

# A Novel Game-Theoretic Framework for Modeling Interactions of ISPs Anticipating Users' Reactions

Ioanna Papafili, Sergios Soursos, George D. Stamoulis

Department of Informatics, Athens University of Economics and Business, Athens, Greece  
{iopapafi, sns, gstamoul}@aub.gr

**Abstract**—Effective management of overlay traffic is crucial for ISPs, due to the high interconnection costs incurred. In this paper, we investigate the interactions among ISPs that manage effectively overlay traffic but also take into account users' reactions. We introduce an innovative game-theoretic framework that employs separately two metrics, quantifying the ISPs' interconnection costs and the effects of their actions to users' QoE with the permissible strategies at each state being “memory-based”, i.e., depending on the payoffs of the previous states. We present the details of this framework and justify why it fits nicely our problem. Furthermore, we study two games that model the adoption of ISP-driven locality promotion and of ISP-owned caches that intervene in the overlay. We formulate these games and investigate their evolution and equilibria by means of two theoretical models, one of which is introduced here for quantifying the effects of promoting locality, and simulations.

**Keywords**—Game-Theoretic Framework, Two Metric Game with Memory, Economic Traffic Management, Performance, Cost.

## I. INTRODUCTION

Overlay networks such as Peer-to-Peer (P2P) networks (e.g. BitTorrent [1]) already generate large volumes of traffic in the Internet. Also, as forecast in [2], although the percentage of P2P traffic is expected to decrease in the future as a percentage of the overall Internet traffic, i.e., only 23% in 2015, instead of 30% today, its absolute volume is expected to increase by an annual growth rate of 20-25%. These huge traffic volumes constitute a crucial problem for ISPs, since they result in high charges for inter-domain traffic. Most of the traffic management mechanisms in the literature promote traffic locality by employing alternative neighbor selection, based either on some proximity metric provided by the Internet Service Provider (ISP) ([3], [4], [5]), or on proximity vectors acquired by processing regular DNS queries, thus exploiting an existing public infrastructure such as Domain Name System (DNS) [6]. Moreover, other works propose the provision of additional resources to the overlay ([7], [8]).

Individual optimization of the overlay or the underlay results in traffic oscillations and sub-optimal Quality of Experience (QoE) for all parties [9]. In this context, Economic Traffic Management (ETM) proposes an incentive-based approach, developed by EU funded project SmoothIT (www.smoothit.org), which employs economic concepts and mechanisms to deal with overlay traffic in a way that is mutually beneficial for all stakeholders of the Internet, namely end users, Overlay Providers and ISPs, i.e. to enable

“TripleWin”. Under the ETM paradigm, these stakeholders are considered to be interacting, self-interested players, making their individually preferable choices as permissible by the mechanisms. This implies that certain interactions among these players can be analyzed by means of Game Theory. Our work provides intuition on such interactions, yet by means of an innovative game-theoretic framework. Although Game Theory has been used in related works to analyze such interactions, most of the models proposed in literature use a single metric for each ISP's payoff, which is inadequate to capture both ISP's monetary benefits and user-related effects.

Our contribution lies in the investigation of the ISPs' interactions under a more realistic approach that involves *two metrics*; namely, the monetary cost of the ISPs (players) and the QoE of their end-users, for which we argue that they should be considered separately rather than in a combined way as in [16]. We also argue that it is meaningful to assume that ISPs have *memory* of previous states' payoffs and do make decisions on their strategy based on their own and on their users' payoff at both the current and previous time steps.

In particular, the main contributions of this paper include: a) the definition of a novel game-theoretic framework to investigate ISPs' interactions when employing ETM mechanisms, b) the development of the decision making, taking into account two separate metrics, i.e. payoffs of current and previous states, in a prioritized way, c) the modeling of the impact of the users' reactions on the decision making; and the modeling of the loss of customers as infinite cost for ISPs, d) the application of the proposed framework to investigate ISPs' dynamics in two cases: i) application of ISP-driven Locality mechanism, and ii) the insertion of ISP-owned Caches in the overlay, and e) the development of a simple yet realistic model for the study of the ISP-driven Locality .

## II. RELATED WORK

Over the last years, investigation of the interactions among players performing overlay traffic management is carried out by means of Game Theory in several works. The work of [9] addresses the interaction between the underlay physical network, e.g. ISPs, and the overlay, e.g. P2P networks. The game between these two players is formulated as a game that converges to an inefficient Nash Equilibrium Point for both underlay and overlay due to the Information Asymmetry between them. This work advocates for the necessity of cross-layer optimization. Furthermore, [11] classifies possible

cooperation schemes between ISPs employing Traffic Engineering (TE) and Content Providers (CPs) constituting content distribution networks based on the information that needs to be shared among the players. In [12], a two-stage multi-leader-follower game is formulated where ISPs (leaders) set the prices for access provisioning and self-interested users (followers) respond to the prices by selecting ISP. In [13], two game-theoretic models for ISPs, one considering ISPs as self-interested and the other as altruists, aiming at reducing their bandwidth costs deploying cooperating caches are formulated. Additionally, in [14], ISPs' dynamics in locality adoption are investigated. A major conclusion drawn in this study is that initially the largest AS starts to promote locality, and then the other ones are one-by-one forced to follow this strategy too.

Moreover, in [15], a game-theoretic framework is proposed for the development of techniques to promote cooperation (in the form of mutually beneficial resource sharing) among ISPs in P2P streaming platforms. The derived strategies aim at minimizing the inter-ISP traffic, though the implication of ISPs' actions on users' performance is not considered. In [16], a game between ISPs and Content Distribution Systems is formulated and the existence of equilibria is proven; note that in [16], a combined metric is employed, instead of a single metric, as used in each of the previous approaches. However, as also explained in [16], even such a metric may not be able to capture the efficiency losses in all cases. Therefore, in this paper, we introduce the use of two metrics in a separate and prioritized way instead of a combined one.

The memorization of the payoffs of previous states in the selection of the current strategy has been addressed before, yet in different formulations and contexts than ours. For example, in [17], when the strategies for playing the iterated Prisoner's Dilemma are considered, the outcomes of the three previous moves are used to make a choice in the current move. The strategy selection for the *noisy* iterated Prisoner's Dilemma is studied in [18], where it is shown that the minimum memory required to find good strategies is 4, i.e., the strategy should take into account the action of both players in the previous two time steps. Moreover, in [19], memory-based strategies are evolved for the generalized Hawk-Dove game, and are proven to be essential for social stability.

In [20], a two player game with two choices A and B is formulated. In this framework, each agent (player)  $k$  is considered to select his strategy based on events kept in his memory  $M_k$ , where each event  $m$  is memorized as a tuple including the time of the event, the strategy of the player at that time and the reward received (payoff). In a different setup, in [21], the authors consider the minority game, which is a repeated game with odd number of players that much choose between two alternatives at each step. The action of each player is also here determined based on information, called memory, on which action provided higher payoff at the last  $m$  time steps. In our work, we consider the use of memory of the payoffs of the previous two states explicitly in conditions defining the set of permissible strategies, and we also extend our study to multi-player and multi-strategy games.

### III. ECONOMIC TRAFFIC MANAGEMENT

Economic Traffic Management (ETM) proposed, developed and evaluated by EU FP7 SmoothIT project ([www.smoothit.org](http://www.smoothit.org)) constitutes an innovative approach to manage application traffic flows in overlay networks. Its main objective is to achieve the co-operation between the overlay and the underlay, resulting in traffic patterns that optimize the use of network resources according to multiple given criteria. This is attained by means of ETM mechanisms that aim to be beneficial to all players, by combining traditional mechanisms of traffic management with economic incentives of the involved stakeholders. Below, two of the main ETM mechanisms studied within SmoothIT are briefly described.

The *ISP-driven Locality* ETM mechanism [5] is designed to incorporate the information exchange between the two layers, i.e. the overlay and the underlay in a mutually beneficial way. ISP-driven Locality employs BGP information, in order to characterize resources (e.g. servers or peers) and promote locality in the overlay's choices. In particular, overlay peers provide to an ISP-provisioned service lists of other peers possessing the content requested by each of them. The ISP service replies to the querying peer by providing ratings of these peers using rating functions that make use of BGP values. Then, the requesting peer decides on whether to select peers based on the ratings provided by the ISP service or not. It is to the ISP's best interest to provide such ratings so that the user experiences an improved download time and the ISP lower inter-domain traffic charge.

An ISP-owned Peer (IoP) [7] is an entity that also aims at increasing the level of traffic localization within an ISP and at improving the performance enjoyed by the users of the overlay application. The IoP is a resourceful entity in terms of access bandwidth (mainly in the uplink) and storage capacity. It acts as a regular peer, running the native overlay protocol, and simultaneously serves as a cache server storing selected (e.g. popular) content files and serving them to regular peers. An IoP serves local peers, and by choice remote ones. Additionally, an IoP can be totally transparent to the regular peers, or advertised by the ISP itself or the overlay provider (i.e. the tracker, if a suitable agreement is established between the ISP and the tracker). In the rest of the paper, the terms IoP and cache are used interchangeably.

### IV. DEFINITION OF A GAME-THEORETIC FRAMEWORK

In this section, we define a novel game-theoretic modeling framework for studying the dynamics of self-interested ISPs' interactions for different ETM mechanisms under which ISPs make related choices that affect each other's payoff as well as of their users. Games under the proposed framework are *non-zero sum*, *multi-player* games with *incomplete information* with *two metrics* considered *separately* and in a *prioritized* way (which is an innovative feature) that involve *memory* of two steps in order for the player to select his next strategy.

In order to model the *Information Asymmetry* among players in a realistic way, we assume that each player knows

neither the payoff matrices, nor the state of his opponent(s). Thus, an ISP plays according to his own expected payoffs only, choosing his *best response* strategy given the previous and current strategies of the other player(s) and their consequences. We only assume that each player can indeed observe and learn such consequences (not necessarily the actions) as well as estimate the expected payoffs of his own actions, e.g. by means of trial-and-error. Thus, ISPs interactions in this framework are modeled as a sequence of *best-response dynamics* with incomplete information but *known* own payoffs for each player. Of course, after several iterations we could reasonably assume that each player has a good knowledge of the payoff matrix of his opponent(s). However, the games considered may reach their equilibrium in a few steps only, thus rendering this feature irrelevant.

It should be noted that games within the proposed framework are in general non-zero sum, in the sense that both players may improve/deteriorate their states simultaneously, as well as the improvement of one player's state is not necessarily done at the expense of the detriment of the other(s); thus "win-win" may be achieved. However, we do not consider or aim at promoting cooperation between players as in [13] and [15].

The strategy space in the proposed framework can be either discrete (e.g. application of an ETM mechanism of not)  $S = \{No\ ETM, ETM_1, \dots, ETM_{S-1}\}$ , or continuous (e.g. selection of the capacity of a cache from a continuous range). In the sequel, we focus on the discrete strategy space, where in the simplest case, each ISP has to choose from two strategies, i.e. an ISP can employ or not a specific ETM mechanism, i.e.,  $S = \{No\ ETM, ETM\}$ . More complex cases include an ETM mechanism with the options of multiple variations and/or different permissible values of associated parameters etc.

Within this modeling framework, the behavior of the players  $P_i \in P, i=\{1, 2, \dots, N\}$  (i.e., ISPs) is studied based on specific metrics quantifying the consequences of their actions, which are expressed by the pair  $(C_i, Q_i)$  for each player  $P_i$ . In particular,  $P_i$  is interested in minimizing his monetary cost  $C_i$ , for which we focus on interconnection costs, i.e. cost for the inter-domain traffic crossing the link connecting ISP  $P_i$  with his transit ISP. Simultaneously,  $P_i$  is pursuing to attract more customers and to maintain existing ones. Hence,  $P_i$  is also interested in providing his users with services of satisfactory QoE, i.e. maximizing  $Q_i$ . In the proposed game-theoretic framework, the two aforementioned metrics are considered *separately* and in a *prioritized* way as analyzed below.

As the game considered is played in the form of best-response dynamics, we consider a series of discrete steps in each of which one player decides to modify his strategy because he is inclined to do so. The strategy of player  $P_i$  actually playing at time step  $m, S_i(m)$ , is selected on the basis of two important factors: a) memory of the payoffs of player  $P_i$  at the previous two steps  $m-1, m-2$ , of the game, i.e. on  $(C_i(m-1), Q_i(m-1)), (C_i(m-2), Q_i(m-2))$ , taking also into account b) the impact of the action of the opponent player  $P_j$  that played at step  $m-1$  on the payoff metrics of  $P_i$ .

The reason why we choose to keep in memory exactly the two previous states are the following: In the case of applications with elastic QoS requirements, e.g. file-sharing or non-real-time streaming media, there is no prespecified "hard" QoS guarantee. However, if an ISP's customers get used to nice features, e.g. good QoE, then they cannot accept a subsequent QoE deterioration. Thus, memorization of only one state is inadequate, since then no comparison between the payoff of the state to be and the payoff of the state before the opponent's action can be performed and therefore the player could not make a decision on his current strategy. Furthermore, memorization of three or more states is unnecessary, since then there would be no change in the decision making, since users will not accept QoE deterioration; even if their new QoE level is better than the respective one e.g. three steps earlier. On the other hand, in the case of applications with inelastic QoS requirements, each player needs to compare his customers' performance payoff to a specific QoE threshold; therefore, memorization of previous performance values is not required.

Regarding the decision making criteria, we now argue that the two metrics *should not* be combined under a single metric, as in [16], both for applications with elastic and inelastic QoS requirements. Combining the two metrics under a single one may lead the system to undesirable effects where the interconnection cost  $C$  is significantly improved for an ISP but together with noticeable deterioration of  $Q$ . The deterioration of users' performance (for elastic applications), or its dropping below the permissible threshold (for inelastic ones), imply loss of customers in the long-term for this ISP, which ultimately leads to very high monetary cost that should be definitely avoided (modeled in our framework as "infinite" cost). Hence, such a situation should be excluded from the set of feasible equilibria for the ISPs' interactions, which is achieved by considering the two metrics in a prioritized way.

In particular, if the action of the opponent of player  $P_i$  that played the last (i.e., at step  $m-1$ ) led to deterioration of  $P_i$ 's performance metric, then in the long-term this implies infinite cost for  $P_i$  due to dissatisfaction of his users by the performance they experience and their subsequent migration to another ISP. Thus, if  $Q_i(m-1) < Q_i(m-2)$ , then the player  $P_i$  must improve his performance metric at any cost, if possible. If the set of strategies that lead to performance improvement such that the new performance value is equal to or greater than the performance value prior to the opponent's move is non-empty, i.e.,  $S_Q \neq \emptyset$ , where the set  $S_Q$  is defined as  $S_Q \subseteq S, S_i \in S_Q$ , if  $Q_i(S_i, m) \geq Q_i(m-2)$ , then the player  $P_i$  chooses among those strategies that lead in improvement of his performance metric the one that results in lower cost, i.e.,:

$$S_i^*(m) = \operatorname{argmin}\{C_i(S_i, m)\}, \text{ for } S_i \in S_Q,$$

else if none of the strategies lead to performance "restoration" but the set of strategies that lead to some performance improvement is non-empty, i.e.,  $S'_Q \neq \emptyset$  (where  $S_Q$  was defined above) and  $S'_Q \subseteq S$ , which is defined as  $S'_Q \subseteq S$  and  $S'_i \in S'_Q$ , if  $Q_i(S'_i, m) \geq Q_i(m-1)$ , then the player  $P_i$  chooses among the

latter set of strategies the one that leads to the highest possible improvement of his performance metric, i.e.,:

$$S_i^*(m) = \operatorname{argmax}\{Q_i(S'_i, m)\}, \text{ for } S'_i \in S'_Q.$$

If there is no strategy that leads to any improvement of his performance metric, then the player  $P_i$  maintains his current strategy, i.e.,

$$S_i^*(m) = S_i^*(m-1).$$

Therefore, it should be noted that the performance metric  $Q$  is considered of *higher priority* compared to the cost metric  $C$  in case of a previous performance deterioration.

On the other hand, if the action of the opponent of player  $P_i$  that played the last (i.e. at step  $m-1$ ) led to improvement or no deterioration of  $P_i$ 's performance metric, i.e., if  $Q_i(m-1) \geq Q_i(m-2)$ , then, player  $P_i$  tries to improve his cost metric without deteriorating his performance metric, if possible, or otherwise he maintains his current strategy. In particular, if the set of permissible strategies for this case is non-empty, i.e.,  $S_C \neq \emptyset$ , which is defined as  $S_C \subseteq S$  and  $S_i \in S_C$ , if  $C_i(S_i, m) < C_i(m-1)$  and  $Q_i(S_i, m) \geq Q_i(m-1)$ , then the player  $P_i$  chooses among this set of strategies the one that leads to the highest possible improvement of his cost metric, i.e.,

$$S_i^*(m) = \operatorname{argmin}\{C_i(S_i, m)\}, \text{ for } S_i \in S_C,$$

else player  $P_i$  maintains his current strategy:

$$S_i^*(m) = S_i^*(m-1).$$

In the discussion above, we have presented the criteria under which a player  $P_i$  has the incentive to modify his strategy after the action of one of his opponents. If multiple players have such incentives, then we take that a randomly selected one acts first in the way presented above. Then, all other players reassess their strategy after observing the consequences of this last action.




Figure 1. Strategy decisions w.r.t. the previous state payoff.

We further illustrate our approach in Fig. 1, for a game with two players and two strategies. We depict two payoff matrices, one for inter-domain cost  $C$  and the other for users' performance  $Q$  at the application level, for either an elastic application such as file-sharing, or one with stricter requirements such as video streaming. These metrics are *not* treated symmetrically in order to determine the equilibrium point and the achieved monetary payoffs. In particular, the game's initial state is (No ETM, ETM) and ISP2 decides his strategy. If he plays No ETM, his users will experience performance degradation (depicted in red in the upper left cell of the performance-payoff matrix), although this would benefit

him more (as depicted in green in the upper left cell of the cost-payoff matrix) than the current state in terms of inter-domain traffic cost. Thus, ISP2 maintains his current strategy. However, ISP1 changes his strategy due to the improvement of both his users' performance and his costs (bottom right cells in performance (yellow) and cost (green) payoff matrices), thus leading the system to the new state (ETM, ETM).

## V. LOCALITY PROMOTION GAMES

In this section, we provide a generic theoretical framework for modeling BGP-Locality promotion in overlay networks, as an effect of the ETM deployment. The purpose is to study a game between two ISPs and gain insight on how each ISP reacts on the locality-promoting decisions of the other, considering both the cost savings and the user performance effects, as mentioned earlier. Modeling such an environment is quite complex since there are several parameters to consider, falling under two main categories: i) network configuration (i.e., capacity of access and core links, topology, costs and inter-domain SLAs, background traffic, overlay traffic) and ii) overlay instantiation (i.e., distribution of peers, churn rate, swarms, mix of seeds and leechers, file size). In this section we provide a simplified analysis, keeping some of the above parameters fixed, so as to have a first insight on the outcome of such a game, while in Section 5, we consider a specific ETM mechanism that leads to traffic localization.

### A. Notation

The two ISPs are considered to be of Tier 3, both customers of the same Tier 2 ISP, thus being neighbor ISPs. The Tier 2 ISP connects both of them to the rest of the Internet, as shown in the topology of Fig. 2. Also, we assume that the intra-domain topology is a star topology, where the hub is the gateway to the upper tier ISP. Based on the topology, we identify 3 types of paths: the intra-domain path,  $p_{intra}$ , the path from one tier 3 domain to the neighboring domain,  $p_{neigh}$ , and the path from a domain to the rest of the Internet,  $p_{inter}$ . We do not consider core and access links to reach a peer residing on an average domain of the rest of the Internet, due to simplicity reasons, and the assumption that possible bottlenecks do not lie there. The links are considered to be symmetric, with reserved capacity for each direction.

The number of peers in each domain is denoted by  $n_j$ , where  $j \in \{0, 1, 2\}$  and domain 0 stands for the rest of the Internet. A random peer is not aware of the entire population of peers in all domains, but only of a fraction of them, denoted by  $a_k$ , i.e., the percentage of known peers (the index "k" stands for "known"). Note that to simplify the analysis, we assume that a peer in domain  $j$  knows  $a_k \cdot n_j$  other peers in its own domain, rather than  $a_k \cdot (n_j - 1)$ . Furthermore, a peer starts the exchange of data with another peer, if it gets unchoked. The probability  $p_u$  of a peer  $P$  to get unchoked when trying to connect to another peer  $R$  depends on: the number of unchoking slots available to peer  $R$ , as well as on whether the domain where the peer  $P$  and/or  $R$  belong to promote locality, and the population of peers in the various domains. The

number of unchoking slots is taken as equal to 4 [1]; though the model also applies to other values for this parameter. Since we don't consider a specific locality-promoting ETM mechanism but rather aim to develop a generic model, we refrain from addressing operational details of the mechanism itself, but rather introduce a parameter expressing its effect. Indeed, let  $p_l$  denote the percent reduction of the inter-domain traffic when a domain promotes locality.




Figure 2. The topology under study.

### B. Assumptions

We consider the case of a single swarm and all peers are assumed to be leechers. In the analysis that follows, we consider only the effects of Tit-for-Tat (T4T) principle employed by BitTorrent, given the specific number of available regular unchoking slots, and neglect the optimistic unchoking. This implies that each peer exchanges content with a fixed-size set of neighbors, either initially provided by the application provider, due to explicit agreement with the ISP, or emerged from the overlay procedures of neighbor selection.

The model works on the flow level. One flow is defined as the transfer of one or more chunks between two peers. This transfer happens after a peer is unchoked by another peer. Since a peer has 4 unchoking slots, 4 upload flows can be established during a transfer period, which we take, for simplicity, as having fixed duration and denote as 'round'. By the same token, the average number of download flows per peer is also 4. Once unchoked by another peer during a round, a peer can download as many chunks as possible according to the download rate achieved. Hence, the total download rate of a peer is used as a proxy for performance, since it can be employed to calculate the expected download time. For each flow, the corresponding rate depends on the number of flows accommodated in the bottleneck link of its path. This is related to the number of peers too, since the number of flows from one domain to another depends on the peer population in each domain and on whether a domain employs locality or not.

As a consequence of the T4T rule, we can specify the number of peers (either local or remote) that an arbitrary peer knows when it joins the swarm (see example below):

a peer in domain 1 knows  $a_k n_2$  peers in domain 2  
a peer in domain 2 knows  $a_k n_1$  peers in domain 1  
a peer in domain 1 is known by  $a_k n_2$  peers in domain 2

This implies that the T4T rule applies in the domain level too, i.e., the number of flows originating from domain 1 and are destined to domain 2 is the same with the number of flows originating from domain 2 and are destined to domain 1. Additionally, the last bullet makes the T4T rule applicable in the peer level as well.

Finally, the preservation of the number of the known overlay neighbors is a fundamental assumption for our model. Each peer initially knows a set of local and remote peers and opens a specific number of connections to a subset of those peers. However, the deployment of a locality-promoting mechanism either in one or both domains results in denying some of these connections which, in turn and due to the overlay procedures for neighbor selection, result in a new set of known peers, with the same size but under different distribution across the domains. This assumption will be later used to calculate how the origins of flows are distributed among the different domains, depending also on which of the domains employ locality promotion.

### C. Modeling Traffic Flows

We now present the model for analyzing the traffic flowing between the two ISPs, taking into account the population of peers in each domain in the following three cases: 1) no ISP employs locality, 2) only one of the ISPs promotes locality, and 3) both ISPs promote locality.

#### 1) No locality promotion in any domain.

First, we assume that no ISP promotes locality. If we focus on domain 1, we need to specify the incoming (download) traffic to a peer of this domain. The sources of the flows with destination a peer in domain 1 can have three origins: domain 1 (intra-domain traffic), domain 2 (neighbor traffic) or the Internet (Internet traffic). The aggregation of neighbor and Internet traffic gives the (total) inter-domain traffic.

First, we calculate the probability  $p_u$  of a peer to get unchoked when trying to connect to a contacted peer. In any of the domains this probability is given by the total number of unchoking slots of the peer contacted divided by the number of peers trying to connect with the specific peer, which due to T4T equals the total number of known peers

$$p_{u-no\_loc} = \frac{4}{a_k(n_1 + n_2 + n_0)},$$

Then, we calculate the traffic (number of flows) per cache, as specified earlier. For the intra-domain traffic (flows with source and destination peers in domain 1), we have that the  $n_1$  peers know  $a_k \cdot n_1$  other peers in domain 1 and they connect to them with probability  $p_{u-no\_loc}$ . Hence, the average total number of intra-domain flows is equal to  $n_1 \cdot a_k \cdot n_1 \cdot p_{u-no\_loc}$ . Following the same analysis, we obtain that the average total number of flows from the neighbor ISP (domain 2) and from the rest of the Internet is equal to  $n_1 \cdot a_k \cdot n_2 \cdot p_{u-no\_loc}$  and  $n_1 \cdot a_k \cdot n_0 \cdot p_{u-no\_loc}$ , respectively. From these, we later calculate the throughput per peer for all three domains (due to symmetry) considering also the bottleneck in the various path types.

### 2) Locality promotion only in domain 1.

Let us now assume that ISP 1 decides to deploy a locality promoting mechanism, while ISP 2 does not. As a result, and compared to the previous case of no locality promotion, the distribution of neighboring peers will be different since a percentage  $p_1$  of the inter-domain traffic will be denied/rejected. Hence, a peer in domain 2 will now know more peers from domain 2 and the rest of the Internet so that the total number of neighboring peers is preserved. Following the same approach as before, the probability that a peer in domain 1 unchokes an incoming connection is given by:

$$p_{u-loc\_1} = \frac{4}{a_k((1+x)n_1 + (1-p_1)n_2 + (1-p_1)n_0)},$$

where  $x$  denotes the percentage of new known peers in domain 1 that the new neighbor distribution provides due to the percentage decrease of  $p_1$  in the inter-domain traffic. For a peer in domain 2, it will hold that

$$p_{u-no\_loc\_2} = \frac{4}{a_k((1-p_1)n_1 + (1+y)(n_2 + n_0))}.$$

The factor  $(1-p)$  is inserted due to the T4T rule and  $y$  is the percentage of new peers from domain 2 included in the new neighbor set to compensate for the reduction of domain 1 peers known. Also, in the rest of the Internet, there will hold:

$$p_{u-no\_loc\_0} = \frac{4}{a_k((1-p_1)n_1 + (1+y)n_2 + n_0)}.$$

According to the neighborhood size preservation assumption, it should hold that

$$a_k((1+x)n_1 + (1-p_1)(n_2 + n_0)) = a_k(n_1 + n_2 + n_0),$$

$$a_k((1-p_1)n_1 + (1+y)(n_2 + n_0)) = a_k(n_1 + n_2 + n_0),$$

from which we obtain that

$$x = p_1(n_2 + n_0)/n_1, \text{ and } y = p_1 n_1/(n_2 + n_0).$$

Similarly to C.1 and based on the new distribution of peers, we can calculate the number of incoming/outgoing flows for both domains.

### 3) Locality promotion in both domains.

The number of incoming/outgoing flows for domains 1 and 2 when both of them employ the locality promotion mechanisms can be studied similarly to C.1 and C.2.

### D. Download Rate Calculation

We can now calculate the utilization of the involved links, given their capacities. From the links' capacity and the number of flows per link, we will specify the bottleneck for each type of path and hence estimate the average download rate experienced by a peer in a given domain. Note in the previous section we have been dealing with the total traffic per domain, while in this section we consider the rate with which a single peer in a domain downloads the file, and more precisely, we consider the average total download rate per round, i.e., to

download chunks of the file from all peers by which a given peer has been unchoked.

Hence, the total download rate considering the three different types of traffic and paths (intra-domain, neighbor, Internet) and the bottlenecks of each type of path, is given by  $r_t = z_{intra} \cdot r_{t_{intra}} + z_{neigh} \cdot r_{t_{neigh}} + z_{inter} \cdot r_{t_{inter}}$ , where  $z_x$  denotes the total number of flows for each type of traffic flow  $x$  (as computed in previous subsection) and  $r_{t_x}$  denotes the average rate achieved per flow in the bottleneck of the specific type of path. In order to find the average rate achieved in each type of path, the following algorithm is considered:

1. Compute the average number of flows each link.
2. Compute the fair share of the bandwidth per flow in each link, employing the average number of flows and the link capacities.
3. Find the bottleneck of each path, i.e. the link with the smallest fair share.
4. Compute the average path rate for each flow, by taking the fair bandwidth share of flows at the bottleneck link.

### E. Game Analysis

Below, we provide a set of simple numerical examples that will illustrate better the findings of the previous subsections. We first decide on the capacity, in terms of the number of flows supported per unit of time (i.e. round) of the different types of network links. More specifically, we have that  $l_1 = 10$ ,  $l_2 = 5(n_1 + n_2)$ ,  $l_3 = 2(n_1 + n_2)$ ,  $l_4 = 5$ , and  $l_0 = 12(n_1 + n_2)$ , where  $l_j$  denotes the capacity of link  $j$  as denoted in Fig. 2.

Note that we have dimensioned the links in such a way that the domain 1 is considered a high-speed domain, while domain 2 is a lower-speed domain, in order to study the effects of locality promotion in domains with heterogeneous access speeds. Also, we have related the capacity of the inter-domain links to the number of peers in both domains, so as to be able to study the locality effects under varying population sizes and distributions, keeping at the same time the total number of peers, and the capacity of all inter-domain links fixed. For the experiments conducted we assumed that  $n_0 = 1000$ ,  $n_2 = 500 - n_1$ ,  $n_1 \in [50, 450]$ . In the analysis that follows we consider the "symmetric" case as the case where the two neighboring domains 1 and 2 have the same number of peers, i.e., where  $n_1 = n_2 = 250$ , and the "asymmetric" case when the peer populations are unequal.

We first consider the symmetric case by examining the benefits for the end users that will reveal the feasible set of solutions for the ISPs. The performance payoff matrix of Table I depicts the performance achieved in each state (i.e. strategy-pair) in terms of download rate; the higher the rate, the better the performance. We employ as a proxy for monetary costs the inter-domain traffic incurred by each of the ISPs, contained in the other payoff matrix of Table II; the higher the traffic, the higher the associated costs, and the less profitable a state for the ISP. We assume that initially neither of the ISPs employs locality. If ISP1 decides to do so, both his performance and costs will improve; the same applies for

ISP2, both before and after the action of ISP1. Thus, the two ISPs will be led to the equilibrium  $(Loc, Loc)$ .

Our analysis for the asymmetric case, where  $n_1 = 100, n_2 = 400$ , leads to the same equilibrium. In fact, this is the case for all analyzed variations of peer allocation in the two domains. This suggests that both ISPs will employ the locality promoting mechanism, leading to improvement of both performance and inter-domain traffic compared to the initial state, and thus achieving win-win for each ISP and his users. Note that once this state is reached, none of the ISPs deviates by selecting *No Loc*, despite the resulting reduction of his cost, because of the associated performance deterioration.

TABLE I. PERFORMANCE PAYOFF MATRIX FOR THE SYMMETRIC CASE

ISP2 ISP1	<i>No Loc</i>	<i>Loc</i>
<i>No Loc</i>	9.13	4.83
<i>Loc</i>	9.40	4.86

ISP2 ISP1	<i>No Loc</i>	<i>Loc</i>
<i>No Loc</i>	9.44	5.00
<i>Loc</i>	9.45	5.00

TABLE II. COST PAYOFF MATRIX FOR THE SYMMETRIC CASE

ISP2 ISP1	<i>No Loc</i>	<i>Loc</i>
<i>No Loc</i>	833.33	602.78
<i>Loc</i>	602.78	635.19

#### F. Discussion of locality effects

Fig. 3 and Fig. 4 summarize the outcome of the game for different peer populations for the two ISPs. The x-axis denotes the allocation of peers in each domain, with the first number representing the peers in domain 1 and the second the peers in domain 2. (Recall that the total number of peers is fixed, namely 500.) For clarity reasons, we only compare the strategy-pair where no locality promotion takes place in the two domains, i.e. *NN*, with the strategy-pair of promoting locality in both domains, i.e. *LL*. Each curve is also indexed with the domain it refers to (i.e. *High* or *Low*). From Fig. 3 and Fig. 4, we can deduce that locality promotion seems to have a uniform effect on the inter-domain traffic, regardless of the peer distribution, which is explained due to the T4T rule as well. Thus, we observe similar effects on both the high-speed and the low-speed domains.

However, this does not hold for the download rates. From a first look, we can observe what is reasonably expected: the peers of the high-speed domain achieve higher download rates than the peers of the low speed domain. For the peers of the high-speed domain, in the *LL* case, as their population increases, they achieve higher rates since they start to benefit from the increased intra-domain exchanges and the still high number of low-speed peers in the other domain. In the *NN* case, this increasing trend leads the peers of the *High* domain to attain download rates that is only slightly lower than when locality is employed. This is observed for only one case close

to the symmetric point. When the number of high-speed peers increases beyond some point the increasing trend is reversed, due to the fact that high-speed peers now rely on a small number of low-speed peers, which now become the bottleneck. However, this detrimental effect is avoided when locality promotion is deployed, since low-speed remote peers are now replaced by the additionally discovered local high-speed peers. It should also be noted that, overall, in both Fig. 3 and Fig. 4, all curves corresponding to *NN* are dominated by the respective ones for the *LL*, thus always leading to win-win.




Figure 3. Inter-domain traffic for different peer allocations.




Figure 4. Download rates for different peer allocations.

In the low-speed domain, when no locality is employed, the download rate initially increases, as the population of its peers decreases, due to the fact that lower congestion is observed inside the domain while more remote higher-speed peers are contacted. This increasing trend stabilizes after a point, due to the fact that the peers' low speed access link becomes now the bottleneck. When locality is employed, the rates achieved are initially higher and increasing (as the number of low-speed peers decreases), and then stabilize at the same value with the one achieved in the no-locality case, due to the same reason as before.

## VI. CACHE INSERTION GAMES

In this section we analyze and investigate ISPs dynamics, when resourceful caches such as IoPs [7] and their variations can be inserted by one or more ISPs in the overlay. An IoP is an in-network cache controlled by the ISP and installed within his premises. Its purpose is to assist the content delivery by pre-fetching content from remote locations and delivering it to local peers. The IoP can be either deployed to be either

transparent, or “advertised” to the users by the ISP. In the first case, the IoP is preferred by regular peers due to its abundant resources and due to the T4T principle employed by BitTorrent; while in the latter case, the users are aware of the IoP's existence and free to decide whether they want to be served by it or not. The insertion of IoPs constitutes an ETM mechanism that aims at traffic localization and increase of the system's upload capacity, and therefore at inter-domain cost minimization for the ISP and possibly at performance improvement for the users.

To investigate the ISP dynamics when inserting IoPs, extensive evaluations have been performed both by means of: a) the theoretical model of [10], which extends the model of [22] for BitTorrent performance evaluation to include also caches, and is combined with a simple model for inter-domain traffic, and b) simulations, using the SmoothIT Simulator [23], in complex, more realistic scenarios where also business agreements, i.e. peering, can be considered between ISPs. Note that we have also investigated the dynamics of three competing ISPs by means of (a), however the analysis didn't provide any further insight; thus we confine our forthcoming discussion to the 2 ISPs case.

#### A. Investigation by means of the theoretical model

We consider the case of two competing ISPs, where ISP2 has 3 times more peers than ISP1 (asymmetric case); peer arrivals in ISP1 and ISP2 follow the Poisson distribution with rates  $\lambda = \{6, 18\}$ , respectively. The ISPs are interconnected by means of a common transit provider that has no peers; background traffic in the transit links is ignored. After peers finish downloading, they serve as seeders with mean seeding time  $1/\theta = 100s$ . Peers' access bandwidth is 16/1 Mbit/s, while the IoPs are assumed to have symmetric access links, taken as 50 Mbit/s. Finally, we assume a single swarm where a video of 150 MB size is disseminated.

In our setup, each ISP (player) can follow three strategies: a) no cache insertion (*no IoP*), b) transparent cache insertion (*IoP*), and c) advertised cache insertion (*adIoP*). We use *peers download rate* as performance payoff metric, and *ISP incoming inter-domain traffic* as a monetary cost metric. The relevant results are presented in Tables III and IV.

Initially neither of the ISPs employs an IoP, i.e. the initial state is (*no IoP*, *no IoP*). ISP1 is considered to play first. Although *adIoP* results in higher performance improvement than *IoP*, when compared to ISP1's previous state, i.e. *no IoP*, the *IoP* strategy is preferred since it leads in *higher* cost reduction than *adIoP* as depicted by arrow (1) in Table IV. The new strategy played by ISP1 benefits also ISP2 in terms of cost, but it implies performance deterioration for him. Therefore ISP2 needs to definitely change his strategy to improve his performance, if possible. Similarly to step 1, although *adIoP* implies higher performance improvement for ISP2, *IoP* leads in a significant performance improvement too but with a higher cost reduction; thus ISP2 follows strategy *IoP* too (arrow (2)). Since ISP1's performance metric is deteriorated by ISP2's strategy compared to the exact previous

state, ISP1 now changes his strategy to *adIoP* (arrow (3)) to improve his performance even at the expense of inter-domain traffic increase. This action by ISP1 implies deterioration of ISP2's performance metric compared again to the exact previous state. Therefore, ISP2 changes his strategy to *adIoP* too (arrow (4)), to improve his performance metric, again at the expense of inter-domain traffic increase. Thus, after four steps, the system converges to (*adIoP*, *adIoP*) state, which is an equilibrium, since none of the players can change his strategy without 'harming' his users' performance.

TABLE III. PERFORMANCE PAYOFF MATRIX FOR 2 ISPs

ISP1 \ ISP2	no IoP	IoP	adIoP
no IoP	20 6.66	17.56 7.18	17.48 7.26
IoP	20.49 4.25	18.14 4.68	18.07 4.75
adIoP	20.72 4.02	18.42 4.40	18.35 4.47

TABLE IV. COST PAYOFF MATRIX FOR 2 ISPs

ISP1 \ ISP2	no IoP	IoP	adIoP
no IoP	5	5.32	5.38
IoP	5.59 (1) ↓	4.54 (2) →	4.52
adIoP	3.12	4.25 (3) ↓	4.23 (4) →

Note that if the game were played taking as single payoff the performance metric, then the system would lead to the same equilibrium, i.e. (*adIoP*, *adIoP*). On the other hand, if the game were played taking as single metric the cost, then the system would be led to (*IoP*, *no IoP*). Although that in this state both ISPs would achieve higher reduction of their cost metrics (compared to two metric application), the respective performance metric of ISP2 is deteriorated. Practically, ISP2 would exploit the positive impact of ISP1's ETM mechanism on his inter-domain traffic (*freeriding*), but would perform no action to countermeasure the negative impact on his users' performance. The utilization of the two separate metrics allows us to avoid exactly such situations.

Let us now investigate qualitatively the dynamics in the aforementioned setup, when we also consider the amortized ETM deployment cost as part of the cost metric of the players. In particular, the transparent IoP deployment, i.e. *IoP*, (e.g., the cost of installing a few cache servers equipped with high bandwidth VDSL lines taking into account that storage cost is in general minimal nowadays) may have insignificant cost for the ISP compared to the inter-connection cost savings. On the other hand, the cost of advertising these caches, i.e. *adIoP*, to multiple users either by developing an alternative version of a BitTorrent client and diffusing it among peers or by establishing some agreement with an overlay provider, i.e. a BitTorrent tracker, might have a significant impact on the



ISP's total cost. Therefore, it is expected that the consideration of the deployment cost would lead the system to new equilibria, particularly if cache advertisement is an option. Of course, a detailed cost-benefit analysis would be required to quantify such deployment cost and its impact on the ISP dynamics; this is left for future investigation.

### B. Investigation by means of simulations

Simulation results for the case the simple case of two competing ISPs just analyzed lead to similar conclusions with the numerical ones, and are omitted due to space limitations. Next, we consider the case of two competing ISPs yet in a more complex and realistic topology of totally 9 ASes depicted in Fig. 5. We assume again peers' bandwidth equal to 16/1 Mbps (down/up), and IoP capacity equal to 50 Mbps (down/up). The content file disseminated is a video of 150 MB size. The players' allowable strategies are: *no IoP* (no IoP insertion), *IoP* (insertion of IoP without policy) and *IoPUP* (insertion of IoP with policy, i.e. the IoP serves only peers located in the same AS). Also, we assume that the overlay tracker is unaware of the caches' existence, when inserted.




Figure 5. 9-AS topology under study.

In order to calculate costs, we consider the 95-th percentile of the traffic volume per 5 minute intervals. Due to the fact that both ISP1 and ISP2 are tier-2 ISPs, we need to take into account not only the incoming traffic from the upper tier ISP (i.e. A1), but also revenues due to the outgoing traffic towards same or lower tier ASes. Note that a peering agreement is considered between ISP2 and B3; when the peering ratio (i.e.  $T_{ISP2-B3}^{95}/T_{B3-ISP2}^{95} \geq pR$ ) is violated, we assume that ISP2 charges B3 for this extra incoming traffic of B3 from ISP2, which is a source of revenue for ISP2. Also, we assume that ISP2 charges tier-3 ISP C3 for the traffic flowing from ISP2 to C3; similarly, ISP1 charges tier-3 ISPs C1 and C2. Therefore, taking into account all the above payments and revenues of ISPs 1 and 2, and assuming for simplicity that the per unit price for traffic is the same in all links, we obtain the following cost expressions per 5 min slot:

$$C_{ISP1} = T_{A1-ISP1}^{95} - T_{ISP1-C1}^{95} - T_{ISP1-C2}^{95},$$

$$C_{ISP2} = T_{A1-ISP2}^{95} - T_{ISP2-C3}^{95} - D_{ISP2-B3},$$

$$\text{where } D_{ISP2-B3} = \max \left\{ T_{ISP2-B3}^{95} - pR \cdot T_{B3-ISP2}^{95}, 0 \right\}.$$

The latter means that in case of violation of the peering ratio, B3 is charged from ISP2 for the exceeding traffic, otherwise

B3 is not charged at all. In Table V, we present the average download duration of the entire content file by peers in each player's domain, and in Table VI, the players' cost values are average values of the costs estimated as explained above.

We assume that the initial state of the game is  $\{no\ IoP, no\ IoP\}$ , from which both ISPs have the incentive to deviate. We take that ISP1 plays first, and improves both his performance and cost by playing *IoP*, depicted by arrow (1) in Table VI, and leading the game to the state  $(IoP, no\ IoP)$ . Although, ISP1 action has lead to improvement ISP2's performance metric compared to the previous state, ISP2 checks whether another feasible strategy can lead him to further cost reduction. Indeed, both strategies, *IoP* and *IoPUP*, are feasible (i.e. they lead to further performance improvement for ISP2) and moreover, they both lead to cost reduction. Therefore, ISP2 will choose the action that leads to the highest cost improvement, i.e. *IoP* (arrow (2)). Since ISP1's cost increased after ISP2's action, ISP1 will change his strategy. Since *noIoP* is not feasible for ISP1, because performance is degraded by 3,5 times, ISP1 will play *IoPUP* (arrow (3)). Although, ISP2's cost has been further improved by ISP1's action, his users' performance has been slightly deteriorated; therefore ISP2 will change his strategy to *IoPUP* too (arrow (4)). Note that *noIoP* is not feasible due to the large performance deterioration that it implies for ISP2. Finally, none of the ISPs has an incentive to change his strategy in this state, i.e.,  $\{IoPUP, IoPUP\}$ .

TABLE V. PERFORMANCE PAYOFF MATRIX FOR 2 ISPs

ISP1 \ ISP2	no IoP	IoP	IoPUP
no IoP	650 649	162 512	157 525
IoP	507 162	146 145	144 147
IoPUP	514 158	149 144	146 146

TABLE VI. COST PAYOFF MATRIX FOR 2 ISPs

ISP1 \ ISP2	no IoP	IoP	IoPUP
no IoP	23.8	16.0	16.7
IoP (1)	15.5 0.3	13.9 15.7 (2)	14.0 16.0
IoPUP	18.5	2.14 4.4 (3)	1.77 15.4 (4)
	1.90	2.57	

In this investigation we considered a simple cost metric based only on the 95th percentile of the traffic passing through the links attached to the tagged player assuming that the price per bit is the same in all inter-domain links. In a more complex analysis, where also different prices per bit would be considered, the ISPs' dynamics could considerably differ. For example, assume that the price per bit of the outgoing traffic from a tier-2 ISP towards a customer tier-3 ISP is larger than the respective price from a tier-1 ISP to that tier-2 one. As long as his users' QoE is not negatively affected, then the tier-

2 ISP has incentive *not* to employ a restrictive policy (e.g. as in *IoPUP*), in order to intentionally increase his outgoing traffic towards his tier-3. Further investigating such cases, as well as the interplay of the game studied with the choice of prices in various inter-domain links is left for future research.

## VII. CONCLUSIONS

In this paper, we studied the dynamics of ISPs that deal with overlay traffic by employing ETM mechanisms anticipating users' reactions. We introduced a novel game-theoretic framework that employs separately two metrics, quantifying the effects of ETM mechanisms to the ISP and his users, as well as memory of the previous two states' payoffs to assess current strategies. We studied games that model, respectively, the adoption of ISP-driven locality, for which we developed a special model to quantify its effects, and of ISP-owned Peers that intervene in the overlay, by means of a theoretical model presented in literature and simulations.

In the case of locality games, we have shown that, under specific assumptions for the overlay processes as well as for the topology under study, in the equilibrium both providers employ locality promotion, independently of the access speeds they offer to their customers and the distribution of the peers in the two domains. We have also shown that for the simple case of two strategies, both user-related and ISP-related metrics are compatible, i.e., they are both improved when the ISP adopts locality promotion.

In IoP game, we showed that indeed the decision making under the proposed framework can lead the system to an equilibrium that is different than the one to be reached if users' reactions were not anticipated. However, by following our approach, ISPs manage to achieve both cost reduction and performance improvements for their users at the same time, thus leading again to a win-win situation for each ISP and his users. Also, for tier-2 ISPs, we have argued for the need to take into account both the outgoing traffic towards lower tier ISPs and the charging schemes in the various inter-domain links, since these may considerably affect the game outcome.

The proposed game-theoretic framework is more generally applicable. For example, future research can focus on games modeling dynamics of multiple ISPs and for other types of overlay traffic, e.g. traffic generated by cloud applications or by social networks; moreover, different metrics could be also considered, either separately or in combination with those already employed, such as metrics related to energy efficiency.

## ACKNOWLEDGMENT

This work has been partly accomplished in the framework of the SmoothIT project. The authors would like to thank all SmoothIT partners for useful discussions on the subject of the paper, and Dr. Ev. Markakis and the anonymous reviewers for their valuable feedback. I. Papafili has been co-financed by the EU (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong

Learning" of the NSRF (National Strategic Reference Framework) - Research Funding Program: Heracleitus II.

## REFERENCES

- [1] B. Cohen, Incentives build robustness in BitTorrent, In Proceedings of the First Workshop on the Economics of Peer-to-Peer Systems, Berkeley, CA, USA, June 2003.
- [2] Cisco Visual Networking Index: Forecast and Methodology, 2010-2015, June 2011.
- [3] V. Aggarwal, A. Feldmann, C. Scheideler: Can ISPs and P2P users cooperate for improved performance?, ACM SIGCOMM Computer Communication Revue, Vol. 37, No. 3, 2007.
- [4] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz: P4P: Provider Portal for Applications, ACM SIGCOMM, Seattle, WA, USA, August 2008.
- [5] P. Racz, S. Oechsner, F. Lehrieder: BGP-based Locality Promotion for P2P Applications, ICCCN 2010, Zürich, Switzerland, August 2010.
- [6] D. R. Choffnes, F. E. Bustamante, Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems, ACM SIGCOMM Comput. Commun. Oct. 2008.
- [7] I. Papafili, S. Soursos, G. D. Stamoulis, Improvement of BitTorrent Performance and Inter-domain Traffic by Inserting ISP-owned Peers, ICQT'09, Aachen, Germany, May 2009.
- [8] R. Swaha, V. Ishakian, A. Bestavros, Angels In The Cloud, A Peer-Assisted Bulk-Synchronous Content Distribution Service, Technical Report BUCS-TR-2010024.
- [9] Y. Liu, H. Zhang, W. Gong, D. Towsley, On the Interaction Between Overlay and Underlay Routing, Proc. IEEE INFOCOM 2005.
- [10] F. Lehrieder, G. Dán, T. Hoßfeld, S. Oechsner, V. Singeorzan, The Impact of Caching on BitTorrent-like Peer-to-peer Systems, P2P 2010, Delft, the Netherlands, August 2010.
- [11] W. Jiang, R. Zhang-Shen, J. Rexford, M. Chiang, Cooperative Content Distribution and Traffic Engineering in an ISP Network, ACM SIGMETRICS, Seattle, USA, 2009.
- [12] J.H. Wang, D.M. Chiu, J.C.S. Lui, A Game Theoretic Analysis of the Implications of Overlay Network Traffic on ISP Peering, Computer Networks, 2008.
- [13] G. Dán, Cache-to-Cache: Could ISPs Cooperate to Decrease Peer-to-peer Content Distribution Costs?, IEEE Transactions on Parallel and Distributed Systems, to appear.
- [14] SmoothIT Deliverable D2.4: Performance, Reliability, and Scalability Investigations of ETM Mechanisms, March 2011.
- [15] V. Bioglio, R. Gaeta, M. Grangetto, M. Sereno, S. Spoto, A game theory framework for ISP streaming traffic management, in Journal of Performance Evaluation (Elsevier), 68, (2011).
- [16] D. DiPalantino, R. Johari, Traffic Engineering vs. Content Distribution: A Game-Theoretic Perspective, IEEE INFOCOM, 2009.
- [17] R. Axelrod, The Evolution of Strategies in the Iterated Prisoner's Dilemma, in Genetic Algorithms and Simulated Annealing (L. Davis, ed.), London, Pitman, 1987
- [18] K. Lindgren, Evolutionary Phenomena in Simple Dynamics, in Artificial Life II (C. G. Langton, C. Taylor, J. D. Farmer, S. Rasmussen, ed.) Redwood City, California: Addison-Wesley, 1991
- [19] P. H. Crowley, Dangerous Games and the Emergence of Social Structure: Evolving Memory-based Strategies for the Generalized Hawk-Dove Game, Behavioral Ecology, Vol. 12, 1996.
- [20] J. E. Kittock, Emergent Conventions and the Structure of Multi-Agent Systems, in Proceedings of the 1993 Santa Fe Institute Complex Systems Summer School, 1993.
- [21] K. Izumi, T. Yamashita, K. Kurumatani, Control of Learning Types for Efficient Resource Allocation, Proc. North American Association for Computational Social and Organizational Science (NAACSOS) Conference 2004.
- [22] D. Qiu, R. Srikant, Modeling and Performance Analysis of BitTorrent-like peer-to-peer networks, Proc. ACM SIGCOMM Conference on Applications, 2004.
- [23] SmoothIT, Del. D2.2: ETM Models and Components (Initial Version).