$ for the symmetric edges with an increased value of s_j , $s_{j+1} > s_j$. By the symmetry of the network, repeating the phase construction an unbounded number of times, we will create an unbounded number of packets in the network.




Fig. 1 The network

From the inductive hypothesis, initially, there are s_j packets (that constitute the set of packets S) in the queues f_{4l-9}', f_{4l-6}' requiring to traverse the edges e_0, f_1 . In order to prove the induction step, it is assumed that the set S has a large enough number of $|S| = s_j$ packets in the initial system configuration.

During phase j the adversary plays l rounds of injections as follows:

- **Round 1:** It lasts $|T_1| = s_j$ time steps.

Adversary's behavior. During this round the edge f_1 has high slowdown D , while all the other edges have unit slowdown. The adversary injects a set X_1 of $|X_1| = \rho |T_1|$ packets in the queue e_0 wanting to traverse the edges $e_0, f_2, f_3, f_6, f_7, f_{10}, \dots, f_{4l-9}, f_{4l-6}, e_1, f_1'$. These injections satisfy the load condition because the edges of the assigned path have unit slowdown.

Evolution of the system configuration. The packets of the set S delay the packets of the set X_1 in the queue e_0 because they are longer time in the system than the packets of the set X_1 . At the same time, the packets of the set S are delayed in f_1 due to the high slowdown of the edge f_1 . At the end of this round, the remaining packets of the set S in f_1 are $|S'| = |S| - |T_1|/D$.

- **Round 2:** It lasts $|T_2| = |S'|$ time steps.

Adversary's behavior. During this round the edge f_2 has high slowdown D , while all the other edges have unit slowdown. The adversary injects a set X_2 of $|X_2| = \rho |T_2|$ packets in the queue f_1 requiring to traverse the edges $f_1, f_3, f_6, f_7, f_{10}, \dots, f_{4l-9}, f_{4l-6}, e_1, f_1'$. These packet injections satisfy the load condition because the assigned path consists of edges that have unit slowdown during this round.

Evolution of the system configuration. The packets of the set X_2 are delayed by the packets of the set S' in the queue f_1 because the packets of the set S' are longer time in the

system than the packets of the set X_2 . At the same time, the packets of the set X_1 are delayed in the queue f_2 due to its high slowdown D . Therefore, the remaining packets of the set X_1 in the queue f_2 are $|X_1| - |T_2|/D$.

- **Round 3:** It lasts $|T_3| = |X_1| + |X_2| - |T_2|/D$ time steps.

Adversary's behavior. During this round the edge f_6 has high slowdown D , while all the other edges have unit slowdown. The adversary injects a set X_3 of $|X_3| = \rho |T_3|$ packets in the queue f_3 requiring to traverse the edges $f_3, f_5, f_7, f_{10}, \dots, f_{4l-9}, f_{4l-6}, e_1, f_1'$. These packet injections satisfy the load condition because the assigned path consists of edges that have unit slowdown during this round.

Evolution of the system configuration. The packets of the sets X_1, X_2 delay the packets of the set X_3 in the queue f_3 because they are longer time in the system than the packets of the set X_3 . At the same time, the packets of the sets X_1, X_2 are delayed in f_6 due to the high slowdown of the edge f_6 . Therefore, the remaining packets of the sets X_1, X_2 in the queue f_6 are $|X_1| + |X_2| - |T_2|/D - |T_3|/D$.

- **Round l :** It lasts $|T_l| = \sum_{i=1}^{l-1} |X_i| - \sum_{i=2}^{l-1} |T_i|/D$ steps.

Adversary's behavior. During this round the edge f_{4l-6} has high slowdown D , while all the other edges have unit slowdown. The adversary injects a set X_l of $|X_l| = \rho |T_l|$ packets in the queue f_{4l-9} requiring to traverse the edges $f_{4l-9}, f_{4l-7}, e_1, f_1'$. These packet injections satisfy the load condition because the assigned path consists of edges that have unit slowdown during this round.

Evolution of the system configuration. The packets of the sets X_1, \dots, X_{l-1} delay the packets of the set X_l in the queue f_{4l-9} because they are longer time in the system than the packets of the set X_l . At the same time, the packets of the sets X_1, \dots, X_{l-1} are delayed in f_{4l-6} due to the high slowdown of the edge f_{4l-6} . Therefore, the remaining packets of the sets X_1, \dots, X_{l-1} in the queue f_{4l-6} are $\sum_{i=1}^{l-1} |X_i| - \sum_{i=2}^l |T_i|/D$.

Thus, the number of packets in the queues f_{4l-9}, f_{4l-6} requiring to traverse the edges e_1, f_1' at the end of this round is

$$s_{j+1} = \sum_{i=1}^l |X_i| - \sum_{i=2}^l |T_i|/D \quad (2)$$

But,

$$\sum_{i=1}^l |X_i| = \sum_{i=1}^l \rho |T_i| \quad (3)$$

Thus,

$$s_{j+1} = \rho s_j + \left(\rho - \frac{1}{D}\right) \sum_{i=2}^l |T_i| \quad (4)$$

Moreover,

$$\sum_{i=3}^l |T_i| = \left(\rho + \frac{D-1}{D}\right) \sum_{i=3}^{l-1} |T_i| + \left(2\rho - \frac{\rho+1}{D} + \frac{1}{D^2}\right) |T_l| \quad (5)$$

Solving (5), we take

$$\sum_{i=3}^l |T_i| = \left(2\rho - \frac{\rho+1}{D} + \frac{1}{D^2}\right) |T_l| \frac{1 - \left(\rho + \frac{D-1}{D}\right)^{l-2}}{\frac{1}{D} - \rho} s_j \quad (6)$$

From (4) and (6) we take

$$s_{j+1} = \rho s_j + \left(\rho - \frac{1}{D}\right) \frac{D-1}{D} s_j + \left(\rho - \frac{1}{D}\right) \left(2\rho - \frac{\rho+1}{D} + \frac{1}{D^2}\right) \frac{1 - \left(\rho + \frac{D-1}{D}\right)^{l-2}}{\frac{1}{D} - \rho} s_j \quad (7)$$

In order to have instability, we must have $s_{j+1} > s_j$. From (7) it suffices to hold that

$$\rho + \left(\rho - \frac{1}{D}\right) \frac{D-1}{D} + \left(\rho - \frac{1}{D}\right) \left(2\rho - \frac{\rho+1}{D} + \frac{1}{D^2}\right) \frac{1 - \left(\rho + \frac{D-1}{D}\right)^{l-2}}{\frac{1}{D} - \rho} > 1 \quad (8)$$

If we let $\rho = 0.0057$, $D = 1000$ and $l = 1000$, the inequality holds. Thus, for $\{D, l\} > 1000$ the inequality holds, too. This argument can be repeated for an infinite number of phases ensuring that the number of packets in the system at the end of a phase will be larger than at the beginning of the phase. ■

Corollary 1. If the parameters D and l that are related to the high link slowdown and the number of network queues correspondingly tend to infinity then there is a network N_l and an adversary A of rate $\rho > 0$ such that the system $\langle N_l, A, LIS \rangle$ is unstable.

Proof: From Theorem 1, we can construct an adversary that leads the system $\langle N_l, A, LIS \rangle$ to instability. It suffices (8) in order to guarantee instability. If $D \rightarrow \infty$, it holds that $\frac{1}{D^k} \rightarrow \infty$ for all $k \geq 1$. Then, (8) becomes

$$\rho > \frac{1}{2(\rho+1)^{l-2}} \quad (9)$$

But, if $l \rightarrow \infty$ and $x > 0$, it holds that $(1+x)^{l-2} \rightarrow \infty$. Therefore, when the parameters D and l tend to infinity (10) holds for $\rho > 0$. Note that if we have a sequence of equations $f_{(D,l)}(\rho)$ and there exists the limit $\lim_{(D,l) \rightarrow \infty} f_{(D,l)}(\rho) = f_\infty(\rho)$, then it holds fundamentally by the theory of function limits that if $\rho(D,l)$ is the root of $f_{(D,l)}(\rho) = 0$, then $\lim_{(D,l) \rightarrow \infty} \rho(D,l)$ is the root of $f_\infty(\rho)$. Therefore, for $\rho > 0$ the system is unstable. ■

IV. CONCLUSION

In this work, we studied how the dynamic changing of the network link slowdowns affects the instability properties of the LIS contention-resolution protocol using an extension of the adversarial model that was first initiated by Borodin *et al.* in [8], the Adversarial, Quasi-Static Slowdown Queueing Theory Model. In particular, we proved that the LIS protocol can be unstable at arbitrarily low injection rates due to dynamic link slowdowns.

However, a lot of problems remain open. Our results suggest that, for every unstable network, its instability bound

in the model of quasi-static slowdowns may be lower than for the classical adversarial queueing model or other dynamic adversarial models. Proving (or disproving) this remains an open problem. Studying the impact of dynamically changing link slowdowns on other greedy protocols and networks is another interesting problem. Finally, it is worth to receive attention the study of the stability behavior of networks and protocols in environments where the adversary controls the movement of the network nodes.

REFERENCES

- [1] C. Alvarez, M. Blesa, J. Diaz, M. Serna and A. Fernandez, "Adversarial Models for Priority-Based Networks," *Networks*, Vol. 45, No. 1, pp. 23-35, January 2005.
- [2] C. Alvarez, M. Blesa and M. Serna, "A Characterization of Universal Stability in the Adversarial Queueing model," *SIAM Journal on Computing*, Vol. 34, No. 1, pp. 41-66, 2004.
- [3] C. Alvarez, M. Blesa and M. Serna, "The Impact of Failure Management on the Stability of Communication Networks," in *Proceedings of the 10th International Conference on Parallel and Distributed Systems*, pp. 153-160, July 2004.
- [4] M. Andrews, B. Awerbuch, A. Fernandez, J. Kleinberg, T. Leighton and Z. Liu, "Universal Stability Results for Greedy Contention-Resolution Protocols," *Journal of the ACM*, Vol. 48, No. 1, pp. 39-69, January 2001.
- [5] E. Anshelevich, D. Kempe and J. Kleinberg, "Stability of Load Balancing Algorithms in Dynamic Adversarial Systems," in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pp. 399-406, May 2002.
- [6] B. Awerbuch, P. Berenbrink, A. Brinkmann and C. Scheideler, "Simple Routing Strategies for Adversarial Systems," in *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pp. 158-167, October 2001.
- [7] A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan and D. Williamson, "Adversarial Queueing Theory," *Journal of the ACM*, Vol. 48, No. 1, pp. 13-38, January 2001.
- [8] A. Borodin, R. Ostrovsky and Y. Rabani, "Stability Preserving Transformations: Packet Routing Networks with Edge Capacities and Speeds," in *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 601-610, January 2001.
- [9] H. Chen and D. D. Yao, *Fundamentals of Queueing Networks*. Springer-Verlag, 2000.
- [10] J. Diaz, D. Koukopoulos, S. Nikolettseas, M. Serna, P. Spirakis and D. Thilikos, "Stability and Non-Stability of the FIFO Protocol," in *Proceedings of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures*, pp. 48-52, July 2001.
- [11] D. Koukopoulos, M. Mavronicolas, S. Nikolettseas and P. Spirakis, "On the Stability of Compositions of Universally Stable, Greedy, Contention-Resolution Protocols," in *Proceedings of the 16th International Symposium on Distributed Computing*, LNCS 2508, pp. 88-102, October 2002.
- [12] D. Koukopoulos, M. Mavronicolas, S. Nikolettseas and P. Spirakis, "The Impact of Network Structure on the Stability of Greedy Protocols," in *Proceedings of the 5th Italian Conference on Algorithms and Complexity*, LNCS 2653, pp. 251-263, May 2003. Also, accepted in the *Journal Theory of Computing Systems*.
- [13] D. Koukopoulos, S. Nikolettseas and P. Spirakis, "Stability Issues in Heterogeneous and FIFO Networks under the Adversarial Queueing Model," in *Proceedings of the 8th International Conference on High Performance Computing*, LNCS 2228, pp. 3-14, Dec. 2001. Invited Keynote Address.
- [14] D. Koukopoulos, M. Mavronicolas and P. Spirakis, "Instability of Networks with Quasi-Static Link Capacities," in *Proceedings of the 10th International Colloquium on Structural Information and Communication Complexity*, pp. 179-194, June 2003.
- [15] P. Tsaparas, "Stability in Adversarial Queueing Theory," *M.Sc. Thesis*, Computer Science Department, University of Toronto, 1997.