Simulations, Models, and Testbeds

A Mutual Catalysis

Ralf Steinmetz, André König
Multimedia Communications Lab (KOM)
Technische Universität Darmstadt, Darmstadt, Germany
{*Ralf.Steinmetz, Andre.Koenig*}*@KOM.tu-darmstadt.de*

Abstract

Recent developments in the area of decentralized and infrastructure-less systems opened avenues for novel applications. Along with these new technologies, new questions regarding their operational bounds in terms of e.g. scalability and security arose. Due to the sparse presence of real-world decentralized and infrastructure-less systems, new protocols and applications have to be scrutinized by means of simulation, in (small-scale) testbeds, and by analytical models. In this article, we discuss challenges of evaluating security mechanisms for mobile ad hoc networks and peer-to-peer systems. We focus on harmonizing predictions of analytical models and results obtained from simulation studies and testbed experiments.

Keywords:
security, mobile ad hoc networks, peer-to-peer systems, simulation studies, testbed experiments, analytical models

1. Introduction

Recent developments in the area of decentralized systems and infrastructure-less networks opened avenues for novel applications. A prominent example are tools for supporting communication and coordination of on-site units in large-scale emergency response scenarios. Here, a combination of peer-to-peer systems and mobile ad hoc networks forms a communication substrate which operates without a preexisting infrastructure, thus offering en-

hanced communication services beyond the 'traditional' infrastructure-based client/server world. Along with these new technologies, new questions regarding their operational bounds in terms of e.g. scalability and security arose. Due to the yet sparse presence of real-world decentralized and infrastructure-less systems, new protocols and applications have to be scrutinized by means of simulation and in (small-scale) testbeds. Additionally, mathematical models must be developed both to verify the validity of results obtained from simulation and testbed studies and to provide tools for the online adaptation of relevant system parameters in future real-world scenarios.

In this article, we recapitulate parts of our latest research on decentralized systems and infrastructure-less networks, concentrating on evaluation challenges. We present recently finished and ongoing work of German research projects such as SicAri [1] and G-Lab [2] that focus (amongst other objectives) on the emergency response application scenario. We present two examples taken from our research on novel security mechanisms for (1) mobile ad hoc networks and (2) peer-to-peer systems. We briefly outline the research questions and present our proposed solutions, the evaluation methodologies we applied and the evaluation results we obtained. We then focus on challenges in harmonizing predictions of analytical models and results obtained from simulation studies and testbed experiments.

2. Challenges in Evaluating Security Mechanisms for Mobile Ad Hoc Networks

In this section, we present the challenges we were facing while evaluating a novel, location-based intrusion response mechanism for mobile ad hoc networks. The evaluation was based on an analytical model combined with simulation studies. We used a mutual feedback between model and simulation to obtain a reasonable evaluation methodology covering the parameter space of the system as far as possible. The following is based on prior work presented in [3] and [4].

2.1. Security in Mobile Ad Hoc Networks

Wireless data transmission and lack of a communication infrastructure turn mobile ad hoc networks into a challenging environment regarding security and, in particular, availability. Various attack vectors that exist in mobile ad hoc networks in addition to known attacks from infrastructure-based systems were identified. For survey information, we refer to [5] and

[6]. To deal with these new types of misbehavior, preventive as well as reactive security measures that are tailored to the conditions of mobile ad hoc networks were proposed. Well known preventive security mechanisms are secure routing protocols such as SAODV [7] or Ariadne [8]. Based on means of cryptography, these protocols were designed to prevent false routing information from being injected into the network as well as correct routing information from being tampered with. Though designed meticulously, vulnerabilities for both protocols were found recently in [9, 10, 11]. For the case of compromised preventive security mechanisms, reactive measures can be taken to establish a second line of defense. In general, reactive security measures consist of intrusion detection systems combined with intrusion response systems.

Intrusion detection systems for mobile ad hoc networks are well investigated; for survey information we refer to [12]. Yet, only sparse attention has been paid to intrusion response. In most intrusion response schemes, e.g. in [13, 14, 15], an address-based response is performed. Here, a misbehaving node is identified by its (network) address and transmissions to/from this address are blocked. Since nodes in a mobile ad hoc network are beyond the control of a central instance, changing addresses is possible with little effort. Therefore, an address-based intrusion response system may have significant drawbacks when deployed in mobile ad hoc networks. As an alternative, we proposed a location-based intrusion response system in [3]. Here, to exclude a misbehaving node from the network, a geographical quarantine zone is established around the node's location. Transmissions into or out of the quarantine zone are blocked. Thus, as shown in Figure 1, routes in the mobile ad hoc network are established bypassing quarantine zones. This way, communication is kept away physically from areas affected by misbehavior.

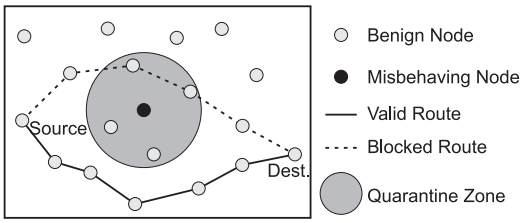

Figure 1: Schematic representation of the location-based intrusion response

In [3] we compared the performance of the location-based intrusion response system with an address-based solution when confronted with a com-

bination of a black hole and a Sybil attack. Based on simulation studies for selected scenarios, we showed that a location-based intrusion response may perform better than an address-based approach. Yet, in general, this trade-off depends on various factors such as the particular configuration of the mobile ad hoc network, the attack mechanisms, and the intrusion detection system as well as the intrusion response system. Since this parameter space is too large to be scrutinized completely by means of simulation (or in testbeds), we presented a generalized analytical model of the attack mechanism and the intrusion response system in [4]. We followed an elementary approach based on a combined geometric and stochastic description of the routing protocol and the location-based intrusion response system.

In this section we recapitulate the challenges we experienced in harmonizing model predictions and simulation results. For this, we focus on modeling and evaluating the packet loss that is caused by misbehaving nodes. To set our results in the context of related work, we summarize existing work on attacks and security mechanisms for mobile ad hoc networks focusing on the evaluation methodologies applied in the next section. We then briefly describe the relevant details of the mobile ad hoc network and the location-based intrusion response system that are subject of our evaluation in Section 2.2.1. As the main part, we present the entire chain of model development (Section 2.2.2), simulation studies (Section 2.2.3), and refining the model based on simulation results (Section 2.2.4). To conclude this part on challenges in evaluating security mechanisms for mobile ad hoc networks, we present the lessons learned during our work in Section 2.3.

2.1.1. Related Work

In this section we present an overview on research that motivated our work on security mechanisms for mobile ad hoc networks. We focus on attacks, intrusion detection and response, location-based routing approaches and highlight the particular evaluation methodology applied.

Attacks on Mobile Ad Hoc Networks. A mobile ad hoc network operates based on the cooperation of nodes that are usually not within the administrative domain of a service provider, but controlled directly by the end user. This makes it easy for an attacker to tamper with the behavior of a node in order to launch an attack on the network. Attack vectors for mobile ad hoc networks exist on each layer of the ISO-OSI communication model. A comprehensive review can be found e.g. in [5] and [6].

An attack that is often cited in the context of mobile ad hoc networks is the black hole attack [5]. In our work, we use the black hole attack to exemplify misbehavior due to the severe and easily quantifiable effects on network availability. The attack can be classified as an active attack on the network/routing layer. Comparable to the effect of a black hole in terms of astronomy, a black hole in a mobile ad hoc network attracts and 'absorbs' network traffic. The detailed mode of operation of the attack depends on the particular routing protocol deployed. The negative effects of black holes on the availability of a mobile ad hoc network were studied e.g. in [16]. Regarding the evaluation methodology, the work is closely related to ours since an analytical model of misbehavior was developed and compared with results from simulation studies. However, [16] focuses on the route-length distribution as a metric only and does neither include modeling of packet loss nor of security mechanisms.

Intrusion Detection and Response in Mobile Ad Hoc Networks. An intrusion detection system in combination with an intrusion response system as a reactive security measure can be deployed as a second line of defense for the case of subverted preventive security measures. One of the first approaches for intrusion detection and response in mobile ad hoc networks was presented in [17]. Detection of misbehavior is based on a packet loss metric. For this, an intrusion detection component (Watchdog) on each node monitors the forwarding behavior of its neighbors. If a certain threshold of dropped packets is exceeded, Watchdog sends a notification to the source of the packets. Intrusion response is performed by the Pathrater component, which collects information from the Watchdogs, calculates ratings for each node, and aggregates these node ratings to a rating for the route to the intended receiver. If multiple routes to a receiver exist, Pathrater chooses the one with the best rating for transmission. By doing so, Pathrater prevents the inclusion of detected malicious nodes in a route. The evaluation of the Watchdog/Pathrater approach presented in [17] is based on simulation studies. An analytical description of the approach was not part of the paper. Experiments in testbeds were not conducted.

Subsequent approaches for intrusion detection were developed with the goal of overcoming limitations of the Watchdog/Pathrater approach, taking into account constraints of devices and harnessing the characteristics of mobile ad hoc networks.

The intrusion detection system proposed in [18] uses a clustering approach

to reduce the system load caused by intrusion detection tasks. For this, one node out of a set of neighboring nodes (a cluster) is selected as the clusterhead. In each cluster, only the clusterhead is responsible for performing tasks for intrusion detection. For reasons of load balancing, the role of being clusterhead is switched periodically between nodes belonging to the same cluster. While the metrics used for intrusion detection are described analytically, the evaluation of the intrusion detection system is based on simulation studies. An analytical evaluation or a testbed-based validation were not performed.

In [19], a cooperative intrusion detection system was presented. Here, an exchange of intrusion detection information among nodes is used to enhance detection performance. With this, it is possible to lower false positive (a detection when there is no attack) and false negative (no detection when there is an attack) rates. The evaluation was performed by means of simulation. An analytical description of the system or testbed experiments were not presented.

Comprehensive approaches that take into account intrusion detection as well as intrusion response are e.g. CONFIDANT [13], CORE [20], and OCEAN [15]. For each of these, intrusion detection is performed by a component whose core idea is comparable to the Watchdog as proposed in [17]. The information that is gathered by the Watchdog component is collected and managed in a reputation system that is either based on locally available information only (OCEAN) or global information that is exchanged between nodes (CONFIDANT, CORE). The assignment of reputation values to nodes is based on the addresses of nodes as identifiers in all three approaches. Consequently, intrusion response is the exclusion of misbehaving nodes with a bad reputation from the network based on their addresses. In general, no data is sent to/via or received from/via addresses that are assigned a bad reputation. Regarding the evaluation methodology, CONFIDANT and OCEAN were evaluated in simulation studies (a testbed-based analysis of attacks on mobile ad hoc networks as basis for attack detection/reputation determination is presented in [13] in the context of CONFIDANT, yet an analysis CONFIDANT itself in the testbed was not part of [13]). CORE was evaluated both by means of simulation and analytically based on game theory. Yet, the goals of the evaluation were different for the simulation and the analytical part. A mutual validation of evaluation results obtained analytically and in simulation studies was not presented.

Recently, game theory became more and more popular in the context of not only evaluating but also designing new mechanisms for mobile ad hoc

networks. From an abstract point of view, game theory is a tool for modeling give-and-take interactions of two parties. In [21], the authors propose game theoretic views on issues such as transmission power control, medium access strategies, and packet forwarding. A model based on game theory for misbehavior in mobile ad hoc networks is the focus of [22]. Closely related to this is the work in [23]. Here, a reputation mechanism used to decide on whether to interact with a node or not (i.e. relaying its packets or using it as a relay) is developed based on game theory.

Location-based Routing in Mobile Ad Hoc Networks. Several routing mechanisms for Mobile Ad Hoc Networks that harness knowledge on geographical location of nodes were proposed. An overview can be found e.g. in [24]. Two of these protocols which are related to our work are LAR [25] and DREAM [26]. Both protocols use location information to restrict the propagation of broadcast messages in order to reduce the routing overhead generated.

For LAR, two schemes to improve the overhead of the route discovery phase of on-demand routing protocols were proposed. For the first scheme, a Request Zone is specified such that it contains the initiator of a route request and an area in which the destination is expected to be located. Route request messages are only forwarded by nodes that are located in the Request Zone. The second scheme is based on the distance between the destination and nodes that could potentially forward a route request. Here, a node only forwards a route request if it is closer (within certain bounds) to the destination than the node it received the request from. LAR was evaluated by means of simulation. Analytical studies or testbed-based evaluations were not presented.

DREAM was developed as a robust protocol in which each message (not only route requests) is sent as a restricted broadcast. To restrict the broadcast, DREAM determines the direction from sender to receiver based on their geographical location. Only nodes that are (within certain bounds) located in this direction forward messages. At this point, it is necessary to mention that the main objective of DREAM was the efficient dissemination of the location information of nodes throughout a network. However, a detailed description of this process is beyond the scope of this article. Regarding evaluation methodology, the functionality of DREAM was verified in simulation studies. Analytical models or experiments in testbeds were not presented in [26].

2.2. Evaluation

We now present relevant system details and the evaluation of the location-based intrusion response system for mobile ad hoc networks. The functionality of the location-based intrusion response system is based on excluding benign nodes from the network in order to maintain overall network availability. Further, quarantine zones established by the location-based intrusion response system obviously have to be (at least) of the size of the transmission range of nodes. Thus, the approach is applicable for large-scale mobile ad hoc networks only and an evaluation in a testbed that is beyond a basic proof of concept is hardly possible. We therefore base our evaluation on a combination of an analytical model and simulation studies to cover the parameter space of the system as far as possible.

2.2.1. Specification: System Properties

In this section we shortly describe the mobile ad hoc network routing protocol, our implementation of the black hole attack, the intrusion detection system, and the intrusion response system. The model we develop is valid for an abstract reactive routing protocol that selects routes based on age and a distance metric and uses an expanding ring search in order to reduce network load. We exemplify this routing behavior using the AODV protocol [27], since AODV will be used in our simulation studies for model verification.

The Routing Protocol. Reactive routing protocols establish a route between a source and a destination when it is required for data transmission. To establish a route, AODV starts a route discovery process initiated by the source by sending a route request (RREQ) message for the intended destination. The RREQ is disseminated as a broadcast. Each intermediate node that forwards an RREQ remembers the preceding node from which the RREQ was received, thus establishing a reverse route. Upon receiving an RREQ at the destination node, a route reply (RREP) message is sent along the reverse route back to the source. Upon receiving the RREP at the source, the route is established. In case of receiving multiple RREPs for the same destination, AODV chooses the most recent (determined by a sequence number) and shortest (determined by a hop count) route, whereas newer routes are preferred over shorter ones.

Since disseminating RREQs as broadcast messages causes a high network load, AODV may use an expanding ring search to mitigate this effect. In this process, route discovery is first restricted to the neighborhood of the source.

For this, the time to live (TTL) of RREQ messages is set appropriately. [27] proposes to search routes consecutively in the 1-, 2-, 3-, 5-, and 7-hop neighborhood of the source. If a route is not found in these steps, the search is extended to the full network diameter which is assumed to be 35 hops in [27].

The Black Hole Attack. Comparable to the effect of a black hole in terms of astronomy, a black hole in a mobile ad hoc network attracts and 'absorbs' network traffic. For attacking AODV, this is done by issuing RREP messages with falsified age and/or distance metrics. To obtain a worst-case behavior, we consider a black hole that eavesdrops sequence numbers and answers every RREQ received with the current sequence number of the intended destination incremented by one and a hop count set to one (regardless of whether the intended destination really is a neighbor of the black hole). With this, the route that is offered by the black hole appears to be both newer and shorter than the route that is offered by the real destination. Thus, the route offered by the black hole is selected by the source for data transmission according to the specification of AODV. To complete the 'absorbing' effect, the black hole does not forward any data packets after the route is established.

The Intrusion Detection System. Since our goal was not to develop a new intrusion detection system for mobile ad hoc networks, we implemented a lightweight intrusion detection system that is tailored to the black hole behavior described above while providing a realistic detection performance that is comparable to that of other intrusion detection systems proposed in literature. The intrusion detection system works based on local (per node) information only, similar to OCEAN [15]. Attack detection is performed in two steps. During a monitoring interval t_{mon}, a node X keeps track of the forwarding behavior of its neighbors (we call node Y a neighbor of node X if it is located within the transmission range of X). For each of its neighbors, X maintains a counter n_{rec} for packets that Y received for forwarding. A second counter n_{forw} is maintained for packets that Y forwarded correctly. After each monitoring interval, X calculates a rating R_Y describing the forwarding behavior of each of its neighbors. R_Y is defined as

$$R_Y \quad = \quad R'_Y \cdot \frac{n_{rec}}{w_{bal}} - n_{forw}$$

In this definition, R'_Y denotes the rating of the previous monitoring interval and $w_{bal} \geq 1$ is a weighting factor to balance n_{rec} and n_{forw}. If R_Y

exceeds a threshold $thres_{black}$, X classifies Y as a black hole.

The Location-based Intrusion Response System. If a node Y is classified as a black hole by the intrusion detection system of a node X, X establishes a quarantine zone with radius r_{quar} around Y. As long as X is located within this quarantine zone, it will not forward any messages. Further, all active routes which X is part of become invalid. Since X will not forward any messages while it is located in a quarantine zone, subsequent route request messages will not reach the black hole Y, as shown in Figure 2. Thus, we prevent Y from being part of newly established routes while it is quarantined.

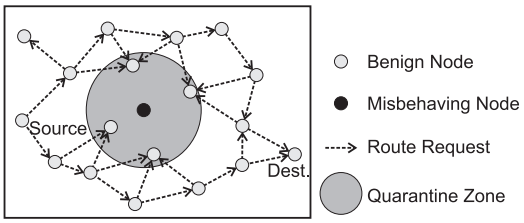

Figure 2: Mode of operation of the location-based intrusion response

We assume that the position of a node is not observable within a quarantine zone. Therefore, updating the quarantine zone if Y moves is not possible. For this reason, a revocation of quarantine zones is performed after the interval t_{reset}.

2.2.2. Expectation: Analytical Model

Within this section we develop the analytical model describing the effects of the black hole attack on the packet loss experienced in the mobile ad hoc network. After defining basic assumptions we start with modeling the expanding ring search behavior of the routing protocol which leads to the probability for a black hole being included in a route. Based on this, we describe the packet loss that can be charged to black holes in a mobile ad hoc network without intrusion detection and response. We continue with modeling the packet loss caused by black holes with activated intrusion detection and intrusion response subject to the time the intrusion detection system requires to detect a black hole and the mobility of nodes. The notations used for building the model are listed in Table 2.2.2.

Assumptions. In the following, we assume

Notations of formulae

t_{mon}	Monitoring interval of the intrusion detection system in seconds
n_{rec}	Number of packets a node received during t_{mon}
n_{forw}	Number of packets a node forwarded during t_{mon}
w_{bal}	Factor to balance n_{rec} and n_{forw}
R_Y	Rating for node Y
$thres_{black}$	Threshold of R_Y for classification as black hole
r_{quar}	Radius of quarantine zones
t_{detect}	Time needed to detect a black hole
t_{reset}	Time after which quarantine zones are revoked
r_{trans}	Transmission range of nodes
A_{net}	Size of the network
l	Side length of the network area
d_{hop}	Average distance per hop
n_{total}	Total number of nodes in the network
n_{black}	Number of black hole nodes in the network
n_s	Number of nodes reached in step s of the ring search
ρ	Network density in nodes per area

- a square network area A_{net} with side length l,
- a circular transmission range with radius r_{trans},
- quarantine zones to match the transmission range of the corresponding misbehaving node at the time the quarantine zone is established,
- a geometrically uniform distribution of benign nodes and black holes within the network area,
- randomly selected sources and destinations of traffic,
- a random distribution of traffic patterns among all nodes, i.e., the network load is constant and nodes can not be distinguished by their communication,
- a connected network, i.e., a route between any two nodes can be established at any time.

Expanding Ring Search. For modeling the packet loss we need to describe the expanding ring search in terms of the number of new nodes an RREQ message reaches in each step as shown in Figure 3. For this, let A_h be the circular area covering all nodes located in a distance of at most h hops to the source of the RREQ, where $A_0 = 0$ and $\Delta A_h = A_h - A_{h-1}$. Depending on the actual network configuration, a constant average geometric per-hop routing progress $0 < d_{hop} < r_{trans}$ can be assumed. For the remainder of this paper we use $d_{hop} = \frac{r_{trans}}{\sqrt{2}}$ which was obtained experimentally in [28] for a realistic mobile ad hoc network configuration. This leads to a first estimate of ΔA_h defined as

$$\begin{aligned} \Delta\Lambda_h &= \Lambda_h - \Lambda_{h-1} = \pi(h \cdot d_{hop})^2 - \pi((h-1)d_{hop})^2 \\ &= (2h-1)\pi \cdot d_{hop}^2 \end{aligned}$$

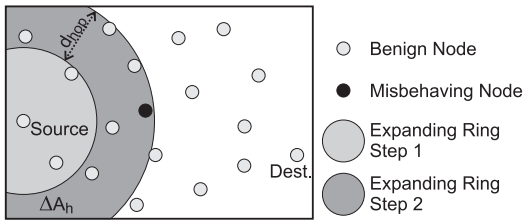

Figure 3: Area covered by expanding ring search subject to d_{hop}

This description holds if the ring search does not leave the simulation area. To include edge effects for nodes being located at the borders of the simulation area we use a four-step approximation. For a ring radius $r = h \cdot d_{hop} < \frac{l}{2}$ of less than half the side length l of the simulation area, the ring can be contained fully in the simulation area in the best case for a node being located at the center. In the worst case, for a node being located at a corner, only one quarter of the ring reaches into the simulation area. Since nodes are assumed to be distributed uniformly, both cases and any in between are equally likely. Thus, for $r = h \cdot d_{hop} < \frac{l}{2}$ we obtain $\Delta A_h = \frac{1}{2}\left(\Lambda_h + \frac{1}{4}\Lambda_h\right)$. In the same way, we approximate the cases for a ring radius between half side length and side length, between side length and diagonal length, and larger than diagonal length. Altogether, we obtain

$$\Delta A_h = \begin{cases} \frac{1}{2}\left(\Lambda_h + \frac{1}{4}\Lambda_h\right) & \text{if } h \cdot d_{hop} \leq \frac{l}{2} \\ \frac{1}{8}\Lambda_h & \text{if } \frac{l}{2} < h \cdot d_{hop} \leq l \\ \frac{1}{16}\Lambda_h & \text{if } l < h \cdot d_{hop} \leq \sqrt{2l^2} \\ 0 & \text{otherwise} \end{cases}$$

The expanding ring search is performed in consecutive steps in each of which the time to live of an RREQ is incremented. Let h_s be the time to live in hops for step s. With the default parameters specified in [27], h_s is given by the following table (we set $h_0 = 0$ for reasons of simplification).

s	0	1	2	3	4	5
h_s	0	1	3	5	7	35

For the description of the number n_s of nodes that are reached in step s of the expanding ring search let n_{total} be the total number of nodes and A_{net} be the total area of the network. Because nodes are assumed to be distributed uniformly, the geometric density of nodes is given by $\rho = \frac{n_{total}}{A_{net}}$. We get

$$\begin{aligned} n_0 &= 0 \\ n_s &= n_{s-1} + \Delta n_s \text{ where } \Delta n_s = \sum_{i=h_{s-1}+1}^{h_s} \Delta A_i \cdot \rho \end{aligned}$$

Packet Loss due to Black Holes in a Defenseless Mobile Ad Hoc Network. To model the packet loss caused by black holes, we start with describing the probability $p_{black}(s)$ for an RREQ reaching at least one black hole in steps 1 to s of the expanding ring search. If we assume n_{black} black holes in n_{total} total nodes, we have $n_{total} - n_{black}$ benign nodes in the network. Thus, the number of possible combinations to select the n_s nodes that are reached exactly in step s of the ring search only from benign nodes is $x = \binom{n_{total}-n_{black}}{n_s}$. The number of combinations for selecting the n_s nodes from n_{total} total nodes is $y = \binom{n_{total}}{n_s}$. Now, $\frac{x}{y}$ describes the probability that an RREQ reaches only benign nodes in steps 1 to s of the ring search. For the probability $p_{black}(s)$ that a RREQ reaches at least one black hole in steps 1 to s of the ring search, we have to consider the special case of $n_s > n_{total} - n_{black}$, i.e., the number of nodes reached in step s exceeds the number of benign nodes. Certainly, the RREQ reaches a black hole in this case. Altogether, we get

$$p_{black}(s) = \begin{cases} 1 & \text{if } n_s > n_{total} - n_{black} \\ 1 - \frac{\binom{n_{total}-n_{black}}{n_s}}{\binom{n_{total}}{n_s}} & \text{otherwise} \end{cases}$$

We assume a random selection of traffic patterns as well as of source and destination nodes. Hence, the probability for a data packet to get lost

due to a black hole equals the probability for a black hole being part of the corresponding route between source and destination. To determine this, we need to calculate the probability $p_{dest}(s)$ for reaching the destination node in step s of the ring search. Since destinations are assumed to be chosen randomly, this probability correlates to Δn_s. We obtain $p_{dest}(s) = \frac{\Delta n_s}{n_{total}}$. With this, the probability for the RREQ reaching the destination in step s and at least one black hole in steps 1 to s of the ring search is given by $p_{dest}(s) \cdot p_{black}(s)$. For the overall probability $p_{loss,black}$ for data packets to get lost due to a black hole, we need to consider all steps of the expanding ring search. We get

$$p_{loss,black} = \sum_{i=1}^{5} p_{dest}(i) \cdot p_{black}(i)$$

The upper bound of the sum depends on the maximum number of steps performed for the ring search. For our model, we assume a maximum of 5 steps as specified in [27].

Packet Loss due to Black Holes with Intrusion Detection and Intrusion Response. Although the location-based intrusion response system excludes black hole nodes from the network, the influence of black holes can only be mitigated and not be thwarted completely. Thus, the probability $p_{loss,defended}$ for packet loss despite active security measures can be described as

$$p_{loss,defended} = p_{loss,black} \cdot p_{IRSfail}$$

Here, $p_{IRSfail}$ denotes the probability that the intrusion response system fails to prevent a black hole from dropping packets. As the main influences on $p_{IRSfail}$ we identified the detection time of the intrusion detection system and mobility of nodes. Both lead to independent parts $p_{IRSfail,detect}$ and $p_{IRSfail,move}$ of $p_{IRSfail}$. Thus, we obtain

$$p_{IRSfail} = p_{IRSfail,detect} + p_{IRSfail,move}$$

To detect ongoing misbehavior, the intrusion detection system has to monitor the suspicious node for a certain time. In our scenario, the black hole may continue dropping packets during this time. To describe the resulting probability $p_{IRSfail,detect}$ (as a part of $p_{IRSfail}$) for packet loss during the time the intrusion detection system needs to identify misbehavior, we start

by modeling the detection time. For our intrusion detection system, the detection time is given as $t_{detect} = n_{mon} \cdot t_{mon}$ where n_{mon} denotes the number of monitoring intervals needed until a black hole is detected and t_{mon} denotes the duration of a monitoring interval.

As described, a node Y is classified as a black hole, if the rating r_Y exceeds the threshold $thres_{black}$. Since a black hole node in our case does not forward any packets, R_Y is defined as $R_Y = \frac{n_{rec}}{w_{bal}}$ per monitoring interval t_{mon}. If we consider a steadily loaded network, the traffic during the detection time is (nearly) constant at a rate λ. Thus, $n_{rec} = \lambda \cdot t_{mon}$. If a black hole is detected, $n_{mon} \cdot \frac{\lambda \cdot t_{mon}}{w_{bal}} > thres_{black}$ holds. The number n_{mon} of monitoring intervals needed to detect a black hole is then given by $n_{mon} = \frac{thres_{black} \cdot w_{bal}}{\lambda \cdot t_{mon}}$. Thus, for the detection time of the intrusion detection system, we obtain

$$t_{detect} = \frac{thres_{black} \cdot w_{bal}}{\lambda}$$

After a black hole is detected, the corresponding quarantine zone excludes the black hole from the network for the time t_{reset}. Altogether, a black hole can be active for the time t_{detect} as part of the total detection-protection-period $t_{detect} + t_{reset}$. Thus, for the probability $p_{IRSfail,detect}$ of losing packets due to n_{black} black holes during the detection time of the intrusion detection system, we obtain

$$p_{IRSfail,detect} = n_{black} \cdot \frac{t_{detect}}{t_{detect} + t_{reset}}$$

We assume that quarantine zones can not be adapted directly when a node moves since tracking quarantined nodes is not possible. Thus, as shown in Figure 4, mobility of a black hole leads to a newly affected area A_{affect}. Nodes in this area are not aware of the black hole and will forward RREQ messages without restrictions.

The probability $p_{IRSfail,move}$ for packet loss due to node mobility can be modeled based on the number n_{affect} of nodes located in A_{affect}. Each of these nodes has to perform a detection of the black hole which leads to a corresponding multiple of $p_{IRSfail,detect}$ as described in the previous section.

The angle α, as shown in Figure 4, can be determined by $\alpha = 2 \cdot \arccos\left(\frac{d}{2r_{trans}}\right)$ where d denotes the distance between the position where the black hole was first detected and its new location. With this, we can determine the area

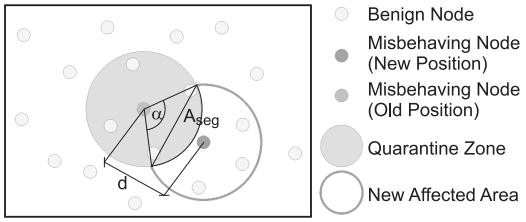

Figure 4: Influence of black hole mobility on the intrusion response system

A_{seg} of the circular segment defined by the quarantine zone and A_{affect}. We get $A_{seg} = \frac{r_{trans}^2}{2} \cdot (\alpha - \sin(\alpha))$, thus

$$A_{affect} \quad = \quad \pi r_{trans}^2 - 2 \cdot A_{seg}$$

Since nodes are distributed uniformly within the network area, we get $n_{affect} = A_{affect} \cdot \rho$. Thus, the loss due to node mobility can be described as

$$p_{IRSfail,move} \quad = \quad n_{affect} \cdot n_{black} \cdot p_{IRSfail,detect}$$

2.2.3. Observation: Simulation Studies

To validate the model, we compare the model predictions with (parts of the) results of the simulation study we presented in [3]. To visually demonstrate the accuracy achieved, we show the model predictions as curves together with the 95% confidence bars taken from the simulation results. The simulation results were obtained from a scenario consisting of $n_{total} = 1000$ nodes with $n_{black} \in \{1, 2, 3, 5, 10\}$ black holes on a square simulation area A_{net} with $l = 4750m$ side length. The size of the simulation area was chosen such that a connected network is typically achieved, i.e., a route between arbitrary nodes exists with a high probability. The nodes moved continuously according to a random waypoint mobility model (which can be considered as a worst-case scenario regarding predictability) with a speed between 1 and 2 meters per second. This leads to an average distance $d = \frac{t_{reset} \cdot \text{meters}}{6 \cdot \text{second}}$ from the location at which a black hole was detected to its new location when the quarantine zone is revoked. We used a constant bitrate traffic pattern with a network load of 20 streams in parallel consisting of packets with a size of 512 byte at a rate of 2048 kbytes per second resulting in a moderately loaded network. With a detection threshold $thres_{black} = 10$ determined as an optimum in preliminary simulations for the scenario given, the resulting detection time of the intrusion detection system is $t_{detect} = 1s$. The reset intervals for quarantine zones were taken from $t_{reset} \in \{15, 30, 45, 60, 90, 120, 180, 300, 420, 600\}$

in seconds. Each scenario was simulated for one hour simulated time split up in 6 parts of 10 minutes each to reduce unwanted side effects of the random waypoint mobility model. The simulations were performed with a modified version of the JiST/SWANS simulation tool for mobile ad hoc networks [29] running on a CONDOR distributed computing cluster [30].

The comparison of model prediction and simulation results for the loss caused by black holes in a defenseless network is shown in Figure 5. For the 2, 3, and 5 black hole scenarios, the model prediction is within the confidence interval of the simulation results. The small deviation for the setups with 1 and 10 black holes can be explained by the inaccuracy of the heuristic we used to model edge effects of the expanding ring search.

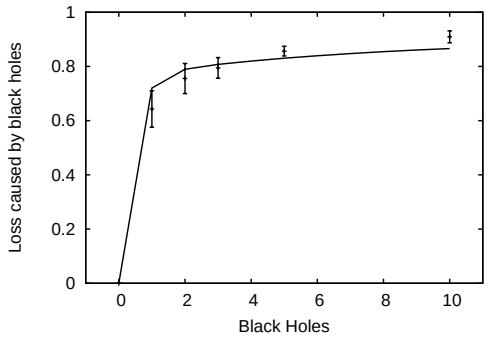

Figure 5: Model prediction (curves) compared with simulation results (confidence bars) of packet loss caused by black holes without intrusion response

Figure 6 shows the comparison of model predictions and simulation results for the loss caused by black holes in a network with intrusion detection and intrusion response. Please note that for reasons of readability only the results for the 3, 5, and 10 black hole scenarios are presented. While prediction and confidence intervals (i.e. simulation results) match well for the 1, 2, and 3 black hole scenarios, it stands out that for the 5 and 10 black hole setups the prediction and the simulation results differ strongly (note that the curve matching the 10 black hole confidence bars belongs to the 5 black hole prediction).

2.2.4. Adaptation: Matching Model Predictions and Simulation Results

The predictions made by the analytical model and the simulation results match reasonably for the effect of black holes on a mobile ad hoc network without security mechanisms. Small deviations can be explained by the

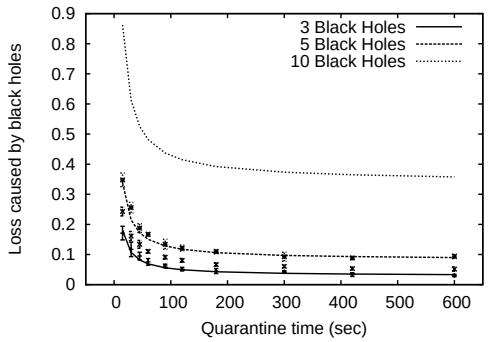

Figure 6: Model prediction (curves) compared with simulation results (confidence bars) of packet loss caused by black holes with location-based intrusion response

inaccuracy of the model due to the heuristic used to model edge effects. We therefore assume that the model describes the effect of undefended black holes within proper bounds.

Further, model predictions and simulation results match for the 1, 2, and 3 black hole setups in a mobile ad hoc network with intrusion detection and location-based intrusion response. Only predictions made for scenarios with 5 and 10 black holes do not correspond to the results obtained by simulation. Thus, our assumption was that the model itself described the behavior of the intrusion detection system and the location-based intrusion response system reasonably but might be instantiated incorrectly. Rerunning and tracing the simulations, we found that in the scenario with 5 black holes only 4, and in the scenario with 10 black holes only 5 are active at a time. We attribute this to overlapping-effects of quarantine zones. If we instantiate the model accordingly (4 / 5 black holes in the model for the 5 / 10 black hole simulation setups), we obtain a match for all scenarios as shown in Figure 7. We, thus, need to extend the model to cover the probability of a black hole being active in scenarios with multiple black holes and active location-based intrusion response.

2.3. Lessons Learned

In this first part of the article, we presented an analytical model that describes the effects of black hole attacks and location-based countermeasures on mobile ad hoc networks. Based on a combined geometric and stochastic approach, we developed a light-weight model for the routing process as well as for the packet loss caused by misbehavior.

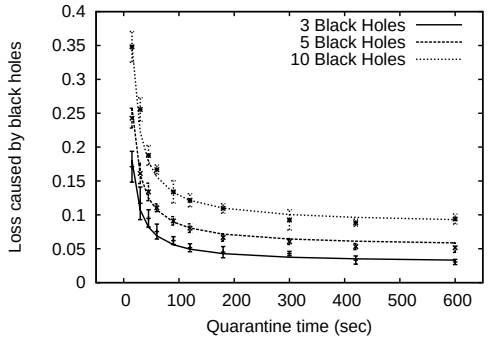

Figure 7: Corrected model prediction (curves) compared with simulation results (confidence bars) of packet loss caused by black holes with location-based intrusion response

To evaluate effects of the black hole attack and the location-based intrusion response, we needed both an analytical model and simulation studies. The analytical model was required to cover the parameter space of the system as far as possible. To validate (and correct) the model predictions, we had to perform simulation studies for selected parameters. We found that, despite following an elementary approach, on the first attempt we did not consider all factors that affect the interaction of the black hole attack and the location-based intrusion response system. To achieve an appropriate evaluation of the location-based intrusion response system for mobile ad hoc networks, it was necessary to combine both the analytical and the simulation perspective and to have a mutual feedback between both. We only presented the feedback from simulation to model in this section. Yet, a feedback in the opposite direction was used during implementation of the attacks and the intrusion detection system as well as the location-based intrusion response. This way, implementation errors leading to strong deviations of model predictions and simulation results could be corrected beforehand.

3. Challenges in Evaluating Security Mechanisms for Peer-to-Peer Systems

In this second part of the article, we present challenges we faced in evaluating cooperative security mechanisms for peer-to-peer systems. As evaluation methodology, we used an analytical model and testbed experiments having a mutual feedback between model and testbed. The following is based on prior work presented in [31] extended by recent results of testbed experiments.

3.1. Security in Peer-to-Peer Systems

Peer-to-peer systems enable enhanced communication services in environments where client/server-based solutions can not be established due to e.g. time and/or cost constraints. As for mobile ad hoc networks, large scale emergency response scenarios are a prominent application domain. Related projects that apply peer-to-peer technology in large scale emergency response scenarios, but do not consider cooperative security mechanisms, are e.g. DUMBO [32] and SoKNOS [33].

In an emergency response scenario, the exchange of information and services between aid organizations can offer a considerable benefit for the coordination of search and rescue or reconstruction efforts. However, information and services should not be accessible in an unrestricted way. By means of authentication and access control, the admission to services and information can be governed. However, contemporary means for authentication and access control such as Kerberos [34] are based on central trusted instances. Thus, these mechanisms can not be transferred directly from the client/server domain to a peer-to-peer environment. In the absence of central trusted instances, security objectives such as authentication and access control can be implemented by threshold cryptography. In threshold cryptography, the cooperation of (at least) a certain number of peers is required to perform cryptographic operations. No single (possibly compromised) peer is able to e.g. sign and issue certificates that grant access to restricted services (which will be the example we use in this section). This way, a security level comparable to that of centralized solutions can be achieved.

The applicability and performance of threshold cryptography in peer-to-peer systems has been studied comprehensively. However, only little attention has been paid to the fact that (cooperative) security-relevant decisions require a well-defined set of regulations, if they have to be made automatically. In the application scenario outlined, any interaction desired can hardly be foreseen. Thus, the availability of predefined security policies can not be assumed. To deal with this, we consider the case of authorized users being involved directly in security-relevant decisions. In this case, performance issues of the network as well as of the threshold cryptography schemes deployed are negligible compared to the delay that is introduced by the users themselves. Rather, the number of users involved per decision has to be minimized to keep the approach feasible for real-world deployment. Nevertheless, the minimization has to take into account users that do not provide their decision in a reasonable amount of time. To deal with this, a certain redundancy

has to be introduced regarding the number of users requested with respect to the number of users that have to cooperate as specified by the threshold cryptography scheme deployed.

In the following, we recapitulate our work presented in [31], describing how security-relevant decisions can be performed if neither a central trusted instance nor predefined security policies are available. We shortly discuss variants of threshold cryptography that are able to deal with the resulting challenges and present related work on applying threshold cryptography in peer-to-peer systems focusing on the evaluation methodology applied. Having laid these basics, we describe possible interaction schemes between the peer that requests a decision and the peers that (potentially) contribute to the decision. We provide a stochastic analysis of the different interaction schemes that allows for the real-time minimization of the number of users involved in a joint decision. We then validate the model with selected results obtained recently in a series of testbed studies. We explain deviations observed between model predictions and testbed results and present a harmonized instantiation of the model where model predictions and testbed results match reasonably.

3.1.1. Related Work

In this section, we present related work that has motivated and influenced our research. We focus on basics of threshold cryptography and on studies on the performance of threshold cryptography in peer-to-peer environments highlighting the evaluation methodology applied.

The idea of threshold cryptography was initially presented in [35]. The mode of operation is based on shares of a secret key that are generated by choosing a polynomial $p(x)$ of degree n such that the shared secret key equals $p(0)$. The keyshares are calculated as $p(1), ..., p(m)$ where $m > n$. By Lagrange interpolation, the secret key can be reconstructed if $n + 1$ keyshares are combined. In [35], reconstruction is performed by a trusted device that collects the keyshares and, subsequently, may perform cryptographic operations. In peer-to-peer systems, the availability of a trusted device that is accessible for all peers is not given. In threshold cryptography approaches developed for such distributed and decentralized environments, the keyshares $p(1), ..., p(m)$ are distributed to the peers. Also the generation of the keyshares can be performed cooperatively and in a distributed way, without the need for a central instance. With the keyshares, the peers are able to produce partial signatures (e.g. for a certificate that grants access

to restricted resources within the peer-to-peer system). The computation of the full signature is again based on Lagrange interpolation.

Due to the mode of operation of Lagrange interpolation, most of the threshold signature schemes proposed require that the indexes of all keyshares with which the partial signatures shall be generated (i.e., the x of $p(x)$) must be known in advance, before the partial signatures can be generated. A partial signature is then valid only for the specified set of keyshares. In the context of peer-to-peer systems, this results in multiple rounds of communication required between the peers that contribute partial signatures. In peer-to-peer systems in general, peers may go offline at any time during the signature process. For our application scenario in particular, we also have to consider mobile devices that are connected wirelessly and, thus, may be subject to more frequent disconnections and also energy constraints. If one member of the set of peers providing partial signatures does not provide its partial signature in time and has to be replaced, all other partial signatures have to be discarded and the process has to be restarted from the beginning resulting in an increased time required to produce a signature and increased energy consumption. The threshold scheme deployed should therefore be able to deal with signers that do not provide a signature without having to discard partial signatures that have been provided already. To make the system as reliable as possible, we build upon a non-interactive threshold cryptography scheme that does not need any communication between the peers providing partial signatures. The scheme proposed in [36], as an enhancement of [37], meets this requirement.

In [38], the authors compare threshold signature schemes with respect to their performance in controlling access to closed user groups in peer-to-peer systems. User interactions have not been considered. Non-interactive signature schemes were not part of the evaluation. Regarding evaluation methodology, performance is measured in terms of basic operation costs (the time needed to produce partial signatures) and join time (the amount of time a new peer needs to join a closed user group). The evaluation is performed in a small-scale local testbed consisting of four machines running multiple virtual peers each. An analytical model used for evaluation was not presented. The work presented in [38] was extended to additional signature schemes in [39]. Here, the evaluation was based on a testbed consisting of ten nodes running multiple peers each. An analytical model was not presented.

A non-interactive mechanism for access control in mobile ad hoc networks has been proposed in [40] and [41]. In contrast to [37], the protocol

proposed in [40] and [41] is not based on a cryptographic key that is shared among multiple parties, but on bivariate polynomials that can be used to establish pairwise shared secret keys. User interactions have not been considered. Because 'standard' signed certificates as required in our scenario can not be produced with this approach, we have not taken [40] into account as a possible cryptographic mechanism for our application scenario. Regarding evaluation methodology, a performance evaluation subject to metrics describing the basic performance of the threshold cryptography scheme deployed in terms of computational costs and energy consumed was performed. For this, measurements in a testbed for a mobile ad hoc network consisting of five nodes were conducted. An analytical model for system evaluation was not presented.

3.2. Evaluation

We now present relevant system properties and the evaluation of user-based cooperative decisions for peer-to-peer systems. While common Internet-based peer-to-peer systems for e.g. voice over IP communication or filesharing usually consist of millions of peers, for our emergency response application scenario, we assume the network size to be in the order of hundreds to thousands of peers. Thus, a testbed evaluation of the user-based cooperative decision process with a network size as expected in a real-world system is possible. We therefore base the evaluation on the combination of an analytical model and testbed experiments to achieve a high degree of realism.

3.2.1. Specification: System Properties

A peer that requests a security-relevant decision has to send this request to a set of peers holding keyshares. Each of these may take part in the decision process by issuing a partially signed certificate. The strategy according to which the requests are disseminated within the peer-to-peer overlay directly affects the number of users requested and the probability of receiving enough partially signed certificates to be able to interpolate a full signature. In the following, we describe different interaction schemes between requesting peers and peers that contribute to the decision process. We take into account different levels of knowledge about which peers hold keyshares. We shortly explain how the interaction schemes can be realized in a Pastry peer-to-peer overlay [42] which will be the basis for the testbed evaluation. Yet, the interaction schemes as proposed abstractly in the following and the an-

alytical model describing the interaction schemes are independent from the particular implementation of the peer-to-peer overlay.

Interaction Scheme for Unknown Shareholders. If the peers holding keyshares are unknown to the peer that requests a decision, broadcasting the request within the entire peer-to-peer overlay would be a straightforward interaction scheme. In this case, all users holding keyshares are contacted. Thus, in case the decision is positive, a broadcast would result in the highest probability for a decision request to be successful (i.e., for receiving enough partially signed certificates to be able to interpolate a fully signed certificate). Yet, due to the high number of users involved that is caused by a broadcast in the peer-to-peer overlay, the applicability of a broadcast is limited in our scenario. Instead, a multicast approach is reasonable. We assume that the multicast is initiated by the requesting peer and that the requesting peer has no knowledge about which peers are holding keyshares. In this case, as shown in Figure 8, a multicast can be realized by sending requests to a set of peer-IDs that are selected randomly. Regarding the implementation on top of a Pastry peer-to-peer system, if an ID is not used, Pastry will route the request to the peer with the ID closest to the one selected.

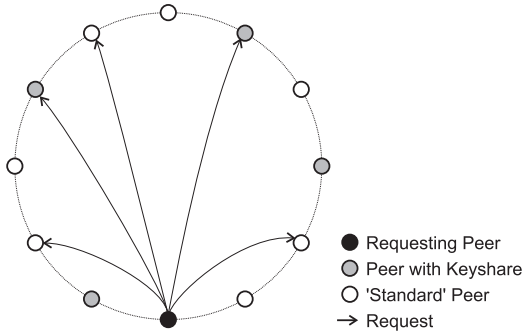

Figure 8: Schematic representation of the multicast with unknown signers

Interaction Scheme for Known Shareholders. In our application scenario, we assume administrative restrictions that limit the number of peers which are authorized to take part in security-relevant decisions. Thus, a random selection of the peers to which a request is sent may not reach enough peers holding keyshares. To increase reliability, it is reasonable to base the dissemination of requests on some knowledge about which peers hold keyshares

(and about the status of their users). This approach can be based e.g. on a peer that acts as mediator for the decision process. The mediator keeps track of the peers holding keyshares and accepts and relays requests appropriately, as shown in Figure 9.

While introducing a mediator can be done independently from the particular peer-to-peer overlay, a way to implement a multicast with knowledge about the distribution of keyshares on top of a Pastry peer-to-peer overlay is to make use of Scribe multicast groups [43]. Peers holding keyshares (and with users ready to contribute to a decision) subscribe to a corresponding multicast group. Requests can be sent to this group along with the requested number of partial signatures.

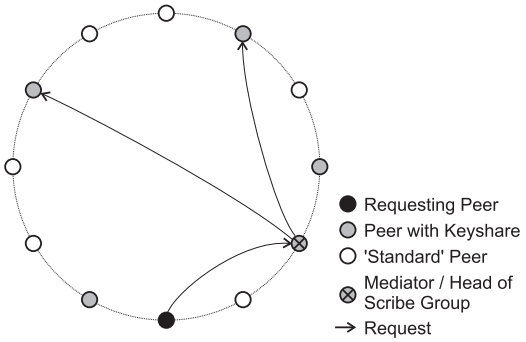

Figure 9: Schematic representation of the multicast with known signers

3.2.2. Expectation: Analytical Model

In this section we develop stochastic models that describe the success probability p_{succ} of the different interaction schemes. A request is considered successful if a sufficient number of partially signed certificates is received such that a fully signed certificate can be interpolated. The number of partially signed certificates received is sufficient if it is greater or equal to the threshold n_{thres} that is defined by the threshold cryptography scheme deployed. Table 1 provides an overview of the notation we use in the following.

Model of the Interaction Scheme for Unknown Shareholders. For the interaction scheme for unknown shareholders as outlined in Section 3.2.1, we assume a random limitation of requested peers with respect to the distribution of keyshares and to the status of peers. That is, the scheme does not consider whether a peer to which a request is sent holds a keyshare or

Table 1: Notations of formulae

n_{thres}	Number of partially signed certificates required for computing a full signature
p_{rep}	Probability with which a single peer answers a request
n_{total}	Total number of peers in the peer-to-peer overlay
n_{keys}	Number of peers holding keyshares
n_{mult}	Number of peers to which a request is sent
n_{rep}	Number of replies received for one request
$p(n_{rep})$	Probability for receiving n_{rep} replies from one request
p_{succ}	Probability for receiving a sufficient number (i.e., more than n_{thres}) of replies from one request

whether the user is currently able to answer within a reasonable time. This random restriction can be modeled by a hypergeometric random variable. In our case, the hypergeometric variable describes the intersection of the set of peers to which a request is sent and the set of peers that would potentially reply to a request received. Thus, for the probability $p(n_{rep})$ of receiving a certain number n_{rep} of replies, we get

$$p(n_{rep}) = \frac{\binom{n_{keys} \cdot p_{rep}}{n_{rep}} \binom{n_{total} - (n_{keys} \cdot p_{rep})}{n_{mult} - n_{rep}}}{\binom{n_{total}}{n_{mult}}}$$

For a request to be successful it is not important to receive a certain number of replies, but to receive *at least* enough partially signed certificates to enable interpolation of a fully signed certificate. That is, a request is successful if the number n_{rep} of replies received is greater than or equal to n_{thres}. The success probability $p_{succ}(n_{thres})$ with respect to n_{thres} thus can be described as the sum of the probabilities for receiving a certain number n_{rep} of replies starting from n_{thres}. The upper bound of the sum is given by the number n_{mult} of peers to which a request is sent. We obtain

$$p_{succ}(n_{thres}) = p(n_{rep} \geq n_{thres}) = \sum_{n_{rep}=n_{thres}}^{n_{mult}} p_(n_{rep})$$

Model of the Interaction Scheme for Known Shareholders. For the interaction scheme with knowledge about which peers hold keyshares as presented in Section 3.2.1, the probability $p(n_{rep})$ of receiving a certain number n_{rep} of

replies to a request can be modeled as a binomial random variable, resulting in

$$p(n_{rep}) = \binom{n_{mult}}{n_{rep}} p_{rep}^{n_{rep}} (1 - p_{rep})^{n_{mult} - n_{rep}}$$

The success probability $p_{succ}(n_{thres})$ with respect to n_{thres} again can be described as the sum of the probabilities of receiving a certain number n_{rep} of replies starting from n_{thres}. The upper bound of the sum is given by the number n_{mult} of peers to which a request is sent. We get

$$p_{succ}(n_{thres}) = p(n_{rep} \geq n_{thres}) = \sum_{n_{rep}=n_{thres}}^{n_{mult}} p(n_{rep})$$

3.2.3. Observation: Testbed Studies

We validate the models for the different interaction schemes by experiments in the PlanetLab [44] and G-Lab [2] testbeds. PlanetLab is a worldwide testbed currently consisting of about 1000 nodes distributed at about 500 sites worldwide that are connected over the Internet. Although minimum requirements regarding hardware and connectivity of nodes are specified, PlanetLab is a heterogeneous environment. This together with large variances in the computational load of nodes reflects the conditions of Internet-based peer-to-peer systems. G-Lab, on the other hand, is a German national testbed that currently consists of about 150 nodes at 6 universities. Compared to PlanetLab, hardware and connectivity is homogeneous and the overall load of nodes is moderate. G-Lab uses PlanetLab software for resource sharing and access control of its nodes. Thus, our experimental setups can be executed on PlanetLab and on G-Lab without modifications.

In the following, we present selected results of our testbed studies that show all particularities we observed. As experimental setup (and, thus, as instantiation of the models), we used a peer-to-peer system consisting of $n_{total} = 100$ peers of which $n_{keys} = 50$ hold keyshares. The threshold (i.e., the number of partially signed certificates required for interpolation of a fully signed certificate) is set to $n_{thres} = 10$; the probability with which a single peer provides an answer to a request received in a reasonable time is defined as $p_{rep} = 0.5$ (random numbers following a shifted standard normal distribution are drawn to determine whether a reply should be provided). The desired success probability is set to $p_{succ} = 0.95$. To achieve this, the model demands to send $n_{mult} = 49$ requests using the interaction scheme for unknown shareholders and $n_{mult} = 26$ requests using the interaction scheme for

known shareholders. In order to validate the model not only at this singular point, but for any success probability $0 \leq p_{succ} \leq 1$, we varied the threshold in $1 \leq n_{thres} \leq 30$ for the results presented in the following. All plots are given with 95% confidence intervals obtained from at least 40 request-reply iterations.

Experimental Results for the Interaction Scheme with Unknown Shareholders. Figure 10 shows the predictions of the model instantiated as described above as well as results from both the G-Lab and the PlanetLab testbed for the interaction scheme with unknown shareholders. Phenotypically, the model provides a reasonable prediction of the success probability p_{succ} subject to the threshold n_{thres}. Yet, the success probability observed in both testbeds is always significantly below the model prediction. Also the results of the two testbeds show a significant deviation. The success probability measured in PlanetLab is always below the results measured in G-Lab. The deviations can be explained by lost requests and replies as well as by peers holding keyshares that go offline during our experiments. The model also does not consider that sending requests to randomly selected peer-IDs might result in peers receiving two requests concurrently. In this case, our implementation drops these duplicate requests since two partially signed certificates generated with the same keyshare are of no use for the interpolation of a fully signed certificate. Due to the inherent heterogeneity in hardware, load, and connectivity of PlanetLab compared to the homogeneous nodes of G-Lab, the effects of packet loss and churn are more severe in PlanetLab than in G-Lab.

Experimental Results for the Interaction Scheme with Known Shareholders. The model predictions and the results obtained from both testbeds for the interaction scheme with known shareholders is shown in Figure 11. We present the results for an implementation based on a dedicated peer that keeps track of all peers that hold keyshares and acts as mediator during the decision process. The interaction scheme based on Scribe multicast groups is shown to be more susceptible to packet loss and churn.

Again, the model predictions of the success probability subject to the threshold match the success probability measured in the testbeds phenotypically. It stands out that the model predictions match the results measured in G-Lab reasonably, without significant deviations. This can be explained by the relatively low packet loss observed in G-Lab and the missing effect of

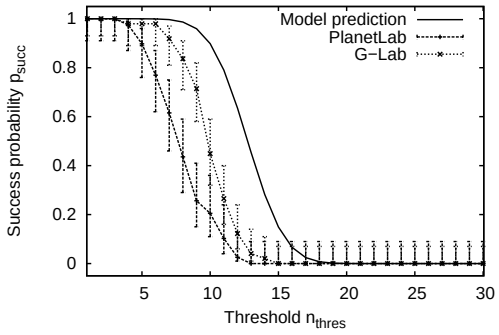

Figure 10: Success probability subject to threshold for the interaction scheme with unknown shareholders. Comparison of model prediction and results from PlanetLab and G-Lab.

duplicate requests for the interaction scheme with known shareholders, where the mediating peer relays requests appropriately. Due to the higher packet loss and churn of PlanetLab, the success probability measured in PlanetLab is again below the model prediction and the results measured in G-Lab.

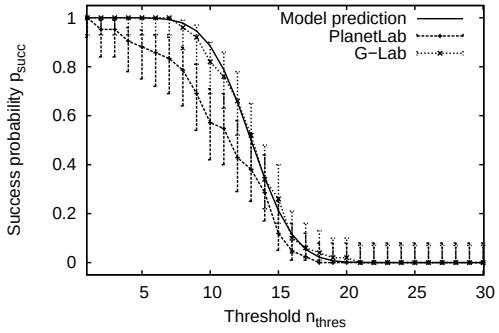

Figure 11: Success probability subject to threshold for the interaction scheme with known shareholders. Comparison of model prediction and results from PlanetLab and G-Lab.

3.2.4. Adaptation: Matching Model and Testbed

We now compare the model predictions with the results obtained from both testbeds if we adapt the instantiation of the model according to the conditions we observed in the testbeds. Please note that, if not specified otherwise, the instantiation remains as described above.

Adapted Model Parameters for the Interaction Scheme with Unknown Shareholders. During the experiments for the interaction scheme with unknown shareholders in the PlanetLab testbed, we measured an average total number of peers $n_{total} = 92$. The difference to the original instantiation can be explained by churn (i.e., by peers going offline) during our experiments. Due to the same reason, we observed $n_{keys} = 39$ online peers holding keyshares. Due to packet loss and duplicate requests, the probability with which a single peer provides a reply to a request received decreased to $p_{reply} = 0.39$. The number of requests sent decreased to $n_{mult} = 43$. If we instantiate the model accordingly, the predictions match the results measured in PlanetLab without significant deviations as shown in Figure 12

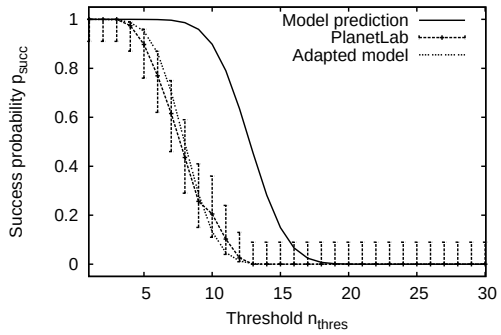

Figure 12: Success probability subject to threshold for the multicast with unknown shareholders. Comparison of model prediction, results from PlanetLab, and adapted model.

In G-Lab, we observed no effects of churn, but effects of packet loss and duplicate requests resulting in $p_{reply} = 0.43$ and $n_{mult} = 44$. If we instantiate the model with these parameters, we obtain a match of model predictions and success probability measured in G-Lab without significant deviations as shown in Figure 13.

Adapted Model Parameters for the Interaction Scheme with Known Shareholders. For the interaction scheme with known shareholders, we observed $n_{total} = 90$, $n_{keys} = 35$, $p_{reply} = 0.5$, and $n_{mult} = 21$. The better values of p_{reply} and n_{mult} compared to the original instantiation and the deviations observed during the experiments for the interaction scheme with unknown shareholders can be explained by the missing effect of duplicate requests for the interaction scheme with known shareholders. If we instantiate the model

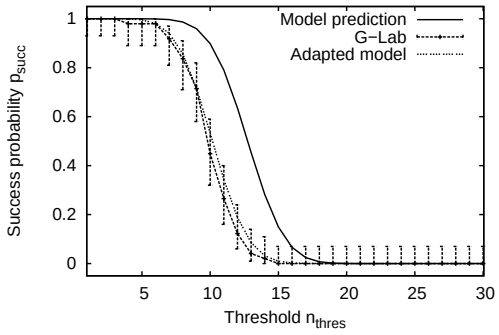

Figure 13: Success probability subject to threshold for the multicast with unknown shareholders. Comparison of model prediction, results from G-Lab, and adapted model.

accordingly, we obtain a better fitting model prediction. Yet, model prediction and simulation results still show differences. These remaining deviations of model predictions and simulation results can be explained by experiments during which the peer that acts as mediator in the decision process is affected by churn. If this peer, as a single point of failure, leaves the peer-to-peer system during the phase in which requests have to be relayed to peers holding keyshares, the decision process most probably is not successful.

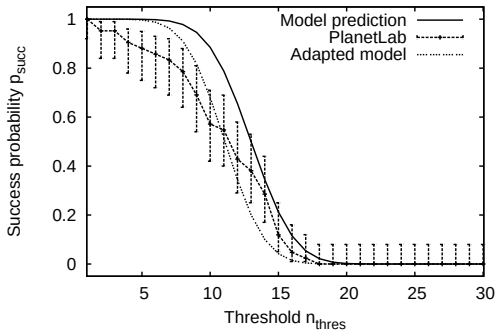

Figure 14: Success probability subject to threshold for the multicast with known shareholders. Comparison of model prediction, results from PlanetLab, and adapted model.

During the experiments conducted in G-Lab for the interaction scheme with known shareholders, we only observed a minor deviation of the probability with which a single peer provides a reply to a request received, resulting in $p_{succ} = 0.49$. The result of the adapted model instantiation is shown in

Figure 15. Since we only had the minor difference of p_{succ} with respect to the original instantiation, original model prediction, adapted prediction, and testbed results do not show significant deviations.

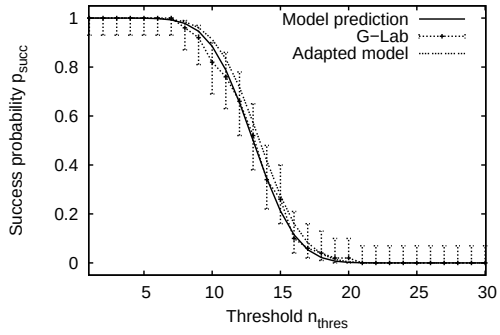

Figure 15: Success probability subject to threshold for the multicast with known shareholders. Comparison of model prediction, results from G-Lab, and adapted model.

3.3. Lessons Learned

In this second part of the article, we discussed user-based cooperative decisions in peer-to-peer systems as a means to counterbalance missing central trusted instances and predefined security policies in order to achieve basic security objectives such as authentication and access control. We described different interaction schemes for the cooperative decision process. For each interaction scheme, we outlined a stochastic model that describes the performance of the scheme.

To validate the model, we performed experiments in two testbeds which have different inherent conditions regarding heterogeneity of hardware and connectivity to prove the general validity of the analytical model. As for our work on security mechanisms for mobile ad hoc networks, we needed a two-dimensional view of the system, i.e., a mutual feedback between testbed experiments and analytical model to achieve an appropriate evaluation methodology. Although we only presented the feedback from testbed to model in this section, we used the feedback in the opposite direction to correct errors in the implementation beforehand.

4. Conclusions

In this article, we presented selected parts of our recent research on security mechanisms for decentralized systems and infrastructure-less systems. In particular, we discussed a location-based intrusion response mechanism for mobile ad hoc networks and interaction schemes for user-based cooperative decisions in peer-to-peer systems. We set the focus on the evaluation methodology we applied, showing results from the entire chain of analytical models, simulation studies, and testbed experiments. We especially highlighted the process of harmonizing the different perspectives, i.e., matching results obtained from the different evaluation methodologies by providing mutual feedback between them.

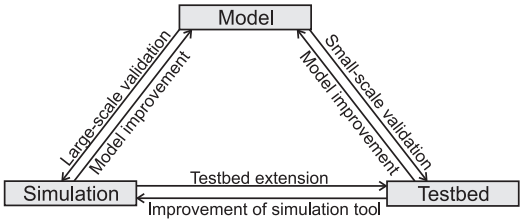

Figure 16: Mutual interaction of models, simulation tools, and testbeds for comprehensively evaluating protocols and mechanisms for communication networks

As a concluding remark, we want to emphasize that, from our point of view, it is necessary to combine perspectives of models, simulation tools, and testbeds during the evaluation of novel protocols and mechanisms for communication networks. The mutual feedback process, which is underlined also in related work as [45], is shown in Figure 16. This way, we think it is possible to obtain a reasonably comprehensive view on the system scrutinized, thus minimizing the 'gap' between science and reality.

References

[1] IT Transfer Office, Technische Universität Darmstadt, SicAri Project Homepage, http://www.ito.tu-darmstadt.de/projects/sicari/index_en_html.

[2] G-Lab Consortium, G-Lab Project Homepage, http://www.german-lab.de.

[3] A. König, M. Hollick, T. Krop, R. Steinmetz, GeoSec: quarantine zones for mobile ad hoc networks, Security and Communication Networks (Wiley SCN) Published Online. doi:10.1002/sec.68.

[4] A. König, D. Seither, M. Hollick, R. Steinmetz, An Analytical Model of Routing, Misbehavior, and Countermeasures in Mobile Ad Hoc Networks, in: Proceedings of the IEEE Global Communications Conference (GLOBECOM 2009) (Accepted for Presentation), Honolulu, Hawaii, USA, 2009.

[5] B. Wu, J. Chen, J. Wu, M. Cardei, A Survey on Attacks and Countermeasures in Mobile Ad Hoc Networks, in: Y. Xiao, X. Shen, D.-Z. Du (Eds.), Wireless/Mobile Network Security, Vol. 17 of Network Theory and Applications, Springer, 2006, Ch. 12.

[6] D. Djenouri, L. Khelladi, N. Badache, A Survey of Security Issues in Mobile Ad Hoc and Sensor Networks, IEEE Communications Surveys & Tutorials 7 (2005) 2–28.

[7] M. G. Zapata, N. Asokan, Securing Ad hoc Routing Protocols, in: Proceedings of the 3rd ACM Workshop on Wireless Security (WiSE '02), ACM Press, New York, NY, USA, 2002, pp. 1–10. doi:http://doi.acm.org/10.1145/570681.570682.

[8] Y. C. Hu, A. Perrig, D. B. Johnson, Ariadne: A Secure On-Demand Routing Protocol for Ad Hoc Networks, in: Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MobiCom '02), ACM Press, New York, NY, USA, 2002, pp. 12–23. doi:http://doi.acm.org/10.1145/570645.570648.

[9] Y. C. Hu, A. Perrig, D. B. Johnson, Rushing attacks and defense in wireless ad hoc network routing protocols, in: Proceedings of the 2nd ACM Workshop on Wireless Security (WiSe '03), ACM Press, New York, NY, USA, 2003, pp. 30–40. doi:http://doi.acm.org/10.1145/941311.941317.

[10] G. Acs, L. Buttyan, I. Vajda, Provable Security of On-Demand Distance Vector Routing in Wireless Ad Hoc Networks, in: R. Molva, G. Tsudik, D. Westhoff (Eds.), Proceedings of the 2nd European Workshop on Security and Provacy in Ad hoc and Sensor Networks (ESAS '05), Vol.

3813/2005 of Lecture Notes in Computer Science, Springer, 2005, pp. 113–127.

[11] G. Acs, L. Buttyan, I. Vajda, Provably Secure On-Demand Source Routing in Mobile Ad Hoc Networks, IEEE Transactions on Mobile Computing 5 (11) (2006) 1533–1546. doi:http://dx.doi.org/10.1109/TMC.2006.170.

[12] T. Anantvalee, J. Wu, A Survey on Intrusion Detection in Mobile Ad Hoc Networks, in: Y. Xiao, X. S. Shen, D.-Z. Du (Eds.), Wireless Network Security, Signals and Communication Technology, Springer US, 2007, pp. 159–180.

[13] S. Buchegger, Coping with Misbehavior in Mobile Ad-hoc Networks, Ph.D. thesis, École Polytechnique Fédérale de Lausanne (February 2004).

[14] P. Michiardi, R. Molva, Core: A Collaborative Reputation Mechanism to Enforce Node Cooperation in Mobile Ad Hoc Networks, in: Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, Kluwer, B.V., Deventer, The Netherlands, The Netherlands, 2002, pp. 107–121.

[15] S. Bansal, M. Baker, Observation-based Cooperation Enforcement in Ad hoc Networks, Tech. rep., Stanford University (July 2003).

[16] M. Hollick, J. Schmitt, C. Seipl, R. Steinmetz, On the Effect of Node Misbehavior in Ad Hoc Networks, in: Proceedings of the IEEE International Conference on Communications (ICC '04), Vol. 6, Paris, France, 2004, pp. 3759–3763.

[17] S. Marti, T. J. Giuli, K. Lai, M. Baker, Mitigating Routing Misbehavior in Mobile Ad Hoc Networks, in: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MobiCom '00), ACM Press, New York, NY, USA, 2000, pp. 255–265. doi:http://doi.acm.org/10.1145/345910.345955.

[18] Y. A. Huang, W. Lee, A Cooperative Intrusion Detection System for Ad Hoc Networks, in: Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN '03), 2003.

[19] Y. Zhang, W. Lee, Y.-A. Huang, Intrusion Detection Techniques for Mobile Wireless Networks, ACM/Kluwer Journal on Wireless Networks 9 (5) (2003) 545–556. doi:http://dx.doi.org/10.1023/A:1024600519144.

[20] P. Michiardi, Cooperation enforcement and network security mechanisms for mobile ad hoc networks, Ph.D. thesis, Ecole nationale supérieure des télécommunications (December 2004).

[21] V. Srivastava, J. Neel, A. B. MacKenzie, R. Menon, L. A. DaSilva, J. E. Hicks, J. H. Reed, R. P. Gilles, Using Game Theory to Analyze Wireless Ad Hoc Networks, IEEE Communications Surveys and Tutorials 7 (4) (2005) 46–56.

[22] G. Theodorakopoulos, J. S. Baras, Malicious Users in Unstructured Networks, in: Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM 2007), Anchorage, USA, 2007, pp. 884–891.

[23] J. J. Jaramillo, R. Srikant, DARWIN: Distributed and Adaptive Reputation mechanism for WIreless ad-hoc networks, in: Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking (MobiCom '07), ACM Press, New York, NY, USA, 2007, pp. 87–98. doi:http://doi.acm.org/10.1145/1287853.1287865.

[24] M. Mauve, J. Widmer, H. Hartenstein, A Survey on Position-Based Routing in Mobile Ad Hoc Networks, IEEE Network Magazine 15 (6) (2001) 30–39.

[25] Y.-B. Ko, N. H. Vaidya, Location-Aided Routing (LAR) in Mobile Ad Hoc Networks, in: Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98), ACM Press, New York, NY, USA, 1998, pp. 66–75. doi:http://doi.acm.org/10.1145/288235.288252.

[26] S. Basagni, I. Chlamtac, V. R. Syrotiuk, B. A. Woodward, A Distance Routing Effect Algorithm for Mobility (DREAM), in: Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98), ACM Press, New York, NY, USA, 1998, pp. 76–84. doi:http://doi.acm.org/10.1145/288235.288254.

[27] C. E. Perkins, E. M. Belding-Royer, S. R. Das, Ad hoc On-Demand Distance Vector (AODV) Routing, IETF RFC 3561.

[28] M. Hollick, Dependable Routing for Cellular and Ad hoc Networks, Ph.D. thesis, Multimedia Communications Lab (KOM), Technische Universität Darmstadt (November 2004).

[29] R. Barr, An Efficient, Unifying Approach to Simulation using Virtual Machines, Ph.D. thesis, Cornell University (2004).

[30] D. Thain, T. Tannenbaum, M. Livny, Distributed Computing in Practice: The Condor Experience, Concurrency and Computatio: Practice and Experience 17 (2-4) (2005) 323–356.

[31] A. König, M. Hollick, R. Steinmetz, A Stochastic Analysis of Secure Joint Decision Processes in Peer-to-Peer Systems, in: Proceedings of the 44th IEEE International Conference on Communications (ICC '09), 2009.

[32] K. Kanchanasut, A. Tunpan, M. A. Awal, D. K. Das, T. Wongsaardsakul, Y. Tsuchimoto, A Multimedia Communication System for Collaborative Emergency Response Operation in Disaster-affected Areas, Tech. Rep. TR_2007-1, Internet Education and Research Laboratory (intERLab), Asian Institute of Technology (AIT), Bangkok, Thailand (January 2007).

[33] SoKNOS Consortium, SoKNOS Project Homepage, http://www.soknos.de.

[34] C. Neuman, T. Yu, S. Hartman, K. Raeburn, The Kerberos Network Authentication Service (V5), IETF RFC 4120.

[35] A. Shamir, How to Share a Secret, Communications of the ACM 22 (1979) 612–613.

[36] R. Gennaro, S. Halevi, H. Krawczyk, T. Rabin, Threshold RSA for Dynamic and Ad-Hoc Groups, in: N. Smart (Ed.), Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '08), Vol. 4965 of Lecture Notes in Computer Science, Springer, 2008, pp. 88–107.

[37] V. Shoup, Practical Threshold Signatures, in: Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques (EUROCRPYT 2000), Vol. 1807 of Lecture Notes in Computer Science, Springer, Bruges, Belgium, 2000, pp. 207–220.

[38] M. Narasimha, G. Tsudik, J. H. Yi, On the Utility of Distributed Cryptography in P2P and MANETs: The Case of Membership Control, in: Proceedings of the 11th IEEE International Conference on Network Protocols (ICNP '03), IEEE Computer Society, Washington, DC, USA, 2003, p. 336.

[39] N. Saxena, G. Tsudik, J. H. Yi, Admission Control in Peer-to-Peer: Design and Performance Evaluation, in: Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN '03), ACM Press, 2003, pp. 104–113.

[40] N. Saxena, G. Tsudik, J. H. Yi, Efficient Node Admission for Short-lived Mobile Ad Hoc Networks, in: Proceedings of the 13TH IEEE International Conference on Network Protocols (ICNP'05), IEEE Computer Society, Washington, DC, USA, 2005, pp. 269–278. doi:http://dx.doi.org/10.1109/ICNP.2005.14.

[41] J. H. Yi, Energy-Efficient and Non-interactive Self-certification in MANETs, in: Proceedings of the 8th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2006), Vol. 4280 of Lecture Notes in Computer Science, Springer, Dallas, TX, USA, 2006, pp. 533–547.

[42] A. Rowstron, P. Druschel, Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems, in: R. Guerraoui (Ed.), Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms (Middleware 2001), Vol. 2218 of Lecture Notes in Computer Science, Springer, 2001, pp. 329–350.

[43] M. Castro, P. Druschel, A.-M. Kermarrec, A. Rowstron, Scalable application-level anycast for highly dynamic groups, in: B. Stiller, G. Carle, M. Karsten, P. Reichl (Eds.), Proceedings of the 5th COST 264 International Workshop on Networked Group Communications (NGC 2003), Vol. 2816 of Lecture Notes in Computer Science, Springer, Munich, Germany, 2003, pp. 47–57.

[44] PlanetLab Consortium, PlanetLab Homepage, http://www.planet-lab.org.

[45] P. Tran-Gia, T. Hoßfeld, M. Menth, R. Pries, Emerging Issues in Current Future Internet Design, e&i Elektrotechnik und Informationstechnik, Special Issue 'Future Internet', ISSN: 0932-383X (print), ISSN: 1613-7620 (online) 07/08.