

Topic 4: Caching and Replication

1 Guidelines

An efficient Peer-to-Peer system can't be completed without caching and replication techniques. By caching data at clients, the system can significantly reduce the cost of bandwidth since queries can be answered locally, and hence avoid overload at nodes, which are destination of popular queries. On the other hand, data replication helps to increase the availability of data, which is really essential in Peer-to-Peer environment, where the system is dynamic and unstable due to the fact that nodes can join and leave the system any time. Additionally, both caching and replication can help to improve the performance of query processing since they can help to shorten the routing path of a query because the result data may be located near the query node. In this topic, you should focus on discussion about caching and replication techniques.

In details, caching is a technique in which data results returned from a query can be cached at the query node for subsequence queries. Since the size of cache is limited and the interest of users may change after time, the challenge in caching is which data should be cached and when the cache is full which data in the cache should be removed. These problems should be discussed. In addition, you should also discuss the way Peer-to-Peer systems can be used to improve caching.

On the other hand, replication is a technique in which data can be stored at more than one nodes in the system. The benefits of replication is that it can improve availability of data in case of failure, avoid hotpot problem due to many queries targeting to the same node, and improve performance of query processing since a replica may be nearer to the query node than its original data. Nevertheless, the benefits do not come for free. Replication requires some costs for creating and maintaining replicas. Since it is costly if we blindly replicate data especially in Peer-to-Peer systems, where the amount of data is huge, a challenge in replication is how to determine which object should be replicated, how it can be replicated, how many replicas is enough, and how to maintain consistency of replicas with their original data. A detail view of all above problems should be given.