Research Statement
Yi Lin

Overview

As the rapid growth of the Internet, computer applications transit from desktops to network. Network applications target to support users in different geographical sites. The transition introduces many challenges. For example, how to enable users to access data fast even though they do that through the Internet? For another example, how to scale a system if there are more and more users? How to ensure the availability of a system in case of disasters like 911?

My research aims to provide answers to the above challenges for data-oriented applications. In general, my approach is to use database replication. Database replication replicates data to database servers which are distributed in different geographical sites. Thus, users can access data fast by only contacting local database server or servers with the smallest network distance. Since there are several copies of databases, availability of the data will be guaranteed as long as one of the database servers survives. By adding more servers, more clients should be able to be handled by the whole replicated database system.

Current research

One-Copy-Snapshot-Isolation in Replicated Databases

Since there are several copies of databases in a replicated database system, updating data in one database will lead to inconsistency of data in other databases. The challenge of database replication is how to preserve data consistency in case of update. For a long time, one-copy-serializability has been considered as standard consistency level for transactions in distributed databases, corresponding to serializability, the highest transaction isolation level in centralized databases. However, in recent years snapshot isolation has been used extensively in both commercial and opened-source products such as Oracle, PostgreSQL, SQL server 2005. Hence, most of exiting replication protocols providing one-copy-serializability are not applicable to these databases.

In cooperation with researchers at University Politecnica de Madrid, I proposed a new consistency level, one-copy-snapshot-isolation, for distributed databases, corresponding to snapshot isolation for centralized databases. According to snapshot isolation reads will neither be blocked by writes nor block writes, whereas reads and writes will block each other according to serializability. Since reads happen more frequently than writes in many database applications, one-copy-snapshot-isolation is more attractive than one-copy-serializabilty. This work is published in SIGMOD 2005[1] and has been cited more than 30 times according to Google scholar.

Middleware-based Data Replication in Wide Area Networks

In this research, I analyzed existing replication protocols and found that most existing replication protocols were designed for Local Area Networks (LANs) and did not work well in Wide Area Networks (WANs). The reason is that they require several message rounds within the response time of one transaction. Due to large message delay in WANs, this will lead to large response time for WANs applications. Moreover, the existing protocols put severe restrictions to applications.

Taking these restrictions and number of messages into account, I proposed two new protocols which require only one WAN message per update transaction and provide one-copy-snapshot-isolation. These two protocols are well suited for two typical WAN network topologies, i.e., a set of individual sites connected through a WAN and a set of clusters connected through a WAN, respectively. They do not put any restrictions on their applications. Moreover, the protocols are middleware-based and work for different databases. Part of this work is published in Euro-Par 2005 [2] and part of it is under review in WWW 2007[3].

Applying Database Replication to Multiplayer Online Games

Multiplayer Online Games (MOGs) are currently very popular distributed applications with billions dollar market. MOG requires strong system supports in terms of performance, scalability, and availability. These requirements are exactly the benefits provided by database replication. However, there is no existing work applying database replication into MOGs.

In this project, I investigate the feasibility and challenges of applying database replication into MOGs. I designed a small game for the proof of our concept. The game ran on top of a replicated database system built in previous projects. Many experiments were conducted and the results were satisfactory. Some pitfalls and challenges of applying database replication were discovered in the project, and further investigation is desirable. This work is published in Netgames 2006[4].

Future Research

As web has become an essential part of our life, data available in the web explode dramatically which leads to the increase of complexity in managing the data. Internet users expect faster and more reliable services of data management as data size is getting larger. Database replication will play a very important role in the future. I am thrilled by the potential of the research area.

Very Large Web Data Management with Caching and Replication

Multi-tier architectures are the most common infrastructures for current web-based applications. A standard way to reduce data flow between client and server, and between tiers is to cache data. One of challenges of caching data is how to preserve data consistency between cache and servers once the data in servers are changed. Weak consistency schemes allow cached data to be out-of-date while strong consistency keeps the cached data close to consistent with the data at the server.

I am interested to investigate how to manage very large web data so that queries can be answered quickly and with data as fresh as possible. For systems with huge amount of caches, I would like to investigate how caches can cooperate with each other in a peer-to-peer manner to answer queries.

Data consistency between caches and servers should be also considered. I have proposed one-copy-snapshot-isolation, a relatively strong consistency guarantee. I would like to look at the possibility of relaxing one-copy-snapshot-isolation for caching in very large web data management. From there, I will consider data consistency in caches and servers at the same time to provide a suitable or adaptive consistency guarantee.

Peer-to-peer data management

Peer-to-Peer (P2P) emerges as a new generation of network infrastructure. It distinguishes itself from traditional client-server architecture in large scalability, high dynamism and heterogeneity. Data management in P2P has attracted researchers in database areas recently. Due to large number of peers, data in the whole system cannot be kept consistent at all time. Due to high dynamism, data in some peers might be unreliable. My interest in this research is how to make data in P2P systems consistent and reliable as much as possible without losing performance too much. Data replication is desirable in P2P for high availability of data because peers leave and join systems dynamically. But data consistency has to be relaxed due to large number of peers. I plan to investigate how to move data replication from client-server to P2P.

Database Replication in Multiplayer Online Games

I learned from the previous project that MOGs are different from general database applications in that they are write intensive instead of read intensive. Currently most of database replication protocols do not scale well with write intensive applications. I plan to look at the possibility of designing a specific replication protocols for MOGs with massive number of users. Moreover, from the architecture point of view, I have looked at MOGs with client-server architecture. However, peer-to-peer MOGs are more interested because of massive number of users. I am interested in looking at the peer-to-peer data management support of MOGs.

Besides, there are some other interesting research topics such as distributed transactional memory. By applying database replication to MOGs, we modeled the synchronization of conflicting distributed objects by means of distributed transactions. Programmers only need to identify beginning and ending of transactions when their applications access atomic blocks. Distributed transactional memory applies transaction to concurrency control in memory. Since MOGs are multithread applications which are intensive in concurrency control, I am interested to investigate distributed transactional memory in MOGs.

Other: Airflow Modeling and Temperature Distribution Simulation

As you may have seen in my resume, I have multi-discipline background in both computer science and mechanical engineering. I have conducted research in airflow modeling and temperature simulation, collaborating with colleagues in Massachusetts Institute of Technology before I switched to computer science. I proposed a numerical simulation model, POMA, based on thermodynamics and fluid dynamics, to predict indoor airflow pattern and temperature distribution. The outcomes of POMA can be used as guidelines or suggestions to ventilation design. Expensive experimental costs can be saved by simply running the simulation in computers. This area of research is interdisciplinary between mechanical engineering and computer science. In the future, I would like to collaborate with mechanical engineering faculty in this research if possible.

[1] Y. Lin, B. Kemme, M. Patino-Martinez, and R. Jimenez-Peris, "Middleware based Data Replication providing Snapshot Isolation", SIGMOD'05, Baltimore, USA, 2005

[2] Y. Lin, B. Kemme, M. Patino-Martinez, and R. Jimenez-Peris, "Consistent Data Replication: Is it feasible in WANs?" Euro-Par'05, Lisbon, Portugal, 2005

[3] Y. Lin, B. Kemme, M. Patino-Martinez, and R. Jimenez-Peris, "Moving Data to the Edge: A New Approach to Delivering Dynamic Content in Edge Computing", submitted to the World Wide Web Conference 2007

[4] Y. Lin, B. Kemme, M. Patino-Martinez, and R. Jimenez-Peris "Applying Database Replication to Multi-player Online Games", Netgames 2006, Singapore