Choosing Appropriate Peer-to-Peer Infrastructure for Your Digital Libraries

Hao Ding and Ingeborg Sølvberg

Dept. of Computer and Information System, Norwegian Univ. of Sci. &Tech.
Sem Sælands vei 7-9, NO-7491, Trondheim, Norway
{hao.ding,ingeborg} @idi.ntnu.no

Abstract. Peer-to-Peer (P2P) overlay network aims to be a feasible platform for building federated but autonomous digital libraries. However, due to a plethora number of P2P infrastructures and corresponding functionalities, it is often not easy to choose appropriate candidates for specific applications. This paper is devoted for this issue by comparing some typcial P2P systems widely used in digtal library or databbase communities and extending an open discussion on how to determine proper infrastructures according to specific system requirements.

1 Introduction

Peer-to-Peer (P2P) overlay network becomes a substantial research topic in recent digital libraries applications. One fact is that a special working group has been settled in DELOS [1] focusing on constructing highly scalable, customizable and adaptive digital libraries, where building digital libraries over P2P overlay network is a major research activity. However, due to a plethora number of P2P infrastructures available, a core requirement exists in determing appropriate P2P infrastructure for specific digital library applications. In this paper, we study some representative P2P infrastructures by comparing several key features which are critical for system functionalities and performances. Section 2 presents a comparsion on selected P2P systems or infrastructures, followed by discussions concerning key issues in deciding appropriate P2P solutions. Conclusions and recommended future work are in Section 3.

2 P2P Systems Summary and Discussion

2.1 Comparing P2P Systems

Although many criteria must be taken into consideration in comparing P2P systems, we believe that information searching will be one of the most significant factors. Hence, the comparison will be centered mainly on this issue. Table 2.1 illustrates the comparisons over selected systems in aspects of markup schema, hash table usage, semantic routing style, query forwarding support, semantic query support and system topology.

Gnutella (http://www.limewire.com/), Napster (http://www.napster.com), and Freenet [2] are ancestors in P2P computing. They all support keyword-based search. Gnutella is a representative instance for query flooding which can not scale well. Napster goes

Table 1. Typical P2P Systems Summary

System	Markup-Scheme	Hash Table /Usage	Semantic Routing	Query Forwarding	Semantic Query
Gnutella	Keyword	No	No	Yes	No
Naspter	Keyword	No	No	No	No
Freenet	Keyword	Yes(binary)	Serial	Yes	No
Routing Indices	Keyword	No	Serial	Yes	No
Chord	Keyword	Yes	Parallel	Yes	No
CAN	Keyword	Yes	Parallel	Yes	No
pSearch	Keyword	Yes	No	Yes[1]	No
Piazza	Database	No	No	Yes	relational+XML
HyperCup	Keyword	Yes	Separate HyperCube	Yes	Yes
JXTA Search	XML	No	Parallel	Yes	No
Edutella	RDF	No	Parallel	Yes	Yes (regional)
Bibster	RDF/DAML+OIL	No	Parallel	Yes	Yes (global)
OAI-P2P(ongoing)	RDF	No	Parallel	Yes	Yes (regional)
RDFPeers	RDF	Yes	Parallel	Yes	Yes (global)
P-Grid	Keyword	Dist. Search Tree	Serial	Yes	No

in the opposite direction and adopts central servers to maintain a centralized directory from which connected peers can register their profiles/expertises and also retrieve a list of peers of user's interest. In Freenet, each file/document is identified by a binary key which is generated using some hash function; each peer maintains a local routing table which keeps information about neighbouring peers and the keys are a sequence of *(file key, node address)* pairs used for retrieval.

Crespo [3] uses Routing Indices (RIs), created and maintained by each peer, to forward queries to neighbours that are more likely to have answers. Any peer's joining or leaving can lead to a cascade of updates in RIs. And this is the overhead generated for the sake of efficient query forwarding in RIs instead of random flooding.

Distributed Hash Table (DHT) is probably the most widely used algorithm in P2P computing. Generally, DHT systems assign each entity (e.g. file names) a key generated by a hashing algorithm, then map the key to the node which also has an ID (e.g. hashed IP address). Normally this ID is the one closest to the key. In consequence, the storage and lookups of keys are distributed among multiple hosts. The performance of all DHT algorithms has been justified pretty good [4]. For instances, each node maintains information only about $O(logN)$ other nodes, and a lookup requires $O(logN)$ messages. Chord [5], CAN [6] and pSearch [7] are representative DHT systems. One requirement in applying DHT is that all participating libraries (peers) must be highly coupled, and moreover, an uneven distribution of document pointers may be expected over the peers and the global state is required beforehand as a support for the algorithm.

Piazza [8] is a peer data management system that enables sharing heterogeneous data in a distributed and scalable way. The assumption in this system is that the participants have similar content to share within other peers. Then, pair-wise mappings are defined between their schemas and users can formulate queries over their preferred schema. Piazza also creates a query answering system for expanding recursively any mappings relevant to the query, retrieving data from other peers.

JXTA [9] is a P2P interoperability framework created by Sun Microsystems. All peers can publish their profiles (i.e., content summary) in way of 'advertising'. One peer in JXTA can thus discover other peers by discovering posted 'advertisements' and then join favorite peer groups. Communications between peers are conducted by 'pipes' specifically generated by them. Typical systems include Edutella [10] and Bibster [11]. Both of them support metadata search within P2P networks while the former focuses on educational domain and the latter on bibliographic records respectively. JXTA itself can be regarded as a super-peer network consisted of many 'rendezvous' peers [9]

HyperCup [12] proposes a graph topology which allows very efficient broadcast and search which intend to reach all peers in the network with the minimum number of messages possible. The number of messages generated when peers leave and join the network is $O(log_b N)$ (b is the base of the hypercube), which can be more efficient than DHT algorithm. Moreover, a global ontology is proposed to determine the organization of peers in the graph topology, allowing for efficient concept-based search.

The ongoing OAI-P2P project [13] aims to design a P2P network for open archives, where data providers form a P2P network which supports distributed search over all connected metadata repositories. In this scenario service providers can be removed from this network and make the data repositories more up-to-date.

RDFPeers [14] is a very interesting approach by extending DHT to support searches over RDF triples. Basically, RDFPeers becomes a scalable distributed RDF repository that stores each triple at three places in a multi-attribute addressable network by applying globally known hash functions to its subject, predicate, and object. Such an approach is very suitable to search through highly distributed RDF repositories.

P-Grid [15] is a kind of Semantic Overlay Network (SON) [16], which differs from other approaches such as Chord, CAN, etc. in terms of practical applicability (especially in respect to dynamic network environments), algorithmic foundations (randomized algorithms with probabilistic guarantees), robustness, and flexibility. The most important properties of P-Grid are: complete decentralization;self-organization; decentralized load balancing; data management functionalities (update);management of dynamic IP addresses and identities; efficient search[15].

2.2 Discussion

In constructing specific P2P-based digital libraries, different institutes may have different requirements in constructing P2P networks. Some critical requirements demanding special considerations are listed as follows:

- Degree of autonomy: does your library accept arbitrary incoming queries? Or can you support a common shared schema? It will force you to convert queries before sending them to connected P2P system.
- Keyword-based search or metadata/ontology-based search;
- Multiple (heterogeneous) metadata schemas support: e.g., LOM, DC, etc.
- Metadata records harvesting: if it is not necessary to keep data up-to-date, consistency issue must be considered.
- Authentication: must the library users be authenticated?
- Peer Selection/Discovery: do you need to locate specific libraries or just let system to find them dynamically?

As to applications quering few metadata fields, such as music file sharing which may request only file names, keyword-based searching over a query flooding P2P environment (e.g., Gnutella) or a centralized server-based P2P network(e.g., Napster) is sufficient. Moreover, if libraries can be highly coupled, DHT-based solutions can be used to achieve more efficient and effective performance. One issue which needs further clarification here is that DHT-based solution can only release the impact of frequent requests for some information. It can not release the impact of data hotspots due to key collisions which may be caused by too much entities/data being associated with one key. Recent approaches in super-peer based topology [17] or SON can be considered as alternatives to improve query efficiency. These approaches can be contributed for requirements when many digital library systems take *autonomy* as a central value since these approaches can support a more flexible mechanism for loosely coupling among peers. It is dissimilar with the rigid infrastructure as in DHT, although the latter makes it easier to locate content later on.

The rest discussion is devoted to the issues which shall be taken into consideration in applying different information searching methods, namely traditional information retrieval (IR)-based and metadata-based (semantic) search. Basically, either search method may be found more suitable than the other in some application scenarios. On one hand, many collections in participating libraries may have various metadata schemas which involve multiple fields, such as title, author, publication, etc. That is, searching over collections can be roughly regared as a matchmaking procedure on related fields recursively. However, when more and complex metadata elements get involved, such as Bibtex metadata with up to 100 metadata entries [11], an advance mechanism for supporting more complex queries is then required. Edutella and Bibster, in this concern, demonstrate the possibility to conduct complex queries over metadata records, by using RQL [18] alike query language and RDF-based database management systems - Sesame [19]. One weakness here is that few approaches have been conducted to support queries across heterogeneous but semantically related metadata records. However, approaches are being conducted in this direction [20–22] and we can foresee more researches to come. On the other hand, if the textual parts, such as description and abstract, take up a large proportion in metadata records, or just simply, any textual documents are available, we may need a conventional full-text IR method to conduct search. However, it becomes complicated in context of P2P overlay network. For example, when using Space Vector Model (SVM) (c.f. [23]), inverted document frequencies (IDF) may not be easy to maintain simply because of the dynamic nature of P2P systems. In order to keep such global statistic information, a huge index may then be maintained which will correspondingly take up a large bandwidth. So as long as there are thousands of peers in a P2P system, it would be problematic to collect the intact information about all the document collections on all peers. Additionally, even if it is possible to get the global statistic parameters by gathering information from all the peers involved, we are still faced with a problem that a peer would join or leave the system at any time. In this case, the collected global statistic information would be out of date and must be updated when new peers joins and old peers leaves. An alternative solution is routing indices (RI) which can avoid the overhead of constant index updates, but due to its local nature, it is in turn difficult to obtain necessary global information. Fortunately, substantial approaches are conducted in this direction [24, 25]. Shen et.al [25] combines Latent Semantic Index (LSI) (c.f. [23]) model to search semantic relevant documents in P2P network, by comparing users query and documents at the concept level, not just matching the keywords. Balke et.al [24] still uses SVM method, but create a novel indexing technique that allows to query using collection-wide information with respect to different classifications.

As a summary, Table 2 illustrates a preliminary result of our discussion. It is not a complete one but can be served as a stepping stone for P2P infrastructures selection.

Table 2. Preliminary Results

Scale	metadata elements	Semantic support	Autonomy	Adaptable P2P Network	Info. Srching Technique
small	few	No	high	pure P2P, RI	Information retrieval (IR)
small	few	No	low	pure P2P, Central server-based P2P, DHT	IR
small	many	No	high	pure P2P, RI	XML-based IR, RDF database
small	many	Yes	high	pure P2P, RI	RDF database
small	many	No	low	pure P2P, Central server-based P2P	XML-based IR, database
large	few	No	high	Super-Peer, SON	IR
large	few	No	low	DHT, Central server-based P2P	IR
large	many	No	high/low	Super-Peer, SON	XML-based IR, database
large	many	Yes	high	Super-Peer, SON	RDF database + RQL
large	many	Yes	low	Super-Peer, SON, DHT + logical layer	RDF database + RQL

3 Conclusion and Future Work

In summary, determining appropriate infrastructures for P2P-based digital libraries needs a consolidated guideline. This paper compares some representative P2P systems and aims to clarify advantages and weaknesses in applying different topologies. A discussion based on information searching is conducted and leads to preliminary results which are highly necessary for the future research. The paper can serve as a stepping stone for deciding architectures and techniques in the context of P2P-based digital library applications.

References

1. : DELOS Digital library architecture (WP1)- Cluster Objectives. http://www.delos.info/WP1.html (2004)
2. Clarke, I., Sandberg, O., Wiley, B., Hong, T.W.: Freenet: A distributed anonymous information storage and retrieval system. In: International Workshop on Design Issues in Anonymity and Unobservability, LNCS 2009. Volume 2009 of Lecture Notes in Computer Science., Berkeley, CA, USA, Springer (2000) 46–66
3. Crespo, A., Garcia-Molina, H.: Routing indices for peer-to-peer systems. In: Proceedings of the 22 nd International Conference on Distributed Computing Systems (ICDCS'02). (2002)
4. Jain, S., Mahajan, R., Wetherall, D.: A Study of the Performance Potential of DHT-based Overlays. In: USENIX Symposium on Internet Technologies and Systems 2003, Seattle, Washington, USA (2003)
5. Ion Stoica, Robert Morris, D.K., Kaashoek, F., Balakrishnan, H.: Chold: A scalable peer-to-peer lookup service for internet applications. In: SIGCOMM, California, USA (2001)
6. Ratnasamy, S., Francis, P., Handley, M., Karp1, R., Shenker, S.: A scalable content-addressable network. In: SIGCOMM, San Diego, California, USA (2001)

7. Tang, C., Xu, Z., Mahalingam, M.: psearch: Information retrieval in structured overlays. In: Proceedings of HotNetsI02, Princeton, New Jersey, USA (2002)
8. Halevy, A.Y., Ives, Z.G., Mork, P., Tatarinov, I.: Piazza: data management infrastructure for semantic web applications. In: WWW. (2003) 556–567
9. Oaks, S., Traversat, B., Gong, L.: JXTA in a Nutshell. O'Reilly & Associates, Inc. (September, 2002)
10. Nejdl, W., Wolf, B., Qu, C., Decker, S., Sintek, M., Naeve, A., Nilsson, M., Palmer, M., Risch, T.: EDUTELLA: a P2P networking infrastructure based on RDF. In: International World Wide Web Conferences (WWW). (2002) 604–615
11. Haase1, P., Broekstra3, J., Ehrig, M., Menken, M., et.al.: Bibster:a semantics-based bibliographic peer-to-peer system. In: The 3rd International Semantic Web Conference (ISWC2004), Japan (2004)
12. Schlosser, M.T., Sintek, M., Decker, S., Nejdl, W.: Hypercup - hypercubes, ontologies, and efficient search on peer-to-peer networks. In: AP2PC. Volume 2530 of Lecture Notes in Computer Science., Springer (2002) 112–124
13. Ahlborn, B., Nejdl, W., Siberski, W.: OAI-P2P: A Peer-to-Peer Network for Open Archives . In: 2002 International Conference on Parallel Processing Workshops (ICPPW'02), Vancouver, B.C., Canada (2002)
14. Cai, M., Frank, M.: Rdfpeers: A scalable distributed rdf repository based on a structured peer-to-peer network. In: WWW. (2004) 17–22
15. Aberer, K., Cudré-Mauroux, P., Datta, A., Despotovic, Z., Hauswirth, M., Punceva, M., Schmidt, R.: P-grid: a self-organizing structured p2p system. SIGMOD Record **32** (2003) 29–33
16. Crespo, A., Garcia-Molina, H.: Semantic overlay networks for p2p systems. Technical report, Computer Science Department, Stanford University (2002)
17. Yang, B., Garcia-Molina, H.: Designing a super-peer network. In: IEEE International Conference on Data Engineering, 2003. (2003) http://dbpubs.stanford.edu:8090/pub/showDoc.Fulltext?lang=en&doc=2003-33&format=pdf&compression=. (Available: 2004.07).
18. Greg Karvounarakis, e.: The RDF Query Language (RQL). http://139.91.183.30:9090/RDF/RQL/ (2003)
19. : Sesame: Rdf schema querying and storage. http://www.openrdf.org (2005)
20. Calvanese, D., Damaggio, E., Giacomo, G.D., Lenzerini, M., Rosati, R.: Semantic data integration in p2p systems. In: DBISP2P. (2003) 77–90 http://springerlink.metapress.com/openurl.asp?genre=article&issn=0302-9743&volume=2944&spage=77.
21. Ding, H., Solvberg, I.: Towards the schema heterogeneity in distributed digital libraries. In: 6th International Conference on Enterprise Information System (ICEIS), Portugal (2004)
22. Olmedilla, D., Palm′er, M.: Interoperability for peer-to-peer networks: Opening p2p to the rest of the world. In: WWW'05, Chiba, Japan (2005)
23. Baeza-Yates, R., Ribeiro-Neto, B.: Modern Information Retrieval. Addison-Wesley, USA (1999)
24. Balke, W.T., Nejdl, W., Siberski, W., Thaden, U.: Dl meets p2p - distributed document retrieval based on classification and content. In Rauber, A., Christodoulakis, S., Tjoa, A.M., eds.: ECDL. Volume 3652 of Lecture Notes in Computer Science., Springer (2005) 379–390
25. Shen, H.T., Shu, Y., Yu, B.: Efficient semantic-based content search in p2p network. IEEE Trans. Knowl. Data Eng. **16** (2004) 813–826