

December 2004

On the Formation of Peer-to-Peer Networks: Self-Organized Sharing and Groups

Yung-Ming Li
University of Washington

Yong-Wah Tan
University of Washington

Prabuddha De
Purdue University

Follow this and additional works at: <http://aisel.aisnet.org/icis2004>

Recommended Citation

Li, Yung-Ming; Tan, Yong-Wah; and De, Prabuddha, "On the Formation of Peer-to-Peer Networks: Self-Organized Sharing and Groups" (2004). *ICIS 2004 Proceedings*. 39.
<http://aisel.aisnet.org/icis2004/39>

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ICIS 2004 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

ON THE FORMATION OF PEER-TO-PEER NETWORKS: SELF-ORGANIZED SHARING AND GROUPS

Yung-Ming Li

Business School
University of Washington
Seattle, WA U.S.A.
ymli@u.washington.edu

Yong Tan

Business School
University of Washington
Seattle, WA U.S.A.
ytan@u.washington.edu

Prabuddha De

Krannert School of Management
Purdue University
West Lafayette, IN U.S.A.
pde@purdue.edu

Abstract

In this paper, we investigate the formation of peer-to-peer (P2P) networks with rational participating agents (active peers). In the absence of a central planner, peers choose their own utility-maximizing strategies for coalition and peer formation. P2P networks evolve dynamically through the activities of interactions among individual nodes and group units. We propose a framework for multilevel formation dynamics, including an individual level (content sharing decision and group selection) and a group level (membership admission). The respective utilities of the individual node and the collective player are formulated as functions of operational performance metrics such as expected content availability, search delay, transmission delay, and download delay. We study the impacts of various system parameters on the emergence of self-organized P2P network configuration features such as free-riding level and group size. Furthermore, we investigate the stability and efficiency of P2P networks and propose internal transfer mechanisms that force stable networks to become efficient.

Keywords: Peer-to-peer networks, self-organization, stability, efficiency, incentive mechanism

Introduction

A peer-to-peer (P2P) network is a social network for pooling the resources—such as computing cycle, hard disk storage, network bandwidth, and content—at the numerous edge nodes. Each member participating in various dedicated P2P communities leverages the aggregated commons by resource exchanging and information sharing. Business applications of P2P networks include grid computing (such as SET@home), instant messaging (such as instant messengers of MSN, AOL, and Yahoo!), collaborating (such as Groove networks), and file sharing (such as KaZaA and Gnutella). The emerging popularity of P2P networks has attracted significant attention from public media, and venture capital companies, as well as academia.

P2P technologies are considered to be an efficient way to deliver content because they do not rely on central servers and, therefore, are more scalable. However, there are also drawbacks inherent in P2P networks because of their decentralized structure. First, since peer nodes can decide whether or not to share their contents, the content availability becomes uncertain. Low level of content sharing behaviors will not only reduce the content availability and variety, but may also experience serious congestion at the nodes that allow downloads. Second, the performance of P2P networks strongly depends on peer nodes being

interconnected. Therefore, it is important to examine the impacts of system parameters on network performance and consequently design incentive mechanisms that control P2P players' behaviors—for example, content sharing and request peering—so that a more efficient P2P network structure can emerge.

Various P2P structures differ by their search algorithms. A centralized P2P architecture, such as Napster, has the scale problem because of the difficulty in scaling the central directory server. Pure decentralized P2P architectures, such as Gnutella v0.4, while easily scalable because search is carried out among peer nodes, have to deal with excessive network traffic due to a decentralized broadcast-type search. To alleviate searching inefficiency, newer generations of P2P software, such as KaZaA and Gnutella v0.6, use a combination of centralized and purely decentralized network structures: the peer nodes are grouped and served by different super peers. Various groups are interconnected via super peers so that requests could be forwarded from one group to another, if necessary. In this study, we focus on the analysis of the formation behaviors (such as sharing and grouping decisions) based on this promising network structure involving super peers.

Much of the (technological) research on P2P networks assumes that the operations of P2P networks are centrally governed. The allocation of P2P participants—namely, individual users and the search intermediary—is determined so as to achieve global welfare under centralized control and cooperation. However, the coalition and peer formation of an *ad hoc* P2P network is expected to be incentive compatible. Users join P2P communities because of their self-interests, and the communities are the realization of mutual interests among members. Empirical evidence (Adar and Huberman 2000; Asvanund et al. 2004) shows a very high degree of free riding behaviors in P2P networks. This indicates that P2P network efficiency can be better improved from social and economic incentives, rather than technological mechanisms. In this paper, incentive compatibility is assumed to be the driving force for P2P network formation.

Each active peer node is modeled as an agent. Agents interact with each other as well as with the group (collective agent or delegate). There are also interactions among the delegates of various groups. Strategies are triggered between interacting agents to improve self-interests (utilities). In P2P file-sharing networks, the utility can be described by the value of content availability and costs of activities such as search, transmission, and processing. Evaluating expected utilities through information exchange, the individual and collective (group) units make decisions such as group selection, membership admission, and free-riding behavior. P2P networks are formed and evolve through the realization of local information-based, bottom-up, and self-organizing dynamics. We assume that the information of agent configuration is public knowledge, as its revelation can be achieved through exchanges.

In this paper, we adopt a game-theoretical approach to exploit the emergence of decentralized self-organizing networks. We are especially interested in the characteristics, such as free-riding behaviors and group size, of a stable (incentive-compatible) network structure. A numerical simulation is utilized to illustrate how the system parameters affect the results of P2P network formation. We further investigate the social welfare efficiency loss in the absence of central control, and propose an internal transfer mechanism that induces a stable P2P network to become efficient.

The remainder of the paper is organized as follows. The next section reviews the related literature on P2P networks. A formal description of our model is then presented, the system performance metrics are outlined, and the utility functions are proposed. We next discuss the formation dynamics and the resulting structure, and present some numerical results. The final section concludes our findings and presents future research directions.

Literature Review

There are a number of papers on the technical aspects of P2P networks from the computer science community. These papers focus mainly on developing efficient communication protocols, network topologies, and search algorithms (Ratnasamy et al. 2001; Stoica et al. 2001). Recently, a few researchers have started to explore social and economic aspects of P2P networks such as the free-riding phenomenon and incentive mechanism design (Golle et al. 2001). In addition, many reputation and trust systems are proposed to promote the incentive of cooperation without involving a pricing scheme.

There is an extensive literature on the formation of complexity systems from Physics, Biology, and Computer Science. Most of the group and network formation research are from Political Science, Social Science, and Economics. For example, Jackson and Wolinsky (1996) examine whether efficient (value-maximizing) social networks will form when self-interested individuals can choose to form or sever links. In the Business and Management area, for example, Axelrod et al. (1995) present a theory and an agent-based simulation model to predict the ways businesses develop alliances and sponsor technical standards.

Regarding the formation of P2P networks, Krishnan et al. (2002, 2004) propose a plausible model to analyze the existence of free-riding behaviors in P2P file-sharing networks. However, their framework, assuming a constant sharing cost in the absence of any query forward interconnection, does not explicitly discuss the impacts of system parameters on network structures. Asvanund et al. (2003) propose a content- and physical location-based model for club membership management to improve the performance and avoid network congestion. In their proposed protocol, the peers can determine which club to join, and the clubs manage their membership and determine to which clubs they should be connected. However, the simulation model does not investigate the influence of exogenous factors on the resulting network structure, nor does it address the tension between stability and efficiency of the self-formed P2P networks. Ledlie et al. (2002) develop a hierarchically grouped system that can self organize to overcome unreliability. Khambatti et al. (2002) use attribute-based clustering models to simulate how self-configuring communities are formed. Their simulation results demonstrate that community structures in a random network can be efficiently discovered based on the attribute and link information of peers. However, to the best of our knowledge, little attention has so far been given to quantitative models that exploit the underlying dynamics of P2P network formation and investigate how the system parameters affect the emergence of P2P network features such as sharing level and group size.

The Model

We consider a self-organizing P2P file-sharing network. The peers in the P2P network are categorized as regular peers and super peers. In this structure, only the super peer has the resource-sharing information such as content, bandwidth, and computing capacity of regular peers connected to that super peer. The super peer and a number of regular peers form a group. The super peer provides search service for all of the content requests from users in the same group, maintains up-to-date information on all resources available in the local group, and recommends a service node to each query based on content availability and the bandwidth between any two regular peer nodes (with minimum network transmission delay). Once this information is passed on to the request and service nodes, download occurs directly between these two nodes. If a request cannot be satisfied within the local group, the local super peer will forward the request to all the interconnected super peers in remote groups. The node with minimum transmission delay will be selected as the provision node from all of the nodes that have the content. In our model, we consider that each peer node is connected to only one super peer. A regular node will select the best group to join and decide whether or not it will share its contents for download so as to maximize its own utility. The delegate (super peer) of a group will also evaluate the possibility of admission of a new applicant to improve the group utility.

System Parameters

Since each peer node is a content consumer as well as a content provider, the dynamics of a P2P network are dependent on local peer parameters. The parameters, listed in Table 1, include content provision distribution, content request distribution, and bandwidth distribution (or transmission delay distribution). Given a Poisson request arrival, the search, provision, and receiving processes are described by time-sharing M/G/1/PS queues (Kleinrock 1976). The transmission delay between the content request node and the provision node is assumed to be an *i.i.d.* random variable drawn from a uniform distribution $U[0, \rho]$, where ρ is the upper bound of transmission delay.

Table 1. Model Parameters.

n_i	Number of active P2P users (peers) in group i
γ_i	Proportion of active P2P users who share contents in group i
p	Availability of content at a peer node
μ_{sp}	Service rate of super peer
μ_s	Service rate of sending content
μ_R	Service rate of receiving content
δ	Reduction of receiving rate when a peer allows uploading
ρ	Upper bound of transmission delay
λ	Content request rate by a peer

Performance Metrics

The performance metrics are established based on the benefit and cost of activities occurring during the process of content supply. The activities are request, search, download, and transmission. Hence, the system performance metrics include content availability at the requested peer node, search delay at super peers, provision (download) process delay, and transmission delay on the network.

Content availability. Content availability (or hit rate) is defined as the probability that the content requested can be found in the P2P network. It can be expressed as

$$H(n, \gamma) = 1 - (1 - p)^{n \cdot \gamma}$$

Here, p is the probability that a desired content can be found at the searched node. A more detailed model would require p to be content dependent as it indicates the popularity of that content, which is typically assumed to follow a Zipf-like distribution (Zipf 1929).

Search delay. In a super peer structure, all of the content requests will be forwarded to the group center (super peer). The search delay can be defined as the total system waiting time at the local or external super peer, given the network topology and the level of content availability. Since both free riders and contributors request contents, the search delay is associated with the group size and is not affected by free-riding behaviors. Using the underlying queueing model, we can express

$$S(n) = \frac{1}{\mu_{SP} - \Lambda_{SP}(n)}$$

where $\Lambda_{SP}(n)$ is the aggregated search demand.

Processing delay. The expected processing delays include upload processing delay and download processing delay. The download performance reduction is the primary disincentive for the P2P sharing behavior. Empirical evidence (Feldman et al. 2003) shows that when uploading a single file, the utilization will drop to 80 percent for an Ethernet node (download at 10 Mb/s, upload at 10 Mb/s) and 20 percent for a DSL node (download at 1.5Mb/s, upload at 128 Kb/s). We assume that each user has an upload process rate μ_s and a download process rate μ_r . If a peer allows upload, his downloading process rate will be degraded to $\delta \cdot \mu_r$, where δ is the degradation coefficient and $\delta \in [0,1]$. This parameter depends on the degree of asymmetry between the rates, μ_s and μ_r . Typically, δ increases with the variance of μ_s and μ_r . Explicitly, the download delay, D_{SH} (D_{NS}) for a sharing (not sharing) node, can be written as

$$D_{NS}(n, \gamma) = \frac{1}{\mu_s - \Lambda_s(n, \gamma)} + \frac{1}{\mu_r - \Lambda_r(n, \gamma)};$$

$$D_{SH}(n, \gamma) = \frac{1}{\mu_s - \Lambda_s(n, \gamma)} + \frac{1}{\delta \cdot \mu_r - \Lambda_r(n, \gamma)}.$$

The aggregated demand at the sending node is

$$\Lambda_s(n, \gamma) = n \cdot p \cdot \sum_{k=0}^{n \cdot \gamma - 1} \frac{1}{k+1} \cdot \binom{n \cdot \gamma - 1}{k} \cdot p^k \cdot (1-p)^{n \cdot \gamma - 1 - k} \cdot \lambda = \frac{H(n, \gamma)}{\gamma} \cdot \lambda;$$

and, at the receiving end, it is $\Lambda_r(n, \gamma) = H(n, \gamma) \cdot \lambda$.

Transmission delay. It is difficult to exactly estimate the effective bandwidth and the corresponding transmission latency of download activity. However, using network coordinate-mapping technologies, such as the global network positioning (GNP) approach (Ng and Zhang 2002), the coordinate-based positions of P2P networks can be used to approximately predict the Internet “transmission distance.” The transmission delay is estimated by comparing this transmission distance from each potential

provision node to the request node. The node with the minimum transmission delay to the request node will be selected as the provision node. As a result of order statistics, this delay is given by (Li et al. 2003)

$$T(n, \gamma) = \frac{1 - (1 - p)^{n \cdot \gamma} (1 + n \cdot \gamma \cdot p)}{(1 - (1 - p)^{n \cdot \gamma}) (n + 1) \cdot \gamma \cdot p} \cdot \rho \cdot$$

Utility Functions

The agent utility is evaluated based on the performance metrics described above. Agents wish to maximize their respective gains, while satisfying their demand as quickly as possible. The utility is the value of the expected content availability less the expected total latency when a request is satisfied, and is dynamically dependent on the content-sharing level and group size.

The total delay includes the search delay S at the super peer, download delays at the request node (receiving process delay D_R) and the provision node (sending process delay D_S), and transmission delay T in the network. We assume that the total delay cost is $C(S + D + T)$, where the cost function $C(\cdot)$ is assumed to be convex. Furthermore, we define the utility function as

$$U = V(H) - C(S + D + T).$$

Here, the value function $V(\cdot)$ is concave with content availability.

A regular peer acts as an individual agent, and a super peer serves as an agent for the group. They act according to individual and group utilities, respectively. Given an isolated group g_i , the expected utility for the individuals who share contents is

$$U_{SH}(g_i(n_i, \gamma_i)) = V(H(n_i, \gamma_i)) - C(S(n_i, \gamma_i) + D_{SH}(n_i, \gamma_i) + T(n_i, \gamma_i));$$

while the expected utility for the individuals who do not share contents is given by

$$U_{NS}(g_i(n_i, \gamma_i)) = V(H(n_i, \gamma_i)) - C(S(n_i, \gamma_i) + D_{NS}(n_i, \gamma_i) + T(n_i, \gamma_i)).$$

Hence, the average expected individual utility is

$$U_{AVG}(g_i(n_i, \gamma_i)) = \gamma_i \cdot U_{SH}(g_i(n_i, \gamma_i)) + (1 - \gamma_i) \cdot U_{NS}(g_i(n_i, \gamma_i)).$$

The overall group utility for group g_i can be expressed as

$$U_i(G) = n_i \cdot U_{AVG}(g_i(n_i, \gamma_i)).$$

Formation Dynamics

Based on the individual and group utility functions, multiple-level dynamics are used to simulate the formation of P2P networks and their corresponding structure complexity. The dynamics of the individual level (content sharing decision and group selection) and the group level (membership admission) are based on self interests as well as mutual interests.

Free-Riding Decision

It is obvious that a user has the incentive to be a free rider for better download performance. However, by doing so, the upload performance will degrade. The individual makes a sharing decision to balance the tension between upload and download delays, while also considering the marginal benefits of content hit ratio and transmission delay.

DYNAMIC 1 (SHARING DECISION). Given an isolated group configuration $g_i(n_i, \gamma_i)$

1. a free rider will choose to share his contents if and only if his incentive compatibility constraint (IC) is satisfied, that is,

$$g_i(n_i, \gamma_i) \rightarrow g_i(n_i, \gamma_i^+), \text{ if } U_{SH}(g_i(n_i, \gamma_i^+)) \geq U_{NS}(g_i(n_i, \gamma_i)),$$

$$\text{where } \gamma_i^+ = (n_i \cdot \gamma_i + 1) / n_i;$$

2. a contributor will choose not to share his contents if and only if his IC is satisfied, that is,

$$g_i(n_i, \gamma_i) \rightarrow g_i(n_i, \gamma_i^-), \text{ if } U_{NS}(g_i(n_i, \gamma_i^-)) \geq U_{SH}(g_i(n_i, \gamma_i)),$$

$$\text{where } \gamma_i^- = (n_i \cdot \gamma_i - 1) / n_i.$$

As discussed earlier, the information on group configuration is public knowledge. An individual user can observe the information and make a sharing decision if he can achieve a better utility. It can be shown that there exists a stable cooperation level (Nash equilibrium). To derive managerial insights from the above dynamics, we investigate the properties of the P2P group structure and evaluate the impacts of system parameters on the level of sharing behaviors.

We start with two extreme (boundary) situations. A sharing group $g_i(n_i, 0^+)$ emerges if $U_{SH}(g_i(n_i, 1/n_i)) \geq 0$. Free riding in group $g_i(n_i, 1^-)$ starts to happen when

$$U_{SH}(g_i(n_i, 1)) \leq U_{NS}\left(g_i\left(n_i, \frac{n_i - 1}{n_i}\right)\right).$$

In addition, we introduce two definitions. Group $g_i(n_i, \gamma_{Stable})$ is *stable* if

$$U_{SH}(g_i(n_i, \gamma_{Stable})) \geq U_{NS}(g_i(n_i, \gamma_{Stable}^-)), \text{ and } U_{NS}(g_i(n_i, \gamma_{Stable})) \geq U_{SH}(g_i(n_i, \gamma_{Stable}^+)).$$

The choice of γ_{Stable} is a result of the Nash equilibrium based on individual decision making. On the contrary, a group decision would make group $g_i(n_i, \gamma_{Efficient})$ *efficient*. The value of $\gamma_{Efficient}$ can be found by maximizing the group's total value, that is,

$$\gamma_{Efficient} = \max_{\gamma_i} n_i \cdot U_{AVG}(n_i, \gamma).$$

Numerical results. We develop numerical simulations based on our analytical model. To demonstrate the implications of the network size and system parameters on the resulting network configuration, the value and cost functions are assumed to be linear,

$$U = v \cdot H - c \cdot (S + D + T).$$

These simulations can, however, be easily extended to other concave value functions (such as a logarithm function) and convex cost functions (such as a quadratic function).

The simulations are repeated under different network sizes and various parameters. Typical parameter values used are as follows:

$$n = 100, \delta = 0.8, p = 0.1, \mu_{SP} = 1000, \mu_R = 100, \mu_S = 100, \lambda = 1, \rho = 0.01, v = 400, c = 100.$$

The results (Figures 1 through 4) show that larger networks (n), and higher content availability (p) and higher request rates (λ) at peer nodes, induce higher free-riding ratios. In Figure 1, there exists a critical point in the number of users, beyond which free-

riding behaviors start to emerge. The figures also indicate that the self-formed group is efficient only when the group size is sufficiently small such that all the peers in the group decide to share contents.

Figure 2 shows that a reduction in download capacity (higher δ) results in more file sharing. For example, the sharing behaviors are more common at university campus nodes, which have higher bandwidth capacities, than home nodes with dial-up or DSL connections. Figure 3 shows that free-riding behaviors increase as the content availability at each node becomes larger. This is due to the fact that fewer peers are then required to achieve the desired content availability and that more upload activities occur at each contributor. Also, peers are less willing to share their contents when the content request rate is higher because a higher request rate will result in higher upload congestion at the provision node (Figure 4).

Moreover, we find that the self-organized (Nash equilibrium) group is usually not compatible with the efficient (value-maximizing) group. The sharing ratio in a stable group is, in general, lower compared to that of an efficient group.

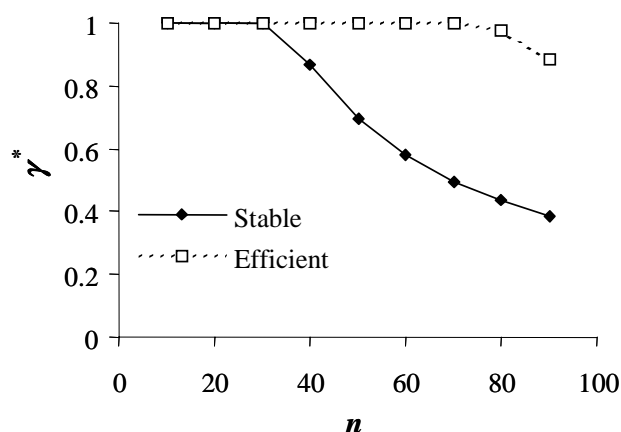


Figure 1. Effect of n on Sharing Ratio γ




Figure 2. Effect of δ on Sharing Ratio γ




Figure 3. Effect of p on Sharing Ratio γ




Figure 4. Effect of λ on Sharing Ratio γ

The complete list of the impacts of system parameters on the sharing ratio is provided in Table 2. The free-riding ratio ($1 - \gamma$) increases with the group size, content availability at each node, content request rate, and cost of waiting time, but decreases with the download capacity reduction, dispersion of the peers' locations, process capacity, and the value of content.

Table 2. Effects of System Parameters on Sharing Ratio γ

$\frac{\partial \gamma^*}{\partial n}$	$\frac{\partial \gamma^*}{\partial p}$	$\frac{\partial \gamma^*}{\partial \delta}$	$\frac{\partial \gamma^*}{\partial \rho}$	$\frac{\partial \gamma^*}{\partial \lambda}$	$\frac{\partial \gamma^*}{\partial \mu}$	$\frac{\partial \gamma^*}{\partial v}$	$\frac{\partial \gamma^*}{\partial c}$
-	-	+	+	-	+	+	-

Internal Transfer Mechanism

As discussed in the last subsection, a stable network is not always efficient. To have a stable *and* efficient network, we need to design an internal transfer mechanism. Users may select to share their contents and are waived any admission fee, or pay an admission fee but opt not to share. The internal transfer payment F_g^* (admission fee) collected by the group can be derived from the modified utility functions

$$U'_{NS}(g_i(n_i, \gamma_i)) = U_{NS}(g_i(n_i, \gamma_i)) - F_g^*$$

and

$$U'_{SH}(g_i(n_i, \gamma_i)) = U_{SH}(g_i(n_i, \gamma_i)).$$

The value of F_g^* is chosen to guarantee that $\gamma_{Stable} = \gamma_{Efficient}$.

Alternatively, the internal transfer can be carried out between contributors and free riders. That is, the collected admission fees from the NS (free riders) are distributed to the SH (contributors) as a bonus. To find the internal transfer payment under this scheme F_p^* , we modify the utilities as

$$U''_{NS}(g_i(n_i, \gamma_i)) = U_{NS}(g_i(n_i, \gamma_i)) - F_p^*$$

and



$$U''_{SH}(g_i(n_i, \gamma_i)) = U_{SH}(g_i(n_i, \gamma_i)) + \frac{1 - \gamma_i}{\gamma_i} \cdot F_p^*.$$

It is straightforward to show that $F_g^* \geq F_p^*$. Notice that the proposed internal transfer mechanisms are utilized as economic incentive mechanisms to encourage (discourage) the sharing (free-riding) behaviors to induce efficient networks without considering any additional cost of incentive mechanism implementation. Similar to P2P software and search intermediaries, incentive mechanisms (accounting and transfer systems) could be parts of the P2P infrastructure investment.

Using the same set of parameter values, we demonstrate the influence of system parameters on the optimal internal transfer. Figure 5 show that, in general, free riders pay a higher fee if there are more peers. It is intuitive, as is shown in Figure 6, that a higher download performance degradation requires more compensation.

Group Formation Decision

In the previous subsections, we investigated the group configuration based on a self-selected sharing decision, given the number of users already admitted in a group. The new user's decision to select a group and the group's willingness to grant membership should mutually improve the utilities of both parties. It is interesting to find the optimal number of active users in a stable and efficient group.

Figure 5. Effect of n on Internal TransferFigure 6. Effect of δ on Internal Transfer

DYNAMIC 2 (GROUPING DECISION). A new user will be interested in joining a group g_j and be admitted to the group if and only if the following conditions are satisfied:

1. incentive compatibility (IC) for peers in the group should hold, that is,

$$(n_j + 1) \cdot U_{AVG}(g_j(n_j + 1, \gamma'_j)) \geq n_j \cdot U_{AVG}(g_j(n_j, \gamma_j));$$

2. individual rationality (IR) for the new user should hold, that is,

$$U_{AVG}(g_j(n_j + 1, \gamma'_j)) \geq 0;$$

3. IC for the new user should hold, that is,

$$j^* = \arg \max_j U_{AVG}(g_j(n_j + 1, \gamma'_j)).$$

The above group formation dynamics describe the situation where participants (regular peers and super peers) make decisions to achieve the goal of individual and group utility maximization, under the assumption that groups operate independently without any interconnections.

If there is only one monopolistic P2P group, the optimal group size can be found as

$$n_G^* = \arg \max_{n_j} n_j \cdot U_{AVG}(g_j(n_j, \gamma_j)).$$

However, if there are competitive P2P groups, the group selection decision is based on whether or not a group will improve the expected individual utility after the new user enters. The optimal group size is now changed to

$$n_{AVG}^* = \arg \max_{n_j} U_{AVG}(g_j(n_j, \gamma_j)).$$

Since the content availability is concave and the delay cost is convex with respect to group size, $U_{AVG}(g_j(n_j, \gamma_j))$ and $n_j \cdot U_{AVG}(g_j(n_j, \gamma_j))$ both are concave functions of group size and $n_G^* > n_{AVG}^*$.

Numerical results. Using the proposed analytical model and numerical simulation, we investigate the influence of system parameters on group formation. Figures 7 and 8 show that a monopolistic group has a larger size than perfectly competitive groups. Both group sizes increase with the capacity of the super peer (Figure 7). As shown in Figure 8, n_{AVG}^* increases with the download performance degradation when $\delta < 80$ percent. We also observe that n_G^* is more sensitive to changes in the capacity of the super peer because the search delay is the bottleneck factor that affects the scalability of structured P2P networks.

Table 3 summarizes the impacts of system parameters on the group size. The group size increases with the capacity of the super peer, dispersion of the peers' locations, and the value of content, but decreases with the content availability at each node, content request rate, and the cost of waiting time.




Figure 7. Effect of μ_{SP} on Group Size n




Figure 8. Effect of δ on Group Size n
($\mu_{SP} = 200$)

Table 3. Effects of System Parameters on Group Size n

$\frac{\partial n^*}{\partial \mu_{SP}}$	$\frac{\partial n^*}{\partial p}$	$\frac{\partial n^*}{\partial \delta}$	$\frac{\partial n^*}{\partial \rho}$	$\frac{\partial n^*}{\partial \lambda}$	$\frac{\partial n^*}{\partial \mu}$	$\frac{\partial n^*}{\partial v}$	$\frac{\partial n^*}{\partial c}$
0	-	+/-	+	-	+/-	+	-

Summary and Conclusions

In this paper, we have developed an analytical model to evaluate the impacts of system parameters on the emergence of self-organized P2P network structures. Utility functions of individual and group decisions are formulated as the benefit of content availability less the delay costs (search delay, upload process delay, download process delay, and transmission delay). We present the multilevel P2P formation dynamics: individual sharing decision and group admission. Numerical results show that the cooperation level (sharing ratio) decreases with group size, content availability, and request rate, but increases with download capacity. The sharing level of a stable group is never higher than that of an efficient group. We also propose an internal transfer mechanism (admission fee for a free-rider) to achieve the compatibility of stability and efficiency.

In our model, we assume that nodes (regular peers and super peers) are symmetric and that information is public knowledge. Investigating the emerging P2P structure under heterogeneous players (peers and super peers) with asymmetric information is a planned future extension. It will also be interesting to study how our results will change if the players of a P2P network are

strategic rather than myopic. Further, utilizing the same performance metrics and utility functions, we plan to extend the model to investigate the interconnection structures among various groups. Based on the principle of incentive compatibility, the interconnection between two groups will be formed if mutual interests exist, that is, if the mutual query forward agreement improves the utilities for both groups.

Besides the (social and economic) incentive-based formation dynamics, an interesting topic for future research would be to study the evolving dynamics of P2P network complexity. With the capacity constraint of the super peers, as the number of peer nodes increases, the formation of a scale-free P2P network should evolve through self-organized group splitting and hierarchical organizing. It would be important and interesting to study how the formation changes over time, and what impacts this might have on the P2P network performance.

References

- Adar, E., and Huberman, B. A. "Free-Riding on Gnutella," *First Monday* (5:10), October 2000 (available online at http://www.firstmonday.dk/issues/issue5_10/adar/index.html).
- Asvanund, A., Bagla, S., Kapadia, M. H., Krishnan, R., Smith, M. D., and Telang, R. "Intelligent Club Management in Peer-to-Peer Networks," *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 5-6, 2003 (available online at <http://www.sims.berkeley.edu/research/conferences/p2pecon/papers/s6-asvanund.pdf>).
- Asvanund, A., Clay, K., Krishnan, R., and Smith, M. D. "An Empirical Analysis of Network Externalities in P2P Music-Sharing Networks," *Information System Research* (15:2), June 2004, pp. 155-174.
- Axelrod, R., Mitchell, W., Thomas, R. E., Bennett, D. S., and Bruderer, E. "Coalition Formation in Standards-Setting Alliances," *Management Science* (41:9), September 1995, pp. 1493-1508.
- Feldman, M., Lai, K., Chuang, J., and Stoica, I. "Quantifying Disincentives in Peer-to-Peer Networks," *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 5-6, 2003 (available online at <http://www.sims.berkeley.edu/research/conferences/p2pecon/papers/s1-lai.pdf>).
- Golle, P., Leyton-Brown, K., Mironov, I., and Lillibridge, M. "Incentives for Sharing in Peer-To-Peer Networks," in *Proceedings of the Second International Workshop on Electronic Commerce*, Heidelberg, Germany, 2001, pp. 75-87.
- Jackson, M. O., and Wolinsky, A. "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory* (71), 1996, pp. 44-74.
- Khambatti, M., Dasgupta, P., and Ryu, K. "Peer-to-Peer Communities: Formation and Discovery," in *Proceedings of the 14th IASTED International Conference on Parallel and Distributed Computing and Systems*, Cambridge, MA, 2002, pp. 166-173.
- Kleinrock, L. *Queueing Systems: Volume 2, Computer Applications*, Wiley, New York, 1976.
- Krishnan, R., Smith, M. D., Tang, Z., and Telang, R. "The Virtual Commons: Why Free Riding Can Be Tolerated in File Sharing Networks," in *Proceedings of the 23rd Annual International Conference on Information Systems*, L. Applegate, R. Galliers, and J. I. DeGross (Eds.), Barcelona, 2002, pp. 807-812.
- Krishnan, R., Smith, M. D., Tang, Z., and Telang, R. "The Impact of Free-Riding on Peer-to-Peer Networks," in *Proceedings of the 37th Hawaii International Conference on System Sciences*, IEEE Computer Society Press, 2004 (available online at <http://csdl.computer.org/comp/proceedings/hicss/2004/2056/07/205670199c.pdf>).
- Ledlie, J., Taylor, J., Serban, L., and Seltzer, M. "Self-Organization in Peer-to-Peer Systems," in *Proceedings of the 10th European SIGOPS (Special Interest Group on Operating Systems)*, Saint Emilion, France, September 2002 (available online at <http://www.eecs.harvard.edu/~syrah/papers/sigops-02/sigops-02.pdf>).
- Li, Y.-M., Tan, Y., and Zhou, Y.-P. "On the Scale of Peer to Peer Networks," in *Proceedings of the 13th Annual Workshop on Information Technologies and Systems*, Seattle, WA, December 2003, pp. 13-18.
- Ng, T. S. E., and Zhang, H. "Predicting Internet Network Distance with Coordinates-Based Approaches," in *Proceedings of the 21st IEEE INFORM INFOCOM (Conference on Computer Communications)*, New York, June 2002, pp. 170-179.
- Ratnasamy, S., Francis, P., Handley, M., Karp, R., and Shenker, S. "A Scalable Content-Addressable Network," in *Proceedings of ACM SIGCOMM (Special Interest Group on Data Communications)*, San Diego, CA, August 2001, pp. 161-172.
- Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., and Balakrishnan, H. "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications," *Proceedings of ACM SIGCOMM*, August 2001, pp. 149-160.
- Zipf, G. K. "Relative Frequency as a Determinant of Phonetic Change," *Harvard Studies in Classical Philology* (15), 1929, pp. 1-95.

