

Stability and Efficiency of Unstructured File Sharing Networks

Honggang Zhang, *Member, IEEE*, Giovanni Neglia, *Member, IEEE*, Don Towsley, *Fellow, IEEE*, and Giuseppe Lo Presti

Abstract—We propose two unstructured file sharing games, *unilateral and bilateral unstructured file sharing games*, to study the interaction among self-interested players (users) of unstructured P2P file sharing applications. In a *unilateral unstructured file sharing game*, players compete for network resources (link bandwidth) by opening multiple connections to each other on multiple paths so as to maximize their individual benefits. A player always allows other players to connect to itself. Multiple concurrent connections are allowed on any path between a pair of players. Per-connection throughput is determined by the transport protocol implemented by users' computers. In a *bilateral unstructured file sharing game*, users adopt a Tit-for-Tat strategy, under which an active connection between two players is set up only when they both find it beneficial. Two players can set up at most one connection between themselves and bottlenecks occur only at upstream access links in a star network. For both games, we prove the existence of an equilibrium, quantify the efficiency losses of equilibria, and demonstrate the dynamic stability of equilibria in best-response or better-response dynamic game playing processes.

Index Terms—game theory, peer-to-peer, BitTorrent, unstructured file sharing, Nash equilibrium, efficiency loss, dynamic stability, Tit-for-Tat, pairwise stability, pairwise equilibrium network.

I. INTRODUCTION

IT IS WELL known that peer-to-peer (P2P) file sharing applications (e.g., BitTorrent [1], Kazaa, eDonkey, and Gnutella [2]) are major contributors of Internet traffic. For example, [3] estimates that P2P generated 60% of all US Internet traffic at the end of 2004. Sprint's IP Monitoring Project [4] shows that in April 2003, 20 – 40% of total bytes corresponded to peer-to-peer traffic on one backbone link. According to [5] P2P traffic accounts for 30% of daytime and 70% of nighttime Internet traffic in Germany. As P2P applications contribute significantly to Internet traffic, it is important to understand the behavior and performance of these applications. The networks formed by such P2P applications can be thought of as unstructured file sharing overlay networks, as users/peers relay

data for each other over the Internet and there are no central routing controls and no well-defined network topologies.

In order to investigate the stability and efficiency of such unstructured file sharing networks, we introduce two unstructured file sharing games to study the strategic behavior of self-interested users (players) on these networks. Henceforth we use the term users and players interchangeably. The stability and efficiency are respectively captured by the equilibria in the games and the network performance at the equilibria.

Specifically, we introduce a *unilateral unstructured file sharing game* and a *bilateral unstructured file sharing game*. By *unilateral*, we mean that a player always allows other players to set up connections or data transfer sessions with itself, and only the player sending a connection request receives data from the connection. In other words, connection setup requests are always accepted and each connection/session is unidirectional. This is representative of P2P networks such as eDonkey [6]. By *bilateral*, we mean that two players set up a connection between each other only when they both find it beneficial and they both receive data through the connection. In other words, each connection is bidirectional and requires mutual consent. Note that the mutual consent captures the Tit-for-Tat strategy implemented in BitTorrent [1] applications.

In addition, in the *unilateral unstructured file sharing game*, players compete for available network resources (link bandwidth or capacity) by opening multiple connections/sessions on multiple paths, and their strategies are the numbers of sessions on available paths. Note that a player can open multiple concurrent connections on each path to another player. The data rates allocated to connections are implicitly determined by the transport protocols implemented by users' computers. The data rate allocation mechanism can be captured by the bandwidth sharing framework proposed by Kelly [7] with TCP networks as special cases [8][9]. In the *bilateral unstructured file sharing game*, we restrict a player to open either zero or one connection to another player on a star network, and we further assume that bottlenecks only occur at upstream access links of players. A player's upload bandwidth is equally shared by other players that download data from this player. Note, in the unilateral file sharing game, proportional fairness is assumed by models on star networks and parallel networks, and link capacities are not assumed to be the same.

We introduce a benefit function to capture the performance that a player attempts to optimize. The benefit is a combination of a utility function (increasing and concave in throughput) and a cost associated with maintaining connections.

The contributions of this paper are summarized as follows. We show the existence of multiple Nash equilibria (NEs) in

Manuscript received on August 15, 2007; revised on March 10, 2008. This paper was presented in part at the 26th Annual IEEE Conference on Computer Communications (INFOCOM), Anchorage, Alaska, USA, 6-12 May 2007.

Honggang Zhang is with the Department of Math and Computer Science, Suffolk University, Boston, MA 02114, USA (e-mail: hzhang@ieee.org).

Giovanni Neglia is with INRIA, Sophia Antipolis - Méditerranée, 06902 Sophia Antipolis, France, and Università degli Studi di Palermo, Italia (e-mail: giovanni.neglia@ieee.org).

Don Towsley is with the Department of Computer Science, University of Massachusetts Amherst, Amherst, MA 01003, USA (e-mail: towsley@cs.umass.edu).

Giuseppe Lo Presti is with CERN, European Organization for Nuclear Research, CH-1211 Geneva 23, Switzerland (e-mail: giuseppe.lopresti@cern.ch).

Digital Object Identifier 10.1109/JSAC.2008.080925.

the unilateral unstructured file sharing game. We then use parallel link and star networks (similar to [10][11][12][13]) to model unstructured file sharing, and prove the existence of an NE. We show that the price of anarchy (the worst efficiency loss of NEs) [14] is upper bounded if users have resource constraints. Note that the efficiency loss of an NE is defined as the ratio of optimal system performance over the system performance at the NE. We show that a dynamic game playing process converges to an NE if users play a variant of the game in a best-response fashion. For the bilateral unstructured file sharing game, we show the existence of *symmetric* pairwise equilibrium networks (PENs) in a homogeneous star network and that the price of anarchy is unbounded (for linear cost functions) even though Tit-for-Tat is used to prevent selfish behavior. We also characterize *asymmetric* PENs and show that all PENs asymptotically converge to symmetric PENs as the number of players increases. We further show that a PEN can always be achieved in a better-response dynamic game playing process when the cost function is linear.

The rest of this paper is organized as follows. A formal description of the unilateral unstructured file sharing game is presented in Section II. Sections III and IV focus on the unilateral unstructured file sharing game on parallel link and star networks respectively. The bilateral unstructured file sharing game is discussed in Section V. Related work is presented in Section VI, and Section VII concludes the paper.

II. UNILATERAL UNSTRUCTURED FILE SHARING GAME

A. Network Formulation

A set of users $\mathbf{R} = \{1, \dots, R\}$ share a network consisting of a set of links $\mathbf{J} = \{1, \dots, J\}$, which provides a set of paths $\mathbf{P} = \{1, \dots, P\}$. Link j has a capacity or bandwidth $C_j > 0$. Each path $p \in \mathbf{P}$ consists of a sequence of links; if link j is used by path p , we denote this by $j \in p$. Each user r has a set of available paths; if path p serves user r , we denote this by $p \in r$. User r can open a number of concurrent connections n_{rp} on path p if $p \in r$. User r 's strategy is defined as $\mathbf{n}_r = (n_{rp})$ with $p \in \mathbf{P}$ and $p \in r$. A composite strategy of all users is given by $\mathbf{n} = (\mathbf{n}_1, \dots, \mathbf{n}_R)$. For a given \mathbf{n} , a data rate allocation mechanism allocates rate y_p to each connection on path p . The rate allocation mechanism is implicitly defined by the data transmission protocol implemented by users. One example in the Internet is Transmission Control Protocol (TCP). We will discuss the rate allocation mechanism in the following section. For now, we simply state that y_p is a function of \mathbf{n} . A rate allocation on all paths can be denoted by vector $\mathbf{y} = (y_p, p \in \mathbf{P})$. Any feasible rate allocation \mathbf{y} must satisfy capacity constraints: $\sum_{r \in \mathbf{R}} \sum_{p: j \in p} n_{rp} y_p \leq C_j, \forall j \in \mathbf{J}$.

Let G_r denote the total data rate or throughput received by user r , written as $G_r(\mathbf{n}_r) = \sum_{p \in r} n_{rp} y_p$. As y_p is a function of \mathbf{n} , G_r is also a function of \mathbf{n} . We assume that user r receives a utility $U_r(G_r)$ when obtaining throughput G_r , where U_r is a continuous, concave, and increasing function of G_r . A cost $\Phi_r(\mathbf{n}_r)$ is associated with opened connections. User r 's objective is to maximize a received benefit, represented as a combination of U_r and Φ_r :

$$B_r(\mathbf{n}_r) = U_r(G_r(\mathbf{n}_r)) - \Phi_r(\mathbf{n}_r). \quad (1)$$

where Φ_r is assumed to be a continuous, convex, and increasing function of \mathbf{n}_r .

One special case of Φ_r is a linear cost function where Φ_r is proportional to the total number of connections opened by user r on all of its available paths, defined as:

$$\Phi_r(\mathbf{n}_r) = \beta \sum_{p \in r} n_{rp}. \quad (2)$$

We assume $\beta \in [0, 1]$, and interpret β as the aggressiveness coefficient. A smaller β is used to characterize a more powerful computation resource. This type of cost is also considered in [15]. A user with a more powerful computational resource puts less weight β on its total number of connections when computing its connection costs.

B. Rate Allocation Mechanism

As mentioned before, the data transmission protocol implemented by users implicitly defines a data rate allocation scheme for connections opened by users. In this paper, we use the α -bandwidth allocation scheme [8][7][9]:

$$\max_{\mathbf{y}} \sum_p w_p n_p^\alpha \frac{(y_p n_p)^{(1-\alpha)}}{1-\alpha} \quad (3)$$

$$\text{s.t.} \sum_{r \in \mathbf{R}} \sum_{p: j \in p} n_{rp} y_p \leq C_j; n_p = \sum_{r: p \in r} n_{rp}. \quad (4)$$

where w_p is the weight of path p . Different values of α yield different rate allocations. For example, as $\alpha \rightarrow \infty$, the resulting allocation is Max-Min fair. Rate allocation in a TCP network is approximated by $\alpha = 2$ and $w_p = 1/T_p^2$, where T_p is the Round Trip Time (RTT) of path p .

If a network consists of only one congested link and user r has n_r flows and different users have different RTTs, then [15] shows that G_r is given by:

$$G_r(\mathbf{n}_r) = \begin{cases} \frac{n_r/T_r}{\sum_{k \in \mathbf{R}} n_k/T_k} C, & \text{if } n_r > 0 \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Note, this expression is only valid for TCP. We refer to this as the *simple rate allocation function*. If all n connections or flows sharing the link have the same RTT, then each flow gets an equal share of C . Note that (5) holds only when all connections share one bottleneck link. Furthermore the links of a path cannot be treated independently, as they all must carry the connections opened on the path. An illustrative example is given in [16]. This requirement makes a user's throughput neither concave nor convex in its numbers of connections. An example of this observation is given in the next section.

C. Unilateral Unstructured File Sharing Game

We now formally introduce the *unilateral unstructured file sharing game* to model the interaction among self-interested users that compete for the total network bandwidth by opening multiple concurrent connections on multiple paths. By *unilateral*, we mean that a user or player always allows other players to set up connections or data transfer sessions with itself, and only the player sending a connection request receives data from the connection. In other words, connection setup requests are always accepted and each connection/session is unidirectional.

Each user r tries to maximize its benefit B_r by adjusting \mathbf{n}_r . More precisely, user r tries to solve the following optimization problem:

$$\max_{\mathbf{n}_r} B_r(\mathbf{n}_r, \mathbf{y}^*(\mathbf{n}_r)) \quad (6)$$

$$\text{s.t. } n_{rp} \in (0, n_{rp}^{max}], \forall p \in \mathbf{P}_r \quad (7)$$

$$\mathbf{y}^* = \operatorname{argmax}_{\mathbf{y}} \sum_p \frac{w_p n_p^\alpha (y_p n_p)^{(1-\alpha)}}{1-\alpha} \quad (8)$$

$$\text{s.t. } \sum_{r \in R} \sum_{p: j \in p} n_{rp} y_p \leq C_j; n_p = \sum_{r: p \in r} n_{rp}$$

where vector \mathbf{n}_r is the decision variable of user r and \mathbf{P}_r is the set of paths available to user r . (8) indicates that the throughput of each connection on a path is the solution of the optimization problem defined in (3). As there is no closed form solution for the rate allocation problem (8) on a general network, we cannot obtain an explicit expression for $B_r(\mathbf{n}_r)$ for a general network. In fact, (6) is a bi-level programming problem, which is NP-hard [17]. However, as shown later, we can obtain an explicit form of $B_r(\mathbf{n}_r)$ for some specific networks such as grid networks. Note that, a related game was studied in the context of proportional games by Hajek and Gopal [18], in which the efficiency loss can be arbitrarily high. But the benefit received by a user in [18] does not include a cost of the user's own actions, only the utility for rate. Please see a more detailed discussion in Section VI.

Let \mathbf{n}_r^* represent the solution to user r 's optimization problem defined above, namely, $\mathbf{n}_r^* = \operatorname{argmax}_{\mathbf{n}_r} B_r(\mathbf{n}_r)$. A Nash equilibrium (NE) is defined as $\mathbf{n}^* = (\mathbf{n}_1^*, \mathbf{n}_2^*, \dots, \mathbf{n}_R^*)$, a vector of connections of all users such that no user can gain by unilaterally deviating from it. In the context of unstructured file sharing networks, an NE represents a stable network state produced by the interaction among users. The network performance at an NE can be quantified by the loss of efficiency of the NE, defined as:

$$L_{eff} = B_{max}/B_{ne} \quad (9)$$

where B_{ne} is the total benefit of all users when the network is at the NE, and B_{max} is the maximum total benefit. The worst efficiency loss is also known as the *price of anarchy* [14].

D. Remarks on Benefit Function and Multiple NEs

It is not necessarily true that throughput $G_r(\mathbf{n}_r)$ always increases in \mathbf{n}_r . Consider the network shown in Figure 1. Suppose user r has three paths: p_1, p_2 , and p_3 . Suppose p_1 consists of links j_1 and j_2 (each with capacity C), p_2 consists of link j_1 , and p_3 consists of link j_2 . If $\mathbf{n}_r = (0, 1, 1)$, then $G_r(\mathbf{n}_r) = 2C$. But if $\mathbf{n}_r = (1, 1, 1)$, then $G_r(\mathbf{n}_r) = 3C/2$. Thus, $G_r(\mathbf{n}_r)$ decreases in n_{p1} .

Next, we illustrate the existence of multiple NEs on a grid network (introduced in [8], and shown in Figure 2(a)). A possible instance of this grid network is a "fish" network, shown in Figure 2(b). A closed form rate allocation based on the α -bandwidth sharing mechanism for such a grid network is given in [8]; See [16] for details.

Consider the unilateral unstructured file sharing game with two players on the grid network. Suppose that player 1 uses route 1 and player 2 uses route 2, and both players have




Fig. 1. A case where the throughput of user r does not increase in \mathbf{n}_r .

the same benefit function. Assume that the players use TCP, namely, $\alpha = 2$ and $w_p = 1/T_p^2$ in (3). Suppose that all vertical and horizontal routes have RTTs of 50ms, and there are ten background flows on each vertical route. In the following, we discuss two examples of the game that differ in users' benefit functions.

First consider an example where β is 0.0005 in the linear cost function given by (2) and a user's utility equals its received throughput. In Figures 3 and 4 we plot the benefit B_1 of user 1 as a function of the number of its connections on its single available route, given that the number of connections of user 2 is 50 and 100 respectively. We observe that B_1 is neither concave nor convex in \mathbf{n}_1 . In addition, B_1 does not always increase in \mathbf{n}_1 . This suggests that existing game-theoretic results on the existence of an NE cannot be directly applied here because those results require the concavity of the benefit function [19][20].

For this game, we can use best response curves to identify NEs. The best response \mathbf{n}_r^* of player r is defined as the solution of r 's optimization problem provided that the strategies of all other players \mathbf{n}_{-r} are fixed. In Figure 5, we plot the best response curves of both players. Note that there are three intersecting points. An intersecting point is an NE because at that point, each user's strategy is the best response to the other user's strategy. Thus, there are three NEs in this game. It is also interesting to note that these two players do not share any common link (Figure 2). Their interaction arises because they share links with other common sessions.

Next, consider the case where players are aggressive in the sense that they are only concerned with total throughput and have no resource limits. That is, $B_r(\mathbf{n}_r) = G_r(\mathbf{n}_r)$. For example, consider again the two users in the previous game. Suppose that user 2 opens 100 connections. From Figure 6, we observe the throughput of user 1 is neither concave nor convex in n_1 , its number of connections on its single available path. However, note that the throughput of user 1 increases in n_1 , which can be verified by checking its first-order derivative. Similarly, user 2's throughput increases in n_2 . Therefore, if both users play the game, there is a unique NE, at which both players open their maximal allowable numbers of connections.

The above examples demonstrate that the interactions among multiple users on a general network topology can be much more complex than the single link TCP connection game in [15]. The existence and uniqueness of NE depend on network topologies and users' utility functions. In what follows, we study this game on two special networks which are used to model P2P unstructured file sharing networks.

III. PARALLEL LINK NETWORK

Parallel link networks can be used to model unstructured file sharing. In such a network, users share a common source node and a common destination node connected by a number of parallel links. For example in eDonkey [6], a peer can download a file from multiple other peers. Peers simultaneously




Fig. 2. (a) is a grid network where squares represent links. (b) is an instance of (a). $A \rightarrow E \rightarrow B$ and $C \rightarrow D \rightarrow A$ correspond to route 1 and 2 in (a). $D \rightarrow A \rightarrow E, E \rightarrow F \rightarrow D, C \rightarrow F \rightarrow B$ correspond to routes 3, 4, 5.

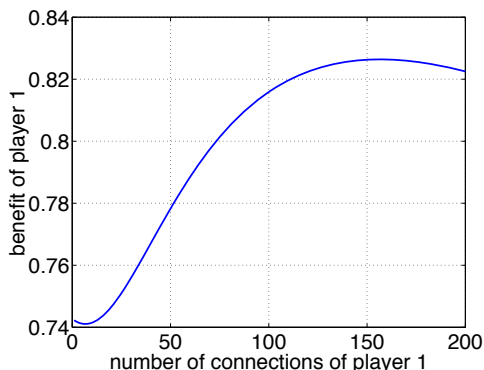


Fig. 3. Benefit of user 1 as a function of the number of connections when user 2 has 50 connections.

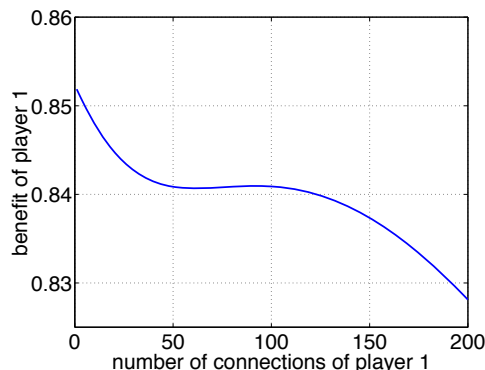


Fig. 4. Benefit of user 1 as a function of the number of connections when user 2 has 100 connections.

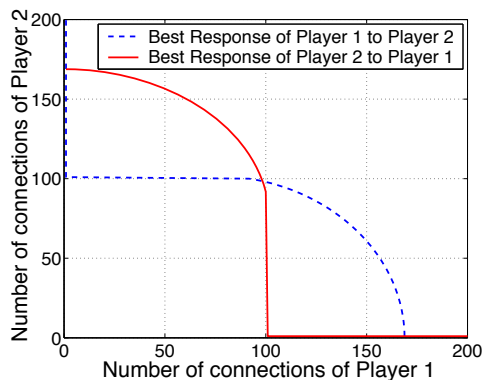


Fig. 5. Best response curves of both player 1 and player 2.

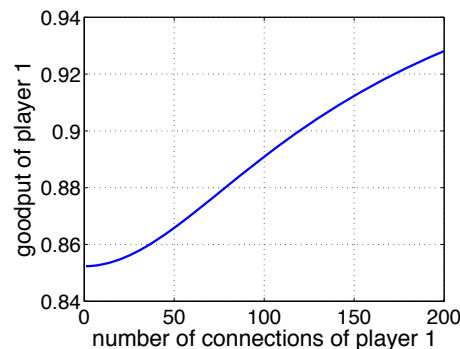


Fig. 6. User 1's data rate as a function of the number of connections on its path.

downloading a common file can be thought of as residing on a common destination node. Each file-providing peer is a “link” or “path” connecting the common destination node with a common virtual file-providing source node. Downloading peers are players competing for those parallel links’ capacities. In this section we show the existence of an NE in the unilateral unstructured file sharing game on parallel-link networks and study its performance.

A. Nash equilibrium and network performance

We first restate the game formulation briefly for the parallel link network. Consider R users sharing L links on a parallel link network. One example is shown in Figure 7. The throughput $G_{r,j}$ obtained by user r on link j is given by the simple rate allocation function introduced in the previous section: $G_{r,j}(n_{r,j}) = C_j n_{r,j} / T_{r,j} / (\sum_{k=1}^R n_{k,j} / T_{k,j})$, where $T_{r,j}$ is the Round Trip Time of user r from source node S to destination node D , C_j is the capacity of link j , and $n_{r,j}$ is

the number of connections of user r on link j . The strategy of user r is a vector of the number of connections on its available paths or links, given as $\mathbf{n}_r = (n_{r,1}, \dots, n_{r,L})$, where $n_{r,j} \in (0, n_r^{max}], \forall j \in \mathbf{L}$ and n_r^{max} is the maximal allowable number of connections for user r . This game is a continuous kernel game [19] as we assume that a user’s strategy is a real-valued vector and users’ benefit functions are continuous.

Our game when applied on parallel links decouples in a manner similar to that in Johari and Tsitsiklis [21]. A key difference is that the social welfare in our model is an aggregate of all users’ benefits (including utility and cost), whereas the social welfare in [21] does not include the cost term. This decoupling does not hold when our game is played on a star network (shown later). It is interesting to note that the utility function in our game when applied on a single link can be transformed into one studied in Maheswaran and Basar [22] in the context of design efficient auction mechanisms for divisible resources. The existence of Nash equilibrium was




Fig. 7. A parallel-link network topology.

also proved in [22]. But, the social welfare defined in [22] differs from ours in that user costs are not included in the computation of social welfare in [22]. It is also interesting to note that Herve Moulin [23] explicitly includes the cost in the efficiency metric in his study on the efficiency and budget balance in cost sharing for infinitely divisible resource. But the utility function in the game model in [23] is completely different from ours. See Section VI for a detailed discussion.

1) *Nash equilibrium*: Recall that the benefit received by user r is: $B_r(\mathbf{n}_r) = U_r(G_r(\mathbf{n}_r)) - \Phi_r(\mathbf{n}_r)$. It is easy to see that G_r is continuous, concave, and increasing in \mathbf{n}_r (see [16] for details). As we assume that $U_r(G_r)$ is continuous, concave, and increasing in G_r , and Φ_r is continuous, convex, and increasing in \mathbf{n}_r , hence, B_r is concave in \mathbf{n}_r . As B_r is also continuous in \mathbf{n} and a user's strategy space is a compact, convex, and nonempty subset of \mathbf{R}^L . Then this game is an n -person concave game which admits an NE [20]. Hence we have the following result; See [16] for a detailed proof.

Theorem 1: A unilateral unstructured file sharing game admits an NE on a parallel link network.

Consider a special case where a user's utility equals its received throughput, and its cost is proportional to its total number of connections (defined in (2)). That is, $B_r(\mathbf{n}_r) = G_r(\mathbf{n}_r) - \beta \sum_{j=1}^L n_{rj}$, where β is the aggressiveness coefficient. We have the following result for this case.

Corollary 1: When a user's utility equals its throughput and all users have the same RTT and the same cost function defined in (2), then the game on a parallel-link network admits a unique NE. In addition, if the network is homogeneous (all links' capacities are the same), the efficiency loss of the NE is bounded, but asymptotically unbounded as the number of users increases.

A detailed proof of Corollary 1 is given in [16]. This corollary gives the number of connections of user r at the NE as $n_{rj}^* = (R-1)C_j/(R^2\beta)$ and user r 's benefit at the NE as $B_r^* = \sum_{j=1}^L C_j/R - \sum_{j=1}^L (R-1)C_j/R^2$. Therefore the total social benefit of all users at the NE is given by $B_{ne} = \sum_{j=1}^L C_j/R$. As the number of users increases, B_{ne} goes to zero.

Note that the efficiency loss of an NE is the ratio between the maximum total benefit and the total benefit at the NE. In this game, the total benefit can be represented as $B = \sum_{r=1}^R B_r = \sum_{r=1}^R \sum_{j=1}^L G_{rj} - \beta \sum_{r=1}^R \sum_{j=1}^L n_{rj}$. Then, we can derive the maximum total benefit as: $B_{max} = \sum_{j=1}^L C_j - \beta N_{min}$. Consider a homogeneous network where all links have the same capacity. Then we have $B_{max} = LC - \beta L$, as we need at least one connection for each link in order to get the bandwidth of that link. Then the efficiency loss of the NE is given by $L_{eff} = B_{max}/B_{ne} = (LC - \beta L)/(LC/R)$. This result indicates that the efficiency loss of the unique NE is bounded. However, as $R \rightarrow \infty$, we have $L_{eff} \rightarrow \infty$. This

suggests that the system performance at the NE can arbitrarily degrade as the number of users becomes large.

B. Stability of NE in best-response dynamics

Best-response dynamics [19][15] is a dynamic game playing process in which users make moves at discrete time steps, following a prescribed sequence. When a user makes a move, it adopts its best response (that maximizes its benefit) to other users' strategies at the previous step. An NE is *globally stable* with respect to a certain moving sequence if the best-response dynamics (in which players make moves following the given moving sequence) converges to the NE no matter where the game starts in the feasible strategy space of the game.

The best response of a user is a solution of its benefit maximization problem given that other users' strategies remain fixed. Consider the special case in the last section where $B_r(\mathbf{n}_r) = G_r(\mathbf{n}_r) - \beta \sum_{j=1}^L n_{rj}$. We can show the best response function of user r is a continuous and concave function of the numbers of connections of other players. Then, based on [15], we can conclude (in the following theorem) that the unique NE in a two-player variant of the game is globally stable. A detailed proof is found in [16].

Theorem 2: In a two-player unilateral unstructured file sharing game on a parallel link network, when a player's utility equals its throughput and all players have the same RTT and the same linear cost function defined in (2), the unique NE is globally stable in the best-response dynamics.

C. Constrained Game

Consider another game where the total number of connections allowed by a user to open is fixed. This game is motivated by the original implementation of BitTorrent [1] where each peer always has exactly five active connections to five different other peers. Note that in BitTorrent, a peer can open up to 80 connections, but only five of them actively send and receive data. We refer to this game as a *constrained game*. Consider the special case in the last section where $B_r(\mathbf{n}_r) = G_r(\mathbf{n}_r) - \beta \sum_{j=1}^L n_{rj}$. We have the following theorem for this game.

Theorem 3: In a constrained unilateral unstructured file sharing game, assume that a user's utility equals its received throughput, a user's cost is proportional to its total number of connections and all users have the same RTT and the same linear cost function. Then, the game admits a *unique symmetric* Nash equilibrium on a parallel-link network.

The basic idea for proving this theorem is to construct Lagrangians for the constrained optimization problems of all users, then one can find the optimal solutions for those Lagrangians, which further gives us the symmetric NE. See [16] for the proof of this theorem. Note that at this symmetric NE, all users have the same number of connections. There may also be *asymmetric* NEs. See [16] for details.

As the total number of connections of user r is fixed at n_r , the maximum total benefit can be derived as $B_{max} = \sum_{j=1}^L C_j - \beta \sum_{r=1}^R n_r$. The optimal system performance is the same as the system performance at the unique symmetric NE.




Fig. 8. An example of star network.

IV. STAR NETWORK

We now investigate the unilateral unstructured file sharing game on a star network which can be used to model a P2P file sharing overlay. As shown in an example in Figure 8, we assume that a user has two access links to the Internet: one downstream link and one upstream link. A user r uses its downstream link to get data from other peers, and it is a “private” link in the sense that this link is only used by user r itself. On the other hand, the upstream link of user r is shared by all peers that download files from user r , so the upstream link of user r is a “public” link from user r ’s point of view. Usually the downstream link has higher capacity than the upstream link. This asymmetric access link assumption is supported in a measurement study in [24], where it is found that most users in current peer-to-peer networks use cable modem or ADSL to connect to the Internet. Similar to [10][11][12][13], we assume that bottlenecks occur only at access links. This is a reasonable approximation of the current P2P file sharing networks such as Gnutella and BitTorrent, where the data throughput of a connection is usually limited by the “last mile” (cable or ADSL or modem) of the connection. Thus, in the example shown in Figure 8, the Internet cloud can be simply represented as a center node in a star network, and all peer nodes (A, B, \dots, I) are connected only to the center node.

In this section, we prove the existence of an NE in the unilateral unstructured file sharing game on a star network, and discuss the best-response dynamics and the loss of efficiency of NE.

A. Nash Equilibrium

Recall that the benefit of user r is given by (1), restated as: $B_r(\mathbf{n}_r) = U_r(G_r(\mathbf{n}_r)) - \Phi_r(\mathbf{n}_r)$. In the following, we first present a lemma and later use it to prove that a user’s benefit function B_r is concave in \mathbf{n}_r . Then following a similar argument for Theorem 1, this game is an n -person concave game which admits an NE [20].

Consider the network in Figure 9. User r is at node D , and there are m paths along which user r can get data from peers S_1, S_2, \dots, S_m . All paths share a common link CD , which can be thought of as the “private” downstream link of user r . On each of those links from S_1, S_2, \dots, S_m to node C , user r competes with other users’ data flows. That is, those links are “public” links, as discussed earlier. A strategy vector of user r is $\mathbf{n}_r = (n_{r1}, n_{r2}, \dots, n_{rm})$ with $m \geq 1$. The following lemma shows that user r ’s throughput G_r is a concave function of \mathbf{n}_r . The proof of this lemma is found in [16]. The outline of the proof of this lemma is the following. First we prove a base case where each user only has two paths. By exploring the structure of the throughput function, we can show it is a concave function of the vector of numbers of connections. Then, this result can be extended to more than two paths.




Fig. 9. A network where a user has multiple paths to get data.

Lemma 1: In a star network, player r ’s throughput G_r is a concave function of its number of connections $\mathbf{n}_r = (n_{r1}, n_{r2}, \dots, n_{rm})$, where m ($m \geq 1$) is the number of paths player r has.

Based on Lemma 1, we can establish the existence of an NE on a star network in the following theorem. The idea of proving this theorem is as follows. Based on the result from Lemma 1, we can show that a user’s benefit function is a continuous and concave function of its vector of numbers of connections. Then the game can be shown as a n -player concave game which admits an NE [20]. See [16] for a detailed proof of this theorem and simulation examples to show the existence of an NE.

Theorem 4: A unilateral unstructured file sharing game admits an NE on a star network.

B. Efficiency Loss of Nash Equilibrium

In reality, most users in P2P networks use cable modem or ADSL to connect to the Internet [24], so downstream links have much higher capacity than upstream links. Consider a case where bottlenecks only occur at upstream access links in a homogeneous network and users are homogeneous. Then we can transform a star network to a parallel link network, then the efficiency loss of an NE can be immediately derived from results given in Section III. This transformation indicates that a variant of the game on a star network can be thought of as a variant of the game on a parallel link network. We illustrate this transformation in the following example.

For example, we can transform the star network in Figure 10 (assume that bottlenecks occur only at upstream access links) into the network in Figure 11. Three users, A, B , and C , are attached to nodes A, B , and C respectively in Figure 10. Note that each user has two download paths with each path consisting of two links. The strategies of users A, B , and C are: $\mathbf{n}_A = (n_{BA}, n_{CA})$, $\mathbf{n}_B = (n_{AB}, n_{CB})$, and $\mathbf{n}_C = (n_{AC}, n_{BC})$. Center node D in Figure 10 is decomposed into six interconnected virtual nodes $D_{Ad}, D_{Au}, D_{Bd}, D_{Bu}, D_{Cd}, D_{Cu}$ in Figure 11. Links between these six virtual nodes have infinite capacity. Node A is decomposed into nodes A_{down} and A_{up} . Link $D_{Ad}A_{down}$ represents the downstream link of node A . Link $A_{up}D_{Au}$ represents the upstream link of node A . Other links have similar interpretations.

Observe that user A at node A_{down} downloads data from node B_{up} and C_{up} via two paths: $B_{up} \rightarrow D_{Bu} \rightarrow D_{Ad} \rightarrow A_{down}$ and $C_{up} \rightarrow D_{Cu} \rightarrow D_{Ad} \rightarrow A_{down}$. The first path includes only one bottleneck link $B_{up}D_{Bu}$. The second one includes only one bottleneck link $C_{up}D_{Cu}$. We also have similar observations for users B and C . This transformation yields a parallel link network. Then the previous results for




Fig. 10. A simple star topology with three users A, B, and C.

the loss of efficiency at NE on a parallel link network can be applied here. That is, when the utility of a user equals its throughput and the cost of a user is proportional to the total number of its connections, the loss of efficiency of any NE is bounded. Also, the loss of efficiency is asymptotically unbounded as the number of users increases.

V. BILATERAL UNSTRUCTURED FILE SHARING GAME

The Tit-for-Tat strategy, an incentive mechanism, was introduced into BitTorrent (BT) [1]. Under this strategy a peer uploads data to a fixed number of peers from which it can download at the highest rate. Note that a peer can open up to 80 connections to other peers in BT, but only five of them actively send and receive data at any given time [1]. We refer to these five connections as *active connections*. In the rest of the paper, a connection in our games corresponds to an active connection in BT. Then under Tit-for-Tat, we say that a connection between two peers is set up only if both peers find it beneficial. This differs from other traditional P2P applications where active connections are always allowed (assume peers have sufficient computation resources), which is assumed in our previous games. In this section, we investigate this Tit-for-Tat strategy by introducing a *bilateral unstructured file sharing game*.

A. Network Model

We again assume that bottlenecks occur only at access links, so an unstructured file sharing P2P network can be modeled as a star network. Each peer is attached to a node and the center node in the star represents the Internet. Peers connect to the center node via access links. We also assume that peers always have demands for data that can be satisfied by others.

Different from previous sections, here we assume that there can only be zero or one connection between a pair of peers. Unlike the last section, we assume that bottlenecks only occur at peers' upstream access links. In other words, we assume that each downstream access link has infinite capacity. For example, a connection or virtual link (i, j) between peers i and j provides peer i with a download and an upload path. Peer i 's upstream access link is the bottleneck for all peers downloading from i . In the mean time, the bottleneck on peer i 's download path from j is the upstream access link of j . This is true for all peers. Note that on the star topology in last section, bottlenecks can appear in the downstream links. The places where bottlenecks can appear depend on the network configuration.

We now look at self-interested user behavior. Peer i obtains a share of the upload bandwidth of peer j through a connection




Fig. 11. Transformation of star network into equivalent parallel link network.

between i and j . In the mean time, other peers may also want to obtain some share of peer j 's upload bandwidth by setting up connections with j . The upload bandwidth of j is equally shared among all connections with other peers. However, if peer j adopts a Tit-for-Tat strategy, it may refuse the connection setup request from peer i if j believes that it cannot benefit from such a connection. One example is that peer i cannot provide a large enough data rate to j (as in BT). Note that a peer may want to obtain bandwidth shares of all other peers' upload bandwidths and wants to maximize its received total bandwidth. When all peers behave in such a self-interested manner, we have a game among the peers, and any stable state of this game is an overlay network consisting of a set of virtual links (or connections) among peer nodes. We refer to this game as the *bilateral unstructured file sharing game* (which will be defined formally in the next section).

The Tit-for-Tat strategy is generally considered robust in the sense that it defeats self-interested behavior and leads the system into an efficient state. To the best of our knowledge, the only analytical support for this belief is found in [12]. Note that in [12], a peer can only establish a fixed number of active connections to other peers. However, in our game, we allow a peer to freely choose the number of other peers that it wants to connect to (this is introduced in some recently developed BT clients such as μ Torrent 1.7.2 [25]). We observe that peers can change the numbers of connections to open in order to improve their performance; see [16] for details. This observation shows that users of BT can effectively cheat by not following the standard protocol in order to gain advantage. Another example is a recently proposed BT client, BitTyrant [13], in which high upload capacity peers can change the default Tit-for-Tat strategy to gain advantage.

B. Game Formulation

Based on the previous network model, we now formally introduce the *bilateral unstructured file sharing game*. We refer to peers as players and connections as links. Note that, by *bilateral*, we mean that two players set up a connection between each other only when they both find it beneficial and they both receive data through the connection. In other words, each connection is bidirectional and requires mutual consent. The mutual consent captures the Tit-for-Tat strategy implemented in BitTorrent [1] applications. Recall that a connection in our game corresponds to an active connection in BitTorrent.

Let $\mathbf{R} = \{1, 2, \dots, R\}$ denote the set of players. Player i 's strategy is the set of *intended* connections it wants to establish, denoted by $s_i = \{s_{i,j} | j \in \mathbf{R} \setminus \{i\}\}$. Here $s_{i,j} = 1$ means player i intends to create a link to player j and $s_{i,j} = 0$ means

i does not intend so. Under *Tit-for-Tat*, both players have to agree on creating a link; hence a link (i, j) is formed if and only if $s_{i,j} = s_{j,i} = 1$. A strategy profile $s = \{s_1, s_2, \dots, s_R\}$ induces an overlay network $g(s) = \{g_{i,j}, i, j \in \mathbf{R}\}$, where $g_{i,j} = 1$ denotes the existence of link (i, j) and $g_{i,j} = 0$ denotes the absence of that link. Given a network g , we use $g+g_{i,j}$ and $g-g_{i,j}$ to denote respectively the network obtained by adding and severing the link (i, j) . We also let $N_i(g) = \{j \in \mathbf{R} : j \neq i, g_{i,j} = 1\}$ be the set of player i 's neighbors in graph g , and let $n_i(g) = |N_i(g)|$. A network is a symmetric or regular graph if $n_i(g) = n_j(g), \forall i, j \in \mathbf{R}$.

The benefit of player i is given by its download rate minus a cost for opening connections: $B_i(n_i) = G_i(n_i) - \Phi_i(n_i) = \sum_{j \in N_i(g)} C_j/n_j - \Phi_i(n_i)$. We assume that Φ_i is a convex function of n_i . The marginal benefit for i to open a new connection with j is given by:

$$\begin{aligned} b_i(n_i(g), n_j(g)) &= B_i(g + g_{i,j}) - B_i(g) \\ &= \frac{C_j}{n_j(g) + 1} - \Phi_i(n_i(g) + 1) + \Phi_i(n_i(g)). \end{aligned} \quad (10)$$

Because of Tit-for-Tat, the concept of Nash equilibrium (NE) is inadequate as it can be seen that an empty overlay network is always an NE. In order to address this issue, a solution concept called *pairwise equilibrium network* (PEN) is proposed by [26] to supplement the idea of NE with the requirement of pairwise stability.

Definition 1: A network g is a *pairwise equilibrium network* (PEN) if the following conditions hold: 1) there is an NE strategy profile that supports g ; 2) for $g_{i,j} = 0$, $B_i(g + g_{i,j}) > B_i(g) \Rightarrow B_j(g + g_{i,j}) < B_j(g)$. Any stable network state in the bilateral unstructured file sharing game is captured by a PEN.

C. Equilibria in Homogeneous Networks

We consider homogeneous networks where all peers have the same upload capacity and benefit function. We characterize PEN overlays, and evaluate the price of anarchy for these overlays.

1) *Existence of Equilibria:* As formulated, the game is a local spillovers game with strategic substitutes property studied in [27]. ‘‘Strategic substitutes’’ simply means that the part of player i 's marginal benefit that comes from another player j 's contribution decreases in the number of connections of j . This is true in our game as the marginal throughput increase to i from setting up a connection with player j is C_j/n_j , which decreases in n_j . In our game $\Phi(\cdot)$ is a convex function. Hence according to the terminology of [27] the benefit function in our game ‘‘satisfies the local spillovers property, concavity in own links and strategic substitutability’’. For a detailed description of why our game is a local spillovers game, see [16].

A key observation is that marginal benefit $b_i(n_i, n_j)$ of player i decreases in n_j and does not increase in n_i . We can transform our games to the games in [27] and modify the derivations in [27], to derive Statements 1, 2 and 3. A detailed description of how to derive these statements can be found in [16].

Statement 1: Symmetric PEN. If the number of players is even, a symmetric PEN always exists. Specifically, if $b(0, 0) \leq 0$, the empty network is a PEN; if $b(r-2, r-2) \geq 0$, the




Fig. 12. Different Pairwise Symmetric Equilibria.

complete network is a PEN; if $b(k, k) \leq 0 \leq b(k-1, k-1)$, the regular graph with degree k is a PEN. When the previous inequalities are strict, the degree of the PEN is unique.

This statement provides a procedure to compute the degree of a symmetric PEN by simply checking the marginal benefit $b(k, k)$. We observe that a symmetric PEN is not necessarily unique. Figure 12 shows two possible PENs with 8 players and degree $k = 2$. In what follows we assume $b(k, h) \neq 0$, namely, we exclude the case where a player is indifferent to creating a connection or not.

Statement 2: Asymmetric PEN. In an asymmetric PEN, there can be at most one singleton player and the rest of the network is a symmetric network with a unique degree. In an asymmetric PEN with no singleton components, if two players i and j with the same number of connections or degree k are connected (or linked) to each other, then the number of players with degrees smaller than k is at most k and any two such players must be mutually connected. In addition, if two players i and j with the same degree k do not have a link to each other, then all players with a degree greater than k are not mutually connected.

This statement implies that if two players i and j with the same degree k are connected, then the number of players with degrees smaller than k is limited (at most k). A related case is described in the first property of Statement 5 (shown below). This statement also implies that a PEN cannot consist of interlinked stars [16] with two or more central players. But this statement does not imply that a star cannot be a PEN. Note that for file sharing purposes, an overlay star topology is inefficient as the operation falls back to the server-client paradigm.

Statement 3: Asymmetric PEN. Let g be a PEN in which $n_i(g) < n_j(g)$. If $\forall u \in N_i(g), \exists v \in N_j(g)$ s.t. $n_u = n_v$, then $B_i(g) < B_j(g)$.

A corollary of this statement is that if player i 's neighborhood is included in player j 's neighborhood ($N_i \subset N_j$), then $B_i(g) < B_j(g)$. Statement 3 shows that players having more connections obtain larger payoffs than other players.

We further characterize an asymmetric PEN in the following theorem.

Theorem 5: Asymmetric PEN. When a network admits both a symmetric PEN with a unique degree h and an asymmetric PEN, then the asymmetric PEN has the following properties: 1) there can be at most h players with a degree smaller than h ; 2) if $b(h-1, h) < 0$ or $b(h, h-1) < 0$, then there can be at most $(h-m)m$ players with a degree greater than h , and their degrees are smaller than or equal to $h+m$, where m denotes the number of players with a degree smaller than h ; 3) if the cost function is linear, no player can have a degree larger than h .

A detailed proof of Theorem 5 is found in [16]. Note that the first property in this theorem states that the number of

players with degrees smaller than h is small. A similar result is described in Statement 2. Based on this theorem, we make the following observation. Let $\sum_{i=1}^R |n_i(g_{PEN}) - h|$ denote the distance between an arbitrary PEN and a symmetric PEN with degree h . We note that the degree of a symmetric PEN h depends only on Φ and C (independent of the number of players R). Then according to Theorem 5, the distance is bounded and becomes less significant as R increases. Formally, we have $\lim_{R \rightarrow \infty} \frac{1}{R} E \left\{ \sum_{i=1}^R |n_i(g_{PEN}) - h| \right\} = 0$. Similarly the average payoff per player in a PEN converges to that of a symmetric PEN, which is also observed in our simulations [16].

2) *Loss of Efficiency*: We define the efficiency loss of a PEN as follows: $L_{eff}(r, C, \Phi) = \frac{\sum_{i \in \mathbf{R}} B_i(g_{opt})}{\sum_{i \in \mathbf{R}} B_i(g)}$, where g_{opt} is a network with maximum total payoff $\sum_{i \in \mathbf{R}} B_i(g_{opt})$, and g is a PEN. Because there is a fixed number of players, the number of possible overlays players can create is finite. Therefore there exists at least one network with highest total payoff. We note that L_{eff} depends on the number of players, the upload capacities, and the cost functions of those players. We show in the following theorem (a proof is in [16]) that L_{eff} may be unbounded even for a class of linear connection cost functions. Therefore, the price of anarchy is infinite. The basic idea of proving this theorem is as follows. We consider the ratio of the total benefit of a symmetric PEN with degree k and a regular network of degree one. This ratio can be shown to approach infinity when the linear coefficient α approaches $1/k$ from the left. The proof of this theorem in [16] shows that the efficiency loss is actually independent of the network size, but it can become arbitrarily high if we change the parameter α in the cost function.

Theorem 6: The efficiency loss of a PEN is unbounded for the class of linear connection cost functions. Specifically, given an even number of players and an upload capacity C , $\forall M \in \mathbb{R}, \exists \alpha^* \in \mathbb{R}^+$ s.t. $L_{eff}(r, C, \Phi^*) > M$, where $\Phi^*(n) = \alpha^* n$.

D. Dynamic Models

We now investigate a discrete-time dynamic game playing process. We focus on linear costs ($\Phi(n_i) = \alpha n_i$). Suppose that the game starts from an empty network. At each time step a player pair (i, j) is randomly chosen. Link (i, j) is created (or kept) if both players find it beneficial. Link (i, j) (if present) is removed if at least one of the two players of that link does not find it useful. We can show that this dynamic process always reach a PEN.

Compared with the best-response dynamics in Section III, the dynamics considered here can be thought of as a *better-response dynamics*. This is because only one randomly chosen link is examined for removing or adding at each step in the better-response dynamics, whereas in the best-response dynamics, the best set of links is chosen from all possible sets to replace the current set at each step.

1) *Convergence to a PEN*: We now introduce some terminology similar to [28]. A network g' is *adjacent* to a network g if $g' = g + g_{i,j}$ or $g' = g - g_{i,j}$ for some pair (i, j) . A network g' *defeats* another network g if either $g' = g - g_{i,j}$ and $B_i(g') > B_i(g)$, or if $g' = g + g_{i,j}$ with $B_i(g') \geq B_i(g)$

and $B_j(g') \geq B_j(g)$ with at least one strict inequality. A game exhibits *no indifference* if for any two adjacent networks, one defeats the other.

In the dynamic process described before, the current network is altered if and only if the addition or deletion of a link can defeat the current network. This process produces an *improving path*, i.e. a sequence of networks g_1, g_2, \dots, g_K where each network g_k is defeated by the subsequent adjacent network g_{k+1} . There are two kinds of improving paths: those exhibiting cycles (which have infinite length) and those terminating with a PEN (a *stable state*). [29] shows that there always exists either a PEN or a cycle in any improving path. Thus, if we can rule out cycles, then the better-response dynamics always achieves a PEN. The following lemma (a theorem in [29]) provides a way to rule out cycles.

Lemma 2: Given G (the set of all possible networks g), if there exists a real valued function $w : G \rightarrow \mathbb{R}$ such that “ g' defeats g ” if and only if “ $w(g') > w(g)$ and g' and g are adjacent”, then there are no cycles. Conversely, if the network game exhibits no indifference, then no cycles implies that there exists a function $w : G \rightarrow \mathbb{R}$ such that “ g' defeats g ” if and only if “ $w(g') > w(g)$ and g' and g are adjacent”.

Based on Lemma 2, we can derive the following theorem. The basic idea of proving this theorem is to construct a function $w(g)$ that satisfies the conditions given in Lemma 2. Its detailed proof is found in [16].

Theorem 7: If the connection cost function is a linear function $\Phi(n) = \alpha n$, the better-response dynamics always achieves a PEN.

Our simulations have verified our analytical results presented up to now. See [16] for details of these simulations. An interesting observation from our simulations is that the average number of iterations to reach a PEN is on order R^2 . Consider this number in the context of BitTorrent (BT) [1]. Each peer in a BT network tries to replace an existing active connection with a new, better connection every 10 seconds. All peers do such replacements asynchronously in parallel, unlike the sequential replacement in our simulations. So R^2 iterations in our simulations corresponds to $10R$ seconds in BT. For a population of 100 peers, the time needed to reach a PEN is on order 17 minutes, which is probably faster than the average time between changes in the population of peers (due to arrivals or departures).

VI. RELATED WORK

The efficiency losses of Nash equilibria was studied in economics [30], and computer science literature [31][14][32] where most of the results focused on network routing and network design.

Johari and Tsitsiklis [21] study a congestion game where users compete for each link independently from other links in the network. But this independence characteristic is not true for our unilateral unstructured file sharing game, because if a user opens a connection on a path, then all links of this path must carry this connection. On the other hand, a special case of our game on the parallel link topology, exhibits this decoupling property. But [21] and our work still differ in how to evaluate the efficiency loss. There is a specific physical meaning of the efficiency loss in our game model as our game models

a specific kind of applications in the Internet. Actually the proportional sharing mechanism adopted in [21] can be traced back to Hajek and Gopalakrishnan [18] where a proportional sharing mechanism was considered in the competition among Internet autonomous systems. Their model can be thought of as a network consisting of only a single link, and a user is allocated a fraction of the link's bandwidth equal to its bid divided by the sum of all users' bids. It is interesting to note that in the context of designing efficient auction mechanisms for divisible resources, Maheswaran and Basar [22] studied a user utility function similar to that in our game on a single link. The existence of Nash equilibrium was also proved in [22]. But, the social welfare defined in [22] differs from ours in that user costs are not included in the computation of social welfare in [22]. Another interesting related work is by Herve Moulin [23], which studied the efficiency and budget balance in cost sharing for infinitely divisible resource. Each user's utility in [23] is its received value minus its cost, both being functions of its strategy, which is its demanded outcome. In our model, a user's received value is a function of its bandwidth share which is then a function of its number of connections (the strategy of this user). The user strategy in our game cannot be mapped to a user's strategy in [23] as the former cannot be interpreted as demand. In addition, the bandwidth share received by a user in our game cannot be mapped to a user's strategy in [23] as the total bandwidth shares of all users equals to the bottleneck link on a path. But the total demands of all users are not fixed in [23]. Another related work is a game studied in [33]. File downloaders compete for the bandwidth of a single file provider (the source node). The utility received by a downloader is a function of its received bandwidth share. [33] shows that there exists a Nash equilibrium where each user's received bandwidth share is proportional to its contribution. This game differs from our model in that in order to achieve a Pareto-optimal bandwidth allocation, it stipulates that each downloader is well-behaved in the sense that its requested bandwidth share is proportional to its contribution value and the source node needs to adopt a special allocation mechanism. Smaragdakis *et al* [34] recently proposed two peer-selection policies used by individual nodes to select their neighbors in order to obtain optimal overlay topologies used by swarming protocols such as BitTorrent. This work differs from our overlay formation in bilateral file sharing game in that different peer-selection policies are considered.

[35] and [15] study the interactions among selfish TCP users competing for a single bottleneck link. [36][37] propose multi-path congestion controllers by which users can coordinate the data transfer sessions on different paths to improve data throughput. A multi-path congestion controller chooses rates at which to send data on all paths available to it. However in our models, all sessions controlled by a single user are independent congestion controllers. [12] studies how Tit-for-Tat affects selfish peers who are able to set their uploading bandwidth. Our work differs from [12] in that we assume that a user can change the number of connections to open. The analytical framework for our bilateral unstructured file sharing game is in [27][28][29].

VII. CONCLUSIONS

Motivated by unstructured P2P file sharing applications such as BitTorrent [1], we introduced unilateral and bilateral unstructured file sharing games to model the interactions among self-interested users in unstructured file sharing networks.

We demonstrated that there exist multiple stable network states (Nash equilibria) in the unilateral unstructured file sharing game on general networks. We further focus on parallel link networks and star networks for modeling unstructured file sharing. We proved the existence of NEs in several variants of the game on both networks. We found that when there are cost constraints, the loss of efficiency of NEs is bounded. In addition, we proved the global stability of NEs in some variants of the game. Furthermore, we studied the Tit-for-Tat strategy through a bilateral unstructured file sharing game, for which the notion of pairwise equilibrium network (PEN) is used to describe an equilibrium. We proved the existence of a symmetric PEN in a homogeneous star network, and showed the convergence to symmetric PENs in a better-response dynamic game-playing process when user cost is a linear function. Although the general belief is that Tit-for-Tat can prevent selfish behavior, we showed that the selfish user behavior we studied can still lead to unbounded loss of efficiency in unstructured file sharing.

ACKNOWLEDGMENT

This research has been supported in part by NSF under grant awards ANI-0085848, CNS-0519998, CNS-0519922, and EIA-0080119, and by Italian MIUR project Famous. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] B. Cohen, "Bittorrent, <http://www.bittorrent.com/>," 2006.
- [2] Wikipedia, "<http://en.wikipedia.org/wiki/peer-to-peer>."
- [3] CacheLogic, "<http://www.cachelogic.com/>."
- [4] Sprint-Labs, "Packet trace analysis, <http://ipmon.sprintlabs.com/>."
- [5] ipoque, "<http://www.ipoque.com/>."
- [6] Wikipedia, "http://en.wikipedia.org/wiki/edonkey_network."
- [7] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *J. Operational Research Society*, vol. 49, pp. 5–12, 1998.
- [8] T. Bonald and L. Massoulié, "Impact of fairness on internet performance," in *ACM SIGMETRICS*, 2001.
- [9] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 556–567, 2000.
- [10] A. R. Bhambe, C. Herley, and V. N. Padmanabhan, "Analyzing and improving a bittorrent network's performance mechanisms," in *IEEE Infocom*, 2006.
- [11] F. L. Piccolo, G. Neglia, and G. Bianchi, "The effect of heterogeneous link capacities in bittorrent-like file sharing systems," in *HOT-P2P*, 2004.
- [12] D. Qiu and R. Srikant, "Modelling and performance analysis of bittorrent-like peer-to-peer networks," in *ACM SIGCOMM*, 2004.
- [13] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "Do incentives build robustness in bittorrent?" in *4th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2007)*, 2007.
- [14] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria," *Lecture Notes in Computer Science*, vol. 1563, 1999.
- [15] H. Zhang, D. Towsley, and W. Gong, "Tcp connection game: A study on the selfish behavior of tcp users," in *IEEE ICNP 2005*.

- [16] H. Zhang, G. Neglia, D. Towsley, and G. LoPresti, "On unstructured file sharing networks," ftp://gaia.cs.umass.edu/pub/Zhang06_UFS.pdf or <ftp://gaia.cs.umass.edu/pub/Zhang06UFS.pdf>, Tech. Rep., 2007.
- [17] L. Vicente and P. Calamai, "Bilevel and multilevel programming: A bibliography review," *J. Global Optimization*, vol. 5, pp. 291–306, 1994.
- [18] B. Hajek and G. Gopal, "Do greedy autonomous systems make for a sensible internet?" in *Conference on Stochastic Networks, Stanford University, CA.*, 2002.
- [19] T. Basar and G. Olsder, *Dynamic Noncooperative Game Theory*. New York: Academic Press, 1998.
- [20] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33(3), pp. 520–534, 1965.
- [21] R. Johari and J. N. Tsitsiklis, "Efficiency loss in a network resource allocation game," *Mathematics of Operations Research*, vol. 29(3), 2004.
- [22] R. Maheswaran and T. Basar, "Social welfare of selfish agents: Motivating efficiency for divisible resources." in *43rd IEEE Conference on Decision and Control*, 2004.
- [23] H. Moulin, "An efficient and almost budget-balanced cost sharing method," *Games and Economic Behavior (submitted)*, August 2007.
- [24] S. Saroiu, P. Gummadi, and S. Gribble, "A measurement study of peer-to-peer file sharing systems," in *Proc. Multimedia Computing and Networking*, 2002.
- [25] μ Torrent 1.7.2, "<http://www.utorrent.com/>," 2007.
- [26] M. Jackson and A. Wolinsky, "A strategic model of economic and social networks," *J. Economic Theory*, vol. 71, no. 1, pp. 44–74, 1996.
- [27] S. Goyal and S. Joshi, "Unequal connections," *Forthcoming in International Journal of Game Theory*, 2006.
- [28] M. O. Jackson, "A survey of models of network formation: Stability and efficiency," in *Group Formation in Economics: Networks, Clubs and Coalitions*. New York: Cambridge University Press, 2004.
- [29] M. Jackson and A. Watts, "The existence of pairwise stable networks," *Seoul Journal of Economics*, vol. 14(3), 2001.
- [30] P. Dubey, "Inefficiency of nash equilibria," *Math. Oper. Res.*, vol. 11, 1–8, 1986.
- [31] C. Papadimitriou, "Algorithms, games, and the internet." in *Proc. 33rd Ann. ACM Sympos. Theory Comput.*, 2001.
- [32] T. Roughgarden and E. Tardos., "How bad is selfish routing?" *J. ACM*, vol. 49(2), 2002.
- [33] R. Ma, S. Lee, J. Lui, and D. Yau, "A game theoretic approach to provide incentive and service differentiation in p2p networks," in *ACM SIGMETRICS Performance Evaluation Review*, 2004.
- [34] G. Smaragdakis, N. Laoutaris, P. Michiardi, A. Bestavros, J. Byers, and M. Roussopoulos, "Swarming on optimized graphs for n-way broadcast," in *IEEE INFOCOM 2008*.
- [35] A. Akella, R. Karp, C. Papadimitrou, S. Seshan, and S. Shenker, "Selfish behavior and stability of the internet: A game-theoretic analysis of tcp," in *ACM SIGCOMM 2002*.
- [36] F. Kelly and T. Voice, "Stability of end-to-end algorithms for joint routing and rate control," *Computer Communication Review*, vol. 35:2, pp. 5–12, 2005.
- [37] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Multipath tcp: A joint congestion control and routing scheme to exploit path diversity on the internet," *IEEE/ACM Trans. Networking*, 2007.

He also received a M.Sc. degree from Purdue University West Lafayette, IN, USA. He is currently an assistant professor in the Math and Computer Science Department at Suffolk University in Boston, MA, USA. His research interest is on computer networking.



Giovanni Neglia Giovanni Neglia received a Master degree in electronic engineering and a Ph.D. degree in telecommunications from the University of Palermo, Italy, respectively in 2001 and in 2005. In 2005 he was a research scholar at the University of Massachusetts Amherst, visiting the Computer Networks Research Group. From June 2005 he is a postdoc at the University of Palermo and from September 2006 an external scientific advisor in the Maestro team at INRIA research laboratory, Sophia Antipolis, France. His research interests include statistical models for IP data traffic, admission control mechanisms, TCP models, peer-to-peer networks, epidemic routing in intermittently connected networks, biologically inspired protocols and algorithms, game theory applications to computer networks.

He currently serves as Editor-in-Chief of *IEEE/ACM Transactions on Networking*, and on the editorial boards of *Journal of the ACM* and *IEEE Journal on Selected Areas in Communications* and has previously served on numerous other editorial boards. He was Program Co-chair of the joint ACM SIGMETRICS and PERFORMANCE '92 conference and the Performance 2002 conference. He is a member of ACM and ORSA.

Don Towsley holds a B.A. in Physics (1971) and a Ph.D. in Computer Science (1975) from University of Texas. He is currently a Distinguished Professor at the University of Massachusetts in the Department of Computer Science. He has held visiting positions at IBM T.J. Watson Research Center, Yorktown Heights, NY; Laboratoire MASI, Paris, France; INRIA, Sophia-Antipolis, France; AT&T Labs - Research, Florham Park, NJ; and Microsoft Research Lab, Cambridge, UK. His research interests include networks and performance evaluation.

He has received the 2007 IEEE Koji Kobayashi Award, the 2007 ACM SIGMETRICS Achievement Award, the 1998 IEEE Communications Society William Bennett Best Paper Award, and numerous conference/workshop best paper awards. Last, he has been elected Fellow of both the ACM and IEEE.



Giuseppe Lo Presti received a Master degree and a Ph.D. in computer science engineering from the University of Palermo, Italy, respectively in 2000 and in 2005. He joined CERN, Switzerland, as Ph.D. student in 2004, for the CMS experiment collaboration, and as postdoc in 2005, for the IT department. From 2008 he is a CERN staff member in the IT department. His research interests include peer-to-peer networking, computer networks management, grid computing, high energy physics computing, large scale hierarchical storage management systems, game theory applications to computer networks.

He also received a M.Sc. degree from Purdue University West Lafayette, IN, USA. He is currently an assistant professor in the Math and Computer Science Department at Suffolk University in Boston, MA, USA. His research interest is on computer networking.

Honggang Zhang holds a Ph.D. in Computer Science from University of Massachusetts Amherst. He received his B.Sc. degree from the Central South University of China, and his M.Sc. degree from Tianjin University of China.