PSON: A Scalable Peer-to-Peer File Sharing System Supporting Complex Queries

Jyoti Ahuja[1], Jun-Hong Cui[1], Li Lao[2], Shigang Chen[3]
jyoti, jcui@engr.uconn.edu, llao@cs.ucla.edu, sgchen@cise.ufl.edu
[1] CSE, UCONN, [2] CS, UCLA, [3] CISE, U. of Florida

ABSTRACT

In this work, we design a P2P file sharing system called PSON to address the important issues of scalability, routing efficiency and complex query support. We propose a semantic overlay network of logical nodes, in which queries are routed on the basis of semantics. A logical node is formed by a cluster of peers that are close to each other in the physical network. Each cluster selects a powerful peer as super peer to support routing in the overlay network. All the super peers are organized in the form of a balanced tree. By exploiting the concepts of hierarchy and semantics, PSON can support complex queries in a scalable and efficient way. Our initial results have shown the promising performance of PSON.

1. INTRODUCTION

In recent years, the peer-to-peer (P2P) computing model has gained tremendous popularity in the research community as well as the industry. One of the key incentives is its high scalability. Among various applications of P2P technology, P2P file sharing systems are probably the fastest growing and the most popular Internet application. It is reported that P2P file-sharing accounts for much more traffic than any other application – including the Web – on the Internet.

A successful P2P file sharing system should achieve the following goals. First, the system should scale to large numbers of peers spreading throughout a wide network across different administrative domains. Second, the system needs to provide an efficient and effective file lookup (or search) scheme. It should return the location of a requested file with minimal communication and computation overhead. Third, the system should support complex (or partial-match) queries. Shared files have attributes that describe their properties, e.g., singer, composer and title for a music file. It is desired that the search mechanism could support partial match queries that contain a subset of the attributes (or even typos).

Many P2P file sharing systems have been proposed in the literature: Napster, Gnutella, KaZaa, BitTorrent, Chord, and CAN, to name a few. Some of them have been implemented and even widely used. However, few systems could achieve the above design goals simultaneously.

To address the issues of scalability, efficiency and complex query support, in this work, we present a file sharing system, called **P**eer-to-peer **S**emantic **O**verlay **N**etwork (PSON). The basic design essentially exploits the concepts of hierarchy and semantics.

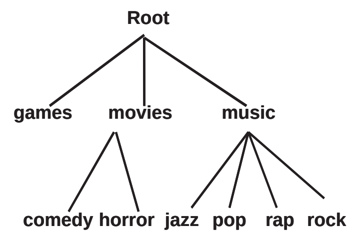

Figure 1: An Example of Semantic Hierarchy

Files shared in PSON are classified into a semantic hierarchy. Fig. 1 gives a simple example where each tree node represents a semantic entry. Any file in the system semantically belongs to one tree node. With a tree structure, semantic entries can be easily compared and ordered, thus we can sort all the files in the system accordingly. For example, we set up several rules to compare the node entries: for any two semantic entries, if they have the same parents, their order is decided lexicographically; if they have different parents, then find their lowest-level uncommon ancestors and compare them lexicographically. Following these, “jazz” is smaller than “pop”, and “comedy” is smaller than “pop”.

Peers joining PSON are categorized into normal peers and super peers. Powerful nodes can serve as super peers. Normal peers connect to a super peer and together form a cluster. If we treat each cluster as a logical node, then all logical nodes form an overlay network (referred to as semantic overlay) in which queries are routed on the basis of semantics. Each of the logical nodes (i.e., clusters) has a content directory to manage. Peers belonging to a cluster share the content directory assigned to the cluster. In this way, when a peer wants to publish a file, it first extracts metadata and generates a location-metadata pair (i.e., a directory item). Then the directory item is distributed into the semantic overlay network, finds a proper “host” cluster, and finally is stored in some peer in the cluster. Similarly, when a peer wants to search some file, a search query is first issued, and then forwarded to a cluster responsible for storing the corresponding directory. By a local search in the cluster, the location for the requested item is obtained.

Clearly, in PSON, semantic overlay construction is the key to the whole system design. In our work, we use red-black tree as the overlay structure, for which an in-order tree walk could yield an encoded semantic tree, which helps to conserve the semantics of the shared content. Fig. 2 gives one possible overlay tree for the semantic tree illustrated in Fig. 1. The in-order walk of the overlay tree is <*games, comedy, horror, movie, jazz, pop, rap, rock, music,*

SIGCOMM’05, Poster Session August 22–26, 2005, Philadelphia, PA, USA.

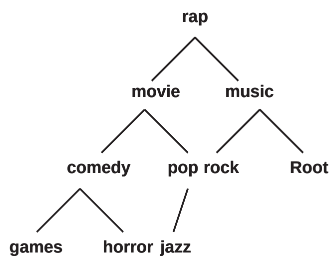

Figure 2: The Structure of the Overlay Tree

Root>, which is exactly the same as the ordering of the semantic entries in the semantic tree. By utilizing a red-black tree overlay structure, the routing (logical) hops can be bounded by $O(log(n))$ (where is n is the number of logical nodes).

PSON is designed as a scalable and efficient file sharing system supporting complex queries. Though the basic idea is not complicated, to make the system work effectively and efficiently, there are many important issues to address.

(1) How to construct a red-black overlay tree in a distributed way? Obviously, the clusters usually do not appear at the same time. Constructing the semantic overlay efficiently while facilitating query routing is desirable.

(2) How to handle the file dynamics and node dynamics? Peers join and leave, and files are inserted and deleted. These dynamics should not affect the quality of the semantic overlay.

(3) How to conserve the robustness of the P2P model? When some peer fails or even super peer fails, the overlay should still function without significant performance penalty.

(4) How to well balance the load of peers in the network? Load balancing is one of the key concerns in P2P system designs. Will the semantic overlay achieve the goal of load balancing when it introduces semantics?

In this work, we address all these design issues of PSON. We construct a balanced binary search tree structure (i.e., red-black tree) for the semantic overlay in an incremental fashion, and efficiently control the tree rotation operations. Centered on the basic semantic overlay structure, we propose solutions for the issues of (node and file) dynamics handling, fault tolerance, and load balancing. We have implemented PSON in NS-2. We conducted simulation studies, and our preliminary results have demonstrated that PSON is capable of supporting complex queries very efficiently. We also show some other aspects of the system performance, such as stability, resilience, and load distribution. The results suggest that PSON is a very promising P2P file sharing system supporting complex queries. Due to space limit, we only show some results with regard to the performance of PSON for complex queries in this extended abstract. Interested readers are referred to the technical report [1] for more details.

2. COMPLEX QUERY PERFORMANCE

In this section, we show the performance of complex queries in PSON. An entry which a complex query searches in PSON could be a leaf as well as a non-leaf node of the semantic tree. A query that corresponds to a leaf node of the semantic hierarchy is characterized as a *leaf query* and the one corresponding to a non-leaf node, as *non-leaf query*. When a non-leaf query is released into the system and is found, the idea is to return the requested node as well as all the semantic descendants of the the node in the semantic hierarchy. For instance, if "music" and two of its semantic children, "jazz" and "pop" are present in the overlay then a query for music will return "music", "jazz" and "pop". This is because a

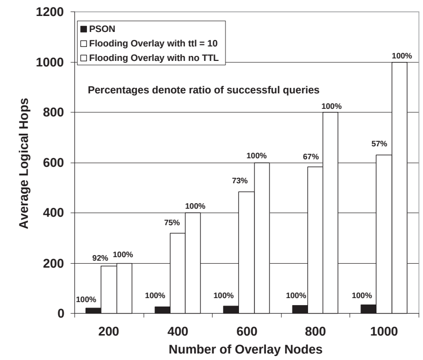

Figure 3: Comparing the complex query performance of PSON and flooding semantic overlay

query for "music" is a kind of complex query (because the query initiator didn't give all of the attributes like which kind of music file it is requesting), and since "jazz" and "pop" are also classified as "music", they should be returned too. On the other hand, if a leaf query is found, just that node is returned.

A complex query is termed as successful if it finds all the nodes who have the requested file. We define "successful percentage" as the ratio of the number of nodes found to the number of nodes with the file. PSON can guarantee 100% successful percentage. In this experiment, we compare PSON's performance with flooding overlay (the only existing semantic based overlay proposed in [2]) for complex queries. We measured the average number of logical hops for PSON and flooding overlay with TTL = 10 and without any TTL limit. We choose to test flooding overlay without any TTL because this gives us the total number of nodes that should respond to this query. We repeated the experiment to see how many of these nodes actually responded in the case when TTL = 10 and measured the successful ratio with the average number of logical hops traversed. Fig. 3 gives the comparison results. From Fig. 3, we can see that for the same query success ratio, PSON utilizes less than five percent of logical hops required by the flooding approach.

3. CONCLUSIONS AND ON-GOING WORK

In this work, we have proposed a powerful P2P file sharing system, PSON. It satisfies all the three design goals of a successful P2P file sharing system: 1) Scalability: PSON is not only scalable to a large number of peers, it is also scalable to a large amount of traffic; 2) Routing Efficiency: By utilizing a red-black overlay tree, PSON could route queries in time bounded by $O(log(n))$; 3) Supporting Complex Queries: By exploiting the concept of semantics, PSON could support various complex queries very effectively. By simulations, we have shown the promising performance of PSON.

On-Going Work To further check the feasibility of PSON in real network environments, we are implementing a prototype PSON system on PlanetLab, which will serve as a powerful vehicle to test the performance of PSON from various aspects.

4. REFERENCES

[1] J. Ahuja, J.-H. Cui, L. Lao, and S. Chen. A scalable peer-to-peer file sharing system supporting complex queries. *UCONN CSE Technical Report, UbiNet TR05-01*, 2005.

[2] A. Crespo and H. Garcia-Molina. Semantic Overlay Networks for P2P Systems. In *Technical Report, Computer Science Department, Stanford University, October*, 2002.