



Window-games between TCP flows[☆]

Pavlos S. Efraimidis^{a,*}, Lazaros Tsavlidis^a, George B. Mertzios^b

^a Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece

^b Department of Computer Science, RWTH Aachen University, Germany

ARTICLE INFO

Article history:

Received 28 December 2008

Received in revised form 14 September 2009

Accepted 23 March 2010

Communicated by P. Spirakis

Keywords:

Algorithmic game theory

Network games

Nash equilibrium

ABSTRACT

We consider network congestion problems between TCP flows and define a new game, the *Window-game*, which models the problems of network congestion caused by the competing flows. Analytical and experimental results show the relevance of the Window-game to real TCP congestion games and provide interesting insight into the respective Nash equilibria. Furthermore, we propose a new algorithmic queue mechanism, called Prince, which at congestion makes a scapegoat of the most greedy flow. We provide evidence which shows that Prince achieves efficient Nash equilibria while requiring only limited computational resources.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Algorithmic problems of networks can be studied from a game-theoretic point of view. In this context, the flows are considered independent players who seek to optimize personal utility functions, such as the bandwidth. The mechanism of the game is determined by the network infrastructure and the policies implemented at regulating network nodes, like routers and switches. The above game-theoretic approach has been used for example in [33] and in several recent works like [22,21,2,30,13,10].

The focus of this work is on congestion problems between competing TCP flows, a problem that has been addressed in [18,2,12]. The Transmission Control Protocol (TCP) belongs to the Internet protocol suite and is the main data transmission protocol of the Internet. TCP flows transmit their data by sending a series of packets. Assume a TCP flow that is ready to send a large volume of data as a sequence of 2^{20} packets, each of size 1024 bytes. How should the flow submit the packets to the network? Each TCP flow has a parameter w , called the congestion window. The flow starts by submitting w packets to the network and then waits until either a packet's arrival is confirmed, normally with an acknowledgement packet (ACK), or a packet is considered lost. As soon as the number of the in-flight packets of the flow is less than w , the flow submits new packet(s); the result is that, at any moment in time, the flow can have at most w packets in flight. So, the size w of the congestion window has a strong impact on the transmission rate of a flow. Actually, the size of the congestion window, to a large degree, controls the speed of transmission of a TCP flow [17]. Consequently, the selection of an appropriate value for w is a very critical task for each flow, and this is where AIMD (Additive Increase Multiplicative Decrease) comes into play. AIMD is the most popular procedure for a TCP flow to constantly adjust its window size to the current network conditions. The basic principle of AIMD is that, for each successful packet delivery the flow increases its congestion window size additively by an amount proportional to a parameter $\alpha > 0$ (usually $\alpha = 1$) and for each lost packet, the flow decreases its congestion window multiplicatively by a parameter $\beta < 1$ (usually $\beta = 1/2$). The values of the α and β parameters have a decisive

[☆] Preliminary results of this work appeared in the Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT 2008), in: Lecture Notes in Computer Science, vol. 4997, Springer, 2008, pp. 95–108.

* Corresponding author. Tel.: +30 25410 79756; fax: +30 25410 79764.

E-mail addresses: pefraimi@ee.duth.gr (P.S. Efraimidis), tsavlid@ee.duth.gr (L. Tsavlidis), mertzios@cs.rwth-aachen.de (G.B. Mertzios).

role on the behavior of the AIMD flow. A large value of α and/or β makes the flow more aggressive, whereas a small value of α and/or β makes it more temperate. A brief description of further TCP related concepts is given in Section 4, while more details can be found in common computer networks books like [31,34].

In this work, we consider games where flows can freely choose their congestion window size or where flows use the AIMD algorithm and can choose its parameters. To this end, we propose a new game, the *Window-game*, which models the problem of network congestion caused by the competing TCP flows. The novelty of the new model lies in the fact that it focuses on the congestion window, a parameter that is in the core of the network algorithms of modern TCP flows. As already noted, the size of the congestion window, to a large degree, controls the speed of transmission of a TCP flow.

The *Window-game* has a router and n flows and is played synchronously, in one or more rounds. Every flow is a player that in each round selects the size of its congestion window. The router (the mechanism of the game) receives the actions of all the flows and decides how the capacity is allocated. Based on how much of its requested window has been satisfied, each flow decides on the size of its congestion window for the next round. The utility of each flow is the capacity that it obtains from the router in each round minus the cost for lost packets.

The main contributions of this work are three. The first is the definition of the *Window-game*, a game-theoretic model for network games between TCP flows. The motivation to address the TCP congestion problem originated from the following question, posed in [18,30]: *Of which game or optimization problem is TCP/IP congestion control the Nash equilibrium or optimal solution?* The *Window-game* is meant to be a natural model that is simple enough to be studied from an algorithmic and game-theoretic point of view, while at the same time it captures essential aspects of the real TCP game. By focusing on the congestion window, the *Window-game* is simpler and more abstract than the model used in [2], while still being sufficiently realistic to model actual TCP games. The *Window-game* allows us to address both repeated versions of the game that resemble actual TCP games (like the games studied in [2]) and one-shot versions of the game with generic (not necessarily AIMD) flows.

The second contribution is the analysis of interesting queue mechanisms from a game-theoretic point of view. In particular, the case of DropTail with non-synchronized packet losses, which we consider the most realistic and interesting version of TCP games, has not been addressed before (to the best of our knowledge). Also new is the analysis of the one-shot versions of the TCP games.

The final contribution is a new queue policy, called Prince, which is simple and efficient enough to cope with the requirements of real network routers. We show that the exemplary punishment of the most greedy flow induces strong incentives on the flows to play in a socially efficient manner. In general, it seems that focusing on the most greedy player proves to be a simple and very effective way to make selfish players that share a common resource behave responsibly. This claim is further supported by the fact that punishing the highest rate flow has also been successful in a different network game model [12] where the flows are Poisson sources. Prince is much simpler than the policy of [12], and it is the first time the “punish the leader” approach is applied to games with TCP flows.

We provide theoretical evidence and experimental results to support the above claims.

Outline. The rest of the paper is organized as follows: Preliminaries and notation are presented in Section 2. The *Window-game* model is described in Section 3. An overview of TCP congestion control concepts is given in Section 4. The Prince mechanism along with common router policies are discussed in Section 5. We consider *Window-games* where the players are AIMD flows in Section 6 and present the corresponding experimental results in Section 7. *Window-games* with generic, non-AIMD, flows are discussed in Section 8 (*Window-games* with complete information) and in Section 9 (*Window-games* with incomplete information). Finally, a discussion of the results is given in Section 10.

2. Preliminaries and notation

A *Nash Equilibrium (NE)* is a strategy profile where no flow can gain by unilaterally changing its behavior. A *Symmetric Nash Equilibrium (SNE)* is a NE where all flows have the same behavior. A game in normal form is *symmetric* if all players have the same strategy set, and the payoff to playing a given strategy depends only on the strategies being played, not on who plays them. The basic notation used in work is summarized in Table 1.

3. The Window-game


The *Window-game* models the interaction between the congestion windows of competing TCP flows. The main entities of a *Window-game* are a router with capacity C and a set of $n \leq C$ flows,¹ as depicted in Fig. 1b. The router uses a queue policy to serve in each round up to C workload. The n flows are the independent players of the game. Unless otherwise specified, the number n is considered unknown to the players and to the router. The game consists of one or more rounds. In each round, every player i selects a size $w_i \leq C$ for its congestion window and submits it to the router. The router collects all requests and applies the queue policy to allocate the capacity to the flows. The common resource is the router’s capacity, an

¹ The constraint $n \leq C$ ensures that the fair share of each flow is at least one. A number of flows $n > C$ would indicate a heavily congested router, where a flow may not be entitled to submit even a single packet. We focus on problems where there is sufficient bandwidth for all flows and each flow tries to identify its optimal number of packets to submit.

Table 1

The main notation used in this work.

TCP	The <i>Transmission Control Protocol</i> of the Internet protocol suite
AIMD	The <i>Additive Increase Multiplicative Decrease</i> algorithm
w_i	The congestion window size of flow i
α_i	AIMD parameter for additive increase of flow i
β_i	AIMD parameter for multiplicative decrease of flow i
g	The cost for a lost packet
C	The router capacity
n	The number of flows in the Window-game
τ	The number of rounds between two congestion rounds in a game
τ_i	The number of rounds between two loss rounds of flow i
G_i	The goodput of flow i : The (average) number of useful packets per round
L_i	The (average) number of packets that flow i loses at loss rounds
N_i	The (average) size of the congestion window of AIMD flow i after a loss round

**Fig. 1.** The Window-game along with the corresponding TCP network model.

abstract concept that corresponds to how much load the router can handle in each round. The objective of each player is to send as many packets as possible. At the same time the player has to avoid packet losses, as each packet loss induces a cost $g \geq 0$ to the player. For games with multiple rounds, each packet loss (and each successful packet) causes the respective flow to adjust the size of its congestion window. The adjustment is determined by the game strategy of the flow and the penalty model.

Definition 1. *One round of the Window-game.*

Mechanism: A router with capacity $C > 0$ and a queue policy

Players: n flows $i = 1, \dots, n$, where $2 \leq n \leq C$

Actions of player i : The window size w_i of i , where $w_i \in \{0, 1, \dots, C\}$

Utility of player i : $u(i) = \text{transmitted}(i) - g \cdot \text{dropped}(i)$, where

g : the cost for each dropped packet

$\text{transmitted}(i)$: number of successfully transmitted packets

$\text{dropped}(i)$: number of dropped packets

Each round of the game is executed independently; no work is pending at the start of a round and no work is inherited to a following round. The window size w_i of a flow i is always $w_i \geq 1$, except for specific rounds where the flow is compelled by its penalty model to refrain from transmitting any packet ($w_i = 0$). The penalty models are discussed in Section 4.

Network routers and flows process massive streams of network packets in real time. Therefore, algorithms used by routers and flows may use only limited computational resources like memory and processing power. In particular, it is considered unrealistic for network routers to keep detailed statistics on a per flow basis. The queue policy of the router should be stateless or use as little state information as possible.²

The Window-game can model not only games between AIMD flows but also games between generic flows. The only requirement is that the packet stream generated by each generic flow can be represented as a sequence of window sizes. More precisely, we consider the following cases:

Window-games with AIMD flows. In these games the players are AIMD flows. Each AIMD flow i selects its parameters (α_i, β_i) once, for the whole duration of the game. The utility for each flow i is the average number G_i of useful (i.e., distinct) packets of the flow that are successfully delivered in each round of the game. G_i is the effective bandwidth of flow i and is called the goodput of flow i . This class of Window-games is strongly related to the game(s) currently played by real TCP flows and exists implicitly in the analysis of [18,2].

² Stateful network architectures are designed to achieve fair bandwidth allocation and high utilization, but need to maintain state, manage buffers, and perform packet scheduling on a per flow basis. Hence, the algorithms used to support these mechanisms are less scalable and robust than those used in stateless routers. The stateless substance of nowadays IP networks allows Internet to scale with both the size of the network and heterogeneous applications and technologies [35,36].

Window-games with generic flows. These are games where the flows can use arbitrary algorithms to choose their congestion window. Depending on the number of rounds and on how well each flow is informed about the existence of other flows, we distinguish the following classes:

- One-shot game with complete information.
- One-shot game with incomplete information with no prior probabilities.
- One-shot game with incomplete information with prior probabilities.
- Repeated game with incomplete information.

The action of a generic flow is to choose the size of its congestion window for each round. In one-shot games, the utility of the flow is the goodput minus the cost for lost packets. In repeated games, the overall utility is the average utility of all the game rounds.

Assumptions. We make a set of simplifying assumptions similar to the assumptions of [2]. The Window-game is symmetric. All flows use the same network algorithms for loss recovery. All packets of all flows are of the same size and have the same Round-Trip-Time (RTT), i.e., the time from the submission of a packet P until the corresponding acknowledgement packet (ACK) is received. Packet losses are caused only by network congestion.

Solution concept. The solution concept for the Window-games is the Nash equilibrium (NE) and in particular the Symmetric Nash equilibrium (SNE). A good reason to start with SNE is that the analysis appears to be simpler compared to general NE. It is noteworthy, that both the analytical and the experimental results show that in many cases there are only symmetric (or almost symmetric) NE. For each Window-game, we study its SNE and discuss how efficient the network operates at them. In certain cases, we search experimentally for general, not necessarily symmetric, NE.

Efficiency and fairness. In the context of Window-games we consider a NE to be efficient if at equilibrium the router capacity is utilized without risking a congestion collapse. At the same time, the packet losses of the flows should be as small as possible. The exact range of acceptable values for a NE to be considered efficient may vary between different Window-games. Some indicative thresholds are that the total throughput should be larger than 80% of the capacity C and the total loss rate should not exceed 10% of the packets. We consider a state of the Window-game (at equilibrium or not) fair, if all flows receive, at least approximately, the same bandwidth. A formal measure of fairness is for example the Chiu-Jain fairness index [7]. For the network model used in this work, the following simple fairness index of [3] is sufficient:

$$\frac{G_{\max} - G_{\min}}{G_{\text{avg}}}, \quad (1)$$

where G_{\max} , G_{\min} and G_{avg} are the maximum, minimum and average flow goodputs, respectively.

Network topology. In this work, we assume a one-hop topology with one router and a set of flows that submit packets to the router. This model, albeit simple, captures important aspects of the real network problem and is commonly used in the analysis of network games and algorithms; for example in [2,33,22], in [19,10,23,12] and also in the original AIMD analysis paper [7].

How realistic is the Window-game? An important difference of Window-games to real TCP/IP games is that Window-games are played in rounds while real TCP/IP network games are clocked at the packet level. One round of the Window-game corresponds to a time-interval roughly equal to the flow's Round-Trip-Time. In this time interval, the flow will receive ACK's for all packets that were pending before P. The number of these pending packets is determined by the window size at the time when packet P was submitted. Therefore, the round-based approach appears to be a natural approximation of the real TCP/IP game. A round-based approach has also been used in the analysis of [2]. A further difference is that the Window-game model does not take into account queuing delay. Because of this, the value of the goodput that are calculated may not always correspond to real goodput value. This, however, is not an obstacle for the applications of the Window-game in this work, since we are mainly interested in the relative goodputs for different strategies of a TCP flow. Finally, we use the term capacity for the router because we want to emphasize that this is the amount of work that the router can handle in one round. The capacity is determined by the bandwidth load that the router can handle.

The novelties of the Window-game model. The Window-game model is a new, general model for TCP games that significantly differs from previous models. The Window-game

- is a general model that can model games between generic Internet flows. The only requirement is that the packet stream generated by each flow can be represented with a sequence of window sizes. In the special case of AIMD flows, it reduces to a model very close to the model of [2];
- is designed to effectively capture games between window-based flows, like TCP flows. To our knowledge, it is the first model with this feature;
- has a plain game structure that admits the definition and analysis of one-shot games between internet flows. This is the first model to serve this purpose. The analysis of one-shot games can provide insight into the nature of Internet congestion games, which in turn can support the analysis of the more realistic repeated Window-games;
- is a model that is both simple and relevant to real Internet congestion problems. The Window-game achieves this by emphasizing on the most critical parameter of window-based flows, the congestion window.

Conclusion. In our view, the Window-game fulfils an important set of requirements that we had set. That is, to capture the interaction of the competing flows, to indicate when a flow can gain more goodput by changing its strategy and to

characterize the efficiency and the fairness of the NE. Additionally, the analytical and experimental evidence show that the Window-game is strongly related to the actual TCP/IP congestion games.

4. Congestion control in TCP

In this section, we provide a short description of the basic concepts of TCP and TCP congestion control.

TCP: The Transmission Control Protocol (TCP) is one of the core protocols of the Internet protocol suite. TCP is a window-based transport protocol; in TCP every flow has an adjustable window, which is called the *congestion window*, and uses it to control its transmission rate [17].

Congestion Window: A TCP flow submits packets to the network and expects to receive acknowledgements (ACKs) for the packets that reach their destination. The congestion window defines the maximum number of outstanding packets for which the flow has not yet received acknowledgements [20]. Essentially, the congestion window represents the sender's estimate of the amount of traffic that the network can absorb without becoming congested. The most common algorithm to increase or decrease the congestion window of a TCP flow is AIMD.

AIMD: The Additive Increase Multiplicative Decrease (AIMD) algorithm [7] can be considered as a probing procedure designed to find the maximal rate at which flows can send packets under current conditions without incurring packet drops [18]. An AIMD flow i has two parameters α_i and β_i ; upon success, its congestion window is increased additively by α_i (in each round), and upon failure, its congestion window is decreased multiplicatively by a factor β_i .

Definition 2. The basic AIMD algorithm:

$$w_i = \begin{cases} w_i + \alpha, & \text{upon successful delivery of all packets of a round,} \\ w_i \cdot \beta, & \text{in case of packet drop.} \end{cases}$$

AIMD is so far considered to be the optimal choice in the traditional setting of TCP Reno congestion control and FIFO DropTail routers. If, however, we consider the developments like TCP SACK and active queue management, AIMD may no longer be superior [3].

Packet loss: When congestion occurs at a router, packets are dropped. AIMD flows assume that a packet has been lost if either the corresponding ACK packet does not arrive within a specific interval (timeout) or if the flow receives a *duplicate ACK (DACK)* that indicates that some packet has not been received.³ An AIMD flow responds to packet loss with a decrease of its congestion window and with a loss-recovery procedure which incurs some cost to the flow. Popular versions of TCP, like TCP Tahoe, TCP-Reno and TCP-SACK, differ in the way they handle a packet loss.

- **TCP Tahoe:** A packet loss is indicated when the timeout occurs before the corresponding ACK is received. The congestion window will then be reduced to 1 packet and the slow start phase will be initiated. If the sender receives three duplicate ACKs (DACKs⁴) for the same packet before a timeout occurs, the sender first retransmits the packet (fast retransmit) and then reduces its congestion window to 1. This reaction of TCP Tahoe corresponds to the severe reaction to packet loss; even a single packet loss causes the flow to drastically reduce its congestion window.
- **TCP Reno:** Has the same mechanism with Tahoe, except that in the case of 3 DACK it retransmits the packet that the receiver waits for and then halves its congestion window (fast recovery). In the case of multiple dropped packets from the same congestion window, a Reno flow tries multiple fast retransmissions until a timeout occurs and then comes in to slow start. The reaction of TCP Reno to packet loss is considered a hybrid reaction.
- **TCP SACK:** This variant of TCP is an extension of Reno. SACK has a selective acknowledgement option that allows receivers to additionally report non-sequential data they have received. When coupled with a selective retransmission policy in TCP senders, an efficient recovery mechanism, which avoids unnecessary retransmissions, is formed. The reaction of TCP SACK to packet loss is considered a gentle reaction.

Since all loss-recovery schemes incur costs to the flow they can be thought of as a penalty on the flows which suspend their normal transmission rate until lost packets are retransmitted.

Penalty Model: We formed a penalty-based model, which is similar to the model of [18,2], to define a flow's behavior when losses occur. Assume that a flow i with current window size w_i has lost $L_i \geq 1$ packets in the last round. Then:

- **Gentle penalty (resembles TCP SACK):** The flow reduces its window to $w_i = \beta \cdot w_i$ in the next round.
- **Severe penalty (TCP Tahoe):** The flow timeouts for τ_s rounds, i.e., it does not send any packet for τ_s rounds, and then continues with a window size of $w_i = \beta \cdot w_i$.
- **Hybrid penalty (TCP Reno):** If $L_i = 1$ the flow applies gentle penalty. If $L_i > 1$ a progressive severe penalty is applied. The flow timeouts for $\min\{2 + L_i, 15\}$ rounds and then restarts with a window equal to w_i/L_i . This penalty is justified by the experimental results of [8,28].

³ A duplicate ACK is sent from the receiver to the sender when an out-of-order packet is received; there are several reasons why an out-of-order packet might be received. In this work, we will assume that an out-of-order packet is received only when its preceding packet (the one that the receiver was expecting) has been lost.

⁴ When a duplicate ACK is received by a real TCP flow, the flow does not know if it is because a TCP segment was lost or simply that a segment was delayed and received out of order at the receiver. If the receiver can re-order segments, it should not be long before the receiver sends the latest expected acknowledgement. Typically no more than one or two duplicate ACKs should be received when simple out of order conditions exist. If however more than two duplicate ACKs are received by the sender, it is a strong indication that at least one segment has been lost.

5. Router policies and the Prince mechanism

The router, and in particular the queue algorithm deployed by the router to allocate the capacity to the flows, defines the mechanism of the Window-game. We examine the common policies DropTail, RED, MaxMin, CHOKe and CHOKe+. The implementations of the above policies in the Window-game aim to simulate the behavior of the real policies on packet streams. We also examine two variants of the proposed Prince policy. Our interest for all policies is on the NE of the corresponding Window-games.

- **DropTail:** The DropTail policy is the simplest and most common queue discipline of network routers. With DropTail the packets are served in FIFO order. If the queue is full, incoming packets are dropped. The Window-game adaptation of DropTail, simulates how DropTail would behave in the real, stream-based case. In each round, if the sum w_{total} of the requested windows does not exceed the capacity C , the router serves all packets. Otherwise, the router drops a random sample of packets of size equal to the overflow, and serves the remaining C packets.
- **RED [9]:** The Random Early Detection (RED) policy is an active queue management algorithm that may drop packets even before an overflow occurs. More precisely, at overflow RED behaves like DropTail. However, for queue occupancies between a low threshold \min_{th} and a high threshold \max_{th} of the queue size, packets are dropped with a positive probability p . The Window-game adaptation of RED works as follows: If the sum w_{total} of the requested windows is $w_{total} < \min_{th}$ all packets are served. If $w_{total} \geq \min_{th}$ a random sample of packets is selected. The selected packets correspond to positions higher than \min_{th} in a supposed queue. Each chosen packet is dropped with a probability proportional to the router's load. The values $\min_{th} = 70\%$ and $\max_{th} = 100\%$ are used for the thresholds.
- **MaxMin:** An allocation of a common resource is MaxMin fair if no share may be increased without simultaneously decreasing another share which is already smaller. Router policies that implement the MaxMin fairness criterion achieve fair allocation of the available capacity to the flows. Given a network with its link capacities and a set of flows with their maximum possible transmission rates, there is a unique set of flow rates that satisfies the max-min conditions. MaxMin Fairness is a stateful, fair queue policy [5] that conceptually corresponds to applying round-robin for allocating the capacity. The router offers its packet transmission slots to its users by polling them in round-robin order. If a flow is offered a chance to use a link slot but has no packets ready, then that same slot is offered to the next flow, until a ready flow is found. In each pass of a link's round robin, a flow may transmit only one packet [15]. The adaptation of MaxMin to the Window-game is straightforward.
- **CHOKe:** CHOKe [29] is a stateless active queue management algorithm that differentially penalizes misbehaving flows by dropping more of their packets. For every packet that arrives at the congested router, CHOKe draws a packet at random from the FIFO buffer and compares it with the arriving packet. If they both belong to the same flow, then they are both dropped, else the randomly chosen packet(s) is admitted into the buffer with a probability that depends on the level of congestion. Packets belonging to a misbehaving flow are both more likely to trigger comparisons and more likely to be chosen for comparison. Therefore, packets of misbehaving flows are dropped more often than packets of well-behaved flows [29]. The Window-game adaptation of CHOKe is similar to the adaption for RED. Every chosen packet that is above the min threshold, is compared to a packet chosen randomly from the packets that have been admitted until that moment. For the lower threshold, the upper threshold and the dropping probability we use the same values as for RED.
- **CHOKe+:** The CHOKe+ policy [2] is a variant of CHOKe that tries to avoid high loss rates and severe under-utilization of the available capacity, when the number of flows is relatively small. For each incoming packet P , the algorithm randomly selects k packets from the queue. Let m denote the number of packets from the k chosen candidate packets that belong to the same flow as the incoming packet. Let $0 \leq \gamma_2 \leq \gamma_1 \leq 1$ be positive constants. If $m \geq \gamma_1 k$, the algorithm drops packet P along with the m matching candidate packets. Otherwise, a drop probability p is calculated in an equivalent RED queue. Suppose that P is to be dropped according to RED. Now, if $\gamma_2 k \leq m < \gamma_1 k$, the algorithm also drops the m matching packets along with P . Otherwise, only packet P is dropped. The Window-game adaption is in the same spirit as the adaptations of CHOKe and RED.

Queue disciplines in use. DropTail is simple and efficient, seems to work reliably in practice and is currently the most widely used queuing discipline at Internet routers. RED is also widely deployed in order to maintain the buffer's utilization within desired levels. Both DropTail and RED, do not actively punish greedy flows.

5.1. The Prince mechanism

We propose a new queue policy, called Prince (of Machiavelli), which, at congestion, drops packets from the flow with the largest congestion window. If the overflow is larger than the window of the most greedy flow, the extra packets are dropped from the remaining flows with DropTail. If the overflow does not exceed the number of packets of the most greedy flow, then flows that do not exceed their fair share do not experience packet loss. Hence a greedy flow does not plunder the goodput of other flows.

The Prince policy is simple and does not require state information on a per-flow basis. This makes the policy appropriate for the requirements of real routers. We show that Prince-based queue algorithms clearly outperform traditional queue policies, like Droptail and RED, with respect to the efficiency of the induced Nash equilibria. Prince also achieves better results

than CHOKe and CHOKe+. In most cases, Prince achieves results comparable to the results of well-proven, but heavyweight, MaxMin-based queue policies.

We define a basic version of Prince that drops packets only in case of overflow and a RED-inspired version of Prince, called Prince-R, which applies its policy progressively starting from a min threshold $\min_{th} = 70\%$.

6. Window-games with AIMD flows

In this section, we analyze three characteristic Window-game scenarios between AIMD flows with gentle penalty, where the flows can choose the value of parameter α . Parameter β is fixed⁵ at its default TCP value $\beta = 1/2$. We admit parameter α to take integer values in the range $[1, C/2]$. In our analysis we use machinery from [2]. Let us start with the following definitions:

Definition 3. A congestion round of the Window-game is a round in which the router has to drop one or more packets. The common cause for a congestion round is that the sum of the requested windows is larger than the capacity of the router. In active queue management policies like RED, a congestion round may also occur when the load exceeds some threshold of the capacity and the router decides to drop one or more packets.

Definition 4. A loss round of flow i is a congestion round in which flow i experiences packet loss.

6.1. A model for AIMD flows

Assume a router with capacity C and a set of n AIMD flows that are active on the router. Let flow i (for fixed i) be one specific of the n flows. Let $(\alpha_i, \beta_i = 1/2)$ be the parameters of flow i . Assume that, when the system has reached steady state, flow i experiences a loss round every $\tau_i + 1$ rounds, that is, τ_i is the number of rounds between two successive loss rounds of flow i .

Let $N_i(k)$ denote the window size of flow i after its k th loss round. Then, as shown in [2], at steady state the window size $N_i(k)$ will converge for $k \rightarrow \infty$ to a fixed value N_i and flow i will have the following behavior. After each loss round, it starts with window size N_i and increases its window size successively to $N_i + \alpha_i, N_i + 2\alpha_i, \dots, N_i + \alpha_i\tau_i$. When the window size reaches $N_i + \alpha_i\tau_i$, the flow experiences packet loss and reduces its window size (in the next round) to $N_i = \beta_i(N_i + \alpha_i\tau_i)$. Thus, the congestion window of flow i takes the following sequence of values:

$$N_i, N_i + \alpha_i, N_i + 2\alpha_i, \dots, N_i + \alpha_i\tau_i. \quad (2)$$

By applying the definition of the AIMD algorithm (Definition 2) we get

$$N_i = \beta_i(N_i + \alpha_i\tau_i) = \frac{1}{2}(N_i + \alpha_i\tau_i), \quad (3)$$

and this gives

$$N_i = \alpha_i\tau_i. \quad (4)$$

If we focus on $\tau_i + 1$ consecutive rounds, starting immediately after a loss round of flow i , then the goodput G_i of flow i is

$$G_i = \frac{1}{\tau_i + 1} \left(N_i + (N_i + \alpha_i) + (N_i + 2\alpha_i) + \dots + (N_i + \alpha_i\tau_i - L_i) \right),$$

where L_i is the average number of packets that flow i loses at loss rounds. This gives

$$G_i = \frac{3}{2}\alpha_i\tau_i - \frac{1}{\tau_i + 1}L_i. \quad (5)$$

In most scenarios we can ignore the last term to get the simplified equation

$$G_i = \frac{3}{2}\alpha_i\tau_i. \quad (6)$$

In the above equations for an AIMD flow i we have assumed that the system has reached steady state. Based on this assumption we derived equations to capture the behavior of the flow. This approach has been introduced in [2]. In an excessively congested state of the system where, for example, a flow may experience packet loss in each round, the periodic

⁵ Experiments with parameter β show that in almost all cases selfish flows will use for β a value close to 1. The same conclusion can be extracted from the results in [2]. A simple argument for this behavior is that parameter β is very important when the flow has to quickly reduce its window size when the network conditions change. The TCP games examined in this work are rather static: Static bandwidth, static number of flows, static behavior of all flows during a game and hence there is no real reason for a flow to be adaptive.

model of Eq. (2) degenerates and the derived Eqs. (3) and (5) may become inaccurate. To avoid such situations, we make the following – reasonable for the scenarios that we consider – assumptions for Window-games with AIMD flows:

$$n \leq \frac{C}{2}, \quad A = \sum_{i=1}^n \alpha_i \leq \frac{C}{2}, \quad \text{and } \forall \text{ flow } i, \tau_i \geq 1. \quad (7)$$

In this work, we choose the above equations to model the behavior of AIMD flows in routers. We will use these equations to study the exact and in some cases the approximate behavior of AIMD flows in Window-games. We deem this model to be an elegant approximation of the real dynamics of AIMD flows. It is simple, it captures the periodic probing behavior of AIMD and, finally, it gives analytic results that are shown to be accurate in interesting scenarios. For these reasons, we prefer it over other models, where for example the flows are assumed to be Poisson sources of packets.

6.2. DropTail router with synchronized packet losses and gentle penalty AIMD flows

The Game: A DropTail router of capacity C serves n gentle penalty TCP flows. The term *synchronized packet losses* refers to the following simplifying assumption: At congestion rounds, all flows experience packet loss, that is, every congestion round is a loss round for all flows. The utility of each flow-player in this game is the goodput of the flow. Each flow can choose its parameter α_i , while parameter β is set to $\beta = 1/2$.

Using the Window-game model it is easy to show that the selfish flows will try to over-utilize the common resource (the network) by using a large value for their parameter α . The outcome will be a network that is heavily congested. A close analogy is the so-called “*Tragedy of the Commons*” [16] problem in economics where each individual can improve her own position by using more of a common resource, but the availability of the resource degrades as the number of greedy users increases.

Since every congestion round in this game is a loss round for all flows, at steady state, the number τ_i of rounds between two successive loss rounds of any flow i is equal to number τ of rounds between two successive congestion rounds of the system, i.e.,

$$\tau_1 = \tau_2 = \dots = \tau_n = \tau. \quad (8)$$

We have AIMD flows, all with a steady number of rounds between loss rounds. So we can apply the model of Section 6.1.

Let $A = \sum_{i=1}^n \alpha_i$ and $A_i = A - \alpha_i$. After each round without packet loss, the total window size of all flows $w_{total} = \sum_{i=1}^n w_i$ increases by A . In congestion rounds w_{total} exceeds the capacity C of the router, i.e., $w_{total} > C$. The number of excess packets in a congestion round is a random integer in the range $1, 2, \dots, A$. We will assume that the average excess is $L = A/2$, i.e., the half of the increase step. Thus, the total window size w_{total} in congestion rounds is $w_{total} = C + A/2$ and the average packet loss for each flow i is $L_i = \alpha_i/2$. After a congestion round, all flows multiplicatively decrease their congestion window. Using Eq. (4), $\beta_i = 1/2$ and $N = \sum_{i=1}^n N_i$ we obtain

$$N = \frac{1}{2}(C + A/2). \quad (9)$$

Using Eq. (3) for all flows and summing up gives:

$$N = \frac{1}{2}(N + A \cdot \tau) \Rightarrow \tau = \frac{C}{2A} + \frac{1}{4}. \quad (10)$$

From Eq. (5) we get that for any flow i the goodput G_i is:

$$G_i(\alpha_i) = \frac{3}{2}\alpha_i \left(\frac{C}{2(A_i + \alpha_i)} + \frac{1}{4} \right) - \alpha_i \left(\frac{C}{A_i + \alpha_i} + \frac{5}{2} \right)^{-1}. \quad (11)$$

We will show that:

Lemma 5. $G_i(\alpha_i)$ in (11) is a strictly increasing function of parameter α_i , for $\alpha_i \leq C$ and $A_i \leq C$.

Proof. The first derivative of G_i with respect to α_i is positive at $\alpha_i = C$, for $A_i \leq C$. In particular, if the derivative at $\alpha_i = C$ is written as a single fraction, its numerator is

$$15A_i^2 + 57A_iC + 14A_i^3C + 15C^2 + 64A_i^2C^2 + 3A_i^4C^2 + 68A_iC^3 + 15A_i^3C^3 + 18C^4 + 24A_i^2C^4 + 15A_iC^5 + 3C^6,$$

and its denominator is

$$4(A_i + C)^2 (3 + A_iC + C^2)^2.$$

The second derivative of G_i with respect to α_i is always negative. More precisely, the second derivative written can be written as a fraction where the denominator is strictly positive and the numerator is the sum of the positive terms

$$12\alpha_i^3C + 36\alpha_i^2A_iC + 36\alpha_iA_i^2C + 12A_i^3C + 4\alpha_i^3A_iC^2 + 12\alpha_i^2A_i^2C^2 + 12\alpha_iA_i^3C^2 + 4A_i^4C^2,$$

and the negative terms

$$\begin{aligned} & -81A_iC - 81\alpha_iA_iC^2 + 4\alpha_i^3A_iC^2 - 81A_i^2C^2 - 27\alpha_i^2A_iC^3 - 54\alpha_iA_i^2C^3 \\ & - 27A_i^3C^3 - 3\alpha_i^3A_iC^4 - 9\alpha_i^2A_i^2C^4 - 9\alpha_iA_i^3C^4 - 3A_i^4C^4. \end{aligned}$$

For $\alpha_i \leq C$, the total sum is always negative, and therefore the second derivative is always negative. Hence, $G_i(\alpha_i)$ is strictly increasing for $\alpha_i \leq C$. Actually, the result holds even for larger values of A_i up to $A_i \leq 15C$, but such values are not relevant to the scenarios that we consider. Recall that for AIMD flows we have assumed the conditions (7). \square

In normal network conditions, the per flow increase parameters α_i and the total increase A in each round should be small compared to C . This means that the term L_i can be neglected in the calculation of the goodput G_i of a flow i and that in congestion rounds w_{total} can be assumed to be $w_{total} = C$. Using these assumptions in Eqs. (9)–(11), the previous analysis gives the simplified expression

$$G_i = \frac{3\alpha_iC}{4(A_i + \alpha_i)}, \quad (12)$$

which is clearly an increasing function of α_i .

From the above analysis we conclude that in this game, every selfish flow i will have incentives to increase its parameter α_i . The result is that the game will operate at a very inefficient state, that resembles a Tragedy of the Commons situation. The above result is not surprising and is in agreement with the results in [2].

6.3. DropTail router with gentle penalty AIMD flows

We consider again a Drop-tail router but this time with non-synchronized packet losses. In this case, at congestion rounds a random set of packets is dropped. The packets to be dropped are randomly selected from the submitted packets and their number is equal to the packet excess. In each congestion round, a flow may or may not experience packet loss. The expected number of packets that it will lose is proportional to its window size. At first glance, this case appears to provide better incentives to flows not to behave aggressively. A flow with a smaller congestion window at congestion rounds, has a smaller probability to experience packet loss, compared to a flow with a larger congestion window. However, we will show that this scenario is a Tragedy of the Commons situation, as well.

This game has not been studied analytically before. The fact that packet losses are not synchronized makes the analysis harder. We provide an analysis for a toy case with $n = 2$ players.

6.3.1. DropTail router with two gentle penalty AIMD flows

The Game: Assume a router with capacity C and $n = 2$ flows with parameters α_1, α_2 , and $\beta_1 = \beta_2 = 1/2$. Each flow i can choose the value of its parameter α_i . *Simplifications:* We assume that at congestion only one flow is randomly selected and the excess is dropped from this flow. Hence only one flow experiences packet loss in each congestion round. The probability that a flow is selected at a congestion round for dropping the excess packets is proportional to the flow's window size. We will also assume that the selected flow has a window size at least as large as the excess. The utility of each flow-player in this game is the goodput of the flow. Each flow can choose its parameter α_i , while parameter β is set to $\beta = 1/2$.

We fix the value α_1 of flow 1 and assume that flow 2 chooses a value $\alpha_2 = z\alpha_1$, for $z > 1$. We will show that flow 2 achieves a higher goodput by increasing its parameter α_2 . We assume that the system has reached steady state and we use the model from Section 6.1 to model the two flows. Note that in this game the number of rounds between loss rounds of a flow i may vary during the game. To overcome this issue, the values τ_i and N_i will be the average values of the corresponding measures at steady state. The use of average values has an impact on the accuracy of the goodput equations. The more the value of τ_i starts to variate, the less accurate equation (6) becomes. Actually, the exact goodput in this case is larger than the value given by the equation. For the case of two flows that we analyze, the error in the goodput estimation by Eq. (6) is bounded (in the experiments it is always less than 5%) and it affects both flows evenly, therefore we will ignore it in the comparative analysis. From Eqs. (4) and (6) we get that $N_i = \alpha_i\tau_i$ and $G_i = \frac{3}{2}\alpha_i\tau_i - \frac{1}{\tau_i+1}L_i$, for $i = 1, 2$.

We can obtain one more expression for the goodput of each flow from a “typical” congestion round of the router. Let $w_{c,i}$ be the average window size of flow i at congestion rounds. We assume that congestion rounds occur periodically, every $\tau + 1$ rounds. In this case, the average total number of packets requested by flow i in $\tau + 1$ rounds, where the last round is a congestion round, is

$$(w_{c,i} - \alpha_i\tau) + (w_{c,i} - \alpha_i(\tau - 1)) + \dots + (w_{c,i} - \alpha_i) + (w_{c,i}). \quad (13)$$

During the above $\tau + 1$ rounds, flow i experiences on average $(\tau + 1)/(\tau_1 + 1)$ loss rounds. At each of its loss rounds, flow i loses on average L_i packets. Combining this with Eq. (13) gives that the average goodput of flow i , calculated from an average congestion round and the previous τ rounds, is

$$G_i = w_{c,i} - \frac{1}{2}\alpha_i\tau - \frac{L_i}{\tau_1 + 1}, \quad (14)$$

and if we ignore the L_i packets that are lost by the flow at loss rounds, the goodput becomes:

$$G_i = w_{c,i} - \frac{1}{2}\alpha_i\tau. \tag{15}$$

Step 1: Let $w_{c,i}(k)$ be the window size of flow i at congestion round k and let S_i be the number of loss rounds of flow i after a large number K of congestion rounds. Let $w_c = w_{c,1} + w_{c,2}$ be the mean total number of packets in congestion rounds. Even though w_c is a mean value, to simplify the analysis we will assume that the total number of packets in each congestion round is exactly w_c . Then, the probability that flow 1 experiences packet loss in congestion round k is $w_{c,1}(k)/w_c$. We can assume for flow 1 and for each congestion round k a binary random variable $X_1(k)$ such that

$$X_1(k) = \begin{cases} 1, & \text{flow 1 experiences packet loss in congestion round } k, \\ 0, & \text{otherwise.} \end{cases}$$

Then, the mean value of $X_1(k)$ is $w_{c,1}(k)/w_c$. The total number Ψ_1 of loss rounds of flow 1 is $\Psi_1 = \sum_k X_1(k)$ and the expected total number of loss rounds of flow 1 after K congestion rounds is $S_1 = E[\Psi_1] = \frac{w_{c,1}K}{w_c}$.

We can use an appropriate Hoeffding–Chernoff to show that for large enough K , the random variable Ψ_1 is arbitrary close to S_1 . Assume a positive constant ϵ , a probability ρ , and a number of rounds K . The following Hoeffding–Chernoff bound is from [14, Page 200, Theorem 2.3].

Theorem 6. *Let the random variables $X_i(1), X_i(2), \dots, X_i(K)$ be independent, with $0 \leq X_i(k) \leq 1$ for $k = 1, \dots, K$. Let $\Psi_i = \sum X_i(k)$ and let $S_i = E[\Psi_i]$. Then*

$$\text{For any } t \geq 0, \quad P[|\Psi_i - S_i| \geq Kt] \leq 2e^{-2Kt^2}.$$

If $K \geq \frac{\ln(\frac{1}{2\rho})w_c^2}{2\epsilon^2w_{c,1}^2}$ then by setting $t = \frac{w_{c,1}K}{w_c}$ in Theorem 6 we obtain that

$$\text{with probability at least } \rho, \quad S_1 \in \left[(1 - \epsilon)\frac{w_{c,1}K}{w_c}, (1 + \epsilon)\frac{w_{c,1}K}{w_c} \right]. \tag{16}$$

A sufficiently large number of rounds K can make the constant ϵ arbitrary small and the probability ρ arbitrary high. Since we consider the system at steady state, we can approximate $S_1 \simeq \frac{w_{c,1}K}{w_c}$ for large K . Thus, we will assume

$$S_1 = \frac{w_{c,1}K}{w_c}. \tag{17}$$

Similarly

$$S_2 = \frac{w_{c,2}K}{w_c}. \tag{18}$$

Let

$$y = \frac{w_{c,2}}{w_{c,1}}. \tag{19}$$

Then, from Eqs. (17) and (18) we get $S_2 = y \cdot S_1$. The frequency $f_{c,i}$ of loss rounds of each flow $i = 1, 2$ at congestion rounds is $f_{c,i} = (\tau + 1)/(\tau_i + 1)$. It follows that $f_{c,2}/f_{c,1} = (\tau_1 + 1)/(\tau_2 + 1)$. Additionally, note that $f_{c,2}/f_{c,1} = S_2/S_1$. Thus

$$y = \frac{\tau_1 + 1}{\tau_2 + 1}. \tag{20}$$

Step 2: Since each congestion round is a loss round either for flow 1 or for flow 2, the frequency of congestion rounds is equal to the sum of the frequencies of loss rounds of the two flows:

$$\frac{1}{\tau + 1} = \frac{1}{\tau_1 + 1} + \frac{1}{\tau_2 + 1}. \tag{21}$$

Using Eqs. (21) and (20) we get

$$\tau = \frac{\tau_1\tau_2 - 1}{\tau_1 + \tau_2 + 2} = \frac{y\tau_2 - 1}{y + 1} = \frac{\tau_1 - y}{y + 1}. \tag{22}$$

It follows

$$\tau_1 = y\tau_2 + y - 1 \quad \text{and} \quad \tau_2 = \frac{\tau_1 - y + 1}{y}. \tag{23}$$

The number of excess packets in a congestion round is a random integer in the range $1, 2, \dots, (\alpha_1 + \alpha_2)$. We will assume that the average excess is the half of the mean increase step. This gives that the average total requested window size w_c in congestion rounds is

$$w_c = C + \frac{\alpha_1 + \alpha_2}{2}. \quad (24)$$

By the definition of the toy case in each congestion round either flow 1 or flow 2 experiences packet loss. We obtain from Eq. (20) that flow 1 experiences packet loss at a congestion round with probability $1/(y + 1)$ and flow 2 with probability $y/(y + 1)$. Therefore, for $i = 1, 2$, the average number L_i of packets that flow i loses at congestion rounds is

$$L_1 = \frac{1}{y + 1} \frac{\alpha_1 + \alpha_2}{2} \quad \text{and} \quad L_2 = \frac{y}{y + 1} \frac{\alpha_1 + \alpha_2}{2}. \quad (25)$$

Step 3: Using $w_{c,1} = \frac{C + \frac{\alpha_1 + \alpha_2}{2}}{y + 1}$ and $\tau = \frac{\tau_1 - y}{y + 1}$ we can calculate τ_1 as follows:

$$\begin{aligned} G_1 &= w_{c,1} - \frac{1}{2} \alpha_1 \tau - \frac{1}{\tau_1 + 1} L_1 \\ \Rightarrow \frac{3}{2} \alpha_1 \tau_1 - \frac{1}{\tau_1 + 1} L_1 &= \frac{C + \frac{\alpha_1 + \alpha_2}{2}}{y + 1} - \frac{1}{2} \alpha_1 \frac{\tau_1 - y}{y + 1} - \frac{1}{\tau_1 + 1} L_1 \\ \Rightarrow \frac{3}{2} \alpha_1 \tau_1 &= \frac{C + \frac{\alpha_1 + \alpha_2}{2}}{y + 1} - \frac{1}{2} \alpha_1 \frac{\tau_1 - y}{y + 1}. \end{aligned} \quad (26)$$

Solving for τ_1 gives

$$\tau_1 = \frac{2C + \alpha_1 + z\alpha_1 + \alpha_1 y}{3\alpha_1 y + 4\alpha_1}. \quad (27)$$

Step 4: We will obtain one more equation for τ_1 by starting this time from the goodput G_2 of flow 2 and then proceeding as in (26):

$$\begin{aligned} G_2 &= w_{c,2} - \frac{1}{2} \alpha_2 \tau - \frac{1}{\tau_2 + 1} L_2 \\ \Rightarrow \frac{3}{2} \alpha_2 \tau_2 - \frac{1}{\tau_2 + 1} L_2 &= \frac{yC + y \frac{\alpha_1 + \alpha_2}{2}}{y + 1} - \frac{1}{2} \alpha_2 \frac{\tau_1 - y}{y + 1} - \frac{1}{\tau_2 + 1} L_2 \\ \Rightarrow \frac{3}{2} z \alpha_1 \frac{\tau_1 - y + 1}{y} &= \frac{2Cy + \alpha_1 y + \alpha_2 y}{2(y + 1)} - \frac{1}{2} \alpha_2 \frac{\tau_1 - y}{y + 1}. \end{aligned} \quad (28)$$

Solving for τ_1 gives

$$\tau_1 = \frac{2Cy^2 + \alpha_1 y^2 + 5z\alpha_1 y^2 - 3z\alpha_1}{4z\alpha_1 y + 3z\alpha_1}. \quad (29)$$

Step 5: Combining Eqs. (27) and (29) to eliminate τ_1 gives the following cubic equation of y :

$$\begin{aligned} f(y) &= (6C + 15z\alpha_1 + 3\alpha_1) \cdot y^3 + (8C + 16z\alpha_1 + 4\alpha_1) \cdot y^2 \\ &\quad - (8Cz + 16z\alpha_1 + 4z^2\alpha_1) \cdot y - (6Cz + 15z\alpha_1 + 3z^2\alpha_1) = 0. \end{aligned} \quad (30)$$

Lemma 7. For $z > 0$, Eq. (30) has exactly one positive real solution $y(z)$ and this solution is a strictly increasing function of z .

Proof. Note that $f(0) = -(6Cz + 15z\alpha_1 + 3z^2\alpha_1) < 0$ and that $f(1 + z) > 0$ (substituting $y = z + 1$ and simplifying the expression gives a sum of strictly positive terms). The derivative of $f(y)$ with respect to y is

$$f'(y) = (18C + 9\alpha_1 + 45z\alpha_1) \cdot y^2 + (16C + 8\alpha_1 + 32z\alpha_1) \cdot y - (8Cz + 16z\alpha_1 + 4z^2\alpha_1). \quad (31)$$

Consider the equation

$$f'(y) = 0. \quad (32)$$

The discriminant of Eq. (32) is

$$\Delta = (16C + 8\alpha_1 + 32z\alpha_1)^2 + 4(18C + 9\alpha_1 + 45z\alpha_1)(8Cz + 16z\alpha_1 + 4z^2\alpha_1), \quad (33)$$

which is strictly positive for $z > 0$. Consequently, Eq. (32) has two real solutions

$$y_{1,2} = \frac{-16C - 8\alpha_1 - 32z\alpha_1 \pm \sqrt{\Delta}}{2(18C + 9\alpha_1 + 45z\alpha_1)}. \quad (34)$$




Fig. 2. Analytical results for y , τ_1 and τ_2 .

From (33) and (34), it is clear that one solution is negative and one is positive. Let y_2 be the positive solution of (29). The function $f(y)$ is decreasing in the interval $(0, y_2)$ and increasing in (y_2, ∞) . Since $f(0) < 0$, it follows that $f(y)$ has exactly one positive real solution $y(z)$, as a function of $z > 0$. Since $f(y_2) < 0$, it follows that

$$y(z) > y_2. \tag{35}$$

Since, by definition, $y(z)$ is a root of function $f(y)$, after replacing y by $y(z)$ in (30), differentiating $f(y(z))$ with respect to z , and solving for the derivative $y'(z)$ we obtain

$$y'(z) = \frac{num(z)}{den(z)}, \text{ where:} \tag{36}$$

$$num(z) = -15\alpha_1(y(z))^3 - 16\alpha_1(y(z))^2 + (8C + 8z\alpha_1 + 16\alpha_1)y(z) + (6C + 15\alpha_1 + 6z\alpha_1)$$

$$den(z) = (18C + 9\alpha_1 + 45z\alpha_1)(y(z))^2 + (16C + 8\alpha_1 + 32z\alpha_1)y(z) - (8Cz + 16z\alpha_1 + 4z^2\alpha_1).$$

For the numerator $num(z)$ note that:

$$f(y(z)) + z \cdot num(z) = 3z^2\alpha_1 + 4z^2\alpha_1y(z) + 4\alpha_1(y(z))^2 + 8C(y(z))^2 + 3\alpha_1(y(z))^3 + 6C(y(z))^3 > 0. \tag{37}$$

Since $f(y(z)) = 0$ and $z > 0$, it follows that $num(z) > 0$. For the denominator $den(z)$ in Eq. (36), note that $den(z) = f'(y(z))$. Now, since $y(z) > y_2$ we get that $den(z)$ is strictly positive. From this we conclude that $y'(z)$ is strictly positive, for $z > 0$. This completes the proof. \square

Since $\alpha_2 = z\alpha_1$, Eqs. (6) and (23) imply that:

$$G_2(z) = \frac{3}{2}\alpha_2\tau_2 = \frac{3}{2}z\alpha_1 \frac{\tau_1 - y + 1}{y}. \tag{38}$$

Now we substitute τ_1 from Eq. (29) and get:

$$G_2(z) = \frac{6Cy + 3\alpha_1y + 3z\alpha_1y + 3z\alpha_1}{8y + 6}. \tag{39}$$

Lemma 8. Function $G_2(z)$ in (39) is increasing for $z > 0$.

Proof. Taking into consideration that y is a function of z , the derivative of $G_2(z)$ with respect to z is:

$$G_2'(z) = \frac{9\alpha_1 + 21\alpha_1y + 12\alpha_1y^2 + 9\alpha_1y' + 3(C - z\alpha_1)y'}{2(3 + 4y)^2}. \tag{40}$$

Since $C \geq z\alpha_1 = \alpha_2$ and $y' > 0$, the proof is complete. \square

In Eq. (38) we can also use the more accurate equation (5). In this case, numerical evaluations show that $G_2(z)$ is still an increasing function of z . However, the analytical expression becomes very complicated. Note that Eqs. (27), (29) and (30) are not influenced if we use the less accurate equation for the goodput, because the term L_i is eliminated in Eqs. (26) and (28).

We used the analytical results on a representative game instance with router capacity $C = 100$ and $n = 2$ flows. The calculations have been performed using Mathematica. In particular, we calculated the positive solution of $y(z)$ from Eq. (30). The results are presented graphically in Fig. 2a. The corresponding plots of τ_1 and τ_2 are given in Fig. 2b.

We used Eq. (5) to calculate the goodputs G_1 and G_2 of flows 1 and 2, respectively. We also used the less accurate equation (6) to calculate $G_{2,approx}$ for flow 2. The results are shown in Fig. 3a. The experimental results for the goodputs G_1 and G_2 are presented in Fig. 3b. The graphs show that the results of the analytical model are very close to the numerical and the experimental results, and confirm that the goodput $G_2(z)$ of flow 2 is an increasing function of z .

The experiments for the toy case with two flows have been conducted with a simple simulator for Window-games with basic AIMD flows, called NetKnackSimple. The router capacity was $C = 100$. To avoid or at least reduce synchronization effects the capacity varied randomly from $C - 1$ to $C + 1$ with an average of C . Parameter α was $\alpha_1 = 1$ for flow 1 and $\alpha_2 = z\alpha_1$ for flow 2. Each game lasted for 100 000 rounds.




Fig. 3. Analytical and experimental results for goodputs G_1 and G_2 as functions of z .

6.4. Prince router with gentle penalty flows

We now consider the case of a Prince router of capacity C and n flows $i = 1, \dots, n$ with parameters $(\alpha_i, \beta_i = 1/2)$. We will assume that the overflow in congestion rounds does not exceed the maximum of all windows $\max_{i=1, \dots, n} w_i$.

We assume that the system has reached steady state and we use the model from Section 6.1 to model the two flows. By the nature of the Prince mechanism, the operation of any flow i (the number τ_i of rounds between successive loss rounds of flow i and the value N_i) does not vary significantly during the course of the game.⁶ This allows us to use the average values for parameters τ_i and N_i . The error that is introduced into Eq. (6) for the goodput is very small and can be safely ignored.

Lemma 9. *Let flows i and k be the flows of minimum and maximum goodput, respectively, and let $\lambda = G_k/G_i$. Then, it holds $\lambda \leq 2$.*

Proof. From Eq. (6) we know that $G_i = \frac{3}{2}\alpha_i\tau_i$ and that the average window size of i at its loss rounds is $w_{c,i} = 2\alpha_i\tau_i$. Similarly, $G_k = \frac{3}{2}\alpha_k\tau_k$ and $w_{c,k} = 2\alpha_k\tau_k$. Note that the congestion window of flow k takes values in the range $\frac{1}{2}w_{c,k}, \dots, w_{c,k}$. Flow i has at its congestion rounds the maximum window among all flows. Thus $w_{c,i} \geq \frac{1}{2}w_{c,k} \Rightarrow \alpha_i\tau_i \geq \frac{1}{2}\alpha_k\tau_k \Rightarrow G_i \geq \frac{1}{2}G_k \Rightarrow \lambda \leq 2$. \square

Lemma 10. *In congestion rounds with total overflow at most $\max_{i=1, \dots, n} w_i$ a flow that does not exceed its fair share C/n , does not lose any packet.*

Proof. At congestion, $\sum_{i=1}^n w_i > C$. Clearly, $\max_{i=1, \dots, n} w_i > C/n$. Thus, a flow i with $w_i \leq C/n$ cannot have the maximum window. \square

We consider first a toy case with $n = 2$ flows and then the more general case with n flows.

6.4.1. Prince router with two gentle penalty AIMD flows

The Game: Assume a Prince router with capacity C and two flows $i = 1, 2$ with parameters α_1, α_2 , and $\beta_1 = \beta_2 = 1/2$. Each flow can choose its parameter α . The utility of each flow is the goodput that it achieves.

We fix α_1 and assume that $\alpha_2 > \alpha_1$. Then, let $\alpha_2 = z\alpha_1$, for $z > 1$. At steady state $G_1 = \frac{3}{2}\alpha_1\tau_1$ and $G_2 = \frac{3}{2}\alpha_2\tau_2$. Assuming λ such that $G_2 = \lambda G_1$, we get

$$z\tau_2 = \lambda\tau_1. \tag{41}$$

If $w_{c,2}$ is the average window size of flow 2 at its loss rounds, then from Eqs. (3) and (4) we know that $w_{c,2} = 2z\alpha_1\tau_2$. Let $w_{s,1}$ be the average window size of flow 1 in rounds that are congestion rounds but not loss rounds for it (flow 1). For a

⁶ At each congestion round, the Prince mechanism selects a flow with the maximum number of packets and drops packets from it. Therefore, a necessary condition for a flow i to experience packet loss at a congestion round is to have a window size larger than C/n . Consequently, the range of window sizes at which a flow i experiences loss rounds is restricted around the average value w_c and the value of parameter N_i is restricted around its average value. For the same reason, parameter τ_i is also restricted around its average value.

round to be a congestion round it must hold $w_1 + w_2 > C$, and therefore for loss rounds of flow 2 we have $w_{s,1} + w_{c,2} > C$. Using $w_{s,1} \geq \alpha_1 \tau_1$ we get $\alpha_1 \tau_1 < C - 2z\alpha_1 \tau_2$. It follows from Eq. (41) that

$$\alpha_1 \tau_1 < \frac{C}{2\lambda + 1}. \quad (42)$$

Furthermore, Prince guarantees that the greatest window size $2\alpha_1 \tau_1$ of flow 1 equals at least its fair share $C/2$, i.e., it holds $\alpha_1 \tau_1 \geq C/4$. Combining this with (42) gives

$$\lambda < \frac{3}{2}. \quad (43)$$

Note that inequality (43) holds for reasonable values of α_1 , i.e., for cases where α_1 is considerably smaller than the router capacity C .

6.4.2. Prince router with *n* gentle penalty AIMD flows

The Game: Assume a router with capacity C and n flows $i = 1, \dots, n$ with parameters $(\alpha_i, \beta_i = 1/2)$. The router uses the Prince policy for its queue. Each flow can choose its parameter α_i .

We provide a theoretical argument for symmetric profiles. Assume that flows $1, \dots, n-1$ use value α_1 for parameter α and flow n uses a value $\alpha_n = z\alpha_1$, for $z > 1$. Due to symmetry, it holds $\tau_1 = \dots = \tau_{n-1}$ and $G_1 = \dots = G_{n-1}$. Let $G_n = \lambda G_1$. Then from Eq. (6)

$$\frac{3}{2}\alpha_n \tau_n = \lambda \frac{3}{2}\alpha_1 \tau_1 \Rightarrow z\tau_n = \lambda \tau_1. \quad (44)$$

Congestion rounds occur on average every $\tau + 1$ rounds. We focus on flow n and consider the last congestion round, before flow n experiences packet loss. In this round the window size of flow n is (on average) $w_n = \alpha_n \tau_n + \alpha_n(\tau_n - (\tau + 1))$. Flow n has not the largest window size in this congestion round, therefore $w_n \leq 2\alpha_1 \tau_1$. Hence

$$\begin{aligned} \alpha_n \tau_n + \alpha_n(\tau_n - (\tau + 1)) &\leq 2\alpha_1 \tau_1 \\ \Rightarrow 2\alpha_n \tau_n &\leq 2\alpha_1 \tau_1 + z\alpha_1 \tau + z\alpha_1 \Rightarrow 2\lambda \tau_1 \leq 2\tau_1 + z\tau + z \\ \Rightarrow \lambda &\leq 1 + \frac{z(\tau + 1)}{2\tau_1}. \end{aligned} \quad (45)$$

Recall that in each congestion round, exactly one flow experiences packet loss. Thus, the sum of the frequencies of loss rounds of all flows is equal to the frequency of congestion rounds at the router:

$$\sum_{i=1}^n \frac{1}{\tau_i + 1} = \frac{1}{\tau + 1}. \quad (46)$$

Since $\tau_1 = \dots = \tau_{n-1}$ it follows from Eq. (46)

$$\begin{aligned} \frac{1}{\tau + 1} &\geq \sum_{i=1}^{n-1} \frac{1}{\tau_i + 1} = \frac{n-1}{\tau_1 + 1} \Rightarrow \tau_1 \geq (n-1)(\tau + 1) - 1 \\ \Rightarrow 2\tau_1 &\geq (n-1)(\tau + 1). \end{aligned} \quad (47)$$

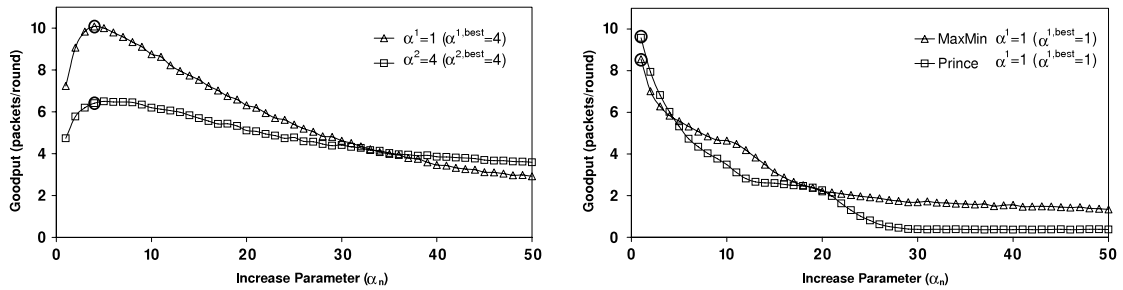
The last step holds from the condition $\tau_1 \geq 1$ of (7). Using (47) in (45) gives:

$$\lambda \leq 1 + \frac{z}{n-1}. \quad (48)$$

The above upper bound on λ shows that an aggressive flow cannot gain more goodput than a bounding factor above its fair share. The upper bound on λ converges to 1 as the number of flows n increases. The flow can increase its parameter z but this incurs costs to the flow. In Eq. (44) we used the approximate type for the goodput, which does not take into account the packet losses of the flow at loss rounds. If an aggressive flow increases its parameters z , then it experiences packet loss with a much higher frequency and hence the packet losses at loss rounds become important. We consider the above results about Prince, strong evidence that Prince will lead the system to efficient and fair Nash equilibria. The experimental results presented in Section 7 support this claim. We intend to work on a more thorough analysis of Prince.

7. Experimental evaluation for AIMD flows

We performed an extensive set of experiments with the network model of Fig. 1a, with $n = 10$ AIMD flows, a router with capacity $C = 100$, numerous queue policies and both the gentle and the hybrid penalty models. For parameter α we used values from the set $\{1, 2, \dots, 50\}$ and for β from the set $\{0.5, 0.51, \dots, 0.99\}$. Our main focus is on the results for varying parameter α when β is fixed at $\beta = 0.5$.



(a) RED router with hybrid penalty flows. The SNE is found in the second iteration.

(b) MaxMin and Prince routers with hybrid penalty flows. For both routers, the SNE is found in the first iteration.

Fig. 4. Finding SNE with methodology M1.

7.1. The methodologies

First, we applied the iterative methodology of [2], we call it M1, to find SNE. Second, thanks to the simplicity of the Window-game, we could perform a brute force search on all symmetric profiles to discover all possible SNE (methodology M2). Finally, we used random, not necessarily symmetric, starting points along with a generalization of the procedure M1 to check if the network converges to general, not necessarily symmetric, NE (methodology M3). The experiments were performed with a simulator for the Window-game, which is called NetKnack [26]. NetKnack is implemented in Java, supports common queue disciplines and flow penalty models and can perform from a simple experiment to massive series of experiments.

Methodology M1 is executed in iterations. In the first iteration, we set $\alpha^1 = 1$ for flows $1, \dots, n-1$ and search for the best response of flow n . Let $\alpha^{1,best}$ be the value α , with which n achieves the best goodput. In the next iteration, flows $1, \dots, n-1$ play with $\alpha^2 = \alpha^{1,best}$ and we search for the best α_n in this profile. If at iteration k , $\alpha^{k,best} = \alpha^k$ then this value, denoted by α_E , is the SNE of the game.

The evolution of methodology M1 is illustrated in Fig. 4a. The first iteration (α^1) shows that $\alpha^{1,best} = 4$, so at the second iteration all flows adopt $\alpha^{1,best}$ and again flow n seeks for the best response to this profile. In the last iteration, the n cannot find a better move than its current strategy and so this profile constitutes a SNE. In some occasions the initial profile is a SNE, where flow n sticks to $\alpha^{1,best} = 1$ (see Fig. 4b).

Methodology M2 applies a brute force search enumerating every possible symmetric profile to detect all SNE. The method consists of a main loop on parameter α . For each possible value k of parameter α , the startup profile for flows $1, \dots, n-1$ is set to $(\alpha_1 = \dots = \alpha_{n-1} = k)$ and measurements are made for all possible values of α_n of flow n . If the best response of flow n is the action taken by all other flows, then we have a SNE. We executed experiments with values for parameter α in the range from 2 to $C/2$.

Methodology M3 is a heuristic that extends methodology M1 to non-symmetric profiles. The first step is to create a random non-symmetric starting point with a random window size for each flow. Then, in every iteration, each flow searches independently for the best response to the current profile, assuming that the other flows do not change their strategy. The simulation reaches a (possibly non-symmetric) NE when in an iteration all flows select as a best response the same action as in the previous profile.

By convention, in all the above methodologies a flow switches to a better value of parameter α only if the average improvement of its utility is at least 2%. We focus on the results for values of parameter α in the range 1 to $C/2$.

Every experiment consists of 2200 rounds. The first 200 rounds are used to allow the flows to reach steady state. To avoid synchronization of the windows of the flows, the capacity C changes randomly with ± 1 steps, in the region 100 ± 5 with an average of 100 packets. Finally, all measurements are averaged over 30 independent executions of each experiment.

7.2. The results

We present graphs with the results of experiments where flows $i = 1, \dots, 9$ use $\alpha = 1$ and flow 10 iterates through all possible values for $\alpha_{10} \in \{1, \dots, 50\}$. The results for gentle penalty flows with different queue policies are shown in Fig. 5. We can see that Prince and MaxMin induce efficient Nash equilibria with small values for parameter α_{10} , while DropTail, RED, CHOKe and CHOKe+ lead flow 10 to use large values for α_{10} . The results for hybrid penalty flows in Fig. 6 show that with Prince and MaxMin the deviator player 10 has clearly suboptimal performance for $\alpha_{10} > 1$. Hence, the symmetric profile $\alpha_i = 1$, for $i = 1, \dots, n$, of the first iteration, is a SNE.

SNE that have been found the methodology M1 are given in Table 2a and b. We would like to note that depending on the parameters of each experiment, like the capacity C and the number of players n , the value α for the NE may differ significantly.

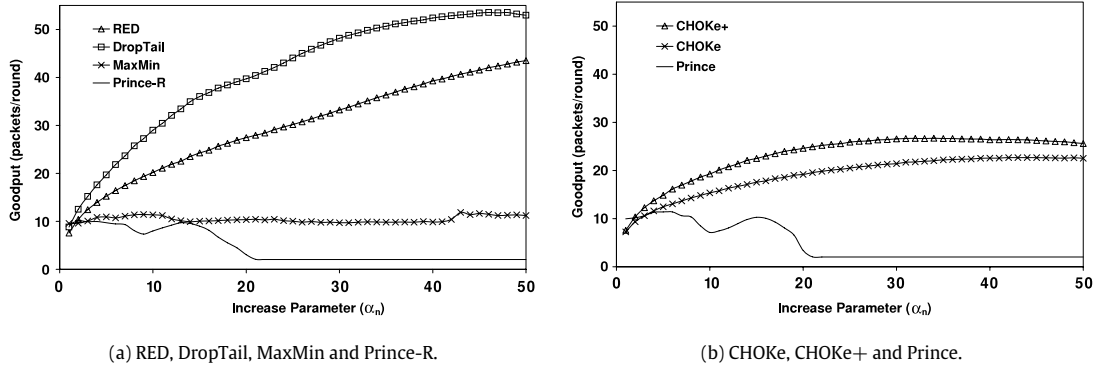


Fig. 5. First iteration of M1 with gentle penalty flows. Flows 1, 2, . . . , n - 1 use $\alpha = 1$, whereas flow n goes through all possible values of α . The graphs show the goodput of flow n for each possible value of its parameter α_n . The graphs show that the Prince-based routers discourage flow n from playing aggressively.

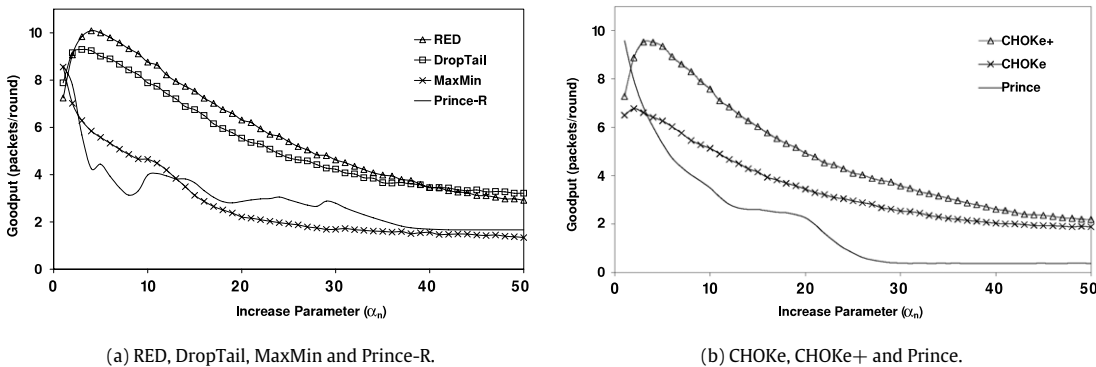


Fig. 6. First iteration of M1, this time with hybrid penalty flows. The graphs show the goodput of flow n for each possible value of its parameter α_n . Like in Fig. 5, flow n is not motivated to play aggressively.

- **Gentle penalty:** In the gentle penalty model, which corresponds to TCP SACK, a flow experiences a mild punishment upon packet loss. This fact motivates aggressive flows to use large values for parameter α (and parameter β). If, additionally, the router’s mechanism does not differentially punish greedy flows the game is driven to inefficient SNE. From Table 2a it is obvious that fully stateless policies like DropTail, RED, CHOKe and CHOKe+ are not able to aim successfully at the aggressive flow, so the game is drawn away to undesirable equilibria. On the contrary, Prince finds and penalizes only the most greedy flow, thus no flow wants to be too aggressive ($\alpha_E = 4$). When flows can vary their increase parameter α , Prince induces efficient equilibria with high goodput and tolerable loss rate. In games where the flows can choose their parameter β , all policies lead the selfish flows to use large values for β .
- **Hybrid penalty:** In the hybrid penalty model, which corresponds to TCP Reno, the cost incurred to flows by packet losses is higher. Therefore, all buffer management policies accomplish more or less to discourage flows from being too aggressive. Again, the predominance of Prince is obvious (Table 2b), having the best achieved goodput and a preferable loss rate. In this scenario, Prince outperforms even MaxMin, a fully stateful policy which keeps a separate queue for each flow. The success of Prince comes up from its targeted principle, namely punishing (dropping packets) only the largest flow, so induces minimal average loss rate. It is noteworthy that only Prince and MaxMin achieve to prevent flows from adopting high values (>0.5) for the decrease parameter β . These policies shield the AIMD algorithm because flows choose to half their window upon decrease, a fact that can lead the repeated window game to a fair equilibrium (equal window sizes).

The brute force search method M2 on DropTail and RED with gentle and hybrid flows, found the SNE that had been found with M1 and revealed additional SNE only for DropTail with hybrid penalty flows. These additional SNE use larger values of parameter α and are less efficient than the SNE found with methodology M1 for the same game.

Finally, the search for non-symmetric NE with methodology M3 for all queue policies and with both gentle and hybrid penalty flows, did not reveal any additional NE. On the contrary, it is noteworthy that this generalized methodology concluded to SNE that had already been found with M1.

8. Non-AIMD flows with complete information.

In the rest of this work we consider the more general class of Window-games where the flows can use an arbitrary strategy to choose their congestion window. The general description of Window-games with non-AIMD flows is that there

Table 2

SNE of games with gentle and hybrid penalty flows. The SNE in the cases where the flows can select parameter α or when they can select parameter β , are shown. For each SNE, the corresponding efficiency (goodput, loss rate) is given.

Queue policy	Parameter α			Parameter β		
	α_E	Goodput (packets/round)	Loss rate (%)	β_E	Goodput (packets/round)	Loss rate (%)
(a) SNE of games with gentle penalty flows.						
DropTail	$\alpha_3 = 50$	5.499	78.8	$\beta_2 = 0.97$	9.841	3.9
RED	$\alpha_2 = 49$	5.485	78.5	$\beta_2 = 0.97$	8.309	4.4
CHOKe	$\alpha_3 = 49$	3.363	86.8	$\beta_2 = 0.94$	7.761	5.4
CHOKe+	$\alpha_3 = 50$	5.431	79.1	$\beta_2 = 0.96$	8.118	4.5
Prince-R	$\alpha_2 = 2$	9.987	7.4	$\beta_3 = 0.94$	9.995	7.3
Prince	$\alpha_2 = 4$	9.111	13.9	$\beta_2 = 0.94$	9.993	7.3
MaxMin	Osc. between $a = 9$ and $a = 12$			$\beta_2 = 0.92$	9.998	4.2
(b) SNE of games with hybrid penalty flows.						
DropTail	$\alpha_5 = 6$	5.863	6.4	$\beta_2 = 0.92$	7.999	2.4
RED	$\alpha_2 = 4$	6.427	4.4	$\beta_2 = 0.96$	7.591	3.0
CHOKe	$\alpha_2 = 2$	6.182	3.6	$\beta_4 = 0.78$	6.674	2.6
CHOKe+	$\alpha_2 = 3$	6.650	3.8	$\beta_2 = 0.93$	7.687	2.9
Prince-R	$\alpha_1 = 1$	8.693	1.1	$\beta_2 = 0.72$	8.759	1.6
Prince	$\alpha_1 = 1$	9.573	2.1	$\beta_1 = 0.50$	9.617	2.1
MaxMin	$\alpha_1 = 1$	8.547	1.8	$\beta_1 = 0.50$	8.503	1.9

is a common resource, every player requests an arbitrary part of this resource and the outcome depends on the moves of all players. Note that even though this game bears some similarity with congestion games [32], it does not fit into the class of (weighted) congestion games (see for example [32,25,24,11]).

In this section, we discuss the case of games with complete information and analyze several one-shot Window-games, while in Section 9 we discuss the case of Window-games with incomplete information. A one-shot Window-game is a Window-game that lasts only one round. In a one-shot Window-game the only consequence for a flow when it loses packets is the cost g induced by each packet loss. There can be no other consequences for packet losses like window reductions or timeouts since there is only one round.

For the one-shot Window-game we assume a router of capacity C and n flows $i = 1, \dots, n$. Let w_i be the window size of flow i and $w_{total} = \sum_i w_i$ the total request of all flows. Furthermore, let each successfully transmitted packet give a profit of 1 and each lost packet incur a cost of $g \geq 0$.

8.1. DropTail

The Game: A DropTail router of capacity C and n flows. The total number of flows n is known to all flows. The game is played in one round. Each flow i chooses its window size w_i .

If there is no congestion, all packets are served. If there is congestion, i.e., $w_{total} > C$, then $w_{total} - C$ packets are randomly selected and dropped. In this case, a flow i will have on average $(w_i C) / w_{total}$ successful packets and $w_i (w_{total} - C) / w_{total}$ lost packets.

Lemma 11. *If $g = 0$, then there is a unique SNE, where each flow requests the maximum possible congestion window $w = C$.*

Proof. Since $g = 0$, the utility of any flow i for $w_{total} = \sum_k w_k$ is:

$$u(i) = \begin{cases} w_i, & \text{if } w_{total} \leq C, \\ \frac{C}{w_{total}} w_i, & \text{if } w_{total} > C. \end{cases}$$

Clearly, in both cases $w_{total} \leq C$ and $w_{total} > C$, the utility of flow i is an increasing function of w_i . Thus, there is a unique SNE, where all flows request the maximum possible value $w_i = C$. The SNE is very inefficient. \square

Lemma 12. *If $g > 0$, then there is a SNE where all flows use the window size*

$$w = \frac{C(1+g)(n-1)}{gn^2}.$$

Proof. Assume that the $n - 1$ flows $i = 1, \dots, n - 1$ use $w_i = y$ and that flow n uses $w_n = x$. In case $x + (n - 1)y < C$ we can increase x at least until $x + (n - 1)y = C$, so we will assume

$$x \geq C - (n - 1)y. \tag{49}$$

The utility of flow n is the average number of successful packets minus g times the average number of lost packets:

$$u_n(y, x) = x \frac{C}{(n-1)y + x} - gx \left(1 - \frac{C}{(n-1)y + x} \right). \tag{50}$$

Table 3

Window sizes at SNE for the one-shot game with complete information. For every queuing policy and different values of the cost factor g , i.e., variations of the utility function, the window sizes of the corresponding SNE are shown.

Queue policy	Utility function			
	$U = \text{passed} - g \cdot \text{dropped}$			
	$g = 0$	$g = 0.1$	$g = 0.5$	$g = 1.0$
DropTail	100	>60	26–30	17,18
RED	100	>70	26–28	20
CHOKe	≈44	≈35	21–22	14–16
CHOKe+	100	>70	29	17–18
Prince-R	10,11	10,11	10	10
Prince	10,11	10,11	10	10
MaxMin	≥10	10	10	10

The only positive root of the partial derivative of $u_n(y, x)$ with respect to x is

$$x = \frac{gy(1 - n) + \sqrt{Cgy(1 + g)(n - 1)}}{g} \tag{51}$$

Setting $x = y$ gives $x = y = \frac{C(1+g)(n-1)}{gn^2}$. We conclude that there is a SNE at

$$w = \frac{C(1 + g)(n - 1)}{gn^2}. \quad \square$$

For example, if $g = 1$, $C = 100$ and $n = 10$, then the SNE is at $w = 18$. Experimental results are presented in Table 3.

A legitimate question is whether there are reasonable conditions, for which the fair share $w_i = C/n$ is a SNE. Using $w \leq C/n$ in Lemma 12 and solving for g gives $g \geq n - 1$. This shows that the larger the number flows, the higher the cost g must be, in order to prevent an undesirable SNE. A sufficient condition on the cost g that does not depend on the number of players n is $g \geq C - 1 \geq n - 1$. However, when g takes such large values the game degenerates. For example, if $g \geq C - 1$, then almost any (not necessarily symmetric) profile with total window size equal to C is a NE.

8.2. MaxMin

The Game: A MaxMin router of capacity C and n flows. The number n is known to all flows. The game is played in one round. Each flow i chooses its window size w_i .

Shenker showed in [33] that a MaxMin policy enforces a fair NE in a network model where the flows are Poisson sources. In the Window-game model as well, the MaxMin policy leads the system to desirable equilibria. More precisely:

Lemma 13. *In MaxMin, there is a SNE where all flows play $w_i = C/n$. If $g > 0$, then this is the only NE of the game.*

Proof. In MaxMin, a flow i does not loose any packet if w_i does not exceed its fair share. Consequently, each flow i request at least its fair share, i.e., $w_i \geq C/n$. Thus, the total request w_{total} will be at least $w_{total} \geq C$ and no flow can gain by playing a value larger than C/n . This implies that there is a SNE where each flow i plays the strategy $w_i = C/n$. If there is a strictly positive cost $g > 0$ for each packet loss, then the above SNE is the only NE of the game. \square

Clearly, the SNE where all flows request $w_i = C/n$ is the optimal solution for the Window-game problem. As already discussed, the disadvantage of MaxMin is that it is a stateful policy that consumes too many computational resources to be applied in practice.

8.3. Prince

The Game: A Prince router of capacity C and n flows. The total number of flows n is known to all flows. The game is played in one round. Each flow i chooses its window size w_i .

The behavior of Prince is very similar to MaxMin. Every flow i can safely request its fair share C/n , so $w_i \geq C/n$. If a flow i requests $w_i > C/n$, then it may very well receive far less than its fair share (unlike MaxMin where it will receive at least its fair share). Clearly, the profile where all flows play $w_i = C/n$ is a SNE. If the cost for packet loss is $g > 0$, then this is the only NE of the game. Hence:

Lemma 14. *If $\sum_i w_i - C \leq \max_i w_i$ then there is a SNE where all flows play $w_i = C/n$. If $g > 0$, then this is the only NE of the game.*

8.4. Experimental results

Results for one-shot games with complete information are presented in Table 3. With MaxMin and both Prince policies, the equilibrium is at the fair share or almost at the fair share, even for low values of the cost factor g . For DropTail, RED, and CHOKe routers, a low value for g leads selfish players to request the maximum possible window size (aggressive behavior). The higher the cost factor g , the more conservative becomes the strategy of selfish flows.

The experiments have been conducted with the simulator NetKnack, with $n = 10$ flows and router capacity $C = 100$. The results for each experiment were calculated from the average of 40 executions. We used methodology M2 of Section 7 to identify the SNE for window sizes in the range $\{1, \dots, 100\}$.

9. Non-AIMD flows with incomplete information

In this section, we discuss Window-games in which the players do not know the total number of players n . These Window-games, where the flows are not fully informed about all the parameters of the game, belong to the class of games with incomplete information. We distinguish three cases of such Window-games.

9.1. One-shot Window-game with no prior probabilities

The Game: A router of capacity C and n players. The players have no prior probabilities on the number of players, that is, they have no distribution information for the unknown number n . The game is played in one round and each player chooses its window size.

Note that in this case, where the players do not have distribution information on the unknown parameters, the Nash equilibrium or the Bayes–Nash equilibrium concepts cannot be applied. Several concepts, like games with no prior probability or pre-Bayesian games [4,1] have been proposed to capture games with no prior probabilities. One possibility is to search for an *ex-post equilibrium*, which is a profile where each player's strategy is a best response to the other player's strategies, under all possible realizations of the uncertain data [1]. In particular, for the Window-games of this Section, each player's strategy would be a best response regardless of the number of participating players. Even though an ex-post equilibrium would be a very attractive solution, the problem is that, in general, ex post equilibria do not exist in incomplete-information games.

Other, more general, related equilibrium concepts are for example the safety-level equilibrium [4], the robust-optimization equilibrium [1], and the minimax-regret equilibrium. The common characteristic of these equilibrium concepts is that the player selects a very conservative action. In this case, the players may, for example, assume that the number of players n is the maximum possible number $n = C$. Hence, the players will play as in the game with complete information with $n = C$. For strategies like Prince and MaxMin, the choice of each player would be $w = 1$. Interestingly, in common TCP implementations, like Reno or Tahoe, flows start with a very conservative initial window size $w = 1$.

9.2. One-shot Window-game with prior probabilities

The Game: A router of capacity C and n players. The number of players n is a random variable and the players know the corresponding probability distribution. The game is played in one round and each player i chooses its window size w_i .

In practice, flows engaged in a TCP-game are likely to have some prior information on the number n or, equivalently, on their fair share. Note that real TCP flows are actually playing a repeated game, which means that the flow will in most cases have an estimation on its fair share from the previous rounds (except of the first round or a round after some serious network change). If the flow has prior information on the distribution of the unknown parameter n , the result is a Bayesian game. We leave the analysis of this and the following case as future work.

9.3. Repeated Window-game with incomplete information

The Game: A router of capacity C and n players. The number of players n is unknown and probability distribution information on n may or may not be given. Each player chooses its window size for each round of the game.

The motivation for this game comes from the fact that the actual game that a TCP flow has to play is a repeated game with incomplete information. The problem has been addressed from an optimization and an on-line algorithm point of view in [18]. One other approach would be to consider this as a learning game; there are C players and nature decides for each player either to participate in the game or to stay idle. The goal of each participating player is to “learn” the unknown number n of players who participate.

We believe that a very interesting class of games that can also be used to model this case of Window-games, is the class of *unknown games* (for example see [6, Section 7.5]). In an unknown game, K players play a game repeatedly, and at each time instance t , after taking an action, player k observes his loss (or gain) but does not know his entire loss (or gain) function and cannot observe the other players' actions. The player may not even know the number of players participating in the game.

Interestingly, even with such limited information there are conditions under which the game converges to a correlated equilibrium or even to a Nash equilibrium.

We may further define the dynamic version of this Window-game, where the number of participating players may change from round to round. In this case, in each round, nature decides with some predetermined probability for every player to switch between the idle and the non-idle state.

10. Discussion

We present a game-theoretic model for the interplay between the congestion windows of competing TCP flows. The theoretical analysis and the experimental results show that the model is relevant to the “real” TCP game. Furthermore we propose a simple queue policy, called Prince, which can achieve efficient SNE despite the presence of selfish flows.

Future work includes the completion of the analysis of Window-games with incomplete information. We also consider the proposed Prince policy of independent interest and intend to further study it. On the one hand, a thorough analysis of Prince presents an interesting problem, related to common resource allocation games. On the other hand, we plan to investigate a realistic adaptation and implementation of Prince, possibly with streaming algorithms, on real TCP networks or with the network simulator [27].

Acknowledgements

We wish to thank Vasilis Tsaoussidis for insightful discussions on TCP congestion control and Lakis Papaschinopoulos for his assistance in the proof of Lemma 7. We thank Evangelia Pyrga for bringing Reference [6] to our attention.

We are grateful to the anonymous reviewers for valuable comments that helped us to improve the presentation of this paper.

References

- [1] M. Aghassi, D. Bertsimas, Robust game theory, *Math. Program.* 107 (1) (2006) 231–273.
- [2] A. Akella, S. Seshan, R. Karp, S. Shenker, C. Papadimitriou, Selfish behavior and stability of the internet: a game-theoretic analysis of TCP, in: *Proceedings of SIGCOMM 2002*, ACM Press, New York, NY, USA, 2002, pp. 117–130.
- [3] A. Akella, S. Seshan, S. Shenker, I. Stoica, Exploring congestion control. Technical Report CMU-CS-02-139, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 2002.
- [4] I. Ashlagi, D. Monderer, M. Tennenholtz, Resource selection games with unknown number of players, in: *Proceedings of AAMAS'06*, ACM Press, New York, NY, USA, 2006, pp. 819–825.
- [5] D. Bertsekas, R. Gallager, *Data networks*. Publication New Delhi, Prentice-Hall of India Pvt. Ltd., 1987.
- [6] N. Cesa-Bianchi, G. Lugosi, *Prediction, Learning, and Games*, Cambridge University Press, 2006.
- [7] D. Chiu, R. Jain, Analysis of the increase/decrease algorithms for congestion avoidance in computer networks, *Comput. Netw. ISDN Syst.* 17 (1) (1989) 1–14.
- [8] K. Fall, S. Floyd, Simulation-based comparison of Tahoe, Reno, and SACK TCP, *Comput. Commun. Rev.* 26 (1996) 5–21.
- [9] S. Floyd, V. Jacobson, Random early detection gateways for congestion avoidance, *IEEE/ACM Trans. Netw.* 1 (4) (1993) 397–413.
- [10] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, P. Spirakis, The structure and complexity of Nash equilibria for a selfish routing game, *Theoret. Comput. Sci.* 410 (36) (2009) 3305–3326.
- [11] D. Fotakis, S. Kontogiannis, P. Spirakis, Selfish unsplittable flows, *Theoret. Comput. Sci.* 348 (2) (2005) 226–239.
- [12] X. Gao, K. Jain, L. Schulman, Fair and efficient router congestion control, in: *Proceedings of SODA'04*, Philadelphia, PA, USA, 2004, pp. 1050–1059.
- [13] R. Garg, A. Kamra, V. Khurana, A game-theoretic approach towards congestion control in communication networks, *SIGCOMM Comput. Commun. Rev.* 32 (3) (2002) 47–61.
- [14] M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, B. Reed (Eds.), *Probabilistic Methods for Algorithmic Discrete Mathematics*, Springer, 1998.
- [15] E. Hahne, Round-robin scheduling for max–min fairness in data networks, *IEEE J. Sel. Areas Commun.* 9 (7) (1991) 1024–1039.
- [16] G. Hardin, The tragedy of the commons, *Science* 162 (1968) 1243–1248.
- [17] V. Jacobson, Congestion avoidance and control, in: *Proceedings of SIGCOMM'88*, ACM, New York, NY, USA, 1988, pp. 314–329.
- [18] R. Karp, E. Koutsoupias, C. Papadimitriou, S. Shenker, Optimization problems in congestion control, in: *Proceedings of FOCS'00*, IEEE Computer Society, Washington, DC, USA, 2000, p. 66.
- [19] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, *Comput. Sci. Rev.* 3 (2) (2009) 65–69.
- [20] R. La, V. Anantharam, Window-based congestion control with heterogeneous users, in: *INFOCOM 2001*, vol. 3, IEEE, 2001, pp. 1320–1329.
- [21] L. López, G. del Rey Almansa, S. Paquelet, A. Fernández, A mathematical model for the TCP tragedy of the commons, *Theoret. Comput. Sci.* 343 (1–2) (2005) 4–26.
- [22] L. López, A. Fernández, V. Cholvi, A game theoretic comparison of TCP and digital fountain based protocols, *Comput. Netw.* 51 (12) (2007) 3413–3426.
- [23] M. Mavronicolas, P. Spirakis, The price of selfish routing, in: *Proceedings of STOC'01*, ACM, New York, NY, USA, 2001, pp. 510–519.
- [24] I. Milchtaich, Congestion games with player-specific payoff functions, *Games Econ. Behav.* 13 (1996) 111–124.
- [25] D. Monderer, L. Shapley, Potential games, *Games Econ. Behav.* 14 (1996) 124–143.
- [26] NetKnack. A simulator for the Window-game. <http://utopia.duth.gr/~pefraimi/projects/NetKnack>.
- [27] NS-2. The network simulator. <http://www.isi.edu/nsnam/ns/>.
- [28] J. Padhye, V. Firoiu, D. Towsley, J. Kurose, Modeling TCP Reno performance: a simple model and its empirical validation, *IEEE/ACM Trans. Netw.* 8 (2) (2000) 133–145.
- [29] R. Pan, B. Prabhakar, K. Psounis, CHOC, a stateless active queue management scheme for approximating fair bandwidth allocation, in: *INFOCOM, 2000*, pp. 942–951.
- [30] C. Papadimitriou, Algorithms, games, and the internet, in: *Proceedings of STOC'01*, ACM Press, New York, NY, USA, 2001, pp. 749–753.
- [31] L.L. Peterson, B.S. Davie, *Computer Networks: A Systems Approach*, fourth edition, Morgan Kaufmann, 2007.
- [32] R. Rosenthal, A class of games possessing pure-strategy Nash equilibria, *International Journal of Game Theory* 2 (1973) 65–67.
- [33] S.J. Shenker, Making greed work in networks: a game-theoretic analysis of switch service disciplines, *IEEE/ACM Trans. Netw.* 3 (6) (1995) 819–831.
- [34] W.R. Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, 1994.
- [35] I. Stoica, S. Shenker, H. Zhang, Core-stateless fair queueing: achieving approximately fair bandwidth allocations in high speed networks, *SIGCOMM Comput. Commun. Rev.* 28 (4) (1998) 118–130.
- [36] I. Stoica, H. Zhang, Providing guaranteed services without per flow management, *SIGCOMM Comput. Commun. Rev.* 29 (4) (1999) 81–94.