

Peer Pressure: Distributed Recovery from Attacks in Peer-to-Peer Systems

Pedram Keyani

Brian Larson

Muthukumar Senthil

Computer Science Department
Stanford University
Stanford, CA 94305

{pkeyani, balarson, msenthil}@stanford.edu

Keywords. Peer-to-peer systems, overlay networks, scale-free networks, fault recovery, malicious attack

Abstract. Peer-to-peer systems such as Gnutella are resilient to failures at a single point in the network because of their decentralized nature. However an attack resulting in the removal of a small percentage of highly connected nodes could cripple such systems. We believe that distributed attack recovery is not simply a reactive process but requires proactive measures by the nodes in the system. We propose a distributed recovery method, where clients proactively detect attacks by monitoring the rate at which their first and second-degree neighbors leave the network and reconfigure themselves to form a topology that is more resilient to attacks when one has been detected. This topology is created and maintained through a new type of node discovery mechanism that is used during normal network operations. The recovery method is able to reconnect the network and deal with any ongoing attacks once one has started.

1. Introduction

P2P systems are becoming more prevalent in our lives as the Internet has gained widespread acceptance. Clearly, before we become more reliant on such systems, they must be shown to be secure, and the next generation of P2P applications must be able to survive malicious attacks to the network. Currently, P2P systems are not the subject of widespread attacks, but it is not difficult to envision scenarios where attacks could occur. As an example, it would not be out of line for a major record label to attack Gnutella because one of their songs becomes freely available.

Gnutella is one of the most widely used peer-to-peer (P2P) protocols for file sharing on the Internet. It is a simple protocol for communication between peers (*servants* in Gnutella) to form an overlay network on top of the Internet topology. This protocol does not specify a recovery mechanism, only a format for clients to communicate [7, 9].

The current topology of Gnutella adheres to a power law distribution in which most nodes have few connections, while a small fraction of the nodes have many connections and hold the entire network together [12]. It is this property that makes the Gnutella topology particularly susceptible to malicious attacks on highly connected nodes. By merely removing a small portion of these highly connected nodes, it is possible to fragment the entire network into many isolated pieces [12].

To make matters worse, current clients do not even attempt to detect attacks on the network. Most clients can and do detect failures of neighboring nodes, but surviving an isolated failure is different from surviving an attack. Furthermore, clients do not have a backup plan that is resistant to attack. This leads to a system extremely vulnerable to malicious attacks and unsuitable for critical applications.

The first step to surviving an attack is detection. It is our position that nodes must be able to detect an attack quickly and without knowledge of the global network topology. Our method for detecting attacks requires nodes to keep track of the rate at which their first-degree-neighbors (direct neighbors) and second-degree-neighbors (neighbors of first-degree-neighbors) leave the network.

It is also our position that nodes must plan for an attack before it happens. Each node needs to discover backup connections during normal network operations and maintain a list of them in case of an attack. Our proposed technique for this requires the addition of a *random discovery ping (RDP)* to the Gnutella protocol. In order to form a more resilient network during an attack, failed neighbors are replaced with the backups discovered using the random node discovery ping during normal operations.

Our results show that this technique reduces fragmentation by more than a factor of 25 times from the standard approach. In addition, our results also show an improvement in querying effectiveness both during and after the attack. Finally, we show that only a small increase in traffic on the network is required to implement our technique.

In section 2 we describe the background to the problem more fully. We describe our proposed solution in section 3. In section 4 we describe the experimental model we used to validate our solution. In section 5 we

describe our simulator and show the results of the experiments we ran. Finally, in section 6 we conclude and talk about future work.

2. Problem Background

There are currently a wide variety of decentralized P2P schemes that are highly resilient to failures of random components [15, 17, 18]. Because these systems expand with no centralized planning, they often result in a type of topology that can be severely fragmented by attacking a few nodes that hold the entire network together [2]. Several programs exist that “crawl” the Gnutella network and report its topology, allowing a person to isolate and attack the nodes that keep the network together [16]. Mounting a distributed denial of service attack against these nodes could effectively kill them in a time-span of a few minutes [11, 13].

It has been shown that if 4% of the most highly connected nodes are removed from Gnutella, the network will severely fragment, rendering it useless [12]. Gnutella’s robustness to random failure and vulnerability to malicious attack is not unique. Indeed, the Internet has similar characteristics; an attack on 5% of nodes would result in the total collapse of the Internet [14].

2.1 Gnutella Network Topology

Many real-world networks fall into a class of inhomogeneous networks called *scale-free networks*, where a few nodes have many connections, but most nodes have only a few connections [3]. Scale-free networks abide by the power-law relationship, $P(k) \sim k^{-a}$, where $P(k)$ gives the probability that a node is connected to k other nodes [2]. It has been shown that networks such as the Internet, the WWW, and Gnutella tend to self-organize into scale-free topologies that abide by the power law [1, 3, 8, 10, 12].

There are two mechanisms that cause the formation of scale-free topologies. First, networks expand continuously by the addition of new vertices, and second, new vertices attach preferentially to vertices that are already well connected [3]. In Gnutella, the first mechanism can be seen by the fact that new nodes are continuously entering and leaving the system, meaning the topology is undergoing constant change and growth. The second mechanism can be seen by the fact that there are only a few hosts that clients initially connect to, due to the way bootstrapping is handled (further discussed in section 4.1). Furthermore, clients preferentially attach to stable, high bandwidth nodes. Hence, the topology of the Gnutella network is scale-free because of its adherence to these two mechanisms.

Failures in P2P systems are viewed as nodes suddenly and unexpectedly dropping out of the network. Scale-free networks such as Gnutella are very robust to failures of random nodes. This robustness is rooted in the inhomogeneous nature of the network. Nodes with few connections will fail with much greater probability than nodes with many connections [2]. Unfortunately, the same inhomogeneous nature of scale-free networks that makes them robust to random failures makes them vulnerable to an attack resulting in the removal of the most highly connected nodes.

2.2 Failure and Attack Detection

Failure detection is the process of detecting isolated random failures of peers on the network. Current Gnutella clients detect a failure when one of their neighbors stops responding to them unexpectedly. The only information that can be derived from this process of detection is that a neighbor has died. There is no indication of how significant the neighbor was in terms of connections to other nodes and how important it was in holding the graph together.

Attack detection, on the other hand, is the process of detecting an attack on the network as a whole. Methods for attack detection could take into account the importance of nodes being lost in the system or the frequency of node loss. Attack detection is necessary for any system where mounting an attack is a relatively easy and inexpensive task.

Responding to failures is very different from responding to attacks because of the nature of the nodes that are lost. If a small percentage of nodes at the core of a network were attacked, clients responding only to the failure of these nodes would not address the severity of the problem. A system that cannot detect an attack has little chance of surviving it.

2.3 Attack Backup Plans

Most Gnutella clients maintain a “cache” of host names, which they have communicated with within a certain time period. These cached hosts are often used as replacements for failures, but they make poor replacements for attacked nodes. If a client does detect an attack in progress, and connects to hosts in the cache as a response, hosts that most recently sent queries through the network would appear in thousands of caches. This means they could be overwhelmed by connection requests, or even become targets of the attack themselves. It becomes very clear that a backup plan must be picked that is resistant to further attacks.

One network topology that is more resilient to attacks is an exponential network [2]. An exponential network is a homogenous network where each node has roughly the same number of links k . Instead of following the power law, $P(k)$ peaks at k and decays exponentially, following the Poisson distribution [2]. Exponential networks are built by connecting random neighbors together with no preference of one node over another. In this configuration, it is improbable that any one node is holding the network together. The homogeneous nature of exponential networks makes them much less vulnerable to attacks on a small number of nodes than scale-free networks.

3. Recovery Method

We would like to have a recovery method that prevents the network from becoming fragmented by an attack and prevents the decrease in querying effectiveness during the attack. By querying effectiveness, we mean the number of neighbors that can be reached with a query.

We would like to be able to reconfigure the topology of the network from scale-free to exponential in order to survive an attack, since exponential networks fare much better than scale free networks under attack on a small percentage of nodes [2]. Our solution is to maintain a virtual exponential overlay network in addition to the active scale free overlay network used by the system. We refer to the network as virtual since connections between nodes on the network will not be made and no query traffic will be sent over the network.

Reconfiguration of the topology to exponential requires nodes to discover other nodes randomly, but this reconfiguration must be done before the network becomes fragmented from the attack. Each node will maintain a list of virtual neighbors in addition to the node’s active neighbors, and they will be selected using a process for random node discovery that we have developed. Because there are no preferential connections, these random connections will constitute an exponential network.

The virtual exponential network will be used in place of the active scale-free network during an attack, which each node in the network is responsible for detecting. When a node detects an attack, the node will begin replacing its dead neighbors with nodes from its virtual neighbor list, thus making the exponential network active as the scale free network disintegrates as a result of the attack. By combining these two types of network topologies, we will have the robustness of scale-free networks to failure coupled with the robustness of exponential networks to attack.

3.1 Random Node Discovery

Our recovery mechanism requires that nodes be able to discover random nodes in the system with little or no preference in order to construct an exponential network. An optimal solution would allow for any node in the network to be chosen with equal probability, a task that could be accomplished by a centralized name authority. In a decentralized system such as Gnutella, maintaining a database of all active nodes in the system would be difficult, if not impossible.

Our solution is for nodes to forward a message randomly through the network for a certain number of hops. This message, which we will call a random discovery ping, will be similar to a Gnutella ping. An RDP contains an originating node, a hop count, and is forwarded from node to node through the network. The final node that receives the RDP, determined when the hop count is decremented to one, will respond with a pong and is thus discovered. This node will then be added to the virtual neighbor list maintained by the originating node.

While standard Gnutella pings are forwarded by a node to all of its neighbors, an RDP is only forwarded to one of the node’s neighbors. This will allow us to use a much larger time to live (TTL), thus getting deeper into the network without overwhelming the network with ping traffic. In our implementation, random discovery pings are created with a TTL of 20. This number is a good approximation of the diameter of the Gnutella network [6], so this

will enable a ping to reach any node in the network. Acknowledgment messages are used to ensure that an RDP is not lost.

Initially a node forwards an RDP to a neighbor N selected randomly with probability scaling linearly with the number of neighbors that N has. This means that a node with 6 neighbors is 3 times as likely to have the ping forwarded to it than a node with 2 neighbors. This strategy is used for the first 10 hops in order to get the ping as far away from the originating node as possible (we want to have the node connect outside its immediate vicinity). If this strategy is not used, in most cases the ping is only forwarded within the immediate vicinity of the originating node, cycling around the originating node's most connected neighbor many times.

For the next 10 hops, the opposite strategy is used. The ping is forwarded to nodes with lower number of neighbors, meaning that a node with 6 neighbors is 1/3 times as likely to have the ping forwarded to it than a node with 2 neighbors. If this strategy is not used, then the nodes with the most neighbors will be selected with a frequency scaling with their connectivity. This is not what we desire because it replicates the preferential attachment property of the scale-free active network.

3.2 Maintenance of Virtual Exponential Network

The process by which each node discovers and maintains a certain number of virtual neighbors globally maintains an exponential network. Nodes are discovered by creating and forwarding an RDP, as outlined in section 3.1. Although this is not specified in the Gnutella protocol, all clients have a range of neighbors $\langle X, Y \rangle$ and must maintain at least X neighbors. Similar to active neighbors, virtual neighbors are maintained through periodic pings, or heartbeats, which must be responded with a pong by the neighbor if it is alive [7, 9]. This allows nodes to detect the liveness of their virtual neighbors. If a node's virtual neighbor is found to be dead or a node has less than the minimum number of virtual neighbors, then the node discovers a new random node.

3.3 Attack Detection.

One of our goals in creating an attack detection mechanism was to minimize the amount of coordination between nodes. Our attack detection scheme leverages the information that a node has of its immediate neighbors (1st degree) and their neighbors (2nd degree), and is based on measuring changes in this local topology.

From the viewpoint of a node in the network, an attack will most likely result in the removal of the most highly connected of its 1st degree neighbors. The removal of highly connected neighbors will be the most detrimental to the node because they contain a disproportionate amount of its 2nd degree neighbors.

For example, consider the node shown in black in figure 1A. This node has three 1st degree neighbors, which are shown in gray. One of the neighbors has 3 connections (not counting the node itself), one has 5, and one has 12. This means the node has a total of 20 2nd degree neighbors, which are shown in white. Figure 1B shows what would happen if the neighbor with 12 connections is removed, such as might happen during an attack on the network. The node will lose 12 of its 20 2nd degree neighbors, or 60%. However, the node only lost 33% of its 1st degree neighbors.

Now consider what happens if one of the neighbors is removed at random, such as in the case of a failure. The average number of 2nd degree neighbors the node will lose is 6.67, or 33%. This is the same as the percentage of 1st degree neighbors lost. Clearly, in an attack a node will lose a greater percentage of its 2nd degree neighbors than 1st degree neighbors. This is different from random failures, where the percentage of 1st and 2nd degree neighbors lost will be roughly equal.




Figure 1. Attack detection example

Our attack detection mechanism requires nodes to maintain the number of 1st and 2nd degree neighbors lost during a specified time period T . These losses will be used to calculate the percentage of 1st and 2nd degree neighbors lost during this period. If the percentage of 2nd degree neighbors lost is greater than the percentage of 1st degree neighbors lost, and greater than a threshold P , an attack is detected. The threshold P is used to filter out false positives where some of a node's neighbors fail at random, but the total percentage of 2nd degree neighbors lost is very small. In our implementation, we used a value of 30 seconds for T and a value of 50% for P .

We have performed some initial experiments varying T and P in order to find optimal values for them but surprisingly have not seen much variation in behavior. We have explored values for P ranging from 90% to 10%. Attack detection rates increase as P decreases, but the detection rate increase is very slight. Varying P anywhere from 80% to 20% makes only about a 5% difference in detection rate, while dropping P below 20% results in increasing amounts of false positives. We have also explored values for T ranging between 10 seconds and 2 minutes. Increasing T from 10 seconds to 1 minute tends to increase attack detection rate a small amount, but increasing T above 1 minute makes nodes increasingly "paranoid" and increases the amounts of false positives. Our initial results show that our values for T and P may not be optimal, but they are within 5% of the best attack detection rates we have been able to generate without seeing false positives.

3.4 Reacting to Attacks

The virtual exponential network will be used in place of the active scale free network during an attack. Upon detecting a malicious attack, a node will begin replacing its dead neighbors with nodes from its virtual neighbor list, thus making the exponential network active as the scale free network disintegrates due to the attack. Also after detection, the node counts the number of neighbors that have failed recently and replaces each of these with a virtual neighbor. The node creates an active connection with the virtual neighbor and removes it from the list of virtual neighbors. During the attack, the node replaces each neighbor that dies with another virtual neighbor, until the node determines that the attack is over, or it runs out of virtual neighbors. Also, it is important that the node not seek out new virtual neighbors to replace its lost virtual neighbors during an attack, as the extra traffic would increase the stress on the already failing network. Once a node no longer detects an attack, it returns to normal operations using its current list of active neighbors, including those added during the attack. As a part of normal operations, it rebuilds a new list of virtual neighbors.

4. Our Experimental Model

Before we propose a solution to making the protocol more robust, it is necessary to describe some assumptions we make about Gnutella and the corresponding models that we built. As it stands, Gnutella is a very free protocol in the sense that it only prescribes the types of messages sent between nodes and leaves the remainder up to the client [7]. Some of the open properties of Gnutella include bootstrapping, the number of neighbors maintained by each node, the up-time distributions of nodes, and the initial topology. We model the first two properties (bootstrapping and number of neighbors maintained by each node) by adopting the method used by the most common Gnutella client,

Bearshare [4]. For the up and down time distributions we use a power law relationship fitted to measurements of node uptimes, and for the initial topology we use data collected through crawling the network [12].

4.1 Bootstrapping

There are many ways that Gnutella clients handle bootstrapping. This used to be done by connecting to a well-known point, Gnutellahosts.com, which maintained a list of servants with high availabilities that were likely to accept incoming connections [12]. On a connection from a newcomer, gnutellahosts.com returned one of these available servers. Bearshare works similarly in the sense that clients connect to a well-known Bearshare server (public.bearshare.net) and receive a list of servants from it [4]. Clients then directly connect to these servants. This preferential connection mechanism is responsible for the construction of a scale-free topology.

4.2 Number of Connections

The number of connections that a node maintains is not specified by the Gnutella protocol, so it is entirely left up to the specific client. Also finding out which client a node is running is not a trivial problem. We believe that a reasonable thing to do here is adopt the default value used by one of the most popular clients, BearShare. This client is set with a default max of 10 and min of 3 neighbors to maintain. We use these max and min values in modeling our nodes in our experiments.

4.3 Initial Topology

All of our experiments are based on a measured Gnutella topology that contains nodes that have been up for a 12-hour period [12]. This data contains roughly 3,000 active nodes and the edges between them. We built an initial topology by adding 17,000 nodes to this base topology using the bootstrapping and uptime models discussed above, resulting in a network with roughly 20,000 nodes alive at any given time.

4.4 Up and Down Time Distributions

A study of Gnutella shows that there are somewhere between 10,000 and 30,000 nodes that are alive at any given time slice [6]. Our initial topology contains 3,000 nodes that are alive during the course of a 12-hour period. We model the uptime distribution using the power law $P(t) = t^{-a}$, where $P(t)$ is the probability that a node is up for time t . The power law is an accurate model to use, because the majority of Gnutella nodes are up momentarily while fewer and fewer are up progressively longer. We conservatively estimate that the shortest period of time a node can be up is 1 minute, since our experience shows this to be the time required to sign on and off from the system.

4.5 Simulation and Data Collection

The results we use to measure the effectiveness of our recovery method were generated using a general P2P network simulator that we have created. The simulator provides a framework for maintaining peers and delivering messages between them. It includes support for peers joining the network, peers leaving the network, and removing peers suddenly from the network. This last feature allows us to simulate an attack by removing the most highly connected peers. To simulate Gnutella, we created two Gnutella clients compatible with the simulator: one based on Bearshare with the base caching and failure recover methods, and another with the additional attack detection and recovery method built into it. We ran our experiments using the models and topology information described above to generate the results presented in the following sections.

4.6 Experiments

We ran two versions of our experiment: a control case with the base recovery method, and a test case with the attack recovery method enabled. Each version was run 10 times and plotted on a 95% confidence interval. In our experiment, we simulate an attack on the network as removal of 5% of the most highly connected nodes. We removed these nodes over a period of 5 minutes, a period measured as a common length for denial of service attacks [11]. Experiments were run for 10 minutes, with the attack starting at 2 minutes.

5 Results

We chose metrics that would represent the fragmentation of the system, and end-to-end performance of the system from the point of view of a user. To measure fragmentation in the system, we calculated the percentage of living nodes in the largest connected component. A connected component is a set of nodes where a path exists between all pairs of nodes [5]. This measure shows whether the majority of nodes are in one connected component, or fragmenting off into other clusters of nodes. Another measure of fragmentation is the number of connected components, which measures the difficulty of rebuilding the network after the attack has done its damage. To measure end user performance, we calculated the average number of nodes reachable within 6 hops of any node as a percentage of the total number of living nodes. We chose 6 hops, because it is roughly the maximum hop count used when performing a search [6]. Because Gnutella is mainly used for searching and downloading files, this metric gives us an end-to-end measurement of the average number of nodes a user will be able reach with a query.

5.1 Connected Component Results

Figure 2 shows the percentage of living nodes that are in the largest connected component. These results show that without our recovery method, this percentage drops significantly from nearly 100% to below 30% in less than 30 seconds after the attack begins. With our recovery method in place, this metric stays almost constant, initially dipping to just above 98% before quickly returning to the 99.9% range. This initial dip is due to the delay between the beginning of the attack and the detection of the attack by individual nodes. Once nodes start detecting the attack, they are able to keep the network from fragmenting and maintain a very high percentage of nodes in the most highly connected component.

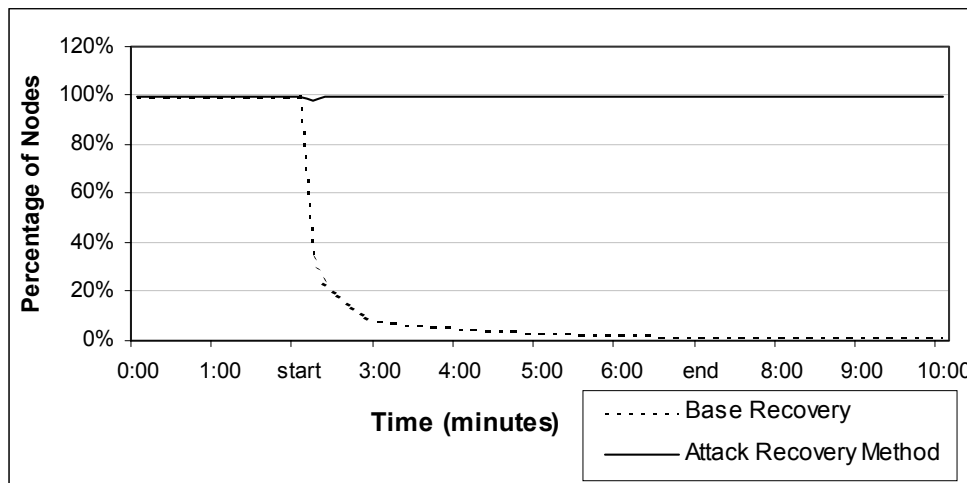


Figure 2. Percentage of nodes in the largest connected component

Figure 3 shows the number of connected components in the system over time. The steady state value measured is close to 5 before the attack, counting the largest connected component and a few nodes that have just come online but have not bootstrapped onto the network yet. These results show that without our recovery method the number of connected components increases very rapidly as soon as the attack begins, yielding hundreds of network fragments. With our recovery method, there is an increase in the number of connected components during the attack as more nodes are temporarily split from the network, but it returns to the initial steady state after the attack.




Figure 3. Number of connected components

5.2 Average Percentage of Neighbors Reachable in 6 Hops

Figure 4 shows the average percentage of nodes reachable within 6 hops in the network. Again, the results show that without our recovery method this percentage drops off significantly from roughly 25% before the attack to under 1% in just 30 seconds. With our recovery method, this percentage drops below 4% initially but begins to recover immediately. The initial drop in nodes reachable within 6 hops is still severe as a result of the most connected parts of the network being removed, but the recovery method is able to increase this measurement to 12% over the length of the attack. Our method shows a return to roughly half of the pre-attack performance of the system in terms of query effectiveness.

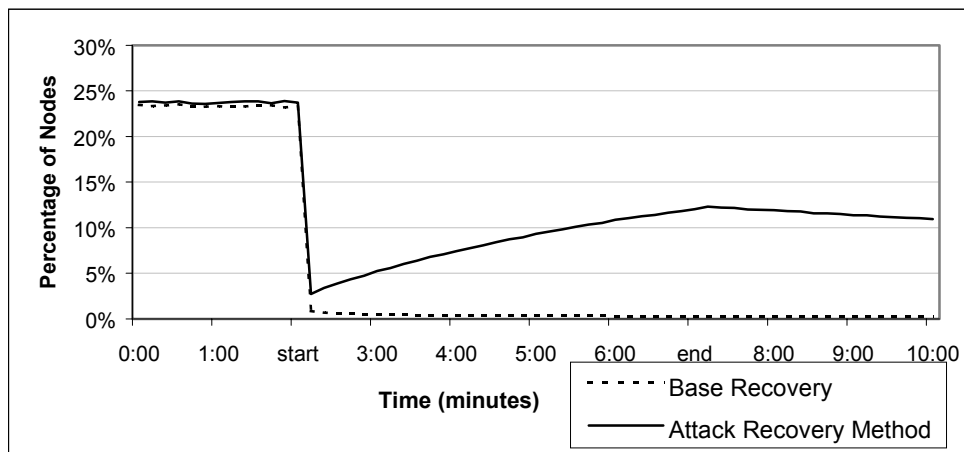


Figure 4. Average Percentage of Nodes Within 6 Hops

5.3 Random Node Discovery

Figure 5 shows the effectiveness of RDPs in discovering nodes in a non-preferential way. The number of times a node is selected, or frequency, is plotted against the percentage of nodes having this frequency. The binomial distribution, showing the results for node selection by random independent trials, is plotted against the distribution measured using RDPs. These results show that our node discovery technique succeeds in selecting nodes in a non-preferential way in order to form an exponential network.




Figure 5. Frequency distribution for number of connections

5.4 Attack Detection Mechanism

Figure 6 shows the effectiveness of our attack detection mechanism at differentiating random failures in the network from true attacks. In this experiment, 5% of the nodes were removed over a period of 5 minutes as discussed in our experimental setup, but the way in which nodes were selected was varied. The solid line shows the attack detection rates if the most highly connected nodes are removed, while the dotted line shows the rates if random nodes are removed. Less than 1% of the nodes detect an attack when removing random nodes, compared to as many as 14% when removing highly connected nodes. Also, it is noteworthy to mention that no attack is detected outside of the attack interval, despite nodes entering and leaving the system constantly.

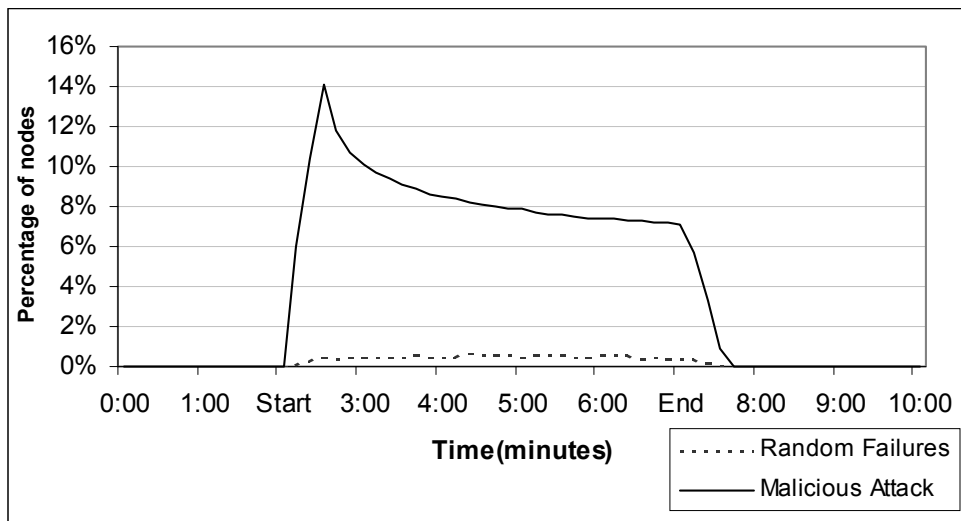


Figure 6. Percentage of nodes detecting attacks

5.5 Messages per Node

Figure 7 represents the average number of messages (pings and pongs) generated per node per second during the experiment. Only messages related to node discovery and connection maintenance are counted, leaving out all query traffic. This graph shows that our recovery mechanism requires only a 20% increase in the amount of traffic used to maintain the system during the steady state before the attack. When compared to query traffic, this increase is less significant. The recovery method increases traffic on the network once the attack begins because of nodes connecting to their exponential backups and seeking new backups with RDPs. Once the attack ends, traffic levels

return to the steady state. As for the standard client, traffic slowly decreases over the length of the attack, as fewer connections must be maintained.




Figure 7. Average number of pings

6. Conclusions

We have shown the current vulnerability of Gnutella to malicious attacks and the steps needed to correct this vulnerability. Peers must be able to detect attacks and create a backup plan for them in advance.

Experimental results show that our recovery method vastly improves the robustness of the Gnutella overlay network to attacks. Fragmentation is all but eliminated, and there is improvement in the effectiveness in querying during an attack. Also, the overhead of the method is small. Unfortunately, our recovery method would need to be adopted by a large percentage of users for it to work effectively. Only clients with our method in place would be able to forward the random discovery ping, meaning that a small percentage of clients running the recover method would not be enough to make any improvement. Without random node discovery, the method cannot work.

Our recovery method has an advantage that there is very little incentive to lie, an important quality given that peers tend to deliberately misrepresent themselves if there is any incentive to do so [12]. The only additional information that peers need to report is the number of connections they have. Reporting this as too high would lessen the likelihood of a peer being connected to a backup neighbor during an attack, while reporting this as too low would give the peer too many backups during an attack and potentially make it a target. Also, this information is something that peers can easily measure if an incentive to lie existed.

7. Future Work

We believe there are certain aspects of our recovery method that can be studied in greater detail. For example, we would like to investigate other techniques for random node discovery. Another area that we would like to look into is bringing the network back to a scale-free topology after an attack is over. Our initial premise was that exponential networks are better at surviving attacks, but a scale-free topology is ideal for handling normal operation including query forwarding and surviving random failures. Finally, we would like to finish our experiments, varying the parameters of the attack detection mechanism in order to find their optimal values.

A potential limitation of our recovery method is that it fails to take into account how suitable a given peer is for the task assigned to it [12]. In our method, this relates to the suitability of a peer selected as a backup to hold the network together during an attack. Much of the recovery work is done by the least highly connected nodes, which may be dialup users who are unable to handle the added stress during attacks, or users who are unwilling to handle the added stress. This limitation is lessened given the high redundancy of the exponential backups but is still a potential problem that we did not model.

Acknowledgments

We would like to thank Stefan Saroiu and Steven Gribble for providing us with topological data they collected on Gnutella. Armando Fox and George Candea are our advisors on this project, and without their advice our work would have been much harder. Finally, we would like to thank Jim Gray and Brendan Murphy for their advice on this project.

References

- [1] W. Aiello, F. Chung, and L. Lu. "A Random Graph Model for Massive Graphs." Symposium of Theory of Computing, 2000.
- [2] R. Albert, H. Jeong, and A. Barabási, "Error and attack tolerance in complex networks," *Nature* 406
- [3] A. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, 286
- [4] <http://www.bearshare.com/>
- [5] F. Buckley and F. Harary, *Distance in Graphs*. Addison-Wesley, New York, 1990.
- [6] www.clip2.com, Gnutella Measurement Project
- [7] <http://www.clip2.com/GnutellaProtocol04.pdf>
- [8] C. Faloutsos, M. Faloutsos and P. Faloutsos, "On power-law relationships of the Internet Topology," Proc. of ACM SIGCOMM, Aug. 1999.
- [9] http://www.gnutelliums.com/linux_unix/gnut/doc/gnutella-prot.html
- [10] A. Medina, I. Matta and J. Byers, "On the Origin of Power Laws in Internet Topologies," *ACM Computer Communication Review*, vol. 30, no. 2, Apr. 2000
- [11] D. Moore, G. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," in Proceedings of the 2001 USENIX Security Symposium.
- [12] S. Saroiu, P. Krishna Gummadi, and S. Gribble, "A measurement study of peer-to-peer file sharing systems," Technical Report UW-CSE-01-06-02, University of Washington, June 2001.
- [13] <http://www.denialinfo.com/>
- [14] R. Cohen, K. Erez, D. ben-Avraham and S. Havlin, "Breakdown of the Internet under intentional attack," *Phys. Rev. Lett.* 86, 3682 (2001)
- [15] B. Yang and H. Garcia-Molina, "Comparing hybrid peer-to-peer systems," In Proceedings of the 27th International Conference on Very Large Databases, September 2001.
- [16] M. Ripeanu, I. Foster and A. Iamnitchi, "Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System", *IEEE Internet Computing Journal special issue on peer-to-peer networking*, vol. 6(1) 2002.
- [17] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, "Tapestry: An infrastructure for fault-resilient wide-area location and routing", Technical Report U. C. Berkeley, April 2001.
- [18] A. Rowstone, P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems", in *Middleware*, 2001.