Heterogenous Networks Can Be Unstable at Arbitrarily Low Injection Rates*

Dimitrios Koukopoulos[1] and Stavros D. Nikolopoulos[2]

[1] Department of Cultural Heritage Management & New Technologies, University of Ioannina, GR-30100 Agrinio, Greece
koukopou@ceid.upatras.gr
[2] Department of Computer Science, University of Ioannina, GR-45110 Ioannina, Greece
stavros@cs.uoi.gr

Abstract. A distinguishing feature of today's large-scale platforms for distributed computation and communication, such as the *Internet,* is their *heterogeneity,* predominantly manifested by the fact that a wide variety of *communication protocols* are simultaneously running over different distributed hosts. A fundamental question that naturally poses itself for such common settings of heterogeneous distributed systems concerns their ability to preserve or restore an acceptable level of performance during link failures. In this work, we address this question for the specific case of stability properties of greedy, contention-resolution protocols operating over a *packet-switched* communication network that suffers from *link slowdowns*. We focus on the *Adversarial Queueing Theory* framework, where an adversary controls the rates of packet injections and determines packet paths. In addition, the power of the adversary is enhanced to include the manipulation of *link slowdowns*. Within this framework, we show that the composition of LIS (*Longest-in-System*) with any of SIS (*Shortest-in-System*), NTS (*Nearest-to-Source*) and FTG (*Furthest-to-Go*) protocols is unstable at rates $\rho > 0$ when the network size and the link slowdown take large values. These results represent the current record for instability bounds on injection rate for compositions of greedy protocols over dynamic adversarial models, and also suggest that the potential for instability incurred by the composition of *two* greedy protocols may be *worse* than that of some *single* protocol.

1 Introduction

Motivation-Framework. Some of the most important features of contemporary large-scale platforms for distributed communication and computation, such as the *Internet,* is their *robustness* and *heterogeneity*. Robustness is the ability of

* This research was co-funded by the European Union in the framework of the program "Pythagoras II" of the "Operational Program for Education and Initial Vocational Training" of the 3rd Community Support Framework of the Hellenic Ministry of Education, funded by national sources and the European Social Fund (ESF).

T. Calamoneri, I. Finocchi, G.F. Italiano (Eds.): CIAC 2006, LNCS 3998, pp. 93–104, 2006.

communication despite network link failures, while heterogeneity comes around in many different flavors. For example, the specifics of how the computers in different parts of the network are connected with each other, and the properties of the links that foster the interconnection, is difficult to characterize uniformly. Moreover, although, conceptually, the Internet uses a unified set of protocols, in practice each protocol has been implemented with widely varying features (and of course bugs) [9, 11]. As the Internet evolves into a ubiquitous communication infrastructure that supports multiple protocols running on different network hosts, 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 heterogeneous distributed systems 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.

Objectives. 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 composition of *protocols* running on different network hosts in order to resolve packet conflicts. The underlying goal of our study is to establish the stability properties of heterogeneous networks when packets are injected by an adversary and the link slowdowns are chosen by the same adversary in a dynamic way.

Model of Quasi-Static Slowdowns. 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 $\{1, \mathsf{D}\}$ for $\mathsf{D} > 1$, 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 and protocols 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)[1]. 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 embark on a study of the impact of heterogeneity of distributed systems on their performance properties if the adversary can determine the paths of packet injections along with the slowdowns of network edges in each time step. More specifically, we wish to pose the general question of which performance properties of heterogeneous *packet-switched* networks (where compositions of protocols are running on different network hosts) are maintained and which are not in the presence of link failures. This subfield of study was initiated by Borodin *et al.* in [8] in the case of networks where a single protocol is responsible for the resolution of packet conflicts. Note that we continue to assume uniform packet sizes.

Stability. Roughly speaking, a protocol P is *stable* [7] on a network $\mathcal{G}$ against an adversary $\mathcal{A}$ of rate ρ if there is a constant B (which may depend on $\mathcal{G}$ and $\mathcal{A}$) 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. Here, we consider four *greedy*, universally stable, contention-resolution protocols under the Adversarial Queueing Theory (Table 1).

Table 1. Greedy protocols considered in this paper (**US**stands for universally stable)

Protocol name	Which packet it advances:	**US**
SIS (*Shortest-In-System*)	The most recently injected packet	$\surd$ [4]
LIS (*Longest-In-System*)	The least recently injected packet	$\surd$ [4]
FTG (*Furthest-To-Go*)	The furthest packet from its destination	$\surd$ [4]
NTS (*Nearest-To-Source*)	The nearest packet to its origin	$\surd$ [4]

Contribution. We define here the *weakest* possible adversary of dynamically changing network *link slowdowns* in the context of Adversarial Queueing Theory (AQM) where the adversary may set link slowdowns to any of two integer values 1 and D ($\mathsf{D} > 1$ is a parameter called high slowdown).[2] 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 (AQSSQM). In this framework, we establish that the

[1] However, infinite link slowdown is only an approximation of link failure, because in a slowdown the packet has left the queue and is being transmitted; however, when a failure occurs, the packet is not being transmitted but stored somewhere, and thus it participates later in the queue scheduling.

[2] In AQM only one slowdown value is available to the adversary.

composition of LIS with any of SIS, NTS and FTG protocols is unstable for arbitrarily low injection rates. We prove that increasing the network size along with dynamic changing of link slowdowns can drop to arbitrarily low values the lower bound on injection rate that guarantees instability for heterogeneous networks. To show this, we provide interesting combinatorial constructions of a *size-parameterized* network where we specify the contention-resolution protocol to be used to each queue. For purpose of completeness and comparison, we summarize, in Table 2, all results that are shown in this work and, also, in [16] (for AQM) and [18] (for the Adversarial Quasi-Static Queueing Model - AQSQM), concerning instability bounds on the injection rate for the composition pairs LIS-SIS, LIS-NTS and LIS-FTG.

Table 2. Instability bounds of the compositions of LIS with any of the SIS, NTS, and FTG protocols in AQM vs. AQSQM vs. AQSSQM

	Instability (AQM)	Instability (AQSQM)	Instability (AQSSQM)
LIS − SIS	$\rho > 0.5$ [16, Thm. 3.1]	$\rho > \sqrt{2} - 1$ [18, Thm. 2]	$\rho > 0$ [Thm. 1]
LIS − NTS	$\rho > 0.5$ [16, Thm. 3.1]	$\rho > \sqrt{2} - 1$ [18, Thm. 2]	$\rho > 0$ [Thm. 2]
LIS − FTG	$\rho > 0.5$ [16, Thm. 3.1]	$\rho > \sqrt{2} - 1$ [18, Thm. 2]	$\rho > 0$ [Thm. 3]

The combinatorial constructions of networks and adversaries that we have employed for showing that certain compositions of universally stable protocols can be unstable for arbitrarily low injection rates when link slowdowns can change dynamically, significantly extend ones that appeared before in [7, 15, 16, 18]. 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 notions of *investing flow* and *short flow*; these are some special cases of packet flows used in our adversarial constructions consisting 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; on the other hand, short flows consist of packets injected on judiciously chosen links of the network and their role is to delay the investing flows.

Related Work. The issue of composing distributed protocols (resp., objects) to obtain other protocols (resp., objects), and the properties of the resulting (*composed*) protocols (resp., objects), has a rich record in Distributed Computing Theory (see, e.g., [20]). For example, Herlihy and Wing [13] establish that

a composition of *linearizable* memory objects (possibly distinct), each managed by its own protocols, preserves linearizability. Robustness has been extensively studied in the context of fault-tolerant distributed systems. A landmark paper on failures in Tandem systems and the techniques to prevent them is [12]. In parallel and even earlier, a mathematical framework was developed in the Operations Research world to manage the robustness and risk in systems composed of various components [5].

Adversarial Queueing Theory [7] received a lot of interest in the study of stability and instability issues (see, e.g., [2, 4, 10, 15, 17, 21]). 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 (Nearest-To-Go) have been proved unstable at arbitrarily low rates of injection in [21]. The subfield of study of the stability properties of compositions of universally stable protocols was introduced by Koukopoulos *et al.* in [15, 16, 17] where lower bounds of 0.683, 0.519 and 0.5 on the injection rates that guarantee instability for the composition pairs LIS-SIS, LIS-NTS and LIS-FTG were presented.

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 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. This work was further extended by Koukopoulos *et al.* in [18, Theorems 2, 3] proving lower bounds of $\sqrt{2}-1$ on the injection rates that guarantee instability for the LIS protocol and its compositions with the SIS, NTS and FTG protocols under dynamically changing link capacities. Also, Koukopoulos in [14] studied the impact of link slowdowns on network stability when a single protocol is used or a forbidden subgraph for universal stability is induced. Moreover, in [1, 3] there have been generalizations of the adversarial queueing theory to networks with dynamic failures. Finally, in [6] it is proposed a generalization of the adversarial queueing theory where the network traffic flow is continuous in time and arbitrary packet lengths, link speeds and link propagation delays are allowed.

2 Preliminaries

The model definitions are patterned after those in [7, Section 3], adjusted to reflect the fact that the edge slowdowns may vary arbitrarily as in [8, Section 2], but we address the weakest possible model of changing slowdowns. We consider that a routing network is modelled by a directed graph $\mathcal{G} = (V, E)$. Each node $u \in V$ represents a communication switch, and each edge $e \in E$ 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 $\mathcal{G}$. 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 $\mathsf{D}_e(t)$ 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 $t + \mathsf{D}_e(t)$ and during this time interval, no other packet can be scheduled on e.

Let $\mathsf{D} > 1$ be an integer parameter. We demand that $\forall e$ and $\forall t$ $\mathsf{D}_e(t) \in \{1, \mathsf{D}\}$. We also demand for each edge e that $\mathsf{D}_e(t)$ stays at some value for a continuous period of time at least equal to $f(\rho, \mathsf{D})s$ time steps, where s is the number of packets in the system at the time of setting the link slowdown to the value and $f(\rho, \mathsf{D})$ 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*. Our model is different from the failure model in [1, 3] because in our model a packet p is delayed after leaving the queue of the edge e, while in the failure model p waits in the queue of e.

Any packets that wish to travel along an edge e at a particular time step, but are not sent, wait in a queue for 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.[3] 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 e lies on the path from its position to its destination at time t.

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, $N(\tau, e) \leq \rho \sum_{t \in \tau} \frac{1}{\mathsf{D}_e(t)}$. We say that a (w, ρ)-adversary injects packets at rate ρ with *window size* w. The assumption that $\rho \leq 1$ ensures that it is not necessary a priori that some edge of the network is overloaded.

In order to formalize the behavior of a network, we use the notions of *system* and *system configuration*. A triple of the form $\langle \mathcal{G}, \mathcal{A}, \mathsf{P} \rangle$ where $\mathcal{G}$ is a network, $\mathcal{A}$ is an adversary and P is the used protocol (or list of protocols) on the network queues is called a system. In every time step t, the current configuration C^t of a system $\langle \mathcal{G}, \mathcal{A}, \mathsf{P} \rangle$ is a collection of sets $\{S_e^t : e \epsilon \mathcal{G}\}$, such that S_e^t is the set of packets waiting in the queue of the edge e at the end of step t.

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

[3] In this work, it is assumed, as it is common in packet routing, that all paths are simple paths where edges cannot be overlapped, while vertices can be overlapped.

will imply instability results for networks with an *empty* initial configuration, as it was established in [4, Lemma 2.9]. For simplicity, and in a way similar to that in [4], 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.

3 Unstable Compositions of Protocols

In this section, we prove that the composition of the LIS protocol with any of SIS, NTS and FTG protocols can become unstable for arbitrarily low injection rates. Before proceeding to the adversarial constructions for proving instability we give two basic definitions.

Definition 1. *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.*

Definition 2. *We denote by S_i the set of packets the adversary injects into the system in the i^{th} round of a phase. These packet sets are characterized as* short flows *because they are injected on judiciously chosen links of the network for delaying investing flows.*

3.1 A Parameterized Network Family $\mathcal{G}_l$

We provide here a parameterized family of networks $\mathcal{G}_l$ (see Figure 1). The motivation that led us to such a parameterization in the network topology is *two-fold*: (a) 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. Also, this structure permits the simultaneous injection of an investing flow on one queue of a pair, and a short flow on the other, without violating the rule of the restricted adversarial model. (b) 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. In LIS-FTG composition, the parameterization, besides the parallel edges, includes additional chains of queues for the exploitation of FTG in blocking investing flows.

3.2 Parameterized Adversarial Constructions

The main ideas of the adversarial constructions we present are: (a) 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, (b) the careful setting of the slowdowns of some edges to D for specified time intervals in order to accumulate packets, and (c) the careful injections of packets that guarantee that the load condition is satisfied.

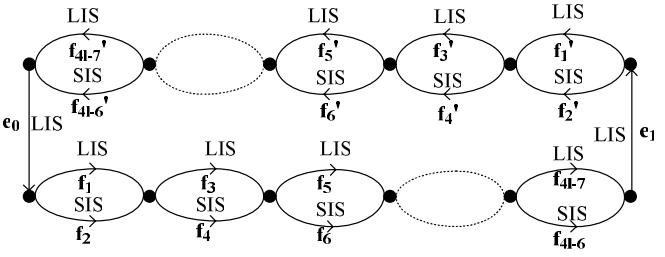

Fig. 1. The network $\mathcal{G}_l$

Theorem 1. *Let $\rho' = 0.0056$. For the network $\mathcal{G}_l$ where $l > 1000$ is a parameter linear to the number of network queues there is an adversary $\mathcal{A}_1$ of rate ρ that can change the link slowdowns of $\mathcal{G}_l$ between the two integer values 1 and $\mathsf{D} > 1000$ such that the system $\langle \mathcal{G}_l, \mathcal{A}_1, \mathsf{LIS}, \mathsf{SIS} \rangle$ is unstable for every $\rho > \rho'$. When $\{\mathsf{D}, l\} \to \infty$ the system $\langle \mathcal{G}_l, \mathcal{A}_1, \mathsf{LIS}, \mathsf{SIS} \rangle$ is unstable for $\rho > 0$.*

Proof. Consider an instance of the parameterized network family (network $\mathcal{G}_l$, see Figure 1). The edges $e_0, e_1, f_1, f_3, f_5, \ldots, f_{4l-7}$ $f_1^{'}, f_3^{'}, f_5^{'}, \ldots, f_{4l-7}^{'}$ of $\mathcal{G}_l$ use the LIS protocol, while the remaining edges of $\mathcal{G}_l$ use the SIS protocol. The construction of the adversary $\mathcal{A}_1$ 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 $\mathcal{A}_1$ 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.

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. 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 $e_0, f_2, f_3, f_6, f_7, f_{10}, \ldots, f_{4l-9}, f_{4l-6}, e_1, f_1^{'}$.

Evolution of the system configuration. The packets of the set S delay the packets of the set X_1 in the queue e_0 that uses the LIS protocol 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 $|S^{'}| = |S| - |T_1|/\mathsf{D}$.

Round 2: It lasts $|T_2| = |S^{'}|$ time steps. 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 $f_1, f_3, f_6, f_7, f_{10}, \ldots, f_{4l-9}, f_{4l-6}, e_1, f_1^{'}$ and a set S_2 of $|S_2| = \rho|T_2|/\mathsf{D}$ packets in the queue f_2 requiring to traverse the edge f_2.

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 that uses the LIS protocol 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 that uses the SIS protocol due to its high slowdown D and the packets of the set S_2 that are shorter time in the system than the packets of the set X_1. Therefore, the remaining packets of the set X_1 in the queue f_2 are $|X_1| + |S_2| - |T_2|/\mathsf{D} = |X_1| + (\rho - 1)|T_2|/\mathsf{D}$.

Round 3: It lasts $|T_3| = |X_1| + |X_2| + (\rho - 1)|T_2|/\mathsf{D}$ time steps. 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 $f_3, f_5, f_7, f_{10}, \ldots, f_{4l-9}, f_{4l-6}, e_1, f_1^{'}$ and a set S_3 of $|S_3| = \rho|T_3|/\mathsf{D}$ packets in the queue f_6 requiring to traverse the edge f_6.

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 that uses the LIS protocol 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 that uses the SIS protocol due to the high slowdown of the edge f_6 and the packets of the set S_3 that are shorter time in the system than the packets of the sets X_1, X_2. Therefore, the remaining packets of the sets X_1, X_2 in the queue f_6 are $|X_1| + |X_2| + (\rho - 1)\frac{|T_2|+|T_3|}{\mathsf{D}}$.

Round l: It lasts $|T_l| = \sum_{i=1}^{l-1} |X_i| - (\rho - 1)\sum_{i=2}^{l-1} |T_i|/\mathsf{D}$ time steps. 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 $f_{4l-9}, f_{4l-7}, e_1, f_1^{'}$ and a set S_l of $|S_l| = \rho|T_l|/\mathsf{D}$ packets in the queue f_{4l-6} requiring to traverse the edge f_{4l-6}.

Evolution of the system configuration. The packets of the sets $X_1, \ldots, X_{l-1}$ delay the packets of the set X_l in the queue f_{4l-9} that uses the LIS protocol 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 $X_1, \ldots, X_{l-1}$ are delayed in f_{4l-6} that uses the SIS protocol due to the high slowdown of the edge f_{4l-6} and the packets of the set S_l that are shorter time in the system than the packets of the sets $X_1, \ldots, X_{l-1}$.

Therefore, the remaining packets of the sets $X_1, \ldots, X_{l-1}$ in the queue f_{4l-6} are $\sum_{i=1}^{l-1} |X_i| + (\rho - 1) \sum_{i=2}^{l} |T_i|/\mathsf{D}$.

Thus, the number of packets in the queues f_{4l-9}, f_{4l-6} requiring to traverse $e_1, f_1^{'}$ at the end of this round is $s_{j+1} = \rho s_j + (\rho + \frac{\rho-1}{\mathsf{D}}) \sum_{i=2}^{l} |T_i|$. Moreover, $\sum_{i=3}^{l} |T_i| = (\rho + \frac{\mathsf{D}+\rho-1}{\mathsf{D}}) \sum_{i=3}^{l-1} |T_i| + (2\rho - \frac{1}{\mathsf{D}} - \frac{\rho-1}{\mathsf{D}^2})|T_1|$. Thus, $s_{j+1} = \rho s_j + (\rho + \frac{\rho-1}{\mathsf{D}}) \frac{\mathsf{D}-1}{\mathsf{D}} s_j + (\rho + \frac{\rho-1}{\mathsf{D}})(2\rho - \frac{1}{\mathsf{D}} - \frac{\rho-1}{\mathsf{D}^2}) \frac{1-(\rho+\frac{\mathsf{D}+\rho-1}{\mathsf{D}})^{l-2}}{\frac{1-(\mathsf{D}+1)\rho}{\mathsf{D}}} s_j$. In order to have instability, we must have $s_{j+1} > s_j$. Therefore, for instability it suffices $\rho + (\rho + \frac{\rho-1}{\mathsf{D}}) \frac{\mathsf{D}-1}{\mathsf{D}} + (\rho + \frac{\rho-1}{\mathsf{D}})(2\rho - \frac{1}{\mathsf{D}} - \frac{\rho-1}{\mathsf{D}^2}) \frac{1-(\rho+\frac{\mathsf{D}+\rho-1}{\mathsf{D}})^{l-2}}{\frac{1-(\mathsf{D}+1)\rho}{\mathsf{D}}} > 1$. If we let $\rho = 0.0056$, $\mathsf{D} = 1000$ and $l = 1000$, the inequality holds. Thus, for $\{\mathsf{D}, l\} > 1000$ the inequality holds, too.

When $\mathsf{D} \to \infty$, it holds that $\frac{1}{\mathsf{D}^k} \to 0$ for all $k \geq 1$. Then, our inequality becomes $2\rho(\rho+1)^{l-2} > 1$. Thus, $\rho > \frac{1}{2(\rho+1)^{l-2}}$. When $l \to \infty$ and $x > 0$, it holds that $(1+x)^{l-2} \to \infty$. Therefore, for $\{\mathsf{D}, l\} \to \infty$ the inequality $\rho > \frac{1}{2(\rho+1)^{l-2}}$ holds for $\rho > 0$. Note that if we have a sequence of equations $f_{\mathsf{D},l}(\rho)$ and there exists the limit $\lim_{\{\mathsf{D},l\}\to\infty} f_{\mathsf{D},l}(\rho) = f_\infty(\rho)$, then it holds fundamentally by the theory of function limits that if $\rho(\mathsf{D}, l)$ is the root of $f_{\mathsf{D},l}(\rho) = 0$, then $\lim_{\{\mathsf{D},l\}\to\infty} \rho(\mathsf{D}, l)$ is the root of $f_\infty(\rho)$. Therefore, for $\rho > 0$ the system is unstable. This argument can be repeated for an infinite number of phases showing that the number of packets in the system increases forever for $\rho > 0$. □

With a similar adversarial construction to Theorem 1, we show that the composition of the LIS and NTS protocols can become unstable for arbitrarily low injection rates considering an instance of the parameterized network family (network $\mathcal{G}_l$, see Figure 1). The network $\mathcal{G}_l$ is also used for proving the instability of the composition of the LIS and SIS protocols. However in this case, the edges $f_2, f_4, f_6, \ldots, f_{4l-6}$ $f_2^{'}, f_4^{'}, f_6^{'}, \ldots, f_{4l-6}^{'}$ of $\mathcal{G}_l$ use the NTS protocol instead of the SIS protocol, while the remaining edges of $\mathcal{G}_l$ use the LIS protocol. Thus, the following theorem, analogous to Theorem 1, holds.

Theorem 2. *Let $\rho' = 0.0056$. For the network $\mathcal{G}_l$ where $l > 1000$ is a parameter linear to the number of network queues there is an adversary $\mathcal{A}_2$ of rate ρ that can change the link slowdowns of $\mathcal{G}_l$ between the two integer values 1 and $\mathsf{D} > 1000$ such that the system $\langle \mathcal{G}_l, \mathcal{A}_2, \mathsf{LIS}, \mathsf{NTS} \rangle$ is unstable for every $\rho > \rho'$. When $\{\mathsf{D}, l\} \to \infty$ the system $\langle \mathcal{G}_l, \mathcal{A}_2, \mathsf{LIS}, \mathsf{NTS} \rangle$ is unstable for $\rho > 0$.*

Similarly, we show that the composition of the LIS and FTG protocols can become unstable for arbitrarily low injection rates considering an instance $\mathcal{G}_l^{'}$ of the parameterized network family $\mathcal{G}_l$ (see Figure 2). The topology of the network $\mathcal{G}_l^{'}$ has a significant difference with the networks that are used for proving Theorems 1, and 2. The network $\mathcal{G}_l^{'}$ contains additional paths, comparing to the other three cases, that start at queues that use the FTG protocol. These paths have sufficient lengths, such that the injected short flows have the same blocking effects over the injected investing flows when they conflict in queues that use FTG, as happens in LIS-SIS and LIS-NTS cases. Thus, the following theorem, analogous to Theorem 1 and Theorem 2, holds.

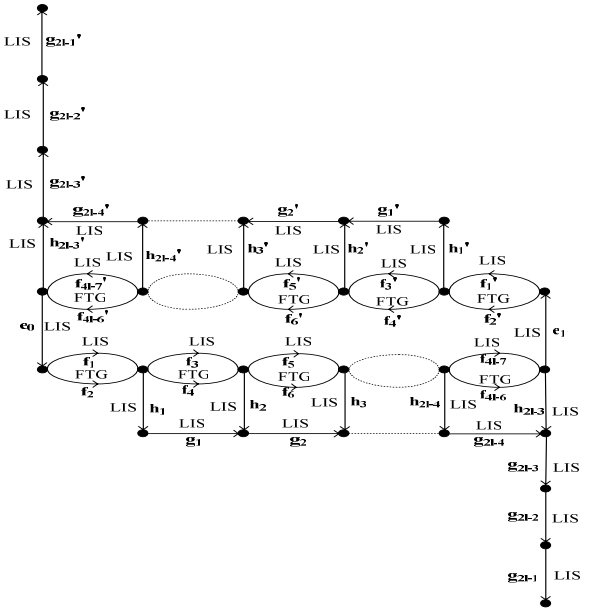

Fig. 2. The network $\mathcal{G}_l^{'}$

Theorem 3. *Let* $\rho' = 0.0056$. *For the network* $\mathcal{G}_l^{'}$ *where* $l > 1000$ *is a parameter linear to the number of network queues there is an adversary* $\mathcal{A}_3$ *of rate* ρ *that can change the link slowdowns of* $\mathcal{G}_l^{'}$ *between the two integer values 1 and* $\mathsf{D} > 1000$ *such that the system* $\langle \mathcal{G}_l^{'}, \mathcal{A}_3, \mathsf{LIS}, \mathsf{FTG} \rangle$ *is unstable for every* $\rho > \rho'$. *When* $\{\mathsf{D}, l\} \to \infty$ *the system* $\langle \mathcal{G}_l^{'}, \mathcal{A}_3, \mathsf{LIS}, \mathsf{FTG} \rangle$ *is unstable for* $\rho > 0$.

4 Conclusions

In this work, we studied how the dynamic changing of link slowdowns affects the instability properties of compositions of contention-resolution protocols that include LIS. However, we do not have any clue what happens with compositions of protocols that do not include LIS. Also, 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 model. Proving (or disproving) this remains an open problem.

References

1. C. Alvarez, M. Blesa, J. Diaz, A. Fernandez, M. Serna, Adversarial Models for Priority-Based Networks, *Proc. of the 28th Int'l Symposium on Mathematical Foundations of Computer Science*, 2003, LNCS. 2747, pp. 142–151.
2. C. Alvarez, M. Blesa, M. Serna, A Characterization of Universal Stability in the Adversarial Queuing model, SIAM Journal on Computing, 34 (2004) 41–66.

3. C. Alvarez, M. Blesa, M. Serna, The Impact of Failure Management on the Stability of Communication Networks, *Proc. of the 10th Int'l Conference on Parallel and Distributed Systems*, 2004, pp. 153–160.
4. M. Andrews, B. Awerbuch, A. Fernández, J. Kleinberg, T. Leighton, Z. Liu, Universal Stability Results for Greedy Contention-Resolution Protocols, Journal of the ACM, 48 (2001) 39–69.
5. R. Barlow and F. Proschan, Statistical Analysis of Reliability and LifeTesting Models, New York: Holt, Rinehart and Winston, 1975.
6. M. Blesa, D. Calzada, A. Fernández, L. López, A. Martínez, A. Santos, M. Serna, Adversarial Queueing Model for Continuous Network Dynamics, *Proc. of the 30th Int'l Symposium on Mathematical Foundations of Computer Science*, 2005, LNCS. 3618, pp. 144–155.
7. A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, D. Williamson, Adversarial Queueing Theory, Journal of the ACM, 48 (2001) 13–38.
8. A. Borodin, R. Ostrovsky, Y. Rabani, Stability Preserving Transformations: Packet Routing Networks with Edge Capacities and Speeds, *Proc. of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2001, pp. 601–610.
9. D. Clark, The Design Philosophy of the DARPA Internet Protocols, ACM Computer Communication Reviews, 18 (1988) 106-114.
10. J. Diaz, D. Koukopoulos, S. Nikoletseas, M. Serna, P. Spirakis, D. Thilikos, Stability and Non-Stability of the FIFO Protocol, *Proc. of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures*, 2001, pp. 48–52.
11. S. Floyd and V. Paxson, Difficulties in Simulating the Internet, IEEE/ACM Transactions on Networking, 9 (2001) 392–403.
12. J. Gary, Why do computers stop and what can be done about it?, *Symposium on Reliability in Distributed Software and Database Systems*, 1986.
13. M. P. Herlihy and J. Wing, Linearizability: A Correctness Condition for Concurrent Objects, *Proc. of the ACM Transactions on Programming Languages and Systems*, 1990, Vol. 12, No. 3, pp. 463–492.
14. D. Koukopoulos, The Impact of Dynamic Link Slowdowns on Network Stability, *Proc. of the 8th Int'l Symposium on Parallel Architectures, Algorithms and Networks*, 2005, pp. 340–345.
15. D. Koukopoulos, M. Mavronicolas, S. Nikoletseas, P. Spirakis, On the Stability of Compositions of Universally Stable, Greedy, Contention-Resolution Protocols, *Proc. of the 16th Int'l Symposium on DIStributed Computing*, 2002, LNCS. 2508, pp. 88–102.
16. D. Koukopoulos, M. Mavronicolas, S. Nikoletseas, P. Spirakis, The Impact of Network Structure on the Stability of Greedy Protocols, Theory of Computing Systems, 38 (2005) 425–460.
17. D. Koukopoulos, S. Nikoletseas, P. Spirakis, Stability Issues in Heterogeneous and FIFO Networks under the Adversarial Queueing Model, *Proc. of the 8th Int'l Conference on High Performance Computing*, 2001, LNCS. 2228, pp. 3–14.
18. D. Koukopoulos, M. Mavronicolas, P. Spirakis, Instability of Networks with Quasi-Static Link Capacities, *Proc. of the 10th Int'l Colloquium on Structural Information and Communication Complexity, Carleton Scientific*, 2003, pp. 179–194.
19. Z. Lotker, B. Patt-Shamir, A. Rosén, New Stability Results for Adversarial Queuing, SIAM Journal on Computing, 33 (2004) 286–303.
20. N. Lynch, Distributed Algorithms, Morgan Kaufmann, 1996.
21. P. Tsaparas, Stability in Adversarial Queueing Theory, M.Sc. Thesis, Computer Science Department, University of Toronto, 1997.