

PEER-TO-PEER COMMUNITIES: FORMATION AND DISCOVERY*

Mujtaba Khambatti, Kyung Ryu, Partha Dasgupta
Arizona State University
Tempe, AZ 85287-5406
U.S.A.

ABSTRACT

Current peer-to-peer systems are targeted for information sharing, file storage, searching and indexing often using an overlay network. In this paper we expand the scope of peer-to-peer systems to include the concept of “communities”. Communities are like interest groups, modeled after human communities and can overlap. They can also exist without anyone knowing about their existence. Communities are created, implicitly when one or more entities claim an interest in the same topic.

Our work focuses on efficient methods to discover the formation of these self-configuring communities. We investigate the behavior of randomly created communities and model the complexity of discovery algorithms. Discovering communities on the fly is essential to being able to perform community directed searching. In addition, efficient discovery algorithms allow us to manage quickly changing community structures (dynamic communities, failures, mobile nodes and so on).

We use some simulations to discover the architecture of randomly created communities and then perform studies on techniques for discovering communities.

KEY WORDS

Parallel/Distributed, Peer-to-peer, Self-Configuring Networks, Search, Communities.

1. INTRODUCTION

Peer to Peer (P2P) systems are distributed systems in which logically distinct computing elements called *peers*, having comparable roles and responsibilities, communicate information, share or consume services and resources amongst each other. These systems have the potential to harness massive amounts of storage with modest investment and no central authority [1, 2]. The emergence of file sharing applications such as Gnutella [1], Freenet [2], and Napster [3] has been the catalyst that drew a lot of attention to P2P systems.

The building block of the current P2P systems is the notion of a peer-group, or a number of nodes that participate with each other for a common purpose. In this paper, we discuss a generalization of the notion of peer

group to a multiplicity of groups (possibly overlapping) called *peer communities*. While a group is a physical collection of objects, a community is a set of active members, who are involved in sharing, communicating and promoting a common interest.

Our concept of peer communities is loosely based on the idea of “interest groups”, for example Yahoo Groups [4] or Usenet Newsgroups. A node in the system claims to have some interests and depending upon the claims of all the peer nodes, the communities are implicitly formed (made up of peers with the same or similar interests). Note that the groups are formed implicitly (i.e. they are self organizing). If a node in New York declares an interest in wombats, and a node in China also declares this interest, then the two of them become part of an implicit, undiscovered community. As is obvious, a node may belong to many different communities and communities may overlap.

In this paper, we provide a motivation for the study of P2P communities and illustrate some scenarios to define and discover the community structure. Using simulated models of communities, we show how communities can be formed and discovered. The simulation results have provided us with an insight on behavior of random networks.

2. MOTIVATION

Psychologists have long shown that people have an affiliation motive [5] and a need for information about the world around us [6]. These are some of the instincts that have driven the formation of human groups. This tradition of group forming is not alien to computer science and we witness examples of it in Usenet groups, web communities [7], yahoo groups [4], chat rooms and so on.

With the exception of web communities that have been shown to be self-organized [7] and the alt.* Usenet groups; almost all other present-day groups are a result of *a priori* planning and implementation or at the very least they require some central control or a central authority through which advertisements can be made.

In contrast to this, P2P systems are defined to be completely de-centralized and can also be dynamic. This makes them very attractive as system solutions to the “little people”, like home users, small-scale networks, ubiquitous computing environments, dynamic coalitions

* This research is partly funded by grants from AFOSR, DARPA, Microsoft and NSF

[8] and so on, who now will have the ability to choose their own policies, roles, and responsibilities and change them autonomously.

2.1 PEER COMMUNITIES

The popularity of schemes to form communities and associations on the web leads to the use of the P2P system structure for realization of underlying community structures. Thus P2P communities are not only a natural extension for arranging distributed systems, but also to enhance the capabilities of each member.

The current crop of P2P systems focuses on global information sharing, searching and replicated file storage. Much of the research is targeted at creating an overlay network (based on node identifiers) and using the overlay network to perform searches on a content addressable space. In our work we extend the above concepts to allow the notion of communities. Communities are useful in structuring the information storage space, discovering resources and pruning the search space. It also aids in better dissemination of useful information. For example, suppose node X belongs to a person interested in Amazonian Biological Catapults (ABC). After X declares this interest, it becomes a member of the community of ABC enthusiasts. Henceforth, all information X wants to share can be placed in a public directory and will be readable/searchable by all members of ABC. This concept can be extended to discover resources, physical devices, network components. It also has some interesting security and access control issues (including trust management).

The above example is interesting in the context of peer communities with no overlapping interests. To enable cross community information exchange we widen the scope of our design, with another example. Consider a digital library. Each node in the library system owns a set of books that it is willing to share with other nodes¹. The subjects of the books owned form the different communities in this peer architecture. We will discuss this type of system further in the next few subsections.

2.2 P2P SEARCHING

Searching for information is one of the mainstays of P2P systems. Centralized searching (as used by Internet search engines such as Google) has the downside that the central authority controls the indexing and presentation of the information. P2P searching allows anyone to put up information in the search index and then cooperatively search the P2P space. P2P searching techniques include flooding, directed flooding, iterative deepening, directed BFS and so on [9].

Considering the digital library again; if the Computer Science and Medicine communities are disjoint, then

searching for medical information by a node belonging to the Computer Science community would not produce any results. If the communities were linked at some point Q (Q belongs to both communities) then medical information would be found, but at a great search expense, as on the average half the Computer Science community would be searched before the Medical community.

To mitigate such problems, we need community-based query propagation. Thus to provide efficient searching, it is better to target a search for one (or more) target communities, irrespective of the current membership of the searching node.

To enable regular searching as well as community-based searching, we must be able to efficiently discover communities, that is, we need searches for community members. This is somewhat complicated by the fact that communities are implicit, self-organizing, dynamic and constantly changing—forming, or breaking down due to changes in the peers.

Our work is targeted to discovery of communities on the fly to enable efficient intra-community as well as inter-community searching. To this end we have studied the community formation characteristics using simulations of large communities and found characteristics that can be exploited to perform discovery and searching. We show that this discovery does not require extensive computation or communication on the part of a peer.

Note that searching is not the focus of this paper, searching is the next step after community discovery algorithms are implemented.

3. RELATED WORK

Contemporary search algorithms are executed on a server or a group of servers that have access to a data repository on behalf of a requesting client machine. The absence of servers within a P2P environment has brought about a number of proposals for efficient search techniques. Some are highly deterministic but can also be very expensive, like flooding, which requires a search query to be sent to every peer existing. An optimization to flooding, called directed flooding, only sends queries to all peers in a certain part of the P2P system based on some knowledge or history. Additionally, [9] has proposed techniques that reduce communication and increase the probability of arriving at a solution to the search query. Examples of these are iterative deepening, which iteratively increases the depth until which to flood; directed BFS, which selects a subset of immediate neighbors heuristically and entrusts them with further propagation of the search query using directed BFS again; and local indices, which requires each peer to maintain a local index of the contents of peers that exist within a pre-determined range.

¹ Assume these are non-copyrighted works.

A considerable amount of research has focused on the analysis of the link structure in collections of objects. Through these analyses, researchers had hoped to derive the procedures for efficient identification and discovery of patterns in the collection. Early attempts to analyze the collective properties of interacting agents have been found in social networks [10], where link structures like cliques, centroids and diameters were studied. The field of citation analysis [11] and bibliometrics [12] seek to identify patterns in collections of literature documents by using citation links. We give below one such notable definition of web communities that uses an analysis of the link structure of web pages to efficiently identify communities.

“We define a community to be a set of web pages that link (in either direction) to more web pages in the community than to pages outside of the community.” [13]

At first glance, the above definition seems just what is needed to identify P2P communities. However a closer look will indicate that if peers are placed in P2P communities based entirely on link analysis, we will not accomplish our goal of allowing peers to simultaneously be members of more than one P2P community.

Perhaps a combination of graph theory and link analysis is needed to correctly identify patterns within collections of peers. For instance, if the links of a peer were classified, as outgoing and incoming links, and if there existed cycles in the directed graph formed by the peers and their links, we would have successfully identified P2P communities. These cycles could be discovered in a depth-first manner by exhaustively traversing the outgoing links to check if they visit a peer twice.

Except for the obvious scalability issue in the above technique, it seemed to guarantee to uncover any pattern that might exist. Yet the procedure would fail for some of the most commonly occurring patterns: the star, where many peers are huddled around a single peer like in a client-server configuration, or the tree, where no cycles would exist. Even in a domain that is purely P2P, one cannot avoid the fact that often many peers will use a small subset of peers or even a single peer as a source of information or for collaborations.

On the topic of discovery, a particularly remarkable solution was proposed by [14]. It offered an approach called HITS that was related to spectral graph partitioning and methods used by the Google search engine [15]. Their recursive definition of *hub* and *authority* web pages work to rank results by a measure of importance and usefulness, thereby identifying key web sites related to some community and also the related websites that might be members of the same community [16]. However their approach is highly dependent on the link topology and therefore cannot effectively aid in the discovery of

communities that are ring-based without any dominating members.

Complementary to the HITS algorithm, [13] requires a “seed” web site as the starting point to begin a focused crawl in order to identify a community. Members of the community are discovered using the maximum flow / minimum cut framework using the two sets called source and sink, initially composed of well-known web sites. When mapped to the P2P domain, this link-based technique has the drawback of not truthfully modeling real world P2P communities where peers can simultaneously be members of more than one community.

In trying to empower peers to discover their community membership, we had to find an efficient solution that would be practical, unlike the NP-complete solutions offered by graph partitioning [17].

4. P2P NETWORK MODELLING

Formation and discovery of peer communities are significantly dependent on how peers declare and use their common “interests”. First we define *attributes* as a method of declaring interests, and then we use these attributes to discover communities (in a simulated model).

4.1 ATTRIBUTE-BASED COMMUNITES

Peer communities are formed based on common interests. In our model, common interests are represented by *attributes*. These are then used to determine the peer communities that a particular peer would participate in. Attributes can be either explicitly provided by a peer or implicitly discovered from past queries. For example, a housewife can express that she is interested in French wines and house decoration. Such expressions are personal declarations. Also, her repeated web search queries to find “K-12 education in Arizona” can be used to provide implicit information about her interests. There are of course privacy and security concerns in using such information, so we divide interest into three classes – *personal*, *claimed* and *group*.

The full set of attributes for a peer is named *personal attributes*. However, for privacy and/or security reasons, all these attributes may not be used to determine community membership. A client may not want to reveal some of the personal attributes. Hence, a subset of these attributes is explicitly claimed public by a peer. We call these the *claimed attributes*. The claimed attributes are a subset of the personal attributes.

In addition to personal and claimed attributes, we introduce the notion of a *group* attribute. The group attribute is location or affiliation oriented and is needed to form a physical basis for communities. Every node belongs to one pre-determined group and has a group attribute that identifies the node as a member of this

group. For example a computer on the campus of Arizona State University belongs to the ASU group. A home computer using the AOL Internet service provider is part of the AOL group. A computer in the office of an IBM employee is in the IBM group. The domain name of an Internet connection may be used as the group identifier.

The group attribute is also considered a personal attribute and it may or may not be one of the claimed attributes. While it is expected that a node will put the group attribute as part of the claimed attribute set, it does not have to.

Now, we formally define a P2P community based on the attributes of each peer.

P2P COMMUNITY: *The non-empty set N of nodes is a peer-to-peer community iff N has a non-null signature.*

SIGNATURE: *Let n be a node and $claim(n)$ be a set containing attributes claimed by n . Consider a non-empty set N of nodes. Then the set resulting from the intersection of $claim(k)$, for all $k \in N$ is called a signature of the set N .*

With this definition, given any collection of peers, we would be able to tell whether the collection is a P2P community or not.

4.2 P2P NETWORK LINKS

P2P communities are attribute-based, that is, attributes (claimed and group) determine the membership of a node in one or more communities. In addition to attributes we also define an overlay network in terms of “links”.

Links are not really needed to form and manage P2P communities. However, they are needed to feasibly run low-cost algorithms for formation and discovery, as it is conceptually and algorithmically simpler to use the notion of a set of “neighbors” when communicating with other peers. First let us motivate a need for a link.

Suppose a node, belonging to a domain abc.com claims the attribute “baseball”. Now this node is essentially isolated, unless it a priori knows about other members of the baseball community or other members of the abc.com community. There is a need for a “seed” to start the community formation and information search needs.

Flooding and querying some central server are two solutions to the isolation problems, the first is expensive and the second violates the self-configuring tenet of the P2P structure. Hence we propose the use of an overlay network based on links.

When a node X is born, it needs to have one or more logical neighbors. If it has three neighbors, A , B and C , then we say it has three links, $X \rightarrow A$, $X \rightarrow B$ and $X \rightarrow C$. Unlike overlay networks used by some P2P systems, our

link based network is not based on node names, but user selected neighbors.

Now we discuss how these links are created. A node X links to a node A if (i) A is a special node (server) designated by the domain for P2P links, (ii) A is a node, known to X that it trusts (iii) A is a node that belongs to many communities X is interested in. For a novice/new node, (i) may be the most appropriate link. As X ages, it finds other nodes and adding these links improves search speed, information access and such. The linkages are similar to friendships in real life, or http links in the Web and are human directed.

4.3 LINKS AND SMALL WORLD NETWORKS

In the real world, peers form relationships (links) with other peers that have something in common with them. While most of these relationships develop between peers that exist within a common local area network, domain, country or some such location specific factors, a small set of random peers would acquire “long-distance” relationships with peers that are considered remote. The former scenario, termed as *regular connections*, is a direct result of the inclination to form groups as discussed earlier. The other type of connection is described as random because they link peers to other randomly selected peers that are not necessarily within any location proximity of each other. *Random connections* are caused by external human factors; like a company employee who uses her computer to occasionally connect to her bank, while simultaneously participating as a member of the P2P community with other company employees.

Previously, such types of semi-random networks have been described by Watts and Strogatz [18] who studied the properties of large regularly connected graphs of nodes that contain a few random long-distance edges between nodes. They modeled this structure and demonstrated that the path-length between any two nodes of the graph is in fact surprisingly small. As a result, they called such semi-random structures as *small-world networks*.

4.4 MODELLING P2P COMMUNITIES

We now present our modeling (simulation) work on forming and discovering P2P communities. For this work, we first create a set of nodes, assign them attributes and then generate links. After that, we run our algorithms for community formation and community discovery and get results on the efficiency of these algorithms.

First, we create a very large number of peers each having a set of *personal attributes* with at least one element. To simplify the process, the attributes are implemented as different letters of the alphabet. Thus we have 26 possible attributes (A-Z), and for each node we pick between 1 to

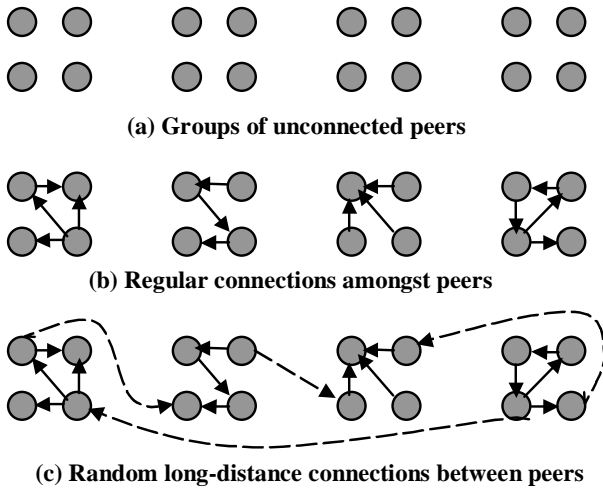


Figure 1: Creation of a P2P small-world network

26 attributes randomly and assign these to the personal attribute set.

Second, from the personal attribute set of each, we randomly select some attributes and make these the claimed attributes of that node.

Third, we randomly divide the set of peers into a set of groups. We assign to each node a group attribute that is same for all nodes belonging to a group.

Finally, we create the links. The link creation is a two-step process; (i) we randomly assign links that are between group members (regular connections) and (ii) we create some lesser number of inter group links (long distance connections).

Thus we generate a set of peers with attributes and links that is representative of a P2P network. We then run further simulations of our algorithms on this network. This process is illustrated in Figure 1.

5. FORMING COMMUNITIES

The community formation algorithm determines the membership of a node in one or more communities.

Using the definitions for the network modeled as described above, it is possible that a set of peers, which are part of the same group, is not identified as a P2P community because all the peers did not claim (make public) their group attribute. The following situation is possible. A P2P community with signature X exists, and a peer that has X in its personal attribute set but has not claimed it, would not be able to join this community and avail of the benefits until it claims X .

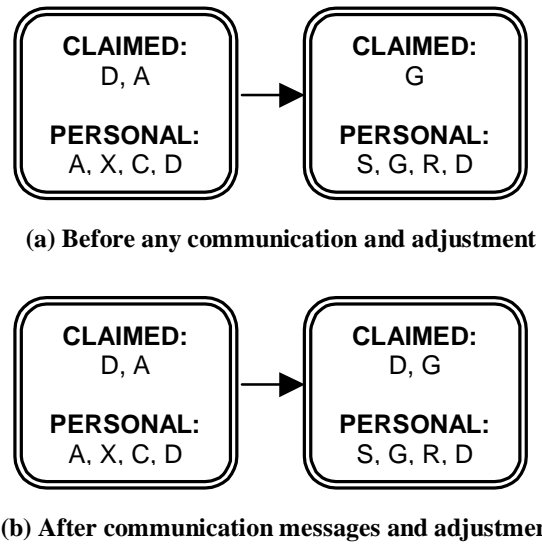


Figure 2: Procedure to escalate attributes (D in this case).
Two peers that are linked (shown by arrow) are considered for the illustration.

Peers therefore need to expose as many attributes as will get them into the maximum possible communities. This can only happen through communication between peers. When no more attributes are being “escalated” by the peers, we say that the communication has achieved its goal and the collection of peers is *stable*.

To achieve it, we provide an algorithm that uses *attribute escalations* as follows:

1. Each peer P communicates with all of its neighbors² and all of its neighbor’s neighbors. The communication involves a transmission of the claimed attribute set by P . The reason we stop after the neighbor’s neighbor is explained later.
2. Each peer Q receiving such a set from one of its incoming links will find the intersection of the set with its own personal attribute set.
3. If an element is present in the intersection set just computed and absent from the claimed attribute set of Q , then it is added to the claimed attribute set of Q .

An example with two peers is illustrated in Figure 2.

Note that the attribute escalation algorithm automatically escalates attributes from the personal set to the claimed set. This may not be desirable in practice and a human may be consulted before such escalation is done. For this paper, we assume automatic escalation.

The above procedure causes a change in each peer and therefore results in a change in the communities identified. Since these changes might happen frequently with peers being created, destroyed, or modified, we are

² Neighbors are directly linked peers.

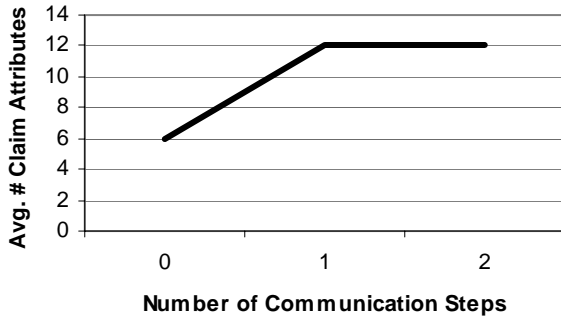


Figure 3: Stabilization of Attribute Escalations

The graph shows the number of times peers needed to communicate before the average number of claimed attributes for all peers becomes constant.

concerned about the stability of the attribute escalation technique given above.

We conducted simulations to measure how many communication steps our algorithm requires to identify all the possible communities. In this simulation, we generate a P2P network using our generation method described in the previous section. The network consists of 10,000 peers and each 100 peers form a group. As described earlier, attribute selections and link formation are random. On the average, we found each simulated peer contains 13 personal attributes and 6 claimed attributes; and is connected to 66 peers through regular connections and 0.1 remote peers through random connections. As shown in Figure 3, the average number of claimed attributes per peer stabilizes *after one iteration* of the attribute escalation algorithm.

Further, the computation at each peer is a trivial set algebra. Hence, each new arrangement of a collection of peers will stabilize quickly and with little overhead.

6. DISCOVERY OF A P2P COMMUNITY

Discovery in the context of P2P communities usually indicates one of the following: discovery of the communities in which the peer is automatically a member, or discovery of the existing communities that a peer might want to join or make use of. In this paper, we focus on the former.

We use the model generated as described earlier and empower the peers to discover two kinds of communities: the communities they are explicitly a part of, by virtue of their common group attribute, and the communities that they become a part of, by virtue of their claimed attribute set after escalations for stability. We run simulations to test the effectiveness of our algorithm in discovering P2P communities.

6.1 COMMUNITY DISCOVERY ALGORITHM

Before describing our discovery algorithm, we define two new terms:

OUTLINK WEIGHT: *The weight given to each claimed attribute based on the percentage of outgoing connections from a peer that can reach, after at most one indirection, other peers claiming the same attribute.*

INLINK WEIGHT: *The weight given to each claimed attribute based on the percentage of incoming connections to a peer that arrive directly from other peers claiming the same attribute.*

The constraint of at most one indirection is necessary to restrict the maximum depth up to which peers will be examined since more than two levels deep resulted in an unacceptably high number of communication messages. Further, our simulations have shown that using only one level of indirection, peers find an average of 8 communities. Probing to two levels of indirection did not increase the average number of communities discovered by a peer and hence is unnecessary.

When introduced into the simulation, these weights can be used in conjunction with a pre-defined threshold, after the stabilization step has been performed, to enable peers in discovering their community membership. If the *outlink* weight is above a threshold, the peer will assume that it is a member of the community indicated by the claimed attribute corresponding to that weight.

We define the membership of a peer for a community in a discovery process as follows:

MEMBERSHIP: *A node n is a member of a peer-to-peer community N with signature set S if S is a subset of $claim(n)$ and every element of the intersection set $I = claim(n) \cap S$ has an outlink weight greater than threshold T .*

The procedural pseudo-code for our membership discovery is shown in Figure 4.

Finally, we define two characteristics of peers that can be derived from the outlink and inlink weight computation. We do not use these definitions in this paper, but give them to hint at further uses for our discovery algorithm.

INVOLVEMENT: *The average of outlink weights for elements of the intersection set I , is directly proportional to the node's involvement in a peer-to-peer community.*

RESPONSIBILITY: *If a node n has a inlink weight associated with any member of the intersection set I , and the weight is greater than threshold B then the node is said to be a responsible member of the peer-to-peer*

```

MEMBERSHIP-DISCOVERY (peer)
/* executes at each peer simultaneously */

for all claimed attributes
  compute outlink weight
  compute inlink weight
end-for

for all subsets S of claimed attributes
set
  for all attributes in S
    if outlink weight > T then
      check next attribute
    else
      break-for
      /* try another subset */
    end-for

  if no break-for was executed
    peer is member of community
    with signature S
  end-for
end-for

```

Figure 4: Community Discovery Algorithm

community. The inlink weight is directly proportional to the responsibility of the node.

The values of the threshold can differ for each community and in this way allows for various types of communities, like *cohesive communities*, with high peer involvement; *supportive communities*, with high peer involvement and high peer responsibility; *sprawling communities*, with low peer involvement, etc.

6.2 SIMULATION RESULTS

We run simulations on peer networks of two different sizes: one with 1,000 peers, and the other with 10,000 peers. The maximum number of personal attributes we allow for each peer is 20, and a peer is explicitly placed into only one group, as described earlier. For both the simulations, we set the outlink threshold to 40%.

The graphs shown in Figures 5 and 6 are the results of our simulations. They demonstrate the performance of using the claimed attribute set and the outlink threshold for the discovery of P2P communities. A majority of peers have discovered communities other than the ones in which they were explicitly placed within (Figure 4). Our simulations have shown that before any communication, the average number of communities known to a peer was 0.5. This number increases to an average of 8.5 after execution of the membership procedure at each peer.

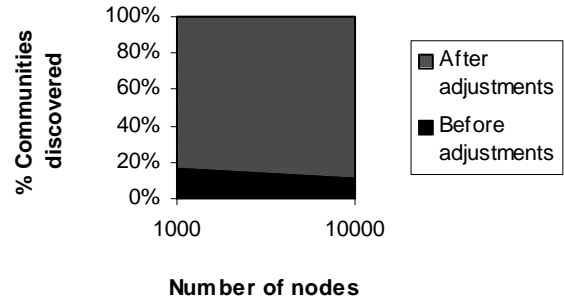


Figure 5: Percentage of communities discovered
The graph shows the average number of communities discovered by all peers before and after weigh calculation followed by attribute escalations. The outlink threshold used was 40%.

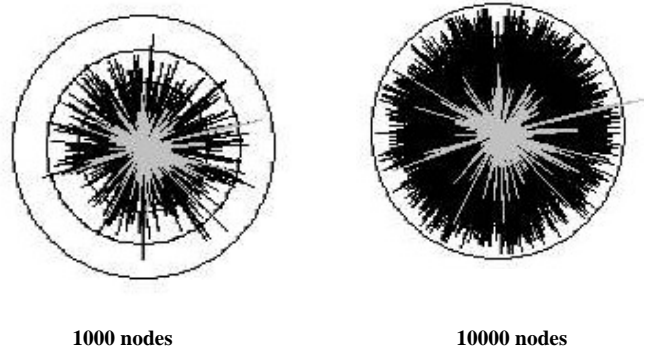


Figure 6: Number of communities discovered
The graphs show the number of communities discovered by the peers. The inner gray area is before and the outer black area is after weight calculation followed by attribute escalations. The outlink threshold used was 40%. The circles have values 0, 5, 10, and 15 from inside to outside.

7. CONCLUSIONS

P2P communities, is a method for arranging large numbers of peers in a self configuring peer relationship based on declared attributes (or interests) of the participating peers. This method is expected to have an impact in sharing of resources and pruning of search spaces based on the interests of the clients.

This paper shows that the attribute based clustering of peers can be made to work, by defining a overlay network, consisting of links. Links are user-directed connections based on experience and can be fine tuned for search and sharing performance. We show that the peer formation algorithm stabilizes in two rounds using the escalation technique. Additionally, we introduced an efficient community discovery procedure using weights and threshold. Our simulations of the discovery procedure have confirmed that peers can quickly discover numerous

memberships in different P2P communities using very little computation and communication messages

We are currently working on extending our results in peer communities to include fast and directed search algorithms and algorithms that handle privacy, security and trust management.

9. REFERENCES

- [1] Gnutella. <http://www.gnutelliums.com/>
- [2] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Workshop on Design Issues in Anonymity and Unobservability*, ICSI, Berkeley, CA, USA, July 2000, 311–320.
- [3] <http://www.napster.com>
- [4] Yahoo Groups. <http://groups.yahoo.com/>
- [5] Murray, H. A., *Explorations in personality* (Oxford University Press, New York, 1938).
- [6] Festinger, L., “Laboratory Experiments: The Role of Group Belongingness,” in J. G. Miller, ed., *Experiments in Social Process* (McGraw-Hill, New York, 1950).
- [7] G.W. Flake, S. Lawrence, C.L. Giles, and F.M. Coetzee, Self-Organization and Identification of Web Communities. In *IEEE Computer*, 35(3), March 2002, 66-71
- [8] K. Jiang and P. Dasgupta, SIMS: A Secure Information Management System for Large-Scale Dynamic Coalitions, In *The 2nd. DARPA Info. and Security Conf. and Expo.* July 2001.
- [9] Beverly Yang and Hector Garcia-Molina. Efficient Search in Peer-to-peer Networks. In *Int’l Conf. On Distr. Computing Sys.*, Vienna, Austria, 2002.
- [10] John Scott, *Social network analysis: a handbook* (SAGE Publications, 1991).
- [11] E. Garfield, *Citation Indexing: Its Theory and Application in Science* (Wiley, New York, 1979).
- [12] H.D. White and K.W. McCain, Bibliometrics, In *Ann. Rev. Info. Sci. and Technology*, Elsevier, 1989, 119-186.
- [13] G.W. Flake, S. Lawrence, and C.L. Giles, “Efficient Identification of Web Communities,” *Proc. 6th Int’l Conf. Knowledge Discovery and Data Mining*, ACM Press, New York, 2000, 150-160.
- [14] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. In *Proc. of the 9th Annual AVM-SIAM Symp. on Discrete Algorithms*, 1998, 668-677.
- [15] L. Page. PageRank: bringing order to the web. *Stanford Digital Libraries working paper* 1997-0072, 1997.
- [16] D. Gibson, J. Kleinberg, and P. Raghavan. Inferring web communities from link topology. In *Proc. 9th ACM Conference on Hypertext and Hypermedia*, 1998.
- [17] M.R. Garey and D.S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness* (W.H. Freeman, New York, 1990).
- [18] D. Watts and S. Strogatz, Collective Dynamics of “Small-World” Networks. *Canadian J. Math.*, vol. 8, no. 3, 1956, 399-404.
- [19] A. Kermarrec, L. Massoulie, and A. Ganesh. Reliable probabilistic communication in largescale information dissemination systems. *MSRTR-2000-105*, Microsoft Research Cambridge.
- [20] A.J. Demers, D.H. Greene, C. Hauser, and J. Larson. Epidemic algorithms for replicated database maintenance. In *Proc of the 6th Annual ACM Symp. on Principles of Distr. Comp.*, Vancouver, British Columbia, Canada, August 1987, 1–12.