

Performance of Scheduling Policies in Adversarial Networks with Non-synchronized Clocks

Antonio Fernández Anta · José Luis López-Presa ·
M. Araceli Lorenzo · Pilar Manzano ·
Juan Martínez-Romo · Alberto Mozo ·
Christopher Thraves

Published online: 16 June 2009
© Springer Science+Business Media, LLC 2009

Abstract In this paper we generalize the Continuous Adversarial Queuing Theory (CAQT) model (Blesa et al. in MFCS, Lecture Notes in Computer Science, vol. 3618, pp. 144–155, 2005) by considering the possibility that the router clocks in the network are not synchronized. We name the new model Non Synchronized CAQT (NSCAQT). Clearly, this new extension to the model only affects those scheduling policies that use some form of timing. In a first approach we consider the case in which although not synchronized, all clocks run at the same speed, maintaining constant differences. In this case we show that all universally stable policies in CAQT that use the injection time and the remaining path to schedule packets remain universally stable. These policies include, for instance, Shortest in System (SIS) and Longest in System (LIS). Then, we study the case in which clock differences can vary over time, but the maximum difference is bounded. In this model we show the universal stability of two families of policies related to SIS and LIS respectively (the priority of a packet in these policies depends on the arrival time and a function of the path traversed). The bounds we obtain in this case depend on the maximum difference between clocks.

A. Fernández Anta
LaDyR, GSyC, Universidad Rey Juan Carlos, Móstoles, Madrid, Spain

J.L. López-Presa
EUITT, Universidad Politécnica de Madrid, Madrid, Spain

M.A. Lorenzo · P. Manzano · A. Mozo
EUI, Universidad Politécnica de Madrid, Madrid, Spain

J. Martínez-Romo
ETSII, Universidad Nacional de Educación a Distancia, Madrid, Spain

C. Thraves (✉)
LaBRI, Université Bordeaux I, domaine Universitaire, 351 cours de la Libération, 33405 Talence, France
e-mail: cbthraves@gmail.com

This is a necessary requirement, since we also show that LIS is not universally stable in systems without bounded clock difference. We then present a new policy that we call Longest in Queues (LIQ), which gives priority to the packet that has been waiting the longest in edge queues. This policy is universally stable and, if clocks maintain constant differences, the bounds we prove do not depend on them. To finish, we provide with simulation results that compare the behavior of some of these policies in a network with stochastic injection of packets.

Keywords Scheduling · Continuous adversarial queuing theory · Adversarial models · Clock skew · Clock drift · Clock synchronization

1 Introduction

Stability is a requirement in a packet switched network in order to be able to provide some quality of service. Stability means that the amount of traffic that is being routed in the network is always bounded. It is well known that an appropriate scheduling of packets is fundamental in order to guarantee stability [3]. A fundamental question is to identify scheduling policies that are able to guarantee stability, and among these, policies that guarantee good quality of service (e.g., latency).

The study of the capability of scheduling policies to guarantee stability under worst-case situations has been done with adversarial models [3, 6–8, 10]. In these models, the arrival of packets to the network is controlled by an adversary which defines, for each packet, the instant and node in the network where it is injected, and very often, its path in the network. To avoid the overload of links in the network, the number of packets that the adversary can inject is bounded. As we said, the main objective of these models is to explore the ability of scheduling policies at the routers to maintain the system stable or provide good quality of service even in the worst conditions.

1.1 Adversarial Models

In this general framework two main adversarial models have been defined. In the Permanent Session Model [7, 8, 10], also known as the (σ, ρ) -regulated injection model, all the traffic in the network is grouped in sessions, whose route and maximum packet injection rate are defined by the adversary. The adversary is restricted on the rates σ_i that it can assign to sessions in the sense that the total rate of the sessions that cross a link does not saturate the link. Additionally, the adversary is in control of the arrival of packets. The adversary always tries to create instability or to increase the maximum latency experienced by packets. There have been many works exploring this model, and proposing stable policies with different guarantees of quality of service (e.g., [7, 8, 10–12]). It is also interesting to observe that, in this model, FIFO (or FCFS), which is the most popular scheduling policy by far, can be made unstable at any constant network load [1].

The Temporary Session Model, commonly known as the Adversarial Queuing Theory (AQT) model [3, 6], relaxes the restriction that packets are assigned to sessions. This relaxation is similar to allowing the adversary to dynamically change the

sessions over time. In this model, each packet is injected with its own path and the only restriction on the adversary is that it cannot overload any link, in an amortized sense. The AQT model assumes that the network evolves in steps, and in each step at most one packet crosses each link. The Continuous AQT (CAQT) model, recently presented in [5], is an extension of the AQT model in which packets can have arbitrary sizes, and links have different bandwidths and propagation delays. This model is closer to reality than AQT, and due to the fact that AQT is a particular case of CAQT, the instability results obtained for AQT remain valid for CAQT. Additionally, it was shown in [5] that many positive results under the AQT model also hold under the CAQT model. The system model we consider in this paper is strongly based on the CAQT model.

Time-Based Scheduling Policies From all the policies that have been shown to be stable in all networks (which we call universally stable) under the AQT and CAQT models, those that seem to provide the lowest end-to-end packet delays are based on timing. In fact, it has been shown by Weinard [14] that for any policy in the family Without-Time-Stamping strategies, there are n -node under-loaded networks in which the delays and queue sizes are $2^{\Omega(\sqrt{n})}$. A policy of this family is assumed to know the network topology, and it assigns the priority of a packet as a function of its path and the number of edges it already crossed. This implies that, in general, it is convenient to use some timing information for scheduling. Unfortunately, the simplest time-based policies, can also suffer of large delays. For instance, for Shortest in System (SIS), the policy that gives the highest priority to the newest packet in the network, there are networks in which the delays and queue sizes are $2^{\Omega(\sqrt{n})}$ [3]. Similarly, for Longest in System (LIS), the policy that gives the highest priority to the oldest packet in the network, there are networks with diameter d in which the delays and queue sizes are $2^{\Omega(d)}$ [2].

The good news are that time-based policies can in fact provide low delay guarantees. For instance, in [3] it is presented a randomized scheduling algorithm that guarantees delays polynomial on the network parameters. This algorithm basically uses a longest-in-system strategy with random permutations. The deterministic scheduling algorithm with polynomial delays presented in [4] uses a similar approach. Additionally, there are simulation studies [13] which show that LIS may in fact behave much better in practice than one may expect from the lower bounds mentioned above.

Clocks, Clock Drifts, and Clock Skews The above mentioned results for time-based scheduling policies are obtained in network models in which it is implicitly assumed that each node in the network has a local clock to provide the time, and that all these clocks are synchronized and provide the same time. However, this latter assumption is not realistic in practice, since the oscillation frequency of each computer timer is different, what produces different *clock drifts*. The consequence of these drifts is that it is not unusual that different clocks provide different times. The differences between the clock times is what we call *clock skews*. In practice, in order to limit the effect of clock drifts, and to bound the clock skews, there are mechanisms, like the Network Time Protocol (NTP), that allow the resynchronization of clocks.

The Non-synchronized CAQT Model In this paper we propose a model of adversary in which clocks do not need to be synchronized. We call the new model *Non-synchronized CAQT* (NSCAQT), since it is basically the CAQT model with this additional generalization. Under this model we will study the behavior of different queue scheduling policies which depend on time and can be affected both by clock skews and clock drifts. To study these policies, we need to make assumptions on how they use the local clocks. For instance, for LIS and SIS we will assume that packets are assigned the local clock value of their injection node at their injection time, value that they carry with them and is used for scheduling.

In this paper we will study three main variations of the NSCAQT model. In the first one we assume that the system has clock skews but no clock has drifts. Hence, in this model clocks do not have the same time, but their time differences remain constant. We call this the *NSCAQT model with constant skews*. The second model we will study is a model in which skews can vary over time, but there is a bound on the maximum difference between the time of the clocks. We call this model the *NSCAQT model with bounded skews*, and is very suited to model a network in which a protocol like NTP is used to periodically resynchronize all the clocks. The third model we explore has constant drifts and no resynchronization mechanism guarantees that the skews are bounded. We call this model the *NSCAQT model with constant drifts*.

1.2 Contributions

The main contribution of this work is the detailed definition of a model in which clocks need not be synchronized. We have not found a model that considers this possibility in any previous adversarial model. Then, the first result we provide is on the NSCAQT model with constant skews. We study under this model scheduling policies that assign priorities to packets based on their injection times and their remaining paths. For these policies we will show how the NSCAQT system can be transformed into a CAQT system by changing the topology of the network and the adversary. As a consequence, we conclude that any such policy that is universally stable under CAQT is also universally stable under NSCAQT with constant skews.

We then explore universal stability under the NSCAQT model with bounded skews. We first define two families of policies that include SIS and LIS, respectively, as special cases. Then we show that all the policies in these families are universally stable in this model. We call these families *Shortest in System considering Path* (SISP) and *Longest in System considering Path* (LISP) respectively. The policies from these families assign packet priorities depending on both the injection time and a function of the packet path and the number of edges it already traversed.

Unfortunately, for the universally stable policies that we identified with the previously mentioned results, all the upper bounds on delays and queue sizes that we could prove depend on the clock skews. In fact, in several cases it can be easily shown that these parameters become larger as the skews grow. Then, the question is whether there are policies whose performance does not depend on the clock skews. We introduce a new policy, *Longest in Queues* (LIQ), which gives priority to the packet that has been waiting in queues the longest. We show that this policy is universally stable in the NSCAQT model with bounded skews. More interestingly, we show that

in the NSCAQT model with constant skews this policy has an upper bound on the end-to-end delay that does not depend on the skews and is close to that of LIS in CAQT.

Finally, we present some simulations which try to shed some light on the behavior of LIS, SIS, and LIQ in a network with stochastic arrival patterns, instead of adversarial, in the NSCAQT model with constant skews. The results show that, as expected by analysis, LIQ is not affected by the clock skews, and presents the best performance from among the three policies.

1.3 Structure

The structure of the rest of the paper is the following. In Sect. 2 we define the NSCAQT model in detail and introduce some notation to be used on the paper. In Sect. 3 we study stability under the NSCAQT model with constant skews. In Sect. 4 we study the policies in SISP and LISP under the NSCAQT model with bounded skews. In Sect. 5 we show the instability of LIS in the NSCAQT model with constant drift. In Sect. 6 we explore the performance of LIQ under the NSCAQT models. Finally, in Sect. 7 we present the simulations that have been done.

2 System Model

Like most previous adversarial network models, the NSCAQT system model has three major elements: an underlying network \mathcal{G} , a scheduling policy used \mathcal{P} , and an adversary \mathcal{A} . With these elements, the evolution of the system can be seen as a game between the adversary, which injects packets in the network trying to create instability, and the scheduling policy, that decides which packets move along their paths in the network, trying to prevent instability. The model of system considered in this paper is a direct extension of that presented in [5].

The Network In this model a network is modeled by a directed graph \mathcal{G} , formed by a set of nodes $V(\mathcal{G})$, representing the hosts and routers, and a set of edges $E(\mathcal{G})$, representing links between the nodes.¹ Each link e of the network has associated a positive finite bandwidth B_e , which determines the transmission speed of the link, and a finite propagation delay $P_e \geq 0$. We use a specific notation for the largest propagation delay and smallest bandwidth as follows: $P_{\max} = \max_{e \in E(\mathcal{G})} \{P_e\}$ and $B_{\min} = \min_{e \in E(\mathcal{G})} \{B_e\}$. (Since \mathcal{G} is finite, these values are well defined.)

The Bounded Adversary In a system with network \mathcal{G} , the adversary \mathcal{A} defines the traffic pattern, continuously deciding which packets are injected. Additionally, for each packet p , the adversary chooses the moment of injection $T_0(p)$, the source node $v_0(p)$, the destination node $v_{d_p}(p)$, and the path the packet has to traverse $\Pi(p) = (e_0(p), e_1(p), \dots, e_{d_p-1}(p))$. (When clear from the context, we may omit the packet p from the notation.) Notice that d_p represents the length of the path packet p has to

¹Throughout the paper we use interchangeably the terms link and edge.




Fig. 1 Path $\Pi(p)$ assigned to a packet p in the network

traverse. We assume that a packet path is edge-simple, i.e. it does not contain the same edge more than once, although it can visit the same node several times.² We denote by d_{\max} the length of the longest path of a packet, which is clearly open bounded by the length of the longest edge-simple path in the network. In Fig. 1 we represent the path assigned to a packet p .

Although the adversary controls the traffic arrival, it is restricted on the load that it can inject to the system. We assume that the injection of a packet is instantaneous. Then, if $N_e(I)$ represents the number of bits of the packets which want to cross edge e injected by \mathcal{A} during an interval I , it must satisfy that

$$N_e(I) \leq r|I|B_e + b = (1 - \varepsilon)|I|B_e + b \quad (1)$$

for every edge e and for every time interval I . We denote by r , $0 < r \leq 1$, the long term rate (load) the adversary can impose on the system. For convenience we sometimes use the notation $r = 1 - \varepsilon$, for $\varepsilon \geq 0$. The parameter b , $b \geq 1$, is the burstiness allowed to the adversary injections, which is the excess of bits allowed to arrive at any time during the complete game. An adversary that satisfies this condition is called an (r, b) -adversary.

Packets, Queues, and Buffers Packets are sequences of bits of possibly different sizes. We denote by L_p the size in bits of a packet p and by L_{\max} the maximum size of a packet. Because of the above restriction (1) on the adversary and the assumption of instantaneous injection of packets, it can be easily observed that $L_{\max} \leq b$. Note that b is also an upper bound on the number of packets that the adversary can inject instantaneously (which is achieved in the improbable case that all packets have size 1).

Packets in the system follow their path traversing one edge after the other toward their destination. As explained in [5], in every node in the network there is a reception buffer for each edge entering the node and an output queue for each edge leaving the node. The output queue of an edge has unbounded capacity and holds the packets that are ready to cross this edge. The scheduling policy of the edge's output queue chooses the next packet to cross the edge from those in this output queue. The reception buffer is used to store the received portion of a packet until it has been completely received. Then, the packet is placed instantaneously by a dispatcher in the corresponding output queue or it disappears from the system if it already reached its destination.

²This assumption does not decrease the generality of the model in terms of universal stability of policies, since it is known that a system in which packets may traverse the same edge several times can be simulated by another system with only edge-simple paths [1].

Note that once a packet p starts to cross edge e , it will spend $\frac{L_p}{B_e} + P_e$ units of time to completely cross it. As parameter of the network we have the greatest amount of time that a packet can spend crossing an edge, denoted D_{\max} , and defined as

$$D_{\max} = \max_{e \in E(\mathcal{G})} \left\{ \frac{L_{\max}}{B_e} + P_e \right\} \leq \frac{b}{B_{\min}} + P_{\max}.$$

Clocks As we said, the main difference between the NSCAQT model we propose and previous models [3, 5, 6, 10] is that we consider here the impact of clocks not being synchronized on the performance. In order to make the model as general as possible, we assume that the output queue of each edge has its own internal clock, called *edge clock* (this is clearly more general than assuming one clock per node). Additionally, we assume there is an external reference clock which is always on time. We refer to this clock as the *real clock* and we say that it provides the *real time*. We assume that the adversary has access to both the edge clocks and the real clock, while the scheduling policy at a given edge has only access to the clock of that edge.

The difference between the real clock and the edge clock of e at real time t is what we call the *clock skew* of e 's edge clock at time t , and is denoted by $\phi_e(t)$. Then, if t_e denotes the value of the edge clock of e at real time t , we have $t_e = t - \phi_e(t)$. If this value changes over time, we say that the edge clock has a *drift*. If an edge clock has no drift we omit the time and denote its skew by ϕ_e . Note that, at any given time, the skew of an edge clock can be positive or negative. However, for convenience we assume that these skews are all non-negative if all edge-clock skews are lower bounded. We can do this freely since the real clock is not available to the scheduling policies and does not interfere in the relation between edge clocks.

We denote by $T_i(p)$, $0 \leq i < d_p$, the real time at which a packet p arrives to the output queue of the edge $e_i(p)$. Due to clock skews, according to edge $e_i(p)$'s clock, the instant when packet p arrives to the output queue is $T_i(p) - \phi_{e_i(p)}(T_i(p))$. Additionally, we denote by $T_{d_p}(p)$ the time at which p is completely received at its destination and leaves the system.

Scheduling Policies As we said above, the scheduling policy is in charge of deciding, whenever a link e is available, which packet from those in the output queue of e must be sent next across e . In this paper we only consider distributed work-conserving time-based scheduling policies. We say that policies are distributed if they do not use the state (and in particular the clock) of other edges to make scheduling decisions. Policies are work-conserving (also called greedy) if they always send a packet across the link as long as the edge's output queue is not empty. Finally, we only consider time-based policies, which are policies that use the edge clocks for scheduling. Note that policies that are not time-based are not affected by clock skews and drifts.

We will only consider in this work systems in which all the queues use the same scheduling policy. The study of systems under the NSCAQT model in which different queues may use different scheduling policies is left for future work.

Two of the most studied distributed work-conserving time-based scheduling policies are Longest in System (LIS) and Shortest in System (SIS). The LIS policy gives the highest priority to the packet that has been in the system for the longest time,

while the SIS policy gives the highest priority to the packet that has been in the system for the shortest time. These definitions do not clearly show the use these policies make of the edge clocks. For that, we need to look at the natural implementation of these policies: upon arrival of a packet p into the system, it is assigned a time-stamp, $TS(p)$, which p carries with it. Then, the LIS and SIS policies only compare the time-stamps of the packets to decide which to schedule next. The edge scheduler in LIS gives the highest priority to the packet with the smallest time-stamp, while in SIS it gives the highest priority to the packet with the largest time-stamp. Note that when clocks are not synchronized, these time-stamps are not accurate, since the time-stamp for a packet p is

$$TS(p) = T_0(p) - \phi_{e_0(p)}(T_0(p)).$$

These two policies have been proved to be universally stable in [5] for the CAQT model, where $\phi_e = 0$ for all $e \in E(\mathcal{G})$.

In addition to these two well-known policies, we will study two families of policies derived from SIS and LIS, that we call Shortest in System considering Path (SISP) and Longest in System considering Path (LISP), respectively. In the policies of these families, packets carry their time-stamp, and (when needed) the packet path and the number of links traversed, so that at its edge $e_i(p)$, packet p is assigned a priority label of the form

$$PL(p, i) = TS(p) + f(\Pi(p), i),$$

where $f(\Pi(p), i)$ is a function which assigns a real number to each pair formed by a packet path $\Pi(p)$ and a value $i \in \{0, 1, \dots, d_p - 1\}$, being i the number of crossed edges by the packet in its path. A policy \mathcal{P}_f is in SISP (resp. LISP) if at each queue it gives the highest priority to the packet p with the largest (resp. smallest) value $PL(p, i)$. Notice that when $f(\Pi(p), i) = 0$ for all $\Pi(p)$ and i , \mathcal{P}_f is equivalent to SIS (resp. LIS). Since the function f is defined over the finite sets of paths and number of edges crossed by a packet, it has a maximum and a minimum values, that we denote by f_{\max} and f_{\min} , respectively.

Finally, we will consider a new policy named Longest in Queues (LIQ). In this policy the highest priority is assigned to the packet that has been waiting the longest in all the output queues it has visited. In our model NSCAQT, in which clocks are not synchronized, we assume that the time in queues is measured locally at each output queue. The time a packet p waits at the edge e 's queue is the difference between the value of the edge clock when the packet arrives and the value when it starts being transmitted, or the current value of the edge clock if it is still waiting. The time used to schedule p is the sum of these waiting times in all the visited queues.

2.1 System Stability

To study stability and performance in packet switching networks, we introduce the concept of a $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ system to represent the game played between an adversary \mathcal{A} and the packet scheduling policy \mathcal{P} over the network \mathcal{G} . In the NSCAQT model, a system $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is stable if the maximum number of packets (or bits) present in the system is bounded at any time by a constant that may depend on system parameters:

the network, the adversary or the policy. A policy \mathcal{P} is *universally stable* if the system $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is stable on every network \mathcal{G} and against every (r, b) -adversary \mathcal{A} with $r < 1$ and $b \geq 1$.

3 Stability of Policies for Constant Clock Skews

In this section we study the case in which all the edge clocks have zero drift, so that ϕ_e is constant. This framework allows to assure the stability in a non-synchronized system if there is stability in a synchronized system for many policies, in particular for those policies that only depend on the injection time and the remaining path of the packets.

We present a proof by transformation of this case. We start from a $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ system with non-synchronized clocks, where \mathcal{A} is an (r, b) -adversary, $r \leq 1$. Then, we vary the network \mathcal{G} and the adversary \mathcal{A} to obtain a synchronized system $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$, where \mathcal{A}' is an (r, b') -adversary, so that if $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is stable, then $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is also stable.

Construction of \mathcal{G}' As we explained in the previous section, since we arbitrarily fix the real clock, we can do it so that all the clock skews are non-negative. Then, let $\mathcal{G} = (V, E)$ be a directed graph in the NSCAQT model with constant skew $\phi_e \geq 0$ for each edge clock. We construct \mathcal{G}' starting from \mathcal{G} as shown in Fig. 2.

For each edge $e \in E(\mathcal{G})$ with $\phi_e > 0$, let $K_e = \lceil b + rB_e\phi_e/2 \rceil$. Then, we add $L_{\max} \times K_e$ edges and $L_{\max} \times K_e$ nodes. New edges and nodes are denoted $e^{l,j}$ and $v^{l,j}$, respectively, where $l \in \{1, 2, \dots, L_{\max}\}$ and $j \in \{1, 2, \dots, K_e\}$. For all l and for all j , we place edge $e^{l,j}$ from node $v^{l,j}$ to the tail of edge e . For every edge $e^{l,j}$, we set the bandwidth to $B_{e^{l,j}} = \frac{2l}{\phi_e}$ and the propagation delay to $P_{e^{l,j}} = \frac{\phi_e}{2}$.

Remark 1 With this construction process, a packet p of size L_p takes $\frac{\phi_e}{2}$ units of time to be fully sent and ϕ_e units of time to be fully received across any edge $e^{L_p,j}$, $j \in \{1, 2, \dots, K_e\}$.




Fig. 2 Basic process to transform \mathcal{G} into \mathcal{G}'

Construction of \mathcal{A}' We now construct the adversary \mathcal{A}' from \mathcal{A} . A packet that is injected by \mathcal{A} at edge e of \mathcal{G} with $\phi_e = 0$ is injected exactly in the same edge at the same time in \mathcal{G}' by \mathcal{A}' . Now, let p be a packet of size L_p that was injected in \mathcal{G} by \mathcal{A} at time t , such that $\phi_{e_0(p)} > 0$. Then, \mathcal{A}' will inject a packet p' in \mathcal{G}' at time $t - \phi_{e_0(p)}$. The size of p' will be $L_{p'} = L_p$ and its path will be

$$\Pi(p') = (e_0^{L_{p,j}}(p), e_0(p), e_1(p), \dots, e_{d_p}(p)),$$

where $e_0^{L_{p,j}}(p)$ is an edge corresponding to the construction described above. This edge must satisfy that no other packet has been injected in it in the previous $\phi_e/2$ time. Since K_e is clearly an upper bound on the number of packets of length L_p that can be injected in any interval of $\phi_e/2$ time, and there are K_e edges for each L_p , there is always a suitable edge to be used.

Under these circumstances, if packet is injected by \mathcal{A} in time t in the queue of edge e , and labeled with an time-stamp of $t - \phi_e$, a corresponding packet injected by \mathcal{A}' will arrive to the same queue at the same time t , labeled with the same time-stamp. Let us now bound the parameters of the adversary \mathcal{A}' .

Remark 2 The packets injected by \mathcal{A} during an interval of size $|I| + \phi_{\max}$ are injected by \mathcal{A}' during an interval I' with maximum size $|I|$. So, since \mathcal{A} is an (r, b) -adversary, \mathcal{A}' is an (r, b') -adversary, where $b' = b + \phi_{\max} B_{\max}$, and $\phi_{\max} = \max_{e \in E(\mathcal{G})} \{\phi_e\}$.

We can now state the main result of this section.

Theorem 1 Let $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ be the synchronized system in the CAQT model obtained from the non-synchronized system $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ through the above process. Let \mathcal{P} be a scheduling policy which considers only the time of injection of the packets and the paths that the packets still have to traverse. Then, $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is stable if and only if $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is stable.

Proof Note that the stability of $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ must occur for any value of the clock skews, including the case in which all skews are zero. Then, it trivially follows that if $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is stable, $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is stable as well.

In the other direction, we first have that the queues of the new edges $e^{l,j}$ in system $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ never present contention since, from Remark 1 and by the construction of the adversary \mathcal{A}' , by the time a packet arrives the previous packet (if any) was already sent. Then, if we observe the output queues of the edges that \mathcal{G}' and \mathcal{G} have in common, we find that similar sets of packets arrive at the same times, with the same time-stamps, and with the same remaining paths to cross in both systems $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ and $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$. Since these are the parameters used by \mathcal{P} to schedule the packets, we have that the behavior of these queues in systems $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ and $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is exactly the same. So, if there is no bound on the number of packets in $(\mathcal{G}, \mathcal{P}, \mathcal{A})$, there is no bound either on the number of packets in $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$. Then we have that if $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is unstable then $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is unstable or, equivalently, if $(\mathcal{G}', \mathcal{P}, \mathcal{A}')$ is stable then $(\mathcal{G}, \mathcal{P}, \mathcal{A})$ is stable. \square

Corollary 1 *The scheduling policies that are universally stable in CAQT and only consider the times of injection and the paths that the packets still have to traverse are universally stable in the NSCAQT model with constant clock skews.*

4 Stability of Policies for Bounded Clock Skews

In this section we will study the case in which clocks may experience drifts. Hence, we assume here that the clock skews are not necessarily constant. However, the maximum difference between real time and any edge clock is bounded. As we said, this model fits naturally with a system in which edge clocks are periodically resynchronized, for instance via NTP.

In this section we will again adapt the real time reference clock, in order to simplify the analysis and the presentation. Like in the previous section, we will assume that all clock skews are non-negative, i.e., for any edge e and any time t , $\phi_e(t) \geq 0$. Additionally, since we assume that skews are bounded, we can safely define $\phi_{\max} = \max_{e,t} \{\phi_e(t)\}$.

Under these assumptions, we show that all policies in the families SISP and LISP are universally stable. Since the model NSCAQT with bounded skews is more general, this also shows that these policies are universally stable in the CAQT model (which was previously unknown in general).

4.1 Universal Stability of SISP

In this subsection we explore the stability of the new family of policies SISP defined in Sect. 2, which is based on the injection time, and the path and number of edges already crossed by a packet. As described there, a policy \mathcal{P}_f in SISP assigns to each packet p at the queue of its edge $e_i(p)$ a priority label $PL(p, i) = TS(p) + f(\Pi(p), i)$, and gives the highest priority to the packet with the largest label.

Now we prove that every policy in SISP is universally stable in the NSCAQT model when the clock skew is bounded by ϕ_{\max} . We start with the following simple lemma.

Lemma 1 *Let p and q be two packets. If $T_0(q) > T_0(p) + f_{\max} - f_{\min} + \phi_{\max}$, then p never has higher priority than q in any queue.*

Proof Let us assume that p and q meet at the output queue of edge $e_i(p) = e_j(q)$. Note that $\phi_{\max} - \phi_{e_0(q)} \geq 0$ and that $f(\Pi(q), j) - f_{\min} \geq 0$. Hence,

$$\begin{aligned} PL(q, j) &= T_0(q) - \phi_{e_0(q)} + f(\Pi(q), j) \\ &> T_0(p) + f_{\max} - f_{\min} + \phi_{\max} - \phi_{e_0(q)} + f(\Pi(q), j) \\ &\geq T_0(p) + f_{\max} \geq PL(p, i), \end{aligned}$$

and then $PL(q, j) > PL(p, i)$ for all i and j . \square

The proof of universal stability of the policies in SISP we present is very similar to that of SIS presented in [5] for the CAQT model. We first recall a lemma proved

there, which limits the time spent by a packet in the queue of an edge e if there are $k - 1$ bits in the system with higher priority to cross that edge. The assumptions and the proof of this lemma do not depend on whether the clocks are synchronized. Hence, it can be applied directly to our model.

Lemma 2 [5] *Let p be a packet that, at real time t , is waiting in the queue of edge e . At that instant, let $k - 1$ be the total size in bits of the packets in the system that also want to cross e and that may have priority over p . Then, p will start crossing e in at most $(k + b)/(\varepsilon B_e)$ units of time.*

Recall that, when a packet p starts crossing an edge e , it spends $P_e + L_p/B_e \leq D_{\max}$ units of time until it crosses it completely. Then, using this and the previous lemma recursively we can prove the following result.

Lemma 3 *Let $k_0 = r(\phi_{\max} + f_{\max} - f_{\min})B_{\max} + b$, and $k_i = k_{i-1} + r(\frac{k_{i-1} + b}{\varepsilon B_{\min}} + D_{\max})B_{\max} + b$, for $0 < i < d_{\max}$. When a packet p arrives to the output queue of edge $e_i(p)$, no more than $k_i - 1$ bits can have priority over it in any edge $e_j(p)$, for $j \geq i$.*

Proof Let us first consider the case $i = 0$. When a packet p , arrives into the system, from Lemma 1, the packets that may have priority over it have been injected at most $\phi_{\max} + f_{\max} - f_{\min}$ units of time earlier than p , because, although they maybe arrived to the system before p , they have a greater time-stamp due to their initial edge's clock skew or the value of the function $f(\cdot)$. The total size of these packets is at most $r(\phi_{\max} + f_{\max} - f_{\min})B_{\max} + b - 1 = k_0 - 1$ bits (since $L_p \geq 1$).

Let us now assume as induction hypothesis that the claim holds for $0 \leq i < d_p - 1$. Then, from Lemma 2, p will arrive at the output queue of edge $e_{i+1}(p)$ at most $(k_i + b)/(\varepsilon B_{\min}) + D_{\max}$ time units after arriving at the output queue of edge $e_i(p)$. During this time at most packets with $r((k_i + b)/(\varepsilon B_{\min}) + D_{\max})B_{\max} + b$ bits are injected that can block p at any edge. Hence packets with at most $(k_i - 1) + r((k_i + b)/(\varepsilon B_{\min}) + D_{\max})B_{\max} + b = k_{i+1} - 1$ bits can block p in any edge $e_j(p)$, $j \geq i + 1$. \square

Using these definitions of k_i and the previous lemma, we can limit the size of the queues and the amount of time that a packet spends in the network as it was done in [5]. The proof of the theorem is verbatim to the final part of the corresponding theorem in [5] and is hence omitted.

Theorem 2 *Let \mathcal{G} be a network and d_{\max} the length of its longest edge-simple directed path, let \mathcal{A} be an (r, b) -adversary with $r = 1 - \varepsilon < 1$ and $b \geq 1$, and let \mathcal{P}_f be a policy in SISF. Then the system $(\mathcal{G}, \mathcal{P}_f, \mathcal{A})$ is stable under the NSCAQT model with bounded clock skews, no queue ever contains $k_{d_{\max}-1} + L_{\max}$ bits, and no packet spends more than*

$$\frac{d_{\max}b + \sum_{i=0}^{d_{\max}-1} k_i}{\varepsilon B_{\min}} + d_{\max}D_{\max}$$

units of time in the system.

Hence the main result of the section.

Corollary 2 *Any policy in SIS, and in particular SIS, is universally stable under the NSCAQT model with bounded clock skew, and hence under the CAQT model.*

4.2 Universal Stability of LISP

As defined in Sect. 2, LISP is a family of policies based on the injection time, and the path and the number of edges already crossed by a packet. A policy \mathcal{P}_f in LISP assigns to each packet p at the queue of its edge $e_i(p)$ a priority label $PL(p, i) = TS(p) + f(\Pi(p), i)$, and gives the highest priority to the packet with the smallest label.

Now we prove that every policy in LISP is universally stable in the NSCAQT model when the clock skew is bounded by ϕ_{\max} . We start with the following simple lemma, which is similar to Lemma 1 in the previous section.

Lemma 4 *Let p and q be two packets. If $T_0(q) > T_0(p) + f_{\max} - f_{\min} + \phi_{\max}$, then q never has higher priority than p in any queue.*

Proof Let us assume that p and q meet at the output queue of edge $e_i(p) = e_j(q)$. Note that $\phi_{\max} - \phi_{e_0(q)} \geq 0$ and that $f(\Pi(q), j) - f_{\min} \geq 0$. Hence,

$$\begin{aligned} PL(q, j) &= T_0(q) - \phi_{e_0(q)} + f(\Pi(q), j) \\ &> T_0(p) + f_{\max} - f_{\min} + \phi_{\max} - \phi_{e_0(q)} + f(\Pi(q), j) \\ &\geq T_0(p) + f_{\max} \\ &\geq PL(p, i), \end{aligned}$$

and then $PL(q, j) > PL(p, i)$ for all i and j . \square

Let p be a packet in the system and let t denote some real time in $[T_0(p), T_{d_p}(p)]$. We denote by $g(t)$ the real injection time of the oldest packet in the system at time t . We define $C_p = \max_{t \in [T_0(p), T_{d_p}(p)]} \{t - g(t)\}$. Notice that C_p represents the age of the oldest packet in the system while p is present. For convenience we use the abbreviation $K = f_{\max} - f_{\min} + \phi_{\max}$.

In the following lemma we start by bounding the amount of time, $T_{i+1}(p) - T_i(p)$, that p takes to move from the queue of e_i to the queue of e_{i+1} , which allows us to bound the time p is in the system.

Lemma 5 *The time packet p is in the system is at most*

$$T_{d_p}(p) - T_0(p) \leq (1 - \varepsilon^d) \left(K + C_p + \frac{D_{\max}}{1 - \varepsilon} \right).$$

Proof Observe that the oldest packet in the system when p arrives at the queue of e_i at time $T_i(p)$ was injected at most at time $T_i(p) - C_p$. Then, from Lemma 4 we

have that p , and all the packets with higher priority than p in the queue of e_i , were injected during the interval $[T_i(p) - C_p, T_0(p) + K]$. Since the packets injected in this interval can have at most

$$(1 - \varepsilon)(T_0(p) + K - T_i(p) + C_p)B_{e_i} + b$$

bits, we know that all of them cross e_i in at most

$$\begin{aligned} & (1 - \varepsilon)(T_0(p) + K - T_i(p) + C_p) + \frac{b}{B_{e_i}} + P_{e_i} \\ & \leq (1 - \varepsilon)(T_0(p) + K - T_i(p) + C_p) + D_{\max} \end{aligned}$$

units of time. Then, we have that

$$\begin{aligned} T_{i+1}(p) & \leq T_i(p) + (1 - \varepsilon)(T_0(p) + K - T_i(p) + C_p) + D_{\max} \\ & = \varepsilon T_i(p) + (1 - \varepsilon)(T_0(p) + K + C_p) + D_{\max}, \end{aligned}$$

and solving the recurrence for $T_{d_p}(p)$ we get

$$T_{d_p}(p) - T_0(p) \leq (1 - \varepsilon^d) (K + C_p + D_{\max}/(1 - \varepsilon)). \quad \square$$

Now, we have bounded the time that a packet spends in the system, but our bound depends on C_p . To finish the proof we need the following lemma.

Lemma 6 *For any packet p we have that*

$$C_p \leq \frac{1 - \varepsilon^{d_{\max}}}{\varepsilon^{d_{\max}}} \left(K + \frac{D_{\max}}{1 - \varepsilon} \right) = C.$$

Proof By contradiction, let us assume there are packets that spend in the system more than C time. Let t be the first time at which a packet has been in the system more than C time, and let p be one packet that satisfies this. Then $T_{d_p}(p) - T_0(p) \geq t - T_0(p) > C$. Note that before t no packet in the system has age older than C . Hence, by Lemma 5 with $C_p = C$ we have that

$$\begin{aligned} T_{d_p}(p) - T_0(p) & \leq (1 - \varepsilon^{d_p}) \left(K + C + \frac{D_{\max}}{1 - \varepsilon} \right) \\ & \leq (1 - \varepsilon^{d_{\max}}) \left(K + C + \frac{D_{\max}}{1 - \varepsilon} \right) \\ & = C - \varepsilon^{d_{\max}} C + (1 - \varepsilon^{d_{\max}}) \left(K + \frac{D_{\max}}{1 - \varepsilon} \right) \\ & = C, \end{aligned}$$

which is a contradiction. \square

Now we can enunciate the final theorem of this section. The proof is by Lemmas 5 and 6, and the definition of K .

Theorem 3 Let \mathcal{G} be a network and d_{\max} the length of its longest edge-simple directed path, let \mathcal{A} be an (r, b) -adversary, with $r = 1 - \varepsilon < 1$ and $b \geq 1$, and let \mathcal{P}_f be a policy in LISP. Then the system $(\mathcal{G}, \mathcal{P}_f, \mathcal{A})$ is stable under the NSCAQT model with bounded clock skew, and no packet spends more than

$$\frac{1 - \varepsilon^{d_{\max}}}{\varepsilon^{d_{\max}}} \left(f_{\max} - f_{\min} + \phi_{\max} + \frac{D_{\max}}{1 - \varepsilon} \right)$$

units of time in the system.

Corollary 3 Any policy in LISP, and in particular LIS, is universally stable under the NSCAQT model with bounded clock skew, and hence under the CAQT model.

5 Instability of LIS in the NSCAQT Model with Constant Drifts

In this section we show the instability of LIS in the NSCAQT model with constant drifts and unbounded skew. To do so, we present a specific network with (only) two links whose skew is not bounded and whose drift is constant. For this network we present an injection pattern (an adversary) such that the number of packets in the network grows forever if LIS is used.


The network \mathcal{G} used in this section is presented in Fig. 3, and is similar to the one shown in Fig. 2 of [3] (sometimes called the “baseball graph”). In this network we assume that no link has skew, except links f'_0 and f'_1 which have no skew at time $t = 0$ and have a constant drift of 1. Then, the skew at time t in these links is $\phi_{f'_0}(t) = \phi_{f'_1}(t) = t$. Hence, any packet p injected in these two links is assigned a time-stamp $TS(p) = 0$,³ and under LIS it will have priority over any other packet in the system injected after $t = 0$ in any other link. Additionally, we assume that all packets p have the same length, $L_p = l$ bits, and that all links e have the same bandwidth $B_e = l$ bits per unit of time, and null propagation delay. With these assumptions, any packet p takes exactly one unit of time to cross each link.

We now present the adversary \mathcal{A} that creates instability in the network \mathcal{G} under LIS. The adversary is described inductively in phases. Initially, at time $t_0 = 0$, \mathcal{A} injects a set S_0 of s_0 packets at node w_1 , of which $s_0/2$ have path f_1, e_0, f_0 and $s_0/2$ have path f'_1, e_0, f_0 . This creates the base case of induction. The induction hypothesis of phase i , for any $i \geq 0$, is the following. Let t_i the time at which phase i starts (which coincides with the time phase $i - 1$ ends). At time t_i , for i even (resp. odd), there is a set S_i of s_i packets at node w_1 (resp. w_0), of which $s_i/2$ have path f_1, e_0, f_0 (resp. f_0, e_1, f_1) and $s_i/2$ have path f'_1, e_0, f_0 (resp. f'_0, e_1, f_1). We will show that at time t_{i+1} there will be more than s_i packets at node w_0 (resp. w_1). We present here only the behavior of even-numbered phases, since odd-numbered phases are perfectly symmetrical.

Then, the sequence of injections and the system behavior during phase t_i , for i even, is as follows:

³Note that this means that the clock of these links has stopped. We consider this situation for simplicity, but is it very simple to modify the proof to use a clock that always advances.

Fig. 3 Network in which LIS is unstable with unbounded skew



1. At time t_i the packets of set S_i , initially at node w_1 , start to be transmitted. The first one will arrive at node v_0 at time $t_i + 1$, and all of the packets of set S_i will arrive at node v_0 in $s_i/2 + 1$ units of time. Then, they will be sent across link e_0 during the time interval $[t_i + 1, t_i + s_i + 1]$.⁴ During this time \mathcal{A} injects in node v_0 a set X of $\lfloor rs_i \rfloor$ packets, whose path is e_0, f'_0, e_1, f_1 . These packets are blocked by the packets in S_i , so that they will be transmitted across e_0 in the interval $[t_i + s_i + 1, t_i + s_i + \lfloor rs_i \rfloor + 1]$.

The packets in S_i are sent across link f_0 in the time interval $[t_i + 2, t_i + s_i + 2]$. During this time \mathcal{A} injects at node w_0 a set Y of rs_i packets which only want to traverse link f_0 . These packets are blocked by the packets in S_i , and will be transmitted across f_0 during the time interval $[t_i + s_i + 2, t_i + s_i + \lfloor rs_i \rfloor + 2]$.

2. During this interval $[t_i + s_i + 2, t_i + s_i + \lfloor rs_i \rfloor + 2]$, the packets in X reach w_0 to cross link f'_0 . In this interval \mathcal{A} injects at node w_0 a set Z of $\lfloor r \lfloor rs_i \rfloor \rfloor$ packets whose path is f'_0 . Due to the skew of link f'_0 and the LIS policy, the packets in Z have higher priority than those in X . Then, at time $t_i + s_i + \lfloor rs_i \rfloor + 2$ there still will be a subset X' of X with $\lfloor r \lfloor rs_i \rfloor \rfloor$ packets in the queue of f'_0 .

During the same time interval, \mathcal{A} also injects in w_0 a set Y' of $\lfloor r \lfloor rs_i \rfloor \rfloor$ packets whose path is f_0, e_1, f_1 . These packets remain in the queue of f_0 at the end of the interval, since they are blocked by those in Y .

Then, at time $t_{i+1} = t_i + s_i + \lfloor rs_i \rfloor + 2$ the following packets will be at node w_0 :

- In the queue of link f_0 , the $\lfloor r \lfloor rs_i \rfloor \rfloor$ packets in set Y' , whose path is f_0, e_1, f_1 .
- In the queue of link f'_0 , the $\lfloor r \lfloor rs_i \rfloor \rfloor$ packets in set X' , whose path is f'_0, e_0, f_1 .

These two sets form the set S_{i+1} that satisfies the induction hypothesis of phase $i + 1$. Now, for instability the size $s_{i+1} = 2 \lfloor r \lfloor rs_i \rfloor \rfloor$ of set S_{i+1} must be a larger than s_i . For this condition we can find a rate r and an initial value of s_0 such that queues get larger at each phase, making the system unstable.

⁴The last packet is sent at time $t_i + s_i + 1$ and arrives completely at w_0 at time $t_i + s_i + 2$.

Theorem 4 *The system $(\mathcal{G}, LIS, \mathcal{A})$ is unstable, where \mathcal{A} is the (r, b) -adversary described above for the network \mathcal{G} of Fig. 3, with $s_0 \geq 4$, $r \geq \sqrt{0.5 + 2/s_0}$, and $b \geq s_0 l$.*

Proof From the above, it is enough to show that $s_{i+1} = 2\lfloor r \lfloor r s_i \rfloor \rfloor > s_i$ for all $i \geq 0$. Since $\lfloor x \rfloor > x - 1$, we have that $\lfloor r \lfloor r s_i \rfloor \rfloor > r(r s_i - 1) - 1 > r^2 s_i - 2$. Hence, instability follows from $r^2 s_i - 2 \geq s_i/2$, which yields a lower bound of $r \geq \sqrt{0.5 + 2/s_i}$. It is enough that this holds for $i = 0$, since then $s_{i+1} > s_i$ and the bound only decreases. Hence, $s_0 \geq 4$ to allow $r \leq 1$. The bound on b follows trivially. \square

Corollary 4 *LIS is not universally stable in the NSCAQT model with constant drifts and unbounded skew.*

6 Universal Stability of LIQ

In previous sections we have shown how several policies are universally stable in the NSCAQT models with constant and bounded clock skews, respectively. Unfortunately, the bounds on end-to-end packet latencies we derived were dependent on the maximum clock skew that can occur in the system. This means that in a system with high maximum skew, the latencies can be very high. It is easy to construct examples for policies like SIS and LIS in which this can be observed.

In this section we study a new policy named LIQ, which gives the highest priority to the packet that has been waiting in output queues for the longest time. We prove that LIQ is universally stable in the NSCAQT model with bounded clock skews. The bad news is that in this case the end-to-end latency bound we obtain depends also on the maximum skew. The good news is that for the NSCAQT model with constant clock skews LIQ is universally stable, and the bound does not depend on the maximum skew, and it is similar to that obtained with LIS in a synchronized system.

As in previous sections, we assume that $\phi_e(t) \geq 0$ for all e and t . Then, we define $\phi_{\min}(e) = \min_t \{\phi_e(t)\}$ and $\phi_{\max}(e) = \max_t \{\phi_e(t)\}$. Finally, let $\Delta\phi = \max_e \{\phi_{\max}(e) - \phi_{\min}(e)\}$. Observe that in the model of constant skews, $\Delta\phi = 0$.

Let $W_{i,t}(p)$ be the amount of real time which packet p has waited in output queues when at time t it is at the queue of edge $e_i(p)$. We have that

$$t = T_0(p) + W_{i,t}(p) + \sum_{k=0}^{i-1} \left(\frac{L_p}{B_{e_k(p)}} + P_{e_k(p)} \right)$$

and, hence,

$$t - T_0(p) - d_{\max} D_{\max} \leq W_{i,t}(p) \leq t - T_0(p). \quad (2)$$

Recall that measuring the waiting time of a packet at a queue is done by taking the local time when the packet arrives and the local time when it leaves (or the current local time if it is still in the queue). Then the measured time at one queue can have an error of up to $\pm\Delta\phi$. Hence, the measured waiting time of a packet p that at time t is in queue $e_i(p)$, denoted $M_{i,t}(p)$, is a value in the interval $[W_{i,t}(p) - i\Delta\phi, W_{i,t}(p) + i\Delta\phi]$.

The proof we have is similar to the proof for the LISP family of policies. We first prove a lemma analogous to Lemma 4.

Lemma 7 *Let p and q be two packets. If $T_0(q) > T_0(p) + d_{\max}(2\Delta\phi + D_{\max})$, then q never has higher priority than p in any queue.*

Proof Let us assume that p and q meet at the output queue of edge $e_i(p) = e_j(q)$ at time t . We need to show that $M_{j,t}(q) < M_{i,t}(p)$ is satisfied for any i, j , and t . The largest value $M_{j,t}(q)$ can take is

$$M_{j,t}(q) \leq W_{j,t}(q) + j\Delta\phi \leq W_{j,t}(q) + d_{\max}\Delta\phi.$$

Similarly, $M_{i,t}(p) \geq W_{i,t}(p) - d_{\max}\Delta\phi$, and therefore we are left with the problem of showing that $W_{j,t}(q) < W_{i,t}(p) - 2d_{\max}\Delta\phi$. Then, by using (2) and the assumption of the lemma, we have

$$\begin{aligned} W_{j,t}(q) &\leq t - T_0(q) \\ &< t - T_0(p) - d_{\max}(2\Delta\phi + D_{\max}) \\ &\leq W_{i,t}(p) - 2d_{\max}\Delta\phi \end{aligned}$$

which completes the proof. \square

Now, defining $K = d_{\max}(2\Delta\phi + D_{\max})$, we have that Lemmas 5 and 6 are also valid in this case. Then, we can enunciate the following theorem.

Theorem 5 *Let \mathcal{G} be a network with d_{\max} the length of its longest edge-simple directed path, and let \mathcal{A} be an (r, b) -adversary with $r = 1 - \varepsilon < 1$. Then*

1. *The system $(\mathcal{G}, LIQ, \mathcal{A})$ is stable under the NSCAQT model with bounded clock skews, and no packet spends more than*

$$\frac{1 - \varepsilon^{d_{\max}}}{\varepsilon^{d_{\max}}} \left(d_{\max}(2\Delta\phi + D_{\max}) + \frac{D_{\max}}{1 - \varepsilon} \right)$$

units of time in the system.

2. *The system $(\mathcal{G}, LIQ, \mathcal{A})$ is stable under the NSCAQT model with constant clock skews, and no packet spends more than*

$$\frac{1 - \varepsilon^{d_{\max}}}{\varepsilon^{d_{\max}}} \left(d_{\max}D_{\max} + \frac{D_{\max}}{1 - \varepsilon} \right)$$

units of time in the system.

Corollary 5 *LIQ is universally stable under the NSCAQT model with bounded clock skews, and hence under the CAQT model.*

7 Simulations

In order to partially evaluate the theoretical results we have developed several simulation experiments. All the experiments in this article have been carried out using the J-Sim discrete event simulator [9]. J-Sim has been designed to simulate network behaviors in a realistic way, including propagation delays, packet processing times, etc. The J-Sim package has been modified in several ways, mainly to adapt it to our model. First, the traffic generator has been modified in order to ensure that destinations are uniformly distributed over all nodes in the network. Then, the sink monitor has also been changed in order to log several parameters that are not stored by J-Sim by default (e.g., the mean and the variance of the packet delay and the queue size, and samples of these values chosen at random). Also, the routing algorithm has been replaced and some scheduling policies discussed in this paper have been implemented.

The network topology used in all the experiments is an 11×11 torus, in which every node is, at the same time, router, source, and sink of packets. Each node periodically generates new packets, whose destination is chosen randomly and uniformly among the nodes of the network. The routing is deterministic, so that the traffic is balanced among all the links. We have adjusted the average load of the network to 99% in the simulations because we are interested on the response of the network with high load levels. In our experiments all the queues are of unbounded size, links have no propagation delays, the link bandwidth is set to 100 Kbps and the packet size is 105 bytes. The simulation experiments have been run for 6000 seconds, ignoring the first 1000 seconds in the analysis of the results.

We assign to local node clocks (all output edge clocks in a node are the same) different skews following a normal distribution with a mean value of 0 seconds and a standard deviation of up to 10^5 milliseconds. Before starting the experiment, each node randomly chooses a constant skew for its clock from the above distribution.

Figure 4 shows the mean and maximum latencies experienced by packets that cross 10 links when, as said before, the distribution of clock skews have standard deviations ranging from 0 to 10^5 milliseconds. As expected, LIQ is not affected by clock skews, since it does not consider injection time (which could be affected by clock skews), but waiting time, which is always correctly computed since all clocks run at the same speed (there are no drifts). It is also noticeable that the mean and the maximum latencies of LIQ are very low, which is not the case for the other policies, especially when clock skews grow. At first sight, it seems a bit paradoxical the fact that the mean latency with SIS decreases when skews grow. However, this behavior may be attributed to the fact that increasing skews randomizes the behavior of the policy. Note that the maximum latency with SIS does not seem to be significantly affected by the skew variation. LIS suffers from increasing clock skews, since its effectiveness relies on the accuracy of clocks. When skews grow, LIS clearly degrades its performance. Finally, we want to emphasize the great distance between the mean and the maximum in the case of SIS, and in the case of LIS for large skews.

Figure 5 shows how the number of hops a packet needs to reach its destination affects the latency. Here we see that LIS when all clocks are synchronized (no skews) and LIQ give analogous results, and behave quite uniformly on the number of hops. It is again noticeable that the mean and the maximum are much closer in the cases

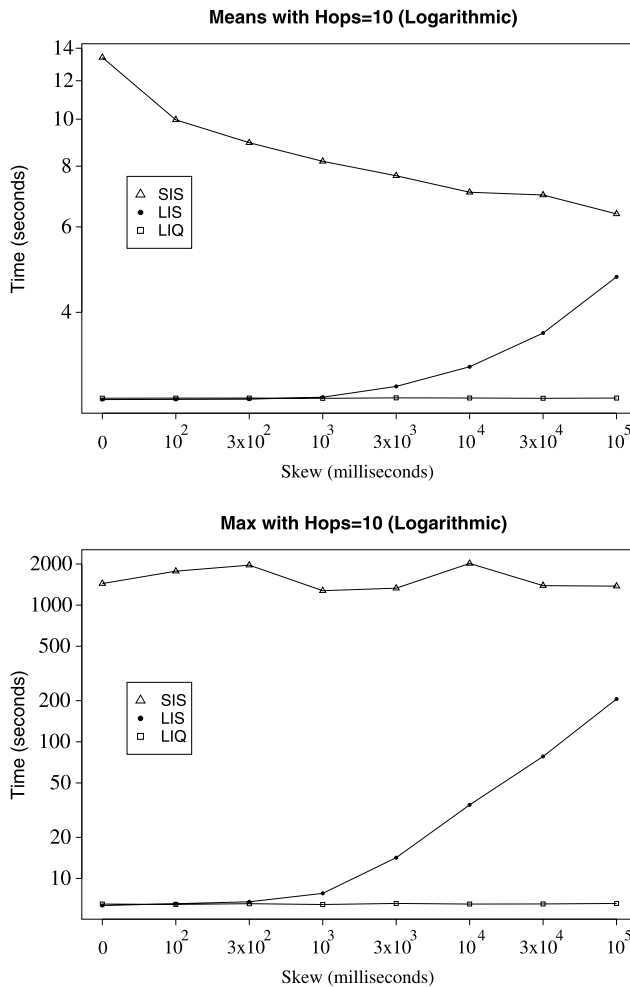


Fig. 4 Latency experienced by packets that cross 10 links with policies LIS, SIS, and LIQ under distribution of skews with different standard deviations

of LIS with no skews and LIQ, than in the other cases (between one and two orders of magnitude). Observe that, while with LIS the slope of the curve increases with the skew, with SIS the slope decreases.

8 Conclusion

We considered the continuous version of the well-known Adversarial Queuing Theory (AQT) model as scenario. It was generalized by considering the possibility that the router clocks in the network are not synchronized. The model was called Non-Synchronized CAQT (NSCAQT). We have shown that in NSCAQT when, although

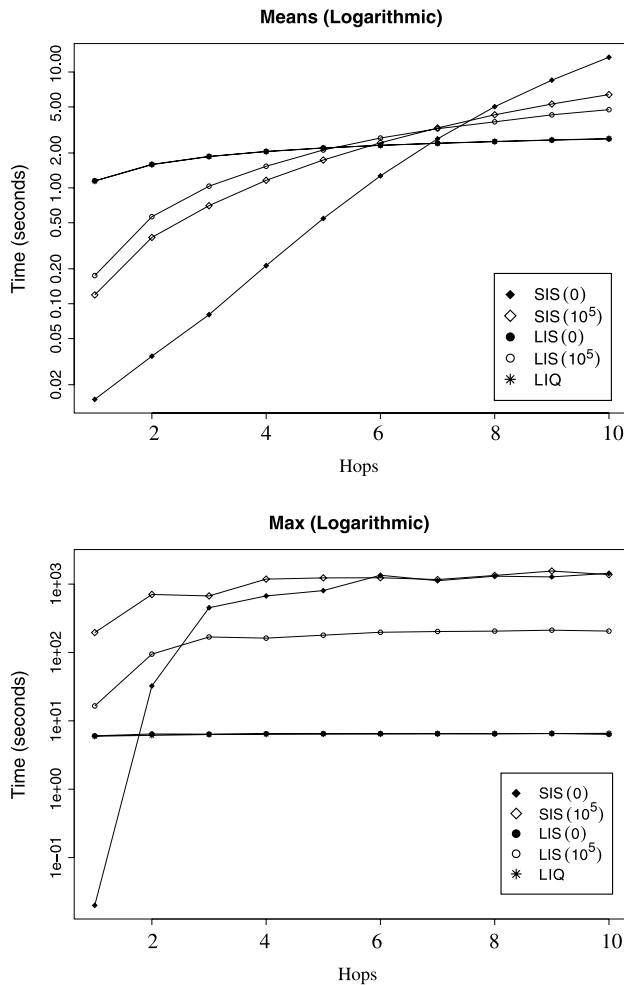


Fig. 5 Latency experienced by packets with policies LIS, SIS, and LIQ under normal distribution of skews with standard deviations of 0 and 100000

not synchronized, all clocks run at the same speed; all universally stable policies in CAQT that use the injection time and the remaining path to schedule packets remain universally stable. In a second approach, we studied the case in which clock differences can vary over time, but the maximum difference is bounded. Under this framework, we introduced two families of policies called LISP and SISP based on the well-known policies LIS and SIS. We have shown that both LISP and SISP families are universally stable. The bounds that we showed in these cases depend on the maximum difference between clocks. This is a necessary requirement, since we also have shown that LIS is not universally stable in systems without bounded clock difference. We introduced a new universally stable policy called Longest In Queues. In the case when clocks maintain constant differences, the bounds we proved do not depend on

them. To finish, we have provided simulation results that compare the behavior of some of these policies in a network with stochastic injection of packets.

We believe that the policy proposed LIQ could be an interesting alternative to other popular policies. We believe that further study of LIQ is required. In particular, we would like to know whether it is stable in NSCAQT under unbounded skew but constant drift. Along these lines, it would be interesting to devise a policy that is stable in this model.

Acknowledgements The authors would like to thank Anna Puig, Agustín Santos and Juan Céspedes for fruitful discussions. This work was partially supported by the Spanish MEC under grants TIN2005-09198-C02-01 and PR2008-0015, the Spanish MICINN under grant TIN2008-06735-C02-01, and the Comunidad de Madrid under grant S-0505/TIC/0285. The last author gratefully acknowledges the support of Universidad de Chile, Facultad de Ciencias Físicas y Matemáticas via a postgraduate fellowship, Proyecto Mecesp UCH0009, CONICYT via Anillo en Redes ACT08 and FONDAP in Applied Mathematics.

References

1. Andrews, M.: Instability of FIFO in the permanent sessions model at arbitrarily small network loads. In: Bansal, N., Pruhs, K., Clifford, S. (eds.) SODA, pp. 219–228. SIAM, Philadelphia (2007)
2. Andrews, M., Zhang, L.: The effects of temporary sessions on network performance. *SIAM J. Comput.* **33**(3), 659–673 (2004)
3. Andrews, M., Awerbuch, B., Fernández, A., Kleinberg, J., Leighton, T., Liu, Z.: Universal stability results and performance bounds for greedy contention-resolution protocols. *J. ACM* **48**(1), 39–69 (2001)
4. Borodin, A., Fernández, A., Goel, A., Zhang, L.: Source routing and scheduling in packet networks. *J. ACM* **52**(4), 582–601 (2005)
5. Blesa, M.J., Calzada, D., Fernández Anta, A., López, L., Martínez, A.L., Santos, A., Serna, M.J., Thraves, C.: Adversarial queuing model for continuous network dynamics. *Theor. Comput. Syst.* **44**(3), 304–331 (2009)
6. Borodin, A., Kleinberg, J., Raghavan, P., Sudan, M., Williamson, D.P.: Adversarial queuing theory. *J. ACM* **48**(1), 13–38 (2001)
7. Cruz, R.L.: A calculus for network delay, part I: network elements in isolation. *IEEE Trans. Inf. Theory* **37**(1), 114–131 (1991)
8. Cruz, R.L.: A calculus for network delay, part II: network analysis. *IEEE Trans. Inf. Theory* **37**(1), 132–141 (1991)
9. <http://j-sim.org/>
10. Le Boudec, J.-Y., Thiran, P.: Network Calculus: A Theory of Deterministic Queuing Systems for the Internet. Lecture Notes in Computer Science. Springer, Berlin (2001)
11. Parekh, A.K., Gallager, R.G.: A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Trans. Netw.* **1**(3), 344–357 (1993)
12. Parekh, A.K., Gallager, R.G.: A generalized processor sharing approach to flow control in integrated services networks: the multiple-node case. *IEEE/ACM Trans. Netw.* **2**(2), 137–150 (1994)
13. Santos, A., Fernández Anta, A., López, L.: Evaluation of packet scheduling policies with application to real-time traffic. In: Actas de las V Jornadas de Ingeniería Telemática, JITEL (2005)
14. Weinard, M.: The necessity of timekeeping in adversarial queuing. In: Nikolettseas, S.E. (ed.) WEA. Lecture Notes in Computer Science, vol. 3503, pp. 440–451. Springer, Berlin (2005)