

IS250 - Mini Project

Fnu Renaldi (#20682424)

February 15th, 2004

Gnutella

I. Intro to P2P

A peer-to-peer (or P2P) computer network is a network that relies on the computing power of the participants in the network rather than concentrating it in centralized, relatively low number of servers (Wikipedia). In contrast to a client-server system where the roles are clearly defined, every node in P2P system typically can both act as a client and a server. Because of this nature, a P2P system generally exhibits the following properties: uncoordinated, resilient to attack, and large collection of resources (Kangasharju 10).

One popular example that we will be focus on this paper is a protocol for a distributed search on a P2P network called Gnutella. Gnutella is a fully distributed system, which means that everyone participates equally in the network. Each Gnutella node is connected to a small number of other Gnutella nodes. When a peer tries to search for a file in the network, it will send a request to its neighbors. If any of the neighbors has the resource, that particular neighbor will send a query hit message to the sender. Otherwise, the request will be forwarded to all neighbor's connected nodes.

II. Gnutella Network Implementation

This section provides an overview of how a new peer joins the Gnutella network and searches for a file within the network; this is according to "From P2P to Web Services and Grids" by Taylor.

Below are the steps to join the network:

1. Supposedly user A wants to join the network, he must first discover a peer in the network by contacting GnuCache to get other peers' addresses.

2. After having this list of addresses, user A can choose from it and try to contact another peer (user B) by sending a "Gnutella Connect".
3. When user B accepts and returns a "Gnutella OK" to user A, user A will be connected to the Gnutella network.
4. To get more connections, user A will announce its presence by sending a Ping message to its neighbor (user B). User B will forward this to user X, who will forward it to user C.
5. If user C returns a Pong message which also contains its address, user A can now connect to user C by using that address.
6. User A will repeat step 4 until he gets to the maximum connection count, which typically set to 4.

Below are the steps in searching within Gnutella network:

1. Supposedly user A wants to search for a file, he will send the request to all its direct neighbors.
2. When A's neighbors receive the *query request*, they will search their local files. If they do not find the requested file, they will forward the *query response* to all its direct neighbors. Otherwise, they return a *query response* containing all the matches it finds.
3. *Query responses* follow the reverse path of *query request* to reach the initiating peer.

Other important implementations:

To prevent *query requests* and *query responses* flooding and congesting the network, each query contains a TTL (Time To Live) field. Every time a request is forwarded, its TTL is decremented. If a peer receives a request with a zero value of TTL, then the request will not be forwarded anymore.

III. Strengths and Weaknesses:

Knowing how the Gnutella works, we can now analyze the advantages and disadvantages of it. Gnutella is more robust than a centralized model because it eliminates the single node failures threat. Simultaneously, it will also be harder to bring Gnutella down because of the very large number of peers in the system. As long as a node is not fully isolated, that node still works. From the developers' perspective, it is considerably easy to write for clients from all platforms because Gnutella is an open protocol.

However, there are also some problems and disadvantages associated with Gnutella. While searching for a file, the request is flooded to the network. Thus, it wastes a lot of network and peer resource. Since your machine is part of this network, you give up some amount of your bandwidth to handle requests from other users. Because search only reaches a subset of peers, there is also no guarantee that the file you search is on any of those machines you can reach (Kangasharju 24).

IV. Conclusion

Although it is very tempting to use a P2P system like Gnutella, at the end, there are some considerations that you still need to take. P2P is good when budget is limited, resources have wide spread interest and relevance, trusts among participants are high, and criticality is low. Some of the industries that use P2P network are file sharing industries (LimeWire, etc), online game industries and iTunes.

References

Kangasharju, Jussi. Class Lecture. *Peer-to-Peer Network*. University of Helsinki, Finland. 14 Jan 2008

"Peer-to-peer" Wikipedia, *The Free Encyclopedia*. Wikimedia Foundation, Inc. 15 April 2010. Web. 12.

Feb. 2004. <<http://en.wikipedia.org/wiki/Peer-to-peer>>

Taylor, Ian J. *From P2P to Web Services and Grids*. London: Springer, 2005