
Detecting and Excluding Misbehaving Nodes in a P2P Network

Détection et Exclusion de Nœuds Malveillants dans un Réseau Pair-à-Pair

François Lesueur — Ludovic Mé — Valérie Viet Triem Tong

ABSTRACT. *Given their fully distributed architecture, P2P networks allow the design of low cost and high availability systems but also pose new security problems. In these collaborative networks, security properties need to be ensured by the participants themselves. In this paper, we propose to detect and exclude misbehaving nodes to allow honest participants to enforce security properties. The proposed scheme is fully distributed.*

RÉSUMÉ. *De par leur nature totalement distribuée, les réseaux pair-à-pair permettent de mettre en place des systèmes à faible coût et haute disponibilité mais posent également de nouveaux problèmes de sécurité. Dans ces réseaux collaboratifs, les propriétés de sécurité doivent être garanties par les participants eux-mêmes. Dans cet article, nous proposons de détecter et d'exclure des nœuds malveillants afin de permettre aux nœuds honnêtes de garantir des propriétés de sécurité. Le système proposé est complètement distribué.*

KEYWORDS: *P2P, Security, Behavior, Detection, Exclusion*

MOTS-CLÉS : *Pair-à-Pair, Sécurité, Comportement, Détection, Exclusion*

1. Introduction

P2P networks have been widely used for the last few years to design low cost and high availability systems. Indeed, P2P networks are fully distributed systems composed of many nodes (ordinary PCs) provided with replication mechanisms, allowing high availability. File sharing applications are well known (Gnutella [Cli00]), but these networks can also provide distributed file systems (CFS [DKK⁺01]), publish/subscribe applications (Meghdoot [GSAA04]), multicast (Split-Stream [CDK⁺03]) or Voice over IP (P2PSIP [BLJ05]).

There are two types of P2P networks: unstructured ones and structured ones. In unstructured P2P networks (Gnutella [Cli00]), requests are broadcasted or routed through random walks. In structured ones (Chord [SMK⁺01]), requests are routed efficiently using generated routing tables. We consider here structured P2P networks security.

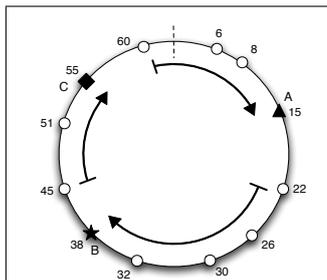


Figure 1: Simplified representation of Chord with identifiers going from 0 to $2^6 - 1 = 63$. ▲, ◆ and ★ are nodes (PC) and ○ are resources (files). Node A, which *nodeId* is 15, is responsible for resources 60, 6 and 8.

Structured P2P networks provide a virtual space called *overlay*. Each node (PC) is uniquely identified by a *nodeId* \in *KeyId*s; in the same way, each *resource* (file, ...) is uniquely identified by a key identifier *keyId*, *keyId* \in *KeyId*s (for a file, the key is usually its SHA1 fingerprint). Nodes and resources thus share the same identifier space *KeyId*s, which is finite but supposed large enough (often 2^{160} elements due to the size of SHA1 hashes). Each node is responsible for the management of a part of the resources. The overlay provides routing facilities with a logarithmic cost for a node to access a specific identifier.

In Chord for instance, the identifiers space goes from 0 to $2^{160} - 1$ and each node is responsible for all the resources which identifiers are between its identifier and the one of its preceding node in the overlay. This organization is illustrated in Figure 1.

To the contrary of classical networks in which security is provided through a centralized entity, P2P networks are collaborative networks in which each member is responsible for a part of the system. To resolve security problems, impact of malicious nodes has to be minimized.

We propose thus a scheme to identify users and then detect and exclude misbehaving members in a P2P network. The first part, presented in [LMT08c], allows users to obtain a unique provable identity and protects against the sybil attack [Dou02]. An accepted node obtains a unique membership certificate. The second part, presented in this paper, is a distributed misbehavior detection and node exclusion system, allowing to exclude misbehaving nodes from the network through the agreement of a static ratio of the nodes and without an administrative authority. Exclusion is materialized by the revocation of the membership certificate and the revocation process is handled using the distributed certification we presented in [LMT08b] (revocation is just a special case of certification). In this system, the normal behavior is dynamically determined by each node using static specifications defining how to compare different nodes behaviors: a node is detected as misbehaving if it shows a minority behavior. We thus make the (reasonable) assumption that there is a majority of honest nodes in the network.

We present in Section 1 the related work. In Section 2, we present the distributed certification we use to revoke membership certificates. In Section 3, we overview the detection and exclusion of misbehaving nodes. In Section 4, we detail the detection part and in Section 5, we precise the exclusion part. In Section 6, we analyze experimental results and we finally conclude and suggest some future work.

2. Related Work

In this section, we present work related to the evaluation of nodes in P2P networks. We first present work on admission control to a peer group and we then introduce reputation systems.

2.1. Admission Control to a Group

Kim *et al.* proposed in [KMT03] to limit the number of attackers in peer groups (P2P networks but also ad-hoc networks) through admis-

sion control to the network. Admission control is based on a *Group Charter* defining conditions to enter the network. A newcomer's access is then granted through a voting system among the members.

However, defining accepting conditions in P2P groups seems unrealistic since P2P systems are naturally open and large networks. Since open, everyone should be able to access the network; since large, we cannot rely on a significant part of the members knowing each newcomer. A newcomer can only be known by a very slight fraction of the members, and so relying on this very slight fraction of members to accept this newcomer would allow malicious users to easily join the network. Limiting the number of attackers in P2P networks can thus only be a reaction to misbehaving members, rejecting such members after bad behavior, and not a prevention.

2.2. Reputation Systems

Precisely, reputation systems [JIB05] assess nodes behavior and then apply relevant rules. Each node attributes each other node a given reputation, computed from local observations and recommendations. Reputation systems allow to favor good members over selfish ones.

However, even using recommendations, reputation systems are not tailored to exclude attackers after only one attack. If a honest node N detects an attacker A , N can stop communicating with A but other nodes will continue to communicate with A . In fact, other nodes cannot be sure that N is honest and so need other testimonies. Thus, an attacker attacking only a few nodes can be seen as honest by most of the nodes, whereas these few nodes might be in fact the replicas of an attacked resource. Some actions such as fake responses should lead directly to the attacker's exclusion, as soon as a single node detects it. We propose thus in this paper a complementary system which allows to directly and globally exclude some types of attackers.

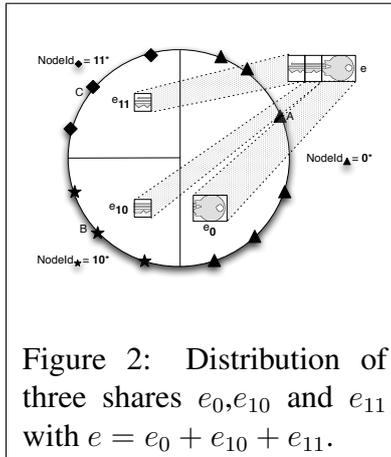
3. Distributed Certification

As stated in the introduction, we use the distributed certification scheme we proposed in [LMT08b] to revoke membership certificates

(revocation is a special case of certification) and we thus briefly describe relevant parts of this scheme. In this system, signing some data requires the collaboration of a fixed ratio t of the number of nodes. This ratio is enforced using a fully distributed scheme which tolerates misbehaving nodes and thus complies with the P2P basics. To our best knowledge, this is the only fully distributed scheme based on the cooperation of a fixed ratio of nodes instead of a fixed number of nodes, which is mandatory in varying size P2P networks.

The network is characterized by an RSA public/secret key pair (P, S) , $P = (d, m)$ being publicly known and $S = (e, m)$ being shared among the nodes (no node knows S entirely). This key pair is originally generated by founding members using a distributed algorithm as Boneh and Franklin proposed in [BF97]. The network is decomposed in s sharing groups, each group knowing one share e_i of S such that the sum of all the different shares equals e . In this case, the RSA signature of some data d is $d^e[m] = d^{\sum_{i=1}^s e_i}[m] = \left(\prod_{i=1}^s d^{e_i}[m] \right)[m]$. In other words, the RSA signature of d with S , which is $d^e[m]$, is equal to the product of the partial signatures with each share modulo m .

Signing a certificate requires then every share and thus the collaboration of one node of each sharing group. Given g_{min} and g_{max} the minimal and maximal sizes of sharing groups, the ratio t of nodes needed to sign a certificate verifies $\frac{1}{g_{max}} < t < \frac{1}{g_{min}}$ and thus, sharing groups can split (resp. merge) when nodes join (resp. leave) with only local knowledge to enforce this ratio t .



Each share is uniquely identified by a binary *shareId* and is only known by nodes which identifiers are such that $nodeId = shareId*$ in binary form (i.e., *shareId* is a binary prefix of *nodeId*). Splitting a group knowing e_i (resp. merging two groups knowing e_{i0} and e_{i1}) creates two groups knowing respectively e_{i0} and e_{i1} (resp. one group knowing e_i) such that $e_i = e_{i0} + e_{i1}$. This distribution is illustrated in Figure 2.

When a node A which identifier starts with 0 wants to sign some data d with $S = (e, m)$, A recursively partially signs d with (e_0, m) and asks a node B which identifier starts with 1 to partially sign it with (e_1, m) . B then partially signs it with (e_{10}, m) and recursively asks C to partially sign it with (e_{11}, m) . A node stops the recursion when it is asked for its own share, in which case it returns $d^{e_i}[m]$; otherwise, the node returns the product of the two recursive calls. A finally gets the product of the partial signatures with all shares which is the signature with S .

In this scheme, colluding attackers may eventually obtain every share of S or control every node in a sharing group, in which case the network is considered broken. Previous analysis showed that this scheme can resist such attacks. For instance, in a 10,000 nodes network with $g_{min} = 20$ and $g_{max} = 40$, the probability to obtain every share is infinitesimal for less than 20% of attackers and the probability to control every node in a sharing group is infinitesimal for less than 60% of attackers.

4. Detection and Exclusion

We propose to identify users and then to detect and exclude misbehaving members. Each node is first accepted in the network through a sybil-protection mechanism as the one we proposed in [LMT08c]. An accepted node obtains a unique certificate proving its membership. Then, misbehaving nodes are automatically excluded to limit their number. Nodes are thus monitored through the messages they send and those showing a bad behavior are excluded. In the following, we present the detection and exclusion of misbehaving nodes.

The detection of misbehaviors is based on Observable Behavior Specifications (OBS) (see Section 5). We assume that the majority of the nodes are well-behaving and so misbehaviors are detected by difference from other nodes. Since misbehaving nodes are only a minority, they have a different behavior from the majority. Normal behaviors are deduced dynamically according to observed behaviors and OBS.

We propose then to prove misbehaviors to lead to node exclusion. When a node detects a misbehavior, it generates a proof of this misbehavior. Such a proof must be verifiable by any other node in the network

and should thus not contain local unverifiable information. Moreover, a misbehavior proof leading directly to node exclusion, an attacker should not be able to forge a proof. In the proposed system, proofs are based on exchanged messages.

This proof is then used to globally exclude the misbehaving node through the collaboration of a given ratio of the nodes (see Section 6). The certificate of the node, which is its membership proof, is revoked using the distributed certification algorithm presented in Section 3.

5. Detecting Misbehaviors

In this section, we present the detection of misbehaving members. We first detail the precautions needed to prove misbehaviors. Then, we define the Observable Behavior Specifications (OBS) used to compare different nodes behaviors and we illustrate them on a proposed attack. Finally, we present the detection architecture.

5.1. Precautions to Consider

In our system, misbehaviors are proved and since a proof directly leads to the exclusion of the concerned node, an attacker should not be able to forge proofs. First, messages are signed in order to be unrepudiable and unforgeable. Then, we present in this section precautions to be able to judge the validity of a proof. We present *timeframes* allowing to check that observed behaviors are comparable and *decorellation* preventing an attacker from forging proofs with the help of accomplices.

5.1.1. Timeframes

Since P2P networks are dynamic, responses to the same request may be different over time. For instance, after the refresh of a share identified as *shareId*, the partial signature of given data with *shareId* changes.

To distinguish such legitimate changes from misbehaviors, nodes are provided with a common clock. The needed precision depends on the resources dynamics, but if changes are not very frequent (refreshes are not so frequent for instance), a precision of a few minutes is sufficient.

Nodes can initially synchronize their local clocks using NTP servers, an already deployed clock synchronization mechanism, and then rely solely on these local clocks.

Each response contains then two dates: the date since when this response has been valid (for instance the date when the share was last refreshed) and the date when the response has been sent. Dates are here used in the wide sense and contain days but also hours and minutes. The recipient accepts the message if and only if the emission date is the current date, tolerating a shift of a few minutes. Dates allow then to compare messages which should be the same.

5.1.2. *Decorellation*

Since nodes detect misbehaviors by the comparison of several responses to the same request, an attacker could forge a misbehavior proof by requesting some accomplices and one honest node: this honest node would exhibit a minority behavior and so would be detected as misbehaving.

To prevent this attack, the set of compared nodes is constrained: nodes used to generate a misbehavior proof must have uniformly distributed node identifiers among all the nodes eligible for the request. Since attackers have truly random identifiers (otherwise, this would be a sybil attack and the identification of users to enter the network prevents the sybil attack), the ratio of attackers in the constrained set of nodes is the same as in the network, hence a minority.

5.2. *Format of Observable Behavior Specifications*

The behavior of a node is reduced to its observable part: the messages it sends. Moreover, the misbehaviors are detected by the comparison to other nodes behaviors. OBS define thus how to compare messages of several nodes to detect misbehaviors.

We consider here that messages are composed of the *nodeId* of the sender node, an operation, a payload and a timeframe. The payload itself is composed of several fields depending on the type of message and is independent of the sender (all honest nodes reply with the same pay-

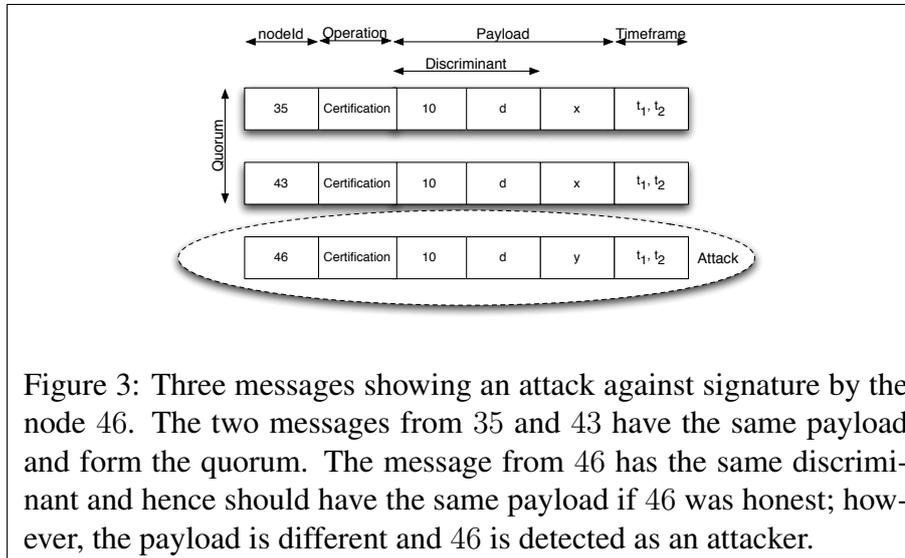


Figure 3: Three messages showing an attack against signature by the node 46. The two messages from 35 and 43 have the same payload and form the quorum. The message from 46 has the same discriminant and hence should have the same payload if 46 was honest; however, the payload is different and 46 is detected as an attacker.

load for a given request). The proposed scheme compares thus the payloads of the messages to detect misbehaviors. The timeframe is composed of two dates defining since when this response has been valid and when this response has been sent. This format of message is illustrated in Figure 3 which represents more precisely an attack against signature (see Section 5.3).

An OBS first specifies the attacked *operation*. In Figure 3, the three messages concern a Certification operation.

An OBS then describes the *direction* of attack which defines if the attack is in the request for an operation or in the response.

The *discriminant* defines the fields of the payloads which must be equal to consider that the payloads should be the same, if all the nodes were honest. The discriminant is composed of fields on which attackers cannot cheat, for instance representing the request they are answering: attacks consist thus in biasing the other fields of the payload. In Figure 3, the discriminant is composed of two fields defining the partial signature which was asked and the three presented messages have the same discriminant: these messages should be the same if nodes were honest.

The *quorum* defines the number of identical payloads needed to consider that the contained value is the valid one. In Figure 3, if we con-

sider an OBS with a quorum of 2 and the proposed discriminant, then the messages from nodes 35 and 43 prove that the valid payload is composed of 10, d and x . Then, the different payload sent by 46 is considered as an attack (same discriminant but different payload).

Finally, the *scope* defines the set of eligible nodes for the request. This set can be the whole network, a set of nodes (for instance k replicas) or the nodes sharing a given prefix (for a partial signature in distributed certification). This *scope* allows to define the set of identifiers in which the compared nodes representing the *quorum* have to be uniformly distributed and enforces then the decorellation.

5.3. Detecting an Attack against Signature

We plan on detecting attacks on any application. However, we start by detecting attacks against the distributed certification since the exclusion mechanism presented in this paper relies itself on this distributed certification. In this paper, we focus precisely on the attack against signature; the interested reader can refer to [LMT08a] for two other detected attacks against the distributed certification. We propose then in future work to detect attacks against other P2P services such as DHT.

In the distributed certification presented in Section 3, an attacker can return a wrong partial signature when asked to sign with its own share or return a wrong product otherwise. To prevent this attack, each node asks $nbAsks$ nodes instead of one and uses the majority result; however, it is crucial to limit these attacks to improve the performance (asking several nodes is expensive) and to prevent attackers from becoming a majority, in which case the certification would not work anymore. The payload of a response contains three fields: the share identifier which was asked ($shareId$), the data which has been signed ($data$) and the partial signature returned ($psig$). This attack is the one illustrated in Figure 3 with $shareId = 10$, $data = d$ and $psig = x$ or y .

To detect the attack against signature, we propose to compare the $nbAsks$ responses ($nbAsks = 3$ in the figure). This attack consists in sending a fake response to a partial signature request and so the *operation* is Certification and the *direction* is response. This attack resides

in replying a fake *psig* (y instead of x in the figure) for given *shareId* and *data* and so the fields which have to be equal among compared payloads (*discriminant*) are *shareId* and *data*. The *quorum* is fixed to $\frac{nbAsks+1}{2}$ which means that a payload is considered valid as soon as the majority of the nodes returned it. The eligible nodes for a request are all the nodes which identifiers are prefixed by the share identifier asked and the *scope* is thus all these nodes.

When some attackers have been excluded, honest nodes automatically detect that all the partial signatures are valid and can thus dynamically reduce the number of nodes asked for each partial signature, in order to balance the correctness and the efficiency of the requests.

5.4. Detection Architecture

We first overview the detection mechanism, composed of the *classifier* and of the *checker*. We then detail the classifier and the checker.

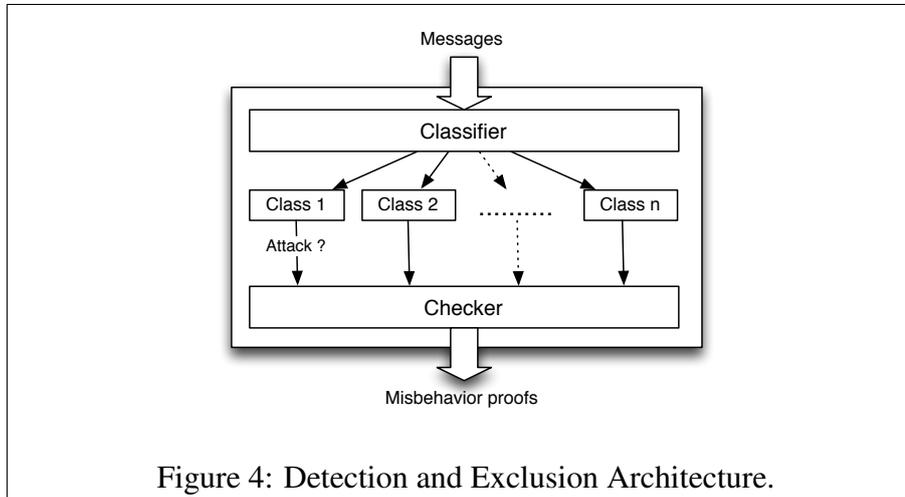
5.4.1. Overview

To detect misbehaviors, each node monitors all the messages it sends and receives. This observer is placed between the applicative layer, represented here by the distributed certification mechanism, and the P2P overlay. The detection of misbehaviors is split in two parts: the *classifier* and the *checker* (Figure 4).

First, the *classifier* (see Section 5.4.2) associates each sent or received message to other related messages in a *class*, according notably to discriminants. For instance, Figure 3 corresponds to a class of messages. Behaviors are then checked among messages of the same class.

When a message is added to a class, the *checker* (see Section 5.4.3) verifies if this class corresponds to a misbehavior. If a misbehavior is detected, a misbehavior proof accusing one or several nodes is created from this class.

This misbehavior proof is then used to exclude the misbehaving node (see Section 6).



5.4.2. Classifier

The *classifier* associates in a *class* messages which should be compared to detect misbehaviors. Comparable messages are the ones corresponding to a given OBS, having the same discriminant and sent during the same timeframe. Such messages should be identical and so a different message in a class indicates a misbehavior, detected by the *checker*.

Each new message is labeled by a set of triplets $\{(OBS, Discriminant, Timeframe)^*\}$ which is initially empty. First, a triplet is added to this set for each OBS defined on the operation of this message. Then, the discriminant of each of these triplets is set to the values found in the message, w.r.t. the discriminant of the related OBS. Finally, the timeframes are set to the timeframe of the message. For instance, the three messages of Figure 3 are labeled by the set $\{(Signature\ OBS, (10, d), (t_1, t_2))\}$.

Each new message is then associated in a class to every other message having a common triplet in its label. Two timeframes are considered common as soon as their intersection exists, i.e., they have a common period of validity. The three messages of Figure 3 form thus a class and should be identical if all the nodes were honest.

Each class is finally erased after a fixed amount of inactivity time.

5.4.3. *Checker*

When a message is added to a class, the *checker* applies the OBS to detect attacks in this class, attacks which correspond to differences in the payloads. If a misbehavior is detected, the checker creates the proof accusing one or several nodes.

A misbehavior proof is a minimal class proving the misbehavior. A proof contains the minimal number of messages representing the normal behavior, i.e., the quorum of the OBS triggered, plus the message(s) of the misbehaving nodes.

However, a misbehaving node may try to fool the system by sending a fake response at a date t and announcing this response is valid since $t - \epsilon$, ϵ being a short time. Given the imprecise clock used, it is not possible without further investigations to detect whether this is an attack or whether the valid response has just changed. In such a case, the checker waits a bit and then re-asks the nodes composing the quorum. If the suspected node is honest then the novel responses also take into account the change of the valid value; if the suspected node is misbehaving then the novel responses include the date t in their validity period and the misbehavior proof includes thus these new messages as well as the old message of the attacker.

The node which has detected the misbehavior uses then the proof to exclude the misbehaving node through the mechanism presented in the following section.

6. Excluding Attackers

The exclusion of an attacker is materialized by the revocation of its certificate. We first present how a certificate is revoked. Then, we explain the publication of the revocation and we finally detail what happens after an exclusion.

6.1. *Certificate Revocation*

A certificate revocation is of the form $\{Cert, revoked\}_s$, where $Cert$ is the membership certificate to revoke and $revoked$ is a flag indi-

cating the revocation, all that being signed with the network secret key S . Certificate revocation uses then the same algorithm as distributed certification to obtain the signature with S . The node detecting the misbehavior provides the proof of the misbehavior to each node involved in the revocation.

Each involved node locally validates the proof. Each node checks the proof as if it was a local class of messages and also verifies that the quorum is constituted of nodes uniformly distributed in the scope of the operation (otherwise, an attacker would be able to ask some accomplices to bias the normal behavior). However, these nodes cannot be truly uniformly distributed since node identifiers represent only a small part of the identifiers space: chosen nodes are in fact the ones responsible for the identifiers which are uniformly distributed. During the revocation, each node has thus to check whether the proposed nodes are really responsible for these resources.

To be able to decide if a given node is close enough to a resource to be responsible for it, each node calculates the density of nodes in its neighborhood. Node identifiers being uniformly distributed using a hash function, this density is roughly the same in the whole overlay and a node can thus estimate the density of nodes in the overlay. Even if not so precise, this heuristic allows to limit the set of nodes used in a proof.

If the proof is valid, this node proceeds with the distributed certification and the initiator node finally obtains the signed revocation if a given ratio of the nodes validated the proof. It distributes then this revocation to other nodes.

6.2. Revocation Publication

To effectively exclude the misbehaving node, the revocation must be notified to all nodes. To limit the communications, revocations are stored in the DHT and each node checks for their presence before communicating with an unknown node (however, nodes mostly communicate with a small and quite stable subset of the nodes). If the identifier of the attacker is $nodeId$ then this revocation is put at $h(nodeId)$ (putting it at $nodeId$ would transiently put this revocation under the responsibil-

ity of the attacker). Given the replication mechanisms of a DHT, this revocation is finally stored on some successors of $h(nodeId)$, which guarantees its availability even if one of the replicas is malicious.

Revocations are also directly sent to nodes already communicating with the attacker. These nodes are the members of the attacker's sharing group, the nodes which have the attacker in their routing table, the nodes which are in the attacker's routing table and possibly some nodes connected at the applicative level. Members of the sharing group of the attacker are easily found since the prefix of their identifiers is known (the attacker has the same prefix). The way to find nodes which have the attacker in their routing tables and nodes which are in the attacker's routing table depends on the overlay used. In Chord for instance, there are 2^{160} identifiers and the attacker identified by $nodeId$ has in its routing table the nodes succeeding the identifiers $nodeId + 2^{159}$, $nodeId + 2^{158}$, etc. Nodes having the attacker in their routing table can be found in the same way. Finally, finding nodes connected to the attacker at the applicative level depends on the application. The distribution of revocations could also be handled through a *publish/subscribe* mechanism, each node interested in the hypothetical revocation of a given node subscribing to this feed and the revocation being published on this feed. All that being done, the attacker cannot communicate anymore with any node in the overlay.

6.3. *Actions after an Exclusion*

The exclusion of a node leads to modifications of the P2P overlay and to maintenance of the shared secret key of the network S . Modifications of the overlay are the same than when a node leaves the network: the attacker is replaced in routing tables and another node takes control of the resources it was responsible for.

Then, the share of the secret key which was known by this attacker must be invalidated. Indeed, this share being known by an attacker, it is considered exposed. This share is thus *refreshed* as detailed in [LMT08b], operation in which this share is mixed with another one rendering old version of both shares useless. Basically, two groups agree on a random value Δ , one group adds it to its share and the other one

subtracts it: both shares change yet their sum remains constant. Finally, since the group of the attacker may become composed of less than g_{min} members after the exclusion, this group may trigger a merge operation with the adjacent group.

7. Simulations

In this section, we present experimental results. These results are presented in the case of the attack against signature which seems the most dangerous. Indeed, this attack quickly decreases the performance of the certification system, certification system which is used to revoke misbehaving members. In the first part, we study the reliability of the detection system. In the second part, we show the impact of node exclusion in the lifetime of the system.

7.1. Detection System Reliability

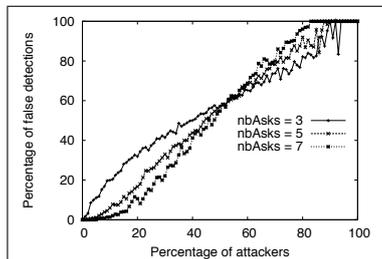


Figure 5: Percentage of false detections in function of the percentage of attackers. Each experiment contains 5000 nodes.

Figure 5 shows the percentage of false detections (honest nodes detected as misbehaving) in function of the percentage of attackers. This figure shows a symmetry which corresponds to the change of the normal behavior. In the first half of the curve, attackers are a minority and the normal behavior is the honest one; in the second half, attackers are a majority and the normal behavior is the misbehaving one, yielding the detection of honest nodes as misbehaving.

However, we consider that a ratio of 1 attacker for 10 nodes is already quite high. In such a case, the percentage of false detections is less than 5% when each node asks each partial signature to 5 nodes ($nbAsks = 5$). Irregularities with more than 90% of attackers are explained by the fact that attackers collude and never participate in the detection of one of them, whereas honest nodes may testify among them.

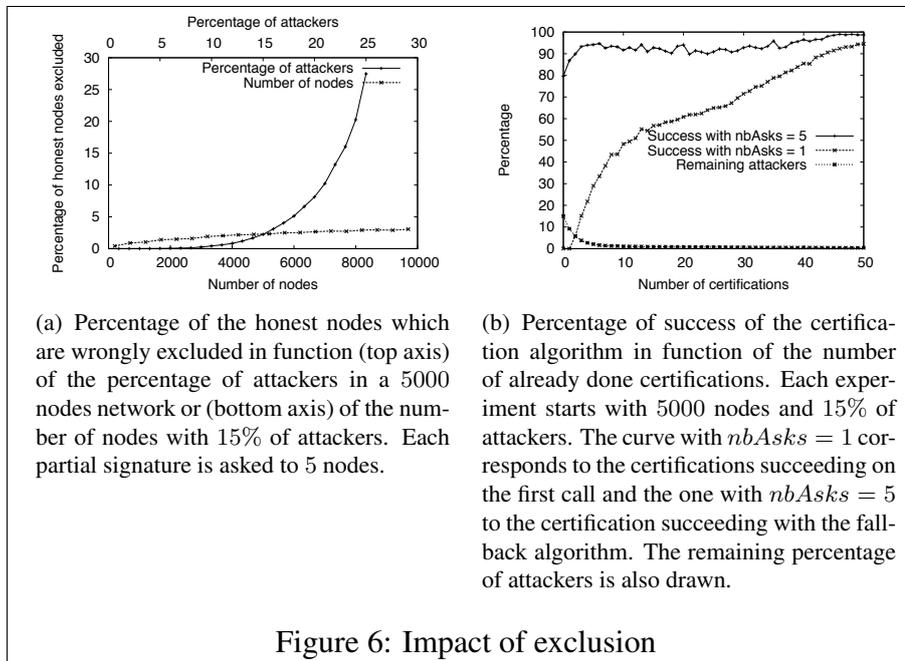


Figure 6: Impact of exclusion

7.2. Impact of Exclusion

Figure 6(a) shows the percentage of the honest nodes which are wrongly excluded with $nbAsks = 5$. Honest nodes can be wrongly excluded if attackers are enough to return a (wrong) majority response. Even if some misbehaviors are not detected the first time, all attackers are finally excluded in the simulated configurations and the number of honest nodes which have been excluded is calculated when all attackers have been excluded. The first curve is drawn for networks composed of 5000 nodes with a percentage of attackers varying from 0 to 25%; the second curve is drawn for networks composed of 0 to 10000 nodes with 15% of attackers. The number of honest nodes excluded is low since only 2% of them are wrongly excluded in a 5000 nodes network with initially 15% of attackers.

In Figure 6(b), we illustrate the impact of exclusion on the success of certification operations. Each node asks first only 1 node and then 5 if the obtained signature is invalid. This figure shows the success of successive certifications and shows that with initially 15% of attackers in a

5000 nodes network, the algorithm with $nbAsks = 1$ is quickly usable after having used $nbAsks = 5$ a few times, this last value allowing to exclude misbehaving nodes. Success with $nbAsks = 1$ quickly grows after only a few certifications. This figure also shows the remaining percentage of misbehaving nodes after each certification. The percentage of misbehaving nodes drops from 15% to 1% in 10 certifications.

8. Conclusion

In this paper, we proposed to detect and exclude misbehaving nodes in a P2P network. The detection is based on the comparison of nodes behavior and the exclusion is handled through distributed algorithms. Based on an attack against signature, we precisely described the mechanisms to detect misbehaviors, to prove misbehaviors to other nodes and finally to exclude misbehaving nodes through revoking their certificates. The mechanism is extensible to detect attacks on other applications. We finally experimentally studied the performance of the proposed system. These simulations showed that our system performs well in networks composed of up to 10% of attackers, which is a high percentage since these attackers are progressively excluded from the network.

In this proposition, each node monitors and temporarily saves many messages. It will thus be interesting to study the impact if each node only monitors a subset of message classes. Since several nodes may see the same attack, it should be possible to attain comparable performance. We are now interested in evaluating the benefits of this security mechanism on a P2P application. In a DHT for instance, an attacker can return a wrong value for an asked resource. Since resources are replicated on several nodes, an expensive countermeasure is to ask several nodes, through different routes, for the same resource. With our detection and exclusion system, it is possible to detect the attackers which returned a wrong value, exclude them, and finally reduce the number of nodes asked for the same resource without sacrificing reliability.

Acknowledgements

The authors thank Orange Labs for partially funding this work and especially Hervé Debar for his involvement in the project and his advice and valuable comments.

References

- [BF97] Dan Boneh and Matthew Franklin. Efficient generation of shared RSA keys. In *Proceedings of the 17th Annual International Cryptology Conference (CRYPTO)*, volume 1294 of *Lecture Notes in Computer Science*. Springer, 1997.
- [BLJ05] David A. Bryan, Bruce B. Lowekamp, and Cullen Jennings. SOSIMPLE: A serverless, standards-based, P2P SIP communication system. In *Proceedings of the International Workshop on Advanced Architectures and Algorithms for Internet Delivery and Applications (AAA-IDEA)*, 2005.
- [CDK⁺03] Miguel Castro, Peter Druschel, Anne-Marie Kermarrec, Animesh Nandi, Antony I. T. Rowstron, and Atul Singh. Splitstream: high-bandwidth multicast in cooperative environments. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, volume 37, 5 of *Operating Systems Review*, pages 298–313. ACM Press, 2003.
- [Cli00] Clip2. The gnutella protocol specification v0.4. http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf, 2000.
- [DKK⁺01] Frank Dabek, M. Frans Kaashoek, David Karger, Robert Morris, and Ion Stoica. Wide-Area cooperative storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, volume 35, 5 of *Operating Systems Review*, pages 202–215. ACM Press, 2001.
- [Dou02] John R. Douceur. The sybil attack. In *Proceedings of the International Workshop on Peer-to-Peer Systems (IPTPS)*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer, 2002.
- [GSAA04] Abhishek Gupta, Ozgur D. Sahin, Divyakant Agrawal, and Amr El Abbadi. Meghdoot: Content-based publish/subscribe over P2P networks. In *Proceedings of the ACM/IFIP/USENIX International Middleware Conference*

(*Middleware*), volume 3231 of *Lecture Notes in Computer Science*, pages 254–273. Springer-Verlag, 2004.

- [JIB05] Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. In *Decision Support Systems*, 2005.
- [KMT03] Yongdae Kim, Daniele Mazzochi, and Gene Tsudik. Admission control in peer groups. In *Proceedings of 2nd IEEE International Symposium on Network Computing and Applications (NCA)*, pages 131–139. IEEE Computer Society, 2003.
- [LMT08a] François Lesueur, Ludovic Mé, and Valérie Viet Triem Tong. Detecting and excluding misbehaving nodes in a P2P network. Technical Report TR-2008-04, SUPELEC, 2008. http://www.rennes.supelec.fr/ren/perso/flesueur/publis/LMV08_tr_04.pdf
- [LMT08b] François Lesueur, Ludovic Mé, and Valérie Viet Triem Tong. A Distributed certification system for structured P2P networks. In *Proceedings of the 2nd International Conference on Autonomous Infrastructure, Management and Security (AIMS)*, Lecture Notes in Computer Science. Springer-Verlag, 2008.
- [LMT08c] François Lesueur, Ludovic Mé, and Valérie Viet Triem Tong. A sybil-resistant admission control coupling Sybil-Guard with distributed certification. In *Proceedings of the 4th International Workshop on Collaborative Peer-to-Peer Systems (COPS)*. IEEE Computer Society, 2008.
- [SMK⁺01] Ion Stoica, Robert Morris, David R. Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the ACM SIGCOMM Conference (SIGCOMM)*, Computer Communication Review, pages 149–160. ACM Press, 2001.