An evolutionary game-theoretic approach to congestion control

D.S. Menasché^a, D.R. Figueiredo^b, E. de Souza e Silva^{a,*}

^a Computer Science Department, IM and COPPE, Federal University of Rio de Janeiro, 21945-970 Rio de Janeiro, Brazil

^b Department of Computer Science, University of Massachusetts, Amherst, MA 01003, USA

Available online 19 August 2005

Abstract

This paper investigates a system where a set of users sharing a bottleneck link must choose the transmission rate at which multimedia traffic is received. Users are assumed to be self-regarding and make their decisions with the sole goal of maximizing their perceived quality. We are interested in the dynamic process by which users adapt their data rates and the convergence of this process to equilibria. We propose a novel two-layer model to represent this system: the upper layer is an evolutionary game-theoretic model that captures how users adapt their rates; the lower layer model captures the network performance and the quality perceived by the users. Using the model proposed, we demonstrate analytically and numerically several interesting properties of the system equilibria. In particular, we establish the relationship between system states that have non-negligible steady state probabilities and Nash equilibria of the induced game.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Congestion control; Game theory; Markov chain; VoIP; Audio streaming; UDP

1. Introduction

Congestion control is one of the most fundamental problems in computer networks and has been widely studied for decades. In essence, the problem consists of determining how bandwidth is shared among

* Corresponding author.

E-mail addresses: sadoc@land.ufrj.br (D.S. Menasché), ratton@cs.umass.edu (D.R. Figueiredo), edmund@land.ufrj.br (E. de Souza e Silva)

different data flows traversing a congested link. In today's Internet, where the end-to-end principle has been strongly adopted, congestion control is performed solely by the end-hosts. In particular, end-hosts run TCP, a window-based congestion control mechanism that regulates the data rate entering the network. TCP is also the dominant transport protocol used by applications for transferring data, which guarantees the stability of the Internet.

However, it is generally accepted that TCP is not suited for transporting delay sensitive data such as voice and video traffic. With the increase in the demand for such multimedia applications, transporting this type of data effectively becomes an important issue. The UDP protocol provides a viable alternative and has been widely adopted as the default transport protocol by many real-time applications and protocols, such as RTP. Although UDP does not perform congestion control, researchers have advocated that applications or protocols that use UDP should also implement some form of congestion control [8]. Researchers have worried that the lack of any congestion control could lead to unfairness issues among different applications and even to congestion collapse.

The deployment of congestion control mechanisms by multimedia application developers that adopt UDP as their transport protocol is completely voluntary. In fact, competing developers are likely to implement protocols that best suit their own needs and that yield the best possible media quality to their users. It is reasonable to assume that applications will dynamically regulate their sending rates with the single goal of maximizing the media quality perceived by their users. Note, however, that this rate control may implicitly provide a form of congestion control, as arbitrarily increasing the data rate may not provide the best media quality to a user. One possibility is not to implement any congestion control mechanism and transfer the responsibility of regulating the data rate to the users. In the context of voice and video applications, for example, tools such as VivaVoz [7], and encoders such as Speex [11], already permit users to freely choose the rate at which data is sent. The dynamic rate adjustment is accomplished by allowing users to change, during a session, the amount of redundancy sent by the application, the encoder being used, and associated parameters, such as the number of layers being sent in a multilayered video or audio stream.

The idea of allowing applications or users to determine their data rates as an implicit mechanism for congestion control constitutes a broad field of research. In this scenario, users have an utility function which depends on performance characteristics of the data transfer (e.g., throughput and delay) and are assumed to be *self-regarding*. This latter assumption means that applications are only interested in maximizing the utility provided to their respective users (a self-regarding application is one which is dynamically configured, automatically or in an ad-hoc fashion, by a self-regarding user, in order to improve the user's QoS). Note that applications are competing for a shared resource and decisions made by one application 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.

In this paper, we consider the scenario where a number of self-regarding applications share a common bottleneck link and dynamically adapt their data rates to maximize the media quality delivered to their respective users. In particular, we are interested not only in the equilibria of such system, but also on the dynamic process by which such applications adapt their rates. To characterize this process, we propose a two-layer dynamic game-theoretic model based on evolutionary models and performance models. The evolutionary models determine how applications' data rates evolve with time, while the performance models are used to quantify the instantaneous media quality. We provide both analytical and numerical evaluation of our model. Analytically, we show that, in steady state, there is a strong relationship between the states that have non-negligible probabilities and

the Nash equilibria. Numerically, we apply the well-studied E-model to illustrate properties of the equilibria.

The main contribution of this paper is the modeling framework we introduce which combines evolutionary models from game theory with performance models that allows us to understand the behavior of self-regarding applications (or users). Although we focus on the problem of multimedia applications sharing a common bottleneck link, our model is general enough and it could be applied in other contexts, such as path switching [28].

The remainder of this paper is organized as follows. In the next section we provide a brief introduction to game theory and evolutionary games, and also present and discuss the related work on this area. Section 3 describes in detail the proposed model. Section 4 provides the analytical results concerning properties of the equilibria of the model and in Sections 5 and 6 we present a numerical evaluation of the model. Finally, Section 7 summarizes the paper.

2. Background and related work

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 options. In game theory nomenclature, the 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. More formally, we have [25]:

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 $u_i : A = \times_{j \in N} A_j \rightarrow \Re$ (the utility function of player i).

In general, one is interested in determining the choices that players will make when faced with a particular game, which is sometimes referred to as the *solution* of the game. We will adopt the most common *solution* concept, known as Nash equilibrium: 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 of choices $\mathbf{a}^* = (a_1^*, \dots, a_N^*) \in A = \times_{j \in N} A_j$, one for each player, such that for each $i \in N$, $u_i(\mathbf{a}_{-i}^*, a_i^*) > u_i(\mathbf{a}_{-i}^*, a_i)$ for all $a_i \in A_i$ [where $\mathbf{a}_{-i} = (a_1, a_2, \dots, a_{i-1}, a_{i+1}, \dots, a_N)$ is a vector of choices, one for each player, except player i].

One of the core issues in game theory is the fact that the Nash equilibrium frequently does not coincide with a point that yield high utility values to all players. This last observation motivates the definition of a Pareto equilibrium: a set of choices such that there is no other 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 of choices $\mathbf{a}^* = (a_1^*, \dots, a_N^*) \in A = \times_{j \in N} A_j$, one for each player, such that there is no $\mathbf{b} \in A$ that satisfies $u_i(\mathbf{b}) > u_i(\mathbf{a}^*)$ for all $i \in N$.

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 assumption is too stringent because of the high complexity to determine the Nash equilibrium and that multiple Nash equilibrium can co-exist.

Evolutionary game theory has emerged as an alternative perspective to classical game theory. It considers a dynamic scenario where players are constantly interacting with one another and adapting their choices based on the utility they receive. 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. Besides that, game 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. An evolutionary game, in its simplest form, consists of players that have identical sets of choices and symmetric utility functions.

Definition 4. A symmetric **evolutionary game** $\langle N, |M|, (A), (u_j), \sigma(t), \boldsymbol{\sigma}(t), \mathcal{D} \rangle$ models the interaction of players over time, and is composed of the following: (i) a finite set N of players, who are matched in groups of size $|M|$; (ii) a non empty set A (the set of available choices for each player), with $k = |A|$; (iii) for each player j , an utility function $u_j : A^{|M|} \rightarrow \mathbb{R}$ (the utility function of each player); (iv) a vector of values between 0 and 1, $\boldsymbol{\sigma}(t) = (\sigma_1, \dots, \sigma_k)(t)$, where $\sigma_j(t)$ is the fraction of the population selecting the choice $a_j \in A$ at time t (the population state at time t); (v) an ordinary game $\mathcal{A}(t) = \langle M, (A), (u_j) \rangle$ (the “stage game” at time t); (vi) a dynamic process of strategy adjustment, which is a function (possibly stochastic) $\mathcal{D} : (\boldsymbol{\sigma}(t), \mathcal{A}(t), t) \rightarrow \boldsymbol{\sigma}(t + \Delta t)$ (the dynamic selection process).

Players are repeatedly matched in groups of $|M|$ in order to play the stage game $\mathcal{A}(t)$ at time t . The stage game is played in accordance with $\boldsymbol{\sigma}(t)$, the fraction of the population selecting each choice. Given the utilities obtained, players adjust their choices according to the dynamics $\mathcal{D}$, which can be probabilistic, based for instance on a Markov chain. In this case, the state of the MC is usually given by a vector representing the number of players making a given choice [18,29].

The idea of approaching the congestion control problem in computer networks using a game-theoretic formulation is not new. Nagle was one of the first to suggest that game theory could be applied to congestion control [24], and there is a growing literature on this relationship. In a landmark paper, Shenker characterized the Nash and Pareto equilibria of a bottleneck link under different service disciplines [27].

Altman and Başar [1] consider the scenario where each user requests a desired quality of service (QoS) to the network, before the game starts. Then, the authors propose a congestion control scheme in order to regulate the rates of the users in accordance to the QoS requested. In this paper, however, we do not propose an optimal control mechanism. Instead, we study a scenario which resembles that suggested by Key and McAuley [20, Section 4], completely asynchronous and characterized by the lack of agreements, where greedy users select their rates solely based on the perceived *application level* QoS [28].

The problem of handling data flows that are not regulated by any congestion control mechanism (e.g., unresponsive flows) has also been recently investigated. For example, Garg et al. [9] propose a new service discipline to punish unresponsive flows, while Chandrayana and Kalyanaraman [3] propose a policing mechanism at the network edge to manage uncooperative flows. In this line of work, it is usually assumed that users’ utility functions are such that there is no benefit in reacting to network congestion. This is not the case considered in this paper. Here, the perceived QoS depends on multiple network characteristics and can increase when a user reacts to congestion by reducing its sending rate.

When dealing with any game-theoretic formulation, the problem of *equilibria selection* naturally emerges. This problem addresses the following two important issues: which equilibrium is reached (in case the formulation at hand yields multiple equilibria); how can the equilibrium be reached. Shenker has addressed this issue in the context of a specific service discipline, describing a hill climbing algorithm that converges to the sole Nash equilibrium of the game [27]. More generally, Greenwald et al. has addressed the problem of equilibria selection in the context of dynamic network games [13]. They use extensive simulations of simple dynamic learning algorithms to understand convergence on such games. However, the congestion control problem is only briefly discussed.

Evolutionary game-theoretic models have been widely studied in the economics literature. These models capture the dynamic process by which players of a game adapt their strategies. Two prominent models are the works of Kandori et al. [18] and Young [29] (generalized by Samuelson [26]), which provide strong analytical results on the Nash equilibria selection problem. Although these two models share a common basis, they differ considerably in the details. Both consider a finite player population, a discrete time model and assume that players can make mistakes when changing their strategies. However, their dynamics of the strategy selection are significantly different.

We propose a novel model based on evolutionary game-theoretic models to capture the strategy selection process of the data rate chosen by applications (or users) that share a bottleneck link. A performance model for the bottleneck link is employed to determine the user-level quality of the media being received. Although our model for dynamic strategy selection draws on ideas exposed in [18,29], there are some fundamental differences, making it 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 bidimensional stage game. Fourth, the strategy selection process we consider is very different from those in the previous models. Section 3 describes our model in more detail.

3. The model

We consider the problem where a finite number of self-regarding applications (or equivalently, users) share a common bottleneck link to transfer multimedia data. Users can dynamically select the rate at which data is received from a discrete set of available rates. We also assume that users have a well-defined utility function which depends on performance characteristics of the received data (e.g., throughput and delay). This utility function is determined by the user-level QoS and can drastically vary with the link performance. The shared bottleneck link imposes performance characteristics on the data transfer, which in turn determine the quality of the media being received. The goal of each user is to maximize his own QoS by selecting a specific data rate. Fig. 1(a) illustrates the scenario considered here, with a bottleneck link shared by N users.

It is important to note that we are interested not only in the final equilibria, but also in the actual dynamic process by which users adapt their data rates. To investigate this phenomena, we focus on a modeling framework that captures this dynamic process using an evolutionary game-theoretic model coupled with

Fig. 1. (a) System overview; (b) model snapshot.

a performance model. In the following, we describe this model in detail, first presenting the stage game and then the dynamic process.

Given the scenario discussed above, users are faced with the following symmetric strategic game: $\mathcal{G} = \langle N, (A), (U_i) \rangle$, where N is the set of users sharing the bottleneck link and A the set of data rates available to the users. The utility function $U_i : \times_{j \in N} A \rightarrow \mathbb{R}$ determines the QoS perceived by user i at each possible state of the game. In the jargon of evolutionary game theory, this is the system's "stage game".

We propose a continuous time Markov chain (MC) $\mathcal{X} = \{X(t) : t \geq 0\}$ to model the dynamic process representing how users change their choices of data rate as a function of time. Let $k = |A|$ denote the number of strategies available to the users. The model has a finite state space $\mathcal{S}$ and each state is described by the number of users adopting each of the available strategies. Thus, $s_i = \langle n_1, \dots, n_k \rangle \in \mathcal{S}$ is a state of the model where n_l , $1 \leq l \leq k$ represents the number of users choosing strategy l . Every state in $\mathcal{S}$ induces a performance model that will determine the characteristics of the shared bottleneck link. This model yields the appropriate performance metrics, such as throughput and delay, which are then used to calculate the QoS perceived by the users. These QoS values will also be used to determine the transition rates among the states of the MC. Fig. 1(b) illustrates our two-level model for the case with three users ($N = 3$) and two different data rates ($k = 2$).

The state transition rates of the MC will determine the dynamic process by which users change their data rate. There are several important features of this dynamic process and our model captures the following characteristics: (i) users change their strategy to improve their QoS (Darwinian adjustments); (ii) when a user changes strategy, this triggers a local improvement in his QoS (myopia hypothesis); (iii) there is a small chance that users change their strategy to a state where the perceived QoS worsens (perturbation hypothesis); (iv) not all users react instantaneously to their environment (inertia hypothesis); (v) users are concerned only with their own QoS when choosing a strategy (self-regarding hypothesis).

The Darwinian adjustments means that users change strategy to improve their QoS and that changes that yield a higher QoS are more likely to occur. This dynamic process may gradually move the system into a state in which no single user can unilaterally improve his QoS by making further changes. This state is considered a Nash equilibrium of the model. The myopia hypothesis means that users have a limited reasoning capability and do not take into consideration the long run implications of their choices.

The perturbation hypothesis means that users have a non-zero probability of changing to a strategy that yields a lower QoS than the current strategy. For instance, this could occur if users do not fully know the impact of changing their current data rate to a new rate. Therefore, at each state, each user can not only select a strategy that will improve his QoS, but can also select a strategy that will decrease his QoS. However, we consider that the probability of making such mistakes is small compared to changing to a strategy that yields a higher QoS.

The inertia hypothesis means that users react asynchronously when changing their strategies. This is appropriate for the scenario we consider, as users can not easily coordinate their changes. Thus, in our model, at most one user can change strategies at any given instant of time. Finally, the self-regarding hypothesis means that users are oblivious to other users' QoS and to the performance of the bottleneck link. The characteristics reflected in our model were also discussed in [18,29].

We now present the details of the dynamic process of strategy adjustment, which is determined by the state transition rates of the MC introduced above. Recall that each state of the MC defines an outcome for the stage game. Thus, each state yields k different payoffs which correspond to the QoS value that each user will receive when playing a given strategy in that state. Since we consider symmetric utility functions [10], we write $U(l, s)$, where $l \in A$ and $s \in \mathcal{S}$, to denote the QoS perceived by a user that chooses strategy l when the state of the game is given by s . Let $s_i = \langle n_1^{(i)}, \dots, n_l^{(i)}, \dots, n_m^{(i)}, \dots, n_k^{(i)} \rangle$ and let $s_j = \langle n_1^{(i)}, \dots, n_l^{(i)} - 1, \dots, n_m^{(i)} + 1, \dots, n_k^{(i)} \rangle$ be two states of the MC, where $n_l^{(i)}$, $1 \leq l \leq k$ represents the number of users playing strategy l in state s_i . The process transitions from s_i to s_j when a user changes from strategy l to strategy m . The transition rate from s_i to s_j is a function of the difference in the QoS received in these two states:

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

We consider a scenario where users have imperfect information about the *application level* QoS received (analog to the replicator dynamics [10, chapter 9]). So, the larger the gain envisioned by a user from a certain change of strategy, more likely the user will effectively modify the strategy in that way. Indeed, Eq. (1) captures the fact that the better is a choice, the greater is the priority given to it.

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. We assume that ϵ is smaller than the greatest difference of QoS perceived by a user. As a consequence of introducing perturbations, it can be shown that the resulting MC is ergodic. Thus, there is a unique stationary distribution and as $t \rightarrow \infty$ the state distribution along any sample path approaches the stationary distribution.

Our model is hierarchical (Fig. 1). On the upper layer, we model the behavior of users and their strategy selection, while in the lower layer, we model the effect of network performance metrics on QoS. This resembles the modeling paradigm employed in *performability* analysis [14].

4. Model analysis

In this section we investigate the problem of equilibria selection. Note that the underlying N player game associated with the model may have multiple Nash and Pareto equilibria. Since the model is ergodic and thus has a unique steady state behavior, we would like to understand the relationship between the Nash equilibria of the underlying game and the steady state solution of the model.

In particular, we would like to characterize this relationship when users become less likely to make mistakes, that is, as the perturbation parameter in the model, ϵ , goes to zero. In what follows we show that all Nash equilibria have a non-zero probability in the steady state distribution while a certain set of

states has zero probability. However, different Nash equilibria can have different steady state probabilities, which reflects the preference of the user population over the different equilibria.

Definition 5. Let ϵ be the perturbation rate which is one of the parameters of the dynamic process $\mathcal{D}$, 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 $o(\epsilon)$.¹

Note that, since our model is ergodic, it does not contain absorbing sets. However, it does contain quasi-absorbing sets. 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 [26]. 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.

Let $\mathcal{P} = \{S_1, S_2, \dots, S_n, T\}$ 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 (T necessarily does not contain quasi-absorbing subsets). By minimal, we mean that no proper subset of S_i is a quasi-absorbing set.

Proposition 6. When $\epsilon \rightarrow 0$, a state s is a Nash equilibrium of our proposed game $\mathcal{G}$ if and only if it is a singleton quasi-absorbing set.

Proof. By definition, a Nash equilibrium is a state where no agent has unilateral incentives to deviate from his current strategy. From Definition 5 and Eq. (1) a singleton quasi-absorbing set also characterizes a state such that, when $\epsilon \rightarrow 0$, no agent has incentive to change strategy. So, the definitions of Nash equilibrium and singleton quasi-absorbing set coincide. $\square$

The **support or carrier** of a probability distribution π consists of the states receiving positive probability on π : $C(\pi) = \{s_i \in \mathcal{S} | \pi_i > 0\}$. For convenience, let $C_P(\pi)$ be the support of the steady state solution π of the MC with transition probability matrix P . Now, we will consider a system with state space $\mathcal{S} = S \cup T$ and transition probability matrix P , where each principal submatrix $(P_{11}, P_{22}, \dots, P_{M-1M-1}, P_{MM})$ corresponds to the particular subset of states $(S_1, S_2, \dots, S_{M-1}, T)$. The following theorem establishes that when $t \rightarrow \infty$ and $\epsilon \rightarrow 0$, taking the limits in this order, the states in T will not be on the support of the steady state solution of the system.

Lemma 7. Consider an ergodic MC with state space partitioned according to $\mathcal{P}$ (Fig. 2(a)). For each state $t \in T$ there is at least one state $s \in S_1 \cup \dots \cup S_{M-1}$ and one path θ from t to s such that θ does not contain any transition which depends on ϵ .

Proof. By contradiction, assume that the Lemma does not hold for a maximal set of states $T_0 \subseteq T$. Then, from each state $t_{T_0} \in T_0$ all the transitions to states outside T and to states inside $T \setminus T_0$ are equal to $o(\epsilon)$ (if there were transitions different from $o(\epsilon)$, t_{T_0} would not be in T_0). Hence, we conclude that T_0 is a quasi-absorbing set. But this contradicts the fact that T does not contain quasi-absorbing subsets. $\square$

Theorem 8. Consider an ergodic MC with transition probability matrix P_ϵ and state space $\mathcal{S}$ partitioned according to $\mathcal{P}$ (Fig. 2(a)). Assume that each strictly positive element of every block P_{Mj} ($j \neq M$) is not $o(\epsilon)$. In addition, all the elements of P_ϵ , except those in the diagonal, are either $o(\epsilon)$ or do not depend on ϵ at all. Then, when $\epsilon \rightarrow 0$, the states characterized by the block P_{MM} will not be on the support of the steady state solution.

¹ Unlike the standard definition, we say that $f(\epsilon)$ is $o(\epsilon)$ if $\lim_{\epsilon \rightarrow 0} f(\epsilon) = 0$.

Fig. 2. (a) Sample MC partition; (b) another illustrative MC.

Proof. Let π_ϵ be the steady state solution of the MC $\mathcal{X}_\epsilon = \{X(t) : t \geq 0\}$ with transition probability matrix P_ϵ and π_0 be the solution of a related MC with transition probability matrix P_0 obtained from P_ϵ by setting all the $o(\epsilon)$ entries to zero (the MC in this case is not ergodic). We are interested in $\lim_{\epsilon \rightarrow 0} \pi_\epsilon$. Note that $\lim_{\epsilon \rightarrow 0} \pi_\epsilon \neq \pi_0$ since the latter depends on the initial state of the MC, while the former does not. For ϵ small and positive, and well placed so the MC is irreducible even when $\epsilon \rightarrow 0$, Karr [19] showed that the stationary distribution is a continuous function of the transition probabilities. However, our limiting chain (when $\epsilon \rightarrow 0$) is not ergodic, so Karr's result does not hold.

We use aggregation/disaggregation theory [23] to prove that the states of T are not on $\lim_{\epsilon \rightarrow 0} C P_\epsilon(\pi_\epsilon)$. We aggregate the MC $\mathcal{X}_\epsilon$ as follows: all the states in $S_1, \dots, S_{M-1}$ are aggregated into a single state s_s , and the remaining states are aggregated into s_t . Let p and q be the transition probabilities from s_s to s_t and from s_t to s_s , respectively. We construct another MC $\mathcal{X}'_\epsilon$ with two states, s'_s and s'_t , such that $p' \geq p$ and $q' \leq q$ (Fig. 2(a)). So, the fraction of time in s'_t upper bounds that of s_t .

Let $\pi^s(\pi^t)$ be the conditional steady state probabilities of states in $S = S_1 \cup \dots \cup S_{M-1}$ (T). Also, let P_{ST} be submatrix of P_ϵ which gives the transitions from elements in S to elements in T . From aggregation theory, p is equal to $\pi^s P_{ST} \mathbf{e}$, where $\mathbf{e}$ is a vector with all elements equal to one. Since P_{ST} is a finite matrix consisting of elements that are equal to 0 or $o(\epsilon)$, and the sum of the elements of π^s is 1, then p is a weighted sum of elements equal to 0 or $o(\epsilon)$. Therefore, p is upper bounded by $o(\epsilon)$. Since $o(\epsilon) \geq p$, we may choose $p' = o(\epsilon) \geq p$.

Let $\Gamma = \{\gamma \in T | \exists s \in S \text{ s.t. } p_{\gamma s} \neq o(\epsilon) \text{ and } p_{s\gamma} \neq 0\}$ and $P_{\Gamma S}$ be the transition probability matrix from states in Γ to states in S . So, $P_{\Gamma S}$ is such that all its lines have at least one strictly positive element that is not $o(\epsilon)$. Let $\pi^t = \langle \pi^t_{\bar{\Gamma}}, \pi^t_{\Gamma} \rangle$, where $\bar{\Gamma} = T \setminus \Gamma$ is the complement of Γ .

Now we prove that there is at least one element $s_l \in \Gamma$ such that π^t_l is not $o(\epsilon)$. Independently of ϵ , we have $\sum_i \pi^t_i = 1$. So, there is at least one element k in π^t that is different from $o(\epsilon)$ – i.e., there is a state s_k such that $\pi^t_k \neq o(\epsilon)$. If $s_k \in \Gamma$ the proposition is trivially true. If $s_k \in \bar{\Gamma}$ then, from Lemma 7, there is a state $s_l \in \Gamma$ such that s_l may be reached from s_k through a finite path θ with transitions different from $o(\epsilon)$. Analyzing the flow balance equations for state s_k and its successors in θ , it is clear that all the elements of the path θ , including $s_l \in \Gamma$, are such that $\pi^t_m \neq o(\epsilon) (\forall s_m \in \theta)$.

Let $\gamma_{\min}^t = \min_{i|s_i \in \Gamma, \pi^t_i \neq o(\epsilon)} \{\pi^t_i p_{is}\}$ where $p_{is} = \sum_{s \in S} p_{is}$. Then, $q' = \gamma_{\min}^t$ is a lower bound for $q = \pi^t P_{ts} \mathbf{e}$ that is not $o(\epsilon)$. The steady state solution of the MC $\mathcal{X}'_\epsilon$ is $(\pi'^s, \pi'^t) = (q'/(q' + o(\epsilon)), o(\epsilon)/(q' + o(\epsilon)))$. So, when $\epsilon \rightarrow 0$, $\pi'^t_i \rightarrow 0$. Hence, the states in T are not on $\lim_{\epsilon \rightarrow 0} C P_\epsilon(\pi_\epsilon)$. This occurs because the probabilities of these states is upper bounded by π'^t . $\square$

Corollary 9. *Suppose game $\mathcal{G}$ with dynamic model $\mathcal{D}$ has at least one Nash equilibrium. When $\epsilon \rightarrow 0$ let state s be contained in the support of the steady state distribution of $\mathcal{D}$. 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.*

Proof. This follows directly from Theorem 8 and Proposition 6. Clearly, by Eq. (1), a consequence of the dynamics $\mathcal{D}$ is that all the transitions in the model have rates which are either equal to $o(\epsilon)$ or do not depend on ϵ at all. Also from Eq. (1), there will be no transitions equal to $o(\epsilon)$ from states in T to states outside T .

Notice that Theorem 8 refers to a Discrete Markov Chain, while our dynamics $\mathcal{D}$ is based on a Continuous Markov Chain. However, it is possible to transform a CTMC into an “equivalent” Discrete Time Markov Chain (DTMC) subordinated to a Poisson process using uniformization [5,16]. Let Λ be the uniformization rate, i.e. any value greater than the largest absolute value in the diagonal of $\mathbf{Q}$. Dynamics $\mathcal{D}$ implies that the output rate from any state to one of its neighbors is at most equal to $N\Delta$, where Δ is the greatest payoff difference perceived by a user and N is the number of users in the system [see Eq. (1)]. Therefore, we may set $\Lambda = N^2\Delta + \delta$, where δ is an arbitrary positive value. Since Λ does not depend on ϵ , Theorem 8 can be applied to $\mathbf{P} = \mathbf{I} + \mathbf{Q}/\Lambda$, where $\mathbf{I}$ is the identity matrix. And the result follows for the process characterized by $\mathbf{Q}$. $\square$

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

Proof. Consider, for instance, MC of Fig. 2(b). The steady state solution of this system is $[\alpha/(\alpha + \beta + o(\epsilon)), o(\epsilon)/(\alpha + \beta + o(\epsilon)), \beta/(\alpha + \beta + o(\epsilon))]$. If $\alpha \gg \beta$, one Nash equilibrium will concentrate almost all the probability mass. $\square$

Proposition 10 indicates that even though game $\mathcal{G}$ 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 $\epsilon \rightarrow 0$, even when more than one Nash equilibrium is present. Corollary 9 and Proposition 10 concerning our game dynamics are in accordance with results from other dynamical models of evolutionary game theory. Similar lemmas have been proved, for instance, on the context of the replicator dynamics. Gintis [10] 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.

5. A performance and QoS model

In order to solve the proposed model numerically, we must determine a specific performance model for the bottleneck link and a specific user-level QoS model. To illustrate, we will use the M/M/1/K model [21] to yield the performance metrics from the bottleneck link and the ITU E-model [15] to yield the user-level QoS. Note that other models could be easily coupled to our framework.

The M/M/1/K model will provide the average loss ratio and the average delay for each user in each state of the model. We will assume that the bottleneck link has a buffer capacity of K bits and a link bandwidth of μ bps. The arrival rate at the bottleneck link is given by the aggregate data rates chosen by the users. Note that each state of the model determines an aggregate arrival rate to the link. Let $\lambda(l)$ be the data rate (in bps) associated with strategy l . Thus, the aggregate data rate at state s is given by

Table 1

Codecs and their respective E-model parameters ([6] and [28, Tables I and II])

Codec	Bit rate (kbps)	γ_1	γ_2	γ_3	D_{codec} (ms)
G.723.1.A + VAD-6.3	6.4	15	30	17	47.5
G.729	8.0	10	47.82	18	35
G.723.1.B-5.3	5.4	15	30	17	47.5

$\gamma_s = \sum_l n_l^{(s)} \lambda(l)$. Let $\rho_s = \gamma_s / \mu$. Let $L_s(l)$ and $D_s(l)$ denote the average loss ratio and the average delay of a user that chooses strategy l in state s , respectively. Using the M/M/1/K formulae, we obtain:

$$L_s(l) = \frac{\rho_s^K (1 - \rho_s)}{1 - \rho_s^{K+1}}, \quad Q_s(l) = \rho_s \frac{(1 + K\rho_s^{K+1}) - (K+1)\rho_s^K}{(1 - \rho_s)(1 - \rho_s^{K+1})}, \quad D_s(l) = \frac{Q_s(l)}{\gamma_s(1 - L_s(l))} \quad (2)$$

where $Q_s(l)$ is the average queue size in state s . Note that $L_s(l)$ and $D_s(l)$ do not depend on l due to the PASTA property (so we drop index l). Thus, all users perceive the same loss probability and the same delay in a given state of the model.

The ITU E-model [15] provides a method of estimating the quality of the voice perceived by a user receiving voice data. Although the E-model is very general and has many parameters, which vary according to specific voice codecs and the network performance, default values are usually used for most of the parameters. In our case, we are interested only in the effects of network performance, namely loss ratio and delay, on the voice quality. Voice quality is typically measured by the *mean opinion score* (MOS), a subjective quality score that ranges from 1 (unacceptable) to 5 (excellent). The MOS measure is related to the E-model R -factor through the following equation: $\text{MOS} = 1 + 0.035R + 7 \times 10^{-8}R(R - 60)(100 - R)$. Following the work of Cole and Rosenbluth [4] the R -factor can be defined by the following equation $R = 94.2 - I_e - I_d$ where I_d and I_e cover the impairments caused by delay and losses, respectively.

I_e accounts for impairments caused by both encoding and transmission losses. Some codecs may be more sensitive to losses than others, so I_e may be expressed as a function of e , the overall packet loss probability, as follows [4]: $I_e = \gamma_1 + \gamma_2 \ln(1 + \gamma_3 e)$ where γ_1 , γ_2 and γ_3 depend on the codec. In this paper, we will focus on codecs G.729 and G.723.1 (Table 1 [6]). The overall packet loss probability is the sum of the network loss probability and buffer underflow probability, i.e., $e = E_{\text{network}} + E_{\text{playout buffer}}$. We will assume that $E_{\text{network}} \gg E_{\text{playout buffer}}$ and so, e is equal to L_s in (2).

The impairments caused by the delay can be modeled as follows [4]: $I_d = 0.024d + 0.11(d - 177.3)I(d - 177.3)$ where $I(x) = 0$ if $x < 0$ and 1 otherwise. The end-to-end delay is: $d = D_{\text{network}} + D_{\text{codec}} + D_{\text{playout buffer}}$. We consider $D_{\text{network}} + D_{\text{codec}} \gg D_{\text{playout buffer}}$. D_{codec} is shown in Table 1, and D_{network} is given by D_s in (2).

Finally, the formulae above provide a user-level QoS for each user choosing a given strategy in each state of the model. This fully determines the utility function $U(l, s)$, which yields the MOS score for strategy l in state s .

6. Numerical examples

In this section we numerically evaluate some examples to show the applicability of our model. We are interested in the steady state behavior of the model and its relationship to Nash equilibria. We also illustrate how the user-level QoS (e.g., MOS) varies for different scenarios and highlight some tradeoffs

Table 2

MOS values for a scenario with $N = 16$, two strategies, $\mu = 160$ Kbps and buffer capacity 10K

State	Rate	MOS	Prob.
1: 16 0	102.4	(3.8460 0.0000)	9.6e–146
2: 15 1	104.0	(3.8339 3.9867)	2.1e–131
3: 14 2	105.6	(3.8206 3.9667)	2.1e–117
4: 13 3	107.2	(3.8059 3.9446)	1.2e–103
5: 12 4	108.8	(3.7898 3.9202)	4.6e–90
6: 11 5	110.4	(3.7722 3.8933)	1.1e–76
7: 10 6	112.0	(3.7531 3.8638)	1.9e–63
8: 9 7	113.6	(3.7323 3.8315)	2.1e–50
9: 8 8	115.2	(3.7098 3.7963)	1.5e–37
10: 7 9	116.8	(3.6856 3.7581)	6.6e–25
11: 6 10	118.4	(3.6596 3.7168)	1.4e–12
12: 5 11	120.0	(3.6318 3.6722)	1.0e–00
13: 4 12	121.6	(3.6021 3.6244)	5.6e–13
14: 3 13	123.2	(3.5706 3.5733)	6.0e–26
15: 2 14	124.8	(3.5373 3.5190)	2.4e–39
16: 1 15	126.4	(3.5021 3.4613)	4.3e–53
17: 0 16	128.0	(0.0000 3.4006)	2.7e–67

The bold state indicates the unique Nash equilibrium.

between the model equilibria and the parameters of the bottleneck link. In all the examples that follow, we take ϵ to be small enough. The steady state solution was obtained using the Tangram-II performance analysis tool [2] applying the GTH solution method [12].

We first consider the scenario where users can choose between codecs G.723.1 and G.729, which correspond to bit rates of 6.4 and 8.0Kbps, respectively (Table 1). Table 2 illustrates the MOS value for the scenario with 16 users, $\mu = 160$ Kbps and buffer capacity of 10K. Each state in the system is represented by the number of users choosing each of the available bit rates (6.4 and 8.0Kbps, respectively). In this example, the model has a single Nash equilibrium which is given by state (5, 11), as shown in Table 2. Note however that this Nash equilibrium is not the best state. State (16, 0) is Pareto optimum and yields a higher MOS for all users in that state. But this state cannot be an equilibrium since a self-regarding user would benefit to change its rate to 8.0Kbps and receive a higher MOS (see Table 2). Note that the Nash equilibrium state concentrates the steady state probability mass, while the Pareto optimum state has negligible probability. This example shows what is sometimes referred to as “loss of efficiency” due to miscoordination among the users and the lack of a central authority [17].

We now focus on the impact of the bottleneck link parameters and number of players on the equilibrium of the model. For the examples that follow, the system under consideration had a single Nash equilibrium which received very high probability in steady state. We considered this state the equilibrium point of the system. For the evaluation, we varied individually the buffer capacity, the link capacity and the number of users in the system. The starting scenario consists of 10 users, $\mu = 80$ Kbps and buffer capacity of 120K.

6.1. Impact of buffer capacity

Fig. 3(a) illustrates the behavior of the system equilibrium as a function of the buffer capacity. Note that when the buffer size is small, users choose the lower data rate (i.e., 6.4Kbps). There is no incentive

Fig. 3. System equilibrium as a function of (a) buffer capacity, (b) link capacity and (c) number of users.

to choose the higher data rate as this would cause a larger loss rate. Note that the loss impairment, I_e , has a significant impact contribution in the overall MOS. However, as the buffer capacity increases, our model indicates that users will start choosing the higher data rate, being more aggressive. This occurs until the buffer capacity is augmented to 90K. At this point, we see a subtle increase on the number of users choosing 6.4Kbps. This occurs because at this point the delay impairment, I_d , starts to contribute significantly to the overall MOS. Finally, even if the buffer capacity is further increased, users will not choose to send at the higher data rate. This is due to the well known fact that although a large buffer capacity yields a lower loss rate, it also yields larger delays.

The MOS perceived by the users in this example is illustrated in Fig. 4(a). The figure shows the MOS for users choosing 6.4Kbps, 8.0Kbps and the average MOS across the users. Note that when the buffer capacity varies from 40 to 90K, the average MOS perceived by the users in the equilibrium state decreases continuously. When the buffer capacity equals 90K we have a discontinuity on the MOS. This occurs exactly when the MOS of users adopting 8.0Kbps becomes lower than that received by users adopting 6.4Kbps. At this point a user choosing 8.0Kbps has an incentive to change to 6.4Kbps. Note that this favored not only the user making the change, but also the entire user population.

More insights about this issue may also be gained from Fig. 5(a). In Figs. 5(a) and (b), darker and lighter cells represent lower and higher average MOS, respectively. The y-axis indicates the state of the system, ordered by their aggregate data rate (state 0 representing the case when all users choose 6.4Kbps and state 10 the case when all users choose 8.0Kbps). As we can observe, there is a dark region near the top of the matrix, when users are being more aggressive when choosing their data rates. The region becomes

Fig. 4. User MOS in the equilibrium state as a function of (a) buffer capacity, (b) link capacity and (c) number of users.

Fig. 5. Average user MOS in the equilibrium state as function of (a) buffer capacity and (b) link capacity.

darker as the the buffer capacity increases. This reflects the fact that the larger the buffer capacity, the larger is the delay impairments for users adopting aggressive strategies.

The main insight shown in this example is that when the buffer capacity is small, users choose a lower data rate since choosing the higher data rate would generate a large loss rate, reducing the MOS. When the buffer capacity is large, users do not prefer the higher data rate since this would generate a large delay. In fact, as the buffer capacity increases, the system equilibrium converges to a state where some users choose the lower data rate while others choose the higher data rate.

6.2. Impact of link capacity

Consider now what happens when the bottleneck link capacity varies. Fig. 3(b) shows the equilibrium state as a function of the link capacity. Note that when the link capacity is small, all users choose the lower data rate; when the link capacity is large, all users choose the higher data rate. This transition between all users choosing the lower and the higher data rates is also smooth, as illustrated in Fig. 3(b).

Fig. 4(b) shows the average MOS at the equilibrium state as a function of the link capacity, μ . Note that the average MOS (middle curve) has an increasing trend with the increase in μ . It varies from the lower extreme, when all users choose 6.4Kbps, to the higher extreme, when all users choose 8.0Kbps. On two extremes, the average MOS across the population equals the MOS for population.

Fig. 5(b) shows the average MOS values for different link capacities and different states of the model. Note that as the link capacity increases (from left to right in the x-axis), the average MOS for a given state also increases. For each link capacity, consider the largest state with an average MOS value of 3.5. Note that as the link capacity increases the largest state also increases. In fact, this relationship is linear (i.e., the frontier between the regions with a MOS of 3.5 and 2.5). This linear relationship is related to the linear transition between all users choosing 6.4 and 8.0Kbps in Fig. 3(b).

6.3. Impact of number of users

In this scenario we vary the number of users sharing the bottleneck link — while the other parameters are fixed. Fig. 3(c) shows the system equilibrium as the number of users increases. Note that when there

are few users in the system, the equilibrium is given by a state where all users choose the higher data rate. This reflects the fact that there are plenty of resources available in the system and users behave more aggressively in order to obtain a higher MOS. This behavior continues up to nine users. After this point, the number of users choosing a lower data rate increases with the population size. Note that users reduce their choice of data rate to obtain a larger MOS. As the population size increases, the number of users choosing the higher data rate goes to zero, and the equilibrium is given by the state where all users choose the lower rate. By increasing the population even further, the model degenerates, as the MOS is bounded below by 1. This occurs because the bottleneck link becomes overloaded even when all users choose the lower data rate.

As Fig. 4(c) shows, when the number of users in the system increases, the average MOS perceived by the users decreases. Note that the average MOS is equal to the MOS of users choosing 8.0 and 6.4Kbps when the number of users in the system population is small and large, respectively. We also notice a sharp decrease in the average MOS as the population increases. This represents the fact that the performance of the bottleneck link (e.g., loss rate and delay) quickly degrades as the number of users increases.

6.4. Three strategies case

We now consider a case where users have three data rates to choose from, each corresponding to a different voice codec (Table 1). This example leads to a system that has more than one Nash equilibrium. Table 3 shows the MOS value for a scenario with four users ($N = 4$). The state in this example corresponds to the number of users choosing 6.4, 8 and 5.4Kbps, in this order. Note that the states where all users choose 5.4Kbps (state 1) and where all users choose 6.4Kbps (state 7) are Nash equilibria. However, the steady state probability of state 7 is negligible while the probability mass of the system is concentrated in state 1. This indicates the preference users have over the two Nash equilibria. In fact, the Nash equilibrium characterized by state 1 yields a much higher MOS score for all users when compared to state 7. This

Table 3
MOS values for the scenario with 3 strategies

State	MOS	Prob.
1: 0 0 4	(0.000 0.000 2.724)	1.0e+00
2: 1 0 3	(2.619 0.000 2.421)	5.6e−56
3: 2 0 2	(2.332 0.000 2.105)	1.4e−111
4: 0 1 3	(0.000 1.770 1.917)	2.2e−57
5: 3 0 1	(2.054 0.000 1.796)	2.2e−113
6: 1 1 2	(1.898 1.453 1.623)	8.6e−113
7: 4 0 0	(1.800 0.000 0.000)	5.5e−59
8: 2 1 1	(1.662 1.211 1.371)	3.6e−114
9: 0 2 2	(0.000 1.108 1.246)	2.5e−170
10: 3 1 0	(1.461 1.057 0.000)	2.5e−115
11: 1 2 1	(1.359 1.007 1.090)	2.7e−171
12: 2 2 0	(1.220 1.000 0.000)	4.5e−172
13: 0 3 1	(0.000 1.000 1.000)	1.4e−228
14: 1 3 0	(1.071 1.000 0.000)	9.7e−229
15: 0 4 0	(0.000 1.000 0.000)	3.3e−285

Bold states indicate Nash equilibria.

Fig. 6. Steady state probability of each state as a function of link capacity. In many cases, a single state concentrates the probability mass.

illustrates how the model addresses the problem of equilibria selection (in this example the preferred Nash was selected).

Using this example, we investigate the impact of the link capacity on the steady state distribution of the model. Note that in many cases, more than one Nash equilibrium state is present. Fig. 6 shows a “curtain shape” plot illustrating the steady state probability associated with each state — for each link capacity, we have a different system which has a unique steady state distribution. An impulse of height near 1 at a given state means that the probability mass of the system is concentrated in that state. For many of the link capacities considered, there is a single state (but not the same) that concentrates the steady state probability mass, while other states have negligible probability. This is the case when the link capacity is low (below 29Kbps) and high (above 35Kbps). However, for some link capacities, a few states have non-negligible probability (e.g., at 32Kbps). In this cases, the dynamics are such that the system does not converge to a single Nash equilibrium.

7. Summary

It is a standard procedure in economics literature to apply game-theoretic models to capture human behavior and predict the outcome in situations of conflict [22]. In computer networks literature, this idea has been rapidly growing. However, most of the work within computer networks has focused on investigating *which* decisions are made or *what* decisions should be made. In this work, we bring a new insight into the problem, and propose a dynamic model to investigate *how* decisions are made.

The system considered in this paper assumes that users adopt a *trial and error* mechanism in order to make their choices. This motivates the evolutionary game-theoretic modeling framework [26]. Moreover, there are no guidelines imposed on users when choosing a data rate other than their own perception of quality. We believe this is an interesting approach to modeling the congestion control problem when considering multimedia applications.

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 networks.

Acknowledgements

This research has been supported in part by grants from CNPq (Brazil), FAPERJ and a fellowship from CAPES. D.R. Figueiredo also was supported in part by the NSF under grant awards EIA-0080119, ANI-0070067 and ANI-0085848.

References

- [1] E. Altman, T. Başar, Multiuser rate-based flow control, *IEEE Trans. Commun.* 46 (1998) 940–949.
- [2] R.M.L.R. Carmo, L.R. de Carvalho, E. de Souza e Silva, M.C. Diniz, R.R. Muntz, Performance/availability modeling with the TANGRAM-II modeling environment, *Perform. Eval.* 33 (1998) 45–65.
- [3] K. Chandrayana, S. Kalyanaraman, Uncooperative congestion control, in: *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/PERFORMANCE)*, 2004 258–269.
- [4] R.G. Cole, J.H. Rosenbluth, Voice over IP performance monitoring, *Computer Communication Review* 31 (2) (2001) 9–24.
- [5] E. de Souza e Silva, H.R. Gail, *Transient Solutions for Markov Chains*, Computational Probability, Kluwer, 2000, pp. 44–79.
- [6] L. Ding, R.A. Goubran, Speech quality prediction in VoIP using the extended E-model, in: *Proceedings of the IEEE GLOBECOM*, December, 2003.
- [7] D.R. Figueiredo, E. de Souza e Silva, Efficient mechanisms for recovering voice packets in the Internet, in: *Proceedings of the IEEE/Globecom'99*, 1999, pp. 1830–1837, <http://www.land.ufrj.br>.
- [8] S. Floyd, J. Padhye, J. Widmer, Equation-based congestion control for unicast applications, in: *ACM SIGCOMM*, 2000 43–56.
- [9] R. Garg, A. Kamra, V. Khurana, A game-theoretic approach towards congestion control in communication networks, *ACM SIGCOMM Comput. Commun. Rev.* 32 (3) (2002) 47–61.
- [10] H. Gintis, *Game Theory Evolving*, Princeton University Press, Princeton, NJ, 2000.
- [11] GNU. Speex, 2004, <http://www.speex.org/>.
- [12] W.K. Grassmann, M.I. Taksar, D.P. Heyman, Regenerative analysis and steady state distributions for Markov chains, *Oper. Res.* 33 (1985) 1107–1116.
- [13] A. Greenwald, E. Friedman, S. Shenker, Learning in network contexts: experimental results from simulations, *J. Game Econ. Behav.* 35 (1/2) (2001) 80–123.
- [14] B.R. Haverkort, R. Marie, G. Rubino, K.S. Trivedi, *Performability Modelling: Techniques and Tools*, Wiley, New York, NY, 2001.
- [15] ITU, ITU-T recommendation G.107: the E-model, a computational model for use in transmission planning, Geneva, GE, 2000.
- [16] A. Jensen, Markoff chains as an aid in the study of Markoff processes, *Skandinavisk Aktuarietidskrift* 36 (1953) 87–91.
- [17] R. Johari, J.N. Tsitsiklis, Efficiency loss in a network resource allocation game, *Math. Oper. Res.* 29 (3) (2004) 407–435.
- [18] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993) 29–56.
- [19] A.F. Karr, Weak convergence of a sequence of Markov chains, *Z. Warsch. verw. Geb.* 33 (1975) 41–48.
- [20] P. Key, D. McAuley, Differential QoS and pricing in networks: where flow-control meets game theory, in: *IEE Proceedings Software*, 1999.
- [21] L. Kleinrock, *Queueing Systems, Volume I: Theory*, Wiley/Interscience, New York, NY, 1975.
- [22] A. Mas-Colell, M.D. Whinston, J.R. Green, *Microeconomic Theory*, Oxford University Press, 1995.

- [23] C.D. Meyer, Stochastic Complementation, Uncoupling Markov chains, and the theory of nearly reducible systems, SIAM Rev. 31 (1989) 240–272.
- [24] J. Nagle, On packet switches with infinite storage, IETF Request For Comments (RFC) (1985) 970.
- [25] M.J. Osborne, A. Rubinstein, A Course in Game Theory, The MIT Press, Cambridge, MA, 1994.
- [26] L. Samuelson, Evolutionary Games and Equilibrium Selection, The MIT Press, Cambridge, MA, 1997.
- [27] S. Shenker, Making greed work in networks: a game-theoretic analysis of switch service disciplines, ACM/IEEE Trans. Network 3 (1995) 819–831.
- [28] S. Tao, K. Xu, A. Estepa, T. Fei, L. Gao, R. Guerin, J. Kurose, D. Towsley, Z. Zhang, Improving VoIP quality through path switching, in: Proceedings of IEEE Infocom, 2005.
- [29] H.P. Young, The evolution of conventions, Econometrica 61 (1993) 57–84.

Daniel Sadoc Menasché received the B.S. degree and the M.Sc. degree in Computer Science from the Federal University of Rio de Janeiro, Brazil in 2003 and 2005, respectively. His current interests are in modeling, analysis and performance evaluation of computer systems.

Daniel Ratton Figueiredo received the B.S. degree and the M.Sc. degree in Computer Science from the Federal University of Rio de Janeiro, Brazil in 1996 and 1999, respectively. He recently received the M.Sc. degree and the Ph.D. degree from the Department of Computer Science at the University of Massachusetts, Amherst (in 2005). His research interests include modeling and analysis of computer networks, and in particular, mechanisms used in the Internet.

Edmundo de Souza e Silva received the B.Sc. and M.Sc. degrees in electrical engineering, both from Pontifical Catholic University of Rio de Janeiro (PUC/RJ), and the Ph.D. degree in computer science from the University of California, Los Angeles in 1984.

He has been a Visiting Scientist and Visiting Faculty at the IBM T.J. Watson research Center, a Visiting Scientist at the IBM Tokyo Research Laboratory, a Lecturer with the UCLA Department of Computer Science and Computer Science Department at USC. He has also been a visiting researcher at the Politecnico di Torino, Chinese University of Hong Kong and IRISA/INRIA-Rennes.

He has participated in several technical committees of international conferences. He was a PC co-Chair of the IFIP Third Int. Conf. on Data Commun. Syst. and their Performance (1987), PC co-Chair of IEEE/GLOBECOM'99, PC vice-co-Chair of ITC'2001 and PC co-Chair of ACM/Sigmetrics'2002.

Currently he is a professor at the Federal University of Rio de Janeiro, COPPE and Computer Science Department, and Department chair. He is also vice-chair of the IFIP WG 7.3, and a member of the Board of Directors of ACM/SIGMETRICS. His areas of interest include the modeling and analysis of computer systems and computer communication multimedia networks.