

A Testbed-based Analysis of the Incorrect Lookup Routing Attack on the Pastry DHT

Christian Gottron, André König and Ralf Steinmetz
Multimedia Communications Lab (KOM), TU-Darmstadt, Darmstadt, Germany
{christian.gottron, andre.koenig, ralf.steinmetz}@kom.tu-darmstadt.de

I. INTRODUCTION

Distributed Hashtables (DHT), with Pastry [2] as a prominent representative, are one of the most recent peer-to-peer (P2P) architectures. DHTs scale well to the network size due to a small routing table and a hierarchical routing scheme. In Pastry, the number of routing entries scales logarithmically with respect to the number of participating nodes, because each node maintains route entries to only a few nodes per hierarchical (sub) domain. Therefore, routing requests have to be forwarded by intermediate nodes to the final destination. In order to route successfully, each node has to behave benignly according to the protocol specifications. The *Incorrect Lookup Routing* attack [4] misuses this routing scheme to launch a denial of service attack on the P2P network. While this kind of routing attack was extensively analyzed theoretically and by simulations, to the best of our knowledge, testbed based evaluations have not been performed to date. Such an evaluation may reveal unnoticed influences or (at least) verifies the impact of malicious nodes on a P2P network. Therefore, a testbed evaluation of the *Incorrect Lookup Routing* attack is presented in this paper.

II. BACKGROUND

A. Pastry

Pastry [2] is a structured P2P architecture. The routing algorithm is recursive and based on comparing the prefixes of node-IDs. Whenever a node receives a request, the ID of the requested object is compared with the ID of the node. As objects are maintained by the node with an ID next to the object-ID, the node has to check whether itself or a direct neighbor is the destination of this request. Therefore, the *leafset* routing table provides links to nodes with a node-ID close to the ID of the routing table owner. The *main* routing table has to be used when neither the node itself nor a neighbor is the destination of the request. This routing table is prefix based and, therefore, is capable of providing a link to a node with an ID that is one digit closer to the object-ID. As a result, the request is forwarded in a recursive way to the destination. Whenever two nodes with the same prefix are known to a node, the node providing a lower round trip time (RTT) is stored at the *main* routing table. The routing tables of each node provide only links to a fraction of the nodes available in the network. Therefore, nodes have to cooperate to route requests successfully.

The average number of hops h_{Pastry} required for a lookup request depends on the network size N and the configuration parameter b , which influences the routing table size. Therefore, an upper bound for the number of required hops can be derived by Equation 1 [2].

$$h_{Pastry} \leq \log_2^b(N) \quad (1)$$

B. Incorrect Lookup Routing

Sit and Morris [4] introduced the *Incorrect Lookup Routing* attack during which route requests are dropped or redirected. While dropping requests is a straightforward approach, redirecting messages to other malicious nodes is less conspicuous and harder to detect by countermeasures. Pastry is not capable of detecting either kind of attack and, therefore, we focus on the packet dropping variant of the *Incorrect Lookup Routing*, which is equally efficient but less complex.

Castro et al. [1] described the probability of a successful request by a general equation for all DHT-based P2P architectures (Equation 2). The probability σ of a successful routing depends on the fraction of malicious nodes in the network f and the average hops per request h . The fraction of malicious nodes is defined by the setup of the scenario. The average number of hops depends on the DHT architecture.

$$\sigma = (1 - f)^h \quad (2)$$

III. EVALUATION

A. Testbeds

1) *PlanetLab*: PlanetLab is a worldwide testbed for development and deployment of prototypes in a real-world environment. Larry Peterson (Princeton) and David Culler (UC Berkeley and Intel Research) initiated this project in 2002. Today, PlanetLab consists of more than 1000 nodes distributed around the world. While the hardware and the load of every node differ widely, the PlanetLab software used to manage access to the testbed and the operating system (Fedora based) is homogeneous.

2) *G-Lab*: The G-Lab project started in 2008 and is funded by the German 'Federal Ministry of Education and Research' (BMBF). The major objective of this project is to create a national network to develop and evaluate future Internet technologies. By now, more than 150 nodes distributed over Germany at six Universities (Berlin, Darmstadt, Karlsruhe, Munich, Kaiserslautern, Würzburg) are online. These nodes are equipped with homogeneous hardware and the overall load is low by now. For management, the PlanetLab software is

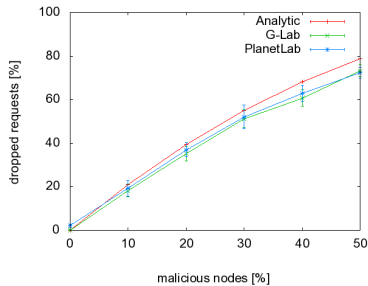


Fig. 1. Effects of the Incorrect Lookup Routing attack on the fraction of dropped lookup requests

deployed such that experiments designed for PlanetLab can be conducted on G-Lab without adaptations.

B. Settings

We used the FreePastry implementation of the Pastry DHT with default settings (parameter $b = 4$, leafset size of 32 nodes) to evaluate the *Incorrect Lookup Routing* attack. To provide object storage and maintenance services, FreePastry was extended by the Past [3] application. We distinguish between two kinds of Pastry nodes. Benign nodes behave correctly according to the protocol specifications. Each benign node requests an object every two minutes on average. Malicious nodes cooperate during routing table maintenance only and drop each received route request.

We evaluated the *Incorrect Lookup Routing* attack in both the PlanetLab and the G-Lab testbed. We selected 100 nodes from each testbed randomly and set up 5 virtual Pastry peers on each physical testbed node. Thus, we had an overall network size of 500 nodes per testbed. We varied the fraction of malicious peers from 0% up to 50% of the overall number of peers and evaluated each scenario for 30 minutes. We conducted 10 iterations per scenario and changed the malicious peers each time, thus equalizing the occasionally strong impact of single peers as described in the following.

C. Evaluation of the Incorrect Lookup Routing Attack

1) *Mathematical model*: In order to determine the Pastry specific equation for the probability of successful lookup requests we combined Equations 1 and 2 as shown in Equation 3. According to this, the lower bound for a successful lookup request depends on the network size, the number of malicious nodes, and the parameter b .

$$\sigma \geq (1 - f)^{\log_{2^b}(N)} \quad (3)$$

2) *Results of the testbed evaluation*: The averaged results of both testbeds are very similar as shown in Figure 1. As Equation 1 provides only an upper bound for the hop length, the theoretical packet delivery rate predicted by Equation 3 is lower compared to the testbed results. The gap between the averaged testbed and theoretical results increases further when the fraction of malicious nodes increases. Yet, the average results of the testbeds are reasonably close to the predicted results by the mathematical model.

On the other hand, analyzing single scenarios reveals effects that are not considered in the mathematical model. We discuss

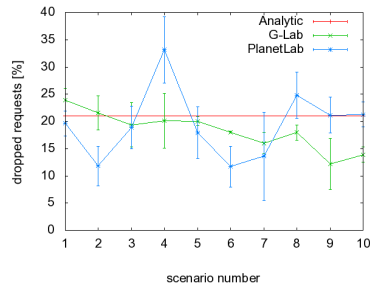


Fig. 2. Influence of the boot order on the fraction of dropped lookup requests

the scenario with 10% malicious peers in which we analyzed the impact of the distribution of malicious peers with respect to the order with which they join the DHT. In Figure 2 the packet loss ratio is shown according to the scenario number. In the first scenario, nodes 1 to 50 in the boot order behave maliciously. In the second scenario, nodes 51 to 100 behave maliciously and so on. The results of the two testbeds differ strongly. As nodes in PlanetLab differ strongly in network load and link quality, Pastry’s RTT-based proximity metric that determines which nodes are listed in the main routing table strongly affects the impact of the routing attack: malicious peers with a good connection are included in routing tables with a higher probability than benign peers with an average connection. In G-Lab, nodes are homogenous regarding load and connectivity. Therefore, the RTT is mainly equal. This results in an increased impact of malicious nodes which boot first during the scenario. As those nodes were distributed in the routing tables at the beginning and not replaced by nodes that provide a better RTT, their impact on the lookup process is increased.

To sum it up, the average testbed and mathematical model results match reasonably. Yet, effects of the proximity metric and boot order are not considered by the model up to now but have a considerable impact on the results of the *Incorrect Lookup Routing* attack.

IV. CONCLUSION

We evaluated the *Incorrect Lookup Routing* attack on the Pastry DHT based on a series of experiments in the PlanetLab and the G-Lab testbed. We compared the results with predictions obtained from an analytical model. Further, we analyzed the impact of the boot order and the connectivity of malicious nodes on the attack.

In our next steps we will evaluate routing attacks in networks with an increased size. Further, we will compare different countermeasures proposed for this attack analytically and in the testbeds.

REFERENCES

- [1] M. Castro et al. Secure routing for structured peer-to-peer overlay networks. *SIGOPS Oper. Syst. Rev.*, 36(SI):299–314, 2002.
- [2] A. Rowstron et al. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proc. of Middleware’01*, 2001.
- [3] A. Rowstron et al. Storage management and caching in past, a large-scale, persistent peer-to-peer storage utility. In *Proc. of SOSP’01*, 2001.
- [4] E. Sit et al. Security considerations for peer-to-peer distributed hash tables. In *Proc. of IPTPS’01*, pages 261–269, 2002. Springer-Verlag.