

Content Pricing in Peer-to-Peer Networks

Jaeok Park

University of California, Los Angeles

Mihaela van der Schaar

University of California, Los Angeles

Abstract

We provide a game theoretic model of content production and sharing in a peer-to-peer (P2P) network. We characterize two benchmark outcomes: Nash equilibrium (NE) without any incentive scheme and social optimum. We show that the P2P network is not utilized at an NE outcome, whereas social optimum in general requires the utilization of the P2P network. In order to obtain a socially optimal (SO) outcome among self-interested peers, we introduce a pricing scheme where downloading peers compensate uploading peers for content provision. For any SO outcome, we can find a pricing scheme with link-dependent linear prices that achieves the SO outcome as an NE. We illustrate our results with several examples. Our illustration shows that the structures of social optimum and optimal prices vary depending on the characteristics of peers such as cost parameters and connectivity.

1 Introduction

In today's Internet, we are witnessing the emergence of user-generated content in the form of photos, videos, news, customer reviews, and so forth. Peer-to-peer (P2P) networks are able to offer a useful platform for sharing user-generated content, because P2P networks are self-organizing, distributed, inexpensive, scalable, and robust. However, as evidenced in experimental studies [1], [2], the free-riding phenomenon prevails in P2P networks, which hinders the effective utilization of P2P networks. In this paper, we analyze a scenario where peers produce content and share produced content in a P2P network. We investigate how content pricing can be used to overcome the free-riding problem and achieve a socially optimal outcome, based on the principles of economics.

Recent work has studied pricing schemes in P2P networks using economics models [3]–[5]. [3] constructs a game theoretic model and proposes a micro-payment mechanism to provide an incentive for sharing. [4] compares different pricing schemes and their informational requirements in the context of a simple file-sharing game. [5] investigates the problem of selecting multiple server peers given the prices of service and a budget constraint. The models of [3]–[5], however, capture only a partial picture of the content production and sharing scenario we are interested in. [3] allows only three sharing levels while production levels are fixed. [4] considers only sharing levels while ignoring production and transfer decisions. [5] does not explain how the prices and the budget are determined.

In our recent work [6], we have formulated a game theoretic model of content production and sharing, where peers make production, sharing, and download decisions in order. We have shown that the P2P network is not

utilized without an incentive scheme whereas social optimum requires all produced content to be shared. We have also shown that an appropriately chosen price of content induces peers to operate at social optimum in their self-interest. However, we have imposed several simplifying assumptions in [6] for analytic convenience, and in this paper, we generalize the model of [6]. In particular, we allow general network connectivity, heterogeneous utility and production cost functions across peers, convex production cost functions, and link-dependent download and upload costs. We show that the main result of [6] continues to hold in the generalized model: There exists a discrepancy between Nash equilibrium and social optimum, and this discrepancy can be eliminated by introducing a pricing scheme. Moreover, by allowing general network connectivity and heterogeneous peers, we obtain richer structures of social optimum and optimal prices.

We briefly mention two related papers [7], [8]. [7] analyzes a price-based exchange model for content distribution. In [7] peers choose their download and upload rates subject to upload capacity and budget constraints, whereas in our model peers choose production, sharing, and download levels while incurring cost of production, upload, and download. [8] studies the problem of information acquisition and link formation, where information acquisition can be interpreted as content production. In [8] information flows freely once a link is established, whereas in our model connectivity is fixed and the transfer of content incurs cost to both uploading and downloading peers.

The rest of this paper is organized as follows. In Section 2, we describe our game theoretic model of content production and sharing. In Section 3, we characterize two benchmark outcomes: Nash equilibrium and social

optimum. In Section 4, we introduce a pricing scheme to show that any socially optimal outcome can be achieved as a Nash equilibrium with appropriately chosen link-dependent linear prices. In Section 5, we consider several examples to illustrate the results obtained in Sections 3 and 4. We conclude the paper in Section 6.

2 Model

We consider a P2P network consisting of N peers. Peers produce content using their own production technologies, and distribute produced content using the P2P network. Let $\mathcal{N} \triangleq \{1, \dots, N\}$ be the set of peers in the P2P network. A peer can download content from a subset of other peers, while it can upload content to a (potentially different) subset of other peers. We use $D(i)$ and $U(i)$ to denote the set of peers that peer i can download from and upload to, respectively, where $D(i), U(i) \subset \mathcal{N} \setminus \{i\}$ for each $i \in \mathcal{N}$. Thus, the connectivity topology of the P2P network can be represented by $\{D(i), U(i)\}_{i \in \mathcal{N}}$. Note that $j \in D(i)$ is equivalent to $i \in U(j)$ by definition.

As in [6], we model content production and sharing in the P2P network as a sequential game consisting of three stages, which is called the content production and sharing (CPS) game.

- *Stage One (Production)*: Each peer determines its level of production. $x_i \in \mathbb{R}_+$ represents the amount of content produced by peer i and is known only to peer i .
- *Stage Two (Sharing)*: Each peer specifies its level of sharing. $y_i \in [0, x_i]$ represents the amount of content that peer i makes available to other peers. Peer i observes $(y_j)_{j \in D(i)}$ at the end of stage two, for example, through a tracker.
- *Stage Three (Transfer)*: Each peer determines the amounts of content that it downloads from other peers. Peer i serves all the requests it receives from any other peer in $U(i)$ up to y_i . $z_{ij} \in [0, y_j]$ represents the amount of content that peer i downloads from peer $j \in D(i)$, or equivalently peer j uploads to peer i .

An allocation of the CPS game is represented by $(\mathbf{x}, \mathbf{y}, \mathbf{Z})$, where $\mathbf{x} \triangleq (x_1, \dots, x_N)$, $\mathbf{y} \triangleq (y_1, \dots, y_N)$, $\mathbf{z}_i \triangleq (z_{ij})_{j \in D(i)}$, for each $i \in \mathcal{N}$, and $\mathbf{Z} \triangleq (\mathbf{z}_1, \dots, \mathbf{z}_N)$. An allocation $(\mathbf{x}, \mathbf{y}, \mathbf{Z})$ is feasible if $x_i \geq 0$, $0 \leq y_i \leq x_i$, and $0 \leq z_{ij} \leq y_j$ for all $j \in D(i)$, for all $i \in \mathcal{N}$.

In stage one, peer i incurs the production cost $k_i(x_i)$ by producing the amount of content x_i , where $k_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the production cost function of peer i . We assume that k_i is continuously differentiable and convex,

for each $i \in \mathcal{N}$. Thus, the content production technology of each peer exhibits decreasing returns to scale. In stage two, there is no direct cost of announcing the level of sharing. In stage three, transferring the amount of content z_{ij} from peer j to peer $i \in U(j)$ induces a cost of $\delta_{ij}z_{ij}$ to peer i (the downloader) and $\sigma_{ij}z_{ij}$ to peer j (the uploader), where $\delta_{ij} > 0$ and $\sigma_{ij} > 0$ are the marginal costs of download and upload, respectively, from peer j to peer i . The cost of download and upload can be considered as cost due to bandwidth usage and energy consumption,¹ and we allow link-dependent download and upload rates. After the transfer of content, peer i consumes content it produces and downloads, (x_i, \mathbf{z}_i) . The utility that peer i receives from consuming content (x_i, \mathbf{z}_i) is measured by $f_i(x_i, \mathbf{z}_i)$, where $f_i : \mathbb{R}_+^{|D(i)|+1} \rightarrow \mathbb{R}_+$ is the utility function of peer i . We assume that f_i is continuously differentiable and concave, for each $i \in \mathcal{N}$. Note that each peer can evaluate content produced by different peers in a different way. The payoff function of peer i in the CPS game can be expressed as

$$\begin{aligned} v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}) &= f_i(x_i, \mathbf{z}_i) - k_i(x_i) - \sum_{j \in D(i)} \delta_{ij}z_{ij} - \sum_{j \in U(i)} \sigma_{ji}z_{ji}. \end{aligned}$$

As an example of the P2P network model formulated above, consider a P2P network where peers need financial data (e.g., earnings of companies, gross domestic products, and interest rates) in order to make forecasts based on which they make investment decisions (e.g., trade stocks and bonds). To obtain financial data, peers can either collect data by themselves or download data shared by other peers. Financial data allow peers to make more informed decisions, and this benefit is captured by the utility functions of the peers.

3 Nash Equilibrium and Social Optimum

We first analyze the non-cooperative outcome of the CPS game. Non-cooperative peers choose their strategies to maximize their own utilities given the strategies of other peers. A strategy for peer i in the CPS game is its complete contingent plan over the three stages, which can be represented by $(x_i, y_i(x_i), \mathbf{z}_i(x_i, y_i, (y_j)_{j \in D(i)}))$. A stage-one strategy for peer i is represented by $x_i \in \mathbb{R}_+$, a stage-two strategy by a function $y_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $y_i(x_i) \leq x_i$ for all $x_i \in \mathbb{R}_+$, and a stage-three strategy by a function $\mathbf{z}_i : \mathcal{I}_i \rightarrow \mathbb{R}_+^{|D(i)|}$ such that $z_{ij}(x_i, y_i, (y_j)_{j \in D(i)}) \leq y_j$ for all $j \in D(i)$, where $\mathcal{I}_i \triangleq \{(x_i, y_i, (y_j)_{j \in D(i)}) | x_i \in \mathbb{R}_+, y_i \in [0, x_i], y_j \in \mathbb{R}_+, \forall j \in D(i)\}$ is the set of all possible information sets for peer i at the beginning of stage three. Nash equilibrium (NE) of the CPS game is defined as a strategy

profile such that no peer can improve its payoff by a unilateral deviation. The play on the equilibrium path (i.e., the realized allocation) at an NE is called an NE outcome of the CPS game. The following proposition characterizes the NE of the CPS game.

Proposition 1. *Suppose that, for each $i \in \mathcal{N}$, a solution to $\max_{x \geq 0} \{f_i(x, 0) - k_i(x)\}$ exists, and denote it as x_i^e . An NE outcome of the CPS game has $x_i = x_i^e$ and $z_{ij} = 0$ for all $j \in D(i)$, for all $i \in \mathcal{N}$.*

Proof. Let $(\mathbf{x}, \mathbf{y}, \mathbf{Z})$ be an NE outcome of the CPS game. Suppose that $z_{ij} > 0$ for some $i \in \mathcal{N}$ and $j \in D(i)$. This implies that $y_j > 0$, and peer j can increase its payoff by at least $\sigma_{ij}z_{ij}$ by deviating to $y_j = 0$, contradicting NE. Therefore, $z_{ij} = 0$ for all $j \in D(i)$, for all $i \in \mathcal{N}$ at any NE outcome. Given that there is no transfer of content at NE, only $x_i = x_i^e$ satisfies the requirement of NE. \square

Peer i can share a positive amount of content at an NE outcome only in a trivial case where no peer demands content produced by peer i . Proposition 1 shows that lack of compensation for upload leads to the collapse of the P2P network. Without an incentive scheme, the free-riding incentive of peers dominates, which prevents the utilization of the P2P network.

We now turn to the socially optimal (SO) outcome of the CPS game. As in [8], we measure social welfare by the sum of the payoffs of peers, $\sum_{i=1}^N v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z})$, and define an SO allocation as an allocation that maximizes social welfare among feasible allocations. The problem of maximizing social welfare among feasible allocations can be written as

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{y}, \mathbf{Z}} \sum_{i \in \mathcal{N}} \left\{ f_i(x_i, \mathbf{z}_i) - k_i(x_i) \right. \\ \left. - \sum_{j \in D(i)} \delta_{ij} z_{ij} - \sum_{j \in U(i)} \sigma_{ji} z_{ji} \right\} \\ \text{subject to } x_i \geq 0, 0 \leq y_i \leq x_i, 0 \leq z_{ij} \leq y_j \\ \text{for all } j \in D(i), \text{ for all } i \in \mathcal{N}. \end{aligned} \quad (1)$$

Using Karush-Kuhn-Tucker (KKT) conditions, we can characterize SO allocations.

Proposition 2. *An allocation $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ is SO if and only if it is feasible and there exist constants μ_i and λ_{ij} for $i \in \mathcal{N}$ and $j \in D(i)$ such that*

$$\frac{\partial f_i(x_i^*, \mathbf{z}_i^*)}{\partial x_i} - \frac{dk_i(x_i^*)}{dx_i} + \mu_i \leq 0, \quad \text{with equality if } x_i^* > 0, \quad (2)$$

$$\sum_{j \in D(i)} \lambda_{ji} - \mu_i \leq 0, \quad \text{with equality if } y_i^* > 0, \quad (3)$$

$$\frac{\partial f_i(x_i^*, \mathbf{z}_i^*)}{\partial z_{ij}} - \delta_{ij} - \sigma_{ij} - \lambda_{ij} \leq 0, \quad \text{with equality if } z_{ij}^* > 0, \quad (4)$$

$$\mu_i \geq 0, \quad \text{with equality if } y_i^* < x_i^*, \quad (5)$$

$$\lambda_{ij} \geq 0, \quad \text{with equality if } z_{ij}^* < y_j^*, \quad (6)$$

for all $j \in D(i)$, for all $i \in \mathcal{N}$.

Proof. Since the objective function of the maximization problem in (1) is concave and the constraints are linear, KKT conditions are both necessary and sufficient. \square

Unlike at an NE outcome, the P2P network may be utilized at an SO allocation. For example, suppose that the utility function of peer i is given by $f_i(x_i, \mathbf{z}_i) = (\alpha_i x_i^\rho + \sum_{j \in D(i)} \alpha_{ij} z_{ij}^\rho)^{1/\rho}$ for some $\alpha_i, \alpha_{ij} > 0$, $j \in D(i)$, and $\rho < 1$. Then (4) cannot be satisfied when $z_{ij} = 0$, and thus we have $z_{ij} > 0$ for all $j \in D(i)$ at an SO allocation.

4 Content Pricing

In order to provide incentives for sharing, we introduce a pricing scheme in the CPS game. We consider a class of pricing schemes under which a peer downloading content makes a payment to the uploading peer. We allow link-dependent prices and use $p_{ij} \in \mathbb{R}_+$ to denote the unit price of content that peer j provides to peer i . In other words, peers i pays peer j the price p_{ij} when peer i downloads a unit of content from peer j . A pricing scheme can be represented by prices $(p_{ij})_{i \in \mathcal{N}, j \in D(i)}$. In the CPS game with a pricing scheme $\mathbf{p} \triangleq (p_{ij})_{i \in \mathcal{N}, j \in D(i)}$, peers take actions over the three stages as in the CPS game described in Section 2 while paying and receiving prices depending on their stage-three choices. Prices are set by the P2P service provider, and peers take prices as given. With pricing scheme \mathbf{p} , peer i pays $\sum_{j \in D(i)} p_{ij} z_{ij}$ to other peers while receiving $\sum_{j \in U(i)} p_{ji} z_{ji}$ from other peers. Thus, the payoff function of peer i in the CPS game with pricing scheme \mathbf{p} is given by

$$\begin{aligned} \pi_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}; \mathbf{p}) \\ = v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}) - \sum_{j \in D(i)} p_{ij} z_{ij} + \sum_{j \in U(i)} p_{ji} z_{ji} \\ = f_i(x_i, \mathbf{z}_i) - k_i(x_i) - \\ \sum_{j \in D(i)} (p_{ij} + \delta_{ij}) z_{ij} + \sum_{j \in U(i)} (p_{ji} - \sigma_{ji}) z_{ji}. \end{aligned}$$

In effect, a pricing scheme increases the cost of download while decreasing the cost of upload. If $p_{ij} > \sigma_{ij}$, peer j receives a net benefit from uploading content to peer i , which provides peer j with an incentive for sharing.

Note that SO allocations are not affected by the introduction of a pricing scheme, because payments are transferred between peers. The following proposition shows that there exists an optimal pricing scheme in the sense that SO allocations are achieved as non-cooperative equilibria of the CPS game with the pricing scheme.

Proposition 3. *Let $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ be an SO allocation and $(\lambda_{ij})_{i \in \mathcal{N}, j \in D(i)}$ be associated constants satisfying the KKT conditions (2)–(6). Then $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ is an NE outcome of the CPS game with pricing scheme $\mathbf{p}^* = (p_{ij}^*)_{i \in \mathcal{N}, j \in D(i)}$, where $p_{ij}^* = \lambda_{ij} + \sigma_{ij}$ for $i \in \mathcal{N}$ and $j \in D(i)$.*

Proof. Choose an SO allocation $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$, and find associated constants μ_i and λ_{ij} for $i \in \mathcal{N}$ and $j \in D(i)$ given in Proposition 2. Consider a strategy of peer i , $(\tilde{x}_i, \tilde{y}_i(x_i), \tilde{z}_i(x_i, y_i, (y_j)_{j \in D(i)}))$, such that $\tilde{x}_i = x_i^*$, $\tilde{y}_i(x_i) = x_i$ if $\mu_i > 0$, $\tilde{y}_i(x_i) = \min\{x_i, y_i^*\}$ if $\mu_i = 0$, and $\tilde{z}_{ij}(x_i, y_i, (y_j)_{j \in D(i)}) = \min\{y_j, z_{ij}^*\}$, for all $j \in D(i)$, for all $i \in \mathcal{N}$.

Consider the maximization problem of peer i in the CPS game with pricing scheme \mathbf{p}^* , given the strategies of other peers:

$$\begin{aligned} \max_{x_i, y_i, z_i} & \left\{ f_i(x_i, \mathbf{z}_i) - k_i(x_i) \right. \\ & \left. - \sum_{j \in D(i)} (\lambda_{ij} + \sigma_{ij} + \delta_{ij}) z_{ij} + \sum_{j \in U(i)} \lambda_{ji} \tilde{z}_{ji} \right\} \\ \text{subject to } & x_i \geq 0, 0 \leq y_i \leq x_i, 0 \leq z_{ij} \leq y_j^* \\ & \text{for all } j \in D(i). \end{aligned} \quad (7)$$

Comparing the KKT conditions for (1) and (7), we can show that (x_i^*, y_i^*, z_i^*) solves (7), which proves that $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ is an NE outcome. \square

In the expression of the optimal prices, $p_{ij}^* = \lambda_{ij} + \sigma_{ij}$, we can see that peer i compensates peer j for the upload cost, σ_{ij} , as well as the shadow price, λ_{ij} , of content supplied from peer j to peer i . Proposition 3 resembles the second fundamental theorem of welfare economics [9] in that it shows that any SO allocation can be achieved in a decentralized manner through prices. However, our model is different from the general equilibrium model in that we consider networked interactions [10] where the set of feasible consumption bundles for a peer depends on the sharing levels of peers from which it can download.

5 Illustrative Examples

In this section, we illustrate the results in Sections 3 and 4 by imposing the following assumptions.




Figure 1: P2P connectivity topologies: (a) fully connected topology, (b) star topology, (c) ring topology, and (d) line topology.

- (Perfectly substitutable content) The utility from consumption depends only on the total amount of content. In other words, for each peer i , there exists a function $g_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $f_i(x_i, \mathbf{z}_i) = g_i(x_i + \sum_{j \in D(i)} z_{ij})$. We assume that g_i is twice continuously differentiable and satisfies $g_i(0) = 0$, $g_i' > 0$, $g_i'' < 0$ on \mathbb{R}_{++} , and $\lim_{x \rightarrow \infty} g_i'(x) = 0$ for all $i \in \mathcal{N}$.
- (Linear production cost) The production cost is linear in the amount of content produced. In other words, for each peer i , there exists a constant $\kappa_i > 0$ such that $k_i(x_i) = \kappa_i x_i$. We assume that $\kappa_i < g_i'(0)$, where $g_i'(0)$ is the right derivative of g_i at 0, for all $i \in \mathcal{N}$ so that each peer consumes a positive amount of content at an SO allocation.
- (Socially valuable P2P network) Obtaining a unit of content through the P2P network costs less to peers than producing it privately. In other words, $\delta_{ij} + \sigma_{ij} < \kappa_i$ for all $i \in \mathcal{N}$ and $j \in D(i)$.

We define g as the average benefit function, $g \triangleq (\sum_{i=1}^N g_i)/N$. By the assumptions on g_i , for every $\alpha \in (0, g'(0))$, there exists a unique $\hat{x}_\alpha > 0$ that satisfies $g'(\hat{x}_\alpha) = \alpha$. We define $g^*(\alpha) = \sup_{x \geq 0} \{g(x) - \alpha x\}$ for $\alpha \in \mathbb{R}$ as the conjugate of g [11].

5.1 Fully Connected Networks with Heterogeneous Peers

We first consider a fully connected P2P network, where each peer can download from and upload to every other peer in the network, as illustrated in Figure 1(a). In this case, we have $D(i) = U(i) = \mathcal{N} \setminus \{i\}$ for all $i \in \mathcal{N}$.

Let $\beta_i \triangleq [\kappa_i + \sum_{j \in D(i)} (\delta_{ji} + \sigma_{ji})]/N$, for $i \in \mathcal{N}$, and let $\beta \triangleq \min\{\beta_1, \dots, \beta_N\}$. Note that β_i is the per capita cost of peer i producing one unit of content and supplying it to every other peer to which peer i can upload, and we call it the cost parameter of peer i . It is SO to have only the most “cost-efficient” peers (i.e., peers with the smallest cost parameter in the network) produce a positive amount, where the total amount of production is given by \hat{x}_β . For simplicity, suppose that the most cost-efficient peer is unique, and without loss of generality index the peer as peer 1. Then at the SO allocation, peer 1 produces the amount \hat{x}_β and uploads it to every other peer. As a result, each peer consumes the amount \hat{x}_β , and the maximum social welfare is $Ng^*(\beta)$. The optimal pricing scheme is given by $(p_{ij}^*)_{i \in \mathcal{N}, j \in D(i)}$, where $p_{ij}^* = g'_i(\hat{x}_\beta) - \delta_{ij}$. The payoff of peer i under the optimal pricing scheme is $g_i(\hat{x}_\beta) - g'_i(\hat{x}_\beta)\hat{x}_\beta$ for peer $i \neq 1$ and $g_1(\hat{x}_\beta) - [\kappa_1 - \sum_{j \neq 1} (g'_j(\hat{x}_\beta) - \delta_{j1} - \sigma_{j1})]\hat{x}_\beta$ for peer 1. In a fully connected P2P network, the role of uploading peers (servers) and downloading peers (clients) at social optimum is determined by the cost parameters of peers.

5.2 Networks of General Connectivity Topology with Homogeneous Peers

We consider homogeneous peers in the sense that the benefit function, g_i , and the cost parameters, κ_i , δ_{ij} , and σ_{ij} , do not depend on $i \in \mathcal{N}$ and $j \in D(i)$. We denote the common respective function and parameters by g , κ , δ , and σ . Since the results about social optimum and optimal pricing schemes depend highly on the specific topology under consideration, we illustrate the results with three stylized network topologies: a star topology, a ring topology, and a line topology, as shown in Figure 1(b)–(d).

5.2.1 Star Topology

Assume $N \geq 3$, and without loss of generality let peer 1 be the center of a star topology. The cost parameters are given by $\beta_1 = [\kappa + (N-1)(\delta + \sigma)]/N = \beta$ and $\beta_j = (\kappa + \delta + \sigma)/2$ for $j \neq 1$. Since peer 1 is more connected than other peers, it is more cost-efficient (i.e., $\beta_1 < \beta_j$ for all $j \neq 1$). Hence, only peer 1 produces a positive amount of content \hat{x}_β and uploads it to every other peer at the SO allocation.² The consumption of each peer is \hat{x}_β , while the maximum social welfare is $Ng^*(\beta)$. The optimal price is given by $p^* = [\kappa + (N-1)\sigma - \delta]/N$, independent of the link, which yields payoff $g^*(\beta)$ to every peer. In a star P2P network with homogeneous peers, the positions of peers determine their roles as servers and clients, while a uniform price suffices to support the SO allocation at NE.

As an extension of our model, suppose that there are N heterogeneous peers and consider the problem of building a P2P network connecting the N peers when it is costly to establish a link. Our analysis implies that forming a star network with the most cost-efficient peer as the center requires the minimum number of links (i.e., $N-1$ links) among networks that realize the maximum social welfare achievable with the N peers.

5.2.2 Ring Topology

In a ring topology, every peer is connected to two neighboring peers, and thus peers have the same cost parameter $\tilde{\beta} \triangleq [\kappa + 2(\delta + \sigma)]/3$. When each peer produces x and shares its content with its neighboring peers, social welfare is given by $N\{g(3x) - [\kappa + 2(\delta + \sigma)]x\}$. Thus, each peer produces the amount $\hat{x}_{\tilde{\beta}}/3$ while consuming $\hat{x}_{\tilde{\beta}}$ at the SO allocation, which achieves the maximum social welfare $Ng^*(\tilde{\beta})$. The optimal price is given by $p^* = (\kappa + 2\sigma - \delta)/3$, yielding payoff $g^*(\tilde{\beta})$ to every peer. In a ring P2P network with homogeneous peers, peers are in a symmetric position, which leads them to play the roles as a server and a client at the same time. Also, the SO amounts of production and consumption do not depend on the number of peers in a ring network, and thus the maximum per capita social welfare is independent of N . On the contrary, in a fully connected or star network with homogeneous peers, the cost parameter β decreases as there are more peers in the network, and thus the SO amounts of production and consumption and the maximum per capita social welfare increase with N .

5.2.3 Line Topology

Index peers as peer 1 through N from the left to the right of a line topology. Then the cost parameters of peers are given by $\beta_1 = \beta_N = (\kappa + \delta + \sigma)/2$ and $\beta_i = \tilde{\beta}$ for all $i \neq 1, N$. Since peers in the end (peers 1 and N) are less cost-efficient than peers in the middle (peers 2 through $N-1$), it is not SO to have peers in the end produce a positive amount of content (assuming $N \geq 3$). Given the number of peers N , we can find SO production levels by equating the marginal social benefit and the marginal social cost. SO allocations for N between 3 and 9 are shown in Figure 2. The optimal pricing scheme has peer-dependent prices, where the price that peer i pays to its neighboring peers is given by $p_i^* = g'(c_i^*) - \delta$, where c_i^* is the consumption of peer i at the SO allocation. For example, when $N = 7$, $p_i^* = g'(2x^*) - \delta$, for $i = 1, 4, 7$, and $p_i^* = g'(3x^*) - \delta$, for $i = 2, 3, 5, 6$, where x^* is determined as shown in Figure 2. When N is a multiple of 3, $c_i^* = \hat{x}_{\tilde{\beta}}$ and thus $p_i^* = (\kappa + 2\sigma - \delta)/3$ for all $i \in \mathcal{N}$. Note that when N is a multiple of 3, the SO operation of the entire P2P network is the same as that

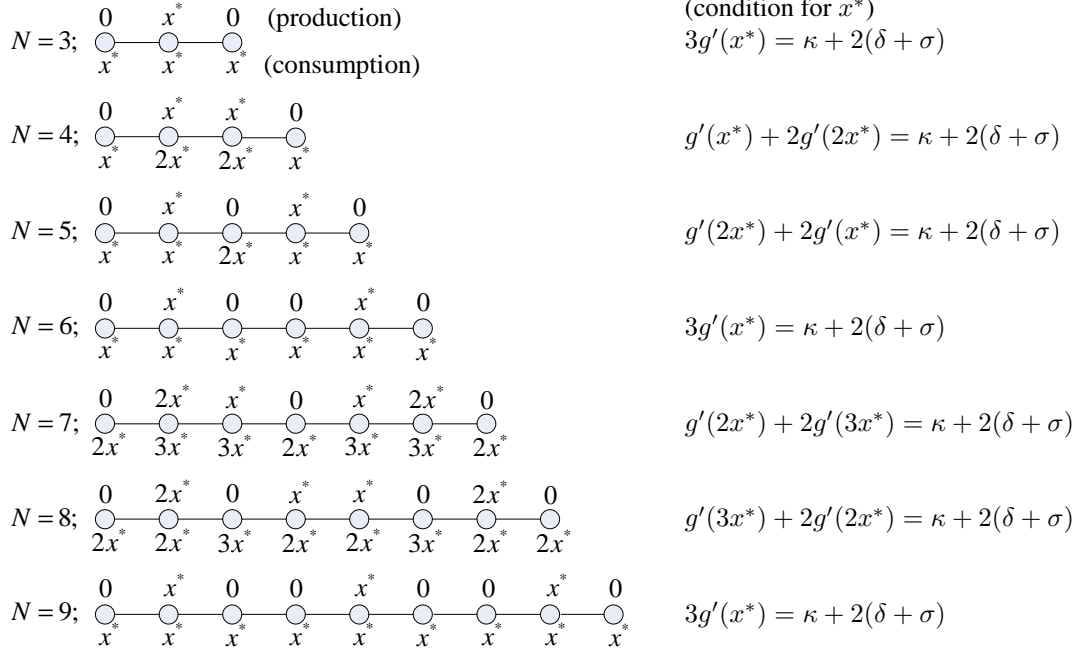


Figure 2: SO allocations for line topology.

of $N/3$ sub-networks, each of which consists of three connected peers. Therefore, the consumption level at the SO allocation and the optimal price do not depend on N . In a line P2P network with homogeneous peers, a peer can become a server, a client, or a server and client at the same time, depending on its position in the topology.

6 Conclusion

We have generalized the model of our previous work [6] and analyzed Nash equilibrium, social optimum, and pricing schemes. Although analytically more complicated, the generalized model offers many interesting insights that could not be captured under the simplifying assumptions in [6]. For example, the cost parameter as well as the position of a peer in connectivity topology determines its role as a server or a client (or both). The CPS game with a pricing scheme can be extended in several directions. First, we can consider a scenario where uploading peers set the prices they receive to maximize their payoffs. Second, we can apply mechanism design in a scenario where prices are determined based on the report of peers on their utility and cost functions. Lastly, we can analyze the problem of link formation by self-interested peers.

References

- [1] E. Adar and B. A. Huberman, "Free riding on Gnutella," *First Monday*, vol. 5, no. 10, 2000.

- [2] S. Saroiu, P. K. Gummadi, and S. D. Gribble, "A measurement study of peer-to-peer file sharing systems," in *Proc. Multimedia Computing and Networking*, 2002.
- [3] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," in *Proc. 2nd Int. Workshop Electronic Commerce (WELCOM)*, 2001, pp. 75–87.
- [4] P. Antoniadis, C. Courcoubetis, and R. Mason, "Comparing economic incentives in peer-to-peer networks," *Comput. Networks*, vol. 46, no. 1, pp. 133–146, Sep. 2004.
- [5] M. Adler, R. Kumar, K. Ross, D. Rubenstein, D. Turner, and D. Yao, "Optimal peer selection in a free-market peer-resource economy," in *Proc. 2nd Workshop Economics Peer-to-Peer Systems*, 2004.
- [6] J. Park and M. van der Schaar, "Pricing and incentives in peer-to-peer networks," in *Proc. INFOCOM*, 2010.
- [7] C. Aperijs, M. J. Freedman, and R. Johari, "Peer-assisted content distribution with prices," in *Proc. CoNext*, 2008.
- [8] A. Galeotti and S. Goyal, "The law of the few," *Amer. Econ. Review*, to be published.
- [9] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.
- [10] M. Jackson, *Social and Economic Networks*. Princeton, NJ: Princeton Univ. Press, 2008.
- [11] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

Notes

¹By assuming linear download and upload cost functions, we are considering a scenario where the download and upload capacity constraints of each node do not bind.

²This statement is true for P2P networks of any topology with homogeneous peers in which peer 1 is the only peer connected to every other peer.