An Evolutionary Game-Theoretic Approach to Congestion Control

Daniel S. Menasché[1] **Daniel R. Figueiredo**[2] **Edmundo A. de Souza e Silva**[1]

[1] COPPE-PESC and Computer Science Dept. – Universidade Federal do Rio de Janeiro

[2]Dept. of Computer Science – University of Massachusetts

`sadoc@land.ufrj.br, ratton@cs.umass.edu, edmundo@land.ufrj.br`

Abstract. *This paper investigates a system where a set of users that share a common network link are free to choose the transmission rate of multimedia data. Users are assumed to be self-regarding and make their decisions with the sole goal of maximizing the QoS they perceive. In order to understand this system we develop an evolutionary game-theoretic model to evaluate equilibria points that are reached by the users. Using our proposed model, we demonstrate analytically and numerically several interesting properties of the system equilibria. In particular, we establish the relationship between states that have non-negligible probabilities and Nash equilibria of the game. Techniques from the theory of aggregation in Markov chains are used to prove this result.*

1. Introduction

Congestion control is one of the most fundamental problems in computer networks and has been widely studied for decades. In essence, congestion control is a resource allocation problem, where bandwidth must be shared among different data flows traversing a congested link. In today's Internet, congestion control is handled solely by TCP, the dominant transport protocol used for transferring data. TCP implements a window-based congestion control mechanism which is responsible for regulating the data rate entering the network. Users are usually oblivious to its functionality and simply rely on its operation.

However, it is generally accepted that TCP is not suited for transferring delay sensitive data such as voice and video. With the increase in the demand for multimedia applications, transferring this type of data effectively becomes important. To address this problem researchers have advocated using the UDP transport protocol coupled with some congestion control mechanism [Floyd et al., 2000, Rejaie et al., 1999]. The use of such mechanisms is important because it can prevent a potential congestion collapse.

The deployment of congestion control mechanisms by multimedia application developers that adopt the UDP protocol is fully voluntary. In fact, competing developers are likely to implement mechanisms that best suit their needs. One possibility is not to implement any congestion control mechanism and transfer the responsibility of regulating the data rate to the users. Moreover, applications can deploy redundancy mechanisms (i.e., forward error correction) to masquerade packet losses in the network and improve media quality. Again, the decision of how much redundancy the application should send can also be transfered to the user. In this scenario, users are free to choose the rate at which data is sent. In the context of voice and video applications for example, encoders already allow users to choose the rate at which data is encoded [GNU, 2004]. Users can then determine the data rates that maximize the perceived quality of the multimedia stream being received.

This idea of allowing users to determine their data rates as a mechanism for congestion control constitutes a broad field of research. In this scenario, users have an utility function which depends on the characteristics of the data transfer (e.g., throughput) and are assumed to be *self-regarding*. This latter assumption means that users are only interested in maximizing their own utility. Note that users are competing for a shared resource and that decisions made by one user can affect the utility of all others. In this context, game theory emerges as a natural framework to model and evaluate the performance of such systems. Classical game-theoretic models have in fact been applied to this problem [Hsiao and Lazar, 1987, Shenker, 1995, Johari and Tsitsiklis, 2004].

In this paper, we develop a dynamic game-theoretic model to evaluate equilibria points that are reached by self-regarding users. Our model is based on evolutionary models where users adapt their data rates based on the perceived quality of service. We consider users' utility functions that are closely tied to quality of service (QoS) metrics for multimedia applications. Using the proposed model, we demonstrate several interesting properties of the equilibria of this system. In particular, we show that in steady state there exists a relationship between the states that have non-negligible probabilities and the Nash equilibria of the game. Techniques from the theory of aggregation in Markov chains are used to prove this result.

The remainder of this paper is organized as follows. In the next section we provide a brief introduction to game theory. In Section 3 we discuss related work. In Section 4 we formally present the problem investigated and the proposed model. Sections 5 and 6 describe the analytical and numerical results obtained from our model, respectively. Finally, Section 7 concludes the paper.

2. Game Theory Review

The main goal of game theory is to understand how players act when confronted with a scenario where there are conflicts of interest. In a given conflict of interest scenario, each player must make a choice from a given set of possible choices. In the game theory nomenclature, this choice is known as the player's strategy and the set of possible choices the strategy set. The joint decision of all players will determine the outcome of the game and each player has some preference over the set of possible outcomes. Classical game theory assumes players have full knowledge of the game and exhibit *rational* behavior[1].

We will introduce definitions and illustrate important concepts in game theory using a simple example. Consider two players, $\mathcal{A}$ and $\mathcal{B}$, that must share a single network link in order to receive multimedia data. Players $\mathcal{A}$ and $\mathcal{B}$ have to their disposal two possible data rates to choose from, a smooth and an aggressive one: $\{\lambda_s, \lambda_a\}$, with $\lambda_s < \lambda_a$. The utility of each player represents the quality of service (QoS) experimented by the player and depends on the outcome of the game.

Tables 1(a) and 1(b) illustrates two different games. The value in the cell of the matrix indicates the QoS received by players $\mathcal{B}$ and $\mathcal{A}$, respectively. For example, if

[1] By rational we mean that a player has consistent preferences and acts to maximize its long run benefits.

player $\mathcal{A}$ chooses λ_s and player $\mathcal{B}$ chooses λ_a, their QoS in the game defined by Table 1(a) will be 5 and 15, respectively. The game defined by Table 1(a) illustrates a scenario

	$\mathcal{A}$ plays λ_s	$\mathcal{A}$ plays λ_a
$\mathcal{B}$ plays λ_s	5,5	5,15
$\mathcal{B}$ plays λ_a	15,5	15,15

(a)

	$\mathcal{A}$ plays λ_s	$\mathcal{A}$ plays λ_a
$\mathcal{B}$ plays λ_s	5,5	1,15
$\mathcal{B}$ plays λ_a	15,1	4,4

(b)

Table 1. Two strategic games.

where a larger data rate yields a larger QoS for each player. In this case, the network link being shared does not become congested when both users choose to receive at λ_a. Table 1(b) illustrates the case where the network link becomes congested when both users choose to receive at λ_a. Note that the QoS experimented by both players in this case is worst than the one experimented formerly when both players choose to receive at λ_a. Having presented an example, we now formally introduce the concept of a strategic game [Osborne and Rubinstein, 1994]:

Definition 1 *A finite **strategic game** is characterized by $\langle N, (A_i), (u_i) \rangle$, which consists of • (i) a finite set N of players • (ii) for each player $i \in N$ a non empty finite set A_i (the set of choices available to player i) • (iii) for each player $i \in N$ an utility function $u_i : A = \times_{j \in N} A_j \to \Re$ (the utility function of player i).*

The fundamental problem in game theory is understanding how players will act when faced with a particular game. In particular, one is interested in determining the choices that players will make, which is sometimes referred to as the *solution* of the game. However, there are several solution concepts defined within game theory. We will adopt the most common one, known as a Nash equilibrium. A Nash equilibrium is a set of choices, one choice made by each player, where no individual player can improve his utility by unilaterally changing his choice. More formally, we have:

Definition 2 *A **Nash equilibrium** of a strategic game $\langle N, (A_i), (u_i) \rangle$ is a vector choices $a^\star = (a_1^\star, \ldots, a_N^\star) \in A = \times_{j \in N} A_j$, one choice $a_i^\star$ for each player i, such that for each $i \in N$, $u_i(a_{-i}^\star, a_i^\star) \geq u_i(a_{-i}^\star, b_i)$ for all $b_i \in A_i$ where $a_{-i}^\star = (a_i^\star)_{i \in N \setminus \{i\}}$ (i.e., $a_{-i}^\star$ is a vector of choices, one choice for each player, except player i).*

In our previous example, the games defined by Tables 1(a) and 1(b) both have a single Nash equilibrium which is given by (λ_a, λ_a). Note that in the game 1(b) both players could obtain a larger QoS if they both chose λ_s. However, (λ_s, λ_s) is not a Nash equilibrium because one player can obtain a higher QoS by unilaterally changing to λ_a. This last observation motivates the definition of equilibria points that yield high utility values to all players. In fact, a Pareto equilibrium is a set of choices such that there does not exist another set of choices where all players receive a higher utility. More formally, we have

Definition 3 *A **Pareto equilibrium** of a strategic game $\langle N, (A_i), (u_i) \rangle$ is a vector choices $a^\star = (a_1^\star, \ldots, a_N^\star) \in A = \times_{j \in N} A_j$, one choice $a_i^\star$ for each player i, such that there is no $b \in A$ that satisfies $u_i(b) > u_i(a^\star)$ for all $i \in N$.*

Returning to our example, the Pareto equilibrium of the game defined by Table 1(a) is given by (λ_a, λ_a), while in the game defined by Table 1(b) is given by (λ_s, λ_s). Note that in the first game the Pareto equilibrium coincides with the Nash equilibrium, while in the second game they differ.

The fact that the Nash equilibrium frequently does not coincide with the Pareto equilibrium constitutes one of the core issues in game theory. Much research has been done in understanding this phenomena and quantifying the difference in utility between the two equilibria. This phenomena is usually associated with two basic premises of classical game theory. First, the assumption that players are perfectly rational and have full knowledge of the game. Second, that the game is played only once and that players make decisions that are definitive.

Another fundamental problem in game theory is determining how players reach their decisions. In particular, the question of how can players arrive at Nash equilibrium is of ultimate importance. Classical game theory usually assumes players are capable of determining the Nash equilibrium that will be played. However, for some games this requirement is too stringent as determining the Nash equilibrium can have a very high complexity cost and multiple Nash equilibrium can co-exist.

2.1. Evolutionary game theory

Evolutionary game theory has emerged as an alternative perspective to classical game theory. One of its main advantages is the dismissal of the assumption that players must be rational. Although players still want to maximize their utility, the reasoning process by which players decide their strategy is usually trivial. Another advantage of evolutionary game theory is that equilibria are usually determined by the result of the evolutionary process. In other words, how players reach and select an equilibrium is part of the modeling framework.

Evolutionary game theory considers a dynamic scenario where players are constantly interacting with one another and adapting their choices based on the utility they receive. The adaptation process is a fundamental component of evolutionary games and is usually interpreted as an evolutionary or learning process. In its simplest form, players are assumed to be symmetric, in that they have identical set of choices and identical utility functions. An evolutionary game can be formally defined as follows:

Definition 4 *A symmetric* ***evolutionary game*** $\langle N, (A), (u), \sigma(t), \mathscr{S}(t), \mathscr{D}\rangle$ *models the interaction of players over time, and is composed of the following: • (i) a finite set* N *of players • (ii) a non empty set* A *(the set of available choices for each player), with* $k = |A|$ *• (iii) an utility function* $u : A = \times_{j\in N} A_j \rightarrow \Re$ *(the utility function of each player) • (iv) a vector of values between 0 e 1,* $\sigma(t) = (\sigma_1, \dots, \sigma_k)(t)$*, where* $\sigma_j(t)$ *is the fraction of the population selecting the choice* $a_j \in A$ *at time* t *(the state of population at time* t*) • (v) an ordinary game* $\mathscr{S}(t) = \langle O, (A), (u)\rangle$ *(the "stage game" at time* t*) • (vi) a dynamic process of strategy adjustment, which is a function (possibly stochastic)* $\mathscr{D} : (\sigma(t), \mathscr{S}(t), t) \rightarrow \sigma(t + \Delta t)$ *(the dynamic selection process).*

In essence, players are repeatedly matched in groups of O, usually randomly, in order to play the stage game $\mathscr{S}(t)$ at time t. The stage game is played in accordance with $\sigma(t)$, the fraction of the population selecting each choice. Given the utilities obtained, players adjust their choices according to the dynamics of $\mathscr{D}$.

The definition of the game dynamics, $\mathscr{D}$, can be deterministic or probabilistic. The initial evolutionary game theory models had a deterministic game dynamics, based on differential equations. A popular dynamics in this context is know as the *replicator dynamics*. More recent models adopt a probabilistic game dynamics, many of which are based on Markov chains. In this case, the state of the Markov chain is usually given by a vector representing the number of players making a given choice [Kandori et al., 1993, Young, 1993].

Returning to our example, we now describe it in the context of an evolutionary game. Let N be the total number of players and assume N is even. We assume that the players interact two at a time and that the stage game $\mathscr{S}$ is given by Table 1(b). Let time be divided into discrete slots and between time t and $t+1$ all players are matched up in pairs randomly to play the stage game and adapt their choices based on the utilities received. Since the stage game considered is bidimensional, the utility obtained by a player when the population state is $\sigma(t)$ is $u(\sigma(t)) = \sum_{j=1}^{k} u(a_j)\sigma_j(t)$ where $u(a_j)$ corresponds to the utility received by a player when the entire population has chosen $a_j \in A$. Finally, we must define $\mathscr{D}$, the process for strategy adjustment. For example, we can stipulate that a player should switch his strategy between time t and $t+1$ if he could have received a higher utility had he played the other strategy. If we assume the initial state of the population is $\sigma(0) = (1, 0)$ – all the players choose λ_s, then at time 1 we have $\sigma(1) = (0, 1)$ – all players choose λ_a.

3. Related Work

Several papers have approached the congestion control problem in computer networks using a game theoretic formulation. The idea that game theory could be applied to congestion control was first noted in [Nagle, 1985]. An extensive analytical study with focus on characterizing Nash and Pareto equilibria under different service disciplines is presented in [Shenker, 1995]. In [Akella et al., 2002], the authors use simulation to evaluate the performance of the network at the equilibria of “TCP games”. However, these prior studies have focused on a static scenario and have not considered the general problem of how Nash equilibria can be attained by network users sharing a FIFO bottleneck queue.

The problem of reaching an equilibrium in the context of dynamic network games was investigated in [Greenwald et al., 2001] through extensive simulations. However, the congestion control problem is only briefly discussed in the context of simple learning algorithms with small number of players. Although there are substantial differences between [Greenwald et al., 2001] and our work, this work is closely related to ours.

Our work is the first to develop and apply evolutionary game-theoretic models to the congestion control problem in the context of multimedia data streams and to consider user-level performance metrics (e.g., QoS). We also provide both analytical and numerical treatment to the our modeling framework.

Evolutionary game-theoretic models have been widely studied in the economics literature. Two prominent models are the works of [Kandori et al., 1993, Young, 1993] (generalized by [Samuelson, 1997]), which provide strong analytical results on the *Nash equilibrium selection* problem. Although these two models share a common basis, they differ considerably in the details. Both consider a finite player population and a discrete time model where at each step players are randomly paired to play a bidimensional stage game. However, the dynamics of the strategy selection process differs. In [Kandori et al., 1993], players change their strategies based on the “best response” to the current population. In [Young, 1993], a more complex structure is considered where players change their strategies according to their short-term memory of previous games played. In both models, it is assumed that players can make mistakes when selecting their strategies.

Although our model draws on ideas exposed by [Kandori et al., 1993, Young, 1993], there are some fundamental differences which make our model more suitable to the congestion control problem considered here. First, our model is asynchronous, which means that users change their strategies one at a time. The two previous models are synchronous, which means that all users can change their strategy at each time step. Second, our model assumes continuous time, which means that at any point in time one (and only one) user can change its strategy. The two previous models assume discrete time. Third, our model considers an N player stage game, where all users simultaneously participate in the game. The previous models assume random pairing of users and a bi-dimensional stage game. Fourth, the payoff of the stage game is given by a performance model which depends on the strategies of all users (in this paper we consider an M/M/1/k queue to obtain performance metrics, but any other performance model could be used as well). Finally, the strategy selection process we consider is also very different from those in the previous models.

4. The model

We consider the problem where a set of self-regarding users transmitting multimedia data must share the bandwidth of a common network link. We assume that no congestion control protocol is deployed and that users are free to choose the bandwidth at which data is transmitted. This problem, for instance, appears when a set of client users access a set of multimedia servers which are capable of delivering data at various rates. This scenario is illustrated in Figure 1(a).

It is well known that despite the fact that multimedia applications can tolerate packet losses and delays [Kurose and Ross, 2004], the quality of the audio/video perceived by the user (e.g., user level QoS) can vary drastically with these variables. We will assume users have a well defined quality of service (QoS) function. Multimedia applications can also adapt the rate at which data is transmitted (i.e., through different coding algorithms), trading smaller bandwidth requirements for lower QoS. We will assume applications offer users a few different data rates to choose from.

In this scenario, users must now determine the application data rate that maximizes their utility function. However, since users share a common network link, decisions of one user affects the quality of service perceived by all others, and consequently the decisions the others will make. Hence, this problem naturally leads to a game-theoretic model where the payoffs represent the QoS that each user perceives.

It is important to note that we are interested not only in the final outcome of the game, but also in the dynamic process by which users change their choices to arrive at a final decision. In fact, in this work we focus on modeling this dynamic process using an evolutionary game-theoretic model that we now describe in detail. Let $\mathcal{X} = \{X(t) : t \geq 0\}$ be a continuous time Markov chain (MC) used to model the dynamics of how users change their choices (i.e., strategies in the game theory jargon) as a function of the time. Let A be the set of data rates that available to the users[2]. Set A is known as the strategy

[2]We assume that all users choose a strategy from the same set A. However, different sets of strategies

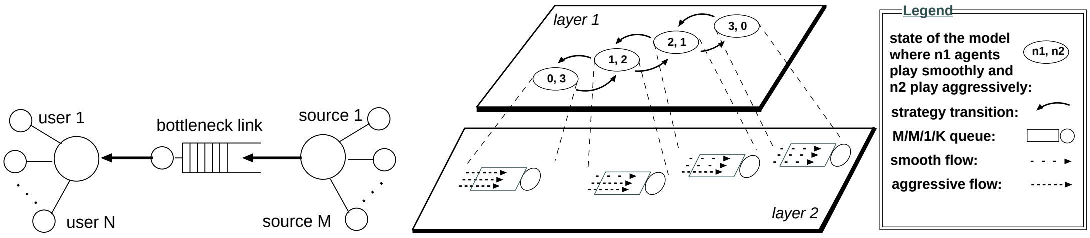

Figure 1. From left to right: (a) The system overview; (b) model snapshot.

set. The model has a finite state space $\mathcal{S}$ and each state represents the number of users adopting each of the available strategies. Every state in $\mathcal{S}$ induces a performance model that will determine the characteristics of the shared network link. This model yields the appropriate performance metrics, such as packet loss and delay, which are then used to calculate the QoS perceived by the users. In the examples that follow we use a simple M/M/1/K performance model to illustrate the basic ideas of our approach. The state transition rates of the MC are then calculated based on the QoS perceived by the users in each of the states. Figure 1(b) illustrates this two-level model.

In our model, users can change their current strategy when they perceive that a change in the data rate will yield a higher QoS for them. This induces "Darwinian" adjustments, in which users gradually move toward a state in which no single user can improve his QoS by making further changes. This state is considered to be a Nash equilibrium. Note that if the model has more than one Nash equilibrium, the one which will be selected in the long run depends on the initial distribution of the population. This is called "path dependence".

We also assume that "perturbations" are possible in our model. Users have a non-zero probability of making a wrong decision and choosing a strategy that yields a lower QoS than the current strategy. For instance, this could occur if users do not correctly infer the impact of changing their data rates. Therefore, at each state, each user can not only select a strategy that will improve its QoS, but can also select a strategy that will decrease its QoS. However, the probability of changing from the present strategy to another is proportional to the QoS gains that result from the change, while the probability of making a mistake is taken to be small. As a consequence of introducing "perturbations", it can be shown that the resulting MC is ergodic. Therefore, the problem of the "path dependence" does not exist since the steady state solution does not depend on the initial conditions.

The dynamic process adopted by users to choose their strategies are based on the following observations: (i) not all users react instantaneously to their environment (inertia hypothesis); (ii) when an user reacts, the reaction is myopic (myopia hypothesis); (iii) there is a small probability that the users change their strategy in the wrong direction ("perturbation" hypothesis); (iv) users are concerned only with their own QoS when choosing a strategy (self-regarding hypothesis). All these observations are reflected in our model and were also discussed in [Kandori et al., 1993]. The inertia hypothesis means that users react asynchronously when changing their strategies. The myopia hypothesis means that users have a limited reasoning capability and do not take into consideration

can be associated with different groups of users as well.

the long run implications of their choices. To summarize, at any point in time users are faced with a strategic game. However, only a single user makes a decision to change his strategy – this yields an asynchronous evolutionary model. For simplicity, we assume that when a user changes his strategy, all other users become aware of the new state of the game, instantaneously.

We now give the details of the evolutionary game-theoretic model, as established in Definition 4. As stated earlier, the dynamic process of strategy adjustment is given by $\mathscr{D}$, a continuous time Markov chain. Let $k = |A|$ denote the number of strategies available to the users. Therefore, each state has k different payoffs which correspond to the QoS value that each user will receive when playing a given strategy. Let $s_i = \langle n_1, \ldots, n_l, \ldots, n_m, \ldots n_k \rangle$ and let $s_j = \langle n_1, \ldots, n_l - 1, \ldots, n_m + 1, \ldots n_k \rangle$ be two states of the MC, where $n_l, 1 \leq l \leq k$ represents the number of users playing strategy l. The process transitions from s_i to s_j when a user changes his strategy from l to m. The transition rate from s_i to s_j is a function of the difference in the QoS received in these two states. Let $U(l, i)$ be the QoS (the utility function) perceived by a user playing strategy l at state s_i, $n_l^{(i)}$ be the number of users in s_i playing strategy l, and $\sigma_l^{(i)} = n_l^{(i)}/N$ the corresponding fraction of users playing l at s_i. The transition rate from s_i to s_j is given by

$$\begin{cases} n_l^{(i)} \times \Big[U(m,j) - U(l,i)\Big] & \text{if } U(m,j) > U(l,i) \\ n_l^{(i)} \times \epsilon & \text{otherwise} \end{cases} \tag{1}$$

Recall that users can make mistakes and transitions from s_i to s_j can occur even if the QoS perceived at s_i is greater than that in s_j. This occurs with rate ϵ, which is a parameter of the dynamic process $\mathscr{D}$. The utility function has the general form $U(l, i) = Q(\text{delay, throughput, loss ratio}, \ldots)$ where Q is an arbitrary function which measures the QoS perceived by a given user. There have been several proposals in establishing an accurate QoS function that maps network performance to the perceived user-level quality when using a multimedia application. For instance, the ITU-T E-Model, assumes that performance measures are additive [ITU, 2003], while other recent work have attempted to obtain different QoS functions [Mohamed and Rubino, 2002]. For simplicity, we will consider additive quantities in our QoS function.

Without loss of generality, we choose only two network performance metrics to influence the QoS perceived by users, one that increases and another that decreases with an increase in the data rate chosen by users: delay and goodput. The utility function is then given by: $U(l, i) = \alpha.\phi_1(d(l, i)) + \beta.\phi_2(g(l, i))$ where $d(l, i)$ is the mean packet delay and $g(l, i)$ is the mean goodput seen by users that play strategy l in state i, respectively; ϕ_1 and ϕ_2 are decreasing and increasing functions, respectively. In other words, the utility function of a user playing strategy l in state i is given by a weighted average of two utility functions. α and β are the weights of the marginal utility functions. Depending on the concavity of ϕ_1 and ϕ_2, we say that the users are risk-averse or risk-loving on a specific QoS parameter [Gintis, 2000]. If ϕ_2 is concave, for instance, the user is risk-averse regarding the goodput achieved. In that case, a good candidate for ϕ_1 and ϕ_2 is $\phi_i = k_0 \times (k_1 + k_2 \log(k_3 + k_4 x))$ where the k_i's are normalization constants. Note that ϕ_1 and ϕ_2 are mainly used to normalize the performance metrics to the $[0 - 1]$ range and to establish how risk-aversive or risk-loving are users when considering each performance metric.

It remains to show how we can obtain $d(l,i)$ and $g(l,i)$. In general, these metrics could be generated via simulation or via an analytical performance model. We adopt the latter strategy. Referring to Figure 1(a) let K be the total buffer space at the bottleneck link, μ the capacity of this link, $\lambda(l)$ the actual transmission rate of associated with strategy l, $\lambda^\star(i) = \sum_l n(l,i)\lambda(l)$ the aggregate transmission rate at state s_i and $\rho(i) = \lambda^\star(i)/\mu$. For the examples in this paper, we use the simple M/M/1/K formulae [Kleinrock, 1975] to obtain the performance metrics of interest. Denote by $p(i)$ and $L(i)$ the loss ratio and mean queue size at the M/M/1/K queue when the MC that represents the system dynamics is in state i. Then $d(l,i) = \frac{L(i)}{\lambda^\star(i)(1-p(i))}$, $g(l,i) = p(i)\lambda(l)$ and:

$$\begin{cases} p(l,i) = \frac{\rho(i)^K(1-\rho(i))}{1-\rho(i)^{K+1}} & L(i) = \rho(i)\frac{(1+K\rho(i)^{K+1})-(K+1)\rho(i)^K}{(1-\rho(i))(1-\rho(i)^{K+1})} \quad \text{if } \lambda^\star(i) \neq \mu \\ p(l,i) = \frac{1}{K+1} & L(i) = \frac{K}{2} \quad \text{otherwise} \end{cases} \tag{2}$$

where $d(l,i)$ does not depend on l in our particular dynamics due to the PASTA property.

A nice analytical property of the proposed ergodic Markovian model consists on the fact that it is hierarchical (see Figure 1). On the upper layer, we model user behavior and their strategy selection, while in the lower layer, we model the effect of network performance metrics on user's QoS. This resembles the modeling paradigm employed in *performability* analysis [de Souza e Silva and Gail, 2001].

Another interpretation for our model is that it describes not actual users but a *distributed congestion control algorithm*. This algorithm attempts only to maximize the QoS that its user will perceive. Each end-host running this algorithm continuously executes a mechanism to detect the QoS perceived by the user as a function of network performance metrics such as packet loss, delay and goodput. The algorithm then chooses different transmission rates in the attempt to obtaining a higher QoS for its user. However, the transmission rate cannot effectively be changed at a high frequency without causing loss of QoS, for instance, without causing flickering in a video. Thus, the probability of making a change in the rate is proportional to the attained increase in QoS. The transitions made at rate ϵ model the measurement inaccuracy or possible decision "errors". These perturbations are also applied in "simulated annealing" in order to prevent the system to remain in a local optimum.

5. Model Analysis

In this section we evaluate the main analytical results obtained for the proposed model. Please, refer to [Menasché et al., 2004] for the proofs.

Definition 5 *Let ϵ be the perturbation rate which is one of the parameters of the dynamic process $\mathscr{D}$, as explained in the previous section. A quasi-absorbing set of states S_a is one that, for each state $s_o \in S_a$, any transition from s_o to a state $s_d \notin S_a$ is equal to ϵ.*

Note that, since our model is ergodic, it does not contain absorbing sets. However, it contains quasi-absorbing sets. According to Definition 2 (see also [Samuelson, 1997]), the Nash equilibrium of a dynamic evolutionary game is a collection of strategies that are optimal when the resulting deviations from the present strategies are evaluated in the closest possible worlds. In the proposed model the notion of closest possible worlds is equivalent to that of adjacency between states. In what follows, we relate quasi-absorbing sets with Nash equilibria.

Definition 6 *The support or carrier $C(q)$ of a probability distribution $q \in \triangle_N$ consists of the states which receive positive probability on $q : C(q) = \{i \in \mathbb{Z} | q_i > 0\}$.*

Let $S_1, S_2, \ldots, S_n$ and T be a partition of the state space of the model where S_i are minimal quasi-absorbing sets, and T is the set of the remaining states.

Proposition 7 *When $\epsilon \to 0$, a state s is a Nash equilibrium of our proposed game $\mathscr{G}$ if and only if it is a singleton quasi-absorbing set.*

Proposition 8 *For each state $t \in T$ there is at least one path from t to a state $s \in S_i$ which does not contain any transition with rate equal to ϵ.*

Proposition 9 *Suppose game $\mathscr{G}$ with dynamic model $\mathscr{D}$ has at least one Nash equilibrium. When $\epsilon \to 0$ let state s be contained in the support of the steady state distribution of $\mathscr{D}$. Then: (i) s is contained in one of the sets S_i (ii) if S_i is a singleton, then s is a pure Nash equilibrium.*

Proposition 10 *Depending on the dynamics $\mathscr{D}$, the Nash equilibria of $\mathscr{G}$ which are on the support of the steady state solution of $\mathscr{D}$ when $\epsilon \to 0$ may receive arbitrarily different probabilities.*

Proposition 9 establishes relations between the Nash equilibria of the game $\mathscr{G}$ with the states on the support of the stationary distribution. Proposition 10 then indicates that even though game $\mathscr{G}$ may admit more than one Nash equilibrium, not all the equilibria necessarily receive high probability in steady state. What determines the probability that a Nash equilibrium will receive in steady state is the dynamics of the system. So, in many cases it is possible to understand the expected behavior in steady state, when $\epsilon \to 0$, even when more than on Nash equilibrium is present. Propositions 9 and 10 concerning our game dynamics are in accordance with results from other dynamical models of evolutionary game theory. Similar propositions have been proved, for instance, in the context of the replicator dynamics. [Gintis, 2000] shows that stable fixed-points (evolutionary equilibria) of the replicator dynamics correspond to Nash equilibria of the characterized games, although not all Nash equilibria will be related to a stable stationary point.

6. Numerical Examples

In this section we present three examples to illustrate a few of the properties of our model and its applicability. The numerical computation was performed with the Tangram-II tool, [Carmo et al., 1998]. The examples analyze the impact of the utility function on the behavior of the agents in steady state, when the error rate ϵ is small.

The first and third examples use as parameters for the utility function the loss ratio and the goodput, while in the second example the utility for each user is given as a function of the delay and the goodput. The values for parameters used in all the examples are shown in Table 2. In this table, QoS parameters with an overline were normalized to the range $(0, 1)$. Consider the scenario characterized in row 1 of Table 2. In this case the utility is a function directly proportional to the goodput and inversely proportional to the loss rate: utility $= \alpha\phi_1(\text{loss ratio}) + \beta\phi_2(\text{goodput})$. The shapes of $\phi_1(\cdot)$ and $\phi_2(\cdot)$ are shown in Figures 2(a) and 2(b) respectively.

Figure 2(c) plots the steady state solution of the model when β varies, for $\alpha = 10$. The impulses in the figure illustrate the values of the probabilities. For instance, when

Table 2. Parameterization of the Examples.

#	K	strategies (λ)	μ	users	QoS parameters i	param.	α	β	ϕ_i i	k_0	k_1	k_2	k_3	k_4
1	10	(7, 8)	80.4	20	1	loss ratio	10	0.4 . . . 2	1	0.14	7	1	1.001	-1
					2	throughput			2	0.14	7	1	0.001	1
2	2	(5, 10, 20)	80.4	4	1	delay	1	0.5 . . . 0.53	1	0.56	-0.6	-1	-1.4	100
					2	throughput			2	0.29	-4.5	1	100	100
3	100	(8, 18.4 64)	4000 (bg traffic of 1400)	40	1	loss ratio	500	1 . . . 180	1	0.14	7	1	1.001	-1
					2	throughput			2	0.14	7	1	0.001	1

Table 3. QoS table for example 1 when $\beta = 0.8$. The bold state is the unique Nash equilibrium.

state	QoS	state	QoS	state	QoS
1: 20 0	(10.7140, ——)	8: 13 7	(9.9785, 10.0947)	15: 6 14	(8.7196, 8.8988)
2: 19 1	(10.6246, 10.7152)	9: 12 8	(9.8452, 9.9672)	16: 5 15	(8.4298, 8.6257)
3: 18 2	**(10.5309, 10.6249)**	10: 11 9	(9.7012, 9.8298)	17: 4 16	(8.0778, 8.2947)
4: 17 3	(10.4325, 10.5302)	11: 10 10	(9.5445, 9.6806)	18: 3 17	(7.6275, 7.8725)
5: 16 4	(10.3288, 10.4305)	12: 9 11	(9.3725, 9.5170)	19: 2 18	(6.9986, 7.2846)
6: 15 5	(10.2192, 10.3252)	13: 8 12	(9.1814, 9.3357)	20: 1 19	(5.9384, 6.2962)
7: 14 6	(10.1027, 10.2136)	14: 7 13	(8.9662, 9.1319)	21: 0 20	(——, 0.8454)

$\beta = 0.4$ Figure 2(c) shows that the steady state probability for state 1 is approximately 1. This indicates that state 1 is the expected steady state Nash equilibrium in the long run for this scenario. Note that all other states in the model have a very small probability, so small that their correspondent impulses are not visible. In this example, we observe that, for any value of β, the probability distribution concentrates most of its mass in a single state.

The example is particularly interesting in order to show the "loss of efficiency" due to miscoordination and the absence of a central authority. Consider the value $\beta = 0.8$. In this case, there is just one Nash equilibrium for this game, namely state 3. However, the Nash equilibrium is extremely inefficient. State 1, which corresponds to the Pareto optimum, gives a better payoff for all the twenty players. But state 1 cannot be an equilibrium point in this non-cooperative scenario because, in this state, a selfish agent would be inclined to change to the aggressive strategy (see Table 3). Therefore, for state 1 to be an equilibrium state, some kind of negotiation between the players (or the use of a central authority) would be required to produce incentives for the players to stay in that state. In summary, the greater the impact of the goodput on the global utility, the greater the number of agents playing aggressively in the long run. As a consequence, due to the absence of a central authority the players tend to play aggressively and do not reach the best strategy for the group.

Consider a second scenario, where four agents share a bottleneck, each with three available strategies to choose from. In this example the utility is a function of the delay and the goodput. Figure 2(d) shows that, as the goodput influence on the utility function increases, state 15 concentrates almost all the probability mass, and this is the state where all agents play aggressively.

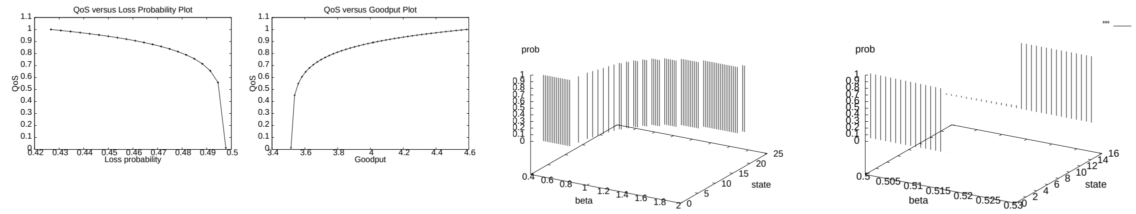

Figure 2. The first 3 figures correspond to case 1 (a) maps loss probability to QoS; (b) maps throughput to QoS; (c) the steady state solution as a function of β; (d) steady state solution for case 2.

Table 4. QoS table for example 2 when $\beta = 0.52$. Bold states are Nash equilibria.

state	QoS	state	QoS	state	QoS
1: 4 0 0	**(1.270, 0.000, 0.000)**	6: 2 0 2	(0.554, 0.000, 0.736)	**11: 0 4 0**	**(0.000, 0.772, 0.000)**
2: 3 1 0	(1.036, 1.126, 0.000)	7: 1 3 0	(0.771, 0.860, 0.000)	12: 0 3 1	(0.000, 0.641, 0.736)
3: 3 0 1	(0.771, 0.000, 0.955)	8: 1 2 1	(0.613, 0.701, 0.796)	13: 0 2 2	(0.000, 0.546, 0.640)
4: 2 2 0	(0.883, 0.972, 0.000)	9: 1 1 2	(0.503, 0.590, 0.684)	14: 0 1 3	(0.000, 0.471, 0.565)
5: 2 1 1	(0.684, 0.772, 0.867)	10: 1 0 3	(0.420, 0.000, 0.600)	**15: 0 0 4**	**(0.000, 0.000, 0.504)**

A state where all agents play the same strategy is called a *convention*. Consider the value $\beta = 0.52$. From Table 4 we can observe that all the conventions are Nash equilibria of the game. However, from Figure 2(d) we see that in this scenario just one state receives a probability close to one in steady state. This observation confirms Proposition 10. Although we have more than one Nash equilibria, the dynamics of the system will determine the state which will receive greater probability in steady state [Young, 1993]. This example serves to illustrate how our model dynamics answers the problem of the Nash Equilibrium selection: if more than one Nash equilibrium is present in the game, which one will be played most frequently in the long run?

Voice over IP applications are becoming common in the Internet. Applications such as Skype, freephone and VivaVoz [Figueiredo et al., 1997, Duarte et al., 2003] and free voice codecs, such as Speex [GNU, 2004], are presently available. Using Speex, for instance, users may choose to transmit/receive compressed voice in the range from 2.4 kbps to 24.8 kbps or choose PCM encoding at 64 kpbs, based solely on the perceived QoS.

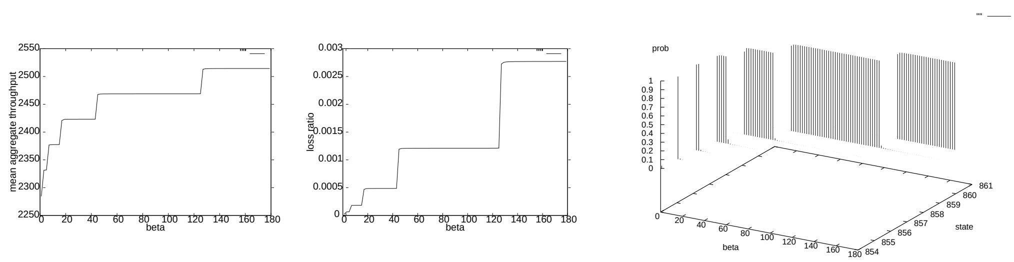

Figure 3. Results for case 3 (a) mean aggregate throughput as a function of β; (b) mean loss ratio as a function of β; (c) steady state solution.

For the third example we consider a bottleneck link with capacity 4 Mbps which is shared by 40 users using a VoIP application. The link is also shared by other applications that generate a background traffic at a rate of 1.4 Mbps (see third line of Table 2). As before, all the parameters values are maintained fixed except for β. Figures 3(a) and 3(b) show the average aggregate throughput and loss rate with β, respectively. Figure 3(c) shows the evolution of the steady state model solution with β. In the figure, the model states are ordered with increasing mean aggregate throughput. Based on these figures, we conclude that, in this scenario, the average aggregate throughput of the Nash equilibrium chosen in the stationary distribution increases monotonically with the weight of the throughput on the users' utility function. However, the loss probability also increases, and the steady state depends on the relative importance of these measures on the QoS.

Finally, one of the interesting properties of our model is the small sensitivity of the equilibrium with respect to small perturbations in the adjustment process of the agents. This is an indication of the robustness of the model. Refer to [Menasché et al., 2004] for details.

7. Conclusion and future work

It is a standard procedure in the literature of economics to try to model the behavior of the human being in order to predict market outcomes [Mas-Colell et al., 1995]. In the literature of networks, this area is rapidly growing. However, until now game theory has been used by the computer networking community mainly as a normative exercise to investigate how decisions *should* be made. In this work, we try to bring a new insight into the problem, and propose a dynamic model to investigate the impact generated by how decisions *are* made. We consider users who make their choices based on the QoS they perceive using *trial and error*. This is why we adopt an evolutionary game theoretic modeling framework [Samuelson, 1997].

We proposed a dynamic model to analyze the scenario where there are no guidelines imposed on users when determining the rate at which data should be transfered, except for their perception of quality. We believe this is an interesting approach to modeling the congestion control problem, specially when considering multimedia applications, such as voice and video. We do not claim that our model is the most accurate representation of the dynamic process by which users choose their rates. However, this is an initial effort in shedding some light into the new research area of behavioral game theory applied to computer networking.

References

Akella, A., Seshan, S., Karp, R., Shenker, S., and Papadimitriou, C. (2002). Selfish behavior and stability of the internet: A game theoretic analysis of tcp. In *ACM SIGCOMM*, volume 32, pages 117–130.

Carmo, R., de Carvalho, L., de Souza e Silva, E., Diniz, M., and Muntz, R. (1998). Performance/Availability Modeling with the TANGRAM-II Modeling Environment. *Performance Evaluation*, 33:45–65.

de Souza e Silva, E. and Gail, H. R. (2001). The uniformization method in performability analysis. In B. R. Haverkort, R. Marie, G. R. and Trivedi, K. S., editors, *Performability Modelling: Techniques and Tools*, chapter 3, pages 31–58.

Duarte, F., de Souza e Silva, E., and Towsley, D. (2003). An Adaptive FEC Algorithm using Hidden Markov Chains. *Performance Evaluation Review*, 31(2):11–13.

Figueiredo, D., Duarte, F., and de Souza e Silva, E. (1997). Implementação de um aplicativo de voz com análise e caracterização de seu tráfego. In *XXIV SEMISH, XVI CTIC*, pages 531–540, Brasília, DF.

Floyd, S., Padhye, J., and Widmer, J. (2000). Equation-based congestion control for unicast applications. In *ACM SIGCOMM*, pages 43–56.

Gintis, H. (2000). *Game Theory Evolving*. Princeton University Press.

GNU (2004). *Speex: a free codec for free speech*. URL: http://www.speex.org/.

Greenwald, A., Friedman, E., and Shenker, S. (2001). Learning in network contexts: Experimental results from simulations. *Journal of Games and Economic Behavior*, 35(1/2):80–123.

Hsiao, M.-T. and Lazar, A. (1987). A game theoretic approach to decentralized flow control of markovian queueing networks. In *Performance '87*, pages 55–74.

ITU (2003). *ITU-T Recommendation G.107: The E-Model, a computational model for use in transmission planning*.

Johari, R. and Tsitsiklis, J. N. (2004). Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435.

Kandori, M., Mailath, G., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kleinrock, L. (1975). *Queueing Systems, Volume I: Theory*. Wiley Interscience.

Kurose, J. and Ross, K. (2004). *Computer Networking : A Top-Down Approach Featuring the Internet*. Addison Wesley, 3rd edition.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.

Menasché, D. S., Figueiredo, D. R., and de Souza e Silva, E. A. (2004). Extended version of this article. URL: http://www.land.ufrj.br/∼sadoc.

Mohamed, S. and Rubino, G. (2002). A study of real-time packet video quality using random neural networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 12(12):1071–.

Nagle, J. (1985). On packet switches with infinite storage. *Request for Comments (RFC970)*.

Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press.

Rejaie, R., Handley, M., and Estrin, D. (1999). Rap: An end-to-end rate-based congestion control mechanism for realtime streams in the internet. In *IEEE INFOCOM*, pages 1337–1345.

Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. The MIT Press.

Shenker, S. (1995). Making greed work in networks: A game-theoretic analysis of switch service disciplines. *ACM/IEEE Transactions on Networking*, 3:819–831.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.