ARTICLE IN PRESS

Available online at www.sciencedirect.com

ScienceDirect

ELSEVIER

Ad Hoc Networks xxx (2006) xxx–xxx

Ad Hoc Networks

www.elsevier.com/locate/adhoc

A localized certificate revocation scheme for mobile ad hoc networks

Geneviève Arboit, Claude Crépeau, Carlton R. Davis *, Muthucumaru Maheswaran

School of Computer Science, McGill University, 3480 University Street, Montréal, Que., Canada H3A2A7

Received 13 June 2005; received in revised form 20 June 2006; accepted 25 July 2006

Abstract

The issue of certificate revocation in mobile ad hoc networks (MANETs) where there are no on-line access to trusted authorities, is a challenging problem. In wired network environments, when certificates are to be revoked, certificate authorities (CAs) add the information regarding the certificates in question to certificate revocation lists (CRLs) and post the CRLs on accessible repositories or distribute them to relevant entities. In purely ad hoc networks, there are typically no access to centralized repositories or trusted authorities; therefore the conventional method of certificate revocation is not applicable.

In this paper, we present a decentralized certificate revocation scheme that allows the nodes within a MANET to revoke the certificates of malicious entities. The scheme is fully contained and it does not rely on inputs from centralized or external entities.

Keywords: MANET security; Public-key cryptography; Trust model; Certificate revocation; Key management; Hash chain

1. Introduction

As MANETs become more ubiquitous, the need for adequate security in these networks is more evident. Security schemes for MANETs generally employ one or more of the following cryptographic technologies: symmetric-key cryptography, digital certificates or threshold cryptography. Each of these cryptographic tools has its particular advantages and drawbacks; for example, security schemes involving symmetric-key cryptography are much less computationally exhaustive than those involving digital certificates or threshold cryptography. Consequently, the use of symmetric-key cryptography has much smaller computational overhead than that associated with digital certificates or threshold cryptography. However, security schemes which are based solely on symmetric-key cryptography, such as [1,2], are less robust and offer lower degree of security than those involving asymmetric key cryptography, owing to the greater probability of the shared (symmetric) keys being compromised.

---

* Corresponding author. Tel.: +1 450 692 6736.

E-mail addresses: garboit@cs.mcgill.ca (G. Arboit), crepeau@cs.mcgill.ca (C. Crépeau), carlton@cs.mcgill.ca (C.R. Davis), maheswar@cs.mcgill.ca (M. Maheswaran).

URLs: http://crypto.cs.mcgill.ca/~garboit (G. Arboit), http://www.cs.mcgill.ca/~crepeau (C. Crépeau), http://www.cs.mcgill.ca/~carlton (C.R. Davis), http://www.cs.mcgill.ca/~maheswar (M. Maheswaran).

1570-8705/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.adhoc.2006.07.003

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

The utilization of threshold cryptography for the design of MANETs security schemes has generated some interest. This approach is based on the work of Shamir [3], who proposed the concept of (k,n) threshold scheme; whereby a secret can be split into n shares, such that for a certain threshold $k < n$, any k components can combine and reconstitute the secret, whereas the combination of $k - 1$ or less shares are incapable of reconstructing the secret. Shamir's work was later extended by [4–6] into verifiable secret sharing, such that the shares can be verified to determine whether or not they are consistent. Robust threshold signature schemes have been developed for both RSA and discrete logarithm-based signature schemes [7,8]. The idea of utilizing threshold cryptography to distribute trust in ad hoc networks was first presented by Zhou and Hass [9]; later extensions of this proposition include [10–14]. Threshold cryptography offers viable security solutions for certain MANETs environments; in that a certificate authority (CA) signing key can be split and distributed to n nodes, such that any k of the n nodes can collaborate and sign digital certificates. In so doing, certificates can be issued on-the-fly without input from external entities. However, the issue of certificate revocation in these distributed environments is still an open problem. To date, the MANET threshold cryptographic security schemes, such as [10,15,16] which explicitly address the issue of certificate revocation, either do not provide protection against certificates being wrongfully revoked through malicious accusations, or they assume—as is the case for [10]—that access to external CAs is available.

Certificates issued via non-threshold cryptographic schemes require the utilization of some sort of trust model. The most commonly used trust models are (a) hierarchical and (b) web-of-trust models. The hierarchical trust model is the more structured approach and the most widely used. In the hierarchical trust model, a root certificate authority (CA) issues certificates to delegated CAs or end users, the CAs in turn issue certificates to end users or to other CAs. Fig. 1 illustrates the hierarchical trust model. The PKI X.509 (PKIX) framework [17] exemplifies this trust model.

The web-of-trust model [18] is the more distributed approach. In this model, there is no distinction between CAs and end users. End users are responsible for all certificate management tasks, such as issuing, storage and revocation of certificates. An end user A issues a certificate to another user B if A trusts B or if a user C that A trusts, vouches for B. Fig. 2 illustrates the web-of-trust model. The web-of-trust model appears attractive for utilization in MANETs security schemes, owing to its distributed nature. However, the web-of-trust model is far more susceptible to infiltration of malicious agents than the more structured hierarchical model, since the latter allows much greater accountability than the former. Consider for example a network where a node A trusts another node B; if B happens to be a malicious agent, B can issue valid certificates to several other malicious agents who would be implicitly trusted by A since B—who A trusts—vouches for these agents. Similarly, if other nodes trust B, these nodes would also implicitly trust the malicious agents B vouches for. Consequently, a number of malicious agents can gain access to the network if a single untrustworthy node happens to convince another node to issue it a valid certificate.

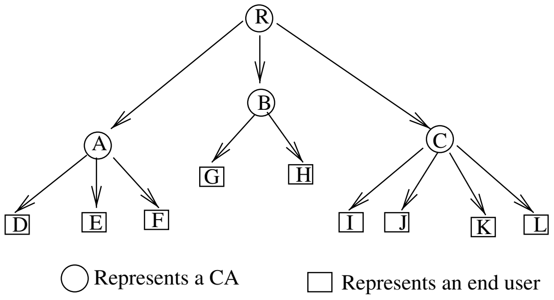

Fig. 1. Hierarchical trust model.

The hierarchical trust model offers greater protection against this eventuality, in that the end users are accountable to the CAs that issue the certificates, and the CAs in turn accountable to other CAs or to the root CA. If a network is compromised, this accountability structure allows the elimination of malicious agents much more readily.

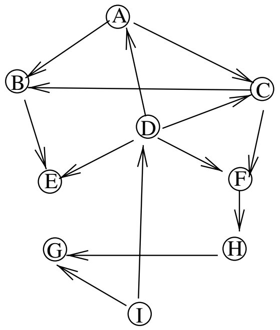

Fig. 2. Web-of-trust trust model.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

Hierarchical trust model is therefore preferable, particularly in environments where higher degrees of accountability and security are required. Security schemes such as [19,20] are viable solutions for some MANET environments; however, owing to the fact that they utilize the less stringent web-of-trust model, they may not be suitable for MANETs environments where high degrees of accountability and security are required.

There are some notable challenges however in utilizing certificates that are based on the more reliable hierarchical trust model in MANETs, owing to the decentralized nature of these networks. One particular challenging problem is the issue of certificate revocation. For various reasons—such as the compromisation of private keys—certificates will need to be revoked periodically, and network peers need to be informed about the revoked certificates in a timely manner. For conventional networks, CAs issue certificate revocation lists (CRLs) [21] which contain information about revoked certificates, at regular intervals. The CRLs are then either broadcast to the relevant nodes, or placed on easily accessible centralized repositories. Alternatively, on-line certificate status protocol (OCSP) [22] can be used to ascertain information about the status of a certificate. These methodologies are not applicable to MANETs, owing to the fact that MANETs do not contain centralized entities, and they typically do not provide on-line access to external entities such as CAs.

In this paper, we present a decentralized certificate revocation scheme for MANETs that allows the revocation of certificates in such a way that protection is provided against wrongful revocation of certificates through malicious accusations. The rest of the paper is organized as follows: Section 2 reviews previous work related to reputation-based systems and certificate revocation in MANETs. Sections 3 and 4 provide an overview and detail, respectively, of our certificate revocation scheme. In Section 5, we present analysis of the scheme; Section 6 contains simulation results, and Section 7 summaries the contributions of this paper.

2. Related work

Most of the proposed ad hoc network security schemes which utilized certificates that rely on hierarchical trust model, do not explicitly address the issue of certificate revocation. Examples of these schemes include [23–26]. Other proposals such as [27,28] make the assumption that periodic access to on-line CAs is available; therefore CRLs can be obtained from the CAs. Then there are proposals such as [29] which make provision for certificate revocation, and do not assume that on-line CAs are accessible; but they do not provide protection against certificates being wrongfully revoked through malicious accusations.

Our scheme can be distinguished from the proposals indicated above, in that it does not assume any accessibility to on-line CAs, and it is specifically designed such that protection against wrongful certificate revocation through malicious accusation is provided. [30,31] contain some preliminary results of this research project.

In [32], Buchegger and Le Boudec proposed the CONFIDANT protocol that is aimed at detecting and isolating misbehaving nodes. It uses reputation systems [33] to rate the nodes. Our work is based on the same principle; however, it can be differentiated from theirs in that we present a methodology for actually computing the trust level or rating of the nodes within a MANET.

A number of reputation systems have been published in research literature. These systems can be divided into two main types: centralized and distributed reputation systems. Centralized reputation systems require central authorities for collecting the rating of participants and derive reputation scores. Examples of these systems are [33,34]: the reputation systems on which eBay[1] forum and Amazon,[2] respectively, are based; and the page ranking scheme [35] developed by the founders of Google.[3] Centralized reputation systems are not suitable for MANETs since MANETs do not have centralized entities. Decentralized systems are more fitting for MANET applications. The majority of proposed decentralized reputation systems are transactional based; that is, they require inputs—such as size of upload or down files, quality, price and upload/download experiences—relating to interactions of providers of services and users of the services. Examples of transactional based reputation systems are [36–40]. The non-transactional based systems previously proposed are not suitable for application in certificate revocation schemes because they are either too complex and have high associated overhead [41,42], or they are based on assumptions such

---

[1] http://www.ebay.com.

[2] http://www.amazon.com.

[3] http://www.google.com.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

as those outlined in [43,44], which are not applicable to certificate revocation schemes.

3. Overview of the certificate revocation scheme

Our scheme stipulates that before entering a network, the MANET nodes must have a valid certificate from a recognized CA, as well as the public keys of the CAs which issued certificates for potential network peers. The certificates can be used for network authentication. The nodes will be able to verify the validity of the certificates, since they have the public keys of the CAs which issued them. The MANET nodes are therefore responsible for all key management tasks except the issuing of certificates. For optimum security, a CA should verify the identity of a node before issuing it a certificate.

Our certificate revocation scheme requires the nodes in a MANET to monitor the behavior of the other nodes. If a node surmises that a given node is behaving suspiciously, it is required to broadcast an accusation against the node in question. Our scheme utilizes the self-healing community approach presented in [45] for disseminating the accusation info via broadcast. Self-healing community approach is based on the observation that in a MANET, any node that is within both node *A* and node *C* transmission range can in principle forward packets from node *A* to *C*. For example, in Fig. 3, nodes *A* and *C* are outside the transmission range of each other. In principle, any of the nodes (*n*1, *n*2, *n*3, *n*4) within the self-healing community can forward packet from *A* to *C*. So, if a malicious or selfish node within a self-healing community chooses not to forward a packet it is asked to forward, any other node within the community can provide the service instead. A self-healing community is functional as long as there is at least one well-behaving node in the community. This approach requires the network interfaces of the MANET nodes to stay in promiscuous reception mode. For further detail and analysis of the self-healing community concept, see [45].

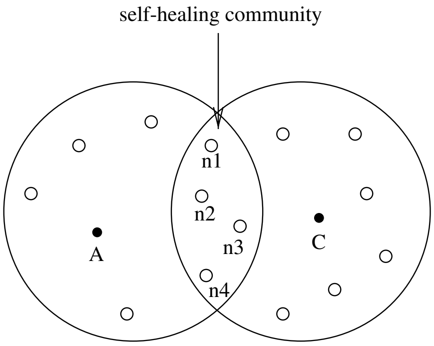

Fig. 3. Self-healing community packet forwarding.

Our certificate revocation scheme requires each participating node to compile and maintain data—based on broadcast accusation info—about all the nodes in the network. The collected data is used to assign a quantitative value for the trustworthiness of a node. Accusations from any given node are weighted based on the trustworthiness of the accuser: the higher the trustworthiness of a node, the greater the weight of its accusations, and vice versa. A node's certificate is revoked if the value of the sum of accusation weights against the given node is greater than a configurable threshold. The protocol aims at providing similar data to each node for computing the trust ratings of the network peers; the end goal being that the nodes have consistent info regarding the status of the certificates of their network peers.

3.1. Cryptographic primitives

For efficiency considerations, rather than relying on digital signatures for message origin authentication and content integrity checks, we mainly use one-way hash chains [46]. One-way hash chains are based on one-way hash functions. A one-way hash function H, maps an input x of any length to an output y of fixed length, such that, given y, it is computationally infeasible to find x, where $H(x) = y$. Two commonly used one-way hash functions are SHA-1 [47]—which produces 160-bit outputs—and MD5 [48], which gives 128-bit outputs.

A one-way hash chain can be created by choosing a random value x of arbitrary length and compute the hash chain values $y_0, y_1, y_2, \ldots, y_{n-1}, y_n$, where $y_0 = x$ and $y_i = H(y_{i-1})$, such that $0 < i \leqslant n$, for a given n. The hash chain values—in order of decreasing subscript i (that is, from right to left in the list above)—at varying point in time can then be used for authentication or as symmetric keys for keyed hashing functions such as HMAC [49]. When the hash chain values are used as keys for keyed hashing functions, for example, y_n can be signed and be distributed to network peers who will use it to authenticate the other y_i values. y_{n-1} can then be utilized with HMAC to generate a message authentication code (MAC) for a message m_1, and appended to m_1 before it is transmitted. After a

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

designated time period, y_{n-1} is released and utilized by the recipient of m_1 to verify the message integrity. Similarly, at a later point in time, y_{n-2} can be used to generate a MAC for another message m_2. The network peers are able to authenticate the y_i values since y_n is signed and they can verify whether $y_{i+1} = H(y_i)$, for all previously seen $i \leqslant n$. Unlike TESLA [50], our protocol does not require time synchronization, owing to the unique way we utilize the hash chains.

4. Detail of scheme

The following assumptions are made regarding to the MANETs and the nodes that constitute the networks:

- The number of malicious or selfish nodes is less than the number of well-behaving nodes.
- The network interfaces of the nodes are capable of operating in promiscuous reception mode.
- Each node has only one valid certificate.

The first duty of a node when it enters a MANET is to compute a series of hash chain values $y_0, y_1, y_2, \ldots, y_{n-1}, y_n$, using an agreed upon hash function H, as outlined in Section 3.1, if they have not been computed a priori; sign y_n and broadcast it along with its certificate to the nodes in the network. Upon receiving a signed y_n and the corresponding certificate, the nodes verify that the certificate is valid. If it is valid and it is not revoked, and the signature on the y_n value is valid, the nodes store both the certificate and y_n; sign their profile tables and their y_n values, and unicast them to the sender of the certificate. Note that if a node has already used any of its y_i values to secure messages, it will sign and send the last y_i it utilized—as its y_n value—to new entrants to the network. A profile table contains information about the behavior profile of the nodes in the MANET.

Upon receiving the profile tables with valid signatures from its network peers, a node is required to compile its own profile table which is initially based on the information contained in the profile tables it received. Transmission of profile tables to new entrants to the network is necessary in order to ensure that the newcomers have up-to-date information regarding the behavior profile of its network peers.

A profile table can be represented as a packet of varied length depending on the number of accusations launched against the nodes. The length ranges from a minimum of 80 bits—when there are no accusations—to a maximum of $97(N-2)+145$, where N is the number of nodes in the network. A profile table contains the following fields:

1. *Owner's ID*: This field is the first 32 bits of the profile table. It contains the certificate serial number of the node that compiled the profile table.
2. *Node count*: This 16-bit field contains a short integer indicating the node perspective regarding the number of nodes in the network.
3. *Peer i ID*: This is a 32-bit field containing the certificate serial number of a node that is accused of misbehavior. This field also serves the purpose of a marker: if it contains zero, it indicates the end of the profile table.
4. *Certificate status*: This field contains 1-bit flag. The bit is set if the certificate is revoked, and unset otherwise.
5. *Accusation info*: The first 32 bits of this 64-bit field contains the certificate serial number of a node that accused peer i of misbehavior. The remaining 32 bits contain the date that the accusation was made.

If field 3 does not contain zero, the profile table continues with the certificate status and accusation info fields; and if there are more than one accusers, it continues with 97-bit blocks containing information about the other accusers. Fig. 4 illustrates the fields of a profile table.

The protocol requires each node to keep track of the following variables, the values of which are obtained from its profile table:

- *Number of accusations against node* (i) (A_i): This is the total number of accusations made against a given node i. When a node receives an authenticated accusation against node i, it updates its profile table, and consequently this variable, if and only if both node i and the accuser certificates are not revoked and no previous accusation by the accuser against node i is recorded.

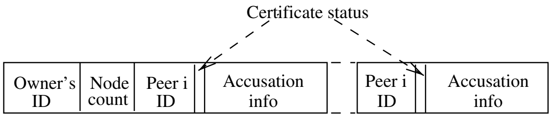

Fig. 4. Fields of a profile table.

- *Number of additional accusations made by node i* (α_i): When a node receives authenticated accusation info from node *i*, it updates its profile table and consequently this variable, if and only if the certificates of both node *i* and the node that is being accused of misbehavior (node *j*) are not revoked and no previous accusation by node *i* against node *j* is recorded. A node is not charged for the first accusation it makes; hence, α_i is actually the total number of accusations node *i* made minus one.
- *Behavior index of node i* (β_i): The behavior index (β_i) of a node *i* is a measure of the trustworthiness of the node *i*. β_i is a real number such that $0 \leqslant \beta_i \leqslant 1$. The greater the value of β_i, the more trustworthy node *i* is perceived to be. β_i is computed as follows:

  $$\beta_i = 1 - \lambda A_i, \tag{1}$$

  where $\lambda = \frac{1}{2N-3}$ and N is the number of nodes in the network.
- *Weight of node i accusation* (ω_i): This is a quantitative value that is assigned to the weight of a node's accusation. It depends on the behavior index of the node and on the number of accusations the node made. ω_i is a real number such that $0 \leqslant \omega_i \leqslant 1$. It is calculated as follows:

  $$\omega_i = \beta_i - \lambda\alpha_i, \tag{2}$$

  where λ is as indicated above.
- *Revocation quotient* (R_j): This real number determines whether the certificate for node *j* should be revoked. A certificate is revoked if R_j is greater than or equal to the revocation quotient threshold R_T. R_T is a configurable parameter whose value depends on the sensitivity of the security requirement. Typical values of R_T are $\frac{1}{2}$, $\frac{1}{3}$ *or* $\frac{1}{4}$. R_j can be computed as follows:

  $$R_j = \sum_{i=1}^{N} \sigma_{ij}\omega_i, \tag{3}$$

  where $\sigma_{ij} = 1$ if node *i* launched a complain against node *j*, and 0 otherwise.
- *Certificate status* (C_j): Indicates whether or not the certificate of node *j* is revoked. As indicated above, a certificate is revoked if $R_j \geqslant R_T$.

4.1. Determining the number of nodes in the network

MANETs are dynamic in nature: nodes may join and leave the networks on frequent basis. Consequently, the number of nodes N in any given MANET will likely not be constant. Our revocation scheme uses the mechanism outlined below for determining the number of nodes in the network at any given time. As outlined earlier, when a node enters a MANET, it is required to broadcast its certificate and the y_n value of its hash chain to all the network nodes. Upon receiving the broadcast, the peers are expected to unicast their certificates along with their hash chains y_n values to the new node. The certificates and the y_n values can be stored using any appropriate data structure. However, our protocol stipulates that each certificate entry should contain a field for storing an associated date. The date, including the time, that the certificate was received should initially be stored in this field.

After broadcasting its certificate, each node is required to broadcast short messages containing its certificate serial number and the date and time that the message was sent, at a configurable time interval of T minutes. The value of T depends on the frequency of the change in the network membership. We called these messages, membership confirmation messages. For message origin authentication and content integrity checks, a MAC of the message should be generated—using an agreed upon secure keyed hashing function and the hash chain value (with the highest subscript) that has not been previously used, as the key—and appended to the message. When a node receives a membership confirmation message m_i, from a node *j*, it stores it in memory or in a temporary file. The next membership confirmation message or accusation info message from node *j*, should contains the y_i value that was used to compute the MAC for the previous message (m_i) from the source. The node should first verify that the y_i value is authentic by ascertaining whether the hash of y_i equals the last previously revealed hash chain value of the source; that is, whether $y_{i+1} = H(y_i)$. If it is authentic, it computes the MAC of the message m_i using y_i as the key; if the MAC is identical to that which was appended to m_i, the node updates the date field associated with the certificate entry for node *j*, with the date indicated in m_i. It should be noted that, as explained in Section 4.2 below, the protocol does not require time synchronization.

If a node does not receive a verified authenticated membership confirmation message from any given node within $1.5T$ min, the certificate entry for the node in question, should be deleted from the node's

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

certificate repository. The number of entries in the certificate repository for any given node, should therefore closely reflect the actual number of nodes in the network.

4.2. Security mechanism

The messages our certificate revocation protocol exchange can be categorized as follows:

1. *Initialization messages*: These messages are sent when there is a new entrant to the MANET. A new entrant broadcasts its digital certificate and its y_n value to the nodes in the network; the MANET nodes in return unicast their y_n values and profile tables to the new entrant. The protocol requires a digital signature scheme for authenticating the y_n values and the profile tables.
2. *Membership confirmation and accusation info messages*: The majority of the messages the protocol exchanges fall in this category. For efficiency considerations, we utilized hash chains for verifying the integrity and authenticity of these messages.

After a node j broadcast its certificate and its hash chain y_n value to its network peers, the next membership confirmation or accusation info message m_i it sends, it uses its hash chain y_{n-1} value to compute a MAC for m_i and appends it to m_i before sending the message. Node j then appends its y_{n-1} value to the next membership confirmation or accusation info message m_{i+1} it sends and in turn uses y_{n-2} to generate a MAC for m_{i+1}. On receiving m_i from node j, the recipients need to wait until they receive m_{i+1} from node j before they can verify the authenticity and integrity of m_i. Membership confirmation messages are sent every T min; T is a configurable parameter. As outlined in Section 4.1, an accusation messages can be sent at anytime. Therefore a node should not have to wait for more than T min to authenticate any given message. If a node does not receive the hash chain value required to verify the authenticity and integrity of a message m_i within $1.5T$ min, the node is required to discard m_i. Time synchronization is not required because the time interval T is a local parameter and as shown below in Section 5.1, it is not necessary to have global consensus on precisely when this interval starts or ends.

5. Discussion

Our certificate revocation scheme allows MANETs' nodes to revoke the certificates of malicious or misbehaving nodes; in so doing the malicious or misbehaving nodes are effectively isolated from a given MANET. The scheme is designed so as to prevent malicious nodes from being able to use wrongful accusations to cause the revocation of the certificates of well-behaving nodes. We elaborate on this issue further in Section 5.1.

The certificate revocation scheme provides a methodology of quantifying the trustworthiness of MANETs' nodes based on the behavior profiles of the nodes. The value of a node's trustworthiness determines the weight of its accusation. The weight of node n_i accusations, depends on the number of accusations made against node n_i, as well as the number of accusations node n_i made. If a number of accusations is made against a node, it is likely that this node in question is malicious or misbehaving. Similarly, if a node made a large number of accusations, particularly if the accusations are not supported by other nodes, it is also likely that this node is malicious. A node is not charged for the first accusation it made. Additionally, when the certificate of a node n_j is revoked, all the nodes that accused node n_j of misbehavior will have one subtracted from the individual total of the number of accusations they made. Similarly, when the certificate of a node n_j is revoked, one is subtracted from the individual total of the number of accusations against all the nodes that node n_j accused of misbehavior. In so doing, the nodes are not permanently charged for legitimate accusations they made; likewise, they are not permanently charged for accusations malicious nodes made against them.

The underline principle of the scheme is that the weight of a node's accusation should be exactly zero if the behavior index (trustworthiness) of the node is the minimum possible value and the node made the maximum number of accusations that is allowed. The maximum number of accusations which can be made against any given node is $N-1$ where N is the number of nodes in the network. Therefore the minimum value for β_i is $1-\lambda(N-1)$. As indicated above, for fairness considerations, a node is not charged for the first accusation it made; hence the maximum number of accusations that any given node can be charged for is $N-2$. Consequently, $\omega_i = 0$ when $A_i = N-1$ and $\alpha_i = N-2$, that is, $\omega_i = 1-\lambda(N-1)-\lambda(N-2) = 0$. So the normalization variable λ, which ensures that the behavior index (β_i) is always within the range of zero and one inclusively, irrespective of the value of N, is equal to $\frac{1}{2N-3}$.

Our revocation scheme requires that new entrants to a MANET be sent the profile tables of the existing members of the MANET. This is necessary to ensure that the newcomers have up-to-date information about the behavior profile of the current members of the MANET. Unlike accusation info and membership confirmation messages, which use message authentication code (MAC) for message origin and integrity checks, profile table messages are authenticated with signatures. The use of signatures eliminate the delay in authenticating the message, in that the recipient of the profile tables do not have to wait for the release of hash chain values to authenticate the message. Profile tables are unicast only when new entrants enter a network; therefore the generation and verification of signatures for profile table messages should have minimal effect on the overall performance of the protocol.

As outlined in Section 3, our certificate revocation scheme utilizes the self-healing community approach presented in [45] for forwarding packets. This approach provides redundancy, in that if a malicious node drops a packet it is expected to forward, a well-behaving node in the community can detect the malicious activity and provide the service of forwarding the packet. If there is no well-behaving node in a self-healing community, adversarial agents may succeed in preventing accusation info from reaching certain nodes. Consequently there may be variations in the profile tables. In cases where there are variations, the new entrant is expected to fill the fields of its profile table with the values in the respective fields of the majority of the profile tables. This may result in differences in the computed β_i, ω_i and R_i values. Hence a certificate may not be revoked on all nodes instantaneously; however within negligible time interval, the certificate of a malicious node should be revoked on enough nodes which participate in the protocol, such that the malicious nodes will be rendered ineffective in perpetuating their adversarial behaviors.

The protocol does not require the cooperation of all nodes in a MANET. Malicious or misbehaving nodes may not adhere to the protocol; furthermore they may attempt to thwart the protocol by not forwarding accusation and membership confirmation messages. There are strong motivations though for well-behaving nodes to participate, since it is within their interest to help eliminate malicious or misbehaving nodes from the network.

5.1. Security analysis

In this section, we analyze the security of our certificate revocation protocol using a game-theoretic approach. In the game, the goals of the adversaries are (i) to disrupt the protocol by preventing accusation info and membership confirmation messages from non-adversarial nodes from reaching their destinations; (ii) prevent the revocation of their certificates; and (iii) cause the revocation of certificates of well-behaving nodes. Whereas the goal of the well-behaving nodes is to revoke the certificates of malicious entities and consequently isolate them from the network. We show below that the probability of adversarial nodes achieving their goals is very low.

Security properties

If the number of well-behaving nodes (k) is sufficiently large, that is, $k \geqslant \frac{2+\sqrt{4+8R_T(2N-3)}}{4}$, where R_T is the revocation quotient threshold and N is the number of nodes in the network, then the protocol is

I. resistant to adversarial attacks;
II. effective in revoking the certificates of adversarial nodes.

Proof sketch of Property (I). The proof utilizes the attack scenarios outlined below to show the following:

1. the effectiveness of the hash chain security mechanism;
2. at least R_T malicious entities are required to cause the revocation of the certificate of a well-behaving node;
3. the probability of malicious nodes succeeding in filtering messages from well-behaving nodes is very small.

1a. *As outlined in Section 4.2 above, there is a delay in verifying the authenticity and integrity of accusation info and membership confirmation messages because the recipients of the messages need to wait until they receive the hash chain values for computing the MAC for the given messages. One possible attack malicious nodes can mount as a result of the delay in verifying the authenticity of a message, is to delay forwarding a message* m_i *until it receives* *the message m_{i+1} which contains the key for computing the MAC for m_i; then modifies m_i and uses the key revealed in m_{i+1} to generate a new MAC for the modified m_i ($\hat{m}_i$), appends it to $\hat{m}_i$, then forwards the modified message.*

If there are functional self-healing communities,[4] the message m_i should get to its destinations before the modified message $\hat{m}_i$. The protocol necessitates that a given y_i hash chain value cannot be used more than once. Therefore on seeing $\hat{m}_i$ been authenticated with the same hash chain value as that utilized to ascertain the authenticity of the previously received m_i, the recipient will discard the modified message $\hat{m}_i$; consequently the attack will not succeed.

1b. *Malicious nodes impersonate other nodes and use the spoofed identities to launch accusations against well-behaving nodes.*

If a malicious entity M spoofed the identity of node j, then prior to sending any accusation message using node j identity, M must prevent membership conformation and accusation messages from j from reaching well-behaving nodes. This is necessary since, as explained in item (1a) above, a hash chain value can only be used once for authenticating a message. If there are functional self-healing communities, this attack will not succeed.

2. *Adversarial entities act in collusion, target one well-behaving node at a time and launch accusations against the targeted node in efforts to cause the revocation of its certificate.*

As outlined in the heuristic argument below, this attack is only possible if the number of malicious nodes is greater than or equal to the revocation quotient threshold R_T. If we assume the worst case scenario where no accusation is made against any of the malicious nodes and the weight of the accusations (ω_i) of each of the malicious nodes is at the maximum value possible; if no accusation is made against any of the malicious nodes, then based on Eq. (1) in Section 4, $\beta_i = 1$ for each of the malicious nodes; and since $\omega_i = 1$ (maximum value), then each of the malicious nodes made only one accusation, which is directed at the victim they targeted (node j). If there are m malicious nodes, based on Eq. (3) in Section 4, $R_j = m\omega_i$, that is, $R_j = m$. A certificate is revoked if $R_j \geqslant R_T$. Therefore if the malicious nodes are to succeed in causing the revocation of a certificate, the minimum requirement is that m must be equal to R_T. If anything other than the worst case scenario is assumed, that is, accusation(s) is/are made against any of the malicious nodes, or any of the malicious nodes made more than one accusations, then m must be greater than R_T for the malicious nodes to succeed in revoking the certificate of a well-behaving node.

3. *Adversarial entities act in collusion and create non-functional self-healing communities; consequently isolate targeted nodes from the rest of the network.*

If colluding adversarial entities form self-healing communities which contain no well-behaving node, they can essentially partition the network and isolate targeted nodes. If this occurs, the adversarial entities can reduce the effectiveness of the protocol; for example, if one or more well-behaving node(s) is/are isolated from the rest of the network, it is possible that the number of un-isolated well-behaving nodes may be less than the number of malicious nodes. If this were to occur, a key assumption on which the protocol is based would not be satisfied. It should be noted however that non-transient non-functional self-healing communities are unlikely considering that malicious nodes typically cannot restrict the movement of non-compromised nodes. Additionally, Kong et al. [45] shows that the probability that an expected area of a self-healing community, $E(A_{\text{heal}})$, contains k honest nodes is given by

$$\Pr[y = k] = \int\int_{E(A_{\text{heal}})} \frac{((1-\theta)\rho_L)^k}{k!} \mathrm{e}^{-(1-\theta)\rho_L}\, \mathrm{d}A,$$

where y is a random variable for the number of honest nodes, L is the number of nodes, θ is the proportion of malicious nodes, and ρ_L is the node density function, which is dependent on the location in space. If $k = 0$, that is, if there are no well-behaving nodes in a self-healing community, this probability becomes

$$\Pr[y = k] = \int\int_{E(A_{\text{heal}})} \mathrm{e}^{-(1-\theta)\rho_L}\, \mathrm{d}A,$$

which is small since the value of the function $\mathrm{e}^{-(1-\theta)\rho_L}$ is small.

Hence, non-transient, non-functional self-healing communities are unlikely. Consequently, the probability of adversarial entities succeeding in filtering messages from well-behaving nodes is low; therefore, by (1a), (1b) and (2) above the protocol is resistant to adversarial attacks. ☐

---

[4] We outline the consequences of non-functional self-healing communities below.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

Proof of Property (ii). Next, we show that the protocol is effective in revoking the certificates of malicious nodes. Recall that from (3) above, non-functional self-healing communities are unlikely.

If there are no non-functional self-healing communities, the following show that malicious entities in a MANET are incapable of preventing the revocation of their certificates provided that the number of well-behaving nodes (k) is greater than or equal to $\frac{2+\sqrt{4+8R_T(2N-3)}}{4}$, where R_T is the revocation quotient threshold and N is the number of nodes in the network. Assume the worst case scenario where each of the $N-k$ malicious nodes made an accusation against each of the k well-behaving nodes. Based on Eq. (1) in Section 4, the behavior index (β_i) for each of the well-behaving nodes would be $\beta_i = 1-\lambda(N-k) = 1-\frac{N-k}{2N-3} = \frac{N+k-3}{2N-3}$. Also, assume that each of the well-behaving nodes made an accusation against each of the $N-k$ malicious nodes; then based on Eq. (2) in Section 4, $\omega_i = \frac{N+k-3}{2N-3} - \left(\frac{N-k-1}{2N-3}\right) = \frac{2k-2}{2N-3}$.

By Eq. (3), the certificate of any misbehaving node j, is revoked if $R_j = k\frac{2k-2}{2N-3} \geqslant R_T$, which implies that $2k^2-2k-R_T(2N-3) \geqslant 0$; that is, $k \geqslant \frac{2+\sqrt{4+8R_T(2N-3)}}{4}$. □

Example. Consider a MANET with 100 nodes, if $R_T = \frac{100}{2}$ then $k \geqslant 70.68$; if $R_T = \frac{100}{3}$, $k \geqslant 57.80$ or if $R_T = \frac{100}{4}$, $k \geqslant 50.13$. These values of k are for the worst case scenario where the malicious nodes choose to accuse all the well-behaving nodes of misbehavior and in so doing, increase the probability of they been more speedily identified as being malicious. If anything other than the worst case is assumed, the values for k would be smaller, that is, a smaller number of well-behaving nodes would be necessary to guarantee that identified malicious nodes are incapable of preventing the revocation of their certificates.

5.2. Computation and communication overhead

Every network security scheme has some associated computation and communication overhead. Our certificate revocation scheme mainly uses message authentication code (MAC)—which can be computed very efficiently—for message origin and integrity checks. Digital signatures are utilized only for authenticating profile table messages and hash chain y_n values when new hash chains are computed. Profile table messages are sent very infrequently: only when a new node enters the MANET; and if the hash chains are made long enough, one or two hash chains per node, that is, one or two y_n value(s) per network session should suffice. Therefore the signing and verification of signatures for profile table messages and y_n hash chain values should have limited effect on the performance of the certificate revocation scheme owing to the infrequency with which these operations occur.

The communication overhead depends on the total number of nodes N in the MANET, the number of misbehaving or malicious nodes, and the value of the configurable time interval T mentioned in Section 4.1. The data the protocol transmit are the profile table and the certificate of each node whenever a new node enters the network. Additionally, each node sends a 64-bit membership confirmation message, plus the 128 or 160-bit MAC every T min, which accounts for bandwidth utilization of approximately $3.4 * N * T$ bits/s. The bandwidth utilize for the broadcast of accusation info depends on the number of malicious or misbehaving nodes in the network.

5.3. Communication complexity

In this section we derive the communication complexity of our certificate revocation protocol. We are interested in knowing how many accusation info messages are required to revoke a certificate. The computation is simple in the case where there is only one adversarial node, say node j. If a well-behaving node i is accused by the adversary, then $A_i = 1$, $\alpha_i = 0$, $\beta_i = 1-\lambda$ and $\omega_i = 1-\lambda$ (recall from Section 4 that A_i is the total number of accusations made against node i, α_i is the number of accusations (minus 1) made by node i, β_i is the behavior index and ω_i is the weight of node i accusation). Similarly, based on Eq. (3) in Section 4, $R_j = \sum_{i\neq j}\omega_i$, since $\sigma_{ij} = 1$. If a malicious node j makes n accusations against the nodes in the set $\mathscr{N}$, then we need N' nodes to accuse node j of misbehavior. Therefore

$$
\begin{aligned}
R_j &= \sum_{i\in\mathscr{N}}\omega_i + \sum_{i\notin\mathscr{N}}\omega_i = a(1-\lambda) + (N'-1-a) \\
&= N'-1-\lambda a \geqslant R_T.
\end{aligned}
$$

Hence, node j certificate is revoked if $N' \geqslant 1+\lambda a + R_T$. In the general case, there is a set $\mathscr{A}$ of $K \leqslant N/2$ adversarial nodes. Let α_{ij} denotes the number of accusations (minus 1) made by well-behaving

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

node i after accusing an adversarial node j. As is the case for the single adversarial node (outlined above), to revoke the certificate of one adversarial node, we need N' such that:

$$\begin{aligned} R_j &= \sum_{i \notin \mathscr{A}, i \leqslant N'} (1 - \lambda A_i - \lambda \alpha_{ij}) \\ &= N' - K - \lambda \sum_{i \leqslant N'} A_i - \sum_{i \notin \mathscr{A}, i \leqslant N'} \alpha_{ij} \geqslant R_T. \end{aligned}$$

The above is obtained by combining Eqs. (1)–(3) in Section 4.

The minimum N' required is

$$N' = K + \lambda \sum_{i \leqslant N'} A_i + \sum_{i \notin \mathscr{A}, i \leqslant N'} \alpha_{ij} + R_T. \tag{4}$$

Since the well-behaving nodes make accusations in random order, we compute the expected value of N'. There are K adversarial nodes such that $K < N/2$, therefore:

$$\sum_{i \leqslant N'} A_i \leqslant (N - K)K \leqslant \frac{N}{2}(N - 1). \tag{5}$$

Since we do not know the total number of accusations that a well-behaving node i will make, we approximate the expected value of α_{ij} to be $\frac{K}{2}$, which is half of the maximum number of accusations it can make, that is:

$$E\left[\sum_{i \notin \mathscr{A}, i \leqslant N'} \alpha_{ij}\right] \approx E[N'] \cdot \frac{K}{2}. \tag{6}$$

Solving for expected value of N' by substituting Eqs. (5) and (6) into (4), we obtain:

$$\begin{aligned} E[N'] &\leqslant \frac{1}{1 - \lambda K/2}\left[K + \lambda \frac{N}{2}(N - 1) + R_T\right] \\ &\leqslant \frac{1}{1 - \frac{1}{4(2 - 3/N)}}\left[\frac{N}{2}\left(1 + \frac{1 - 1/N}{2 - 3/N}\right) + R_T\right] \\ &\approx \text{linear in } N, \end{aligned}$$

where $\lambda = 1/(2N - 3)$.

This implies that a linear number of accusation info broadcasts (which cost order N^2 messages) are sufficient to revoke the certificate of an adversarial node.

6. Simulation setup and results

We simulated the protocol using NS2 network simulator. The aim of the simulation is to determine average case performances of the scheme with regards to its effectiveness in revoking the certificates of identified malicious nodes; and in particular to ascertain the average number of accusations necessary to cause the revocation of certificates for various combinations of number of well-behaving nodes verses number of malicious nodes. The process of identifying malicious nodes is beyond the scope of this paper; however, techniques such as those employed in [51,52] can be utilized. For the purpose of the simulation, we assumed that if a malicious node m_i made less than $\frac{N}{4}$ accusations (where N is the total number of nodes in the network), there is a probability of 0.50 that a given well-behaving node n_j will identify m_i as being malicious when n_j receives an accusation message from m_i; whereas if m_i made more than $\frac{N}{4}$ accusations, there is a probability of 0.75 that n_j will identify m_i as being malicious when n_j receives m_i accusation.

The simulation attempts to balance the following desires of the malicious nodes: (a) Prevent the revocation of their certificates by reducing the weight of the accusations of well-behaving nodes through malicious accusations. (b) Act in collusion with other malicious nodes and cause the revocation of well-behaving nodes' certificates by maliciously accusing targeted nodes. These two eventualities require different approaches. The former is best achieved if each of the malicious nodes launches accusation against all of the well-behaving nodes; whereas the latter needs conservatism regarding the number of accusations a node makes (see Eqs. (1) and (2) in Section 4). We used the following simple heuristic for achieving a balance between these conflicting requirements: When a malicious node m_i receives a message from a well-behaving node n_j, if m_i has not previously accused n_j of misbehavior and m_i made less than $\frac{N}{4}$ accusations and the output from a random number generator (which outputs 0 or 1) is 0, then m_i broadcasts an accusation against n_j. In other words, there is a 0.50 probability that a malicious node m_i will accuse a well-behaving node n_j of misbehavior whenever m_i receives a message from n_j; provided that m_i has not previously accused n_j, and m_i made less than $\frac{N}{4}$ accusations. If m_i however made more than $\frac{N}{4}$ accusations and all else being equal, then the probability that m_i launches an accusation against n_j—when it receives a message from the latter—decreases to 0.25. On the other hand, when a well-behaving node n_i receives an accusation message from a malicious node m_j, if n_i has not previously accused m_j, and m_j made less than $\frac{N}{4}$ accusations, there is a probability of 0.50 that n_i broadcasts an accusation against m_j. Whereas the

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

probability increases to 0.75 if m_j made more than $\frac{N}{4}$ accusations. Regarding the collusion aspect of the malicious nodes, when a malicious node m_i receives an accusation against a well-behaving node n_j from another malicious node, if m_i has not previously accused n_j of misbehavior, m_i immediately launches an accusation against n_j. In so doing, malicious nodes can effectively target non-malicious nodes in attempt to blackmail them and cause the revocation of their certificates.

We simulated a MANET environment running destination sequence distance vector (DSDV) as the routing protocol, and examined the performance of our certificate revocation scheme when the number of malicious nodes varies from 5 to x, where x is less than the revocation quotient threshold (R_T), for R_T values of $\frac{N}{2}$, $\frac{N}{3}$ and $\frac{N}{4}$ when N (number of nodes) equals to 100, 75 and 50.

As expected from intuition, the simulation results indicate that generally, as the number of malicious nodes increases, a slightly larger number of accusations are required to cause the revocation of a malicious node's certificate. The exception being when R_T equals $\frac{N}{4}$ for larger values of N, as is the case for N equals 100 (Fig. 5) and N equals 75 (Fig. 6). Fig. 5 for example, shows that when R_T equals 25.00, only 26 accusations are necessary to cause the revocation of a malicious node's certificate, irrespective of the number of malicious nodes (M) present, as M varies from 5 to 24. The lack of influence of the malicious nodes in this regard can be attributed to the following: with $R_T = \frac{N}{4}$ and the number of malicious nodes being less than R_T, the ratio of well-behaving nodes to malicious nodes is higher as the value of N increases. For example, when N equals to 100, the ratio of well-behaving nodes to malicious nodes (M) ranges from 19 to 3 when M varies from 5 to R_T; whereas when N equals 50, this ratio ranges from 9 to 3 as M varies from 5 to R_T. For lower R_T values, higher ratio of well-behaving to malicious nodes has the effect of diluting the influence of the malicious nodes, since smaller percentages of the available well-behaving nodes are sufficient to cause the revocation of a malicious node's certificate (Fig. 7).

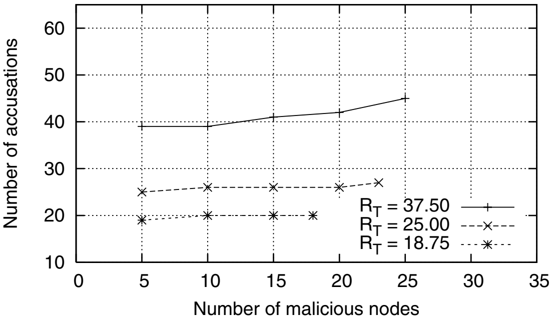

Fig. 6. Simulation results for 75 nodes.

Another deviation in the results from what is expected from intuition is the higher than average increase in the number of accusations required to revoke a certificate when the number of malicious nodes increases from 25 to 30 or from 20 to 25 for N equals 100 or 75 respectively, when R_T equals $\frac{N}{2}$. This can be attributed to the accumulative effect of the increasing number of malicious nodes. Higher R_T values necessitate larger number of accusations to cause the revocation of a certificate. The malicious nodes therefore have more opportunity to accuse well-behaving nodes before their certificates are revoked. Consequently for higher R_T values, as the number of malicious nodes increases, their effect becomes more pronounced.

In summary, the simulation results indicate that the number of accusations in excess of R_T that is necessary to cause the revocation of a malicious

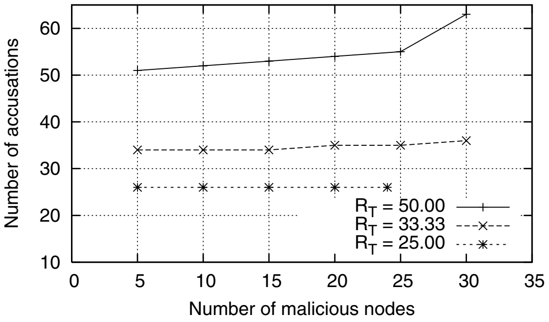

Fig. 5. Simulation results for 100 nodes.

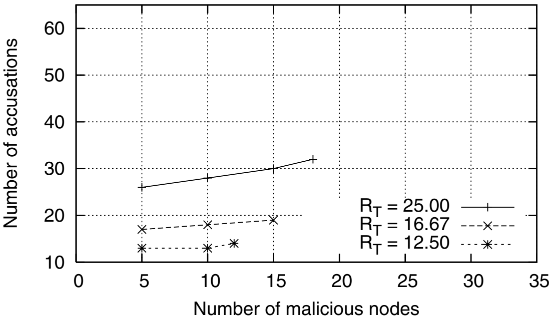

Fig. 7. Simulation results for 50 nodes.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

node's certificate depends on the size of the network (N) and the value of R_T. For lower R_T values, that is, for $R_T \leqslant \frac{N}{3}$, the effect of increasing number of malicious nodes is less pronounced as the size of N increases. However when R_T is greater than $\frac{N}{3}$, the effect of increasing number of malicious nodes is more pronounced for larger networks. In this regard, the simulation results show that when $R_T \leqslant \frac{N}{3}$, $\lceil R_T \rceil + 4$ accusations are sufficient to cause the revocation of a malicious node's certificate irrespective of the number of malicious nodes (k) in the network, provided that $k < R_T$; whereas, when $R_T > \frac{N}{3}$, as many as $\lceil R_T \rceil + 10$ accusations may be required to cause the revocation of a malicious node's certificate. In light of these results, it may be advantageous for R_T to be less than or equal to $\frac{N}{3}$, provided that the number of malicious nodes (k) in the network is expected to be less than this value. If the latter cannot be guaranteed, then R_T should be increased such that it is always greater than k.

7. Conclusion

In this paper, we presented a decentralized certificate revocation scheme which utilizes certificates that are based on the hierarchical trust model. Our scheme delegates all key management tasks—except the issuing of certificates—to the nodes in a MANET; and it does not require any access to on-line certificate authorities (CAs).

Our certificate revocation scheme is based on weighted accusations; whereby a quantitative value is assigned to an accusation to determine its weight. The weight of the accusations from nodes that are considered to be trustworthy are higher than those from less trustworthy nodes. A certificate of a node is revoked when the sum of the weighted accusations against the node is equal to or greater than a configurable threshold (R_T). The scheme mainly uses hash chains for providing data origin and integrity checks and it does not require time synchronization.

We outlined four possible attacks malicious entities can launch against our certificate revocation protocol and examine how the protocol deals with these adversarial activities. We presented communication complexity analysis which shows that order N^2 accusation info messages are sufficient to cause the revocation of a malicious node certificate. Finally, the simulation results indicate that when malicious nodes are identified, their certificates are speedily revoked in such a way that the nodes in the network are cognizant of the certificates revocation information in a timely manner.

Acknowledgments

The authors are grateful to the reviewers for their thoughtful comments and for the several improvements they suggested. This work is supported in part by a research grant from Laboratoire Universitaire Bell.

References

[1] F. Stajano, R.J. Anderson, The resurrecting duckling: security issues for ad-hoc wireless networks, in: Proceedings of the 7th International Workshop on Security Protocols, 2000, pp. 172–194.

[2] N. Shankar, D. Balfanz, Enabling secure ad-hoc communication using contextaware security services, in: Proceedings of Workshop on Security in Ubiquitous Computing (4 UBICOMP), 2002.

[3] A. Shamir, How to share a secret? Communications of the ACM 22 (11) (1979) 612–613.

[4] B. Chor, S. Goldwasse, S. Micali, B. Awerbuch, Verifiable secret sharing and achieving simultaneity in the presence of faults, in: Proceedings of 26th IEEE Annual Symposium on the Foundations of Computer Science (FOCS), 1985, pp. 383–395.

[5] P. Feldman, A practical scheme for non-interactive verifiable secret sharing, in: Proceedings of 28th IEEE Symposium on Foundations of Foundations of Computer Science, 1987, pp. 427–437.

[6] T.P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in: Proceedings of Crypto'91, LNCS, vol. 576, Springer-Verlag, 1991, pp. 129–140.

[7] V. Shoup, Practical threshold signatures, in: Proceedings of Eurocrypt 2000, LNCS, vol. 1807, Springer-Verlag, 2000, pp. 207–220.

[8] R. Gennaro, S. Jarecki, H. Krawczyk, T. Rabin, Robust threshold DSS signatures, in: Proceedings of Eurocrypt'96, LNCS, vol. 1070, Springer-Verlag, 1996, pp. 354–371.

[9] L. Zhou, Z.J. Haas, Securing ad hoc networks, IEEE Network Magazine 13 (6) (1999) 24–30.

[10] J. Kong, H. Luo, K. Xu, D.L. Gu, M. Gerla, S. Lu, Adaptive security for multi-layer ad-hoc networks, in: Special Issue of Wireless Communications and Mobile Computing, Wiley Interscience Press, 2002.

[11] B. Lehane, L. Doyle, D. O'Mahony, Shared rsa key generation in a mobile ad hoc network, in: Proceedings of IEEE Military Communications Conference (MILCOM 2003), 2003, pp. 814–819.

[12] A. Khalili, J. Katz, W.A. Arbaugh, Toward secure key distribution in truly ad-hoc networks, in: Proceedings of 2003 Symposium on Applications and the Internet Workshops, 2003, pp. 342–346.

[13] S. Yi, R. Kravits, Composite key management for ad hoc networks, in: Proceedings of the First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (MOBIQUITOUS 2004), 2004, pp. 52–61.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

[14] G. Xu, L. Iftode, Locality driven key management architecture for mobile ad-hoc networks, in: Proceedings for the 1st IEEE International Conference on Mobile Ad-hoc and Sensor Systems, 2004.

[15] H. Luo, P. Zerfos, J. Kong, S. Lu, L. Zhang, Self-securing ad hoc wireless networks, in: Proceedings of the Seventh International Symposium on Computers and Communications (ISCC'02), 2002, pp. 567–574.

[16] J. Kong, P. Zerfos, H. Luo, S. Lu, L. Zhang, Providing robust and ubiquitous security support for mobile ad hoc networks, in: Proceedings of the 9th International Conference on Network Protocols (ICNP), 2001, pp. 251–260.

[17] S. Chokhani, W. Ford, R. Sabett, C. Merrill, Internet X.509 public key infrastructure certificate policy and certification practices framework, Internet Request for Comments (RFC 3647), November 2003.

[18] P. Zimmermann, The Official PGP User's Guide, MIT Press, 1995.

[19] S. Capkun, L. Buttyan, J.-P. Hubaux, Self-organized public-key management for mobile ad hoc networks, IEEE Transactions on Mobile Computing 2 (1) (2003) 52–64.

[20] J.-P. Hubaux, L. Buttyan, S. Capkun, The quest for security in mobile ad hoc networks, in: Proceedings of ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2001), 2001, pp. 146–155.

[21] R. Housley, W. Polk, W. Ford, D. Solo, Internet X.509 public key infrastructure certificate and certificate revocation list (CRL) profile, Internet Request for Comments (RFC 3280), April 2002.

[22] M. Myers, R. Ankney, A. Malpani, S. Galperin, C. Adams, X.509 internet public key infrastructure online certificate status protocol – OCSP, Internet Request for Comments (RFC 2560), June 1999.

[23] L. Venkatraman, D.P. Agrawal, A novel authentication scheme for ad hoc networks, in: Proceedings of IEEE Wireless Communications and Networking Conference (WCNC), vol. 3, 2000, pp. 1268–1273.

[24] A. Weimerskirch, D. Westhoff, Identity certified authentication for ad-hoc networks, in: Proceedings of the 1st ACM workshop on Security of ad hoc and sensor networks (SASN'03), 2003, pp. 33–40.

[25] T.S. Messerges, J. Cukier, T.A.M. Kevenaar, L. Puhl, R. Struik, E. Callaway, A security design for a general purpose, self-organizing, multihop ad hoc wireless network, in: Proceedings of the 1st ACM Workshop on Security of Ad hoc and Sensor Networks, 2003, pp. 1–11.

[26] S.L. Keoh, E. Lupu, M. Sloman, PEACE: a policy-based establishment of ad-hoc communities, in: Proceedings of the 20th Annual Computer Security Applications Conference (ACSAC), 2004, pp. 386–395.

[27] M.C. Morogan, S. Muftic, Certificate management in ad hoc networks, in: Symposium on Applications and the Internet Workshops (SAINT 2003), 2003, pp. 337–341.

[28] R.R.S. Verma, D. O'Mahony, H. Tewari, Progressive authentication in ad hoc networks, in: Proceedings of the Fifth European Wireless Conference, 2004.

[29] C. Candolin, H. Kari, A security architecture for wireless ad hoc networks, in: Proceedings of IEEE Milcom 2002, 2002.

[30] C. Crépeau, C.R. Davis, A certificate revocation scheme for wireless ad hoc networks, in: Proceedings of 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN 2003), 2003, pp. 54–61.

[31] C.R. Davis, A localized trust management scheme for ad hoc networks, in: 3rd International Conference on Networking (ICN'04), 2004, pp. 671–675.

[32] S. Buchegger, J. Le Boudec, Performance analysis of the CONFIDANT protocol, in: Proceedings of the 3rd ACM International Symposium on Mobile Ad hoc Networking and Computing (MobiHoc'02), 2002, pp. 226–236.

[33] P. Resnick, K. Kuwabara, R. Zeckhauser, E. Friedman, Reputation systems, Communications of the ACM 43 (12) (2000) 45–48.

[34] J. Schneider, G. Kortuem, J. Jager, S. Fickas, Z. Segall, Disseminating trust information in wearable communities, in: Proceedings of 2nd International Symposium on Handheld and Ubitquitous Computing, 2000.

[35] L. Page, S. Brin, R. Motwani, T. Winograd, The pagerank citation ranking: bringing order to the web, in: 7th International World Wide Web Conference (WWW Consortium), 1998, pp. 161–172.

[36] S.D. Kamvar, M.T. Schlosser, H. Garcia-Molina, The eigentrust algorithm for reputation management in p2p networks, in: Proceedings of the Twelfth International World Wide Web Conference, 2003.

[37] L. Xiong, L. Liu, A reputation-based trust model for peer-to-peer ecommerce communities, in: IEEE Conference on ECommerce (CEC'03), 2003, pp. 275–284.

[38] M. Gupta, P. Judge, M. Ammar, A reputation system for peer-to-peer networks, in: Proceedings of ACM 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video, 2003.

[39] C.Y. Liau, X. Zhou, S. Bressan, K.-L. Tan, Efficient distributed reputation scheme for peer-to-peer systems, in: Proceedings of the 2nd International Human.Society@Internet, 2003, pp. 54–63.

[40] T.D. Huynh, N.R. Jennings, N.R. Shadbolt, Fire: an integrated trust and reputation model for open multi-agent systems, in: Proceedings of the 16th European Conference on Artificial Intelligence (ECAI), 2004, pp. 18–20.

[41] B. Yu, M.P. Singh, A social mechanism of reputation management in electronic communities, in: Proceedings of the 4th International Workshop on Cooperative Information Agents, 2000, pp. 154–165.

[42] A. Abdul-Rahman, S. Hailes, Supporting trust in virtual communities, in: Proceedings of Hawaii International Conference on System Sciences HICSS, 2000.

[43] K. Aberer, Z. Despotovic, Managing trust in a peer-2-peer information system, in: Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM01), 2001, pp. 310–317.

[44] Q. Zhang, T. Yu, K. Irwin, A classification scheme for trust functions in reputation-based trust management, in: Proceedings of ISWC Workshop on Trust, Security, and Reputation on the Semantic Web, 2004.

[45] J. Kong, X. Hong, Y. Yi, J.-S. Park, J. Liu, M. Gerla, A secure ad-hoc routing approach using localized self-healing communities, in: Proceedings of the 6th ACM International Symposium on Mobile Ad hoc Networking and Computing (MobiHoc'05), 2005, pp. 254–265.

[46] L. Lamport, Password authentication with insecure communication, Communications of the ACM 24 (11) (1981) 770–772.

[47] N.I. of Standards, Technology, Secure hash standard, Federal Information Processing Standards Publications (FIPS PUBS) 180-1, April 1995.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003

[48] R.L. Rivest, The md5 message-digest algorithm, Internet Request for Comments (RFC 1321), April 1992.

[49] H. Krawczyk, M. Bellare, R. Canetti, Hmac: Keyed-hashing for message authentication, Internet Request for Comments (RFC 2104), February 1997.

[50] A. Perrig, R. Canetti, D. Tygar, D. Song, The tesla broadcast authentication protocol, Cryptobytes (RSA Laboratories, Summer/Fall 2002) 5 (2) (2002) 2–13.

[51] Y. Huang, W. Lee, A cooperative intrusion detection system for ad hoc networks, in: Proceedings of ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN'03), 2003, pp. 135–147.

[52] H. Deng, Q.-A. Zeng, D.P. Agrawal, Svm-based intrusion detection system for wireless ad hoc networks, in: Proceedings of the IEEE 58th Vehicular Technology Conference (VTC 2003-Fall), 2003, pp. 2147–2151.

Geneviève Arboit is a Ph.D. student in the School of Computer Science at McGill University, Montréal, Canada. She works under the supervision of Professor Crépeau. Her research interests are coding theory, cryptology and discrete mathematics. Her current work focuses on backdoors on public key generation.

Claude Crépeau is an Associate Professor in the School of Computer Science at McGill University. He received an M.Sc. degree from the Université de Montréal in 1986 and a Ph.D. degree from M.I.T., in 1990. He served from 1991 to 1995 as an Associate Editor of the Journal of Cryptology. He later served as the Associate Editor for Complexity and Cryptography of the IEEE Transactions on Information Theory from 1995 to 1997. He has also served several times on the program committees of the "Crypto" and "Eurocrypt" conferences. He has worked extensively at the design of cryptographic protocols, including Zero-knowledge protocols, Multiparty Computations, Two-Party Secure Function Evaluation. His major contribution has been to offer alternative (non-computational) assumptions under which such protocols may be implemented using noisy channels and quantum channels.

Carlton R. Davis is a Ph.D. student in the School of Computer Science at McGill University, Montréal, Canada. He works under the supervision of Professor Crépeau and Professor Maheswaran. His research interests are system and network security, with special emphasis on mobile ad hoc network (MANET) security. His current work focuses on the design of security protocols for MANETs.

Muthucumaru Maheswaran is an Assistant Professor in the School of Computer Science at McGill University. In 1990, he received a B.Sc. degree in Electrical and Electronic engineering from the University of Peradeniya, Sri Lanka. He received an MSEE degree in 1994 and a Ph.D. degree in 1998, both from the School of Electrical and Computer Engineering at Purdue University. His research interests include resource management systems, trust management systems, public resource based utility computing systems, toolkits for teaching and learning computer networks, and resource discovery networks. He has authored or coauthored over 50 papers in these and related areas.

Please cite this article as: Geneviève Arboit et al., A localized certificate revocation scheme for ..., Ad Hoc Networks (2006), doi:10.1016/j.adhoc.2006.07.003