Intelligent Club Management in Peer-to-Peer Networks

Atip Asvanund*, Sarvesh Bagla**, Munjal H. Kapadia**, Ramayya Krishnan*, Michael D. Smith*, Rahul Telang*

*Heinz School of Public Policy and Management, CMU

**Information Networking Institute, CMU

{atip, rk2x, mds, rtelang, sbagla, munjalk}@andrew.cmu.edu

May 2003

Abstract

While attention has been paid in the computer science literature to improving efficiency of search algorithms in P2P networks, little attention has been paid to 1) economic incentives that guide users to share content, 2) content-based measures of similarity of interests among users, and 3) cost implications of physical location of peers on the underlying network. Our work draws on ideas from the economics literature of club goods and the information retrieval literature to propose next generation P2P file sharing architectures that rely on content-based, self organizing communities of peers to address these issues. Using the model of ultrapeers and leaf nodes in Gnutella v0.6 architecture as context, we conceptualize an ultrapeer and its network of leaf nodes as a club (in Economic terms). We specify a simple utility-based model for a peer to determine which "clubs" to join, for clubs to manage their membership, and for "clubs" to determine to which other clubs they should be connected. We simulate performance of these models using unique real world dataset collected from the Gnutella v0.6 network. Our preliminary simulations demonstrate that our enhancements result in at least 300% improvement in ultrapeers' ability to satisfy their leaf node's queries, thereby significantly decreasing congestion on the network and enabling more efficient and effective file sharing architectures.

1. MOTIVATION

1.1 Economic Incentives

File sharing networks that use decentralized Peer to Peer (P2P) computing architectures have gained considerable popularity in recent years with the emergence of applications such as Gnutella and Kazaa. New applications (e.g., Groove Networks) are emerging in areas such as knowledge management within large distributed organizations such as the US Department of Defense and global consulting companies. With the growth of these applications, fundamental shortcomings of extant P2P architectures have been widely identified. Of these, the inability to scale or increase the size of the network is one of the major problems. There are two potential reasons for this. One, inefficient search results due to the inability of peers to identify which other peers in the network have relevant content to their queries, and two, the persistence of free riding (i.e., peers that do not provide content to the network) leading to congestion and delay. Recent work on P2P networks has focused on enhancing network scalability and performance through improved indexing schemes, ultrapeers (super nodes), caching, and intelligent linkage promotion based on similarity of interests (Sripanidkulchai et al. 2001; Stoica et al. 2001). In these enhancements, cooperation among peers and their willingness to share their content through replication is taken for granted. However, recent empirical research suggests that cooperation and sharing is not the dominant mode of operation among members of P2P networks, and that this worsens with network size (Adar et al. 2000; Asvanund et al. 2002). Low levels of sharing in peer-to-peer networks limits network scalability and leads to inefficiencies from peers who consume scarce network resources without providing benefits to the network in the form of content, storage, or bandwidth (Krishnan et al. 2002a; Krishnan et al. 2002b).

There are several parallels between the findings in the economic literature and empirical research in the P2P context. The empirical results confirm a prediction of the public goods literature that the private provision of public goods will be socially inefficient either in terms of under-provision (a.k.a. free riding) or over-consumption (a.k.a. "the tragedy of the commons") (Hardin 1968) and that that users have a more incentives to free-ride in larger networks (Olson 1965; Palfrey et al. 1984). In each case, these inefficiencies arise because individuals only take into account their private utility when making consumption and provision decisions even though these decisions affect the utilities of other members on the network. Additionally, P2P networks share many parallels with the club goods literature (Buchanan 1965, Samuelson 1954). P2P networks exhibit characteristics of non-excludability in that access is typically made available to all users of the network. However, excludability could be imposed, as in club goods, through membership rules. In the ideal case, P2P networks also exhibit non-rivalry in demand when a consumer of content becomes a provider of the content, scaling supply to proportionally meet demand (Asvanund et al. 2002). However, in the presence of free riding (when a consumer does not share the files they download), P2P network resources will exhibit rivalry in consumption (i.e., congestion), as in club goods.

Recent work (Golle et al. 2001) attempts to use micropayments as a solution to the free riding problem in P2P networks. While payment has the desirable effect of causing users to internalize the externality they impose on the community, they are impractical in common P2P settings. We explore an alternative approach and make access to club content contingent on content or other resources made available by the user (leaf node) to the community (ultrapeer network). In this way, our approach aligns the incentives of the user with the net utility to the community.

It is important to note that P2P networks differ from the traditional club goods model in the economics literature in several respects. First, club members in P2P networks contribute shared resources rather than monetary payment. Second, a consumer of content also becomes its provider. Third, members may choose to simultaneously connect to more than one club; and finally, clubs can have inter-club relationships as in (Sterbetz 1992) through ultrapeer to ultrapeer connections. We incorporate these differences into both our model and our simulation.

1.2 Sociological and Physical Context

Users can be sociologically characterized by their interest for content. Users with dissimilar interest may inadequately respond to each other's information needs. Recent work (Sripanidkulchai et al. 2001) introduces intelligent linkage promotion. However, it also does not address the content identifiability issue and assumes unique identifiers for content. Unlike books, content in P2P networks do not have unique ID, such as ISBN, and uniquely identifying content is nontrivial. There is the need to quantify similarity in interest – an important metric to determine the capacity of either a single peer or a network of peers (i.e., a club) to satisfactorily respond to queries – and use it to promote the formation of communities of interest. In this paper, we use techniques derived from the Information Retrieval literature to develop such similarity metrics, and these methods do not require content indentifiablity.

P2P networks are an overlay network on top of existing Internet infrastructure (Ripeanu et al. 2002). Users are distributed on several physical networks, and communicating across network boundary is inefficient and costly. A study (Ripeanu et al. 2002) found that most Gnutella users establish connections across network boundaries, which does not take advantage of the underlying infrastructure, and that this results in inefficient routing. A report (Sandvine Incorporated 2002) states that P2P traffic takes up almost 60% of cross network traffic. Thus, a solution should also take into account the cost implications of routing, relaying and satisfying queries over the underlying physical infrastructure.

1.3 Approach

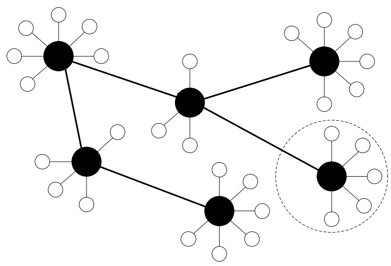

Figure 1: Gnutella v0.6 architecture

We use the Gnutella v0.6 protocol to provide context to our discussion. Gnutella is currently the most popular open protocol (do we have any statistics here?). In Gnutella v0.4, the first public version of Gnutella, all peers are treated equal. Each peer establishes and maintains stateful connections to a few (average of 5) other peers (neighbors), who also recursively establish and maintain connections to other peers. These connections are used for relaying query messages. When a peer issues a query requesting some content, he sends the query message through the connections to all of his neighbors, who in turn recursively forward this query to their neighbors. Query hits then traverse the same path back to the originator. This protocol has proven inefficient because the number of query messages passed overwhelms peers with slower connections. Gnutella v0.6, introduces the concept of ultrapeers. Ultrapeers act as a local hub for its leaf nodes. Leaf nodes are regular peers who maintain stateful connections to their ultrapeers. Ultrapeers are connected to each other in the same fashion as regular peers are connected in v0.4. Ultrapeers forward queries for their leaf nodes and shield their leaf nodes from receiving unnecessary query messages. By knowing the content hashes of their leaf nodes, ultrapeers can determine whether or not to forward a query to a leaf node. By design, peers with high system resources and bandwidth can volunteer to become an ultrapeer, whereas peers with low system resources and bandwidth can remain leaf nodes. Figure 1 illustrates a Gnutella v0.6 topology. Ultrapeers are depicted by dark circles, while leaf nodes are depicted by light circles. Ultrapeer interconnectivity is depicted by thick lines, while leaf nodes' connections to ultrapeers are depicted by thin lines.

We conceptualize the Gnutella network as a collection of clubs operated by ultrapeers who seek to maximize their club value, while leaf nodes (i.e., peers) connect to the right clubs in order to maximize their private utility. Figure 1 depicts a club by a dotted circle, which represents an ultrapeer and its connected leaf nodes. We leverage the ultrapeer architecture to introduce club formation and employ exclusivity to optimize allocation of rivalrous (i.e., shared) club resources. This framework, which is outlined in Section 2 below, imposes an incentive reinforcing structure on the network that encourages users to form the club efficiently, and improve cooperation through sharing and replication.

We use this approach to simulate the performance of P2P networks under a variety of different informational conditions (facilitated through augmentation of the existing protocols). This simulation will be seeded using a unique data set that we have compiled documenting user content, queries, geographical location, and the correlation between the three. Such correlation is critical to accurately simulate club formation. While, we chose Gnutella as a context for our discussion because it was the largest P2P network, and we could obtain real world data from it, our club model is applicable to any decentralized P2P architectures that can take advantage of the ultrapeer structure.

2. MODEL

As noted above, our goal is to treat network participants as economic actors by modeling the utility functions of network participants and, based on these utility functions, to develop a set of incentives such that the participants will choose actions to maximize social welfare. Our model development follows the standard game theoretic specification of each participant's information set, strategy set, and utility function. Following the Gnutella v0.6 architecture, our model participants are leaf nodes and ultrapeers. The information and decision sets we use can be achieved with only minor extensions to the existing Gnutella v0.6 draft protocol.

2.1 Leaf Nodes

Information Set: Each leaf node is provided with an initial set of content. Each leaf node also has its endowed bandwidth and geographical location.

Utility Function: The leaf node's net utility is a sum of the utility provided by each of the ultrapeers to which it is connected. For each ultrapeer *up* that a leaf node *l* is connected to, *l* gains utility <u>directly</u> from that *up* and <u>indirectly</u> from other ultrapeers connected to *up*. Specifically, $U_l(up) = U_{l,d}(up) + U_{l,i}(up)$, where $U_{l,d}$ refers to direct utility and $U_{l,i}$ refers to indirect utility.

To compute direct utility, a leaf node needs to compute the utility it obtains from other leaf nodes connected with this ultrapeer. We use a linear utility formulation such that a leaf node gets some utility from every other leaf node connected to that ultrapeer. Let $\alpha_{l,X}$ for $X \in [1, M]$ represents the *M* linear weights that *l* puts on the utility derived from other *M* leaf nodes. Specifically, $U_{l,d}(up) = \sum_{l' \text{ member of } up} \alpha_{l,1} * v_l(l') - \alpha_{l,2} * c_l(l')$, where v_l is the value gained and c_l is the cost incurred by leaf node l due to a leaf node l', which is also connected to ultrapeer *up*. How do we specify v_l? We use the IR literature to specify v_l as, $v_l(l') = sim(q_l, k_{l'}) * |k_{l'}| * (\alpha_{l,3} * b_{l'} + \alpha_{l,4} * d(l, l'))$, where q_l is l's queries, and $k_{l'}$ is the content provided by l'. Function *sim* is a similarity measure based on information retrieval methods. Furthermore, $b_{l'}$ is the bandwidth of l', and d is a distance measure. Function d assigns higher value to nodes from the same backbone than to nodes from other backbones. In summary, we can think of v_l as a function that assigns higher value to leaf nodes who 1) have high probability to satisfy l' s information needs, 2) provide more content, 3) have high bandwidth, and 4) are located on the same backbone. Simlarly, $c_l(l') = sim(q_{l'}, k_l) * |k_l| * (\alpha_{l,5} * b_{l'} - \alpha_{l,6} * d(l, l'))$. In other words, leaf node l' will also impose costs of lead node l. can think of c_l as a function that assigns higher cost to leaf nodes 1) whose information needs l has a high probability of satisfying, 2) who have high bandwidth, and 3) who are located on a different backbone. Indirect utility is similarly defined. In the interest of space we leave the full specification in our full paper.

Strategy Set: Leaf nodes make the following decisions to maximize their private utility: which pieces of content to share, how many and which ultrapeers to attempt to connect to, and when to drop an ultrapeer connection (and search for a new connection).

For example, we consider a simple decision model where each leaf node l chooses CUP_l, the set of ultrapeers that he is directly connected to, and k_l, the content that l provides (in this simple model, we assume that l provides the same content to all clubs that he is connected to). In this setting, utility of l is the sum of $U_l(up)$ for all $up \in CUP_l$. In each period, l discovers a new set of ultrapeers DUP_l. To optimize his utility, l will add an ultrapeer $up \in DUP_l$ to CUP_l, or will disconnect from a connected ultrapeer to connect to *up* if only the change in CUP_l will result in better utility. Furthermore, each ultrapeer may have a specific requirement about the size of k_l that l will have to offer in order to be accepted by the ultrapeer. Therefore, modifying CUP_l may force l to adjust his k_l as well, and this may affect the overall utility. Thus in each period, each leaf node will be performing a local optimization algorithm on the optimal CUP_l and k_l.

2.2 Ultrapeers

Information Set: Ultrapeers know (or has the ability to estimate) the set of shared content for each connected node (where nodes can include either leaf nodes or other ultrapeers), and other attributes mentioned in the previous section. From this information, the ultrapeer can estimate the value and cost each node provides for other nodes in the network, using utility functions outlined in the previous section.

Utility Function: The utility of an ultrapeer is the weighted sum of the expected utilities provided by each of the connected leaf nodes. In our initial specification weights are equal for each member of the group. We plan to extend this to allow for unequal weights to endogenize the decision to become an ultrapeer.

Strategy Set: Ultrapeers decide which nodes they should connect to (both leaf nodes and other ultrapeers) to maximize the total utility of connected leaf nodes. In this version, an ultrapeer is only concerned with the utility of his connected leaf nodes. When an ultrapeer is at capacity, it can choose to accept an additional connection when the expected utility gained from the requesting node is higher than the lowest utility provided by an existing node. The connected node with the lowest expected utility is then dropped to make room for the new node. The threat of being dropped provides the incentives for nodes to behave in a way that is aligned with the network's interests — both in terms of providing content and consuming scarce network resources.

2.3 Similarity measures using Information Retrieval techniques

A significant part of our club model depends on a peer's and a club's ability to determine who can satisfy its information needs. Although assumed in many previous studies, content in P2P networks is not uniquely identifiable. We use established information retrieval methods to determine similarity by comparing frequency of each word occurrences in a given peer's content naming scheme. These methods allow us to produce a quantifiable estimation of similarity that can be directly used in our club model. We experimented with two popular methods: tf-idf cosine and Jensen-Shannon divergence (Dhillon et al. 2002). Tf-idf is based on the vector-based model, while Jensen-Shannon divergence is based on the language model, which has its roots in Information Theory (Cover 1991).

3. SIMULATION

We simulate our model as a computational game in which leaf nodes and ultrapeers jointly attempt to find their optimal provision of content and optimal club membership respectively (Fudenberg et al. 1998; Gibbens et al. 1995; Ross et al. 1989). Following the standard approach, game participants take their "opponents'" strategies as given and the outcome is expected to converge towards a stable equilibrium in content provision and club formation.

Our simulation uses a lightweight extension of the network and transport layers on top of Javasim's event scheduler. We model backbones as separate network clouds, in which there exists routing delay for communicating across backbones. We implement each Gnutella host as an autonomous agent who performs intelligent decisions for connecting to and disconnecting from other hosts. Our current simulator is capable of including up to 20,000 Gnutella hosts – within an order of magnitude of the actual Gnutella network size.

To parameterize our simulation with real world data we collected data from the Gnutella v0.6 between August 31, 2002 and September 29, 2002. In our data we observed 10,533 unique hosts, 28% of which are ultrapeers and 72% of which are leaf nodes. We find that 42% of all hosts do not provide any files. For hosts that provide files, the average provision is 270, and this follows a long right tail distribution. On average, a host issues 13 queries per hour, and this also follows a long right-tail distribution. We find that 16% of all hosts are distributed among the top 2 backbone providers (att.net and atdn.net), while the remaining are evenly distributed among other significant providers. Most significantly, we find that hosts issue queries with significant similarity to their content. This finding is very important for forming communities and it also allows us to adopt the use of Information Retrieval techniques. We use the correlation between the content, queries, and geographical location to accurately parameterize the hosts in our simulation.

	Utility	**Topology Coverage**	**Club Recall**	**Club Provision**
Initial	21.64	0.0037	0.0045	110.55
Enhanced 5 moves	86.45	0.0037	0.0136	193.39
Enhanced 10 moves	92.47	0.0037	0.0180	223.73

Table 1: Preliminary Results

To demonstrate the effectiveness of our method, we perform a preliminary simulation as shown in Table 1. We include 2,000 hosts and simulate the formation of 100 clubs, each of which has a capacity for 10 nodes. These hosts are randomly chosen from our data collection. We use their real content and query data, and we use a simplification of the model that we proposed to derive utility. Each node can be part of only one club, or not be part of any club at all. In each move we randomly choose a node that will have learnt about 10 random clubs. He will attempt to join a new club if his utility is greater than what he gets from his current club. The club, in turn, will only accept him if by accepting him and removing its worst member, the club would be better off. The initial setup assigns leaf nodes randomly to clubs, and this results in the utility of 21.64, topology coverage of 0.0037, club recall of 0.0045 and club provision of 110.55. Club utility is the combined utility of all users in a club, and is an expected value that we wish to maximize. Topology coverage is the percentage of the network that a leaf node can reach from inside a club. Since we do not change club size and restrict search range of a leaf node to its own club, this value remains stable through out our simulation. Club recall is the percentage of total relevant content in the network that is available in a club. Club provision is the average number of files that club members provide. We can see that after each leaf nodes has moved on an average of 10 times, utility improves to 92.47, club recall improves to 0.0180, and club provision improves to 223.73. Please note that in this set up only half of the 2,000 nodes can belong to a club, thus undesirable nodes (e.g., free riding nodes, and nodes providing undesirable content) are kicked out in each iteration. The improvement in club recall can be attributed to both the increased in provision and similarity of content provided.

4. CONCLUSION AND EXTENSIONS

We have preliminarily demonstrated that our model results in significantly enabling more efficient and effective file sharing architectures by taking into account to economic incentives that guide users to share content and content-based measures of similarity of interests among users. Our model can be extended to address other potential efficiency gains in Gnutella networks. For example, endogenizing the choice to become an ultrapeer, providing more refined incentives through differential quality of service levels for members based on their contribution, and the potential for intentionally malicious leaf nodes, or proposing a model in which an ISP can run its own ultrapeer.

5. BIBLIOGRAPHY

Adar, E. and Huberman, B. A. "Free riding on Gnutella," *First Monday 5(10)*, 2000.

Asvanund, A., Clay, K., Krishnan, R. and Smith, M. "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.

Buchanan, J. "An Economic Theory of Clubs," *Economica*, New Series, Vol. 32, No. 125. (Feb., 1965), pp. 1-14, 1965.

Cover, T. M. and Thomas, J. A. *Elements of Information Theory*, Wiley, 1991.

Fudenberg, D. and Levine, D. *The Theory of Learning in Games*, MIT Press, 1998.

Gibbens, R., Kelly, F. and Key, P. "A Decision-theoretic Approach to Call Admission Control in ATM Networks," *IEEE Journal on Selected Areas in Communications*, 13(6):1101--1114, 1995.

Golle, P., Leyton-Brown, K. and Mironov, I. "Incentives for Sharing in Peer-to-Peer Networks," Working Paper, Stanford University, Palo Alto, CA, 2001.

Gnutella 0.6 Protocol. (available online at http://rfc-gnutella.sourceforge.net/draft.txt; accessed October 13, 2002).

Hardin, G. "The Tragedy of the Commons," *Science*, 162 1243-48, 1968.

Helsley, Robert W. and William C. Strange. 1991. Exclusion and the Theory of Clubs. *Canadian Journal of Economics* 24(4) 888-899.

Dhillon, I., Kumar R. and Manella, S. "Information Theoretic Feature Clustering for Text Classification," TextML-2002.

Krishnan, R., Smith, M., Tang, Z. and Telang, R. "The Virtual Commons: Why Free riding Can Be Tolerated in File Sharing Networks," 23rd Annual International Conference on Information Systems (ICIS'02), Barcelona, Spain, December 15-19, 2002a.

Krishnan R., Smith, M., and Telang, R. "The Economics of Peer-To-Peer Networks," Working Paper, Carnegie Mellon University, Pittsburgh, PA, 2002b.

Ripeanu M., Iamnitchi A., and Foster I. "Mapping the Gnutella Network," IEEE Internet Computing, Jan.-Feb. 2002, pp. 50-56.

Olson, M. *The Logic of Collective Action*, Harvard University Press, Cambridge, MA, 1965.

Palfrey, T. and Rosenthal, H. "Participation and the Provision of Discrete Public Goods: A Strategic Analysis," *Journal of Public Economics 24*, 171-193, 1984.

Ross, K. W. and Tsang, D. H. K. "The Stochastic Knapsack Problem," *IEEE Transactions on Communications*, Vol. 37, No. 7, July 1989.

Samuelson, P. A. The Pure Theory of Public Expenditure. *The Review of Economics and Statistics* 36(4) 387-389, 1954.

Sandvine Incorporated. "Peer-to-Peer File Sharing: The Impact of File Sharing on Service Provider Networks," Industry White Paper, 2002.

Sripanidkulchai, K., Maggs, B. and Zhang, H. "Efficient Content Location and Retrieval in Peer-to-Peer Systems by Exploiting Locality in Interests," SIGCOMM'01, 2001.

Sterbetz, F. P. and Sandler T. "Sharing Among Clubs: A Club of Club Theory", Oxford Economic Papers, 44(1), 1-19, 1992.

Stoica, I., Morris, R., Karger, D., Kaashoek, M. F. and Balakrishnan, H. "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications," SIGCOMM'01, 2001.