

Analysis of Scale Effects in Peer-to-Peer Networks

Yung-Ming Li, Yong Tan, and Yong-Pin Zhou

Abstract – In this paper, we study both positive and negative scale effects on the operations of Peer-to-Peer (P2P) file sharing networks, and propose the optimal sizing (number of peers) and grouping (number of directory intermediary) decisions. Using analytical models and simulation, we evaluate various performance metrics to investigate the characteristics of a P2P network. Our results show that increasing network scale has a positive effect on the expected content availability and transmission cost, but a negative effect on the expected provision and search costs. We propose an explicit expression for the overall utility of a content sharing P2P community that incorporates tradeoffs among all the performance measures. This utility function is maximized numerically to obtain the optimal network size (or scale). We also investigate the impact of various P2P network parameters on the performance measures as well as optimal scaling decisions. Furthermore, we extend the model to examine the grouping decision in networks with symmetric interconnection structures, and compare the performance between random- and location-based grouping policies.

Index Terms – Peer-to-peer networks, network operations and management, performance evaluation, queueing analysis

I. INTRODUCTION

Peer-to-peer (P2P) technologies link social networks into cooperative ventures that share information (audio, video and graphic files), computer resource (computing cycles, hard disk space, and network bandwidth), and communication and collaboration (instant messaging). Members of a P2P community exchange information or other resources directly with each other, with very little or no use of a centralized or dedicated server. Many P2P services exist today, such as file sharing service (Gnutella and Freenet), grid computing service (Popular Power and Distributed Net), instant messaging service (AOL, Yahoo!, and MSN), and online collaboration service (Groove Networks).

Among various P2P applications, file sharing is probably the most popular. P2P file sharing applications accounted for five of the top 10 downloads from the download.com web site in the last week of June 2002, together constituting 4.5 million downloads [18]. In contrast to the traditional web server based content delivery paradigm, this emerging “bottom-up” mode

of information distribution, leveraging the resources on the peer nodes, is considered to be superior [19][20]. P2P file sharing networks have attracted many users and much press attention, along with the ire from media firms who feel threatened by the illegal exchange of digital music and movie files.

P2P technologies have many operational characteristics that make them appealing: First, they rely on peer nodes, not the central servers, to deliver content, and therefore are more scalable. Second, on a large P2P network, it is likely for any node to find another node with the desired content that is “close,” so transmission delay may be lowered as well. There are, however, drawbacks inherent in the P2P networks, due to the same decentralized structure: First, because each peer node can modify its content freely, it may be costly to find desired contents. Second, since P2P users obtain contents from each other, the availability of these contents completely depends on the peer nodes being logged on. So, content reliability may be an issue.

In many ways, the size of a P2P network can impact many of these factors. A large network could alleviate the content reliability problem because the probability of satisfying requested content becomes higher if more peer nodes participate in file sharing activities. It could also reduce transmission delay, on average, as the closest service node will become closer as the network contains more nodes with the same content replica. P2P technologies utilize aggregate bandwidth from edge nodes for content transmission to avoid congestion at dedicated servers. Therefore, the effective bandwidth is scalable with respect to the number of active users. On the other hand, on a large-scale P2P network, the number of queries may cause congestion at directory server (if any) as well as network traffic congestion (one query may be forwarded multiple times before a suitable service node is found), due to limited capacity and network bandwidth. Therefore, determining the “right” network scale is very important for P2P operations.

In this paper, we propose four metrics to evaluate the impacts of network scale on operational performance of a P2P network: content availability, search delay, provision delay, and transmission delay. Using these metrics and balancing all the tradeoffs, we examine the overall scale effect (network externality – the impact of network size, on the individual’s utility) and suggest optimal scale decisions, from a P2P network organizer’s perspective. In particular, we focus on the impacts of local peer parameters, such as P2P participants’

Yung-Ming Li is with Institute of Information Management, National Chiao Tung University, Hsinchu, Taiwan (e-mail: yml@mail.nctu.edu.tw). Yong Tan and Yong-Pin Zhou are with Department of Management Science, University of Washington, Seattle, WA 98195-3200, USA (e-mail: ytan@u.washington.edu; yongpin@u.washington.edu).

computing and bandwidth capacities, local content provision amount, content request pattern and frequency, and sharing propensity.

Adar and Huberman [1] surveyed the Gnutella network and found 70% of peers on the network were free riders. The prevalence of free riders, who take but never contribute, not only reduces the aggregate content availability but also increases the workload at non-free-riding nodes. In this paper, we will also investigate the impacts of dispersion of P2P users, content variety distribution, and content popularity distribution.

Several P2P structures exist and vary by their search algorithms [28]. 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, have to deal with excessive network traffic due to decentralized broadcast-type search. Newer generations of P2P software, such as KaZaA and Gnutella v0.6, use a combination of centralized and pure decentralized network structure: peer nodes are grouped and served by supernodes (or, super-peers). Various groups are interconnected via supernodes to forward requests [27].

In this paper, we analyze the promising supernode-based P2P network structures. In particular, because the scale of a P2P network plays an important role in determining network performance, we investigate two important operational issues of a P2P network: sizing and grouping decisions. Sizing refers to the determination of the optimal size for a P2P community for any given supernode (i.e., the optimal number of peers connected to the same supernode). Grouping refers to the partition of a fixed number of nodes into multiple P2P communities (i.e. the optimal number of supernodes, given the number of peers). An important factor in grouping decision is the interconnection structure among groups. Therefore, we compare performances between random- and location-based grouping decisions, which could be supported by new P2P protocols.

While much of P2P research has been directed to the technological issues such as search algorithms and topology designs, there is little attention paid to the operational aspects of a P2P network. In this paper, we use probabilistic distributions and queueing models to investigate the characteristics of typical P2P system dynamics, and present several main performance metrics for evaluating a P2P network. These metrics allow us to study both positive and negative scale effects on a P2P file sharing network, and suggest optimal scale (sizing and grouping) decisions.

The remainder of the paper is organized as follows. Section II reviews related literatures on P2P networks. Section III gives a formal description of the model, outlines system parameters, and proposes performance metrics. In Section IV, we analyze the effects of scale and other network parameters on the proposed performance metrics and present simulation results. Section V examines the impacts of various system parameters on optimal scale decision. Section VI extends the

model to analyze grouping decision in various interconnection structures. Section VII concludes our findings and offers directions for future research.

II. LITERATURE REVIEW

There are a number of papers on the technical aspects of P2P networks. These papers focus mainly on developing efficient communication protocols, network topologies, and search algorithms [25][33]. Supernode structure is a promising structure of P2P networks developed to improve search efficiency in pure P2P networks. Singh *et al.* [31] present incentive mechanisms for several participants, especially service providers, to deploy a supernode infrastructure. Yang and Garcia-Molina [35] evaluate performances and present practical guidelines for the design of an efficient supernode network. Singh *et al.* [30] present incentive schemes for deploying supernodes and propose a topic-based search mechanism to improve the effectiveness of supernodes.

On the topic of networks scale, Asvanund *et al.* [2] empirically analyze network externality in P2P music sharing networks and suggest that larger networks are not always better. Yang and Garcia-Molina [34][36] design various content-sharing P2P search architectures and compare the maximum number of users that can be served on them. Butler [4] investigates the effect of membership size and communication activity on sustainability of online social structure. The results of this study suggest that networked communication technologies provide benefits to balance the opposing impacts from membership size. These studies provide valuable empirical evidences on scale effect, but they do not present underlying operational metrics for evaluating network performance and for gaining insights on optimal scale decisions.

Regarding the grouping of P2P networks, Asvanund *et al.* [3] propose a scheme for club membership management based on content similarity and physical location. Ledlie *et al.* [17] develop a hierarchically grouped system that can self-organize to overcome unreliability. Khambatti *et al.* [12] use attribute-based clustering models to simulate how self-configuring communities are formed. Their results demonstrate that community structures in a random network can be efficiently discovered based on attribute and link information of peers.

Recently, a few researchers have started to explore the social and economical aspects of P2P free riding phenomenon and incentive mechanism design. For example, Golle *et al.* [9] construct a formal game theoretic model to develop and analyze several payment mechanisms to encourage file-exchange activities. Krishnan *et al.* [13][14] propose a plausible model to analyze the existence of free-riding behaviors in P2P file-sharing networks. However, the 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. While most of researches on P2P networks in technological domains assume that users follow prescribed protocols

without deviation, Shneidman and Parkes [29] advocate a P2P model in which users are rational and self-interested. They develop a new operating mechanism that allows users to behave rationally while still achieving good overall system outcomes. Using economic incentive model, Jackson and Wolinsky [11] examine whether efficient (value maximizing) social networks will form when self-interested individuals can choose to form or sever links.

Additionally, many reputation and trust systems are proposed to provide incentives for cooperation without involving a pricing scheme [7]. For example, Ranganathan *et al.* [24] propose a multi-person prisoner's dilemma model to investigate user behaviors and develop pricing- and reputation-based mechanisms to improve system performance. Wang and Vassileva [37] propose a Bayesian network based model to build reputation that is based on recommendation in P2P network. Kung and Wu [15] present a reputation-based P2P admission system, using eigenvector approach, to allow only those nodes that have made reasonable service contributions to receive services from others.

However, to the best of our knowledge, little attention has been given to the operational aspects of P2P networks so far. Christin and Chuang [6] propose metrics for estimating latency, sharing, routing, and maintaining cost in order to investigate the social optimum structure of P2P networks. In this paper, we focus on scale issues, and develop analytical model to examine how network size and system parameters affect performances of P2P networks and optimal sizing and grouping decisions.

III. THE MODEL

We consider a content-sharing P2P network in which the participants are categorized as regular peer nodes and supernodes. A supernode and a number of regular peer nodes form a community. Only supernode maintains up-to-date information on all resources available in the community. Every content request (query) is generated at one of the peer nodes, and first processed at the local supernode on a first-come, first-served basis. For every query it processes, the supernode recommends a provision node that has the desired file and the lowest expected download delay. For example, many kinds of P2P file-sharing software (such as BitTorrent, KaZaA) provides the information of download speed and estimated time for each qualified provision node. Once this information is passed on to the requesting and provision nodes, download occurs directly between these two nodes. There could be many supernodes but each peer node is connected to only one supernode whenever it logs on. If a query cannot be satisfied from the local community, it will be forwarded to other supernodes.

A. Operating Policy

Figure 1 depicts the operations of a supernode-based P2P network. On a snapshot of a network, a peer node A in community G_1 needing a file that it doesn't own sends a content request to the local community center. The supernode

of community G_1 , SP_1 , searches its directory database and responds with a list of nodes that share the requested content (e.g. nodes B and D), along with the download information (approximate delay). It also recommends the node with the minimum download delay as the provision node (node D). After that, the requesting node A downloads the content directly from provision node D . If the request is not satisfied (i.e., no node shares the requested content in the local community G_1), the query will be forwarded to other interconnected supernodes, SP_2 and SP_3 , based on various peering policies (such as parallel or sequential forward). In the paper, we assume that unsatisfied requests will be broadcasted (forwarded in parallel) to all interconnected supernodes.

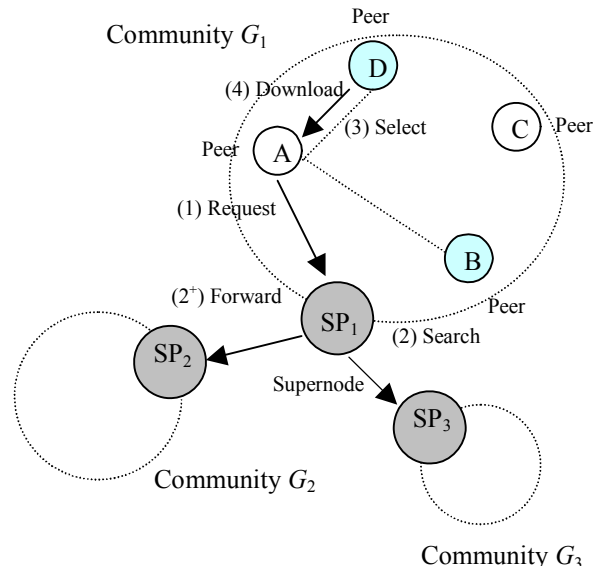


Fig. 1. Operational sequences of hybrid P2P networks (supernode structure).

TABLE I
MODEL PARAMETERS

n	Number of active P2P users (peers)
s	Number of supernodes (groups)
M	Number of file varieties
α	File popularity distribution parameter (Zipf coefficient)
β_i	Average number of files stored in peer node i
$\gamma_{i,m}$	The probability that peer node i share its file m
μ_0	Search service rate of supernode
μ_i	Provision service rate of a regular peer node i
λ_i	Content request rate by peer node i
τ	Upper bound of transmission delay

B. Network Dynamics and Distributions

Since each peer node is a content consumer as well as a content provider, the dynamics of a P2P network are highly dependent on the parameters of local peer nodes. The parameters, listed in Table I, are used to describe network dynamics such as content provision distribution, content

request distribution, and bandwidth (or, transmission delay) distribution.

Content Provision Distribution. More popular contents are assumed to be stored and requested with higher probability. It is interesting to investigate how the variation of content popularity affects the operational performance of a P2P network. We assume that there are M same-size files in a P2P community, denoted by F_1, F_2, \dots, F_M . These files are ranked in descending order by their popularity. Let θ_m be the probability that measures the popularity of file F_m , $\forall m \in \{1, \dots, M\}$. We assume that θ_m follows a Zipf-like distribution: $\theta_m = m^{-\alpha} \cdot \theta_0$, $\forall m \in \{1, \dots, M\}$, where $\theta_0 = 1 / \sum_{m=1}^M m^{-\alpha}$ is the normalization factor, and $\alpha \geq 0$ is a parameter for the relative popularity distribution. The greater the value of α , the larger the variation of popularity among files. When $\alpha = 0$, the distribution is uniform, i.e., each file has the same popularity.

Assume that there are n active peer nodes in a P2P community. We introduce binary random variables, $X_{i,m}$, to indicate the availability of file F_m stored on node i , and $\bar{X}_{i,m}$, the availability of file F_m shared on node i . Explicitly,

$$X_{i,m}(\bar{X}_{i,m}) = \begin{cases} 1, & \text{if node } i \text{ has (shares) file } F_m, \\ 0, & \text{otherwise,} \end{cases}$$

$$\forall i \in \{1, \dots, n\}, m \in \{1, \dots, M\}.$$

Let $P_{i,m} = E(X_{i,m})$ and $\bar{P}_{i,m} = E(\bar{X}_{i,m})$; they denote the probability that node i has (shares, respectively) file F_m . Obviously, $X_{i,m} \geq \bar{X}_{i,m}$ and $P_{i,m} \geq \bar{P}_{i,m}$. Following Chu *et al.* [5] and Gummadi *et al.* [10] who find that content replication follows a Zipf distribution, we assume that $P_{i,m}$ is proportional to θ_m , the probability of popularity of file F_m , i.e., $P_{i,m} = \beta_i \cdot \theta_m$. Here, the value of β_i is limited by $\beta_i \cdot \theta_1 \leq 1$, and β_i represents the intensity of content availability. Alternatively, β_i can be interpreted as the average number of files stored in peer node i . A peer node can decide whether to share its own files (i.e., allow upload). Therefore, the content availability at peer node i becomes $\bar{P}_{i,m} = P_{i,m} \cdot \gamma_{i,m}$, where $\gamma_{i,m}$ is the probability that peer node i decides to share file F_m .

Content Request Distribution. Let λ_i be the request rate for all files by node i , and $\lambda_{i,m}$ be the request rate for file F_m by node i . Obviously, $\sum_{m=1}^M \lambda_{i,m} = \lambda_i$. Sripanidkulchai [32] analyzes the characteristics of queries on Gnutella and finds that the popularity of search strings follows a Zipf-like distribution. Consistent with this finding, we assume that content requests are described by a Zipf-type distribution. Hence, the request rate for file F_m can be written as $\lambda_{i,m} = \theta_m \cdot \lambda_i$. However, a peer node would typically not request a file that it already has; in other words, peers

download the same content at most once. Our queuing analysis corresponds to a ‘‘snapshot’’ of a network, where some peer nodes do have some files stored. Therefore, the actual content request rate for file F_m is $(1 - P_{i,m}) \theta_m \cdot \lambda_i$, which, as one can easily see, does not follow a Zipf-like distribution. This agrees with Gummadi *et al.* [10] who show that the file request from KaZaA network is not Zipf-type since users fetch multimedia objects at most once.

Search and Provision Processes. The two most time-consuming activities during the entire processes of P2P content distribution are search and download. The former occurs at supernode, and the latter occurs between peer nodes. We use standard queuing model to evaluate the delay caused by each activity. Both the service times for search process at supernode and provision process at peer node i are assumed to follow an exponential distribution with rates μ_0 and μ_i , respectively. Previous research [8][21] suggest that Poisson process is valid for modeling arrivals of user-initiated requests. Qiao *et al.* [23] also provide evidence that content request arrivals form a Poisson process in a P2P network. Therefore, we assume that requests follow a Poisson process and the search at supernode can be modeled as an M/M/1 queue. Moreover, because the departure process of an M/M/1 queue is also Poisson (see Corollary 5.6.2 in [26]), the requests that are forwarded to peer nodes for download also follow a Poisson process. Therefore, we can model the provision process at each peer node also as an M/M/1 queue.

Transmission Delay Distribution. In general, it is difficult to exactly estimate transmission latency of download activity. Considering the properties of dynamic uptime and position of peer nodes (for example, users are logged on only for a short time period or use mobile computers), we assume that, at a snapshot, the transmission delay between content request node i and provision node j , $T_{i,j}$, is an *i.i.d.* random variable with its value drawn from a transmission delay density function. For the purpose of analytical tractability, we assume that $T_{i,j}$ is uniformly distributed on $[0, \tau]$, where τ is the upper bound of transmission delay. As is clear in the derivation of expected transmission delay, any other forms of distribution pose no conceptual difficulty, but make it hard to obtain closed-form solutions.

Provision Policy. If more than one node can provide a file, the node with the minimum expected download delay (provision delay plus transmission delay) to the requesting node is selected as the provision node. That is, when node i requests file F_m , the community center suggests the optimal provision node j^* , where,

$$j^* = \arg \min_j (D_j + T_{i,j}) \text{ s.t. } \bar{X}_{j,m} = 1.$$

C. Performance Metrics

The performance metrics are established based on the benefits and costs of each activity (request, search, download, and transmission) during the entire process of P2P content distribution. Specifically, these metrics include content

availability at requested peer node, search delay at supernode, provision (upload) process delay, and transmission delay on the network. Since each peer is simultaneously a content consumer and a content provider, all these metrics are associated with the number of active peers, their behaviors (such as request frequency and sharing decision), and service capacities of the supernode, the provision nodes, and the network. In our analytical model, we assume that all the peer nodes are *statistically* identical (they follow the same set of probability distributions) and the sharing decision distribution for every file is the same. That is,

$$P_{i,m} = P_m, \bar{P}_{i,m} = \bar{P}_m, \lambda_{i,m} = \lambda_m, \mu_i = \mu, \beta_i = \beta, \gamma_{i,m} = \gamma, \\ \forall i \in \{1, \dots, n\}, m \in \{1, \dots, M\}$$

In our simulation studies, this homogeneity assumption is relaxed to allow heterogeneous peer nodes.

Content Availability. Content availability (or hit rate) is an important measure of the quality of content provision. It is defined as the probability that an arbitrary request can be satisfied on a P2P network. Hence, it depends on content stored and shared on peer nodes. Let $H(n, m)$ be the expected content availability of file F_m in a local community with n nodes. We have

$$H(n, m) = 1 - \prod_{i=1}^n (1 - \bar{P}_m).$$

The overall expected content availability can be written as

$$H = E_m(H(n, m)) = \sum_{m=1}^M \theta_m \cdot \left(1 - \prod_{i=1}^n (1 - \bar{P}_m)\right).$$

Search Delay. The cost of waiting occurs at the supernode. In a hybrid supernode P2P network, all content requests are forwarded to the local supernode. The local supernode responds with the download information of the nodes in the same community who have the requested content, or, if none exist, forwards the unsatisfied requests to other interconnected supernodes in remote communities. Therefore, the search delay includes the expected total system waiting times at the local and remote supernodes, for any network topology and request forward protocol. Since content requests from all the peer nodes are independent Poisson processes, the aggregate content request arrival at a supernode is also a Poisson process whose arrival rate, Λ_0 , is the accumulation of all the arrivals:

$\Lambda_0 = n \sum_{m=1}^M (1 - P_m) \theta_m \lambda$. Using a queueing model, we can express the search delay of an isolated community as $S = (\mu_0 - \Lambda_0)^{-1}$, where μ_0 is the service rate of the supernode.

Provision Delay. The provision delay is caused by congestion at the provision node. One provision node may be serving several requests (providing download) simultaneously, so the expected provision delay is estimated from the aggregated content request rate and the process capacity of provision nodes. Specifically, we define a binary random variable, $Z(i, j, m)$, which is 1 if and only if j is selected as

the provision node for file F_m requested by node i . Hence, node j is selected as the provision node for file F_m with probability

$$\Pr(Z(i, j, m) = 1) \\ = \Pr(X_{i,m} = 0; \bar{X}_{j,m} = 1; T_{i,j} < T_{i,k}, \forall k \neq i, j, \text{ s.t. } \bar{X}_{k,m} = 1).$$

Assuming that all nodes are identical, we find

$$\Pr(Z(i, j, m) = 1) \\ = (1 - P_m) \cdot \bar{P}_m \cdot \sum_{k=0}^{n-2} \frac{1}{k+1} \cdot \binom{n-2}{k} \cdot \bar{P}_m^k \cdot (1 - \bar{P}_m)^{n-2-k} \\ = \frac{(1 - P_m)}{(n-1)} \cdot \sum_{k=1}^{n-1} \binom{n-1}{k} \cdot \bar{P}_m^k \cdot (1 - \bar{P}_m)^{n-1-k} \\ = (1 - P_m) \cdot \left(1 - (1 - \bar{P}_m)^{n-1}\right) / (n-1).$$

Therefore, the aggregate request arrival rate at peer node j (given that j shares F_m) can be obtained as

$$\Lambda_j = \sum_{i \neq j} \sum_m \lambda_{i,m} Z(i, j, m) = \sum_{m=1}^M \theta_m (1 - P_m) \left(1 - (1 - \bar{P}_m)^{n-1}\right) \lambda.$$

Finally, the expected provision delay is $D = (\mu - \Lambda_j)^{-1}$.

Transmission Delay. This delay depends strongly on the number of active nodes, because the more the active nodes the more likely it is to find a provision node closer to the requesting node. Let τ be the maximum transmission delay, and denote the expected minimum transmission delay among k nodes by $T_{\min(k)}$. Using order statistics, we have

$$T_{\min(k)} = \int_0^{\tau} t \cdot k \cdot (1 - F(t))^{k-1} \cdot f(t) \cdot dt = \frac{\tau}{k+1},$$

where $f(t)$ and $F(t)$ are the PDF and CDF for transmission delay, which is assumed to follow a uniform distribution. Next, $T_{\min}(n, m)$, the expected transmission delay for file F_m in a network with n nodes, can be evaluated as

$$T_{\min}(n, m) = E(T_{\min(k)} | F_m \text{ is available}) \\ = \frac{1}{H(n-1, m)} \sum_{k=1}^{n-1} T_{\min(k)} \binom{n-1}{k} \bar{P}_m^k (1 - \bar{P}_m)^{n-1-k} \tau \\ = \frac{1}{H(n-1, m) \cdot n \cdot \bar{P}_m} \sum_{k=2}^n \binom{n}{k} \bar{P}_m^k (1 - \bar{P}_m)^{n-k} \tau.$$

After some simplifications, we have,

$$T_{\min}(n, m) = \frac{1 - (1 - \bar{P}_m)^{n-1} \cdot (1 + (n-1) \cdot \bar{P}_m)}{(1 - (1 - \bar{P}_m)^{n-1}) \cdot n \cdot \bar{P}_m} \cdot \tau.$$

We further average the above expression over m to obtain the expected transmission delay

$$T = E_m(T_{\min}(n, m)) = \sum_{m=1}^M \theta_m \cdot T_{\min}(n, m).$$

IV. PERFORMANCE ANALYSIS

All the performance metrics proposed above are inherently determined by the scale of a network. This section is logically divided into three subsections. First, we focus on the impacts of network scale on all performance metrics. Both positive and negative scale effects are discovered (Table II) and their respective characteristics are analyzed. This allows performance bottleneck to be identified. Next, we extend the analysis to examine the effects of other network parameters (Table III). Several interesting observations have been made; for example, the popularity variation parameter (α) is found to have mixed effects on content availability (Figures 2 and 3). Finally, to validate our analytical results, we provide a simulation study that relaxes some restrictive assumptions, such as homogeneity.

A. Analysis of Scale Effects

Using the assumptions described above, we have the following results for the performance of a P2P network when its scale (size) changes. Table II summarizes network scale effects on the performance. Here Δ_n and Δ_n^2 denote the first and second order differences in n , respectively.

TABLE II
EFFECTS OF SCALE ON P2P NETWORK PERFORMANCE METRICS

	H	S	D	T
Δ_n	+	+	+	-
Δ_n^2	-	+	*	+

“+” stands for positive effect

“-” stands for negative effect

*: - if μ or n is large enough

PROPOSITION 1. (CONTENT AVAILABILITY, H) *The expected content availability is an increasing and concave function of the size of a P2P network.*

A larger P2P network improves content availability but the benefit is marginally diminishing. In addition, the average number of file F_m in the community with n nodes can be described as

$$R(n, m) = \sum_{k=1}^n k \cdot \binom{n}{k} \cdot \bar{P}_m^k \cdot (1 - \bar{P}_m)^{n-k} = n \cdot \bar{P}_m.$$

Therefore, the average number of replicas of a file increases with network size.

PROPOSITION 2. (SEARCH DELAY, S) *The search delay is an increasing and convex function of the size of a P2P network.*

A larger P2P network incurs a higher search delay. The search cost displays diseconomy of scale. Given a limited capacity, the search delay at community center is typically the performance bottleneck of a large scale P2P network.

PROPOSITION 3. (PROVISION PROCESS DELAY, D) *The expected provision process delay increases with the size of a P2P network.*

A larger P2P network incurs higher provision delay. Moreover, the cost is marginally decreasing when provision capacity μ , or network size n , is large enough. The aggregated request arrival rate for a certain file F_m at peer node j , $\Lambda_j(n, m)$, is written as

$$\Lambda_j(n, m) = (1 - P_m) \cdot H(n - 1, m) \cdot \lambda_m.$$

This indicates that a larger number of active peers in the network make higher content availability, result in a higher aggregated request arrival rate, and consequently cause higher provision congestion. However, because $\Lambda_j(n, m)$ is bounded by $(1 - P_m) \cdot \lambda_m$, the provision delay is also bounded. This suggests that provision delay is not the critical factor that limits scalability.

PROPOSITION 4. (TRANSMISSION DELAY, T) *The expected transmission delay is a decreasing and convex function of the size of a P2P network.*

A larger P2P network reduces the expected transmission delay, but the benefit is marginally decreasing. We can rewrite $T_{\min}(n, m)$ as,

$$T_{\min}(n, m) = \left(\frac{1}{R(n, m)} - \frac{(n-1) \cdot (1 - H(n-1, m))}{n \cdot H(n-1, m)} \right) \cdot \tau.$$

As investigated in Proposition 1, there will be more content replicas as more peer nodes join in the community. In a traditional client-server network structure, there is no extra benefit if more than one content replica are cached at the same content server. However, in P2P file sharing networks, higher degree of content replicas indicates that better transmission performance may be achieved, by selecting a closer provision node with less transmission time.

B. Analysis of Parameter Effects

The impacts of system parameters on various performance measures are described in Proposition 5 and summarized in Table III.

PROPOSITION 5.

1. (CONTENT INTENSITY) *Higher content intensity (higher β) always yields better performance in content availability, transmission, and search. Higher β improves the provision performance only when the community size is sufficiently large.*

2. (SHARING LEVEL) *Higher sharing ratio (higher γ) improves the performance in content availability and transmission, but increases provision delay.*

3. (REQUEST FREQUENCY) *Higher content request rate (higher λ) results in larger provision and search delays.*

4. (UPLOAD CAPACITY) Higher provision capacity (higher μ) reduces provision delay.

TABLE III
IMPACTS OF SYSTEM PARAMETERS ON P2P NETWORK PERFORMANCE METRICS

numerator\denominator	∂H	∂T	∂D	∂S
r				
$\partial\beta$	+	-	+/-	-
$\partial\gamma$	+	-	+	0
$\partial\lambda$	0	0	+	+
$\partial\mu$	0	0	-	0
$\partial\alpha$	+/-	-	+/-	-

“+” stands for positive effect
“−” stands for negative effect when n is large
“+/-” stands for positive effect when n is small and negative effect when n is large

The impacts of systems parameters can be explained as follows. Higher β yields higher content availability, which increases the expected number of files, and consequently reduces the expected transmission delay. This would induce higher request rate at provision nodes. However, at the same time, higher β also indicates lower probability that a node needs content from other nodes, and consequently reduces request rate at provision nodes. The overall impact on provision delay is therefore determined by these two competing factors. When n is large, the impact of β on content availability becomes less significant, so higher β turns out to reduce provision delay. The opposite is true when the community size is small. Lower probability that a node needs content from other nodes results in lower request rate and lower delay at a supernode.

The more a peer node is willing to share its content (higher γ), the higher content availability and lower expected transmission delay. Since the value of γ is independent of content replica at a peer node, it has no impact on search delay. Higher content availability results in higher request rate, and consequently higher provision delay at a peer node. Higher content request rate (demand) induces higher provision and search delays. Provision delay can be improved by increasing the capacity investment of a peer node.

We have also conducted numerical investigations on the impacts of variation of popularity on performance metrics with respect to different numbers of file varieties and network size. We find that a higher variation of popularity among files, or a higher α , results in better transmission and search performance. However, the effect of popularity variation (α) on content availability and provision delay can be negative or positive, depending on network size and file variety. Figure 2 shows that content availability decreases with α if α is below a threshold value α^* . It can be observed (in Figure 3) that α^* becomes larger when there are more nodes (higher n) or smaller number of file varieties M . Similarly, provision delay is found to decrease with α only when α is greater than a threshold value (which could be different from α^*). This

threshold value decreases with network size but increases with the number of file varieties. Therefore, we can infer that the variety of popularity has positive (negative) effect on content availability and provision delay in a smaller (larger) network or a network with larger (smaller) number of file variety.




Fig. 2. Effect of α on content availability.




Fig. 3. Effect of n and M on critical value α^* .

C. Simulation Validation

The simulation validation is based on our analytical model, applying discrete event simulation methods. We relax the assumption that content distribution across files, or request behavior across users, is identical. The heterogeneity in file sharing behaviors is also introduced, by randomly assigning a specific node to either share (as an altruistic node) or not share (as a selfish node) according to the value of γ . For each simulation run, we generate 10,000 content requests. The simulations are repeated under different network sizes and various parameters such as degree of variety of Zipf content distribution (α), intensity of content availability (β), content request arrival rate (λ), and the number of file varieties (M). Typical parameter values are given as $\alpha=1$, $\beta=1$, $\gamma=1$, $\lambda=1$, $M=10$, $\mu=10$, $\tau=1$.




Fig. 4. Effect of α on content availability $H(n)$.




Fig. 5. Effect of α on transmission delay $T(n)$.




Fig. 6. Effect of λ on provision delay $D(n)$.

The simulation results are consistent with our analytical model. Larger networks have higher content availability (hit rate) and lower transmission delay. However, higher hit rate results in more file transfer activities, and hence increases the expected process delay. Figure 4 plots the effect of Zipf distribution coefficient α . Higher α results in lower content availability when the community size is large, however, content availability increases with α when n is small. The results are consistent with the analytical observations, described in Section IV.B, that the critical value α^* increases with community size, and content availability decreases (increases) with α if α is below (above) α^* . Figure 5 shows higher α always results in lower transmission delay. Figure 6

shows that the expected provision process waiting time is stable with the network size when the request rate λ is low.

V. THE OPTIMAL COMMUNITY SIZE

In the previous section, we investigate the impacts of number of active users in a P2P network on various performance metrics. Our analyses show that these performance measures often display opposite behaviors. This suggests possible tradeoffs, which, if balanced, can determine a network size that optimizes the overall utility of a file sharing P2P community. Next, we examine two operational decisions: optimal sizing (number of peers) for a given supernode in this section, and optimal grouping on multiple supernodes for a given number of peers, which will be discussed in the next section.

From scale effect analysis, we know that search activity at a supernode is the performance bottleneck. Therefore the maximum size for an isolated P2P community is bounded, due to the fact that search delay is convex and increasing with respect to the number of P2P users. The optimal scale is determined so as to maximize the expected utility of a network, which depends on content availability value function and delay cost functions. Let $V(\cdot)$ be the value function of content availability, and $C(\cdot)$ the cost function of waiting times. $C(\cdot)$ is assumed to be convex, while the value function $V(\cdot)$ is concave in content availability. The optimal community size can be obtained as

$$n^* = \arg \max_n V(H(n)) - C(T(n) + D(n) + S(n)).$$

Without loss of generality, the value and cost functions are assumed to be

$$U = v \cdot (H(n))^K - c \cdot (T(n) + D(n) + S(n))^L,$$

where $0 \leq K \leq 1$ and $L \geq 1$.

Numerical Results. In the following, we present some numerical results. Typical parameter values are $K = 1$, $L = 1$, $\alpha = 1$, $\beta = 1$, $\gamma = 1$, $\lambda = 1$, $M = 10$, $\mu = 10$, $\mu_0 = 100$, $v/c = 5$, $\tau = 0.1$. We vary the values of parameters to investigate the impacts of these parameters on the community size decision.




Fig. 7. Optimal size versus K and L .

Figure 7 illustrates the effect of various utility functions on

the decision of optimal community size. The optimal network size increases with coefficient K since a larger value of K generates higher value from content. It is interesting to observe that the optimal network size also increases with coefficient L , which measures the level of tolerance towards delay disutility. A larger L indicates that users are less tolerant to high delay, but at the same time are more tolerant when delay is low (it occurs for this set of parameter values). The opposite of this pattern has been observed as well when delay is high (for example when τ is large).




Fig. 8. Effect of λ and μ_0 on optimal size.




Fig. 9. Effect of γ and β on optimal size.

It is intuitive that the optimal community size n^* increases with the capacities of supernode and peer nodes, but decreases with content request frequency λ (Figure 8). Figure 9 shows n^* decreasing with the sharing level of P2P users. This is because, when users are more willing to share their contents, a smaller number of P2P users are required to achieve the equivalent content availability level. Higher content intensity (the average number of local cached contents) β also results in smaller n^* when γ is high. As investigated in Section IV.B, if content sharing level is low ($\gamma < 0.2$), higher β induces less provision delay, and therefore a larger P2P community can be operated.

Figure 10 shows that the optimal size increases with the variation degree of content popularity α and the number of content varieties M , when α is high (greater than 0.2). The community becomes larger with M when α is less than that

critical value. This observation is consistent with the previous findings of the property of the critical value α^* .




Fig. 10. Effect of α and M on optimal size.

Table IV summarizes the direction of the changes in optimal community size n^* with respect to all the system parameters.

α	β	γ	λ	M	μ_j	μ_0	V	τ
+/-	+/-	-	-	+/-	+	+	+	+

“+” stands for positive effect

“-” stands for negative effect

“+/-” stands for positive effect when n is small and negative effect when n is large

VI. INTERCONNECTION AND GROUPING DECISION

Results from the above performance analysis indicate the size of a P2P network is mainly bounded by the search capacity of the supernode, and a larger network would also result in a higher provision delay. To improve search and consequent overall performances, one can invest to expand the capacity of a centrally operated supernode. However, in reality, search performance is usually enhanced by leveraging multiple decentralized supernodes. Assuming that each user is connected to only one supernode, a larger number of supernodes will segment the entire population of users into groups with smaller sizes. Small group size reduces search delay, but at the same time, lowers content availability and degrades transmission performance. To achieve certain quality level of content availability, unsatisfied content requests are forwarded to interconnected supernodes. Certainly, this will impose extra load on interconnected supernodes. It would be interesting and important to examine various grouping approaches so as to identify the best operational performance.

The performances of P2P networks with multiple supernodes are strongly associated with interconnection structures. To get managerial insights, we analyze several specific symmetric structures. Similar to regular peer nodes, we assume that supernodes are uniformly located in the same domain, with identical capacity. Each node is assumed to be connected to one of the supernodes randomly, such that each supernode will serve equal number of nodes. In the next

section, we will discuss the situation where nodes are connected to the closest supernode.

A. Interconnection Structure

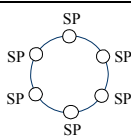
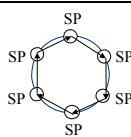
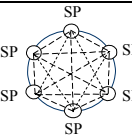
The performance of a network depends on its interconnection structure. Smaller degree of interconnection renders a larger transmission delay, and also requires higher TTL (time-to-live, or number of hops on supernodes each request is allowed) to achieve an equivalent content availability level. Higher degree of interconnection improves the performance of transmission, but at the cost of higher search load imposed on interconnected supernodes. To illustrate interconnection performance, we analyze the operations of three specific symmetric network structures, as depicted in Table V.

Isolated Structure. There is no interconnection among the supernodes. To alleviate search delay, P2P users are partitioned into several isolated groups. The out-degree for each supernode is 0, as is the TTL (time-to live). Isolated structure has best search delay, but the worst hit rate and transmission performance.

Chained Structure. All supernodes are connected, but each supernode can forward requests to exactly one other supernode. Chained structure achieves full content availability. Content requests are searched in one group and, if not satisfied, forwarded to the next. The transmission performance is the same as that of the isolated structure. The out-degree for each supernode is 1, while the TTL is the total number of supernodes deployed less 1.

Complete Structure. All supernodes are fully connected with each other. Content requests are forwarded to *all* interconnected groups at the same time if they are not satisfied by the current group. The complete structure has the best transmission performance, but the highest search load on other supernodes. The out-degree for each supernode is the total number of supernodes less 1, and the TTL is 1.

TABLE V
SYMMETRIC P2P NETWORK STRUCTURES

	ISOLATED	CHAINED	COMPLETE
			
OUT DEGREE	0	1	No. of Supernodes - 1
TTL	0	No. of Supernodes - 1	1

B. Performance Analysis

Utilizing the analytical metrics presented earlier, we formulate the metrics for evaluating the performance of P2P network with multiple supernodes, and investigate the scale effect of grouping with various interconnection structures.

The analytical formulations are summarized in Table VI, where s is the total number of supernodes adopted. It is straightforward to show that the chained and complete structures have the same content availability and provision delay. The transmission delays for chained and isolated structures are also identical.

TABLE VI
PERFORMANCES OF P2P NETWORK WITH MULTIPLE SUPERNODES

	$\bar{H}(n, s)$	$\bar{T}(n, s)$	$\bar{D}(n, s)$	$\bar{S}(n, s)$
ISOLATED	$H(n/s)$	$T(n/s)$	$D(n/s)$	$S(n/s)$
CHAINED	$H(n)$	$T(n/s)$	$D(n)$	$\bar{S}_{\text{chained}}(n, s)$
COMPLETE	$H(n)$	$\bar{T}_{\text{complete}}(n, s)$	$D(n)$	$\bar{S}_{\text{complete}}(n, s)$

In Table VI, the expected search delay for chained structure is

$$\bar{S}_{\text{chained}}(n, s) = \begin{cases} \frac{1 - (1 - H(n/s))^s}{H(n/s)(\mu_0 - (2 - H(n/s))(n/s)\lambda)}, & \text{if } s > 1; \\ \frac{1}{\mu_0 - n \cdot \lambda}, & \text{if } s = 1. \end{cases}$$

Here, each group has a probability of $1 - H(n/s)$ to forward a request to its adjacent group. This will give each group an overall request rate of $(1 + (1 - H(n/s)))(n/s)\lambda$. One request has a probability of $1 - (1 - H(n/s))^s$ to be satisfied. If satisfied, on average, one request is forwarded $1/H(n/s)$ times.

Similarly, for complete structure, the expected search delay can be written as

$$\bar{S}_{\text{complete}}(n, s) = \frac{2 - H(n/s)}{\mu_0 - (1 + (s - 1) \cdot (1 - H(n/s)))(n/s) \cdot \lambda},$$

where $(1 + (s - 1) \cdot (1 - H(n/s)))(n/s) \cdot \lambda$ is the total request rate at each group, and the expected number of search activity is $1 + (1 - H(n/s))$. The expected transmission delay is

$$\bar{T}_{\text{complete}}(n, s) = (H(n/s) \cdot T(n/s) + (H(n) - H(n/s)) \cdot T(n - n/s)) / H(n).$$

Given a satisfied content request, it has a probability of $H(n/s)$ to be downloaded from the local group with an expected transmission delay of $T(n/s)$ and a probability of $(H(n) - H(n/s))$ to be downloaded from the interconnected groups with an expected transmission delay of $T(n - n/s)$.

Numerical Results. We numerically evaluate the impacts of the number of groups, s , on search and transmission performance with various network structures. The parameter values are set as $\alpha = 0$, $\beta = 1$, $M = 10$, $n = 100$, $\lambda = 1$, $\mu_j = 10$, $\mu_0 = 100$, and $\tau = 1$. Figure 11 shows that chained structure

always outperforms complete structure in search, but complete structure always has better transmission performance than chained structure. Interestingly, the expected search delay is a convex function of number of groups for both chained and complete structures. Therefore, there exists an optimal number of groups that would minimize the expected search delay in these two structures. Regarding the expected transmission delay, as expected, small number of groups is preferred for all structures, though for complete structure large number of groups may perform equally well.

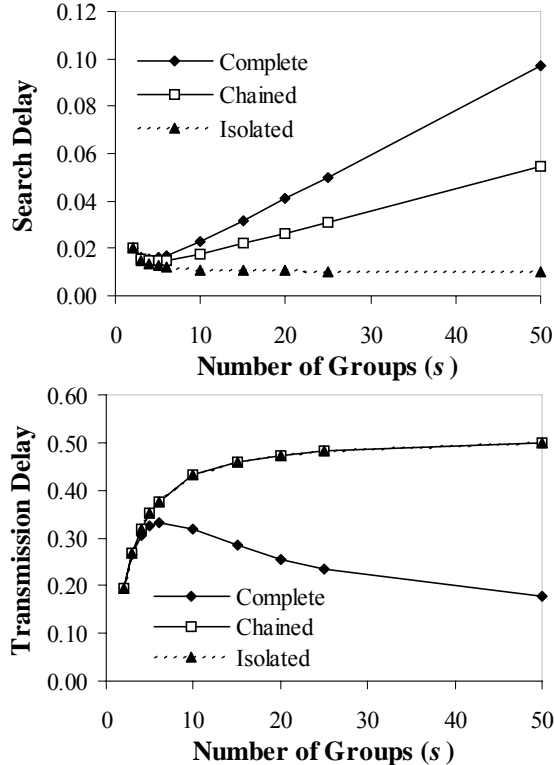


Fig. 11. The impacts of number of groups on expected search and transmission delays.

C. Location Based Grouping (LBG) vs. Random Based Grouping (RBG)

In the previous section, we assume that each peer node is randomly connected to one of the supernodes. This policy is adopted by most of the current P2P technologies, such as Gnutella v0.4 protocol. Network locality-logical distance (latency) techniques, such as the Global Network Positioning (GNP) approach [16] and Lighthouses [22], make it possible to infuse supernodes with locality awareness; and hence better grouping decisions can be made. Here, we analyze a scheme where each peer node is served by the *closest* supernode. This policy can be applied, for instance, in the case where peer nodes are served by a supernode provided by their local ISPs. Intuitively, the advantage of this scheme is a significant improvement of transmission performance as files are downloaded from nearby peer nodes. It is useful to quantitatively examine the impacts of the number of groups on transmission performances of random based grouping (RBG) and location based grouping (LBG). There may be better

performing, than LBG, grouping schemes, however in this paper, we do not attempt to identify them.

Many types of coordinate spaces can be utilized to represent the locality of a peer node and corresponding proximity. To compare the impacts of RBG and LBG grouping technologies on transmission performance with respect to various interconnection structures, we consider a specifically simple circular space with radius $\rho = \tau/\pi$ to model the proximity of the peer nodes, in which an arc between two points in the circle represents the proximity degree (in terms of expected transmission latency) between two peer nodes. A higher ρ indicates greater dispersion of the users. The transmission delay between a content request node and a provision node is an *i.i.d.* random variable uniformly distributed on $[0, \tau]$ when peer nodes are uniformly distributed on the circular space. We assume that there are multiple supernodes, i.e. $s \geq 2$, serving the entire P2P user population. A request node (located at position X) and a provision node (located in position Y) are located in separate groups that are an h -hop distance apart. X and Y are random variables drawn from uniform distributions, $X \sim U(0, 2\rho\pi/s)$ and $Y \sim U(2\rho\pi h/s, 2\rho\pi(h+1)/s)$. If $h = 0$, both nodes are served by the same supernode. The case where $h = 1$ indicates that these two nodes are served by different groups with direct interconnection.

Let $\hat{F}(t|h, s)$ denote the CDF for the expected transmission delay between any two nodes in various groups with h -hop distance. $\hat{F}(t|h, s)$ is defined as $P\{|X - Y| \leq t\}$. Using a convolution of distributions, we have the resulting PDF for transmission delay $\hat{f}(t|h, s)$ as,

$$\hat{f}(t|h=0, s \geq 2) = \frac{s}{\rho\pi} - \frac{s^2 t}{2\rho^2 \pi^2}, \quad 0 \leq t \leq 2\rho\pi/s;$$

$$\hat{f}(t|h \geq 1, s \geq 2) = \begin{cases} 0, & \text{if } 0 \leq t < 2\rho\pi(h-1)/s; \\ \frac{(1-h)s}{2\rho\pi} + \frac{s^2 t}{4\rho^2 \pi^2}, & \text{if } 2\rho\pi(h-1)/s \leq t < 2\rho\pi h/s; \\ \frac{(1+h)s}{2\rho\pi} - \frac{s^2 \cdot t}{4\rho^2 \pi^2}, & \text{if } 2\rho\pi h/s \leq t < 2\rho\pi(h+1)/s; \\ 0, & \text{if } t \geq 2\rho\pi(h+1)/s. \end{cases}$$

The transmission performance (or the expected minimum transmission delay) of P2P network is determined by number of candidate provision nodes. Given that k nodes share the file requested, the minimum transmission delay (using k^{th} order statistics) of LBG is $2\rho\pi/s(2k+1)$, which is strictly superior to the performance achieved in RBG, $\rho\pi/(k+1)$. When the total number of P2P users is sufficiently large, the number of provision nodes can be approximated as $k \approx n \cdot p/s$, where p is the probability that a peer node user share the requested file. The expected minimum transmission delays become

$\rho\pi/(np+s/2)$ for LBG, and $\rho\pi/(np/s+1)$ for RBG respectively. Given a fixed number of P2P users, it can be observed that the transmission performance of RBG worsens as the number of groups increases, while the performance of LBG improves with number of groups.

Next, we numerically investigate the impacts of number of groups on the transmission delay for various interconnection structures, using the same parameter value setting in Section VI.B. Figure 12 shows that, consistent with intuition, the performance of LBG is better than that of RBG, particularly as the number of groups increase. The transmission delays of RBG increase with number of groups for isolated and chained structures. On the contrary, the transmission delay decreases with number of groups when LBG is adopted. This decrease is more significant for isolated and chained structures.

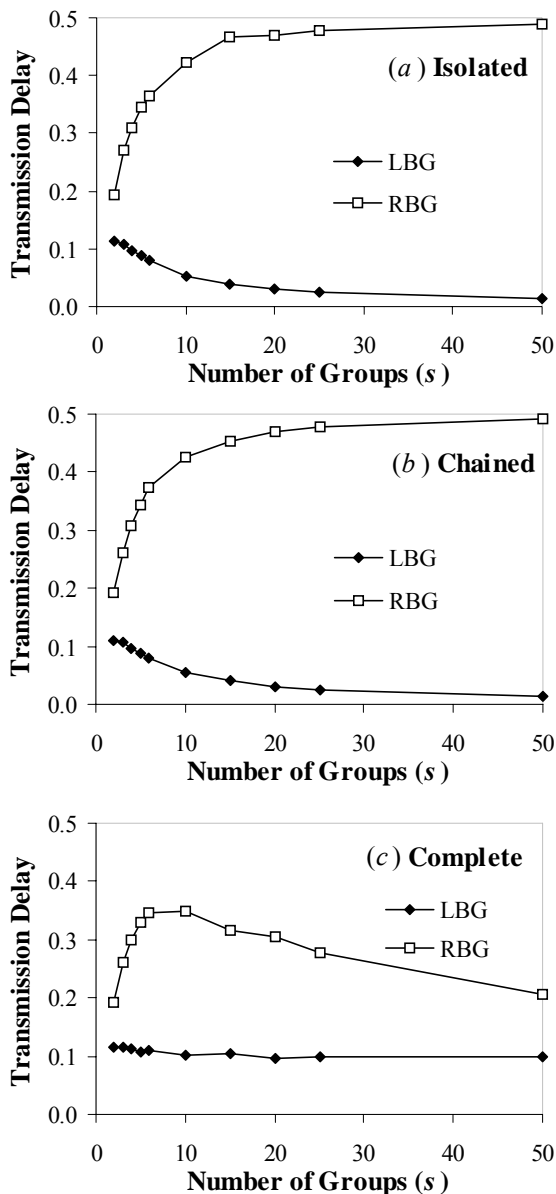


Fig. 12. The impacts of number of groups and LBG on transmission delay.

VII. SUMMARY AND CONCLUSIONS

In this paper, we have developed an analytical model to evaluate the impacts of network scale and system parameters on the performance of P2P networks. Our analytical and simulation results show positive scale effects on content availability and transmission delay, and negative scale effects on provision and search delays. Furthermore, content availability, transmission delay, and provision delay are found to marginally decrease, while search delay is marginally increasing. This suggests that search congestion, rather than provision, is the primary factor that restricts network scalability. Balancing these performance measures, an optimal size can be determined to maximize the overall utility of a P2P community. The optimal community size increases with P2P participant capacity, ratio of content value to waiting cost, as well as location dispersion degree of P2P users, but decreases with content request frequency, and content sharing level of P2P users. Higher degree of variation for content popularity distribution can result in either larger or smaller community size, depending on the number of content varieties.

We further extend the analytical model to examine grouping decision in networks with symmetric interconnection structures. Comparing the performances between chained and complete structures, we find that chained structure has better search performance, but complete structure provides better transmission performance. There exists optimal grouping size that minimizes search delay for both chained and complete structures. The transmission delay in chained structure is found to increase with number of groups. Finally, we compare the impacts of location based grouping (LBG) and random based grouping (RBG) approaches on transmission delay. LBG has better performance, which even improves with number of groups.

In our model, we assume that nodes (regular peer nodes and supernodes) are symmetric and could be coordinated by the central planner. Investigating emerging P2P structure under heterogeneous players (peers and super peers) with incentive compatibility is a planned future extension. It would be interesting to study how our results change if the players of a P2P network are rational to decide sharing, grouping, and interconnection decisions. In this paper, performance metrics are developed from a central operational perspective. It would be important to investigate performance loss due to decentralized decision.

Besides the investigation of factors that affect P2P networks operations, an interesting topic for future research is to study the evolution dynamics of content distribution among peer nodes. As peer nodes exchange contents, popularity of various files changes. It would be important and interesting to study how these changes occur over time, and what impact this may have on P2P network performance.

APPENDIX

A. Proof of Proposition 1

For the expected content availability in a network with n

nodes, $H(n)$, we have,

$$\Delta_n H(n) = H(n) - H(n-1) = \sum_{m=1}^M \theta_m \cdot (1 - \bar{P}_m)^{n-1} \cdot \bar{P}_m > 0;$$

$$\begin{aligned} \Delta_n^2 H(n) &= \Delta_n H(n+1) - \Delta_n H(n) \\ &= -\sum_{m=1}^M \theta_m \cdot (1 - \bar{P}_m)^{n-1} \cdot \bar{P}_m^2 < 0. \end{aligned}$$

B. Proof of Proposition 2

For the expected search delay in a network with n nodes, $S(n)$, we have,

$$\Delta_n S(n) = \frac{\sum_{m=1}^M (1 - P_m) \cdot \theta_m \cdot \lambda}{(\mu_0 - \Lambda_0(n))(\mu_0 - \Lambda_0(n-1))} > 0;$$

$$\Delta_n^2 S(n) = \frac{2 \cdot \left(\sum_{m=1}^M (1 - P_m) \cdot \theta_m \cdot \lambda \right)^2}{(\mu_0 - \Lambda_0(n+1))(\mu_0 - \Lambda_0(n))(\mu_0 - \Lambda_0(n-1))} > 0.$$

C. Proof of Proposition 3

For the expected provision process delay, $D(n)$, we have,

$$\Delta_n D(n) = \frac{\sum_{m=1}^M (1 - P_m) \cdot (1 - \bar{P}_m)^{n-2} \cdot \bar{P}_m \cdot \theta_m \cdot \lambda}{(\mu - \Lambda(n))(\mu - \Lambda(n-1))} > 0;$$

$$\begin{aligned} \Delta_n^2 D(n) &= \sum_{m=1}^M \frac{(1 - P_m) \cdot \bar{P}_m \cdot \theta_m \cdot \lambda \cdot (1 - \bar{P}_m)^{n-2}}{(\mu - \Lambda_j(n+1))(\mu - \Lambda_j(n))(\mu - \Lambda_j(n-1))} \\ &\quad \times (-\bar{P}_m \cdot \mu - (1 - \bar{P}_m) \cdot \Lambda_j(n-1) + \Lambda_j(n+1)). \end{aligned}$$

When μ or n is sufficiently large so that $\mu > (\Lambda_j(n+1) - (1 - \bar{P}_m) \cdot \Lambda_j(n-1)) / \bar{P}_m$, $\Delta_n^2 D(n) < 0$.

D. Proof of Proposition 4

$$\Delta_n T(n, m) = \frac{(1 - \bar{P}_m)^{n-2} (n-1)^2 \bar{P}_m^2 - (1 - (1 - \bar{P}_m)^{n-1})^2}{n(n-1)\bar{P}_m (1 - (1 - \bar{P}_m)^{n-1}) (1 - (1 - \bar{P}_m)^{n-2})} \tau.$$

We need to find the sign of the expression, $(1 - \bar{P}_m)^{\frac{n-1}{2}} (n-1)\bar{P}_m - 1 + (1 - \bar{P}_m)^{n-1}$. It is straightforward to show that this expression is a strictly decreasing function of \bar{P}_m , with a maximum value of 0 at $\bar{P}_m = 0$. Since $\bar{P}_m > 0$, $\Delta_n T(n, m) < 0$. Similarly, we can show that $\Delta_n^2 T(n, m) > 0$.

E. Proof of Proposition 5

1) With respect to β

$$\partial H(n, m) / \partial \beta = n \cdot (1 - \bar{P}_m)^{n-1} \cdot \theta_m \cdot \gamma > 0;$$

$$\frac{\partial T_{\min}(n, m)}{\partial \beta} = - \left(\frac{\theta_m \gamma}{n \bar{P}_m^2} + \frac{n-1}{n} \left(\frac{(n-1) \theta_m \gamma (1 - \bar{P}_m)^{n-2}}{(1 - (1 - \bar{P}_m)^{n-1})^2} \right) \right) \tau < 0;$$

$$\begin{aligned} \frac{\partial \Lambda_j(n, m)}{\partial \beta} \\ &= -\theta_m \left(1 - (1 - \bar{P}_m)^{n-1} \right) + (1 - P_m)(n-1)(1 - \bar{P}_m)^{n-2} \theta_m \gamma. \end{aligned}$$

If $-\left(1 - (1 - \bar{P}_m)^{n-1}\right) + (1 - P_m) \cdot (n-1) \cdot (1 - \bar{P}_m)^{n-2} \cdot \gamma < 0$, which is satisfied when n is sufficiently large or γ sufficiently small, we have,

$$\partial D(n, m) / \partial \beta = (\partial \Lambda_j(n, m) / \partial \beta) (\mu_j - \Lambda_j(n, m))^{-2} < 0$$

$$\partial \Lambda_0 / \partial \beta = -\theta_m^2 \cdot n \cdot \lambda < 0$$

$$\partial S / \partial \beta = (\partial \Lambda_0 / \partial \beta) (\mu_0 - \Lambda_0)^{-2} < 0.$$

2) With respect to γ

$$\partial H(n, m) / \partial \gamma = n \cdot (1 - \bar{P}_m)^{n-1} \cdot P_m > 0;$$

$$\frac{\partial T_{\min}(n, m)}{\partial \gamma} = - \left(\frac{P_m}{n \bar{P}_m^2} + \frac{n-1}{n} \left(\frac{(n-1) P_m (1 - \bar{P}_m)^{n-2}}{(1 - (1 - \bar{P}_m)^{n-1})^2} \right) \right) \tau < 0;$$

$$\partial \Lambda_j(n, m) / \partial \gamma = (1 - P_m) \cdot \lambda_m \cdot (n-1) \cdot (1 - \bar{P}_m)^{n-2} \cdot P_m > 0;$$

$$\partial D(n, m) / \partial \gamma = (\partial \Lambda_j(n, m) / \partial \gamma) (\mu_j - \Lambda_j(n, m))^{-2} > 0.$$

3) With respect to λ

$$\partial \Lambda_j(n, m) / \partial \lambda = (1 - P_m) \cdot (1 - (1 - \bar{P}_m)^{n-1}) > 0,$$

Hence

$$\partial D(n, m) / \partial \lambda = (\partial \Lambda_j(n, m) / \partial \lambda) (\mu_j - \Lambda_j(n, m))^{-2} > 0.$$

It is straightforward to show that

$$\partial \Lambda_0 / \partial \lambda = \sum_{m=1}^M \theta_m \cdot (1 - P_m) \cdot n > 0, \text{ AND } \partial S / \partial \lambda > 0.$$

4) With respect to μ

$$\partial D(n, m) / \partial \mu_j = -(\mu_j - \Lambda_j(n, m))^{-2} < 0.$$

ACKNOWLEDGMENT

The authors thank Professor John Byers (the Associate Editor) and the two anonymous reviewers for their insightful comments and helpful suggestions. Y.T. acknowledges with gratitude the research support by a grant from the Ford Motor Company and by the Neal and Jan Dempsey Fellowship.

REFERENCES

- [1] Adar, E. and B. A. Huberman, "Free-riding on Gnutella," *First Monday* 5(10).
- [2] Asvanund, A., K. Clay, R. Krishnan, and M. D. Smith, "An Empirical Analysis of Network Externalities in P2P Music-Sharing Networks,"

- 23rd Annual International Conference on Information Systems (ICIS'02), Barcelona, Spain, December 15-19, 2002.
- [3] Asvanund, A., S. Bagla, M. H. Kapadia, R. Krishnan, M. D. Smith, and R. Telang, "Intelligent Club Management in Peer-to-Peer Networks," *Workshop on Economics of Peer-to-Peer Systems*, June 5-6, Berkeley, CA, 2003.
- [4] Butler, B. S. "Membership Size, Communication Activity, and Sustainability: A resource- Based Model of Online Social Structure," *Information Systems Research*, Vol. 12, No. 4, December 2001.
- [5] Chu, J., K. Labonte, and B. Levine. "Availability and Locality Measurements of Peer-to-Peer File Systems," *Proceedings of SPIE ITCOM*, 2004.
- [6] Christin, N. and J. Chuang, "On the Cost of Participating in a Peer-to-Peer Networks," *IPTPS*, 2004.
- [7] Dellarocas, C. and P. Resnick, "Online Reputation Mechanisms: A Roadmap for Future Research," *Workshop on Economics of Peer-to-Peer Systems*, June 2003.
- [8] Deng, S. "Empirical model of WWW document arrivals at access links," *Proceedings of the IEEE International Conference on Communication*, June 1996.
- [9] Golle, P., K. Leyton-Brown, I. Mironov, and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," *Proceedings of ACM Conference on Electronic Commerce*, 2001.
- [10] Gummadi, K. P., R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan, "Measure, Modeling, and Analysis of a Peer-to-Peer File-Sharing Workload," *Proceeding of the 19th ACM Symposium on Operating Systems Principles (SOSP-19)*, Bolton Landing, New York, October 2003.
- [11] Jackson, M. O. and A. Wolinsky, "A strategic model of social and economic networks," *Journal of Economic Theory*, 71:44-74, 1996.
- [12] Khambatti, M., P. Dasgupta, and K. Ryu, "Peer-to-Peer Communities: Formation and Discovery," *Fourteenth IASTED International Conference on Parallel and Distributed Computing and Systems*, Cambridge, 2002.
- [13] Krishnan, R., M. D. Smith, Z. Tang, and R. Telang, "The Virtual Commons: Why Free riding Can Be Tolerated in File Sharing Networks," *23rd Annual International Conference on Information Systems (ICIS'02)*, Dec 15-19, 2002.
- [14] Krishnan, R., M. D. Smith, Z. Tang, and R. Telang, "The Impact of Free-Riding on Peer-to-Peer Networks," *Proceedings of the 37th Hawaii International Conference on System Sciences*, 2004.
- [15] Kung, H. T. and C. H. Wu, "Differentiated Admission for Peer-to-Peer Systems: Incentivizing Peers to Contribute Their Resources," *Workshop on Economics of Peer-to-Peer Systems*, June 5-6, Berkeley, CA, 2003.
- [16] Ng, T. S. E. and H. Zhang, "Predicting Internet Network Distance with Coordinates-Based Approaches," *Proc. of the 21th IEEE INFORM INFOCOM*, 2002, pp170-179.
- [17] Ledlie, J., J. Taylor, L. Serban, and M. Seltzer, "Self-organization in Peer-to-Peer Systems," *10th EW SIGOPS*, September 2002.
- [18] Lee, J., "An End-User Perspective on File-Sharing Systems," *Communications of the ACM*, February 2003, Vol. 46, No.2.
- [19] Liebowitz, S., "Will MP3 downloads annihilate the Record Industry? The Evidence so Far," *Advances in the study of Entrepreneurship, Innovation, and Economics Growth*, JAI Press, New York, 2003.
- [20] Oberholzer, F. and K. Strumpf, "P2P Impact on Record Music Sales," *p2pecon*, 2004.
- [21] Paxson, V. and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE /ACM Trans. on Networking*, 3(3):226-244, 1995.
- [22] Pias, M., J.Crowcroft, S. Wilbur, T. Harris, and S. Bhatti, "Lighthouses for Scalable Distributed Location," *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems*, 2003.
- [23] Qiao, Y., D. Lu, F. E. Bustamante, P. A. Dinda, "Looking at the Server Side of Peer-to-Peer Systems," *Proceedings of the 7th workshop on Workshop on languages, compilers, and run-time support for scalable systems*, Houston, Texas, 2004, pp1 - 8.
- [24] Ranganathan, K., M. Ripeanu, A. Sarin, and I. Foster, "To Share or Not to Share: An Analysis of Incentives to Contribute in Collaborative File Sharing Environments," *Workshop on Economics of Peer-to-Peer Systems*, June 5-6, Berkeley, CA, 2003.
- [25] Ratnasamy, S., P. Francis, M. Handley, R. Karp, and S. Shenker, "A scalable content-addressable network," *Proceeding of ACM SIGCOMM*, August 2001.
- [26] Ross, S. M., *Stochastic Processes*, Second Edition, John Wiley & Sons, Inc., 1996.
- [27] Schoder, D. and K. Fischbach, "Peer-to-Peer Prospects," *Communications of the ACM*, Feb. 2003, Vol. 46, No. 2.
- [28] Schollmeier, R., "A Definition of Peer-to-Peer Networking for the Classification of Peer-to-Peer Architectures and Applications," *Proceedings of the IEEE*, 2001.
- [29] Shneidman, J. and D. C. Parkes, "Rationality and Self-interest in Peer to Peer Networks," *Proceedings of the Second International Workshop on Peer-to-Peer Systems (IPTPS'03)*, 2003.
- [30] Singh, M. P., B. Yu, and M. Venkatraman, "Community-Based Service Location," *Communications of the ACM*, Vol. 44 No.4, April 2001.
- [31] Singh, S., S. Ramabhadran, F. Baboescu, and A. C. Snoeren, "The Case for Service Provider Deployment of Supernodes in Peer-to-Peer Networks," *Workshop on Economics of Peer-to-Peer Systems*, June 5-6, Berkeley, CA, 2003.
- [32] Sripanidkulchai, K., "The Popularity of Gnutella Queries and Its Implication on Scalability," Featured on O'Reilly's www.openp2p.com website, February 2001.
- [33] Stoica, I., R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications," *Proceeding of ACM SIGCOMM*, August 2001.
- [34] Yang, B. and H. Garcia-Molina, "Improving Efficiency of P2P Search," *Proceedings of the 28th conference on Distributed Computing Systems*, July 2002.
- [35] Yang, B. and H. Garcia-Molina, "Designing a Supernode Network," *Proceedings of the 19th International Conference on Data Engineering (ICDE)*, Bangalore, India, 2003.
- [36] Yang, B. and H. Garcia-Molina, "Improving Search in Peer-to-Peer Systems," *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS)*, Vienna, Austria, July 2002.
- [37] Wang Y. and J. Vassileva, "Trust and Reputation Model in Peer-to-Peer Networks," *Third International Conference on Peer-to-Peer Computing (P2P'03)*, 2003.



Yung-Ming Li is an Assistant Professor at the Institute of Information Management, National Chiao Tung University in Taiwan. He received his B.S. degree in Computer and Information Science from the National Chiao-Tung University, M.S. degree in Computer Science from the University of Southern California, and Ph.D. in Information Systems from the University of Washington. His research interests include peer-to-peer networks, Internet economics, and business intelligence.



Yong Tan is an Associate Professor of Information Systems and Neal and Jan Dempsey Faculty Fellow at the University of Washington Business School. His research interests include electronic commerce, social networks, software engineering, and economics of information systems. He has published in *Operations Research*, *Management Science*, *INFORMS Journal on Computing*, *IEEE Transactions on Software Engineering*, *IEEE Transactions on Knowledge and Data Engineering*, and *IIE Transactions*.



Yong-Pin Zhou is an Assistant Professor at the University of Washington Business School. He received his Ph.D. in Managerial Science and Applied Economics from the University of Pennsylvania. His research focuses on service operations management, particularly the design and allocation of the capacity of stochastic service processes. His research has appeared in *Operations Research*, *Marketing Science*, *Management Science*, *Manufacturing & Service Operations Management*, and *Queueing Systems*.