Simulations of Chord and Freenet Peer-to-Peer Networking Protocols
Mid-Term Report

Computer Communications and Networking (SC 546)
Professor D. Starobinksi

Brian Mitchell
U09-62-9095
James Nunan
U38-03-0277

Simulations of Chord and Freenet Peer-to-Peer Networking Protocols
Mid-Term Report

In recent years, peer-to-peer (P2P) technology in networking has gained much attention due to file-sharing applications such as Napster and Gnutella. Although P2P has no strict definition, it is generally described as having a structure that is not the traditional server-client model. Each node in the network acts as both client and server, requesting data from neighboring nodes as well as routing and serving data for others. The nature of P2P technology makes it well suited for storing multiple copies of document between several nodes, in turn offering reliable access to data and distributing the load of requests. Additionally, the multiple links between nodes make the system more stable as nodes are added and dropped. All of the features inherent in P2P technology promise a network that is dynamic, scalable and reliable.

The recent attention brought to distributed computing and peer-to-peer technology as well as the potential benefits that P2P networks possess will undoubtedly lead to the development of more applications in the field. Since the technology is based on the interaction between nodes, it is difficult to predict the exact nature of a whole network as the number of nodes grows extremely large. The study of the efficiency of queries (requests of data) within the distributed system as well as the effects of adding and dropping nodes on data retrieval will eventually lead to the design of faster and more robust P2P applications in the future.

This project will compare two different implementations of peer-to-peer technology, Chord and Freenet. Chord was designed to create a network that is reliable, scalable, and decentralized. Chord uses consistent-hashing, a method that evenly distributes hash keys to nodes. Each node contains a finger table of its neighbors and their possible assignments of keys. Nodes are organized in a ring topology, maintaining keys of values that are less than or equal to the assigned node value and greater than the value of the node's predecessor. Unsuccessful queries are forwarded around the ring to successive nodes, which allows each node only track on the order of log N neighbors, where N is the total number of nodes in the network. When a new node enters the network, it randomly assigns itself a node identifier and moves appropriate keys from its

successor. To increase reliability of the system, the values (documents) associated with keys are replicated a configurable number (r) times on the r nodes following the node possessing the associated key. Preceding nodes pointing to the key are also aware that there are r duplicates of the value beyond the key-holding node. If a node associated with a key is dropped from the network, the request is simply forwarded to the next node on the ring. In order for a document to be lost, all r nodes must be lost.

The design goals of Freenet were to create a completely decentralized, scalable peer-to-peer application allowing anonymous input, retrieval and storage of files. Unlike Chord, Freenet assigns hash keys to specific items (files). Like Chord, Freenet nodes contain a table of information about their neighbors. However, Freenet stores information about which keys it suspects neighbors possess. In the case that the node does not have a queried key, and is not certain a neighbor has the key, it will forward the request to the neighboring node that possesses the closest key. Additionally, Freenet makes use of a very aggressive caching scheme. When a document is found, it can be cached at multiple points along the path between the request and the point where the document is located.

For this project, we will use simulation to compare Freenet and Chord. Aspects which will be evaluated are the efficiency of queries in similar networks (same number of nodes and items) and the ability of the networks to adapt to these queries over time. There will also be a comparison on how queries are affected by nodes being added and dropped. The chosen Chord simulator is the one provided with the Chord distribution. There are a variety of potential Freenet simulators, including Serapis, Aurora, and a third included with a simulator for the Neurogrid P2P system. The Freenet simulator(s) chosen to perform the comparison will be dependent on the experiment parameters and the capabilities of the each simulator. The code will be modified to allow output data which can be easily loaded into MATLAB for plotting and comparison.

The timeline for completion of this project is as follows:

- 11/01 – Experiment Parameters and Freenet Simulator Chosen
- 11/08 – Simulation of Chord and Freenet Complete
- 11/15 – Comparison / Graphing Complete
- 11/22 – Presentation Complete
- 12/04 – Website Complete

Because of the differences between the way Freenet and Chord distributes keys, stores items and handles queries, hypothetically the initial number of hops required by a query to locate a document will be significantly higher in Freenet than Chord. However, because of the adaptive nature of nodes and the caching used by the Freenet technology, over time the number of hops between queries will be significantly reduced, possibly becoming even more efficient than Chord. Since both Chord and Freenet store duplicate copies of data it is difficult to predict the effect that dropping nodes will have on the efficiency of queries. In Chord, successive nodes store duplicate data so it will be easy to locate in the case of node failure and a lost key. However, Freenet's aggressive caching scheme may also buffer the system against excessive query hops in the case when a node or section of the network is lost.

References

1. Chord Project website. http://www.pdos.lcs.mit.edu/chord.
2. Freenet Project website. http://freenetproject.org.
3. Neurogrid Project website. http://www.neurogrid.net.
4. Clarke, et al., "Freenet: A Distributed Anonymous Information Storage and Retrieval System," in *Designing Privacy Enhancing Technologies*, Springer-Verlag, Belin, 2001, pp. 46-66.
5. Clarke, et al., "Protecting Free Expression Online with Freenet," in *IEEE Internet Computing,* January-February 2002, pp. 40-49.
6. Ion Stoica, et al., "Chord: A Scalable Peer-to-peer Lookup Protocol for Internet Applications," to appear in *IEEE/ACM Transactions on Networking*.
7. Ion Stoica, et. al., "*Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications*," ACM SIGCOMM 2001, San Deigo, CA, August 2001, pp. 149-160.