

A game theory framework for ISP streaming traffic management[☆]

V. Bioglio, R. Gaeta, M. Grangetto, M. Sereno^{*}, S. Spoto

Dipartimento di Informatica, Università di Torino, Torino, Italy

ARTICLE INFO

Article history:

Available online 3 August 2011

Keywords:

Peer-to-Peer streaming systems

Game theory

Traffic management

ABSTRACT

The overlay/underlay topology mismatch affects the performance of existing P2P platforms that can generate large volumes of unnecessary inter-ISP network traffic. Although recent works have shown the benefits of network awareness P2P solutions, no studies have focused on the investigation of the ISP behavior and their cooperative/non-cooperative attitudes.

This paper proposes a game theoretic framework to help the design of techniques promoting the ISP cooperation in P2P streaming platforms and decreasing unnecessary inter-domain streaming traffic.

We first analyze some simple scenarios to discuss the existence of Nash equilibria, the Pareto optimality, and a fairness criterion to refine the equilibrium points. Moreover, we apply ideas from Evolutionary Game Theory to design a distributed schemata that the ISPs can use to reach “socially acceptable” equilibrium points in a large ISP population. Furthermore, we develop a discrete event simulation to evaluate the effectiveness of the Evolutionary Game Theory framework.

The study presented in the paper shows that the proposed strategies can effectively stimulate ISP cooperation aiming at the minimization of inter-ISP traffic and help to provide reliable P2P streaming service.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

In P2P applications, the participating nodes (or peers) form an overlay network which is largely agnostic on the underlying IP network. This increases the cost associated with P2P traffic for individual Internet Service Providers (ISPs) that face the problem of managing a vast amount of unnecessary inter-domain traffic [1]. Such traffic levels can cause congestion on valuable inter-ISP links and more generally an overall performance degradation in an ISP network.

Prior studies pointed out that network-aware techniques may drastically reduce the inter-ISP traffic volume [2–4]. Moreover, for P2P streaming, the modeling study presented in [5] showed the positive effects of incorporating network awareness in P2P streaming applications.

This paper focuses attention on the role of the ISPs. In particular, we do not consider them as simple providers of network connectivity but as the main players for supporting P2P streaming platforms. To this aim the paper proposes a game theoretic framework to help the design of techniques promoting the ISP cooperation in P2P streaming platforms. With the help of the proposed models we derive strategies that can effectively stimulate ISP cooperation aiming at the minimization of inter-ISP traffic by supporting a reliable P2P streaming service. In particular, the main contributions of the paper can be summarized as follows.

[☆] This work has been partially supported by the COAST project, a research project supported by the European Commission under its 7th Framework Program (FP7-ICT-248036).

^{*} Corresponding author.

E-mail address: matteo@di.unito.it (M. Sereno).

- We propose a game theory scheme to enable cooperation among ISPs to support P2P streaming applications as well as to decrease unnecessary inter-domain streaming traffic.
- We define a utility function of the ISP and discuss the existing equilibrium points and a fairness criterion to select the equilibrium that satisfies a social fairness criterion.
- We derive an Evolutionary Game Theory approach to design a distributed schemata that the ISPs can use to reach “socially acceptable” equilibrium points in a large ISP population.
- Furthermore, we develop a discrete event simulation to evaluate the effectiveness of the Evolutionary Game Theory framework and complement the theoretical analysis.

The remainder of this paper is organized as follows. Section 2 surveys related work while a system description and modeling assumptions are presented in Section 3. Section 4 introduces the ISP game model and analyzes the Nash equilibrium, the Pareto optimality, and a fairness criterion to refine the equilibrium points. The Evolutionary Game Theory (EGT) model of the ISPs game that allows us to investigate large scale systems is presented in Section 5. In Section 6 we present a simulation based investigation that complements the theoretical studies we present. Finally, Section 7 concludes the paper by discussing future research developments.

2. Related works

Recent studies have shown that network-agnostic P2P file sharing applications can generate a large amount of inter-domain network traffic (see for instance [1]). Studies for P2P streaming applications show similar behavior. In particular, in [6] a measurement study of PPLive, reported that the 70%–90% of the traffic observed by tagged peers deployed in North America was originated by peers located in Asia. Furthermore, this study presents some PPLive measurements collected during a popular live event (the 2006 Chinese New Year celebration) that attracted ~200,000 users over the Internet generating an aggregate bit rate of ~100 gigabits/s! Other results pointing out the effects of topology mismatch in P2P streaming applications have been proposed in [7].

The problems related with the poor usage of the underlying IP network originate several different research proposals. In particular, several efforts aiming at a reduction of the inter-ISP traffic. The paper [8] proposes to cache P2P contents by using ISP-controlled proxy-caches. Theoretically this is an interesting solution because it helps to segregate P2P traffic inside the ISP network, but due to legal risks for unauthorized exchange of copyrighted contents, this solution is not applicable. A different approach in [2] proposes the use of an “oracle” provided by the ISP that helps a peer to locate neighbor peers. A different approach has been presented in [3] where the use of an ISP-controlled tracker has been proposed that provides network related information such as topology, traffic engineering information, and so on. Another interesting technique aiming at reducing the inter-ISP traffic in P2P streaming applications has been presented in [9].

We can say that P2P streaming is traditionally a game among two different kinds of players: the content provider and the users/peers. The existing proposals basically share the idea of helping one of these players, or in some case both of them, by providing information related with the underlying IP network. In this paper we move a bit further by introducing a third type of player, i.e., the ISPs. In our study the ISPs become active players of the P2P streaming by shifting their role from passive network providers to active game players that are able to modulate the network resources they require/offer to the P2P streaming platform. In particular, we propose a game theory framework to represent the ISP behavior and then, by using this framework, we investigate the characteristics of the interactions among the ISPs.

The works [10,11] analyze the need to localize P2P file sharing traffic within the ISP boundaries and present some of the architectural ideas exploited in the present study as well. In particular, [10] proposes an overlay structure composed by mTrackers (similar in the spirit to the one described in Section 6) aiming at minimizing transit ISP costs. The paper [11] describes an ISP-friendly variant of BitTorrent to reduce the cross-ISP traffic.

Recently, in the literature, the game theory has gained an important role in modeling strategies used by the peers of a P2P system and several papers on this topic have been published. In [12–14] game theory models have been used for the investigation of peer interactions and to analyze incentive mechanisms for P2P file sharing systems. The work presented in [15] studies the interaction among peers in P2P streaming applications, illustrates the existence of Nash equilibrium points, and introduces refinement criteria to help the selection among these equilibria.

Our proposal and [16] share the same idea to use EGT for the analysis of the ISPs behavior and the implementation of an iterative optimization procedure. In [16] the peers organize into groups, e.g., ISPs, and each peer can act as an agent downloading the content to be shared within the group, leading to the definition of a game with a discrete set of strategies. The main problem of this proposal concerns the stability of the technique (i.e., any perturbation, such as peer churning, might lead the to a new solution). Moreover, the EGT based optimization algorithm in [16] requires the estimation of a payoff function that may produce a significant delay in the optimization. In our proposal the same issue is tackled resorting to a 2-tier architectural design similar, in the spirit, to other 2-tier P2P architectures used for different purposes, such as [17,3].

Another paper that is related with our proposal is [18]. We share with this paper several model assumptions. In particular, we use the same fluid assumption whose consequences are that the peers are able to share their upload bandwidth among all the participant peers. That is, the stream is seen as a fluid that can be split in a continuous manner among all the requesting peers.

Table 1
Notation.

N	Set of peers in the system (with $ N = n$)
r	Video streaming rate (in cps)
c_s	Streaming server upload capacity (in cps)
c_j	j -th peer upload capacity (in cps)
M_i	i -th partition of peers, with $i = 1, \dots, k$, and $ M_i = m_i$
Θ_i	rate at which data are requested by partition i (to other partitions)
Γ_i	rate at which data are received by partition i (from other partitions)
Φ_i	bandwidth resources offered by partition i (to other partitions)
Ψ_i	bandwidth resources provided by partition i (to other partitions)

3. System description and model assumptions

We consider a P2P streaming system composed of a set N of peers ($|N| = n$) that cooperate to stream a video whose rate is r . Each peer j is characterized by its upload capacity c_j that defines its maximum outgoing bit rate; the video originates from a server peer whose upload capacity is denoted as c_s . Upload capacities are expressed in chunks per second (cps).

Each peer j requests data at rate θ_j while it receives data at rate γ_j . We assume an ideal P2P streaming system that does not generate redundant information, e.g., duplicate chunks; this means that the request rate is $\theta_j = r$ and γ_j represents the rate at which peer j receives useful data. In this ideal setting it might be that $\gamma_j \leq \theta_j$ hence the rate at which peer j receives data may be smaller than r resulting in a non-continuous video playback.

At the same time peer j acts as a data server: it offers other peers resources at a rate ϕ_j although resources are actually provided at rate ψ_j . Of course, $\phi_j \leq c_j$, and $\phi_j \geq \psi_j$ i.e., the bandwidth resources offered by a peer are less than or equal to its upload bandwidth and it may happen that not all the offered bandwidth resources are actually used because there are no peers requesting them.

We consider a partition M_1, M_2, \dots, M_k of the set of peers N , i.e., $\forall i, j, 1 \leq i, j \leq k, M_i \subseteq N, M_i \cap M_j = \emptyset, \bigcup_{i=1}^k M_i = N$, and $|M_i| = m_i$. An element M_i of the partition may represent an ISP or an Autonomous Systems and the connection point between M_i and the remaining part of the set of peers can be the link (or the links) connecting the M_i networks to the rest of the Internet. In the following we use the term ISP as a synonym for an element M_i . To simplify the derivations, we assume that the streaming server is the unique (special) peer of ISP M_0 .

Data request rate can be defined at the ISP level. In particular, peer j in M_i requests data from other peers in M_i at rate θ'_j and from peers outside M_i at rate θ''_j . Of course, $\theta_j = \theta'_j + \theta''_j$. The data request rate of ISP M_i is defined as $\Theta_i = \sum_{j \in M_i} \theta''_j$. Data received, offered, and provided rates at the ISP level (denoted as Γ_i, Φ_i , and Ψ_i , respectively) are similarly defined.

To ease the task of the reader Table 1 summarizes the notation used to describe the system under study.

3.1. From peers to ISPs

According to the definition of request rate, Θ_i is the minimum value of the aggregate rate that partition M_i must receive from the outside such that all the peers of M_i receive a streaming rate supporting a continuous video playback. On the other hand, if the rate at which data are actually received, i.e., Γ_i is smaller than Θ_i , there are peers of the ISP that cannot correctly play the stream. The widespread diffusion of P2P streaming applications yields significant implications for ISPs: on the one hand the increasing amount of P2P traffic results in additional costs for the ISP, that is requested to provide a satisfactory level of service to its subscribers (not only the P2P streaming users). On the other hand, these applications represent an important driving force for the adoption of broadband access, which is a significant source of revenue for the ISPs. A successful strategy to manage the traffic generated by P2P streaming applications must address both the ISP perspective of costs and the subscriber perspective of quality of service. From the ISP side the rate at which data are received, i.e., Γ_i , represents a gauge of the offered quality of service. That is, if $\Gamma_i \geq \Theta_i$, the P2P ISP subscribers may correctly receive the streaming. Obviously when $\Gamma_i < \Theta_i$ there are peers belonging to the ISP that cannot correctly play the stream. Hence the ISP aiming at fulfilling its requests tries to get an appropriate Γ_i . Note that the reduction of the ISP requests, i.e., Θ_i can be pursued by using techniques enforcing locality. We define locality awareness as the peer ability to exploit local resources (e.g., chunks) over remote ones whenever possible (see for instance [9] or [3]). Hence we have that an ISP defines its policies for selecting peer neighborhood and this has an impact on the ISP requests, i.e., on Θ_i .

On the other hand the costs due to content uploads can be drastically reduced by setting $\Phi_i = 0$. This can be obtained by dropping all the communications from peers that belong to the ISP towards peers that are outside (e.g., bandwidth throttling). It is clear that this extreme ISP free-riding behavior cannot be socially acceptable from the point of view of the P2P application.

If we assume that the resources offered by an ISP (and by the streaming server) can be used by any other ISP we have that a universal streaming¹ service is possible iff

¹ The system provides universal streaming when all participating peers receive the video at rate r [18]. The concept of universal streaming can be considered a synonym that the P2P streaming platform is providing a good quality service, i.e., the peers of the P2P streaming system are receiving with a rate at least equal to r .

$$c_s + \sum_{i=1}^k \Phi_i \geq \sum_{i=1}^k \Theta_i. \tag{1}$$

In other words, the previous inequality simply states that the continuous playout of the stream is possible if the sum of offered resources (by the ISPs and by the streaming server) is greater than (or equal to) the sum of requested resources.

Furthermore we can say that the previous inequality is the ISP-equivalent of the result presented in [18] (Theorem 1) for the peers of a P2P streaming system, and the underlying assumptions of Inequality (1) are similar to the ones used in [18]. In particular, we are assuming that there is a complete connectivity among the ISPs, and that the resource can be split among the ISPs in a “continuous manner” (in [18] this is called “fluid assumption”). In Section 6.1 we discuss the impact of these simplifying assumptions on realistic P2P streaming systems.

Please note that a peer $j \in M_i$ uses several different criteria to select its neighborhood choosing among peers of M_i and of other ISPs. These impact on the splitting of θ_j between θ'_j and θ''_j . Our goal is the development of a modeling framework for the study of the interactions among ISPs that fits P2P applications with different characteristics (e.g., trees, multi-tree, mesh, pull, push, and so on); therefore we do not want to rely on a specific criterion to select peer neighbors (in the same or other ISPs). We notice that in our framework, the selection criteria change Θ_i , i.e., they impact on the rate at which a ISP requests data to the other ISPs. In other words, we do not account for the selection criteria but we only consider their effects on the ISP requests. In Section 6.1 we show that our framework does not depend on the particular section method but only relies on the ISP requests, i.e., Θ_i .

4. Game theory formulation

In Section 3 we defined the request rate for ISPs (Θ_i , $i = 1, \dots, k$) that depends on several factors, e.g., algorithms for overlay network management, number of peers in ISPs, video bit rate, server upload capacity. Rational ISPs aim at reducing the bandwidth costs by possibly limiting the offered bandwidth Φ_i (e.g., $\Phi_i = 0$). Clearly this free-rider behavior of ISPs would cause the P2P streaming system to work inefficiently.

The main goal of the game theory framework we propose in this section is to avoid this ISP selfish behavior. Here we introduce a game theory model where the players are ISPs. Finally, in Section 5 we propose an analysis of the ISP game using the Evolutionary Game Theory, that allows us to devise a practical algorithm to optimize the interactions among ISPs.

The model is based on the following general utility function

$$u_i(\Phi_1, \dots, \Phi_i, \dots, \Phi_k) = -\text{Cost}(\Phi_i) + \text{Gain} \left(c_s + \sum_{i=1}^k \Phi_i \geq \sum_{i=1}^k \Theta_i \right). \tag{2}$$

where offering resources represents a cost for ISPs while the universal streaming of the P2P system represents a gain. We first provide a mathematical description of this framework by starting with a simple scenario with two ISPs and we later generalize to the case with an arbitrary number of interacting ISPs.

The 2 ISPs game. Inequality (1) states that the system is providing universal streaming (i.e., the P2P streaming service is working properly) if

$$\Phi_1 + \Phi_2 + c_s \geq \Theta_1 + \Theta_2.$$

We can use this condition to define the payoff function of the two ISPs game. Player i ($i = 1, 2$) may receive chunks from both the streaming server and from the other ISP. We define the cost of ISP i equal to Φ_i while if Inequality (1) holds we assume that ISP i receives a gain of g cps.² If (Φ_1, Φ_2) is the action profile, we can refine Eq. (2) and define the payoff as

$$\begin{aligned} u_1(\Phi_1, \Phi_2) &= -\Phi_1 + g \cdot \mathbb{1}(c_s + \Phi_1 + \Phi_2 \geq \Theta_1 + \Theta_2) \\ u_2(\Phi_1, \Phi_2) &= -\Phi_2 + g \cdot \mathbb{1}(c_s + \Phi_1 + \Phi_2 \geq \Theta_1 + \Theta_2), \end{aligned} \tag{3}$$

where

$$\mathbb{1}(\text{condition}) = \begin{cases} 1 & \text{condition holds} \\ 0 & \text{otherwise.} \end{cases}$$

The above payoff function consists of two terms: the first term denotes the cost (expressed in terms of bandwidth) faced by ISP i , while the second term denotes the gain of ISP i .

It is reasonable to assume that $\Phi_1, \Phi_2 < g$, since ISPs will only cooperate with each other if cooperation can benefit both ISPs and gives them positive payoffs. Let $\mathbf{u}(\Phi_1, \Phi_2) = (u_1(\Phi_1, \Phi_2), u_2(\Phi_1, \Phi_2))$ be the payoff profile. Fig. 1 (left plot) shows the set of positive payoff values for Φ_1 and Φ_2 that allow the universal streaming (gray region) if the upload capacity of all peers in ISP i is equal to c_i (in this case, $\Phi_1 \leq c_1 m_1$ and $\Phi_2 \leq c_2 m_2$).

² We use the term g to quantify ISP satisfaction/dissatisfaction.




Fig. 1. Set of payoff values (left plot), and of Φ_1 and Φ_2 (right plot) that allow universal streaming for the two ISPs system.

We can see that all the values of Φ_1 , Φ_2 that are on the segment between $(0, \Theta_1 + \Theta_2 - c_s)$ and $(\Theta_1 + \Theta_2 - c_s, 0)$ in Fig. 1 (right plot) are Nash equilibrium points (NE). In particular, if we assume that (x, y) is a point on this line segment and $\mathbf{u}(x, y)$ is the corresponding payoff profile we have that for any $h > 0$

$$\begin{aligned}
 u_1(x, y) &= -x + g \\
 u_1(x + h, y) &= -x - h + g \quad \text{and} \quad u_1(x, y) > u_1(x + h, y) \\
 u_1(x - h, y) &= -x + h - g \quad \text{and} \quad u_1(x, y) > u_1(x - h, y).
 \end{aligned}$$

Since the same type of derivation can be done for $u_2(x, y)$, then we can conclude that any value (x, y) on the line segment between $(0, \Theta_1 + \Theta_2 - c_s)$ and $(\Theta_1 + \Theta_2 - c_s, 0)$ is a NE point. On the other hand, we must point out that not all the payoff profiles on this line segment are “socially” acceptable. Note that there are points such as $\Phi_1 = 0$ and $\Phi_2 = \Theta_1 + \Theta_2 - c_s$ (and vice versa). In other words the NE condition does not guarantee a fair contribution of the two ISPs. As a first effort we try to refine these equilibrium points by using the Pareto optimality [19,20]. In particular, we know that a payoff profile $\mathbf{u}(\Phi_1, \Phi_2)$ is Pareto optimal if there is no other payoff profile $\mathbf{u}(\Phi'_1, \Phi'_2)$ such that $u_1(\Phi'_1, \Phi'_2) \geq u_1(\Phi_1, \Phi_2)$ and $u_2(\Phi'_1, \Phi'_2) \geq u_2(\Phi_1, \Phi_2)$. Pareto optimality means that no one can increase his/her payoff without degrading other’s. It is clear from Fig. 1 (right plot) that all the values on the line segment between $(0, \Theta_1 + \Theta_2 - c_s)$ and $(\Theta_1 + \Theta_2 - c_s, 0)$ are Pareto optimal. Hence, not even the Pareto optimality allows us to discriminate among the values on the segment $(0, \Theta_1 + \Theta_2 - c_s)$ and $(\Theta_1 + \Theta_2 - c_s, 0)$. What we need in this case is a criterion that permits us to discriminate among the values on this segment. In particular, we are interested in a fairness criterion that allocates the Φ_i -s proportionally to the amount of required resources. A similar fairness criterion has been proposed in [21] where it has been called *effort fairness*. A version of this criterion for our game provides us the following constraint

$$\frac{\Phi_1}{\Theta_1} = \frac{\Phi_2}{\Theta_2}.$$

With some simple algebraic manipulation we can derive that the values of Φ_1 and Φ_2 satisfying the previous condition are

$$\Phi_1^* = \frac{\Theta_1(\Theta_1 + \Theta_2 - c_s)}{\Theta_1 + \Theta_2} \quad \text{and} \quad \Phi_2^* = \frac{\Theta_2(\Theta_1 + \Theta_2 - c_s)}{\Theta_1 + \Theta_2}.$$

The k ISPs game. Let $\Theta_1, \dots, \Theta_k$ be the incoming aggregate rates of the k ISPs game. In this case the Nash equilibriums can be found on the hyperplane satisfying the following relation:

$$\sum_{i=1}^k \Phi_i = \sum_{i=1}^k \Theta_i - c_s.$$

Furthermore, the effort fair points $(\Phi_1^*, \dots, \Phi_k^*)$ are the solution of the following system of equations

$$\begin{cases}
 \sum_{i=1}^k \Phi_i^* = \sum_{i=1}^k \Theta_i - c_s \\
 \frac{\Phi_i^*}{\Theta_i} = h \quad \text{for } 1 \leq i \leq k,
 \end{cases}$$

where h is a constant. From the second equation we can derive that

$$\Phi_j^* = \Theta_j \frac{\Phi_i^*}{\Theta_i} \quad \forall 1 \leq i, j \leq k.$$

By substituting this in the first equation we obtain that

$$\Phi_i^* + \sum_{j=1, j \neq i}^k \Theta_j \frac{\Phi_i^*}{\Theta_i} = \sum_{j=1}^k \Theta_j - c_s.$$

From this equation with simple algebraic manipulation we can derive that

$$\Phi_i^* = \Theta_i \left(1 - \frac{c_s}{\sum_{j=1, j \neq i}^k \Theta_j} \right) \quad i = 1, \dots, k. \tag{4}$$

5. Evolutionary analysis of streaming games

In this section we present a procedure that ISPs can use to compute the effort-fair values Φ_i^* . Note that these values can be computed by using Eq. (4) but this requires the knowledge of all the values Θ_j for $j = 1, \dots, k$, with $i \neq j$, i.e., besides its bandwidth request Θ_i each ISP has to know the Θ_j values of all the other ISPs. We assume an ISP does not disclose (to the other ISPs) information related with the services it offers to the subscribers (i.e., Θ_i).

Therefore, to avoid dissemination of its private information we use a two tier structure of trackers composed of a *super-tracker* and a collection of *ISP-trackers*, one for each ISP. This kind of architecture is inspired by the tracker architectures proposed in [9,22].

In our proposal the ISP- i 's tracker sends to the super-tracker the values Θ_i and the current value for Φ_i . Then the super-tracker is able to compute the indicator function $\mathbb{1}(\cdot)$ and it can send the output to the ISPs. In this manner, ISPs do not disseminate their Θ_i -s and Φ_i -s but they only send these values to a (trusted) super-tracker. Obviously the implementation of this architecture must ensure that the ISP-trackers and the super-tracker mutually trust each other. Furthermore, this architecture could also be enhanced with measurement tools that help to avoid cheating behaviors. The two tier tracker architecture can be viewed as a sort of "oracle" that ISPs can use to know if system is providing universal streaming or not.

By exploiting the two tier tracker infrastructure and by using ideas from Evolutionary Game Theory (EGT) [23] we devise a distributed strategy to allow ISPs to compute the effort fairness point.

In EGT each player has local information of the game. In particular, in case of the ISP game the local information are bandwidth request Θ_i , the current value for Φ_i , and a super-tracker to know if the system is providing universal streaming or not. The player learns which moves are better by trying them and observing how its utility function changes. The evolution of the system of players is ruled by equations called *replicator dynamic* [23]. The replicator dynamic considers the evolution of each players behavior and shows the temporal evolution of players moves. In this way, if the game is repeated the replicator could reach an equilibrium, that is the equilibrium of the game.

The common EGT approach is to consider games with a finite number of strategies. On the other hand the game we propose has infinite strategies because Φ_i can take values over a continuous interval. For this reason we cannot use the classical equation of replicator dynamic, but we create a new dynamic with the same kind of properties.

In the following we illustrate the method for the two ISPs game. In this case we look for a replicator dynamic that has an equilibrium in (Φ_1^*, Φ_2^*) . The equation of ISP- i should depend only on Φ_i and Θ_i and from the knowledge of being in the universal streaming region or not. In particular, if the ISP is in the universal streaming region it should decrease its Φ_i to increase its utility function. If we denote by Φ_i^s the value of the Φ_i at the s -th iteration step and Φ_i^0 the initial values of the Φ_i (for $i = 1, \dots, k$), we use a replicator dynamic described by the following equations

$$\Phi_i^{s+1} = \begin{cases} \Phi_i^s(1 - \beta) & \text{if universal streaming} \\ \Phi_i^s + \beta\Theta_i & \text{otherwise} \end{cases} \tag{5}$$

where $0 < \beta < 1$ is a parameter that tunes the speed of convergence. We can assume that the ISP- i decreases Φ_i proportionally to the same value computed in the previous iteration. If the ISP has a high Φ_i it is probably far from the bound of the region, so it could decrease Φ_i a lot. Instead, if the ISP has a low Φ_i it is probably close to the bound, so it should pay attention on decreasing Φ_i . On the other hand, if the ISP is not in the universal streaming region it would try to increase its Φ_i faster to reach that region. In this case, it is reasonable that the ISP increases Φ_i depending on its request, i.e., Θ_i . We can see the evolution of the dynamic system with the help of Fig. 2.

If we start the iterative method from a point (Φ_1^0, Φ_2^0) that is in the universal streaming region, then the system will move out the region in a radial way, i.e., following the line to the point $(0, 0)$. On the other hand, if we start from a point that is out of the universal streaming region then we will move along the vector $(\beta\Theta_1, \beta\Theta_2)$. This means that in this case we always move in the same direction, i.e., all the movements are parallel to the line $t^* : \Phi_1\Theta_2 = \Phi_2\Theta_1$. In this manner there is no need need to know the bandwidth request of the other ISPs to choose its move (i.e., the Θ_j -s (for $j \neq i$)).




Fig. 2. (a) Vector field of the dynamic system (the symbol \blacklozenge denotes the point (Φ_1^*, Φ_2^*)); (b) Evolution of the replicator dynamic.

We must point out that with a similar iteration strategy the process will oscillate around the effort-fair points and the size of this fluctuation interval depends on β . That is, a small value of β yields smaller fluctuation interval but with lower speed of convergence, and a large value of β yields higher speed of convergence but larger fluctuation interval. To overcome this problem we can adopt a variable β , i.e., a value of β that is function of s and of the stability of the convergence process. We empirically, derive a suitable updating strategy for the values of β .

In the following we prove that the iterative method converges to the effort-fair points.

Proposition 1. Each ISP converges to its own Φ_i^* .

Proof. For the sake of readability we only present the proof for the case $k = 2$ (the proof for $k > 2$ follows the same ideas but with the use of a cumbersome notation).

First of all, we show that the line $t^* : \Phi_1\Theta_2 = \Phi_2\Theta_1$ is a fixed line of the replicator dynamic, i.e., if the starting point of the iteration, i.e., $p^{(0)} = (\Phi_1^0, \Phi_2^0)$ is on the line t^* then for any other iteration step s we have that $p^{(s)} = (\Phi_1^s, \Phi_2^s) \in t^*$. If $p^{(0)} \in t^*$ then

$$p^{(0)} = \left(\Phi_1^0, \frac{\Theta_2}{\Theta_1} \Phi_1^0 \right).$$

Now there are two cases:

- $p^{(0)}$ is in the universal streaming region and then

$$p^{(1)} = p^{(0)} - \beta p^{(0)} = \left(\Phi_1^0(1 - \beta), \frac{\Theta_2}{\Theta_1} \Phi_1^0(1 - \beta) \right),$$

so $p^{(1)} \in t^*$.

- $p^{(0)}$ is not in the universal streaming region. In this case we have that

$$p^{(1)} = p^{(0)} + \beta(\Theta_1, \Theta_2) = \left(\Phi_1^0 + \beta\Theta_1, \frac{\Theta_2}{\Theta_1}(\Phi_1^0 + \beta\Theta_1) \right),$$

and again $p^{(1)} \in t^*$.

We can conclude that t^* is a fixed line of the replicator dynamic.

We prove now that the line t^* is an attractor of the dynamic system. Let $p = (x, y)$ be a generic point (see Fig. 2) in the universal streaming region. In this case the system will move in a radial direction along the line t until it falls out of the universal streaming region, e.g., in the point $p' = (x', y')$. This point is not in the universal streaming region, so the system will move along a new line, parallel to t^* , until it returns into the universal streaming region. Let $p'' = (x'', y'')$ be the new point reached by the system. The system will move along a new radial line t' , that is closer to t^* with respect to the previous line t . Going on we have that the radial line tends to t^* , so we can conclude that t^* is an attractor of the system.

Finally, we have that the system oscillates around the effort-fair point (Φ_1^*, Φ_2^*) . Indeed, each Φ_i^s will oscillate into the interval $[\Phi_i^* - \beta\Phi_i^*, \Phi_i^* + \beta\Phi_i^*]$. This means that using an opportune function for decreasing β we obtain for each ISP- i the convergence to its effort-fair point Φ_i^* . \square




Fig. 3. Numerical experiment with $k = 1000$, $r = 50$ cps, $c_s = 50 \cdot r$, $m_i \in \text{Unif}[100, 1000]$, and $\theta_i \in \text{Unif}[r, m_i \cdot r]$.

Numerical experiments. The convergence of the EGT scheme in case of large systems is investigated with a set of numerical experiments with the following parameters: we consider $r = 50$ cps to be the streaming rate, $c_s = 50 \cdot r$ the server upload capacity, and $k = 1000$ the number of ISPs. The number of peers in the i -th ISP, is derived in the following manner: $m_i = \text{Unif}(100, 1000)$, and the value for θ_i is derived as follows $\theta_i = \text{Unif}(r, m_i \cdot r)$, where $\text{Unif}(a, b)$ is a random number uniformly selected in the interval $[a, b]$. From the large set of experiments we performed we derived similar results in terms of speed of convergence. Fig. 3 summarizes the results of one of these experiments. In particular, the figure shows the evolution of three values of Φ_i^* derived by applying our EGT based approach. For each value of the Φ_i^* -s the figure also shows the corresponding theoretical values computed by using Eq. (4).

The experiments show that in all the cases the convergence is achieved with a few number of iterations (e.g., between 30 and 40 iterations).

6. An ISP streaming traffic management algorithm

In Section 4 we introduced a simple game theoretic framework but in all the models we presented the equilibrium points can only be achieved via common knowledge. In other words, each ISP needs to know the streaming server upload capacity and the θ_i -s of all the other ISPs. Obviously this hypothesis is not realistic in practical systems. Nevertheless, by using the ideas presented in Section 5 we can define a distributed algorithm that implements the EGT version of the game without the constraints that the ISP must be aware of the θ_i -s of all the other ISPs.

The proposed distributed algorithm could be implemented by means of a two tier structure of trackers (similar in the spirit to the tracker architecture proposed in [9,22,17]). At the top level there is the *super-tracker* and at the lower level there is a collection of *ISP-trackers* (one for each ISP). Each ISP-tracker must be able to measure the current values of (i) the number of active peers in the ISP, (ii) the number of peers that receive the video at rate r , and (iii) the values for θ_i , and Φ_i . Furthermore, the ISP-tracker is able to modify the bandwidth offered to the other ISPs (e.g., by using bandwidth throttling).

Each ISP-tracker periodically (e.g., each δ sec) sends the collected measures to the super-tracker that can use them to verify whether the system is in a universal streaming region or not. The super-tracker replies to the ISP-trackers with a *universal-streaming* or *no-universal-streaming* message. In turn, the ISP-trackers are required to update the values of Φ_i -s according to Eq. (5).

This two-tier architecture represents a practical solution for using global knowledge at the ISP level. Obviously, the ISPs shall trust the common super-tracker, that is actually acquiring ISP-trackers private information but it is not disclosing such data to the ISPs. In fact the super-tracker only discloses the result of the universal streaming indicator function. In this paper we neglect a number of open issues that are left for future research and practical implementation. Among such extensions we mention the need for incentive mechanisms to enforce cooperation among the ISPs, the implementation control techniques to guarantee ISP honesty, and the confidentiality of the communications between the super-tracker and ISP-trackers. Finally it is important to notice that the two-tier architecture, composed of the super-tracker and the collection of ISP-trackers, only aims at ISP traffic management and does not replace or interfere with the architecture of the adopted P2P application. In other words, if the P2P streaming architecture uses a tracker similar to the BitTorrent protocol to assist the streaming distribution (e.g., to provide peer addresses) this tracker does not interfere with the ISP traffic management.

6.1. Experimental results

To investigate the characteristics of the EGT scheme we develop a discrete event simulator that implements in an abstract manner most of the features of a pull-based P2P streaming application that interacts with our EGT algorithm.

Simulator description. The simulator models a P2P streaming system where the peers are grouped into k different ISPs. The simulator is based on the implementation of three agents, namely the peer, the ISP-tracker and the super-tracker. The

interactions among the peers, and among the ISPs are governed by two different dynamics: the peer dynamic and the proposed EGT algorithm.

The simulator assumes a chunk based P2P system and hence all the download and upload capacities are discretized in terms of chunk size. Each peer belongs to a certain ISP i and it is characterized by its upload bandwidth. Without loss of generality, we assume that all the peers in the i -th ISP share the same maximum upload capacity c_i (in chunk). The goal of the peer is to connect to other peers in the swarm in order to download enough chunks per seconds to play the video. In the simulator this is represented by r download connection slots that the peers try to fill with resources offered by other peers. In conclusion, the peers can be represented as nodes in a graph that models the swarm, where each peer has c_i outgoing stubs and r incoming stubs.

The peer dynamic is represented by means of k independent $M/M/\infty$ (or more generally $M/G/\infty$) stochastic processes, where λ_i and μ_i represent the arrival and the departure rates of peers for the i -th ISP. When a new peer joins the P2P swarm it selects (with the help of a P2P tracker) the candidate neighboring peers according to a given selection strategy and attempts to establish at least r connections. If all the required incoming connections cannot be created because of lack of resources offered in the system the peer keeps searching for more neighbors with a granularity of T seconds, that in our simulation has been fixed at 10 s.

An important issue of our simulator concerns the neighborhood selection criteria. We implemented two very different neighboring management policies so as to show that the proposed EGT solution does not depend on a particular choice. According to the first criterion each peer selects a neighbor a peer in its ISP with a fixed probability p , (from other ISPs with probability $1 - p$). The second approach is based on a random choice over all the connected peers; this latter criterion represents the most diffused implementation using a central P2P tracker that logs all the peers in the swarm and randomly selects a set of neighbors upon request by a given peer. The latter criterion yields a variable probability p_i of neighbor selection in the i -th ISP equal to $p_i = \frac{m_i}{\sum_{i=1}^k m_i}$.

The other agents in the system are the k ISP-trackers and the super-tracker that implement the proposed EGT algorithm. The i -th ISP-tracker has two roles: the first is to evaluate the values of Θ_i and Φ_i , by counting all the bandwidth resources requested/offered by the peers in the ISP, and to communicate such figures to the super-tracker every δ seconds. The second task is to react to the super-tracker messages by controlling the offered resources Φ_i according to the policy in Eq. (5). When necessary the ISP-trackers reduce the offered resources by simply pruning some of the inter-ISP connections. In turn, the peers will try to establish new links according to the P2P application policy that is unaware of the on-going EGT algorithm. The role of the super-tracker is to collect the values of Θ_i and Φ_i so as to verify if the system is in the universal streaming regime or not, using Eq. (1). Finally, the super-tracker communicates the result of the evaluation to all the ISP-trackers. The δ parameter, which determines the frequency of the EGT adaptation round, is not very critical. A small value for δ yields more control over the peer dynamics. On the other hand, a frequent action of the EGT would rise frequent bandwidth throttling operations that have a negative impact on the continuity of service offered to the end users.

The simulator allows us to test the performance of the proposed EGT and its interactions with the maintenance policies of the P2P application. The behavior of the EGT is monitored by logging the values of Θ_i , Γ_i , Φ_i and Ψ_i . Moreover, we introduce a performance index of the quality of service offered by the i -th ISP QoS_i , which is defined as the ratio between the number of peers that are enjoying the video and the number of active peers in the same i -th ISP.

Experiments. The discrete event simulator has allowed us to test the proposed EGT algorithm and analyze its performance as a function of the most important system parameters, namely the peer's upload capacity, the neighbor selection policy and the ISPs average population size. In all the following simulations the EGT is played with a frequency $\delta = 30$ s. The same simulator has been used to obtain the reference performance of a P2P application that does not enforce any control on the inter-ISP traffic by disabling the EGT algorithm.

Table 2 summarizes the parameters of the investigated scenario. We simulate a streaming server equipped with an upload capacity that is 50 times the video rate r . As already stated we assume that the peer dynamics are represented by k independent $M/M/\infty$ processes. In the following we present results for $k = 3$ ISPs with different peers' arrival and departure rates so as to test the proposed optimization in the presence of ISPs with different average sizes. Moreover, we consider ISPs that are heterogeneous also in terms of the upload capacity c_i . In all the simulated cases the average sojourn time of a peer in the system is kept constant and equal to $\bar{w} = 15$ min. The total number of peers in the P2P swarm is on average $n = 1000$. As shown in Table 2, ISP-1 has the largest upload capacity ($5 \times r$) but covers only the 10% of the peers, ISP-2 represents a small group of peers with limited upload whereas ISP-3 contains 80% of the P2P overlay but offers the smallest upload.

The first goal of the simulative study is to prove the feasibility and the effectiveness of the proposed EGT-based optimization. To this end we analyze the dynamic of the requested, offered and provided bandwidth resources at ISP level (Θ_i , Φ_i and Ψ_i , normalized with respect to r). In the following, we show average values computed over 32 simulation trials so as to achieve good accuracy that has been estimated using 95% confidence interval. All the numerical results reported in the following exhibit a margin of error below 1%. The transient period before the $M/M/\infty$ process driving the peer dynamic achieve the steady state regime has been removed from all averages.

Table 3 shows the average value of m_i , Θ_i , Φ_i and Ψ_i when the P2P application uses a fixed probability of connecting inside the same ISP $p = 1/3$. All the bandwidth resource values are expressed in chunks and normalized with respect to the video rate r . We can observe that without EGT all the peers offer to the other peers all their spare bandwidth resources, as

Table 2
System parameters for the first set of experiments.

k	3		
n	≈ 1000		
r	35 (cps)		
c_s	1750 (cps) ($50 \times r$)		
\bar{w}	15 min		
c_i	175 ($5 \times r$)	42 ($1.2 \times r$)	35 ($1.0 \times r$)
m_i	$\approx 0.1 \times n$	$\approx 0.1 \times n$	$\approx 0.8 \times n$

Table 3
Simulation experiment with fixed probability neighbor selection $p = 1/3$.

	Fixed probability neighbor selection					
	Without EGT			With EGT		
	ISP ₁	ISP ₂	ISP ₃	ISP ₁	ISP ₂	ISP ₃
m_i	100.63	98.13	794.48	106.69	96.81	799.02
Θ_i	50.31	49.06	397.24	53.34	48.40	399.51
Φ_i	503.14	117.75	794.48	48.30	43.59	361.95
Ψ_i	255.10	55.81	71.80	48.47	43.68	85.66
QoS _i	0.9831	0.9844	0.9825	0.9830	0.9822	0.9823

Table 4
Simulation experiment with variable probability neighbor selection $p_i = m_i / \sum_{j=1}^3 m_j$.

	Variable probability neighbor selection					
	Without EGT			With EGT		
	ISP ₁	ISP ₂	ISP ₃	ISP ₁	ISP ₂	ISP ₃
m_i	104.03	102.60	805.96	98.18	100.07	805.36
Θ_i	93.27	92.12	164.97	88.49	90.00	159.51
Φ_i	520.14	123.12	1128.35	75.99	76.20	136.99
Ψ_i	81.86	61.97	127.29	70.74	62.67	124.07
QoS _i	0.9826	0.9828	0.9823	0.9827	0.9827	0.9824

evident from the values of Φ_i ; such resources are exploited in a random way by the other parties using the fixed probability criterion for neighbor selection. Recalling that in Table 3 the reported values are averages computed when the P2P system reaches its equilibrium, one notes that the P2P application provides an *unfair* resource usage. In fact ISP-1, that has the largest upload, provides on average $\Psi_1 = 255.10 \times r$ chunks regardless of the fact it requests only $\Theta_1 = 50.31 \times r$ chunks to serve the 10% of the peers in the swarm. On the other hand ISP-3, that is the largest in terms of size but has a limited upload, requests $\Theta_3 = 397.14 \times r$ chunks whereas it contributes only with $\Psi_3 = 71.80 \times r$ chunks. The effect of EGT can be appreciated on the right hand side of Table 3. One observes that the EGT algorithm is able to equalize the requested and offered resources (see the values of Θ_i versus Φ_i) pushing the system towards the optimal point identified by Eq. (4), that clearly corresponds to a criterion of *effort fairness*. Also in this case the offered resources are exploited according to the rules of the P2P application determining the values of the upload that are actually provided. It turns out that by means of EGT, ISP-1 is able to reduce its upload towards the other ISPs down to $\Psi_1 = 48.47 \times r$ chunks. To counteract this ISP-3 is asked to increase the resources it provides from about $72 \times r$ to $86 \times r$ chunks. Finally, it is worth pointing out that the EGT optimization does not incur any penalty in terms of the offered service leading to values of QoS_i larger than 0.98 in all the simulated cases.

Table 4 shows the results obtained by changing the behavior of the P2P application, that in this case implements the variable probability neighbors selection criterion ($p = m_i / \sum_{j=1}^3 m_j$). The obtained results are quite similar to the one we have just commented except for the absolute values of the requested and provided resources, that clearly depend on the efficiency of the P2P application. As an example, in absence of EGT, the variable probability criterion turns out to be less unfair for the ISP-1 ($\Psi_1 = 81.86 \times r$ chunks) with respect to the results in Table 3. Nonetheless, the EGT is again able to equalize the resource usage leading to a lower $\Psi_1 = 70.74 \times r$ chunks without impacting on the quality of service.

From previous analysis of the experimental results we can conclude that *the EGT improves the resource usage fairness but does not affect the quality of the service*, provided that enough resources can be allocated to the service by the ISPs so as to guarantee service continuity. Moreover, we have shown that *the EGT does not depend on the neighbor selection criterion*, thus turning out to be useful independently of the approach adopted by the P2P application to build and maintain the streaming overlay.

Table 5

Experiment with variable probability neighbor selection with $c_3 = 33, 32, 31, 30, 29$, that is, $c_3 = 0.9428 \times r$, $c_3 = 0.9142 \times r$, $c_3 = 0.8857 \times r$, $c_3 = 0.8571 \times r$, and $c_3 = 0.8286 \times r$ (the table only reports the values for QoS_i).

	No-EGT			EGT		
	ISP ₁	ISP ₂	ISP ₃	ISP ₁	ISP ₂	ISP ₃
$c_3 = 33$	0.9829	0.9831	0.9828	0.9834	0.9836	0.9727
$c_3 = 32$	0.9829	0.9832	0.9829	0.9823	0.9829	0.8560
$c_3 = 31$	0.9832	0.9835	0.9836	0.9824	0.9829	0.7378
$c_3 = 30$	0.9837	0.9839	0.9836	0.9819	0.9823	0.6583
$c_3 = 29$	0.9836	0.9843	0.9842	0.9826	0.9829	0.6054




Fig. 4. Simulation traces for the scenario with $c_3 = 32$ ($c_3 = 0.9142 \times r$), left plot with $\alpha = 0$, right plot with $\alpha = 0.3$.

Table 6

Simulation experiments with $c_3 = 32$ using different values of α .

	$\alpha = 0$			$\alpha = 0.3$		
	ISP ₁	ISP ₂	ISP ₃	ISP ₁	ISP ₂	ISP ₃
m_i	92.69	98.37	823.34	98.35	101.33	805.74
ϱ_i	84.16	88.78	155.59	88.64	91.06	160.50
Φ_i	70.53	73.50	135.35	102.77	105.20	184.99
QoS_i	0.9823	0.9829	0.8560	0.9833	0.9835	0.9831

For the next experiment we slightly modify the system parameters of Table 2. In particular, we set $c_3 = 33$, $c_3 = 32$, $c_3 = 31$, $c_3 = 30$, and $c_3 = 29$ so that the resources offered by ISP-3 are lower than the minimum required. In other words, the peers in ISP-3 exhibit a (light) free-rider behavior. Table 5 compares the values of QoS_i obtained using and not using the EGT. For conciseness we present only the results for the variable probability neighbor selection criterion. In Table 5 it is important to note that the EGT, increasing the effort fairness, penalizes the free-rider ISP (see the values for QoS_3 in the two cases). Therefore we can conclude that *the EGT can be used to discourage free-riding at the ISP level*.

In this last part of the experimentation we use EGT and the super-tracker for a different purpose, namely to reserve a margin of spare resources in the system to counteract abrupt dynamics due to peer churning or flash crowd or simply to improve the overall quality of service. The super-tracker driving the EGT can also be used to force the ISPs to offer more than what is strictly necessary in order to globally improve the quality of service. Such an approach could be employed to limit the effects of bandwidth resource fluctuations that may arise in case of high peer churning. This goal can be achieved by modifying the universal-streaming test performed by the super-tracker as follows:

$$c_s + \sum_{i=1}^k \Phi_i \geq \left(\sum_{i=1}^k \Theta_i \right) \times (1 + \alpha), \quad (6)$$

where $0 \leq \alpha \leq 1$ is a parameter used to assume that the ISPs needs are larger than what they are instantaneously reporting by their Θ_i values.

The effects of Eq. (6) is shown in Fig. 4 that reports two simulation traces with $\alpha = 0.0$ (left) and $\alpha = 0.3$ (right). Each plot shows two 200 s traces representing the overall amount of requested resources (solid line) and the overall amount of received resources (dashed line) for the peers of the ISP-3. These quantities have been normalized with respect to the streaming rate r . Clearly, the time intervals during which the received resources are less than required (the dashed line lies below the solid one) correspond to periods when the video service cannot be assured to all peers in the ISP. It is evident

that $\alpha = 0.3$ guarantees a better quality of service to the peers. This qualitative evidence is also supported by the average indexes reported in Table 6. It is worth noting that the values of Θ_i remain almost constant with varying α , whereas the offered and provided resources Φ_i and Ψ_i increase with α .

7. Conclusions and future work

In this paper we have developed a game theory framework that can be used for the design of ISP strategies aiming at providing efficient and reliable support for P2P streaming application. Our framework allows us to illustrate the existence of equilibrium points, the role of possible strategies to refine these points, and how to use ideas from the evolutionary game theory to derive techniques, that a player of the game can use, to computed the (operational) points assuming only limited knowledge on the state of the other players.

The future developments of this research follow several directions. First of all, we are going to deepen the simulation based analysis show the relation among peer dynamics, the effects of these dynamics of the fluctuation of the functions Θ_i -s and Φ_i -s, and possible strategies to smoothen these variations.

Another direction we are developing concerns the use of more sophisticated payoff functions that include information on the size of the ISP and on the subscription fees it applies to its subscribers. This could be obtained by using in Eq. (3) a term g that depends on the size of the ISP, and on the subscription fees. In other words we could use a term g_i , i.e., function of the peculiarities of the ISP- i . This model enhancement could explore other different interactions/competitions among the ISPs, e.g., competition for attracting users, and competition on subscription fees. Furthermore, we can also implement other fairness criteria, based for instance, on the function g_i , i.e., that accounts for ISP rewards.

References

- [1] T. Karagiannis, P. Rodriguez, K. Papagiannaki, Should Internet service providers fear peer-assisted content distribution?, in: IMC'05: Proceedings of the 5th ACM Sigcomm Conference on Internet Measurement, 2005.
- [2] V. Aggarwal, A. Feldmann, C. Scheideler, Can ISPs and P2P users cooperate for improved performance? SIGCOMM—Computer Communication Review 37 (3) (2007).
- [3] H. Xie, Y.R. Yang, A. Krishnamurthy, Y.G. Liu, A. Silberschatz, P4p: provider portal for applications, in: Proceedings of Sigcomm, 2008.
- [4] D.R. Choffnes, F.E. Bustamante, Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems, in: Proceedings of Sigcomm, 2008.
- [5] L. Dai, Y. Cui, Y. Xue, On scalability of proximity-aware peer-to-peer streaming, in: Proceedings of IEEE Infocom, 2007.
- [6] X. Hei, C. Liang, J. Liang, Y. Liu, K.W. Ross, A measurement study of a large-scale P2P IPTV system, IEEE Transactions on Multimedia 9 (2007).
- [7] Y. Liu, L. Guo, F. Li, S. Chen, A case study of traffic locality in Internet P2P live streaming systems, in: Proc. of IEEE ICDCS '09, 2009.
- [8] S. Guobin, W. Ye, X. Yongqiang, Y.Z. Ben, Z. Zhi-Li, HPTP: relieving the tension between ISPs and P2P, in: Proc. of IPTPS, 2007.
- [9] F. Picconi, L. Massoulie, ISP friend or foe? making P2P live streaming ISP-aware, in: Proceedings IEEE International Conference on Distributed Computing Systems ICDCS'09, 2009.
- [10] V. Reddy, Y. Kim, S. Shakkottai, A.L. Narasimha Reddy, Designing ISP-friendly peer-to-peer networks using game-based control, Technical Report, The Computing Research Repository, CoRR, February 2009. <http://arxiv.org/abs/0912.3856>.
- [11] M.H. Lin, J.C.S. Lui, Dah-Ming Chiu, An ISP-friendly file distribution protocol: analysis, design and implementation, IEEE Transactions on Parallel and Distributed Systems 21 (9) (2010).
- [12] C. Buragohain, D. Agrawal, S. Suri, A game theoretic framework for incentives in P2P systems, in: Proceedings of the 3rd International Conference on Peer-to-Peer Computing P2P'03, 2003.
- [13] M. Feldman, K. Lai, I. Stoica, J. Chuang, Robust incentive techniques for peer-to-peer networks, in: Proceedings of the 5th ACM Conference on Electronic Commerce EC'04, 2004.
- [14] M. Feldman, C. Papadimitriou, J. Chuang, I. Stoica, Free-riding and whitewashing in peer-to-peer systems, in: Proc. of the ACM PINS'04, 2004.
- [15] W.S. Lin, H.V. Zhao, K.J.R. Liu, Incentive cooperation strategies for peer-to-peer live multimedia streaming social networks, IEEE Transactions on Multimedia 11 (3) (2009) 396–412.
- [16] Y. Chen, B. Wang, W.S. Lin, Y. Wu, K.J.R. Liu, Cooperative peer-to-peer streaming: an evolutionary game-theoretic approach, IEEE Transactions on Circuits and Systems for Video Technology 20 (10) (2010) 1346–1357.
- [17] J. Seedorf, S. Niccolini, M. Stiernerling, E. Ferranti, R. Winter, Quantifying operational cost-savings through ALTO-guidance for P2P live streaming, in: Proceedings of the Third International Conference On Incentives, Overlays, and Economic Traffic Control, ETM'10, Springer-Verlag, 2010.
- [18] R. Kumar, Y. Liu, K.W. Ross, Stochastic fluid theory for P2P streaming systems, in: Proceedings of IEEE Infocom. 2007.
- [19] M.J. Osborne, A. Rubinstein, A Course in Game Theory, The MIT Press, 1994.
- [20] G. Owen, Game Theory, Academic Press, 1995.
- [21] D.A. Eckhardt, P. Steenkiste, Effort-limited fair (ELF) scheduling for wireless networks, in: Proceedings of IEEE Infocom, 2000.
- [22] N. Magharei, R. Rejaie, ISP-friendly P2P streaming, IEEE Multimedia Communications Technical Committee E-Letter (2009).
- [23] J. Weibull, Evolutionary Game Theory, The MIT Press, 1995.



V. Bioglio was born in Torino in 1984. He took his Laurea degree in Mathematics in July 2008 at the University of Turin, Italy. He currently is a Ph.D. student at the Computer Science Department, University of Torino. His main research interests include game theory and error correcting codes, with a particular interest in rateless codes.



R. Gaeta received his Laurea and Ph.D. degrees in Computer Science from the University of Torino, Italy, in 1992 and 1997, respectively. He is currently Associate Professor at the Computer Science Department, University of Torino. He has been recipient of the Best Paper award at the 14-th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2006) and at the 26th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (PERFORMANCE 2007). His current research interests include the design and evaluation of peer-to-peer computing systems and the analysis of compressive sensing and coding techniques in distributed applications.



M. Grangetto (S'99-M'03-SM'09) received his Electrical Engineering degree and Ph.D. degree from the Politecnico di Torino, Italy, in 1999 and 2003, respectively. He is currently Associate Professor at the Computer Science Department, University of Torino. His research interests are in the fields of multimedia signal processing and networking. In particular, his expertise includes wavelets, image and video coding, data compression, video error concealment, error resilient video coding unequal error protection, and joint source channel coding. Dr. Grangetto was awarded the Premio Optime by Unione Industriale di Torino in September 2000, and a Fulbright grant in 2001 for a research period with the Department of Electrical and Computer Engineering, University of California at San Diego. He has participated in the ISO standardization activities on Part 11 of the JPEG 2000 standard. He has been a member of the Technical Program Committee for several international conferences, including the IEEE ICME, ICIP, ICASSP, and ISCAS.



M. Sereno was born in Nocera Inferiore, Italy. He received his Laurea degree in Computer Science from the University of Salerno, in 1987 and his Ph.D. degree in Computer Science from the University of Torino, in 1992.

He is currently Full Professor at the Computer Science Department, University of Torino.

His current research interests are in the area of performance evaluation of computer systems, communication networks, peer-to-peer systems, compressive sensing and coding techniques in distributed applications, game theory, queuing networks, and stochastic Petri net models.

Dr. Sereno is the recipient of the Best Paper award at the 26th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (PERFORMANCE 2007).



S. Spoto was born in Andria, Italy in 1976. He received his Laurea degree in Computer Science from the University of Catania, in 2007, and his Ph.D. degree in Computer Science from the University of Torino, in 2011. He currently is a temporary researcher at the University of Torino, Computer Science Department.

His research interests include peer-to-peer systems and communication networks.