Available online at www.sciencedirect.com

SciVerse ScienceDirect

Procedia Technology 3 (2012) 61 – 68

Procedia Technology

The 2012 Iberoamerican Conference on Electronics and Computer Science

Evaluating P2P Networks against Eclipse Attacks

Francisco de Asís López-Fuentes*, Iñaki Eugui-De-Alba, Otoniel M. Ortíz-Ruiz

Department of Information Technology, Universidad Autónoma Metropolitana-Unidad Cuajimalpa, Av. Constituyentes 1054, Col. Lomas Altas, México, D. F., 11950, México

Abstract

Due to its distributed nature, the P2P networks are subject to more intricate attacks than client-server networks. Therefore, the security issues on a P2P network represent a great research challenge. Most P2P networks do not have a central management, and causes them can be endangered by its own nodes if enough of them decide to behave maliciously. In this paper, we evaluate the performance of P2P networks against an Eclipse attacks. Specifically, we evaluate the performance of the unstructured P2P networks and super-peer P2P systems against the Eclipse attack. To this end, our simulations are realized using schemes based on Gnutella and KazaA. Both networks are evaluated in terms of number of affected peers and their performance. Our results show that an Eclipse attack may have different effects on different kinds of P2P networks.

Keywords: Peer-to-Peer networks, security, overlay network, Eclipse attack

1. Introduction

Peer-to-Peer (P2P) networks have become a popular concept during the last years. Networks such as BitTorrent [1] and eMule [2], make it easy for people to find what they want and share what they have. However, file sharing with anonymous and unknown users on the Internet could go against the basic rules of security. P2P networks were not originally designed to withstand an adversary attack and can be easily

* Corresponding author. Tel.:+52-55-91776650; fax:+52-55-91776650.
E-mail address: fcoasis@correo.cua.uam.mx.

doi:10.1016/j.protcy.2012.03.007

compromised. Due to their distributed nature, P2P networks are subject to more intricate attacks than client-server networks. Therefore, security issues on a P2P network represent an important research challenge. Most P2P networks do not have central management, and can therefore be endangered by its own nodes if enough of them decide to behave maliciously. A malicious node may lie about its characteristics misdirect requests from other nodes requesting its assistance or take advantage of its role within a P2P network. Thus, P2P network security represents a rapidly growing research field.

A series of attacks are aimed at P2P networks consistency. These attacks are mainly focused on exploiting the networks's Distributed Hash Tables, and their goal is to disrupt the communications among the network nodes. The most potentially harmful for the P2P networks are the Sybil and Eclipse attacks [3]. In a Sybil attack, a malicious entity creates multiple false identities of itself and uses them to influence the system behavior. Douceur [4] introduced a formal model to explain this attack type. A system affected by the Sybil attack can provide false information to other legitimate nodes, and create a majority of colluding malicious nodes in the overlay network. Another common attack in a P2P network is the Eclipse attack. In this attack, an entity uses multiples identities to conspire and to cut off traffic going to and from a particular legitimate node. Here, the goal is to eclipse legitimate nodes from the network. An Eclipse attack is more general than a Sybil attack, but the attackers can use a Sybil attack to launch an Eclipse attack. However, defense against Sybil attacks does not prevent Eclipse attacks because attackers may manipulate the overlay network to mount an Eclipse attack [5]. Thus, a malicious entity can use both attacks to facilitate the execution of many other attacks. For example, many malicious identities in the system can control the majority of the replicas for a given key by polluting the routing tables of honest nodes [6].

In this paper, we evaluate the performance of P2P networks under adverse attacks. Specifically, we use unstructured P2P networks, because these systems are designed more specifically for heterogeneous and distributed environments. An Eclipse attack is simulated for two different unstructured P2P schemes in this work. To this end, the schemes used in our simulations are based on Gnutella [7] and KazaA [8]. The remainder of this paper is organized as follows. We introduce the general background knowledge about the P2P structures and analyze their vulnerabilities in Section 2. Subsequently, in Section 3, we focus on the simulation of the Eclipse attacks in two P2P structures and evaluate the damages. Finally, our conclusions about the overall paper are given in Section 4.

2. Background

A P2P communication network is formed by a group of nodes located in a physical network. These nodes build a network abstraction on top of the physical network known as an overlay network, which is independent of the underlying physical network. P2P networks are classified mainly into two categories: structured and unstructured. This classification is based on how the nodes in the overlay structure are connected to each other. Structured P2P networks maintain a close coupling between the network topology and the location of data via a hash table (DHT). On the other hand, an unstructured P2P network is formed when the logical links among participating nodes are established randomly. Unstructured P2P networks can be further divided in: centralized P2P, pure P2P and hybrid P2P [9]. In this section we present the unstructured P2P schemes evaluated in this paper.

2.1. Pure P2P networks

Pure P2P systems are used more often for heterogeneous and distributed environments [10], where maintaining strict restrictions on control data placement and the network topology is not possible. In a pure P2P network, the queries have to be flooded through the network, which causes a high amount of traffic in the

network. In these systems all peers are equal and no peer holds any permanent information about which objects are stored where. No directory with the data of the peers which are a part of the network exists. Figure 1 shows this scenario.

Examples of pure P2P architecture are Gnutella 0.4 [7] and FreeNet [11]. Gnutella was designed by Justin Frankel and Tom Papper in early 2000, and it is still the top P2P network on the Internet with an estimated market share of more than 40%.

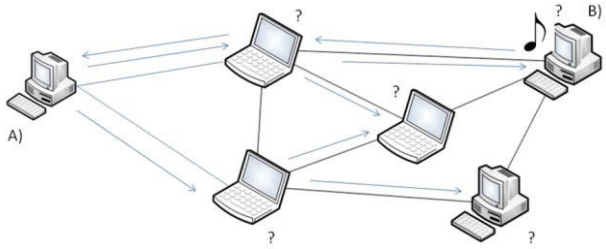

Fig.1. A pure P2P network

2.2. Hybrid P2P networks

A hybrid architecture attempts to strike a balance between the accuracy of the centralized architecture and the lower load of the pure architecture. An example of hybrid P2P structure is super-peer P2P systems. Yang and Garcia-Molina [12], state that a super-peer is a node in a peer-to-peer network that operates both as a server to a set of clients, and as an equal in a network of super-peers. Thus, the super-peer networks strike a balance between the inherent efficiency of centralized search, and the autonomy, load balancing and robustness to attacks provided by distributed search. Furthermore, each super-peer takes advantage of the heterogeneity of capabilities across peers such as bandwidth and processing power. Figure 2 shows an example of a hybrid P2P network.

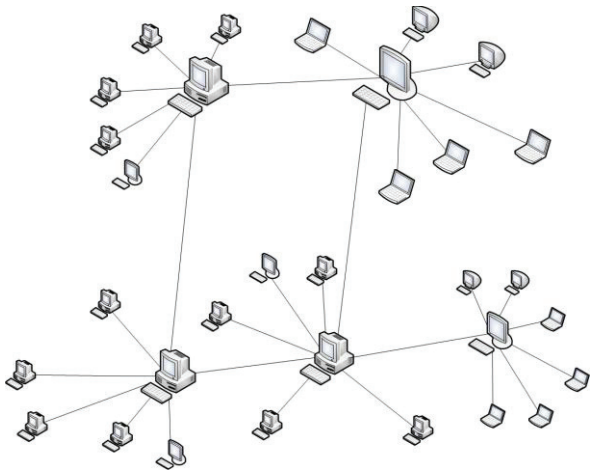

Fig.2. A hybrid P2P network

A P2P system that has adopted the super-peer model in its design is KaZaA [12], which was introduced in March 2001. KazaA [8] is a digital music service, which allows instantly unlock millions of top songs on demand. A detailed description of how KaZaA works is given in [13]. KazaA has become one of the largest distributed systems ever to be deployed in the Internet.

2.3. Eclipse Attack

An import problem in the P2P networks is the security. Due to its distributed nature, the P2P networks are subject to more intricate attacks than client-server networks [14]. Main attacks are aimed at P2P networks consistency. A general attack in P2P overlay networks is the Eclipse attack [3].

In an Eclipse attack, an attacker controls a large part of the neighbors of a good node. Thus, the attacker may manipulate the overlay network and control the majority of the honest peers. In this situation, the union of malicious peers can work together to fool the good node by writing their addresses into the neighbor table of a good node [5]. By using an Eclipse attack, an attacker can control the most significant part of overlay network. In addition, large scale malicious nodes can eclipse more good nodes to control the entire overlay network. Figure 3 shows an Eclipse attack in a hybrid P2P overlay network. The overlay nodes cannot forward messages correctly and the whole network cannot be managed.

During an Eclipse attack, the in-degree of attacker peers in the P2P network must be much higher than the average in-degree of good peers in the P2P network [5]. Here, in-degree refers to the number of direct routes coming into a node, while out-degree refers the number of direct routes going out of a node. Good peers can choose the neighbors whose in-degrees and out-degree are below a certain threshold. But malicious nodes can also make a poisoning attack by consuming all the in-degree of good peers. Hence, the good peer will never choose a neighbor.

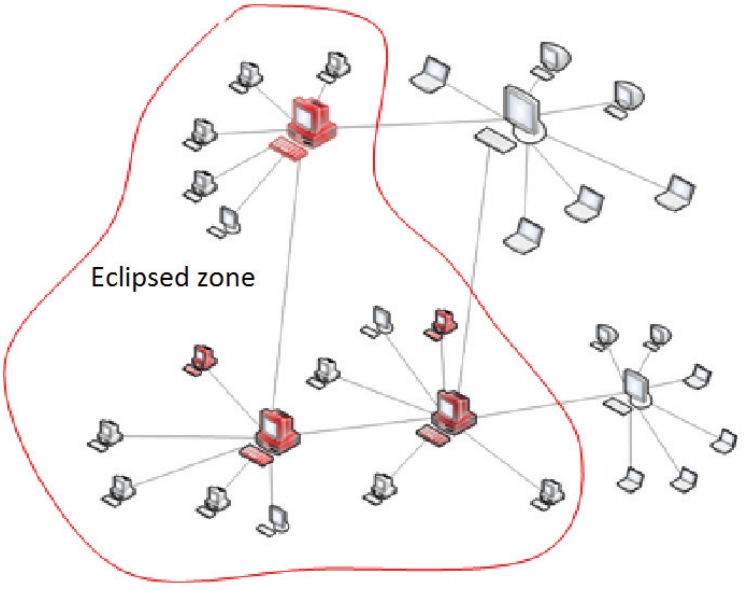

Fig.3. An Eclipse attack

3. Evaluation

In this section, we evaluate two unstructured P2P architectures against an Eclipse attack. Specifically, our evaluated structures are based on the Gnutella and KazaA schemes. We try to answer the following questions:

- How do client nodes of a super-peer P2P network are affected by an Eclipse attack?
- How super-peers are affected by an Eclipse attack?
- How does an Eclipse attack affect the performance of a pure P2P network or a super-peer P2P network?

3.1. Experimental setup

Our simulation scenario is implemented on the Peersim P2P simulator [15], which has been designed to support extreme scalability and dynamicity. Peersim is written in Java and it is composed of two simulation engines: a cycle-based one and event-based one. Both engines are supported by many simple, extendable, and pluggable components, with a flexible configuration mechanism. The experiments deployed in this paper were developed in the simulator based on cycles. For our experiments, we use architectures based on super-peer [16] and Gnutella [17].

To simulate the Eclipse attack against the Gnutella network, we use a structure with 100,000 peers, which are attacked by 100 malicious peers. During the Eclipse attack the malicious peers are inserted in the network. Once these malicious peers have established links with other peers, these malicious peers recommend to other malicious nodes in order to distribute throughout the network. This attack tries to take the routing control of the network and then eclipse the network. Once links have been established, the malicious nodes blocked all communication, preventing the information reach its destination. This way the network capacity is affected. The Eclipse attack is evaluated in two parts. The first attack is present when the network is initialized (start-up) and the second attack takes place when the network is working. It has been designed this way in order to observe the time at which the network is most vulnerable.

To simulate the Eclipse attack against the super-peer network, we follow a procedure similar to the one previously described.

3.2. Eclipse attack to pure P2P network

We evaluate the impact of an Eclipse attack on a Gnutella overlay network in this section. For this case, we conducted two experiments. The first experiment measures the number of peers that are neighbors of malicious peers, in order to observe how the good peers can be affected by the malicious peers. We simulated an Eclipse attack when the network is starting-up and when it is operating. Results obtained were similar in both cases. This is because in a pure architecture the changes are unpredictable due to clients joining and leaving the network constantly. The fig. 4 shows how with 100 malicious peers almost other 700 neighbor peers are affected. Although this is a significant harm, it is not very high. This is because the Eclipse attack is very difficult to implement in a pure P2P architecture since the network structure is unknown and the most vulnerable peers cannot be identified.

Our second experiment evaluates how the network capacity is affected when an Eclipse attack occurs. Like the other experiments this is done when the network is starting and when it is operating. In both cases the network is affected in a similar way. Figure 5 shows how the capacity of the network is affected slightly. That occurs because the attacker does not know how the network is constructed. There are no peers with special functions that let it affect other peers. Thus, the Eclipse attack can only reach a limited number of nodes.

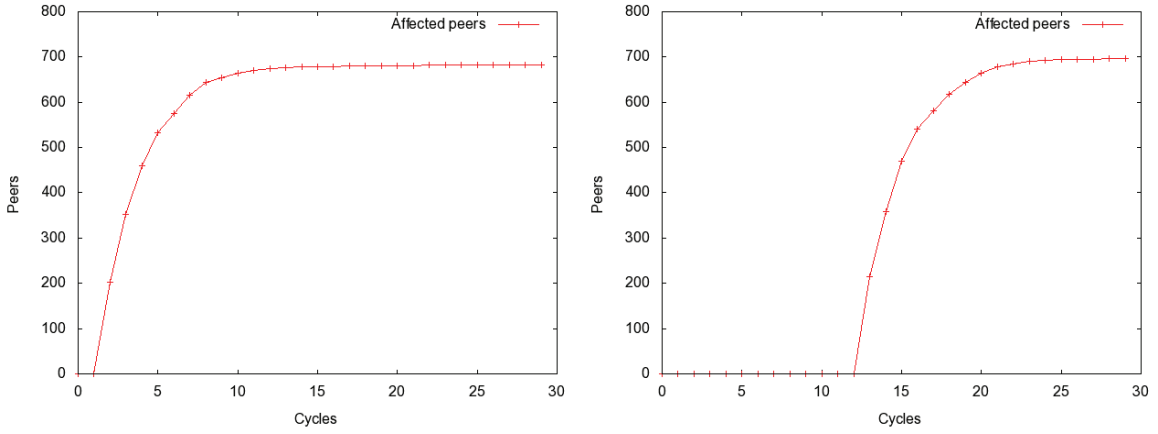

Fig. 4. Number of peers affected by malicious peers. a) network initialization, b) network operation

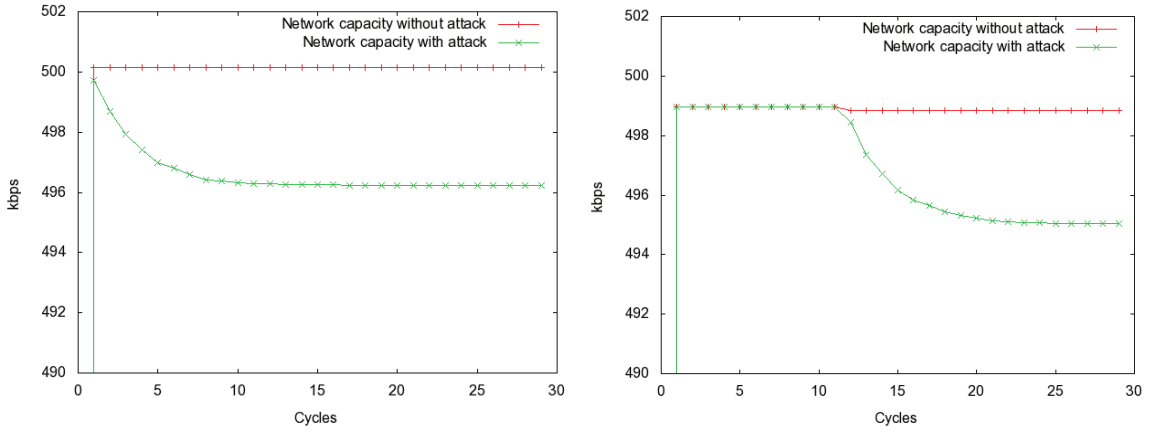

Fig. 5. Network capacity affected by an Eclipse attack. a) network initialization, b) network operation

3.3. Eclipse Attack to a super-peer P2P network

In this section we evaluate the impact of an Eclipse attack against a super-peer P2P overlay network. In this case, first we evaluate how many malicious peers could be super-peer in two different moments: one when the network is initialized and other when the network is operating normally. Our simulations indicate that when a network is in the initialization phase around 90% of the malicious peers could become super-peers. On the other hand, if the Eclipse attack occurs when the P2P network is already operating, around a 40% of the malicious peers could become super-peers.

Our second experiment analyzes the number of affected clients during an Eclipse attack. Figure 6 depicts this scenario. Results show that the number of affected clients is considerably superior when the network is starting than when it is already operating.

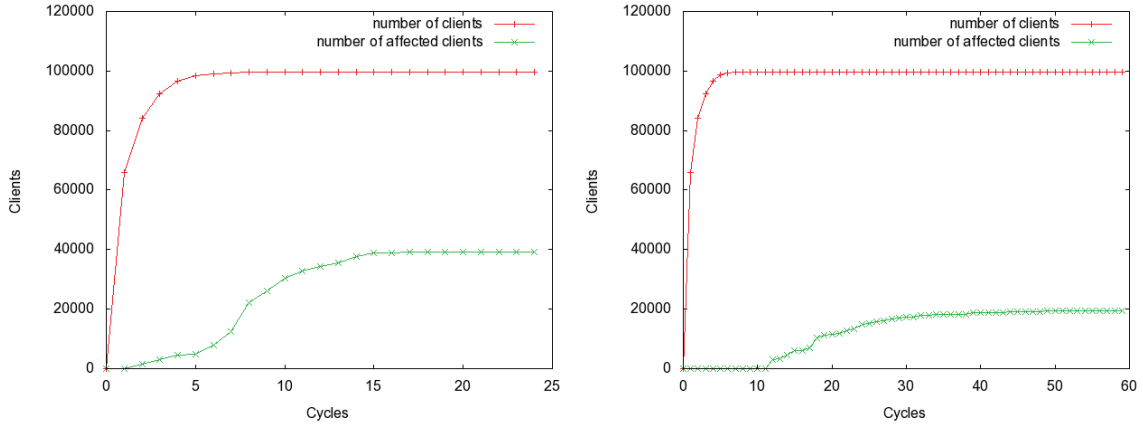

Fig. 6. Affected peers during an Eclipse attack. a) network initialization, b) network operation

Finally, we evaluate the overall performance of the network during an Eclipse attack. In both cases the network is affected significantly. However, the network capacity is more affected when the Eclipse attack occurs during the network initialization. Figure 7 shows this scenario. This happens because when the network is being constructed there are more malicious peers as super-peers and more clients can be blocked. Contrary to that, when the network is operating there are less malicious nodes as super-peers and fewer clients can be blocked. Thus, 0.1% of malicious peers can affect around 40% of the network capacity.

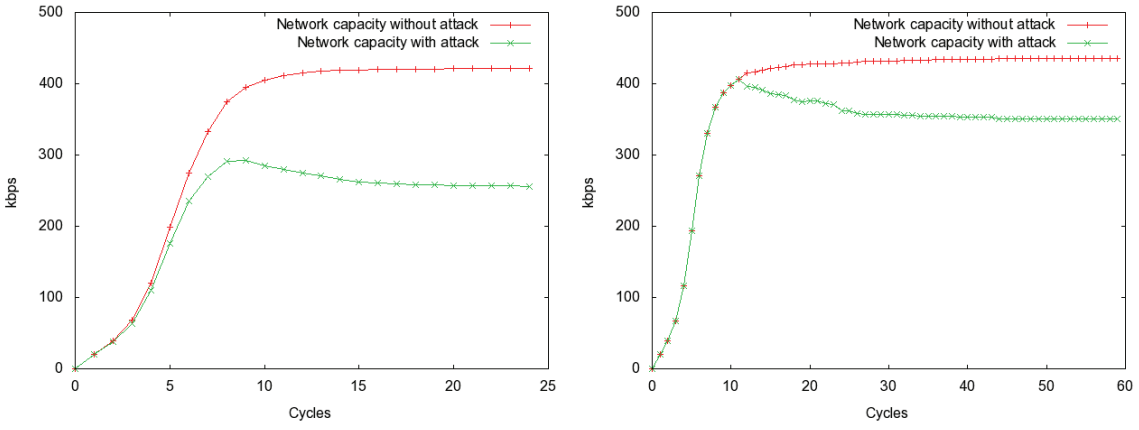

Fig. 7. Network capacity affected by an Eclipse attack. a) network initialization, b) network operation

4. Conclusions

An Eclipse attack is a general attack in overlay P2P networks. In this paper we analyzed Eclipse attacks against two type of unstructured network: pure P2P and super-peer P2P networks. An Eclipse attack can occur when the network is starting and when it is operating. Our results show that in both cases the pure P2P network is affected in a similar way. This is due to the attacker not knowing how the pure P2P network is

constructed. Contrary to that, an Eclipse attack has a different effect in a super-peer P2P network, if it occurs when the network is initialized or during its operation. Our work can be extended in different directions as future work. First, we can study which of the two unstructured P2P architectures (pure P2P and Super-peer) are more tolerant to Eclipse attacks. Then, we can design a mechanism to avoid or reduce the negative impact of this attack against this type of network. Second, we can extend our study about the Eclipse attack to structured P2P networks.

References

[1]. Cohen, B. Incentives build robustness in BitTorrent. *1st P2PECON*, Berkeley, CA, USA, 2003.
[2]. E-mule-project homepage: http://www.emule-project.net/, seen on November 2011.
[3]. Singh A, Castro, M, Druschel, P and Rowstron, A. Defending Against Eclipse Attacks on Overlay Networks. *11th ACM SIGOPS European Workshop*, Leuven, Belgium, 2004.
[4]. Douceur, JR. The Sybil Attack. 1st IPTPS*, LNCS*, Springer-Verlag 2002, 2429, 251–260.
[5]. Singh, A, Ngan, TWJ, Druschel, P. and Wallach, DS. Eclipse Attacks on Overlay Networks: Threats and Defenses, *25th IEEE INFOCOM*, Barcelona, Spain, 2006, p.1-12.
[6]. Urdaneta, G, Pierre G and Van Steen, M. A Survey of DHT Security Techniques. *ACM Computing Surveys*, 2011, 43(2), p.8-8.
[7]. Gnutella 0.4, Gnutella homepage: http://rfc-gnutella.sourceforge.net/developer/stable/index.html, seen on November 2011.
[8]. KaZaA file sharing network. KaZaA homepage. http://www.kazaa.com/, seen on November 2011.
[9]. Eberspächer, J and Schollmeier, R. First and Second Generation of Peer-to-Peer systems. *P2P Systems and Applications, LCNCS*, Springer-Verlag, 3485, 2005, p. 35-58.
[10]. Li, J and Vuong, S. An Efficient Clustered Architecture for P2P Networks. *18th AINA*, Fukuoka, Japan, 2004, p. 278-283.
[11]. The Free Network Project. FreeNet homepage: http://freenetproject.org/, seen on November 2011.
[12]. Yang, B and Garcia-Molina, H. Designing a Super-peer Network. *19th ICDE*, Bangalore, India, 2003, p. 49-60.
[13]. Liang, J, Kumar, R and Ross, K. Understanding Kazaa, *Technical Report*, Polytechnic University Brooklyn, NY, USA, 2004.
[14]. Seedorf, J. Security challenges for peer-to-peer SIP, *IEEE Network*, 20(5), 2006, p. 38-45.
[15]. Montresor, A and Jelasity, M. PeerSim: A scalable P2P simulator, *9th IEEE P2P*, Seattle, WA, USA. 2009, p. 99-100, PeerSim software downloaded on March 25, 2011, from http://peersim.sourceforge.net/
[16]. Montresor, A. A Robust Protocol for Building Super-peer Overlay Topologies. *4th IEEE P2P,* Zurich, Switzerland, 2004, p. 202-209.
[17]. Gnutella-peersim software downloaded on March 25, 2011, from http://code.google.com/p/gnutella-peersim/.