



A mathematical model for the TCP Tragedy of the Commons[☆]

Luis López^{a,*}, Gemma del Rey Almansa^a, Stéphane Paquelet^b,
Antonio Fernández^a

^aLADyR, GSyC, Universidad Rey Juan Carlos, C/Tulipán S/N, 28933 Móstoles, Madrid, Spain

^bMitsubishi ITE, 1 Avenue de Beaulieu 35000, Rennes, France

Abstract

This paper presents a novel mathematical model for the TCP Tragedy of the Commons, using Game Theory concepts. This tragedy may appear in a TCP/IP-based network when hosts do not respect the protocol rules and try to monopolize the shared network resources by using a selfish strategy. Our model quantifies the effects of this evil behavior in a simple and standard network topology and allows to obtain some interesting results which we prove formally. Finally, we validate the model results by comparing its predictions with a set of extensive simulations carried out using the NS Network Simulator.

© 2005 Elsevier B.V. All rights reserved.

Keywords: TCP; Game Theory; Tragedy of the Commons

1. Introduction

Many important communication systems in our days are based on the principle of sharing a common resource among different users. Some of the most relevant examples are the Aloha protocol [1], where the shared resource is a radio link bandwidth, the original Ethernet system [9], inspired in Aloha, in which all the emitters share the communication capacity

[☆] Partially supported by the Spanish MCyT under grant TIC2001-1586-C03-01, the Comunidad de Madrid under grant 07T/0022/2003, and the Universidad Rey Juan Carlos under grant PPR-2004-42.

* Corresponding author. Tel.: + 34 91 488 81 31.

E-mail address: luis.lopez@urjc.es (L. López).

of a wire, and the conventional TCP/IP networking scheme [19,16], where all the users in a route share the same communication links and buffering capabilities of the network routers. One of the main objectives of the communication protocols in these systems is to establish a number of rules guaranteeing that the common resources are fairly shared among all the users. When a protocol is well designed and all the competitors respect its rules, usually the communication capabilities are adequately split. Nevertheless, when users do not cooperate and present evil conducts by not respecting the protocol rules, it is possible that unfair or unstable behaviors emerge in the system. One of the main problems in relation with this concern is that networks and protocols usually provide no means for guaranteeing or verifying whether users respect the rules or cheat.

Many authors have already noticed that the conventional TCP end-to-end congestion control scheme is voluntary in nature and critically depends on end-user cooperation [7,5,2]. This is mainly due to the fact that TCP was originally designed in the context of an academic network where all users tend to collaborate. The fundamental objective of TCP was to supply a reliable delivery of packets between two applications in different hosts. Then, the congestion control mechanism used by the protocol is mainly based on two principles, which might be roughly sketched as follows:

- (1) Each time a packet is lost in a connection, TCP reduces the sending rate of that connection (approximately to half).
 - (2) TCP tries to avoid the duplication of packets in the network. That is, TCP does not repeat a transmitted packet until it is (almost) sure that the packet has been actually lost.
- When all hosts comply with these principles, TCP guarantees that the communication resources are efficiently used and fairly shared. Nevertheless, both rules are based on a voluntary acceptance of the rules by the users, without any verification from the network. Nothing prevents a particular user from presenting an evil behavior, and maintain or even increase its sending rate if there are packet losses, or to repeat a particular packet as many times as it wishes until the appropriate acknowledgment has been received.

The problem that arises in this situation is that, as some authors have already observed [11,10,2], most Internet routers use a drop-tail FIFO (First In First Out) scheduling discipline. Then, each host can obtain more network bandwidth by transmitting more packets per unit of time, because this policy gives the most resources to the sender of the most packets. Nevertheless, if the hosts overload the network, it may be demonstrated that the total throughput drops [11]. Thus, the optimal strategy for each host is strongly suboptimal for the network as a whole. Using a Game-Theory view [6,13], these problems arising in TCP/IP networks are strongly related to the stability of multi-player games. In particular, the well-known Prisoner's Dilemma problem is an example of a game with this property. But a closer analogy is the, so-called, Tragedy of the Commons [8] problem in economics, where each individual can improve her own position by using more of a free resource, but the total amount of the resource degrades as the number of greedy users increases. Historically, this analysis was applied to the use of common grazing lands; it also applied to such diverse resources as air quality and time-sharing systems. In general, experience indicates that multiple-player systems with this type of instability tend to get into serious trouble.

To understand precisely what a Tragedy of the Commons is, we first need to observe that, in the context of Game Theory, players choose their strategy in a selfish way trying to maximize their benefit (yield). If the system gets into a state where no player has an incentive

to unilaterally change its strategy we say that the system has reached a Nash equilibrium. In this context, a game is a Tragedy of the Commons when the following two conditions hold:

- (1) Condition 1. There is always an incentive for a new host to become evil. This guarantees that the Nash equilibrium is reached when all hosts are evil.
- (2) Condition 2. The final yield for evil hosts in the Nash equilibrium is under the initial yield of fair hosts when all players collaborate.

These conditions guarantee the essential ingredients of a Tragedy of the Commons: if players behave in a selfish way, the Nash equilibrium will be reached, and hence, the yield of the defectors will always be less than the initial reward of the fair player. Hence, all players lose.

This problem of the TCP protocol has already been addressed in the networking literature using a game-theoretic perspective. Some of the most remarkable works in this field are the ones carried out by Nagle [11,10] and Garg et al. [7]. Both of them propose a solution based on creating incentive structures in the systems that discourage evil behavior. The first one suggests replacing the single FIFO queue associated with each outgoing link with multiple queues, one for each source host, which are served in a round-robin fashion. The second introduces a novel and sophisticated sample service discipline called rate inverse scheduling (RIS) that punishes evil behavior and rewards cooperation, in such a way that the resulting Nash equilibrium leads to fair allocation of resources. Both solutions require a (sometimes huge) per-packet processing, which might be impractical in many realistic applications (in Internet core routers, for example). There are other interesting proposals related to problems similar to this [21,18,17,15,4], and in all cases, these works show the potential applications of Game Theory within the problem of congestion control and routing in packet networks.

As it can be observed, the available networking literature on this topic mainly concentrates in proposing possible solutions to these problems, but we believe that there is a lack of appropriate mathematical models allowing a better understanding of the TCP Tragedy of the Commons, and a quantification of its effects. One of the few works in this direction is the one by Nagle [11], who showed this tragedy to be really tragic by proving that a packet network with infinite storage, FIFO scheduling, and a finite packet lifetime will, under overload, drop all packets. This is a very remarkable work, but it is based on somewhat very restrictive assumptions, which drive the model to highly pessimistic results. Besides this, it only studies the Nash equilibrium state, but it does not explicitly contemplate the effect of the tragedy when only a fraction of the users present evil behavior.

More recently, an excellent analysis of TCP behavior in the context of Game Theory has been carried in an attempt by Akella et al. [2]. In this work, a combination of analyses and simulations is carried in an attempt to characterize the performance of TCP in the presence of selfish users. The conclusions obtained in this work are similar to the ones presented here. Besides, their analyses cover different variations of TCP (Reno, SACK, etc.) and buffer management policies (Drop Tail, RED, etc.) They even propose a stateless solution based on the CHOKe scheduling policy [14], which can minimize the impact of the tragedy. Nevertheless, again their interest is concentrated exclusively on the Nash equilibrium itself. The model we present here is more simple and, thanks to this, it allows to analyze all the intermediate states of the network until we get to the final equilibrium in which the Tragedy of the Commons may be found. Moreover, for all these intermediate states we offer an

estimation of both evil and fair throughputs. This allows the computation of the incentive to become evil for any host in any given state of the game.

In this direction, our objective in this paper is to develop a novel mathematical model to quantify the effect of the Tragedy of the Commons in the TCP protocol. With this purpose, we use an approach completely based on Game Theory, in an attempt to analyze the problem as an N -player game, where N rivals compete for a limited resource. Then, we define a model of a system with a very simple network and a streamlined version of a protocol similar in spirit to TCP. We analyze the behavior of this system, and get to obtain the expression of the effective transmission rate for each kind of player, fair and evil. Then, from this we obtain conditions to have a Tragedy of the Commons. We do this first for a simplified version of the system, in which we assume that there will never be duplicated packets. Then, we do it for a more realistic model in which we also consider the effect of duplicated packets in the system. Finally, we perform simulations of a system similar to the one proposed but with the real TCP protocol (used by fair players) and a modified selfish version of it (used by evil players). We show that the simulation results are qualitative, the same as those of the analysis.

The rest of the paper is structured as follows: In Section 2 we present the theoretical model of the system we will analyze, and obtain some basic parameters of this model. In Section 3 we conduct a first simple analysis without duplicated packets. In Section 4 we conduct a more complete analysis with duplicated packets. In Section 5 we present the results of the simulations and compare them with those of the analyses. Finally, in Section 6 we present concluding remarks.

2. The model

The model we propose in this paper is based on the simplest network topology we could imagine, which is the one depicted in Fig. 1. In this network, we have N hosts (the players), which compete for sending packets to a remote host, which we call the *Destination*. Packets from the N players are routed through a single switching device, which we call the *Router*. This Router is assumed to have a finite queue of size m packets, which is scheduled using a FIFO discipline. The buffer management policy is drop tail, which means that when the queue is full, all new incoming packets are dropped. Hence, all competitors share the same buffering capabilities and the same communication link to reach the Destination.

Once we have described the interconnection topology, we may introduce the network dynamics of our model, which is somehow inspired in the one proposed by Borodin et al. [3]. In this direction, we assume the network to be a directed graph where all links have the same capacity. Time proceeds in discrete steps, so that packets travel atomically consuming one time step to traverse any of the links. We consider that, when traversing a link, a packet leaves the node at one end of the link at the beginning of a step and arrives to the node at the other end of the link at the end of that step. Then, a buffer position may be occupied by one packet (outgoing) at the beginning of a time slot and by another different packet (incoming) at its end. One novelty we introduce is that, given that the buffer size is limited to m packets and the queue policy is fixed to drop tail, all packets trying to access the buffer in a particular time slot have the same probability of entering into it. In the simple case

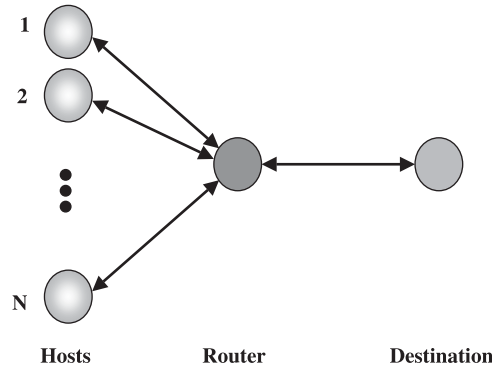


Fig. 1. The interconnecting topology of the network we use in the proposed model. N hosts try to access a common communication link through a router. This router has a finite buffer which drops packets when it is full. All lines have a capacity of one packet per time slot (or step).

where $m = 1$, this means that if n packets try to access this buffer in a particular time step, the probability of any of them to occupy it is $1/n$. As none of the hosts is privileged, we say that our queuing policy is fair.

Modeling a full-featured TCP implementation is far beyond the scope of this work. Nevertheless, for our model, we assume a simplified version, which captures its essential aspects and behaviors. In this direction, we assume that our hosts implement a trivial TCP (tTCP) protocol, fulfilling two main postulates:

- (1) The tTCP protocol optimizes the use of the common resource for reliable transmission of packets.
- (2) The common resource is fairly shared among all the players (the hosts).

We are not interested in the mechanism used by the protocol to achieve these objectives. We just accept that, in a steady state, it manages to find, for each host, the appropriate input rates and time schedules to accomplish the two postulates. Observe that a real-world TCP implementation also has these two objectives, in the sense that it tries to order the behavior of the hosts to obtain the maximum profit in a communication network while equitably sharing the resources among all the users. A very important point is that we make our tTCP protocol reliable by assuming that a host is always sending the same packet until it properly arrives at its destination. For simplicity, we do not contemplate any explicit acknowledgment mechanism; we just assume that when the destination receives a packet, the originator of that packet is informed automatically. Although this may seem to be a very naive assumption, it may be showed that it does not introduce any remarkable modification on the behavior of the hosts. These simple mechanisms allow us to claim that our model is connection oriented, because by using them, it is assured that all packets will, sooner or later, arrive at their destinations and they will do so in the correct order.

Once we know all the details about the proposed model, we may define the possible strategies that a player may adopt to participate in this game. For simplicity, we assume that a host has only two possibilities: to be fair and collaborate, or to be evil and cheat. A fair player is one complying with the rules imposed by the tTCP protocol we have defined. On the other hand, an evil player is one following a selfish strategy, and hence breaking these

rules. In order to be able to perform a mathematical analysis of the model, we adopt a more precise definition of these terms. The following paragraphs are devoted to this.

To define the fair behavior, consider a situation where all the players cooperate and respect the tTCP protocol. Following our initial hypothesis, after a transient period, the tTCP internal mechanism guarantees that the resources are being used optimally and fairly among all the users. In this situation, it is easy to understand that the most optimal and fair solution is to share the network capabilities using a round-robin cyclic service algorithm. In this case, the proposed protocol divides the available time into a sequence of groups of N slots. Each player owns one exclusively reserved time slot in each of these groups to send its packet. Once a packet has been sent, the host must wait $N - 1$ time slots to be able to send another. We assume fair players are greedy in the sense that they always use their time slot to send a packet. With these ideas, it is possible to formally define the behavior of a fair player in the following way. Given a host $i \in \{1, \dots, N\}$, the tTCP protocol assigns to it an exclusive time slot $s_i \in \{1, \dots, N\}$ in each round (the same for all rounds). We say that this host is fair when the probability of host i sending a packet in slot j of a round is

$$p_i(j) = \begin{cases} 1 & \text{if } j = s_i, \\ 0 & \text{if } j \neq s_i. \end{cases} \quad (1)$$

That is, we assume that fair players use only one out of the N available time slots in a round, independent of whether any of the other $N - 1$ remaining players presents evil behavior.

On the other hand, we have evil hosts, which do not respect the protocol rules. We may define the evil behavior in many different ways, but one of the simplest is the following. Given a host $i \in \{1, \dots, N\}$, we say this host is p_i -evil when it forgets about the time slot that the protocol has assigned to it and sends a packet in any time slot with probability p_i . That is,

$$p_i(j) = p_i, \quad \text{for all } j. \quad (2)$$

For simplicity, we assume that all evil hosts share the same probability p . Hence, $p_i = p$ for all i . Without loss of generality, we may assume that N_e hosts p -evil behavior and N_f exhibit fair behavior. Since all hosts have only one of these two possibilities, it must be satisfied that $N_e + N_f = N$.

Once we have defined the set of possible strategies in the network, our objective is to evaluate the transmission rate for any of the hosts as a function of its own strategy and of the strategy of the other hosts. The transmission rate measures the average number of useful packets that a given host sends on a time slot. Hence, it is measured in *packets/slot*. We will denote by $R_e(N, N_e, p)$ the transmission rate for evil hosts in a network of N hosts, when N_e of them exhibit p -evil behavior. In the same way, we will denote by $R_f(N, N_e, p)$ the transmission rate of fair hosts in a network of N elements when N_e of them exhibit p -evil behavior.

To evaluate $R_e(N, N_e, p)$, it is necessary to introduce some new concepts. Note that our tTCP protocol divides the available time in a round into N different slots. Observe that a particular time slot has only two possibilities:

- (1) To be assigned to a host that behaves with evil behavior (*E*-slot).
- (2) To be assigned to a host that behaves with fair behavior (*F*-slot).

Without loss of generality, we assume the simple case where the buffer size $m = 1$. (Observe that in terms of accessing probabilities, this is equivalent to a situation where the buffer is full and on each time slot there is only a single free buffer position that can be accessed by the hosts.) Thus, in an E -slot, the players competing for the shared buffer position are N_e evil hosts. Remember that we have assumed that evil hosts send a packet with probability p and that the queue is fair. Hence, given a particular evil host, this host has probability 1 of success when it is the only one sending a packet in that slot, an event that occurs with probability $p(1-p)^{N_e-1}$. In the same way, the host has probability $1/2$ of succeeding when only another host and itself are sending packets, which occurs with probability $p\binom{N_e-1}{1}p(1-p)^{N_e-2}$. We could continue with this reasoning until we get to the situation where all evil players send a packet ($p\binom{N_e-1}{N_e-1}p^{N_e-1}$). In this case, the probability of success is $1/N_e$. Hence, we can conclude that the probability of a particular evil host to really transmit a packet through the shared resource in an E -slot is the sum of all these different contributions

$$p_E^e(N_e, p) = \sum_{i=1}^{N_e} \frac{1}{i} \binom{N_e-1}{i-1} p^i (1-p)^{N_e-i}. \quad (3)$$

This sum can be easily evaluated using the binomial theorem just by observing that $\frac{1}{i} \binom{N_e-1}{i-1} = \frac{1}{N_e} \binom{N_e}{i}$

$$p_E^e(N_e, p) = \frac{1 - (1-p)^{N_e}}{N_e}. \quad (4)$$

Following a similar argument, it is possible to obtain an analogous expression for an F -slot. The main difference is that, in this kind of slot, there is, with probability 1, a packet coming from the fair host assigned to that slot. Then, it can be easily shown that the probability of an evil host to really transmit a packet on an F -slot is given by

$$p_F^e(N_e, p) = \sum_{i=1}^{N_e} \frac{1}{i+1} \binom{N_e-1}{i-1} p^i (1-p)^{N_e-i}. \quad (5)$$

This expression can also be evaluated using a procedure similar to the one of Eq. (4) to a value of

$$p_F^e(N_e, p) = \frac{(N_e+1)p - 1 + (1-p)^{N_e+1}}{N_e(N_e+1)p}. \quad (6)$$

We may proceed in a similar way to evaluate the transmission probability for a fair host. In this case, the analysis is easier because this type of host can only transmit in a particular F -slot which is assigned to it by the protocol. The number of packets competing to enter the shared buffer depends on the behavior of the N_e players also trying to access the common resource. It may be easily shown that, in this situation, the probability of the fair host to really transmit its packet in its F -slot is

$$p_F^f(N_e, p) = \sum_{i=0}^{N_e} \frac{1}{i+1} \binom{N_e}{i} p^i (1-p)^{N_e-i}. \quad (7)$$

Like in the preceding cases, this sum can also be simplified to a value

$$p_F^f(N_e, p) = \frac{1 - (1 - p)^{N_e+1}}{(N_e + 1)p}. \quad (8)$$

These probabilities are essential for the evaluation of the transmission rates. Nevertheless, to accomplish this task, it is also necessary to define an appropriate framework describing the behavior of the hosts in terms of a set of states characterizing the evolution of packets within the network. The following section has the objective of introducing the simplest way of doing it.

3. The stateless approach

In this section we propose an extremely simplifying assumption which consists in considering that the network is stateless. This assumption may be justified in a situation where all the packets entering the router buffer are useful (that is, duplicates are not allowed). As this network must have very small latency, we may assume that the buffer size is fixed to $m = 1$. Although these assumptions may seem too simplistic, they allow us to introduce some important concepts that are essential for the understanding of the rest of the paper.

In this situation, given that the number of slots in a round is N and that there are exactly N_e type E slots and $N - N_e$ type F slots in the round, we may prove that the average transmission rate for an evil host can be expressed in the following way:

$$R_e(N, N_e, p) = \frac{N_e p_E^e(N_e, p) + (N - N_e) p_F^e(N_e, p)}{N}. \quad (9)$$

We may proceed in a similar way for evaluating the transmission rate for fair hosts. In this case, there is only one single slot in the round where the host may send a packet. Hence, it is straightforward that the average transmission rate for a fair host is

$$R_f(N, N_e, p) = \frac{p_F^f(N_e, p)}{N}. \quad (10)$$

In the framework of the Game Theory, the preferred method for quantifying the effects of non-cooperation in N -player games is by evaluating the reward as a function of the number of defectors. In our case, these rewards are the transmission rates R_e and R_f , which indeed depend on the number of defectors N_e but also on the probability p associated with the evil hosts. It can be easily shown that different types of games are possible in our model depending on the value of p . We are especially interested in the case where the game takes the form of a Tragedy of the Commons. We have the following theorem.

Theorem 1. *In the stateless model, $p \in (\frac{2}{N+1}, 1)$ is a necessary and sufficient condition for the game to be a Tragedy of the Commons.*

Proof. To prove the claim, observe that the following must hold:

- (1) Condition 1. $R_e(N, n + 1, p) > R_f(N, n, p)$ for all $n \in \{0, \dots, N - 1\}$. (The Nash equilibrium is reached when all hosts are evil.)

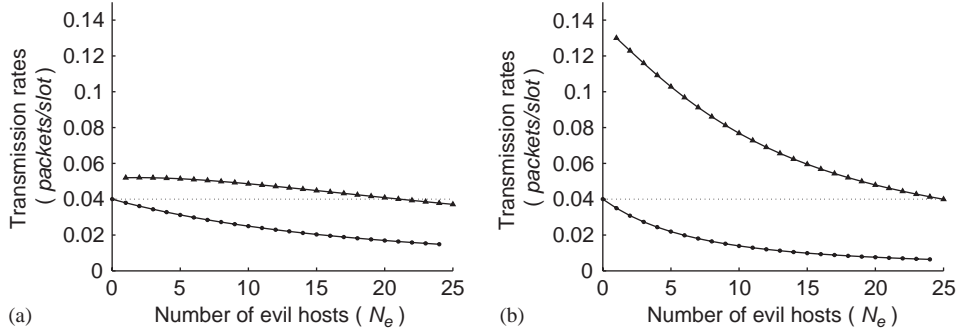


Fig. 2. The transmission rates for the stateless model as a function of the number of evil hosts, N_e , in a network with $N = 25$ players. (a) has been evaluated using an evil probability $p = 0.1$. (b) $p = 0.25$. The line with solid triangles represents the evil rate $R_e(N, N_e, p)$. The line with solid circles represents the fair rate $R_f(N, N_e, p)$. The dotted line indicates the yield of the all-fair situation, $R_f(N, 0, p)$. As can be observed, in both cases the game is a Tragedy of the Commons because there is always an incentive for a new player to become evil and the final rate of evil hosts in the Nash equilibrium, $R_e(N, N, p)$, is always under the initial yield of fair players $R_f(N, 0, p)$.

(2) Condition 2. $R_f(N, 0, p) > R_e(N, N, p)$. (The reward in the Nash equilibrium is below the yield obtained when all hosts are fair.)

Then, we split the proof into the following three claims.

Claim 2. *In the stateless model, $p \in (\frac{2}{N+1}, 1)$ is a necessary condition for the game to be a Tragedy of the Commons.*

Proof. If the game is a Tragedy of the Commons, Condition 1 above must hold. Hence, $R_f(N, n, p) < R_e(N, n+1, p)$ for all $n \in \{0, \dots, N-1\}$. Then, $R_f(N, 0, p) < R_e(N, 1, p)$. Using Eqs. (3), (5), (7), (9), (10), we may write the inequality as

$$\frac{1}{N} < \frac{p + (N-1)\frac{1}{2}p}{N}, \quad (11)$$

which in turn is equivalent to $p > \frac{2}{N+1}$. Now, observe that if $p = 1$, Condition 2 would not be satisfied. The claim is proved. \square

Claim 3. *In the stateless model, $p > \frac{2}{N+1}$ is a sufficient condition for Condition 1 to hold.*

Proof. Let us assume that $p > \frac{2}{N+1}$. Then, it follows that $p > \frac{i+1}{Ni+n+1}$ for all $i \in \{1, \dots, N\}$ and for all $n \in \{0, \dots, N\}$. This can be seen by assuming a worst case ($n = 0$), and by observing that the inequality holds for $i = 1$ and that the expression on the right strictly decreases with i . The inequality can also be written as

$$\frac{n+1}{i} + \frac{N-n-1}{i+1} > \frac{1}{ip}. \quad (12)$$

Now, we define $f_i = \binom{n}{i-1} p^i (1-p)^{n+1-i}$. Observe that f_i is always positive. In this situation, we can multiply Eq. (12) by f_i without changing the inequality. Hence, summing all the inequalities for all i yields

$$\sum_{i=1}^{n+1} \frac{n+1}{i} f_i + \frac{N-n-1}{i+1} f_i > \sum_{i=1}^{n+1} \frac{1}{ip} f_i. \quad (13)$$

Observe that substituting f_i , making a change of variables in the second part of the inequality, dividing by N and recovering the original expressions of p_E^e , p_F^e and p_F^f from Eqs. (3), (5), (7) yield

$$\frac{(n+1)p_E^e(n+1, p) + (N-n-1)p_F^e(n+1, p)}{N} > \frac{p_F^f(n, p)}{N}, \quad (14)$$

which using Eq. (9) is equivalent to $R_e(N, n+1, p) > R_f(N, n, p)$. \square

Claim 4. *In the stateless model, $p < 1$ is a sufficient condition for Condition 2 to hold.*

Proof. The proof is immediate using Eqs. (9), (10). Observe that $R_f(N, 0, p) = \frac{1}{N}$ and $R_e(N, N, p) = \frac{1-(1-p)^N}{N} = \frac{1}{N} - \frac{(1-p)^N}{N}$. As we assume $p < 1$, it clearly follows that $R_f(N, 0, p) < R_e(N, 1, p)$. Then, Condition 2 is also satisfied. \square

The combination of the above three claims trivially proves the theorem. \square

The Tragedy of the Commons can be observed graphically in Fig. 2, where we show the evolution of R_e and R_f as a function of the number of defectors, N_e , in a game with $N = 25$ players. Then, we may conclude that our simple model presents a Tragedy of the Commons phenomenon. Nevertheless, there is an aspect of these results that requires a more exhaustive analysis. From Fig. 3, it is possible to observe that, in the Nash equilibrium, the $R_e(N, N, p)$ increases when p grows. Moreover, we may easily prove that when $p = 1$, $R_e(N, N, p) = R_f(N, 0, p) = 1/N$. This means that, by increasing the greediness of the evil players, it is possible to minimize the tragedy and even make it disappear.

This conclusion is somehow unnatural, because it suggests that the proposed communication protocol inspired in TCP is as efficient as a totally disordered situation based only on the absolute selfishness of the users. If these results were true in the real Internet, there would be no need for any congestion control mechanism other than encouraging the users to send as many packets as possible. This odd behavior appears because when we assumed a stateless approach, and withdrew a very important effect that is present in real networks but which is not exhibited here. As some authors have already remarked [11], one of the main problems in packet networks is the presence of duplicated packets. The stateless model does not contemplate this possibility. For this reason, we are going to improve the model by assuming a more realistic situation, assuming that large buffers are possible and that the acknowledgment of the packets only occurs when they arrive at their final destination. The following section is devoted to this objective.

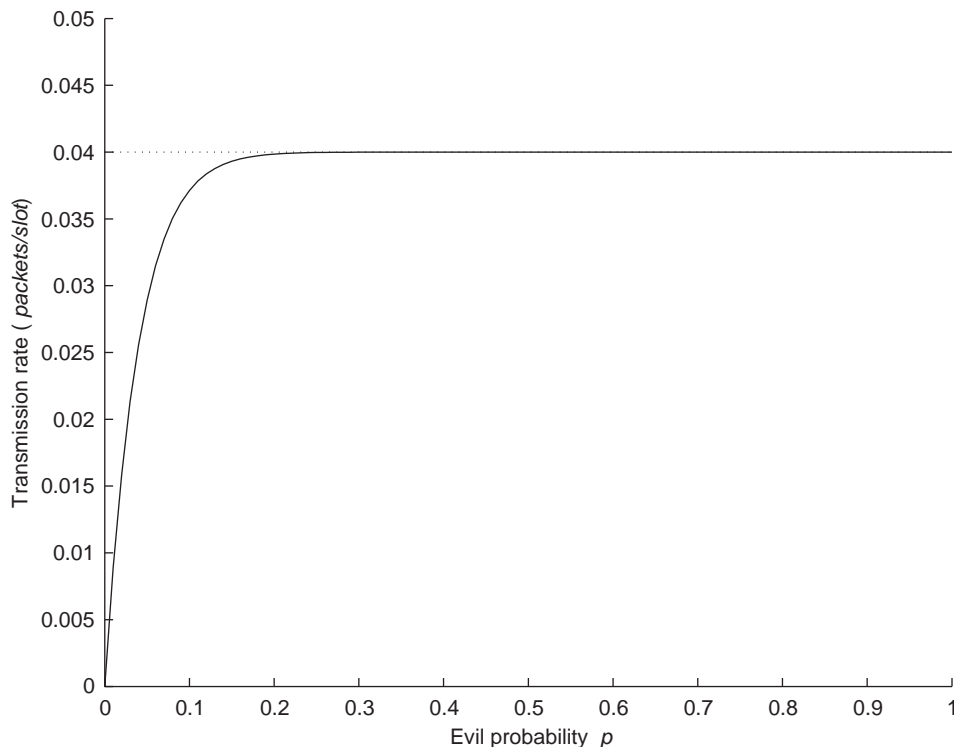


Fig. 3. Increase in the final evil rate in the Nash equilibrium, $R_e(N, N, p)$, with p . The dashed line represents the initial fair rate $R_f(N, 0, p)$. The simulation has been performed in a network of $N = 25$ hosts.

4. The Markovian approach

In this section, we consider a modification of the model presented previously and assume that the router buffer has a drop-tail queue of m packets, where m may be different from 1. In this situation, it is necessary to track the position of the packets within the buffer for both evil and fair hosts. If we start with fair hosts, we note that these players send only one packet in each round of N slots. In order to simplify the analyses, we formulate a reasonable hypothesis and assume that $N \geq m$ (it may be shown that for $N < m$ we obtain similar conclusions but the analyses becomes more complicated). Hence, fair hosts never have duplicates, because the time between two successive packet injections is smaller than their maximum lifetime in the network. For this reason, we may claim that all the packets from fair hosts that get to the router buffer are useful. This means that the new m -buffer transmission rate $R_f(N, N_e, p, m)$ is equal to the one evaluated in Eq. (10).

Nevertheless, when we consider evil hosts, things are very different because they try to access the buffer in all slots with probability p . Hence, the presence of duplicate packets is possible and not all the packets accessing the router buffer contain useful information. To

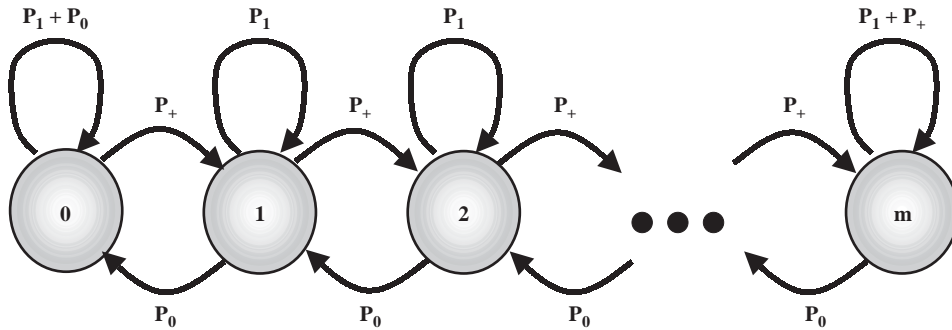


Fig. 4. The Markov Chain proposed to model the occupation of the router buffer. State i corresponds to a situation where there are exactly i packets in the buffer, p_0 is the probability of 0 packets being issued by hosts in a given time slot, p_1 is the probability of 1 packet being issued and p_+ is the probability of more than one packet trying to get into the buffer in that time slot. As an approximation, we consider that the buffer occupation can only increase in steps of one packet. The probability of the buffer being full corresponds to the probability of state m .

be able to mathematically analyze this situation we need to establish a number of assumptions. First, we assume that the network has reached its steady state. In this situation, we concentrate our interest in the case where congestion occurs and the buffer queue is always full. Hence, in each time slot, only one buffer position is liberated when the first packet is served and the rest advance one position.

Note that the assumption of permanent congestion is not necessarily true in any possible scenario. Nevertheless, in a Game Theoretic perspective the congested state is the most interesting because the presence of selfish greedy users always drives the system to congestion, and the congested state will be the most common state in which the system exists. To conclude this, we need to formulate some novel definitions. Let us denote the probability that 0 new packets try to access the router buffer on a given time slot by p_0 , the probability of exactly 1 packet doing so by p_1 , and the probability of more than one packet doing so by p_+ . To simplify the analyses we assume that the distribution of fair and evil slots in a given round is not fixed, and they occur randomly with probabilities $\frac{N_f}{N}$ and $\frac{N_e}{N}$, respectively. In these conditions, it is easy to show that

$$p_0 = \frac{N_e}{N}(1-p)^{N_e}, \quad p_1 = \frac{N_f}{N}(1-p)^{N_e} + \frac{N_e}{N}N_e p(1-p)^{N_e-1} \quad \text{and} \\ p_+ = 1 - p_0 - p_1.$$

Then, the dynamics of the buffer occupation is driven by these probabilities. Observe that, in any given slot, the number of packets in the buffer decreases (if it is not empty) with probability p_0 , remains constant with probability p_1 and increases (if it is not full) with probability p_+ . As a first approximation, we will consider that the buffer occupation can only increase in steps of one packet. In that case, the buffer behavior may be described by the Markov chain presented in Fig. 4. Observe that, in reality, the number of packets in the buffer might increase more rapidly than described by this chain, because of our approximation, but that it cannot decrease at a higher rate.

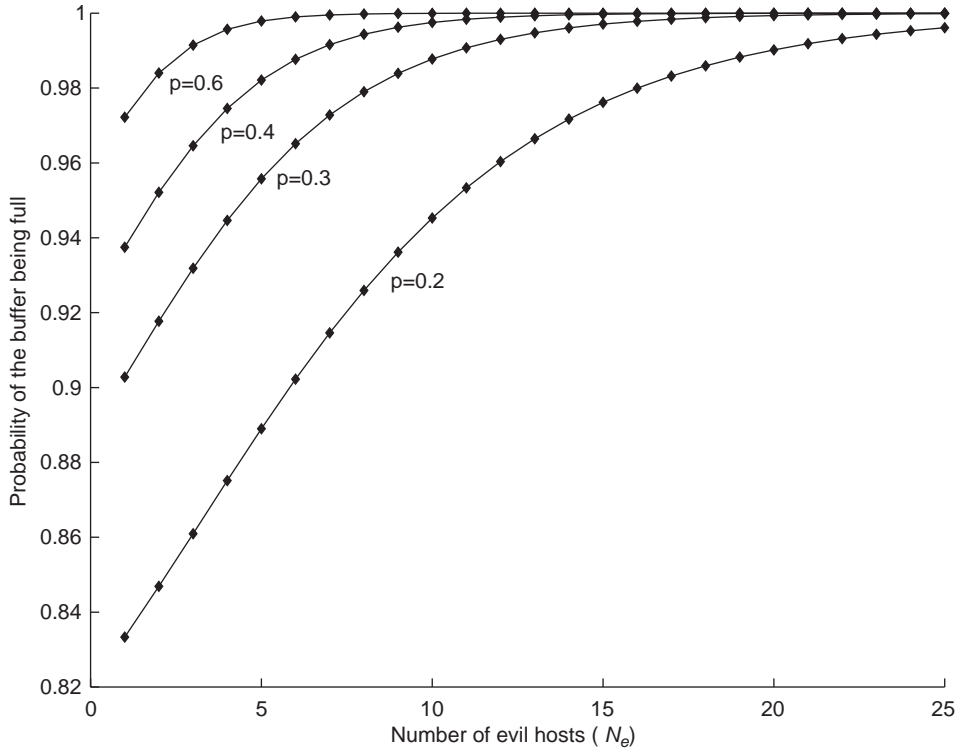


Fig. 5. The stationary probability X_m of the buffer being full as a function of the number of evil hosts N_e for 4 different values of the evil selfishness p . The size of the network has been fixed to $N = 25$. As can be observed, even for moderate values of p , X_m rapidly goes to 1.

The stationary probabilities of the chain of Fig. 4 can be calculated as the normalized eigenvector associated to the first eigenvalue of the transition matrix of the chain to a value:

$$X_i = \frac{1 - \rho}{1 - \rho^{m+1}} \rho^i, \quad (15)$$

where $\rho = \frac{p_+}{p_0}$. Notice that the probability of the buffer being full in a given time slot is X_m . As can be observed in Fig. 5, using a reasonable value of $N = 25$, the value of X_m rapidly goes to 1, even for moderate values of the evil selfishness p . Then, we can conclude that it is reasonable to assume permanent congestion.

Now, to simplify the analysis of the congested system with duplicated packets we maintain the assumption that the ordering of slots in a round is not established a priori. Then, the problem can be thought of as a switching process in which the current slot is an F -slot with probability $p_{SF} = \frac{N_f}{N}$ and an E -slot with probability $p_{SE} = \frac{N_e}{N}$. To mathematically analyze this situation, we have associated a Markov Chain with each evil host as depicted in Fig. 6. In this chain, the states are defined as a function of the position in the buffer

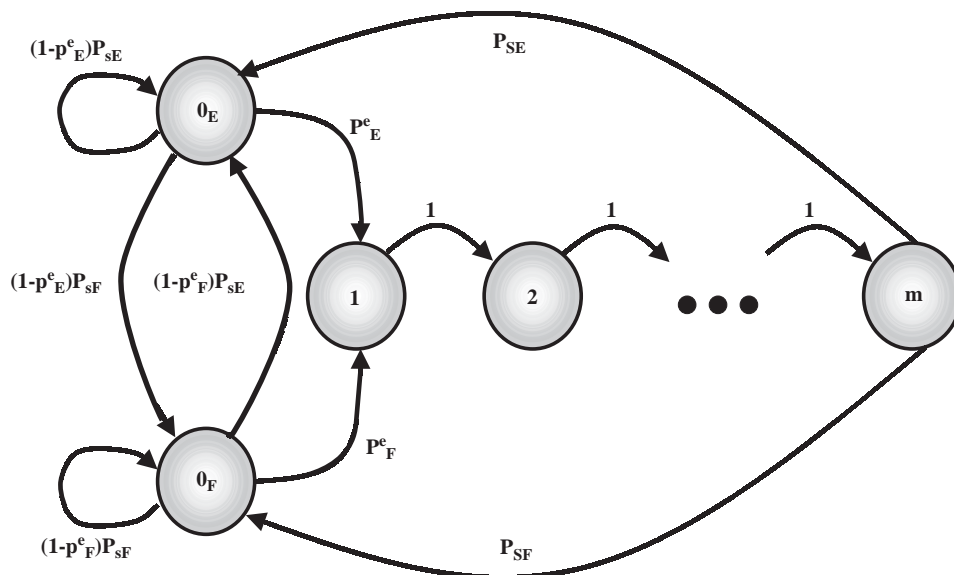


Fig. 6. A schema of the Markov Chain proposed to track the evolution of packets from a given evil host in the router buffer. In states 0_E and 0_F there are no useful packets in the buffer and the host is trying to inject one. The difference between these two states is in the type of current slot: E in the former case and F in the latter. In the rest of the cases, the state reflects the position of the useful packet within the buffer. Observe that $p_{sF} = \frac{N_f}{N}$ is the probability of the next slot to be fair while $p_{sE} = \frac{N_e}{N}$ is the probability of the next slot to be evil. Once the packet is in the FIFO buffer, its probability to pass to the next state is always 1. When the packet is delivered (state m), the host is immediately notified and it starts sending a new packet in states 0_E or 0_F depending on whether the subsequent slot is evil or fair.

of the oldest useful packet of the considered host. In states 0_E and 0_F , no useful packet of this host is present in the buffer. Observe that 0_E corresponds to a situation where the host is trying to access the buffer in an E -slot, while in 0_F the slot is an F -slot. In state 1 there is only one useful packet which occupies the last position of the buffer (following the FIFO buffer policy, it will be served within m time slots). In state 2, this packet has moved to the next buffer position. Note that another packet of the same host could occupy the preceding buffer place. We do not differentiate this case because the packet can only be a duplicate of the previous one, which has not yet arrived at its destination. Following the chain, we get to state m , in which the router sends the packet to the destination. Remember that, in our model, we have assumed that there exists an immediate hidden acknowledgment mechanism that allows the host to be aware of the arrival. Given this, the chain goes back to the zero states, because the host stops repeating the same packet and starts sending a new one which is not yet in the network. Observe that all the preceding duplicate packets existing in the network are not really useful for this reason, we do not consider them.

We may calculate the stationary probabilities of the proposed Markov chain just by obtaining the normalized eigenvector associated with the first eigenvalue of the transition

matrix of the chain. By doing so, we obtain the stationary probabilities:

$$\begin{aligned}\Pi(0_E) &= \frac{PSE}{1 + mq}, \\ \Pi(0_F) &= \frac{PSF}{1 + mq}, \\ \Pi(i) &= \frac{q}{1 + mq} \quad \text{for } i = 1, \dots, m,\end{aligned}\tag{16}$$

where $q = \frac{N_e}{N} p_E^e + \frac{N_f}{N} p_F^e$. Since a useful packet is only actually transmitted when the state m is reached (when the packet is sent out by the router to its destination), the probability of a useful packet to be issued in a time slot is given by the probability of state m , which is $\frac{q}{1+mq}$. Hence, we can conclude that the useful transmission rate for evil hosts is

$$R_e(N, N_e, p, m) = \frac{q}{1 + mq} = \frac{(N_e/N)p_E^e + (N_f/N)p_F^e}{1 + m(N_e/N)p_E^e + m(N_f/N)p_F^e},\tag{17}$$

Different types of games are possible in this situation. In particular, we have the following theorem (see Appendix A for the proof), which shows that the Tragedy of the Commons must appear for values of $p \in (\frac{2}{N-m}, 1]$.

Theorem 5. *In the Markovian model, for any $N > 1$, and $m < N - 2$, $p \in (\frac{2}{N-m}, 1]$ is a sufficient condition for the game to be a Tragedy of the Commons.*

The proof shows that for $p \in (\frac{2}{N-m}, 1]$, it holds that $R_e(N, i+1, p, m) > R_f(N, i, p, m)$, which guarantees that there is always an incentive for a host to break the ordered state and start cheating. Besides, $R_e(N, N, p, m) < R_f(N, 0, p, m)$, which guarantees the final yield to be below the performance of the ordered situation. Observe that, in the memoryless model, the interval was open, while, in this case, it is closed and the value $p = 1$ also produces a tragedy. The reason for this behavior arises from the fact that, in this case $R_e(N, N, p, m) < R_f(N, 0, p, m)$ even for $p = 1$, so, the two conditions for the game to be a Tragedy of the Commons hold. On the other hand, in the memoryless situation, $R_e(N, N, p) = R_f(N, 0, p)$ for $p = 1$, and the second condition of the tragedy is broken.

Furthermore, there is another interesting conclusion that can be obtained from Eq. (17). When the value of m increases, the reward obtained by evil hosts decreases. This effect is shown in Fig. 7, where we show the evil and fair rates as a function of the number of evil hosts, N_e , for $N = 25$. As can be observed, the effect of the tragedy becomes more considerable when m grows, because the evil rate in the Nash equilibrium decreases with this parameter. This phenomenon is produced by the effect of evil duplicates, which occupy more and more of the shared resources, but which do not provide any advantage in terms of transmission rate neither to evil nor to fair players.

Another extremely interesting conclusion that can be obtained from the model is that, in agreement with some previously proposed models [2], in the Nash equilibrium, the global performance of the network does not necessarily drop to zero (the network does not drop all packets). In fact, it can be easily shown using Eqs. (5), (7), (10), (17) that for completely greedy evil players ($p = 1$), the worst network performance (achieved when all players

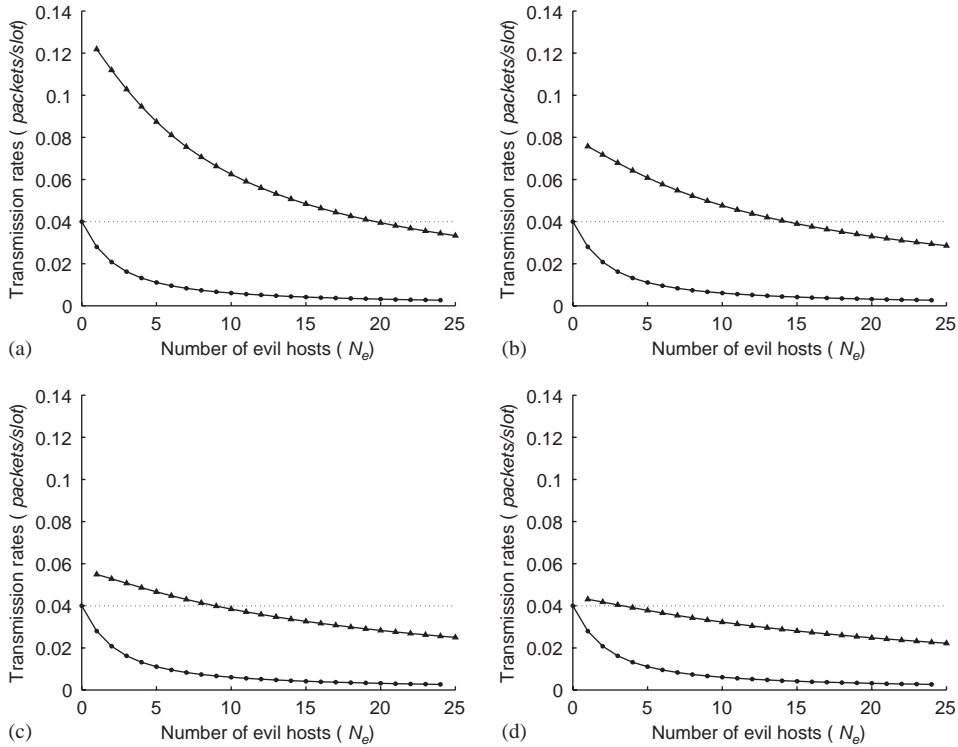


Fig. 7. Transmission rates for the Markovian model as a function of the number of evil hosts, N_e , in a network with $N = 25$ players. The evil probability is fixed to $p = 0.6$ in all graphs. Curve (a) has been evaluated using a buffer size $m = 5$. In (b), $m = 10$. In (c) $m = 15$ and in (d), $m = 20$. The evil rate $R_T(N, N_E, p, m)$ is represented by solid triangles. In circles we have the fair rate $R_F(N, N_e, p, m)$. The dotted line indicates the initial yield of the all-fair situation, $R_F(N, 0, p, m)$. As can be observed, in all cases the game is a Tragedy of the Commons, but its effect is more remarkable when the buffer size, m , increases.

are evil) falls to

$$R_f(N, N - 1, 1, m) = \frac{1}{N^2}, \tag{18}$$

for fair players (remember we assume $N \geq m$), and to

$$R_e(N, N, 1, m) = \frac{1}{N + m}, \tag{19}$$

for evil players. Hence, in a network full of evil greedy players, a fair host shows a performance that decreases in proportion to the square of the number of hosts, while an evil host shows a performance that decreases roughly in proportion to the number of hosts.

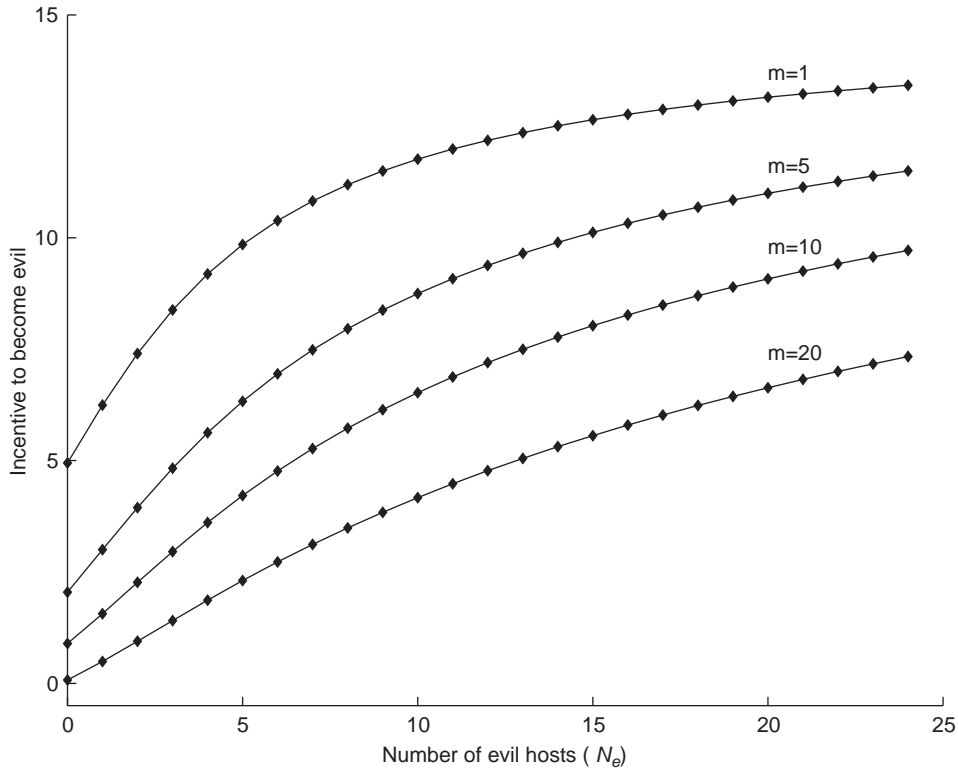


Fig. 8. The incentive to become evil as a function of the number of evil hosts in a network with $N = 25$ hosts. As can be observed, the incentive decreases when m grows. In the simulations the evil probability has been fixed to a value $p = 0, 6$.

Another interesting conclusion can be obtained by investigating how the incentive for a new host to become evil evolves. This incentive can be defined as

$$\Upsilon(N, N_e, p, m) = \frac{R_e(N, N_e + 1, p, m) - R_f(N, N_e, p, m)}{R_f(N, N_e, p, m)}. \quad (20)$$

This measures the relative increase in yield that a host would achieve by becoming evil. As can be observed in Fig. 8, in the Markovian model Υ decreases with m . This is a very interesting conclusion, because it shows that when the effect of the tragedy is more important (m high), the incentive to cheat is lower.

5. Simulations

The tTCP protocol we have defined imitates TCP, but using extremely simplified assumptions. Our objective in this section is to show that this simplification does not overlook the essential aspects of the real protocol. With this purpose, we have performed an extensive

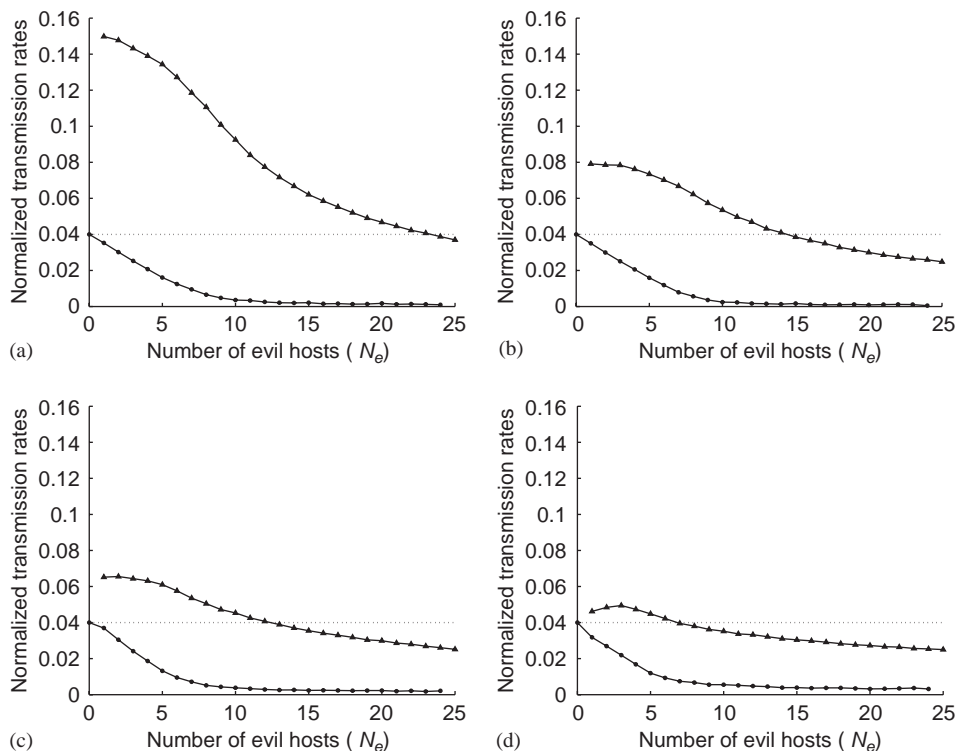


Fig. 9. The curves show the normalized transmission rates obtained from the simulations as a function of the number of evil hosts, N_e , in a network with $N = 25$ players. Curve (a) has been evaluated using a buffer size $m = 2$. In (b), $m = 10$. In (c) $m = 30$ and in (d), $m = 100$. The line with solid triangles represents the normalized evil rate. In solid circles we have the normalized fair rate. The dotted line indicates the initial yield of fair hosts. As can be observed, in all cases, the behavior of the simulations is similar to the predictions of the theoretical model.

set of simulations to show that the behavior of the conventional TCP protocol is very close to the predictions of our model.

The simulations have been carried out using the popular NS2 Network Simulator [12]. We have created a single bottle neck simulation scenario as the one depicted in Fig. 1, where all the communication lines have a fixed capacity of 100 Kbps and a delay of 10 ms. In all cases, the simulated time that we have used is of 150 seconds and the numerical results have been averaged over 50 different repetitions.

The fair hosts have been implemented using BayFullTCP agents, which comply with the TCP conventional standard. To implement the evil behavior we have developed a new type of agents based on the BayFullTCP code, but which obeys a strategy based on selfishness. The EvilBayFullTCP agents that we have created have all the necessary constituents of the protocol to work appropriately, but some of its pieces have been modified to emulate the greediness of this kind of player. In particular, the TCP congestion control algorithm [20], which calculates the size of the congestion window, has been altered so that this window never decreases. In practice, this means that after a short transient, the congestion window

in evil hosts reaches the maximum allowed size. Furthermore, in EvilBayFullTCP, we have also modified the retransmission timeout, which is no longer calculated using the Jacobson algorithm [20]. Evil hosts impose a fixed timeout, $T_e = 200$ ms, which is shorter than the round trip time (RTT) of the packets in the network when congestion occurs. Observe that T_e determines p , the probability an evil host attempts to retransmit, in the sense that a low timeout leads to a high value of p .

The graphs depicted in Fig. 9 represent the normalized evil and fair transmission rates obtained in the simulations. Here, normalized means that we have divided the transmission rates by the effective capacity of the communication links. Hence, the transmission rate for fair hosts when no evil hosts are present is always $1/N$, where N is again the number of players. Thanks to this normalization, we can easily compare these pictures with the theoretical results depicted in Fig. 7. As can be observed, the qualitative behavior obtained through the simulations is identical to the one predicted by the mathematical Markovian model that we have presented. Another interesting conclusion is that, in the simulations, we need larger buffer sizes to obtain the same kind of evolution. This is produced by the presence of a congestion window, which allows hosts to send new packets without having received the previous acknowledgments. Finally, note that the mathematical model overestimates the rate of fair hosts. This is also logical because the real protocol reacts to the presence of evil hosts decreasing the sending rate of fair hosts, which is not a feature considered in the model. In any case, we can conclude that the model qualitatively reproduces the main characteristics of the real TCP Tragedy of the Commons and it is an appropriate starting point to develop more complete analytical models which could allow a better understanding of this problem.

6. Conclusions

We have established a mathematical model for the TCP Tragedy of the Commons based on a simplifying assumption which consists in considering that the effect of the protocol is to render an ordered state to the network. In this state, we obtain an optimal collective strategy, where the resources are shared fairly among all hosts participating in the game. We have also defined an evil behavior driving to a strategy based on selfishness, which breaks the ordered state supplied by the protocol. In these circumstances, we have proved that the model predicts that there exists a threshold for the evil greediness over which the game takes the form of a Tragedy of the Commons.

Moreover, we have also shown that the effect of the tragedy is more remarkable when the network latency increases. We have demonstrated that this is due to the effect of duplicates. We have also proved that, assuming a worst case, the model predicts that the performance decays with the square of the number of hosts, for fair players, and proportionally to this number, for evil hosts. Besides, we have shown that when the incentive of hosts to become evil is higher, the tragedy becomes less significant, and when the tragedy is more prominent, the incentive decreases.

We have confirmed the validity of the mathematical model by performing an extensive set of simulations using NS [12]. In all cases the main conclusions of the model have been ratified by the simulations.

Acknowledgements

We would like to thank the anonymous referees for their useful comments and suggestions.

Appendix A.

A.1. Proof of Theorem 5

To prove the theorem, we need to formulate a couple of definitions. Let us define the pessimistic evil rate and the optimistic evil rate, respectively, as $R_e^{\text{pes}} = \frac{p_F^e}{1+mp_F^e}$ and $R_e^{\text{opt}} = \frac{p_E^e}{1+mp_E^e}$. Now observe that, in the Markovian model, the rate can be written as $R_e = \frac{q}{1+mq}$, which is strictly increasing for $q \in [0, 1]$ (its first derivative is positive). Since $p_E^e < p_F^e$ in all cases, it follows that $R_e^{\text{pes}} \leq R_e \leq R_e^{\text{opt}}$. That is, R_e^{pes} is in fact a pessimistic approximation of R_e , and R_e^{opt} is an optimistic approximation of R_e .

In these conditions, observe that proving Theorem 5 is equivalent to proving the following two conditions.

- (1) Condition 1. There is always an incentive for a new host to become evil. Observe that we ensure this condition to be true when the pessimistic evil rate is $R_e^{\text{pes}}(N, n+1, p, m) > R_f(N, n, p, m)$ for all $n \in \{0, \dots, N-1\}$. This indicates that the Nash equilibrium is reached when all hosts are evil ($N_e = N$).
- (2) Condition 2. The final yield for evil hosts in the Nash equilibrium is below the initial yield of fair hosts when all players collaborate. This condition is guaranteed when the optimistic approximation verifies $R_f(N, 0, p, m) > R_e^{\text{opt}}(N, N, p, m)$.

The proof may be split into two parts. Assuming $N > 1$ and $m < N-2$, they can be written as follows:

Lemma A.1. *In the Markovian model, $p > \frac{2}{N-m}$ is a sufficient condition for Condition 1 to be satisfied.*

Lemma A.2. *In the Markovian model, $p \leq 1$ is a sufficient condition for Condition 2 to be satisfied.*

A.2. Proof of Lemma A.1

We start by proving Lemma A.1. For this, note that the real rate is never smaller than the pessimistic approximation. Then, we can guarantee that Condition 1 is satisfied when

$$\begin{aligned} R_e^{\text{pes}}(N, n+1, p, m) &= \frac{p_F^e(n+1, p)}{1+mq_{\text{pes}}(n+1, p)} > \frac{p_F^f(n, p)}{N} \\ &= R_f(N, n, p, m). \end{aligned} \quad (\text{A.1})$$

Recovering the expressions of Eqs. (6), (7), this can be written as

$$\frac{(n+2)p-1+(1-p)^{n+2}}{(n+1)(n+2)p+m((n+2)p-1+(1-p)^{n+2})} > \frac{1-(1-p)^{n+1}}{p(n+1)N}. \quad (\text{A.2})$$

This equation has the form $\frac{A}{B+mA} > \frac{C}{D}$, which can also be written as $\frac{D}{C} - \frac{B}{A} > m$. Hence, the problem is reduced to proving that

$$\frac{p(n+1)N}{1-(1-p)^{n+1}} - \frac{(n+1)(n+2)p}{(n+2)p-1+(1-p)^{n+2}} > m, \quad (\text{A.3})$$

which can also be expressed as

$$N \frac{p(n+1)}{1-(1-p)^{n+1}} - \frac{n+1}{1-(1-(1-p)^{n+2})/p(n+2)} > m. \quad (\text{A.4})$$

Now, we define the following function:

$$\varphi_p(n) = \frac{pn}{1-(1-p)^n}. \quad (\text{A.5})$$

Since the inequality (A.4) must hold for any $n \geq 0$, it is equivalent to the following equation with $n \geq 1$ (just carrying out a simple change of variables):

$$N\varphi_p(n) - \frac{n}{1-1/\varphi_p(n+1)} > m. \quad (\text{A.6})$$

Before continuing, we would like to note that $\varphi_p(n)$ is an increasing function with n . To prove this, observe the function

$$f(x) = \frac{px}{1-(1-p)^x}, \quad (\text{A.7})$$

where $x > 0$. This function has a derivative equal to

$$f'(x) = \frac{p}{(1-(1-p)^x)^2} (1-(1-p)^x(1-x \ln(1-p))), \quad (\text{A.8})$$

where the sign of the derivative is the sign of the second factor. Then $f'(x) > 0 \Leftrightarrow 1 > (1-p)^x(1-x \ln(1-p)) \Leftrightarrow g(x) = (1-p)^{-x} + x \ln(1-p) > 1$. It is immediate that

$$g'(x) = -\ln(1-p)(1-p)^{-x} + \ln(1-p). \quad (\text{A.9})$$

Hence, $g'(x) > 0 \Leftrightarrow (1-p)^{-x} > 1 \Leftrightarrow (1-p)^x < 1$, which is always true. Then $g(x)$ is increasing, and as $g(0^+) = 1^+$, we may conclude that $g(x) > 1$, so $f'(x) > 0$ and $f(x)$ strictly increases with x .

Now, we will focus on the following inequality, which may be used to prove Eq. (A.6)

$$\frac{2}{p}\varphi_p(n) - \frac{n}{1-1/\varphi_p(n+1)} \geq 0, \quad (\text{A.10})$$

where $0 < p \leq 1$. After some simple operations, we may show that this is equivalent to

$$1 + (1-p)^n \geq \frac{2(1-(1-p)^{n+1})}{p(n+1)}. \quad (\text{A.11})$$

If we let $u = 1-p$, this is equivalent to

$$\psi(u) = (n-1) - (n+1)u + (n+1)u^n - (n-1)u^{n+1} \geq 0. \quad (\text{A.12})$$

As $\psi''(u) = (n+1)n(n-1)u^{n-2}(1-u) \geq 0$ in $[0, 1]$, ψ' is an increasing function in that interval. As $\psi'(1) = 0$, ψ' is negative in $[0, 1]$, and ψ is decreasing. So, we may conclude that $\psi(u) \geq \psi(1) = 0$. Hence, inequality (A.10) holds.

Coming back to the main problem, observe that the left-hand side of Eq. (A.6) can be written as

$$\begin{aligned} N\varphi_p(n) - \frac{n}{1 - 1/\varphi_p(n+1)} \\ = \left(N - \frac{2}{p}\right)\varphi_p(n) + \frac{2}{p}\varphi_p(n) - \frac{n}{1 - 1/\varphi_p(n+1)}. \end{aligned} \quad (\text{A.13})$$

Combining (A.10) with (A.6) it follows that

$$N\varphi_p(n) - \frac{n}{1 - 1/\varphi_p(n+1)} \geq \left(N - \frac{2}{p}\right)\varphi_p(n) \geq \left(N - \frac{2}{p}\right)\varphi_p(1). \quad (\text{A.14})$$

As we assume $p > \frac{2}{N-m} > \frac{2}{N}$, $N - \frac{2}{p} > m$. Besides, $\varphi_p(1) = 1$. This allows to conclude that

$$N\varphi_p(n) - \frac{n}{1 - 1/\varphi_p(n+1)} > m. \quad (\text{A.15})$$

This expression is identical to the one of Eq. (A.6), which is what we wished to prove. Hence, Lemma A.1 is proved. \square

A.3. Proof of Lemma A.2

Given that the real rate of the game is always below the results obtained using the optimistic approximation, Lemma A.2 can be proved by showing that $R_e^{\text{opt}}(N, N, p, m) < R_f(N, 0, p, m)$. In this sense, remember that $g(x) = \frac{x}{1+mx}$ is strictly increasing for values of $x \in [0, 1]$. In the same way, it is easy to show that $h(p) = p_F^f(N, N, p, m)$ is strictly increasing with $p \in [0, 1]$. The composition of two strictly increasing functions is also an increasing function, so

$$R_e^{\text{opt}}(N, N, p, m) = \frac{p_E^e(N, p)}{1 + mp_E^e(N, p)}, \quad (\text{A.16})$$

increases with p . For $p = 1$, we may easily evaluate the preceding expression as $R_e^{\text{opt}}(N, N, 1, m) = \frac{1}{N+m}$. So, given that the function is strictly increasing and that $R_f(N, 0, p, m) = 1/N$, it is evident that for any $p \leq 1$ and any $m \geq 1$

$$R_e^{\text{opt}}(N, N, p, m) \leq \frac{1}{N+m} < \frac{1}{N} = R_f(N, 0, p, m), \quad (\text{A.17})$$

Hence, Lemma A.2 is proved. \square

References

- [1] N. Abramson, The Aloha system—another alternative for computer communications, Fall Joint Computer Conference, AFIPS Conference Proceedings, Vol. 37, 1970, pp. 281–285.

- [2] A. Akella, S. Seshan, K. Karp, S. Shenker, C. Papadimitrou, Selfish behavior and stability of the internet: a game-theoretic analysis of TCP, ACM SIGCOM, August 2002.
- [3] A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, D. Williamson, Adversarial queueing theory, J. ACM 48 (1) (2001) 13–38.
- [4] D. Dutta, A. Goel, J. Heidemann, Oblivious AQM and nash equilibria, in: Proc. Conf. on Computer Communications (IEEE Infocom), April 2003.
- [5] S. Floyd, K. Fall, Promoting the use of end-to end congestion control in the internet, IEEE/ACM Trans. Networking 7 (4) (1999) 458–471.
- [6] D. Fudenberg, J. Tirole, Game Theory, MIT Press, Cambridge, MA, October 1991.
- [7] R. Garg, A. Kamra, V. Khurana, A game-theoretic approach towards congestion control in communication networks, AC; Comput. Comm. Rev. 32 (3) (2002) 47–61.
- [8] G. Hardin, The Tragedy of the Commons, Science 162 (1968) 1243–1248.
- [9] R.M. Metcalf, D.R. Boggs, Ethernet: distributed packet switching for local computer networks, Comm. ACM 19 (7) (1976) 395–404.
- [10] J. Nagle, Congestion control in IP/TCP internetworks, Request for Comments 896, Internet Engineering Task Force, January 1984.
- [11] J. Nagle, On packet switches with infinite storage, IEEE Trans. Comm. 35 (4) (1987) 435–438.
- [12] NS, UCB/LBNL/VINT Network Simulator—ns (version 2), <http://www.isi.edu/nsnam/ns>.
- [13] G. Owen, Game Theory, third ed., Academic Press, New York, 1995.
- [14] R. Pan, B. Prabhakar, K. Psounis, CHOKe, a stateless active queue management scheme for approximating fair bandwidth allocation, Proc. Conf. on Computer Communications (IEEE Infocom), March 2002.
- [15] C. Papadimitriou, Algorithms, games and the internet, in: Proc. 33rd Annual ACM Symposium on the Theory of Computing, 2001, pp. 749–753.
- [16] J. Postel, Transmission control protocol, Request for Comments 793, Internet Engineering Task Force, September 1981.
- [17] T. Roughgarden, E. Tardos, How bad is selfish routing?, J. ACM 49 (2002) 236–259.
- [18] S. Shenker, Making greed work in networks: a game-theoretic analyses of switch service disciplines, in: Proc. SIGCOMM, London, UK, September 1994, pp. 47–57.
- [19] W.R. Stevens, The Protocols, TCP/IP Illustrated, Vol. 1, Addison-Wesley, Reading, MA, 1994.
- [20] A.S. Tanenbaum, Computer Networks, third ed., Prentice-Hall, Englewood Cliffs, NJ, 1996.
- [21] H. Yaiche, R.R. Mazumdar, C. Rosenberg, A game theoretic framework for bandwidth allocation and pricing in broadband networks, IEEE/ACM Trans. Networking 8 (5) (2000) 667–678.