PeerDB: A P2P-based System for Distributed Data Sharing

Wee Siong Ng[1]   Beng Chin Ooi[1]   Kian-Lee Tan[1]   Aoying Zhou[2]

[1]Dept Computer Science
National University of Singapore
3 Science Drive 2, Singapore 117543

[2]Dept Computer Science and Engineering
Fudan University
Shanghai, China, 200433

Abstract

In this paper, we present the design and evaluation of PeerDB, a peer-to-peer (P2P) distributed data sharing system. PeerDB distinguishes itself from existing P2P systems in several ways. First, it is a full-fledge data management system that supports fine-grain content-based searching. Second, it facilitates sharing of data without shared schema. Third, it combines the power of mobile agents into P2P systems to perform operations at peers' sites. Fourth, PeerDB network is self-configurable, i.e., a node can dynamically optimize the set of peers that it can communicate directly with based on some optimization criterion. By keeping peers that provide most information or services in close proximity (i.e, direct communication), the network bandwidth can be better utilized and system performance can be optimized. We implemented and evaluated PeerDB on a cluster of 32 Pentium II PCs. Our experimental results show that PeerDB can effectively exploit P2P technologies for distributed data sharing.

1 Introduction

Peer-to-peer (P2P) technology, also called peer computing, is an emerging paradigm that is now viewed as a potential technology that could re-architect distributed architectures (e.g., the Internet). In a P2P distributed system, a large number of nodes (e.g., PCs connected to the Internet) can potentially be pooled together to share their resources, information and services. These nodes, which can both consume as well as provide data and/or services, may join and leave the P2P network at any time, resulting in a truly dynamic and ad-hoc environment. The distributed nature of such a design provides exciting opportunities for new killer applications to be developed.

Many domain specific P2P systems have already been deployed [19], e.g., Freenet [7], Gnutella [8], Napster [17], ICQ [11], Seti@home [20] and LOCKSS [15]. However, most of these P2P systems are limited in several ways. First, they provide only file level sharing (i.e., sharing of the entirety of a file) and lack object/data management capabilities and support for content-based search. Unlike existing work on distributed data management, data may be shared without predefined schema! Second, they are limited in extensibility and flexibility. As such, there is no easy and rapid ways to extend their applications quickly to fulfill new users needs. Third, a node's peers are typically statically defined.

In this paper, we present PeerDB, a P2P-based system for distributed data sharing. PeerDB has several distinguishing features. First, each participating node is a full fledge object management system that supports content-based search. Second, in PeerDB, users can share data without a shared global schema. Third, PeerDB adopts mobile agents to assist in query processing. Since agents can perform operations at the peers' sites, the network bandwidth is better utilized. More importantly, agents can be coded to perform a wide variety of tasks, making it easy to extend the capabilities of a PeerDB node. For example, an agent may further manipulate the data retrieved from a node to ship back only summarized data, or filter away uninterested objects. Finally, PeerDB supports mechanisms to dynamically keep promising (or best) peers in close proximity based on some criterion. For example, peers that are most frequently accessed are directly communicable while nodes that are less frequently accessed can be reached through peers. This significantly reduces the response time to queries.

We implemented PeerDB, a prototype P2P distributed object management system that incorporates all the above features. To evaluate PeerDB, we propose a systematic methodology for evaluating P2P systems. Our methodology considers both efficiency

and effectiveness (quality of answers) of P2P systems. We conducted our experiments on a cluster of 32 Pentium II PCs. Our experimental results show the effectiveness of PeerDB for distributed data sharing.

The rest of this paper is organized as follows. In the next section, we discuss what a P2P distributed data management system is. Section 3 gives a quick overview of BestPeer, a P2P platform in which PeerDB is based upon. In Section 4, we present the PeerDB design and its features. Section 5 presents an extensive experimental study to evaluate PeerDB. In Section 6, we review some related works, and finally, we conclude in Section 7 with directions for future work.

2 P2P Distributed Data Management: What is it?

As noted in the introduction, practically all existing P2P systems are designed to support sharing of data at a coarse granularity (e.g., files, documents). In this section, we first distinguish between P2P systems and distributed database systems. We then "define" P2P distributed data management by looking at three examples (due to space constraints) of how P2P technology can be employed for distributed database applications. This will also serve to motivate the need for database technology in P2P systems.

2.1 P2P vs Distributed Database Systems

There are several features that distinguish P2P systems from distributed database systems (DDBS).

1. In P2P systems, nodes can join and leave the network anytime. In DDBS, nodes are added to and removed from the network in a controlled manner, i.e., when there is a need for growth or retirement.

2. In P2P systems, there is usually no predetermined (global) schema among nodes. Queries are largely based on keywords. There are several reasons for this. First, most of the current applications do not require a fixed schema. (Napster is one exception where data is shared with a fixed schema, the one that describes music files.) Second, as nodes can join and leave the network at anytime, a fixed schema does not reflect the actual information that may be available at a single time. In DDBS, nodes are typically stable and have some knowledge of a shared schema.

3. In P2P systems, nodes may not contain the complete data. Further, nodes may not be connected. Thus, answers to queries are typically incomplete. By "completeness", we mean all answers that satisfy a query. In DDBS, one expects and can actually retrieve the complete set of answers.

4. In P2P systems, content location is typically by "word-of-mouth", i.e., a node routes its query to its neighboring nodes, and so on. In DDBS, the exact location to direct the query is typically known.

Based on the above points, we do not consider data integration systems to be P2P distributed data management systems (even if each node has the capabilities to act as middleware and server).

2.2 P2P Distributed Data Management Systems Applications

Instead of formalizing the concept of P2P distributed data management systems, we show with sample applications on what such systems may be like.

Health Care

In a hospital, each specialist has a group of patients that are solely under his care. While some patient data are stored in a centralized server of the hospital (e.g., name, address, etc), other data (e.g., X-rays, prescription, allergy to drugs, history, reaction to drugs, etc) are typically managed by the specialist on his personal PC. For most of these patients, the specialist is willing to share their data, but there are always some cases that he is unwilling to share for different reasons (e.g., part of his research program on a new drug, etc). By making the sharable patient data available to other specialists, it allows them to look for other patients who may have similar symptoms as their own patients, and hence can help them in making better decisions on the treatment (e.g., drugs to prescribe, reactions to look out for, etc).

Here, we can deploy a P2P distributed management system: (1) any specialist can join/leave the network; (2) the answers need not be complete (i.e., missing data from some specialists is not critical), (3) nodes have to search for content as in P2P systems, (4) the schema defined by each specialist may be different, (5) there is a need for data management, and (6) each specialist has something to share and is also interested in others' data.

Genomic Data

The discovery of new proteins necessitates complex analysis in order to determine their functions and classifications. The main technique that scientists use in determining this information has two phases. The first phase involves searching known protein databases for proteins that "match" the unknown protein. The second phase involves analyzing the functions and classifications of the similar proteins in an attempt to infer commonalities with the new protein. While there are several known servers on genomic data (e.g., GenBank, SWISS-PROT and EMBL), there are many more data that are produced each day in the many laboratories all over the world. These scientists create their own local databases of their newly discovered proteins and results, and are willing to share their findings to the world! Clearly, this is an application for P2P distributed data management systems for the same reasons as the health care application.

Data Caching

In the above two examples, each participant is *actively* involved in the process of consuming and supplying data. P2P distributed data management can also be deployed in *passive* nodes: nodes that are used to share resources (storage or computational power) on data that they may or may not be interested. Caching results from earlier queries is one such example - a node may have issued a query to some server (e.g., a data warehouse), the results of the query can be cached on the node (or some other neighboring nodes). In this way, another node that requests for data that overlap the query result can potentially obtain partial answers quickly from this node, and the remainder from the original server. This also lightens the load on the original server. Indeed, Kalnis et. al. have shown how distributed caching can be deployed in P2P environments to speed up OLAP queries [12].

3 BestPeer: An Adaptive Platform for P2P Applications

PeerDB is a database application that is implemented on top of BestPeer [1, 18]. We shall briefly review the features of BestPeer in this section.

BestPeer is a generic P2P system designed to serve as a platform on which P2P applications can be developed easily and efficiently. The network consists of two types of entities: a large number of computers (nodes), and a relatively fewer number of *location independent global names lookup* (LIGLO)

servers. Each participating node runs the BestPeer (Java-based) software and will be able to communicate or share resources with any other nodes (i.e., peers) in the BestPeer network.

BestPeer has several features that distinguish itself from existing P2P systems. First, BestPeer integrates two powerful technologies: mobile agents and P2P technologies. While P2P technology provides resource sharing capabilities amongst nodes, mobile agents technology further extends the functionalities. In particular, since agents can carry both code and data, they can effectively perform any kind of functions. With mobile agents, BestPeer not only provides files and raw data, it also provides processed information (e.g., summaries). Mobile agents can also be dispatched to collect information (e.g., what files/content are sharable, statistics, etc.) on the entire BestPeer network.

Second, BestPeer not only facilitates a finer granularity of data sharing where partial content of a file may be shared, it also shares computational power. The requester sends his/her request for a file together with an algorithm (executable code) that operates on the file. In other words, the requester performs the filtering task at the provider's end! Third, a node in the BestPeer network can dynamically reconfigure itself by keeping peers that benefit it most (subject to individual node's definition of 'most benefit').

Finally, BestPeer introduces a Location-Independent Global Names Lookup Server (LIGLO) to provide each node with a unique global identity. In this way, nodes that may have different IP address can be "recognized" as a single unique entity. LIGLO is a node that has a fixed IP and running Location-Independent Global Names Lookup Server software. It provides two main functions: generates a BestPeer Global Identity (BPID) for a peer and maintains peer's current status, such as the current IP address and whether the peer is currently online or offline (if this information is available).

Due to the security risks posed by such a potentially powerful platform, the basic BestPeer platform provides a secure access to a node's computing resources. Each node comprises two types of data, private data and sharable data. Nodes can only access data that are sharable. This is enforced by a security policy that restricts applications to user-specified locations established during platform initialization. Communications between nodes have also been provided with 128 bit encryption to protect the sensitive data from being eavesdropped and viewed as they travel through the BestPeer network.

4   Peering Up for Distributed Data Sharing

In this section, we will present PeerDB, a prototype P2P-based system for distributed data sharing. PeerDB's P2P-enabling technologies are provided by BestPeer [1, 18]. However, it extended BestPeer in the following ways. First, data in each node is managed by a database system. In other words, PeerDB is a network of database-enabled nodes. Second, data can be shared without a global schema. Third, query processing is assisted by mobile agents. Fourth, each node can reconfigure itself based on some optimization criterion from the answers returned. We shall discuss these features here.

4.1   Architecture of a PeerDB Node

Figure 1 illustrates the internals of a PeerDB node. There are essentially four components that are loosely integrated. The first component is a data management system that facilitates storage, manipulation and retrieval of the data at the node. We have used the MySQL [2] which is a popular Open Source Database as our storage server. Thus, the system can be used on its own as a stand alone DBMS outside of PeerDB. We note that the interface of the data management system is essentially an SQL query facility. For each relation that is created, the associated meta-data (schema, keywords, etc) are stored in a **Local Dictionary**. There is also an **Export Dictionary** that reflects the meta-data of objects that are sharable to other nodes. Thus, only objects that are exported can be accessed by other nodes in the network. We note that the meta-data associated with the Export Dictionary is a subset of those found in the Local Dictionary, and the distinction here is a logical one (as the actual implementation minimizes redundancy). We shall defer the discussion on how the Export Dictionary will be used, and the details on the meta-data when we addressed the query processing strategy.

The second component is a database agent system called DBAgent. DBAgent provides the environment for mobile agents to operate on. Each PeerDB node has a *master agent* that manages the query of the user. In particular, it will clone and dispatch *worker agents* to neighboring nodes, receive answers and present them to the user. It also monitors the statistics and manages the network reconfiguration policies.

The third component is a cache manager. We shall defer the discussion of the cache manager to a later subsection. Here, it suffices for us to know that we are dealing with caching remote data in secondary storage, and the cache manager determines the caching/replacement policy.

The last component is the user interface. This provides a user-friendly environment for user to submit their queries, to maintain their sharable objects, and to insert/delete objects. In particular, users search for data using SQL-like queries.
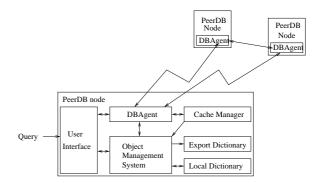

Figure 1. PeerDB node architecture

4.2   Sharing Data Without Shared Schema

One of the main objectives of PeerDB is to allow users to manage their (private and sharable) data using a database management system (DBMS). However, as noted, there is no predetermined and uniform schema that nodes share. Unless users interact with one another somehow, we can expect data to be defined differently by different users even if they may have interests in data from a common domain. For example, in naming a relation, a genome scientist may call his set of protein database by the protein name (e.g, kinases, annexin) while another may name them after the species (e.g, mouse, human, zebrafish). Similarly, at the attribute level, one scientist may call the length of sequences as `length` while another will use the term `len`. A more complicated problem would be for one to create a single "universal" schema, while another may "normalize" his database to multiple tables. Thus, it is difficult to locate data if the traditional method of exact matching of relation names/attributes is used.

To address this issue, we adopt an Information Retrieval (IR) [4] based approach. For each relation that is created by the user, meta-data are maintained for each relation name and attributes. These are essentially keywords provided by the users upon creation of the table, and serve as a kind of synonymous names. (One can think of this as a miniature thesaurus.) Continuing with our examples, for a table of Kinases proteins, while the relation name may be `Kinases`, the keyword `protein` will be useful during searching. Similarly, two users defining length of a sequence as `len` and `length` re-

spectively are likely to have the common keyword `length`. In this way, potentially relevant data can be determined using the following *relation-matching* strategy:

- Consider a query (R, A, C) where R is the set of relations A is the set of target attributes, and C is the set of conditions. (This corresponds to a simple SPJ query in SQL.) Let V denote all attributes that appear in A and C.

- R is searched against keywords for relation names; and V is searched against keywords for attribute names. Note that this search involves looking for matching keywords of R against keywords for other relations; the same holds for attributes. The result of this search process will be a list of relations whose relation name keywords match R (or their keywords) and/or attribute keywords match attribute names in V (or their keywords).

- Given a query Q of the form (R, A, C) and a relation D with attributes T, the degree in which D matches Q can be computed as follows:

$$Match(Q, D) = \frac{(wt_r * r) + \left(wt_a * \underset{match}{N}(A \cup C, T)\right)}{wt_r + (wt_a * N(A \cup C))}$$

where wt_r and wt_a are weights assigned to reflect the importance of matching relation and attribute names respectively. r takes value of 1 or 0; r is 1 if and only if D and R share some common matching keywords, otherwise, r is 0. $\underset{match}{N}(A \cup C, T)$ refers to the total number of matching keywords between attributes involved in Q and those of D. $N(A \cup C)$ indicates the total number of distinct keywords for attributes in Q. The set of relations that potentially contain answers to Q are those that have scores above a certain threshold value.

With the above strategy to locate matching relations, we note that we can share data without explicit sharing of schema. This flexibility is also an important distinction between PeerDB and existing distributed DBMS. Note that the relations and meta-data will be returned to the user first, who will then decide the data that is of interests (see the next section on query processing strategies).

We illustrate the strategy with an example. Suppose we have four peers that share genomic data. Peer P1 defines a relation Kinases(SeqID, length, proteinSeq). Peer P2 defines a relation Protein(SeqNo, len, sequence). Peer P3 defines two relations ProteinKLen(ID, seqLength) and ProteinKSeq(ID, sequence). Peer P4 defines a relation

Protein(name, char). Figure 2 shows the keywords defined for these relations by the various peers. Suppose the user at peer P1 (he knows his own schema but not the schema of other peers) issues the following query to look for kinases sequences that are longer than 30 base pairs:

```
SELECT SeqId, proteinSeq
FROM Kinases
WHERE length > 30;
```

Now, since one of the keyword for Kinases (relation name) is protein, and protein is also a keyword for P2's relation `Protein` and P3's relations `ProteinKLen` and `ProteinKSeq`, these relations match the query relation. Similarly, we find that the attributes `SeqID`, `proteinSeq` and `length` all have matching keywords in P2 and P3. For P3, we note that the query may have to be turned into a join query when evaluated there. For P4, we only have a match in relation name but not in the attributes. Thus, P4 will be ranked lower than P2 and P3. Semantically, we note that P2's data are not actually those that P1 is interested in (since they are not Kinases data). As such, it is important to have the meta-data and additional information returned to the users before fetching the data.

Peer	Names	Keywords
P1	Kinases	protein, human
	SeqID	key, identifier, ID
	length	length
	proteinSeq	sequence, protein sequence
P2	Protein	protein, annexin, zebrafish
	SeqNo	number, identifier
	len	length
	sequence	sequence
P3	ProteinKLen	protein, kinases, length
	ID	number, identifier
	seqLength	length
	ProteinKSeq	protein, sequence
	ID	number, identifier
	sequence	sequence
P4	Protein	protein, kinases, annexin, . . .
	name	name
	char	characteristics, features, functions

Figure 2. Keywords for the relations/attributes names.

4.3 Agent Assisted Query Processing

In PeerDB, we adopt a two-phase query processing strategy. In the first phase, the relation matching strategy is applied to locate potential relations.

These relations are then returned to the query node for two purposes. First, it allows the user to select the more relevant relations. This is to minimize information overload when data may be syntactically the same (having the same keywords) but semantically different. Moreover, this can minimize transmitting data that are not useful to the user, and hence better utilize the network bandwidth. Second, it allows the node to update its statistics to facilitate future search process. Phase two begins after the user has selected the desired relations. In phase two, the queries will be directed to the nodes containing the selected relations, and the answers are finally returned.

PeerDB's query processing is completely assisted by agents. In fact, it is the agents that are sent out to the peers, and it is the agent that interacts with the DBMS. Moreover, a query may be rewritten into another form by the DBAgent (e.g., a query on a single relation may be rewritten into a join query involving multiple relations). To elaborate on the query processing strategy, we shall distinguish two types of queries: *local* query and *remote* query. A query is local to a node if it is initiated there, and remote otherwise.

4.3.1 Processing Local Query

When a user issues a query (SQL-like selection query), a master agent will be created to oversee the evaluation of the query. The following operations are performed by the agent:

Phase I

- The agent "parses" the query to extract the list of relations and attributes names.

- The relation matching strategy is applied on the local dictionary. Promising relations can then be returned to the user immediately.

- At the same time, the master agent will clone *relation matching agents* and dispatch them to all neighbors of the node. Besides the query, the agent also carries with it two other information: (a) IP address of the node that initiates the query; (b) TTL (Time-to-live). The former is needed to allow remote nodes to return answers directly to the query node. The latter indicates the lifetime of an agent. This allows the process of cloning and forwarding to keep on going until the agent lifetime is expired.

- The master agent will wait for the answers (relations schema) from remote nodes. Upon receiving any answers, they will be returned to the users for selection.

- For peers that return multiple relations, we return the individual relations (if their scores on the number of matching keywords exceed the threshold) as well as combinations of relations that are related (e.g., has a key-foreign key relationship). Referring to our earlier example, P3 will produce three answers: `proteinKLen`, `proteinKSeq`, and `proteinKLen ⋈ proteinKSeq`.

Phase II

- For each relation selected by the user, the master agent will clone a *data retrieval agent*. for that relation. One of the first tasks of the agent is to reformulate the query so that it matches the relation name and attributes at the target node. Clearly, it is possible that some attributes may be dropped because the target relation has no such matching attributes. For combination of relations, the data retrieval agent will also rewrite the orginal query into a join query involving the combination of relations.

- If the target relations are found locally, the worker agent will submit a reformulated SQL query to the DBMS to retrieve the data. The data is then returned to the agent, formulated for output and returned to the user.

- If the target relations are on a remote node, then the worker agent will be dispatched with the query node's IP address. Answers will be returned directly from the remote host to the master agent who will then formulate and return the answers to the user.

We note that the two phases are only logical. In fact, as soon as relations are returned, they are shown to the user, and the user can start selecting relations; and as soon as a relation (or combination of relations) is selected, the agent is sent out to retrieve the data. In this way, answers are returned progressively (without long waiting time). Moreover, users could be viewing answers (data) while there may be other agents still searching the PeerDB network for candidate relations.

4.3.2 Processing Remote Query

As mentioned, essentially, for a remote query, it is an agent that arrives at the node.

Phase I: Relation matching agent

- If the agent has not visited the node previously, the TTL value is reduced by one.

- The agent will search the **export dictionary**. Promising relations are then returned to the query node at the IP address provided by the agent.

- If TTL > 0, the agent will clone more relation matching agents and dispatch them to the neighbors of the current node; otherwise, the agent will be dropped.

Phase II: Data Retrieval Agent

- The agent will formulate an SQL query and submit it to the DBMS.

- Once the answers are retrieved, they are returned to the query node directly. If the retrieved data need to be further processed before being returned, then the agent will perform the task (with the code that it carries along) and return the summarized data.

- The agent may then be dropped.

4.4 Monitoring Statistics

One of the tasks of the master agent is to perform the reconfiguration of the network based on a reconfiguration policy selected by the user. The master agent monitors two types of statistics. The first is the relation information obtained from the first phase of the query processing strategy. In particular, the keywords of selected relations may be "exchanged" to update the meta-data. The second is the number of answer objects obtained from the selected relations. This can be used to determine the nodes to be connected directly.

PeerDB also extended BestPeer with a temporal locality based reconfiguration policy that favors nodes that have most recently provided answers. It uses the notion of stack distance to measure the temporal locality. The idea works as follow. Consider a stack that stores all the peers that return results. For each peer that returns answers, move the peer to the top of the stack, and push the existing peers down. The temporal locality of a peer is thus determined by its depth in the stack. The top k peers in the stack are retained as the k directly connected peers, where k is a system parameter that can be set by the node.

4.5 Cache Management

PeerDB supports caching of answers returned from remote nodes in order to reduce the response time for subsequent answers. For every relation that the user retrieved (in phase II of the query processing strategy), we cache the answers. Caching raises
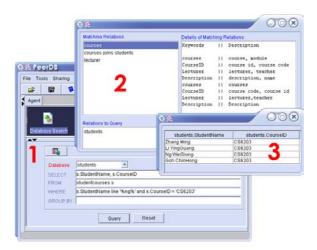

Figure 3. PeerDB interface.

many complicated issues. We looked at three of them here. First, the cached copy may be outdated. To handle this, PeerDB only keeps the answers for a fixed period of time, after which the cache is invalidated. Second, since storage space is limited, we adopt a LRU replacement policy: whenever we run out of disk space, we replace the cache that is least recently used. Finally, in a P2P environment, many PeerDB nodes may be caching the same data. As such, a search may give rise to multiple "copies" of the same data. While this is a semantic issue that is to be left to the user, we attempt to minimize the effort as follows. For each cached relation, We also maintain the information on the BPID of the source node (recall that each node has a unique identifier BPID provided by BestPeer technology). When a node is not the source of a relation, its response to a search will also include the BPID of the source node. All relations, except one, with the same keywords from the same source node will be pruned away during phase I of query processing.

5 A Performance Study

We implemented the PeerDB software with the features discussed in the previous sections. Any node that installs the PeerDB software and registers with a predetermined set of LIGLO servers can participate in the PeerDB network. Figure 3 shows the PeerDB interface - window 1 shows the query interface, window 2 displays the results of matching schemas, and window 3 displays the answer tuples from a selected relation.

In this section, we report an extensive performance study conducted to evaluate PeerDB. We study PeerDB in two aspects. The first studies its relation matching strategy, and the second looks at

the performance of PeerDB.

5.1 On Relation Matching Strategy

In this section, we present the experimental results of PeerDB in the search for matching relations in P2P environment without a global schema.

We generated a large number of relations as follows. First, we created a set of semantically equivalent categories, C. In each category, we have c keywords which are assumed to represent the same semantic meaning, i.e., any two keywords refer to the same meaning. Next, we created a set of relations, and each relation is assigned 2-5 keywords (since users are not expected to enter too many keywords) selected randomly from an arbitrary category picked from C. Each relation has a number of attributes, and each attribute is also assigned 2-5 keywords picked randomly from an assigned category from C.

The following query form has been used in our experiment.

```
SELECT attribute_X
FROM relation_i
WHERE attribute_Y = value_1
      and attribute_Z > value_2;
```

We use the standard precision and recall measures as the performance metrics. Precision measures the purity of search results, or how well a search avoids returning results that are not relevant; recall refers to completeness of retrieval of relevant items. We consider a relation to be relevant to the query if more than k keywords from the relation names match. In our study, we have set k to be 2. This set forms the basis for the computation of precision and recall.

For each relation that is examined, we compute its matching score. We varied the threshold value from 0.1 to 0.9. For all the results returned, we compute its precision and recall. The results are shown in Table 1.

Threshold	Precision	Recall
0.1	0.33	0.85
0.3	0.36	0.78
0.5	0.50	0.57
0.7	1.00	0.28
0.9	1.00	0.21

Table 1. Precision and recall for varying threshold values.

As shown in Table 1, when the threshold value is large, the precision is high as most of the relevant

relations can be identified. In fact, in this experiment, we have 100% precision when threshold is 0.7 and above. However, recall is low because of the large number of irrelevant relations that share some common keywords. These results are consistent with typical IR search results, showing that the proposed strategy is effective.

5.2 On PeerDB Performance

To evaluate PeerDB's performance, we conducted different sets of experiments. We first compare PeerDB against the Client/Server (CS) Architecture. The basic difference between the two models is that in a P2P model the interacting processes can be a client, server or both while in a CS model one process assumes the role of a service provider while the other assumes the role of a service consumer. Our CS model has some flavors of P2P in that a node can be both a client and a server. However, like CS model, the server must return its result to the client - as such the results must be returned along the query path. We study two versions of PeerDB - a static PeerDB where the reconfigurable feature is turned off, and a dynamic PeerDB with the reconfiguration feature turned on. We compare both schemes with the CS architecture. This allows us to see the benefits of the reconfiguration scheme. We shall denote these two schemes as PDMS and PDMR respectively. We also compare PeerDB with pure message-passing based protocol and agent-based protocol. The objective in this experiment is to show the cost and effect of using agents in PeerDB. Before we look at the experiments and findings, we shall propose an evaluation methodology for P2P-based systems.

5.2.1 Evaluation Methodology

Any system has to be evaluated based on its efficiency and effectiveness. The former deals with the performance issue, while the latter deals with the quality of the answers. Unlike existing distributed systems, there is no clear criteria on how P2P systems should be evaluated. Like Internet search engine, the answers to queries depend on the peers that are searched, which may not include every peer in the P2P network. In addition, every query may involve different peers (since peers change over time)!

For purpose of evaluation, a controlled environment is necessary. We propose that the following three scenarios be evaluated. First, different schemes should be evaluated based on a fixed set of nodes. This can be useful for a set of nodes that exploit P2P technology to facilitate collaboration, i.e., it is essentially a traditional distributed environment where all nodes participate in answering a query. Here, we can

study how different P2P protocols or reconfiguration strategies perform.

Second, in a P2P network, the rate at which answers are returned are important. This is because the users have no idea of which peers will be providing the answers to his/her queries, and how many peers will be searched. A long initial waiting time is not likely to be acceptable to the users.

Third, the quality and quantity of the answers returned are important measures too. A node may return answers quickly, but it may return only very few answers or answers that are not very relevant to the query. While quality is based on the semantics of the query, quantity of answers is easy to obtain and use as a performance metrics.

5.2.2 Experimental Setup

The experimental environment consists of 32 PCs with Intel Pentium 200MHz processor and 64M of RAM, and all the PCs are running on WinNT4.0 operating system. The physical network layout is shown in Figure 4.
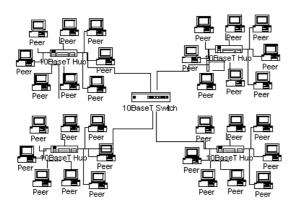

Figure 4. Experimental environment.

There are a total of 10,000 objects, each of which is 10 KB. These data are randomly assigned to nodes, such that each node holds 1,000 objects.

In practice, we expect users to be interested in part of the entire datasets only. There will always be some data that are of no interests to them, and will never be accessed by them. For example, in the case of Napster, while there are always classical music being shared, some user who prefers contemporary music may just dislike them totally. In our experiments, we try to model this by dividing the queries of each user as follows: (a) x% of queries are directed at 'hot' data in the entire dataset. These hot data are also frequently accessed by other users; (b) y% of queries are directed at z% of the cold data. This model the case that individual user may have their own taste on cold data. (c) the remaining queries
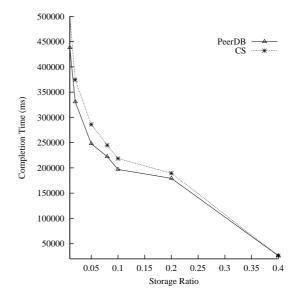

Figure 5. Effect of storage capacity.

are directed at the remaining cold data. As default, we set x to be 80%, y to be 15%, and z to be 20%. Moreover, 20% of the data are hot, and 80% of them are cold.

The experiments were conducted when the machines and the network were fully dedicated. Moreover, each node is "warmed up" to fill out its local storage before we start to collect results on the experiments. The results presented correspond to the average of at least three different executions. The variance across different executions was not significant.

5.2.3 Effect of Storage Capacity on Caching

In the first set of experiments, we compare PeerDB with the CS architecture by varying the storage capacity of each peer. We define the storage ratio to be the size of the storage size of a node to the size of the objects stored at the central server. Figure 5 shows the effect of storage ratio on the response time. First, we observe that as the storage ratio increases, both methods' response time decreases. This is expected as more objects can be found in local and neighboring peers. This also clearly illustrates the benefits of sharing storage resources. Second, we note that PeerDB outperforms the CS model. This is expected as the CS model requires the answers to be returned via the search path.

5.2.4 PeerDB vs CS

In this experiment, we first evaluate the performance of PDMS, PDMR and CS on the rate at which answers are returned. The number of nodes is fixed at
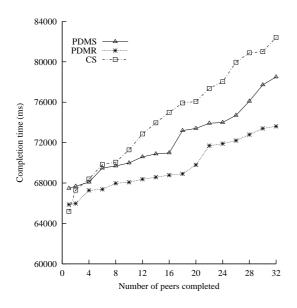
Figure 6. Rate of returning answers.

32. A search query is issued four times, and the average time at which nodes respond are noted. Figure 6 shows the results of the experiment. In the figure, the point (K, T) indicates that K nodes have responded after T time units. We note that it is possible that under different schemes, different nodes respond at different time and with different answers. We shall defer this discussion to the next experiment.

As shown in the figure, PDMR is still the best scheme, outperforming PDMS by virtue of its ability to reconfigure the network. It is able to reach out to more promising nodes directly - after each query, PDMR will reconfigure itself so that the next query can be directed to the more promising nodes first. We note that, except for the first few nodes, CS returns answers much slower than PDMR/PDMS - as it only returns answers along the path that the query has been transmitted.

Having a fast initial response time is not good enough. It is possible that nodes that return answers first provide very few answers. For the earlier experiments that study the initial response time, we also keep track of the number of answers that are provided by each node. Figure 7(a) shows a plot of the result. As shown, it is clear that CS returns the first few answers much faster than PDMS and PDMR. This is expected since the first few directly connected nodes that receive the query can return their answers immediately. For PDMS/PDMR, the overhead of the code-shipping strategy results in a longer initial response time performance. However, as more answers are returned, PDMS/PDMR are superior over CS, demonstrating the superiority of P2P technologies over traditional CS models. We also

note that PDMR is generally better than PDMS. In Figure 7(b), we can further confirm the effectiveness of PDMS/PDMR over CS. By the time PDMS and PDMR have received all the answers (100%), CS has only returned about 40% of the answers. As observed earlier, PDMR can generate more answers quickly by virtue of its ability to keep more relevant peers "closer".

5.2.5 Benefits of Agent-based Querying

In this experiment, we would like to study how much can be gained by using an agent-assisted query processing strategy. Here, we assume that the query requires some functions that is not supported by the DBMS. As such, the operation cannot be pushed down to the DBMS. Instead, the data have to be first retrieved, and the operation performed on the data before the answers to the query can be obtained.

In this experiment we show the cost and effect of using pure message passing protocol and agent based protocol in P2P environment. Here, we assume the query requires only one remote access. The whole process is divided into three phases: sending message (message-passing protocol) or agent (agent-based protocol) to remote host, remote host processes the request, and remote host returns the result to the originator. The answer size is set to be 0.1% of the whole data set. The difference between message based and agent-based protocol is that message-based protocol is a data-shipping strategy, i.e., remote data are transferred to the query node to be processed there. On the other hand, an agent-based protocol is a code-shipping strategy that carries the processing code to the remote site and performs remote execution. Only answers produced by the agent will be returned. The total response time including the cost of data transfer, i.e., message, code and data, and processing time. In Figure 8, we observe that the completion time of the message-based protocol increases exponentially when the data size increased. The overhead of the data-shipping results in a longer response time performance. As a result, when the number of data to be transferred across the network increases, the mobile agent-based protocol is superior.

Most of the network applications (client-server based or P2P based) require more than one communication with another node in order to complete each transaction. Therefore, in this experiment, we looked into the messages overhead in pure message-based protocol vs. agent-based protocol. In this experiment, multiple remote executions are required in order to answer a query Essentially, the query requests for multiple objects from a remote site; however, the query only knows the object to retrieve after
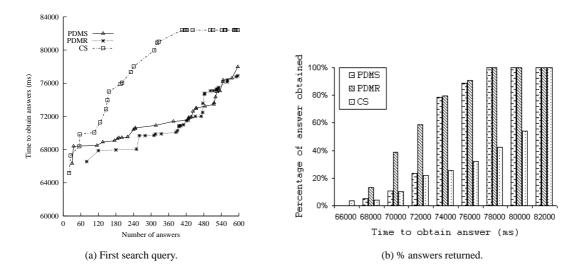
(a) First search query.

(b) % answers returned.

Figure 7. Number of answers returned.

the earlier requests are completed. Thus, we have a chain of queries and computations. In our test, each of the object requested will cause 5MB of data transfer across the network.
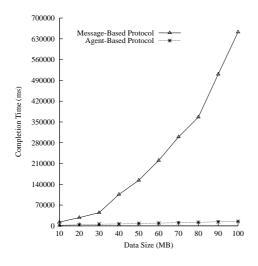

Figure 8. Completion time vs. data size

The result on multiple-communication transaction is shown in Figure 9. Clearly, the agent-based approach is superior. Under the message-passing protocol, the query is transmitted, and the data is returned to the node to be processed on the node. This has to be done before subsequent operations can be issued. On the other hand, in the mobile agent approach, all operations can be performed at remote node once the code is transmitted. Once the agent is constructed at the remote site, thy can interact with
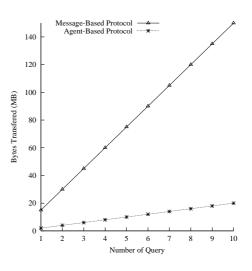

Figure 9. Communication overhead.

the remote node directly until the final result is obtained. Therefore, it optimizes the network resources and bandwidth.

6 Related Works

Agent-based systems have been studied in the literature [16, 5, 13, 14]. However, these technologies provide support for agent collaboration and communication but lack support for peer-to-peer technology. Development of P2P applications based on these platforms would require a longer development effort, which would be costly.

As mentioned, P2P technologies have been deployed in many applications [17, 8, 7, 6, 21, 18].

However, these systems provide coarse granularity of sharing without capabilities for DBMS supoort. In addition, they cannot be easily extended to meet users changing needs. Moreover, their network are statically defined and there is a lack of support for multi-granularity data access.

More recently, the database community has begun to exploit P2P technologies for database applications [9, 10, 12, 22]. In [9], data placement issues were addressed. In [22], the class of "hybrid" P2P systems where some functionality is still centralized is studied. In particular, an analytical model to describe the system performance is developed, and validated against actual hybrid P2P systems. Different architectures such as chained architecture, full replication architecture, hash architecture and unchained architecture were compared. In [12], caching of OLAP queries is addressed in the context of a P2P network. Halevy et. al. address the issue of schema mediation and proposes a language for mediating between peer schemas [10].

7 Conclusion

In this paper, we have presented a P2P-based distributed data sharing system called PeerDB. PeerDB has several nice features. First, it employs a data management system. Moreover, it facilitates data sharing without predetermined shared schema. Second, because its query processing capabilities is assisted by mobile agents, it provides easy extensibility to existing systems. Third, it provides a mechanism to reconfigure a node's peers based on some optimization criterion. Our extensive experimental studies show that PeerDB is a promising system for distributed processing. We plan to extend this work in two directions. First, we plan to make a node more intelligent by allowing it to determine at run-time which strategy to adopt - code-shipping or data-shipping. Second, we have focused on looking for "similar" schema. More recently, keyword-based search engine for relational databases has been developed [3]. We plan to see how such features can be integrated to facilitate keyword-based search in PeerDB.

Acknowledgements

References

[1] Bestpeer. In *http://xena1.ddns.comp.nus.edu.sg/p2p/*.

[2] MySql Home Page. In *http://www.mysql.com/*.

[3] S. Agrawal, S. Chaudhuri, and G. Das. Dbxplorer: A system for keyword-based search over relational databases. In *Proceedings of the 18th International Conference on Data Engineering*, San Jose, CA, April 2002.

[4] R. A. Baeza-Yates and B. A. Ribeiro-Neto. *Modern Information Retrieval*. ACM Press/Addison-Wesley, 1999.

[5] A. Castillo, M. Kawaguchi, N. Paciorek, and D. Wong. Concordia as enabling technology for cooperative information gathering. In *Proceedings of the 31th Annual Hawaii International Conference on System Sciences 1998 (HICSS31)*, 1998.

[6] R. Dingledine, D. Molnar, and M. J. Freedman. The free haven project: Distributed anonymous storage service. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, 2000.

[7] Freenet Home Page. *http://freenet.sourceforge.com/*.

[8] Gnutella Development Home Page. *http://gnutella.wego.com/*.

[9] S. Gribble, A. Halevy, Z. Ives, M. Rodrig, and D. Suciu. What can databases do for peer-to-peer. In *WebDB*, 2001.

[10] A. Halevy, Z. Ives, D. Suciu, and I. Tatarinov. Schema mediation in peer data management systems. In *Proceedings of the 19th International Conference on Data Engineering*, Bangalore, India, March 2003.

[11] ICQ Home Page. *http://www.icq.com/*.

[12] P. Kalnis, W. S. Ng, B. C. Ooi, D. Papadias, and K. L. Tan. An adaptive peer-to-peer network for distributed caching of olap results. In *ACM SIGMOD 2002*, pages 25–36, Madison , Wisconsin, USA, 2002.

[13] G. Karjoth, D. Lange, and M. Oshima. A security model for aglets. *IEEE Internet Computing*, 1(4), 1997.

[14] N. Karnik and A. Tripathi. Agent server architecture for the ajanta mbile-agent systems. In *International Conference on Parallel and Distributed Processing Techniques and Applications*, 1998.

[15] LOCKSS Home Page. *http://lockss.stanford.edu/*.

[16] Mitsubishi Electric. Concordia: An infrastructure for collaborating mobile agents. In *Proceedings of the 1st International Workshop on Mobile Agents (MA '97)*, April 1997.

[17] Napster Home Page. *http://www.napster.com/*.

[18] W. S. Ng, B. C. Ooi, and K. L. Tan. Bestpeer: A self-configurable peer-to-peer system. In *Proceedings of the 18th International Conference on Data Engineering*, page 272, San Jose, CA, April 2002 (Poster Paper).

[19] A. Oram, editor. *Peer-to-Peer : Harnessing the Power of Disruptive Technologies*. O'Reilly & Associates, 2001.

[20] SETI@home Home Page. *http://setiathome.ssl.berkely.edu/*.

[21] M. Waldman and L. Publius. A robust, tamper-evident, censorship-resistant, web publishing system. In *Proceedings of the 9th USENIX Security Symposium*, pages 59–72, 2000.

[22] B. Yang and H. Garcia-Molina. Comparing hybrid peer-to-peer systems. In *VLDB'2001*, 2001.