Computer Networks 53 (2009) 1751–1759

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

The evolution of transport protocols: An evolutionary game perspective

Eitan Altman [a,*], Rachid El-Azouzi [b], Yezekael Hayel [b], Hamidou Tembine [b]

[a] INRIA, Centre Sophia-Antipolis, 2004 Route des Lucioles, 06902 Sophia-Antipolis Cedex, France
[b] University of Avignon, LIA/CERI, 339, Chemin des Meinajaries, Agroparc BP 1228, 84911 Avignon Cedex 9, France

ARTICLE INFO

Article history:
Received 12 December 2007
Received in revised form 11 September 2008
Accepted 16 December 2008
Available online 6 January 2009

Keywords:
Evolutionary games
TCP
Protocols
Replicator dynamics

ABSTRACT

Today's Internet is well adapted to the evolution of protocols at various network layers. Much of the intelligence of congestion control is delegated to the end users and they have a large amount of freedom in the choice of the protocols they use. In the absence of a centralized policy for a global deployment of a unique protocol to perform a given task, the Internet experiences a competitive evolution between various versions of protocols. The evolution manifests itself through the upgrading of existing protocols, abandonment of some protocols and appearance of new ones. We highlight in this paper the modeling capabilities of the evolutionary game paradigm for explaining past evolution and predicting the future one. In particular, using this paradigm we derive conditions under which (i) a successful protocol would dominate and wipe away other protocols, or (ii) various competing protocols could coexist. In the latter case we also predict the share of users that would use each of the protocols. We further use evolutionary games to propose guidelines for upgrading protocols in order to achieve desirable stability behavior of the system.

1. Introduction

When transferring data between nodes, flow control protocols are needed to regulate the transmission rates so as to adapt to the available resources. A connection that looses data units has to retransmit them later. In the absence of adaptation to the congestion, the on going transmissions along with the retransmissions can cause increased congestion in the network resulting in losses and further retransmissions by this and/or by other connections. This type of phenomenon, that had lead to several 'congestion collapses' [15], motivated the evolution of the Internet transport protocol, TCP, to a protocol that reduces dramatically the connection's throughput upon congestion detection.

The possibilities to deploy freely new versions of protocols on terminals connected to the Internet creates a competition environment between protocols. Much work has been devoted to analyze such competition and to predict its consequences. The two main approaches for predicting whether one version of a protocol would dominate another one are

- *Inter-Population Competition (IRPC):* One examines local interactions between connections of different types that interact with each other (by sharing some common bottleneck link). If a connection that corresponds to one version performs better in such an interaction, then the IRPC approach predicts that it would dominate and that the other version would vanish.
- *Intra-Population Competition (IAPC):* In this approach one studies the performance of a version of a protocol assuming a world where all connections use that version. This is repeated with the other version. One then predicts that the version that gives a better world would dominate.

We address the dominance question with the evolutionary game paradigm and provide an alternative answer along with a more detailed analysis of this competition scenario. Our approach predicts whether one can expect

* Corresponding author. Tel.: +33 4 92 38 77 86.
E-mail addresses: Eitan.Altman@sophia.inria.fr (E. Altman), Rachid.Elazouzi@univ-avignon.fr (R. El-Azouzi), Yezekael.Hayel@univ-avignon.fr (Y. Hayel), Hamidou.Tembine@univ-avignon.fr (H. Tembine).
URL: http://www-sop.inria.fr/maestro/personnel/Eitan.Altman/ (E. Altman).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.12.023

one protocol to dominate the other or whether the two protocols can be expected to coexist. It provides with the tools for computing the share of the population that is expected to use each version in case the versions would coexist. Finally, it provides a description of the dynamics of the competition, which may result in a stable behavior that consists of a convergence to some equilibrium, or it may display instabilities and oscillations. By identifying the conditions for a stable behavior, one can provide guidelines for upgrading protocols so as to avoid undesirable oscillating behavior.

1.1. Background on evolutionary games

The evolutionary games formalism studies two concepts: the ESS (for *Evolutionary Stable Strategy*), and the *Evolutionary Game Dynamics*. The ESS, defined in 1972 by the biologist Maynard Smith [20], is characterized by a property of robustness against invaders (mutations). More specifically, (i) if an ESS is reached, then the proportions of each population do not change in time. (ii) at ESS, the populations are immune from being invaded by other small populations. This notion is stronger than Nash equilibrium in which it is only requested that a single user would not benefit by a change (mutation) of its behavior. Although ESS has been defined in the context of biological systems, it is highly relevant to engineering as well (see [26]). In the biological context, the replicator dynamics reflects and models the self organization of behaviors in a population or the self organization of the size of the population(s). In engineering, we can go beyond characterizing and modeling existing evolution. The evolution of protocols can be engineered by providing guidelines or regulations for the way to upgrade existing ones and in determining parameters related to deployment of new protocols and services. In doing so we may wish to achieve adaptability to changing environments (growth of traffic in networks, increase of speeds or of congestion) and yet to avoid instabilities that could otherwise prevent the system to reach an ESS.

1.2. Objectives

Our first objective is to provide a framework to describe and predict evolution of protocols in a context of competition between two types of behaviors: aggressive and peaceful. We compute the ESS for congestion protocols of different degree of aggressiveness. We identify cases in which at ESS, only one population prevails (ESS in pure strategies) and others, in which an equilibrium between several population types is obtained. To study this, we map the problems, whenever possible, into the Hawk and Dove game. We then study the convergence of the replicator dynamics to it.

The second objective of the paper is to provide a framework for controlling evolutionary dynamics (changing or upgrading protocols) through the choice of a gain parameter governing the replicator dynamics. We address the following two design issues concerning this choice:

(i) The tradeoff between fast convergence and stability. We identify a simple threshold on the gain parameter in the replicator dynamics such that the stability is only determined by whether we exceed or not the threshold.
(ii) The stability as a function of delays. We derive new stability conditions for the replicator dynamics in the Hawk and Dove game with non-symmetric delays and apply it to the evolution of the MAC and transport layer protocols.

1.3. Related work

We describe in Section 3 other work on competition between TCP versions. We briefly mention some other work of evolutionary games applied to the networking context. The earliest publication we know of that proposes the use of evolutionary games in networking is [29] that studies through simulations some aspects of competition between TCP users. The ESS concept has later been used in [6] in the context of ALOHA with power control. Some aspects of replicator dynamics (without delays) have been studied in [21] in the context of the association problem between access points in wireless local area networks. Other applications of evolutionary games to wireless networks are [8,14] and references therein as well as [7] on which Section 4 is based. (Ref. [7] is a preliminary version of this paper that does not include the part on competition over wireline.) For an application of evolutionary game to competitive routing see [9].

1.4. Structure

The paper is structured as follows. We first provide in the next section the needed background on evolutionary games. We summarize work on competition between TCP versions in Section 3. We then study the ESS for congestion control protocols (Sections 4 and 5). After that, we investigate the impact of the choice of some parameters in the replicator dynamics on the stability of the system in Section 6. Finally we give some numerical investigations and we conclude with concluding remarks.

2. ESS and replicator dynamics

Consider a large population of players. Each individual needs occasionally to take some action. We focus on some (arbitrary) tagged individual. The actions of some M (possibly random number of) other individuals may interact with the action of the tagged individual (e.g. some other connections share a common bottleneck). In order to make use of the wealth of tools and theory developed in the biology literature, we shall restrict here (as they do), to interactions that are limited to pairwise, i.e. to $M = 1$. This will correspond to networks operating at light loads, such as sensor networks that need to track some rare events such as the arrival at the vicinity of a sensor of some tagged animal.

We define by $J(p,q)$ the expected payoff for our tagged individual if it uses a strategy p when meeting another individual who adopts the strategy q. This payoff is called "fitness" and strategies with larger fitness are expected to propagate faster in a population.

We assume that there are N pure strategies. A strategy of an individual is a probability distribution over the pure strategies. An equivalent interpretation of strategies is obtained by assuming that individuals choose pure strategies and then the probability distribution represents the fraction of individuals in the population that choose each strategy.

2.1. Evolutionary Stable Strategies (ESS)

Suppose that the whole population uses a strategy q and that a small fraction ϵ (called "mutations") adopts another strategy p. Evolutionary forces are expected to select q against p if:

$$J(q, \epsilon p + (1-\epsilon)q) > J(p, \epsilon p + (1-\epsilon)q). \tag{1}$$

A strategy q is said to be ESS if for every $p \neq q$ there exists some $\hat{\epsilon}_y > 0$ such that (1) holds for all $\epsilon \in (0, \hat{\epsilon}_y)$.

In fact, we expect that if for all $p \neq q$,

$$J(q,q) > J(p,q), \tag{2}$$

then the mutations fraction in the population will tend to decrease (as it has a lower reward, meaning a lower growth rate). The strategy q is then immune to mutations. If it does not but if still the following holds,

$$J(q,q) = J(p,q) \quad \text{and} \quad J(q,p) > J(p,p) \quad \forall p \neq q, \tag{3}$$

then a population using q are "weakly" immune against a mutation using p since if the mutant's population grows, then we shall frequently have individuals with strategy q competing with mutants; in such cases, the condition $J(q,p) > J(p,p)$ ensures that the growth rate of the original population exceeds that of the mutants. A strategy is ESS if and only if it satisfies Eqs. (2) or (3), see [27, Proposition 2.1]. The conditions to be an ESS can be related to and interpreted in terms of Nash equilibrium in a matrix game. The situation in which an individual, say player 1, is faced with a member of a population in which a fraction p chooses strategy A is then translated to playing the matrix game against a second player who uses mixed strategies (randomizes) with probabilities p and $1-p$, resp. The central model that we shall use to investigate protocol evolution is introduced in the next subsection along with its matrix game representation. For more on the relation between ESS and Nash equilibria, see [25].

2.2. The Hawk and Dove (HD) game

Consider a large population of animals. Occasionally two animals find themselves in competition on the same piece of food. An animal can adopt an aggressive behavior (Hawk) or a peaceful one (Dove). The matrix in Fig. 1 presents the fitness of player I (some arbitrary player) associated with the possible outcomes of the game as a function of the actions taken by each one of the two players. We assume a symmetric game so the utilities of any animal (in particular of player 2) as function of its actions and those of a potential adversary (in particular of player 1), are the same as those player 1 depicted in Fig. 1. The utilities (i.e. fitness) represent the following:

		Player II	
		H	D
Pl. I	H	0.5 – d	1
	D	0	0.5

Fig. 1. A *H–D* game in matrix form.

> *An encounter D–D* results in a peaceful, equal-sharing of the food which translates to a fitness of 0.5 to each player.
>
> *An encounter H–H* results in a fight in which with equal chances, one or the other player obtains the food but also in which there is a positive probability for each one of the animals to be wounded. Then the fitness of each player is $0.5-d$, where the 0.5 term is as in the *D–D* encounter and the $-d$ term represents the expected loss of fitness due to being injured.
>
> *An encounter H–D or D–H* results in zero fitness to the *D* and in one unit of utility for the *H* that gets all the food without fight.

One can think of other scenarios that are not covered in the original *H–D* game, such as the possibility of a Hawk to find the Dove, in a *H–D* encounter, more delicious than the food they compete over. A generalized version [24] of the HD game given in Fig. 2 is characterized by $A11 < A22 < A12$ and $A21 < A22$. In that case,

(1) If $A11 > A21$ then the pure strategy H is the unique ESS.
(2) If $A11 < A21$ then there is a unique ESS $p = (p_L, p_H)$, it is a mixed strategy given by $p_H = u/(u+v)$ where $Aij = J(i,j),\ i,j \in \{H,D\},\ u = A12 - A22,\ v = A21 - A11$.

Remark 2.1

(i) Note that there are no settings of parameters for which the pure strategy D is an ESS in the *H–D* game (or in its generalized version).
(ii) In case 2 above, the strategies (H,D) and (D,H) are pure Nash equilibria in the matrix game. Being asymmetric, they are not candidates for being an

		Player II	
		H	D
Pl. I	H	A11	A12
	D	A21	A22

Fig. 2. Generalized *H–D* game.

ESS according to our definition. There are however contexts in which one obtains non-symmetric ESS, in which case they turn out to be ESS.

2.3. Evolution: replicator dynamics

We introduce here the replicator dynamics which describes the evolution in the population of the various strategies. In the replicator dynamics, the share of a strategy in the population grows at a rate equal to the difference between the payoff of that strategy and the average payoff of the population. More precisely, consider N strategies. Let $\mathbf{x}$ be the N-dimensional vector whose ith element x_i is the population share of strategy i (i.e. the fraction of the population that uses strategy i). Thus, we have $\sum_i x_i = 1$ and $x_i \geqslant 0$. Below we denote by $J(i,k)$ the expected payoff (or the fitness) for a player using strategy i when it encounters a player with strategy k. With some abuse of notation we define $J(i,\mathbf{x}) = \sum_j J(i,j)x_j$. Then the replicator dynamics is defined as

$$\begin{aligned}\dot{x}_i(t) &= x_i K\left(J(i,\mathbf{x}) - \sum_j x_j J(j,\mathbf{x})\right)\\ &= x_i K\left(\sum_j x_j J(i,j) - \sum_j \sum_k x_j J(j,k) x_k\right),\end{aligned} \quad (4)$$

where K is a positive constant and $\dot{x}_i(t) := \mathrm{d}x_i(t)/\mathrm{d}t$. Note that the right hand side vanishes when summing over i. This is compatible with the fact that we study here the share of each strategy rather than the size of the population that uses each one of the strategies.

2.4. Replicator dynamics with delay

In Eq. (4), the fitness of strategy i at time t has an instantaneous impact on the rate of growth of the population size that uses it. An alternative more realistic model for replicator dynamic would have some delay: the fitness acquired at time t will impact the rate of growth τ time later. We then have:

$$\dot{x}_i(t) = x_i(t)K \times \left(\sum_j x_j(t-\tau)J(i,j) - \sum_{j,k} x_j(t)J(j,k)x_k(t-\tau)\right), \quad (5)$$

where K is some positive constant. (It can be interpreted as a scaling factor of the fitness $J(.,.)$ or as a gain parameter that controls the speed of adaptation. More details will be given in Section 6.1.) The delay τ represents a time scale much slower than the physical (propagation and queueing) delays, it is related to the time scale of (i) switching from the use of one protocol to another (ii) upgrading protocols.

3. Background on competition between congestion control protocols

There are various versions of the TCP protocol among which the mostly used one is New-Reno. The degree of 'aggressiveness' varies from version to version. The behavior of New-Reno is approximately AIMD (Additive Increase Multiplicative Decrease): it adapts to the available capacity by increasing the window size in a linear way by α packets every round trip time and when it detects congestion it decreases the window size to β times its value. The constants α and β are 1 and 1/2, respectively, in New-Reno.

In last years, more aggressive TCP versions have appeared, such as HSTCP (High Speed TCP) [22] and Scalable TCP [16,12,3]. HSTCP can be modeled by an AIMD behavior where α and β are not constant anymore: α and β have minimum values of 1 and of 1/2, resp. and both increase as the window size increases. In other words, the values of α and β increase as new acknowledgements arrive, as long as no congestion is detected. Scalable TCP is an MIMD (Multiplicative Increase Multiplicative Decrease) protocol, where the window size increases exponentially instead of linearly and is thus more aggressive. Versions of TCP which are less aggressive than the New-Reno also exist, such as Vegas [11].

Several researchers have analyzed the performance of networks in which various transport protocols coexist, see [23,10,1,4,17]. In all these papers, the population size using each type of protocol is fixed.

Some papers have already considered competition between aggressive and well behaved congestion control mechanisms within a game theoretic approach. Their conclusions in a wireline context was that if connections can choose selfishly between a well behaved cooperative behavior and an aggressive one then the Nash equilibrium is obtained by all users being aggressive and thus in a congestion collapse [13,18].

We introduce in the next two sections two models of competition between TCP versions, both can be modeled within the framework of the Hawk and Dove game. This will allow us to predict whether a given version of TCP is expected to dominate others (ESS in pure strategies, which means that some versions of TCP would disappear) or whether several versions would coexist. The first model is adapted to competition in wireless networks and the second to wireline networks.

4. Competition in wireless networks

During the last few years, many researchers have been studying TCP performances in terms of energy consumption and average goodput within wireless networks [19,28]. Via simulation, the authors show that the TCP New-Reno can be considered as well performing within wireless environment among all other TCP variants and allows for greater energy savings. Indeed, a less aggressive TCP, as TCP New-Reno, may generate lower packet loss than other aggressive TCP. Thus the advantage of an aggressive TCP in terms of throughput could be compensated with energy efficiency of a more gentle TCP version. (In Section 5, we shall illustrate another consideration that affects the competition between TCP versions.) The goal of this section is to illustrate this point, as well as its possible impact on the evolution of the share of TCP versions, through a simple model of an aggressive TCP.

4.1. The model

We consider two populations of connections, all of which use AIMD TCP. A connection of population i is characterized by a linear increase rate α_i and a multiplicative decrease factor β_i. Let $\zeta_i(t)$ be the transmission rate of connection i at time t. We consider the following simple model for competition:

(i) The RTT (round trip times) are the same for all connections.
(ii) There is light traffic in the system in the sense that a connection either has all the resources it needs or it shares the resources with one other connection. (If files are large then this is a light regime in terms of number of connections but not in terms of workload.)
(iii) Losses occur whenever the sum of rates reaches the capacity C: $\zeta_1(t) + \zeta_2(t) = C$.
(iv) Losses are synchronized: when the combined rates attain C, both connections suffer from a loss. This synchronization has been observed in simulations for connections with RTTs close to each other [2]. The rate of connection i is reduced by the factor $\beta_i < 1$.
(v) As long as there are no losses, the rate of connection i increases linearly by a factor α_i.

We say that a TCP connection i is more aggressive than a connection j if $\alpha_i \geqslant \alpha_j$ and $\beta_i \geqslant \beta_j$. Let $\bar{\beta}_i := 1 - \beta_i$. Let y_n and z_n be the transmission rates of connection i and j, respectively, just before a loss occurs. We have $y_n + z_n = C$. Just after the loss, the rates are $\beta_1 y_n$ and $\beta_2 z_n$. The time it takes to reach again C is

$$T_n = \frac{C - \beta_1 y_n - \beta_2 z_n}{\alpha_1 + \alpha_2},$$

which yields the difference equation:

$$y_{n+1} = \beta_1 y_n + \alpha_1 T_n = q y_n + \frac{\alpha_1 C \bar{\beta}_2}{\alpha_1 + \alpha_2},$$

where $q = \frac{\alpha_1 \beta_2 + \alpha_2 \beta_1}{\alpha_1 + \alpha_2}$. The solution is given by

$$y_n = q^n y_0 + \left(\frac{\alpha_1 C \bar{\beta}_2}{\alpha_1 + \alpha_2} \right) \frac{1 - q^n}{1 - q}.$$

4.1.1. HD game: throughput-loss tradeoff

In wireline, the utility related to file transfers is usually taken to be the throughput, or a function of the throughput (e.g. the delay). It does not explicitly depend on the loss rate. This is not the case in wireless context. Indeed, since TCP retransmits lost packets, losses present energy inefficiency. Since energy is a costly resource in wireless, the loss rate is included explicitly in the utility of a user through the term representing energy cost. We thus consider fitness of the form $J_i = Thp_i - \lambda R$ for connection i; it is the difference between the throughput Thp_i and the loss rate R weighted by the so called tradeoff parameter, λ, that allows us to model the tradeoff between the valuation of losses and throughput in the fitness. We now proceed to show that our competition model between aggressive and non-aggressive TCP connections can be formulated as a HD game. We study how the fraction of aggressive TCP in the population at (the mixed) ESS depends on the tradeoff parameter λ.

Since $|q| < 1$, we get the following limit $\bar{y}$ of y_n when $n \to \infty$:

$$\bar{y} = \frac{\alpha_1 C \bar{\beta}_2}{\alpha_1 + \alpha_2} \cdot \frac{1}{1 - q} = \frac{\alpha_1 \bar{\beta}_2 C}{\alpha_1 \bar{\beta}_2 + \alpha_2 \bar{\beta}_1}.$$

It is easily seen that the share of the bandwidth (just before losses) of a user is increasing in its aggressiveness. Hence the average throughput of connection 1 is

$$Thp_1 = \frac{1 + \beta_1}{2} \times \frac{\alpha_1 \bar{\beta}_2}{\alpha_1 \bar{\beta}_2 + \alpha_2 \bar{\beta}_1} \times C.$$

The average loss rate of connection 1 is the same as that of connection 2 and is given by

$$R = \frac{1}{T} = \left(\frac{\alpha_1}{\bar{\beta}_1} + \frac{\alpha_2}{\bar{\beta}_2} \right) \frac{1}{C} \quad \text{where } T = \frac{\bar{\beta}_1 \bar{\beta}_2 C}{\alpha_1 \bar{\beta}_2 + \alpha_2 \bar{\beta}_1}$$

with T being the limit as $n \to \infty$ of T_n.

Let H corresponds to (α_H, β_H) and D to (α_D, β_D) such that $\alpha_H \geqslant \alpha_D$ and $\beta_H \geqslant \beta_D$. Then, for $i = 1, 2$, $Thp_i(H,H) = Thp_i(D,D)$. Since the loss rate for any user is increasing in α_1, α_2, β_1, β_2 it then follows that $J(H,H) < J(D,D)$, and $J(D,H) < J(D,D)$. We conclude that the utility that describes a tradeoff between average throughput and the loss rate leads to the HD structure.

The mixed ESS is given by the following probability of using H:

$$x^*(\lambda) = \frac{\eta_1 - \eta_2 \lambda}{\eta_3} \quad \text{where}$$

$$\eta_1 = \left(\bar{\mu} \frac{1 + \beta_1}{2} - \frac{1 + \beta_2}{4} \right) C, \quad \eta_2 = \frac{1}{C} \left(\frac{\alpha_1}{\bar{\beta}_1} - \frac{\alpha_2}{\bar{\beta}_2} \right),$$

$$\eta_3 = C \left(\frac{1}{2} - \mu \right) \frac{\beta_1 - \beta_2}{2}, \quad \mu = \frac{\alpha_2 (\bar{\beta}_1)}{\alpha_2 (\bar{\beta}_1) + \alpha_1 (\bar{\beta}_2)}.$$

where $\bar{\mu} := 1 - \mu$. Note that η_2 and η_3 are positive. Hence, the equilibrium point x^* decrease linearly on λ. We conclude that applications that are more sensitive to losses would be less aggressive at ESS (Braess type paradoxes do not occur here).

For more details on this model, including the tradeoff between transient and steady-state behavior, we refer the reader to Altman et al. [7].

5. Competition in wireline networks

5.1. Tradeoff between throughput and fairness

Consider a number of TCP connections that share a common bottleneck. We study the competition between the New-Reno version of TCP and the Scalable version of TCP that has been proposed [16] in the context wireline networks that are characterized by a very high speed and has long distance (and thus large delays). Scalable TCP is then much more aggressive than the New-Reno TCP as it is an MIMD (Multiplicative Increase Multiplicative Decrease) protocol, where the window size increases exponentially instead of linearly at the absence of losses. It is also more aggressive when losses are detected: it decreases its window size (and thus the transmission) rate to 0.875 times the one it had prior to the loss (instead of to halving it as is done in New-Reno). We begin by examining the Inter- and Intra-Population Competition.

5.2. Intra-Population Competition

If the connections are all symmetric and use the same AIMD version, then we know that they share equally the bandwidth [5]. Moreover if the number of connections or the buffer size is large then the connections use all the available bandwidth so that the throughput per connection is the capacity of the bottleneck link divided by the number of connections that share the bottleneck.

The situation is different when a bottleneck link is shared between symmetric MIMD connections. The connections suffer from a high level of unfairness under various scenarios. We briefly summarize some central findings in [4]:

- The case of synchronized losses: if all connections suffer a loss at the same time then the bandwidth share of each connection remains constant in time. This is particularly harmful for new connections that may not be able to grab throughput.
- The case of asynchronous losses: in case that there are some asynchronous random losses (with or without synchronous losses that are due to congestion), the bandwidth is shared more evenly: the average throughputs tend to be the same after a very long time.

Remark 5.1. The situation that we described corresponds to symmetric connections. In case the round trip times are not the same, the throughput share between the connections is inversely proportional to the RTT. In MIMD, in contrast, the connection with larger RTT gets almost no bandwidth.

As for the throughput, MIMD is able indeed to use all the available link capacity (provided that there are buffers; the buffer size needed for full utilization are much smaller than those needed in AIMD. Below we do not consider this aspect.)

5.3. Inter-Population Competition

The way bandwidth is shared between TCP connections of different types (AIMD and MIMD) that share a common bottleneck is complex and depends on various factors. We shall restrict our discussion to links with high bandwidth and large buffering. In those conditions, an MIMD connection gets more throughput than AIMD does. As for fairness, the problems of no fairness (in the sense that the initial share remains the same) or fairness on a very slow time scale, do not occur between two connections of different types (AIMD and MIMD).

5.4. The game formulation

Consider a competition between a large population of users where a fraction of them use an AIMD version of TCP and another MIMD connection sharing a common bottleneck link. We can now use the above description to define the structure of the evolutionary game. There are many pairwise interactions described through the matrix game of the form given in Fig. 2 where H (Hawk) stands for the MIMD TCP and D (Dove) stands for AIMD. Let the fitness of a player that uses an action i that interacts with a player that plays j be given as $F(i,j) = \theta(i,j) - f(i,j)$ where $\theta(i,j)$ is the throughput part corresponding to the share of the bottleneck capacity that the player receives and f is a disutility for lack of fairness. $f(i,j)$ is zero except for the case $i = j = H$; $\theta(i,j) = 0.5$ when $i = j$, where as $\theta(H,D) = 1 - \theta(D,H) = a$ where a is some positive constant smaller than 0.5. Denoting $d = 0.5 - f(H,H)$ we obtain the fitnesses as given in Fig. 3.

5.5. Characterizing the ESS

We note that the conditions defining a Generalized Hawk and Dove game as defined in Section 2.2 are indeed satisfied. We thus conclude the following:

- AIMD is never a dominant strategy (i.e. there is no ESS where all the population uses AIMD). To understand this, we note that if all used AIMD then a shift of a small fraction of the population to MIMD would give this fraction a larger fitness, since MIMD would get a larger throughput when interacting with AIMD but it would almost never suffer from the fairness problem since encounters between MIMD connections would be quite rare.
- MIMD may, on the other hand, be dominant. This occurs if $d > a$. In other words, if users did not care much about the fairness problems of MIMD, then its share in the population could be expected to grow to 1 and AIMD could disappear.
- If $d < a$ then the unique ESS consists of a coexistence between AIMD and MIMD. At ESS, the fraction using MIMD is given by

$$p = \frac{1 - 2a}{1 - 2d}.$$

5.6. Conclusion

We have mentioned in the introduction two simplistic approaches for predicting which population would dominate. The Intra-Population Competition (IAPC) compares

		Player II	
		H	D
Pl. I	H	d	1−a
	D	a	0.5

Fig. 3. MIMD (H) versus AIMD (H).

a world with only one type of behavior and compares the corresponding utility; in our case it would amount to observing that $F(D,D) > F(H,H)$ and thus would wrongly suggest that AIMD would dominate the MIMD and their share would be one, where as we know that this possibility never occurs.

The Inter-Population Competition (IRPC) compares only the interactions that involve H and D and it would predict that MIMD would dominate since it gets better fitness in these interactions. This prediction is again wrong; indeed, we have seen that both population may coexist in spite of having $F(H,D) > F(D,H)$ provided that $d < a$.

6. Architecting evolution

We study in this section numerically Eq. (5); we focus on the choice of two parameters in this replicator dynamics that impact the stability of the evolution process of protocols: the gain parameter K and the delay τ. The standard replicator dynamics (4) appearing in the evolutionary game literature is defined with $K = 1$. K's other than one can be interpreted as if the utilities J were multiplied by a constant. Alternatively, it can be seen as scaling time. The parameter K can thus be used to accelerate the rate of convergence in (4).

6.1. The impact of K and τ on the stability

We consider below the case of two players and two actions. Define

$$\delta_1 = J(B,A) - J(A,A), \quad \delta_2 = J(A,B) - J(B,B),$$
$$\delta = \delta_1 + \delta_2, \quad \theta = \frac{\delta\pi}{2\delta_1\delta_2}.$$

6.1.1. Guidelines for an evolution framework

For $K = 1$, it has been shown in [24] that if the delay τ in (5) satisfies $\tau < \theta$ then the mixed ESS (given in Section 2.2) is asymptotically stable.

For general K, we note that a change of variables can bring Eq. (5) to an equivalent form with $K = 1$. Indeed, let $\xi_i(t) := x_i(s)$ where $s = t/K$. Then,

$$\begin{aligned}\dot{\xi}_i(t) &= \frac{1}{K}\frac{\mathrm{d}x_i(s)}{\mathrm{d}s} \\ &= x_i(s)\left(\sum_j x_j(s-\tau)J(i,j) - \sum_{j,k} x_j(s)J(j,k)x_k(s-\tau)\right) \\ &= \xi_i(t)\left(\sum_j \xi_j(t-K\tau)J(i,j) - \sum_{j,k}\xi_j(t)J(j,k)\xi_k(s-K\tau)\right).\end{aligned} \quad (6)$$

Thus we can use the result of Tao and Wang [24] to conclude that the stability condition for general K is given simply by

$$\tau K < \theta. \quad (7)$$

This provides us with an important guideline for designing evolutionary protocols. In order for such a protocol to be scalable to any delay, the product of the adaptation speed parameter K and delay τ should be $O(1)$. Thus the larger the delay is, the slower we should react to the fitness of a strategy being used.

We note that this type of scaling is quite familiar in other networking contexts: the internet transport protocol TCP has a throughput that scales according to 1/RTT (where *RTT* is the round trip delay). This scaling is obtained by a self clocking mechanism based on ACKs that trigger new transmissions.

An alternative way to obtain the stability condition for general K is to view K as a scaling parameter for the fitnesses. Indeed, we may define $\widetilde{J}(.,.) = KJ(.,.)$ as a new fitness and then the replicator dynamics becomes:

$$\dot{x}_i(t) = x_i(t)\left(\sum_j x_j(t-\tau)\widetilde{J}(i,j) - \sum_{j,k} x_j(t)\widetilde{J}(j,k)x_k(t-\tau)\right). \quad (8)$$

Since the new θ is also scaled by a factor of K we obtain the same stability condition as before.

6.2. Numerical investigation

6.2.1. Impact of gain parameter

Our first numerical experiment studies the behavior of the replicator dynamics for the case of one delay unit as a function of K: we check the speed of convergence and the stability of the replicator dynamics as a function of the gain parameter K. We consider the following fixed parameters: we chose τ to be one unit, i.e. $\tau = 1$. This choice is without loss of generality since one can always rescale the model accordingly.

We further chose $\delta = 2/3$, and let K vary between 0.16 and 15. With this range of K we cover both the stability region where the stability condition is satisfied as well as the one it does not. A unique mixed ESS exists for these parameters, for which the fraction of the population using H is 3/4.

The resulting trajectories of the population ratio using the first strategy, H, as a function of time, is given in Fig. 4 top. For $K = 0.16$, we have stability but the convergence speed is slow. The other extreme is illustrated for $K = 15$ which is seen to be unstable: it oscillates rapidly and the amplitude is seen to grow slowly.

6.2.2. Impact of delay

We now keep K constant and evaluate the stability varying the delay between 0.016 and 15 time units. When $\tau = 0.016$ the system is stable but the rate of convergence to the interior equilibrium is not fast. For $\tau = 15$ the system is unstable, the solution oscillates around the equilibrium $x^* = 3/4$.

6.2.3. Oscillating solution and dependence on the initial state

In Fig. 4 bottom, we display an oscillatory behavior of the population ratios as function of time for two different initial values of $x(0) = 0.03$ and $x(0) = 0.97$ with $K = 15$ and $\tau = 1$. It corresponds to an unstable regime in which the ESS is not attained. The trajectories are seen to converge to periodic ones. The limit trajectories look the same and do not depend on the initial state except for a

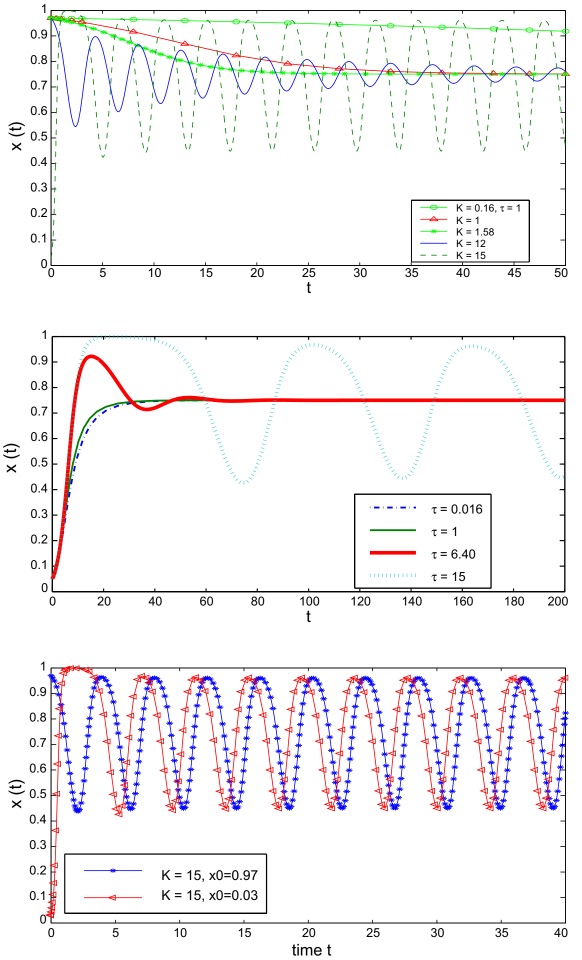

Fig. 4. (top) Effect of K on stability. (middle) Effect of τ on stability. (bottom) Effect of initial state.

dependence through the phase. In this unstable regime, more than one protocol coexist and the ratio of population sizes using the protocols has oscillations with large amplitude.

6.2.4. Validation of stability conditions

In both top and middle parts of Fig. 4, we observe that we have stability when $\tau K < 4\pi \simeq 12.56$. Indeed, in top part of Fig. 4, the parameter $\tau = 1$, hence the stability condition (7) becomes $K < 12.56$. This actually confirms that using $K = 0.16$, 1, 1.56, 12, the system is stable and using $K = 15$, the system is unstable. We observe the same behavior when keeping K constant and varying the delay τ.

7. Concluding remarks

In this paper, we have studied evolutionary aspects of congestion control protocols using the biological paradigm of evolutionary games. We have studied the questions of whether one could expect one type of protocol to wipe away another one or whether we may expect protocols to coexist. In the latter case we provided a quantitative characterization of the share that each protocol could be expected to have in the whole population at equilibrium. We then identified conditions under which there is a convergence to the equilibrium and obtained examined the oscillating behavior that occurs when there is no convergence. The conditions that guarantee convergence can be used a guidelines for deployment of new protocols so that the users upgrades would result in a stable system wide behavior.

Acknowledgements

This work was partly supported by the EuroFGI and EuroNF European Networks of Excellence and partly by the Popeye ARC collaborative program of INRIA, see http://www-sop.inria.fr/maestro/personnel/Eitan.Altman/POPEYE3/home.html.

References

[1] O. Ait-Hellal, E. Altman, Analysis of TCP Vegas and TCP Reno, Telecommun. Syst. 15 (2000) 381–404.
[2] O. Ait-Hellal, E. Altman, D. Elouadghiri, M. Erramdani, N. Mikou, Performance of TCP/IP: the case of two Controlled Sources, in: ICCC'97, Cannes, France, November 19–21, 1997.
[3] E. Altman, K. Avratchenkov, C. Barakat, A.A. Kherani, B.J. Prabhu, Analysis of scalable TCP, in: Proceedings of the Seventh IEEE International Conference on High Speed Networks and Multimedia Communications (HSNMC'04), Toulouse, France, June 30–July 2, 2004.
[4] E. Altman, K. Avrachenkov, B. Prabhu, Fairness in MIMD congestion control algorithms, Telecommun. Syst. 30 (4) (2005) 387–415.
[5] E. Altman, D. Barman, B. Tuffin, M. Vojnovic, Parallel TCP sockets: simple model, throughput and validation, in: IEEE Infocom, 2006.
[6] E. Altman, N. Bonneau, M. Debbah, G. Caire, An evolutionary game perspective to ALOHA with power control, in: Proceedings of the 19th International Teletraffic Congress, Beijing, 29 August–2 September, 2005.
[7] E. Altman, R. El-Azouzi, Y. Hayel, H. Tembine, An evolutionary game approach for the design of congestion control protocols in wireless networks, in: Physicomnet Workshop, Berlin, April 4, 2008.
[8] E. Altman, R. El-Azouzi, Y. Hayel, H. Tembine, Evolutionary power control games in wireless networks, in: Networking, Singapore, 2008.
[9] E. Altman, Y. Hayel, H. Kameda, Evolutionary dynamics and potential games in non-cooperative routing, in: Proceedings of the Workshop Wireless Networks: Communication, Cooperation and Competition (WNC3 2007), Limassol, Cyprus, April 16, 2007.
[10] T. Bonald, Comparison of TCP Reno and TCP Vegas: efficiency and fairness, Perform. Eval. 36–37 (1–4) (1999) 307–332.
[11] L.S. Brakmo, S. O'Malley, L.L. Peterson, TCP Vegas: new techniques for congestion detection avoidance, Comput. Commun. Rev. 24 (4) (1994) 24–35.
[12] R. El-Khoury, E. Altman, Analysis of scalable TCP, in: HET-NETs'04 International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks, West Yorkshire, UK, July 2004.
[13] R. Garg, A. Kamra, V. Khurana, A game-theoretic approach towards congestion control in communication networks, SIGCOMM Comput. Commun. Rev. Arch. 32 (3) (2002) 47–61.
[14] Hamidou Tembine, Eitan Altman, Rachid El-Azouzi, Yezekael Hayel, Evolutionary games with random number of interacting players applied to access control, in: WIOPT, Berlin, April 2008.
[15] Van Jacobson, Congestion avoidance and control, in: SIGCOMM'88, Stanford, CA, August 1988.
[16] T. Kelly, Scalable TCP: improving performance in highspeed wide area networks, Comput. Commun. Rev. 32 (2) (2003).
[17] L. Lopez, A. Fernandez, V. Cholvi, A game theoretic analysis of protocols based on fountain codes, in: IEEE ISCC'2005, June 2005.
[18] L. Lopez, G. Rey, A. Fernandez, S. Paquelet, A mathematical model for the TCP tragedy of the commons, Theor. Comput. Sci. 343 (2005) 4–26.
[19] H. Singh, S. Singh, Energy consumption of TCP Reno, New Reno and SACK in multi-hop wireless networks, in: ACM SIGMERTICS, June 2002.

[20] J. Maynard Smith, Game theory and the evolution of fighting, in: John Maynard Smith (Ed.), On Evolution, Edinburgh University Press, 1972, pp. 8–28.
[21] S. Shakkottai, E. Altman, A. Kumar, The case for non-cooperative multihoming of users to access points in IEEE 802.11 WLANs, invited paper, IEEE J. Select. Areas Commun. 25 (6) (2007) 1207–1215.
[22] E. Souza, D.A. Agarwal, A highspeed TCP study: characteristics and deployment issues, Technical Report LBNL-53215, Lawrence Berkeley National Laboratory, 2002.
[23] A. Tang, J. Wang, S. Low, M. Chiang, Equilibrium of heterogeneous congestion control: existence and uniqueness, IEEE/ACM Trans. Netw. (2007).
[24] Y. Tao, Z. Wang, Effect of time delay and evolutionarily stable strategy, J. Theor. Biol. 187 (1997) 111–116.
[25] Eric van Damme, Stability and Perfection of Nash Equilibria, Springer, Berlin, 1991.
[26] T.L. Vincent, T.L.S. Vincent, Evolution and control system design, IEEE Contr. Syst. Mag. 20 (5) (2000) 20–35.
[27] J. Weibull, Evolutionary Game Theory, MIT Press, 1995.
[28] M. Zorzi, R. Rao, Energy efficiency of TCP in a local wireless environment, Mobile Netw. Appl. 6 (3) (2001).
[29] Youquan Zheng, Zhenming Feng, Evolutionary game and resources competition in the Internet, in: Modern Communication Technologies, 2001, SIBCOM-2001, The IEEE-Siberian Workshop of Students and Young Researchers, 2001, pp. 51–54.

Eitan Altman received the B.Sc. degree in Electrical Engineering (1984), the B.A. degree in Physics (1984) and the Ph.D. degree in Electrical Engineering (1990), all from the Technion-Israel Institute, Haifa. In 1990, he further received his B.Mus. degree in Music Composition in Tel-Aviv University. Since 1990, he has been with INRIA (National Research Institute in Informatics and Control) in Sophia-Antipolis, France. His current research interests include performance evaluation and control of telecommunication networks and in particular congestion control, wireless communications and networking games. He is in the editorial board of the scientific journals: WINET, JDEDs and JEDC, and served in the journals Stochastic Models, COMNET, SIAM SICON. He has been the general chairman and the (co)chairman of the program committee of several international conferences and workshops (on game theory, networking games and mobile networks). He has published more than 140 papers in international refereed scientific journals. More information can be found at http://www.inria.fr/maestro/personnel/Eitan.Altman/me.html.

Rachid El-Azouzi received the Ph.D. degree in Applied Mathematics from the Mohammed V University, Rabat, Morocco (2000). He joined INRIA (National Research Institute in Informatics and Control) Sophia-Antipolis for post-doctoral and Research Engineer positions. Since 2003, he has been a researcher at the University of Avignon, France. His research interests are mobile networks, performance evaluation, TCP protocol, wireless networks, resource allocation, networking games and pricing. His web page is http://www.lia.univ-avignon.fr/chercheurs/elazouzi/.

Yezekael Hayel received his M.Sc. in Computer Science and Applied Mathematics from the University of Rennes 1, in 2003. He has a Ph.D. in Computer Science from University of Rennes 1 and INRIA. He is currently an assistant professor at the University of Avignon, France. His research interests include wireless networks, performance evaluation of networks and bio-inspired technologies.

Hamidou Tembine received the Master degree in Mathematics and Applied Mathematics in 2006, respectively, from University Joseph Fourier, Grenoble, France and Ecole Polytechnique, Palaiseau, France. He is working towards the Ph.D. degree at LIA, Avignon University, France. His current research interests include evolutionary games and networking applications.