

Mapping the Gnutella Network: Macroscopic Properties of Large-Scale Peer-to-Peer Systems

Matei Ripeanu, Ian Foster
{matei, foster}@cs.uchicago.edu

Abstract

Despite recent excitement generated by the peer-to-peer (P2P) paradigm and the surprisingly rapid deployment of some P2P applications, there are few quantitative evaluations of P2P systems behavior. The open architecture, achieved scale, and self-organizing structure of the Gnutella network make it an interesting P2P architecture to study. Like most other P2P applications, Gnutella builds, at the application level, a virtual network with its own routing mechanisms. The topology of this overlay network and the routing mechanisms used have a significant influence on application properties such as performance, reliability, and scalability. We describe techniques to discover and analyze the Gnutella's overlay network topology and evaluate generated network traffic. Our major findings are: (1) although Gnutella is not a pure power-law network, its current configuration has the benefits and drawbacks of a power-law structure, (2) we estimate the aggregated volume of generated traffic, and (3) the Gnutella virtual network topology does not match well the underlying Internet topology, hence leading to ineffective use of the physical networking infrastructure. We believe that our findings as well as our measurement and analysis techniques have broad applicability to P2P systems and provide useful insights into P2P system design tradeoffs.

1. Introduction

Unlike traditional distributed systems, P2P networks aim to aggregate large numbers of computers that join and leave the network frequently. In pure P2P systems, individual computers communicate directly with each other and share information and resources without using dedicated servers. A common characteristic of this new breed of systems is that they build, at the application level, a virtual network with its own routing mechanisms. The topology of this overlay network and the routing mechanisms used have a significant impact on application properties such as performance, reliability, scalability, and, in some cases, anonymity. The topology also determines the communication costs associated with running the P2P application, both at individual hosts and in the aggregate. Note that the decentralized nature of pure P2P systems means

that these properties are emergent properties, determined by entirely local decisions made by individual resources, based only on local information: we are dealing with a self-organized network of independent entities.

These considerations motivate us to conduct a macroscopic study of a popular P2P system: Gnutella (described succinctly in Section 2). In this study, we benefit from Gnutella's large existing user base and open architecture, and, in effect, use the public Gnutella network as a large-scale, if uncontrolled, testbed.

Our measurements and analysis of the Gnutella network are driven by two primary questions (Section 4). The first concerns its connectivity structure. Recent research [1] shows that networks as diverse as natural networks formed by molecules in a cell, networks of people in a social group, or the Internet, organize themselves so that most nodes have few links while a tiny number of nodes, called hubs, have a large number of links. [2] finds that networks following this organizational pattern (*power-law networks*) display an unexpected degree of robustness: the ability of their nodes to communicate is unaffected even by extremely high failure rates. However, random error tolerance comes at a high price: these networks are vulnerable to attacks, i.e., to the selection and removal of a few nodes that provide most of the network's connectivity. We show that, although Gnutella is not a pure power-law network, it preserves good fault tolerance characteristics while being less dependent than a pure power-law network on highly connected nodes that are easy to single out (and attack).

The second question concerns how well (if at all) the Gnutella virtual network topology maps to the physical Internet infrastructure. There are two reasons for analyzing this issue. First, it is a question of crucial importance for Internet Service Providers (ISP): if the virtual topology does not follow the physical infrastructure, then the additional stress on the infrastructure and, consequently, the costs for ISPs, are immense.

This point has been raised on various occasions but, as far as we know, we are the first to provide a quantitative evaluation on P2P application and Internet topology (mis)match. Second, the scalability of any P2P application is ultimately determined by its efficient use of underlying resources.

An orthogonal but important issue concerns the data gathering techniques (Section 3). For the analysis we present here we developed a “crawler” to gather complete topology information on the network. This technique however, is invasive and has limited scalability. (Moreover, recent minor changes in the protocol baffle our ‘crawler’.) We are currently looking at ways to explore and characterize the network by adding limited number of cooperating ‘probes’ (modified nodes) that will monitor the traffic and may insert a small number of control messages.

2. Gnutella Protocol Description

The Gnutella protocol [3] is an open, decentralized group membership and search protocol, mainly used for file searching and sharing. The term Gnutella also designates the virtual network of Internet accessible hosts running Gnutella-speaking applications. Gnutella nodes, called *servents* by developers, perform tasks normally associated with both *SERVERS* and *CLIENTS*. They provide client-side interfaces through which users can issue queries and view search results, accept queries from other servents, check for matches against their local data set, and respond with corresponding results. These nodes are also responsible for managing the background traffic that spreads the information used to maintain network integrity.

In order to join the system a new node/servent initially connects to one of several known hosts that are almost always available (e.g., gnutellahosts.com). Once attached to the network (e.g., having one or more open connections with nodes already in the network), nodes send messages to interact with each other. Messages can be *broadcasted* (i.e., sent to all nodes with which the sender has open TCP connections) or simply *back-propagated* (i.e., sent on a specific connection on the reverse of the path taken by an initial, broadcasted, message). Several features of the protocol facilitate this broadcast/back-propagation mechanism. First,

each message has a randomly generated identifier. Second, each node keeps a short memory of the recently routed messages, used to prevent re-broadcasting and to implement back-propagation. Third, messages are flagged with time-to-live (TTL) and “hops passed” fields.

The messages allowed in the network are:

- *Group Membership (PING and PONG) Messages*. A node joining the network initiates a broadcasted PING message to announce its presence. When a node receives a PING message it forwards it to its neighbors and initiates a back-propagated PONG message. The PONG message also contains information about the node such as its IP address and the number and size of shared files.
- *Search (QUERY and QUERY RESPONSE) Messages*. QUERY messages contain a user specified search string that each receiving node matches against locally stored file names. QUERY messages are broadcasted. QUERY RESPONSES are back-propagated replies to QUERY messages and include information necessary to download a file.
- *File Transfer (GET and PUSH) Messages*. File downloads are done directly between two peers using GET/PUSH messages.

To summarize: to become a member of the network, a *servent* (node) has to open one or many connections with nodes that are already in the network. In the dynamic environment where Gnutella operates, nodes often join and leave and network connections are unreliable. To cope with this environment, after joining the network, a node periodically PINGS its neighbors to discover other participating nodes. Using this information, a disconnected node can always reconnect to the network. Nodes decide where to connect in the network based only on local information, and thus forming a dynamic, self-organizing network of independent entities. This virtual, application level network has Gnutella servents at its nodes and open TCP connections as its links.

In this section we described the original Gnutella protocol (v0.4) as, at the time of our experiment, most nodes complied with this protocol version. We should mention however that a number of protocol changes have been adopted. The most significant result is a switch from the initial, flat, unstructured, peer network toward a two-level network organization: ordinary

nodes link to *SuperPeers* that shield them from some of the traffic. SuperPeers in turn organize themselves into a flat, unstructured, network similar to the original one.

3. Data Collection

We have developed a *crawler* that joins the network as a servent and uses the membership protocol (the PING-PONG mechanism) to collect topology information. The crawler starts with a list of nodes, initiates a TCP connection to each node in the list, sends a generic join-in message (PING), and discovers the neighbors of the contacted node based on the PONG messages that it receives in reply. Newly discovered neighbors are added to the list. We started with a short, publicly available list of initial nodes, but over time we have incrementally built our own list with more than 400,000 nodes that have been active at one time or another.

In order to reduce the crawling time, we developed a client/server crawling strategy. The ‘server’ is responsible with managing the list of nodes to be contacted, assembling the final graph, and assigning work to clients. Given this dynamic behavior of the nodes, it is important to find the appropriate tradeoff between discovery time and invasiveness of our crawler. Increasing the number of parallel crawling tasks reduces discovery time but increases the burden on the application. Obviously, the Gnutella graph our crawler produces is not an exact ‘snapshot’ of the network. However, we argue that the result we obtain is close to a snapshot in a statistical sense: all properties of the network: size, diameter, average connectivity, and connectivity distribution are preserved.

Still, our crawling technique is invasive and has limited scalability. Moreover, recent minor modifications to the protocol changed the PING-PONG mechanism the crawler is based on. These modifications, aimed at reducing the number of messages broadcasted in the network, lead to the widespread deployment of ‘PONG caches’. We are currently looking at separate ways to explore and characterize the Gnutella network by adding a limited number of cooperating ‘*probes*’ (modified nodes) that monitor the traffic and may insert a small number of control messages. While some network properties can only be analyzed using complete graph information, the data gathered by

the ‘*probes*’ is sufficient to estimate a some interesting characteristics: message drop rates, network diameter or its tolerance to attacks (the number of redundant network paths).

While during late 2000 the largest connected network component we found had 2,063 hosts, this grew to 14,949 hosts in March 2001 and 48,195 hosts in May 2001. Recent measurements (www.limewire.com) show that the network in the range of 80-100,000 nodes.

4. Gnutella Network Analysis

We start with a macroscopic analysis of the network and study its connectivity patterns (Section 4.1). We then estimate Gnutella generated traffic volume (Section 4.2), and evaluate the mapping of Gnutella overlay network to the underlying networking infrastructure (Section 4.3).

4.1. Connectivity and Reliability in Gnutella Network. Power-law Distributions.

Recent research [1] shows that many natural networks such as molecules in a cell, species in an ecosystem, and people in a social group organize themselves as so called *power-law networks* (more specifically, in a power-law network the number of nodes with L links is proportional to, L^{-k} where k is a network dependent constant). This structure helps explain why these are generally stable and resilient structures, yet occasional catastrophic collapse does occur [2]. In a power-law network most nodes (molecules, Internet routers, Gnutella servents) have few links, thus a large fraction can be taken away and the network stays connected. But, if just a few highly connected nodes are eliminated, the whole network is broken into pieces. One implication is that these networks are extremely robust when facing random node failures, but vulnerable to well-planned attacks.

Given the diversity of networks that exhibit power-law structure and their properties we were interested to determine whether Gnutella falls into the same category. Figure 1 presents the connectivity distribution in Nov. 2000. Although data are noisy (due to the small size of the networks), we can easily recognize the signature of a power-law distribution: the connectivity distribution appears as a line on a log-log plot. Later measurements (Figure 2) however, show that more recent networks tend to move away from this organization: there are too few nodes

with low connectivity to form a pure power-law network. In these networks the power-law distribution is preserved for nodes with more than 10 links while nodes with fewer link follow a quasi-constant distribution.

An interesting issue is the impact of this new, multi-modal distribution on network reliability. We believe that the more uniform connectivity distribution preserves the network's ability to deal with random node failures while reducing the network dependence on highly connected, easy to single out (and attack) nodes.




Figure 1: Connectivity distribution during November 2000. Each series of points represents one Gnutella network topology we discovered at different times during that month. Note the log scale on both axes. Gnutella nodes organized themselves into a power-law network.




Figure 2: Connectivity distributions during March 2001. Each series of points represents one Gnutella network topology discovered during March 2001. Note the log scale on both axes. Networks crawled during May/June 2001 show a similar pattern.

We speculate that a group of devoted users maintain the small number of Gnutella nodes with the server-like characteristics visible in these power-law distributions. These nodes have a large number of open connections and/or provide much of the content available in the network. Moreover, these server-like nodes have a higher availability: they are about 50% more likely than the average node to be found alive during two successive crawls.

4.2 Estimates of Generated Traffic

We used a modified version of the crawler to eavesdrop the traffic generated by the network. We classified the traffic that goes on randomly chosen links according to message type: 92% QUERY messages, 8% PING messages and insignificant levels of other message types (June 2001). This represents a significant improvement compared to late 2000 when the traffic contained more than 50% overhead messages (PINGs and PONGs).

The topology information collected allows us to analyze the distribution of node-to-node shortest path lengths. Given that 95% of any two nodes are less than 7 hops away (Figure 3), the message time-to-live (TTL=7) preponderantly used, and the flooding-based routing algorithm employed, most links support similar traffic. We verified this theoretical conclusion by measuring the traffic at multiple, randomly chosen, nodes and found 6 Kbps per connection on average. As a result, the total Gnutella generated traffic is proportional to the number of connections in the network. Based on our measurements we estimate the total traffic (excluding file transfers) for a large Gnutella network as 1 Gbps: 170,000 connections (for a 50,000-nodes Gnutella network) times 6 Kbps per connection, or about 330 TB/month. To put this traffic volume into perspective we note that it amounts to about 1.7% of total traffic in US Internet backbone in December 2000 (as reported in [4]). We infer that the volume of generated traffic is an important obstacle for further growth and that efficient use of underlying network infrastructure is crucial for better scaling and wider deployment.

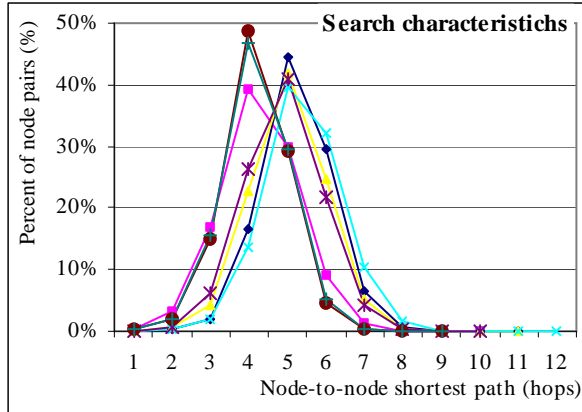


Figure 3: Distribution of node-to-node shortest paths. Each line represents one network measurement. Note that, although the largest network diameter (the longest node-to-node path) is 12, more than 95% of node pairs are at most 7 hops away

4.3. Internet Infrastructure and Gnutella Network

P2P computing brings an important change in the way we use the Internet: it enables computers sitting at the edges of the network to act as both clients and servers. As a result, P2P applications change radically the amount of bandwidth consumed by the average Internet user. Most Internet Service Providers (ISPs) use flat rates to bill their clients. If P2P applications become ubiquitous, they could break the existing business models of many ISPs and force them to change their pricing scheme.

Given the considerable traffic volume generated by P2P applications, it is crucial from the perspective of both their scalability and their impact on the network that they employ available networking resources efficiently. Gnutella's store-and-forward architecture makes the overlay network topology extremely important: the larger the mismatch between the physical network infrastructure and the overlay's topology, the bigger the "stress" on the infrastructure.

Unfortunately, it is prohibitively expensive to compute exactly the mapping of the Gnutella onto the Internet topology, due both to the inherent difficulty of extracting Internet topology and to the computational scale of the problem. Instead, we proceed with a high-level experiment that highlights the topology mismatch: The Internet is a collection of Autonomous Systems (AS) connected by routers. ASs, in turn, are collections

of local area networks under a single technical administration. From an ISP point of view traffic crossing AS borders is more expensive than local traffic. We found that only 2-5% of Gnutella connections link nodes located within the same AS, although more than 40% of these nodes are located within the top ten ASs (10% in the largest). This result indicates that most Gnutella-generated traffic crosses AS borders, thus increasing costs, unnecessarily. A second, similar experiment showed that the node organization does not follow the DNS domain name hierarchical organization either.

5. Summary and Future Work

Despite recent excitement generated by this paradigm and the surprisingly rapid deployment of some P2P applications, there are few quantitative evaluations of P2P systems behavior. The open architecture, achieved scale, and self-organizing structure of the Gnutella network make it an interesting P2P architecture to study. The social circumstances that have fostered the success of the Gnutella network might change and the network might vanish. However, our measurement and analysis techniques can be used for other P2P systems to enhance general understanding of design tradeoffs.

Our analysis shows that, although Gnutella is not a pure power-law network, it preserves good fault tolerance characteristics while being less dependent than a pure power-law network on highly connected nodes that are easy to single out (and attack).

We have estimated that, as of June 2001, the network generates about 330 TB/month simply to remain connected and to broadcast user queries. This traffic volume represents a significant fraction of the total Internet traffic and makes the future growth of Gnutella network particularly dependent on efficient network usage. We have also documented the topology mismatch between the self-organized, Gnutella network and the underlying physical networking infrastructure. We believe this mismatch has major implications for the scalability of this P2P network or for ISP business models. This problem must be solved if Gnutella or similarly built systems are to reach larger deployment.

We see two other directions for improvement. First, as argued in [6], efficient P2P designs

should exploit particular distributions of query values and locality in user interests. Various Gnutella studies show that the distribution of Gnutella queries is similar to the distribution of HTTP requests in the Internet: they both follow Zipf's law (note that, although the Zipf's formulation is widely used, these distributions can also be expressed as power-law distributions).

Other projects [7] try to discover and exploit data sharing patterns emerging at user level for topology optimization and collaborative message filtering.

A second direction of improvement is the replacement of query flooding mechanism with smarter (less expensive in terms of communication costs) routing and/or group communication mechanisms. Several P2P schemes proposed recently fall into the former category: systems like CAN or Tapestry propose a structured application-level topology that allows semantic query routing. We believe, however, that a promising approach is to preserve and benefit from the power-law characteristics that, as shown in this paper, emerge in Gnutella's ad-hoc network topology (and match the underlying, non-uniform resource distribution [5]). A way to preserve the dynamic, adaptive character of the Gnutella network and still decrease resource (network bandwidth) consumption is to use dissemination schemes (e.g., based on epidemic protocols) mixed with random query forwarding.

We have collected a large amount of data on the environment in which Gnutella operates, and plan to use this data in simulation studies of protocol alternatives.

6. Acknowledgements

This work was supported in part by the U.S. National Science Foundation under contract ITR-0086044.

7. References

- [1] R. Albert, A. L. Barabasi, Statistical mechanics of complex networks, *Reviews of Modern Physics*, 74 (47) 2002.
- [2] R. Albert, H. Jeong, A. Barabási, *Attak and tolerance in complex networks*, *Nature* 406, 378, 2000.
- [3] The Gnutella protocol specification v4.0. <http://dss.clip2.com/GnutellaProtocol04.pdf>.
- [4] K. Coffman, A. Odlyzko, *Internet growth: Is there a "Moore's Law" for data traffic?*, *Handbook of*

Massive Data Sets, J. Abello & all editors., Kluwer, 2001.

- [5] S. Saroiu, P. Gummadi, S. Gribble, [A Measurement Study of P2P File Sharing Systems](#), University of Washington Technical Report UW-CSE-01-06-02, July 2001.
- [6] K. Sripanidkulchai, [The popularity of Gnutella queries and its implications on scalability](#), February 2001.
- [7] Locating Data in (Small-World?) P2P Scientific Collaborations, Adriana Iamnitchi, Matei Ripeanu, Ian Foster, in *Proc. of the 1st International Workshop on Peer-to-Peer Systems* Cambridge, MA, March 2002