Small World Architecture For Peer-to-Peer Networks

Lu Liu
Surrey Space Centre
University of Surrey
Guildford, United Kingdom
l.liu@surrey.ac.uk

Stephen Mackin
Surrey Satellite Technology Limited
Surrey Research Park
Guildford, United Kingdom
s.mackin@sstl.co.uk

Nick Antonopoulos
Department of Computing
University of Surrey
Guildford, United Kingdom
n. antonopoulos@surrey.ac.uk

Abstract

Small-world phenomenon has been observed in existing peer-to-peer (P2P) networks, such as Gnutella and Freenet. Due to the similarity of P2P networks to social networks, the previous small-world model proposed by Duncan Watts can be adopted in the design of P2P networks: each node is connected to some neighbouring nodes, and a group of nodes keep a small number of long links to randomly chosen distant nodes. Unfortunately, current unstructured search algorithms have difficulty distinguishing these random long-range shortcuts. This paper presents small world architecture for P2P networks (SWAN) with a semi-structured P2P search algorithm that is used to create and find long-range shortcuts toward remote peer groups. In SWAN, not every peer node needs to be connected to remote groups, but every peer node can easily find which peer nodes have external connections to a specific peer group.

1. Introduction

Small world phenomenon is the hypothesis that everyone in the world can be reached through a short chain of social acquaintances [1]. Small-world phenomena [1] have also been observed in existing peer-to-peer (P2P) networks, such as Gnutella and Freenet [2]. Duncan Watts proposed a mathematical model [3] to analyze the small world phenomenon with highly clustered sub-networks consisting of local nodes and random long-range shortcuts that help produce short paths to remote nodes. Due to the similarity between P2P networks and social networks, where peer nodes are people and connections are relationships, Duncan's theory can potentially be adopted in P2P networks: each peer node is connected to some neighbouring nodes on the P2P topology, and a group of peer nodes keep a small number of long links to randomly chosen distant peer nodes. However, current unstructured search algorithms have difficulty distinguishing among such random long-range shortcuts and finding proper shortcuts for a specific resource search.

This paper presents small world architecture for P2P networks (SWAN) for content discovery in multi-group P2P systems. A semi-structured P2P algorithm of SWAN is used to create and discover long-range shortcuts between different peer groups, which can satisfy the following requirements of design:

- Not every peer node needs to connect to other groups;
- Each peer node needs to know or can easily find which nodes have external connections to which peer groups;
- External links to other peer groups need to be distributed within the peer group and cannot be centralized in one or a few peer nodes.

Existing P2P solutions for resource discovery can be generally classified into two categories: structured and unstructured P2P systems. An observed issue of most DHTs is that the cost of maintaining a consistent distributed index is too high in the dynamic Internet environment [4, 5]. Some structured P2P protocols (e.g. Kademlia [6]) are beginning to seek ways to save the cost of maintaining a consistent index. In contrast, unstructured P2P systems (e.g. Gnutella) do not control data placement and are more resilient in dynamic environments, but current unstructured P2P search techniques tend to either require high search overhead or generate massive network traffic. To address these problems, a semi-structured P2P search method is presented in this paper combining techniques of both structured and unstructured methods. The search algorithm of SWAN does not strictly rely on DHTs which can still find the requested data inside and outside of peer groups with a high probability even though hash functions can not provide accurate information of data locations.

2. Related Work

Maintaining and searching "small world" has been discussed in recent P2P studies. Jon Kleinberg discussed the problem of decentralized search in P2P networks with partial information about the underlying structure in [7]. A study in [8] proposes an enhanced clustering cache replacement scheme for Freenet by forcing the routing tables to resemble neighbour relationships in a small-world acquaintance graph. Semantic Small World in [9] facilitates efficient semantic based search in P2P systems where peers are clustered according to the semantics of their local data and self-organized as a small world

overlay network. Despite the fact that unstructured P2P are more resilient in dynamic environments, the efficiency of these unstructured P2P approaches is still far lower than DHTs. Hierarchical SRDI service of JXTA [10] attempts to use combined techniques with both structured and unstructured search methods. However, it encounters the performance bottleneck of centralized super-peers and maintenance of the distributed index is disordered and redundant. In contrast, our model is built upon a flat P2P overlay network, which is efficient and fault-tolerant for content discovery inside or outside of peer group.

3. Algorithm Description

SWAN is built with the generic group structures similar to Jon Kleinberg's model [7]. The previous studies like [11, 12, 13] have presented the methodologies of building an information sharing system by grouping peer nodes that will not be discussed in this paper. This paper will focus on data publishing and searching algorithms of SWAN with generic group structures. By using a compact representation mechanism (e.g. Bloom Filters [14]), each peer node maintains an inconsistent list about members in the same group and regards other members as "acquaintances." A group of peer nodes keep a small number of long links to distant peer nodes. A semi-structured approach is presented in this section to create long-range links between groups as well as discover the local peer nodes that have specific external connections. An example of SWAN topology is illustrated in Figure 1.
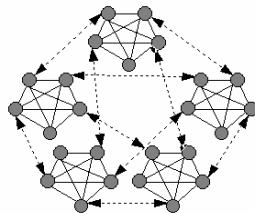

Figure 1. Topology of SWAN.

3.1. Intra-Group Content Searching

Each shared file in SWAN is published by an associated content advertisement that provides the relevant meta-information of the file (e.g. name, node address, description), which is pushed to a target peer node according to the hash value of the name of file, as well as internal neighbours of the target peer node in the member list within a specific distance d to increase the probability of discovery of the advertisement. The advertisement searching process involves two steps: a structured P2P search followed by an unstructured P2P search. The requesting peer node firstly searches the target peer nodes generated from the same hash function (structured P2P search) If the requested advertisement cannot be found in the target peer node (e.g. the target

peer node is offline at the moment), the requesting peer node will continue to search the neighbours of the target peer node in member list within distance d (unstructured P2P search) and can find the requested files with a high probability. Figure 2 illustrates an example of content advertisement publishing and searching. *P1* shares a file with the name *K1*. The publication service on peer *P1* pushes the associated advertisement of *K1* to *P4* according to the hash value of *K1* and the neighbours of *P4 (P3 and P5)* within the distance $d = 1$. Then other peer nodes in the same peer group can easily find the advertisement in a high probability. In this case, *P6* looks for the advertisement by generating the same hash value pointing to *P4* with the same hash function and sends a query to *P4* and find the advertisement with *K1* in *P4*. However, if the requested advertisement cannot be found in *P4*, the requesting peer node will continue to search *P3* and *P5* that are neighbours of *P4* within distance $d = 1$. Publication and searching parameter d is defined based on users' requirements and present rate of peer nodes. Generally, a bigger d is required in a dynamic network with a lower peer present rate.
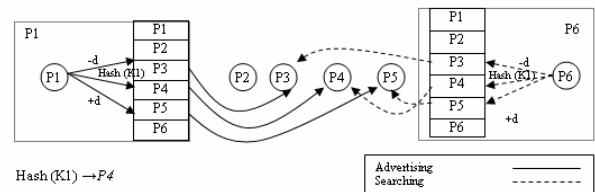

Figure 2. Advertisement publishing and searching.

3.2. Inter-Group Content Searching

In SWAN, a new peer group is advertised by a group advertisement that provides the relevant meta-information about the peer group (e.g. ID, name, contact points, description), which will be multicast through network. Not all the peer nodes in the network will receive the advertisement, but a large percentage of them will. When a peer node receives a peer group advertisement, it will push the advertisement to a target peer node in the peer group according to the hash value of the name of the peer group as well as neighbours of the target peer node within a specific distance d to increase probability of discovery of the advertisement. Similar to the intra-group content searching, if the requesting node cannot find the requested advertisement with the uniform hash function due to high network churn, the requested advertisement will still be found in the neighbours of the target peer node with a high probability by using unstructured P2P searching. Therefore, even though only one peer node is informed, all the peer nodes in the same peer group potentially can find and pull the peer group advertisement. Figure 3 illustrates the process of inter-group link formation. When *P1* receives a group advertisement about group *G2* with contact point *P'3*, it will push the advertisement to the

target peer node *P4* according to hash function as well as its neighbours (*P3* and *P5*) within distance $d=1$. Then *P4* will inform a contact point of group *G2: P′3* with the advertisement of its peer group *G1*. When *P′3* gets the advertisement of *G1*, *P′3* will do the same as *P1* to forward the advertisement of *G1* toward the target peer node *P′5* according to hash function as well as its neighbours within distance *d*. When *P′5* receives the advertisement, *P′5* will do the same as *P4* to send the advertisement of its group *G2* back to *P4*. When *P4* receives it and sends the acknowledgement of inter-group link back to *P′5*, an inter-group link will be built between *G1* and *G2* and be maintained by *P4* and *P′5*. In the same way, more inter-group links will be created and maintained between *P3, P4, P5* and *P′4, P′5, P′6*, in case of $d=1$. Each of them normally keeps 3 (*2d+1*) inter-group links as illustrated in Figure 4 which makes groups connected even in a highly dynamic environment. Inter-group search queries can be propagated toward the requested peer group efficiently via inter-group links and relevant shared files can be found with high probability. The methods of resolving the pair of keyword of clusters and value of shared files have been discussed in [15] and will not be described in detail in this paper.
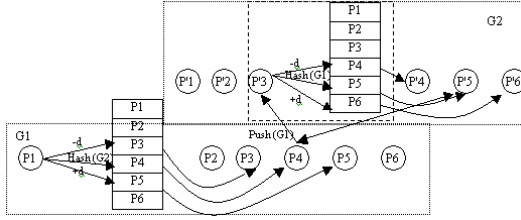

Figure 3. Inter-Peer group link formation.

4. Performance Evaluation

We evaluated the effectiveness of SWAN in the dynamic environments with frequent peer nodes online and offline. In this section, we assume that the requested advertisements have been published successfully to the target peer node as well as its neighbours within a distance *d*. There are *n* peer nodes in the network, each peer node has *k* neighbours in each peer group, peer groups do not overlap and are connected by inter-group links as shown in Figure 1. There are a total of *k+1* peer nodes in each group. The success rate and cost of inter-group content searching will be evaluated with peer present rate *p*.

In SWAN, three conditions must be satisfied to find an advertisement in a different group as shown in Figure 4:

(1) Succeed in finding an advertisement about the requested peer group on the searching peer node itself (case A) or on other members (case B).
(2) Succeed in contacting the requested group (case C).
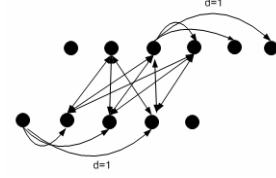(3) Succeed in finding a requested content advertisement in the requested peer group (case D).

Figure 4. P2P searching in dynamic environments.

The probability of failing in finding an advertisement about the requested peer group is:

$$P\left(\overline{A \cup B}\right) = P\left(\overline{A}\right)P\left(\overline{B} \mid \overline{A}\right) = \left(1 - \frac{2d+1}{k+1}\right) \cdot (1-p)^{2d+1} = \frac{k-2d}{k+1}(1-p)^{2d+1}$$

We assume that queries will be forwarded to the requested peer group if the peer nodes on both ends of the inter-group link are online. Recall that the local peer group keeps inter-group links toward *2d+1* peer nodes in a remote peer group. The searching peer node will fail to contact the requested peer group, if all *2d+1* peer nodes are all offline. Therefore, the probability of failing in contacting the requested group is: $P\left(\overline{C} \mid A \cup B\right) = (1-p)^{2d+1}$. The probability of finding an advertisement about the requested group is $P(D \mid C(A \cup B)) = 1 - \frac{k-2d}{k+1}(1-p)^{2d+1}$.

The success rate of finding a content advertisement in the requested peer group is:

$$P_s = P(DC(A \cup B)) = P(D \mid C(A \cup B)) \cdot P(C \mid (A \cup B)) \cdot P(A \cup B)$$

$$= \left[1 - \frac{k-2d}{k+1} \cdot (1-p)^{2d+1}\right]^2 \left[1 - (1-p)^{2d+1}\right]. \tag{1}$$

If *d* is defined as a small value, the success rate is also very low in the network with a low present rate that is the situation need to be avoided in practice. Figure 5(a) shows the minimal values of *d* required to achieve different satisfactory success rates with different peer present rates.

The expected number of messages per query is calculated to judge the traffic cost:

$$\mathrm{E}[N] = E[N(ACD)] + E\left[N\left(AC\overline{D}\right)\right] + E\left[N\left(A\overline{C}\right)\right] + E\left[N\left(\overline{A}BCD\right)\right] +$$

$$E\left[N\left(\overline{A}BC\overline{D}\right)\right] + E\left[N\left(\overline{A}B\overline{C}\right)\right] + E\left[N\left(\overline{AB}\right)\right] = \left(1 - \frac{2d+1}{k+1}\right)\left\{\sum_{i=1}^{2d+1}\left\{p \cdot (1-p)^{i-1} \cdot\right.\right.$$

$$\sum_{j=1}^{2d+1}\left[p \cdot (1-p)^{j-1}\left(\frac{2d+1}{k+1} \cdot (j+i) + \left(1 - \frac{2d+1}{k+1}\right)\sum_{m=1}^{2d+1} p \cdot (1-p)^{m-1} \cdot (m+j+i)\right)\right]\right\} +$$

$$\sum_{i=1}^{2d+1}\left[p \cdot (1-p)^{i-1}\sum_{j=1}^{2d+1} p \cdot (1-p)^{j-1} \cdot \left(1 - \frac{2d+1}{k+1}\right) \cdot (1-p)^{2d+1} \cdot (2d+1+j+i)\right] +$$

$$\sum_{i=1}^{2d+1} p \cdot (1-p)^{i-1}(1-p)^{2d+1} \cdot (2d+1+i) + (1-p)^{2d+1} \cdot (2d+1)\right\} +$$

$$\frac{2d+1}{k+1}\left\{\sum_{j=1}^{2d+1} p \cdot (1-p)^{j-1}\left[\frac{2d+1}{k+1} \cdot j + \left(1 - \frac{2d+1}{k+1}\right) \cdot \sum_{m=1}^{2d+1} p \cdot (1-p)^{m-1} \cdot (m+j)\right] +$$

$$\sum_{j=1}^{2d+1} p \cdot (1-p)^{j-1} \cdot \left(1 - \frac{2d+1}{k+1}\right) \cdot (1-p)^{2d+1} \cdot (2d+1+j)$$

$$+ (1-p)^{2d+1} \cdot (2d+1)\right\}. \tag{2}$$

5. Simulation Results

We evaluated the performance of inter-group content searching by simulation in dynamic environments. In the simulation, we followed the same assumption presented in section 4 that the requested advertisements had been published successfully to the target peer node as well as its neighbours within a distance d. Therefore, if either the target peer node or one of its neighbours within distance d was visited successfully, this search succeeded. In the simulation, each peer group kept 500 peer nodes, all peer nodes in the same peer group were completely connected to each other, and all peer nodes were initialized as present in the network. At the beginning of each search, a set of peer nodes were randomly selected and set as absent according to the parameter of present rate of peer nodes, the searching peer node was randomly selected from the set of present peer nodes and the target peer node with its neighbours were randomly selected from the set of peer nodes regardless of their present situation. In each data search, the searching peer node initials a query that will be passed with the SWAN protocols.

The searching peer node and the target peer node were separated into different groups in the simulation of inter-group content searching. Figure 5(b) and (c) show the results of success rate and average number of messages of each query by the inter-group search (for 1000 queries), in which the theoretical results were generated from Equation (1), (2).
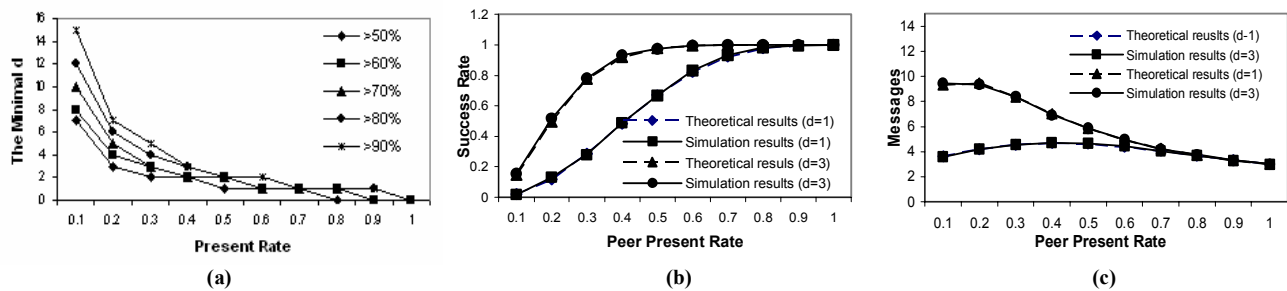

Figure 5. (a) The minimal publishing distance for required success rates. (b) Success rate. (c) Average messages.

6. Conclusion

This paper presented small-world architecture for content discovery in unstructured P2P networks. In SWAN, each node is connected to neighbouring nodes in the same peer group and peer groups are connected by a small number of inter-group links. Not every peer node needs to be connected to remote groups, but every peer node can easily find which peer nodes have external connections to a specific peer group. A semi-structured P2P search method is presented in this paper combining techniques of both structured and unstructured search methods, which can discover the requested data with a high probability even though hash functions can not provide accurate information of data locations. From the simulation results, SWAN potentially has advantages of both structured and unstructured P2P networks and achieved good performance in dynamic environments.

References

[1] S. Milgram, "The Small World Problem," *Psychology Today*, 1967. pp. 60 – 67.

[2] T. Hong, Peer-to-Peer: Harnessing the Power of Disruptive Technologies, O'Reilly, 2001, pp. 203 – 241.

[3] D. Watts and S. Strogatz, "Collective Dynamics of Small-World Networks," *Nature 393*, 1998, pp. 440-442.

[4] B. Yang and H. Garcia-Molina, "Efficient Search in Peer-to-Peer networks," in *Proc. of ICDCS*, 2002.

[5] S. Rhea, D. Gells, T. Roscoe, and J. Kubiatowicz, "Handling Churn in a DHT," Berkeley and Intel Research, 2004.

[6] P. Maymounkov and D. Mazieres, "Kademlia: A Peer-to-Peer Information System Based on the XOR Metric", in *Proc. of IPTPS*, 2002.

[7] Jon Kleinberg, "Small-World Phenomena and the Dynamics of Information," in *Proc.of NIPS*, 2001.

[8] H. Zhang, et. al, "Using the Small-World Model to Improve Freenet Performance," *Computer Networks 46(4)*, 2004, pp 555- 574.

[9] M. Li, W. Lee, and A. Sivasubramaniam, "Semantic Small World: An Overlay Network for Peer-to-Peer Search," in *Proc. of ICNP*, 2004.

[10] B. Traversat, M. Abdelaziz, and E. Pouyoul, "Project JXTA: A Loosely-Consistent DHT Rendezvous Walker," Technical Report, Sun Microsystems, Inc, March 2003.

[11] "Project JXTA 2.0: Java Programmer's Guide," Sun Microsystems, 2003.

[12] M. Khambatti, K.D. Ryu, and P. Dasgupta, "Structuring Peer-to-Peer Networks Using Interest-Based Communities," in *Proc. of DBISP2P*, 2004.

[13] J. Vassileva, "Motivating Participation in Peer-to-Peer Communities," in *Proc. of ESAW*, 2002.

[14] B. Bloom, "Space/time Trade-offs in Hash Coding with Allowable Errors," *Communication of ACM*, 1970, pp.422-426.

[15] N. Antonopoulos and J. Salter, "Improving Query Routing Efficiency in Peer-to-Peer Networks," University of Surrey, Computing Sciences Report, CS-04-01.