

An Attack on the Proactive RSA Signature Scheme in the URSA Ad Hoc Network Access Control Protocol

Stanisław Jarecki, Nitesh Saxena, and Jeong Hyun Yi

School of Information and Computer Science
UC Irvine, CA 92697, USA
{stasio, nitesh, jhyi}@ics.uci.edu

ABSTRACT

Recently, Luo, et al. in a series of papers [17, 14, 13, 18, 15] proposed a set of protocols for providing ubiquitous and robust access control [URSA] in mobile ad hoc networks without relying on a centralized authority. The URSA protocol relies on the new proactive RSA signature scheme, which allows members in an ad hoc group to make access control decisions in a distributed manner. The proposed proactive RSA signature scheme is assumed secure as long as no more than an allowed threshold of participating members is simultaneously corrupted at any point in the lifetime of the scheme.

In this paper we show an attack on this proposed proactive RSA scheme, in which an admissible threshold of malicious group members can completely recover the group RSA secret key in the course of the lifetime of this scheme. Our attack stems from the fact that the threshold signature protocol which is a part of this proactive RSA scheme leaks some seemingly innocuous information about the secret signature key. We show how the corrupted members can influence the execution of the scheme in such a way so that the slowly leaked information is used to reconstruct the entire shared secret.

Categories and Subject Descriptors

C.2 [Computer-Communications Networks]: Security and Protection; C.2.2 [Network Protocols]: Applications

General Terms

Security

Keywords

proactive cryptosystems, threshold cryptosystems, RSA, peer-to-peer, mobile ad-hoc networks, group membership, admission control

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SASN'04, October 25, 2004, Washington, DC, USA.
Copyright 2004 ACM 1-58113-972-1/04/0010 ...\$5.00.

1. INTRODUCTION: BACKGROUND AND MOTIVATION

Threshold and Proactive Signature Schemes. A (t, n) threshold signature scheme [6] enables any subgroup of $t + 1$ members in a group consisting of $n > t$ members, to collaboratively sign a message on behalf of that group. This is achieved by secret-sharing the signature key, e.g. the RSA secret key, among the group members, and allowing them to compute a signature on some message via a distributed protocol in which the members use the shares of the signature key instead of the key itself. The scheme is said to be *t-secure* if any coalition of at most t corrupt members is unable to forge a valid threshold signature on any message which honest members would not sign, and *t-robust* if honest group members can efficiently produce a valid signature even in the presence of at most t malicious members. To achieve *t*-security, a threshold signature scheme must in particular protect the secrecy of the signature key as long as no more than t of the group members are corrupt.

A *proactive* signature scheme [10], based on techniques of proactive secret sharing [20, 11], is a threshold signature scheme which remains secure and robust even if in every *time period*, called “share update interval”, a possibly different set of t group members is corrupted. This is achieved by the members periodically updating their shares of the secret signature key via a distributed share update protocol. Such an update protocol should destroy the correlation between secret shares learned by corrupted members in different time periods, so that the scheme can tolerate any number of corruptions throughout its lifetime as long as in any single time period the number of simultaneously corrupted members does not exceed t .

Application of Proactive Signatures to Peer-to-Peer Group Security. As pointed out by Zhou and Haas [25], proactive signature schemes can be used to implement group access control decisions without relying on a trusted and always accessible group “manager”, who makes all admission and revocation decisions on behalf of the group. In many mobile group settings, such manager may be often inaccessible to some subgroup of members. Moreover, in many applications placing all trust in a single entity creates a security liability.

In contrast, the idea of the proactive signature based access control mechanism for groups is that any large enough set of members can admit a new group member by using the threshold signature protocol to compute the new member's

certificate, and by using a slight variation of the proactive update protocol to give this member his share of the signature key. Similarly, once any large enough set of members agrees to revoke some existing member, the members collectively sign the revocation statement, and trigger the proactive share update protocol to leave the just revoked member without the current share of the signature secret.¹ Motivated by this idea and by the proposals of Luo, et al. (see below), Saxena et al. [22] implemented such access control protocol for ad hoc networks using the proactive DSS signature scheme of Gennaro et al. [9]. Recently, the same authors [23] examined the performance of a more efficient access control protocol based on the proactive BLS signature scheme [4] of Boldyreva [2], which relies on elliptic curve cryptography.

However, the common operation of signature *verification* in DSS, BLS, and in all other discrete-log based signature schemes, is orders of magnitude more computationally intensive than the verification of RSA signatures. Therefore, an efficient provably secure proactive RSA signature scheme would offer an attractive alternative to the above systems.

Problems with Currently Known Provably Secure RSA Proactive Signature Schemes. Unfortunately, the most efficient currently known provably secure proactive RSA signature schemes, two schemes by Frankel et al. [7, 8] and a scheme by Rabin [21], are not easily applicable to securing access control in ad hoc mobile groups by the methods described above. The fundamental reason is that the arithmetic operations involved in the RSA signatures seem more difficult to securely distribute than the DSS or the BLS cryptosystems.² In particular, the proactive RSA scheme of [7] is practical only for small groups, while in the other two provably secure proactive RSA schemes known today [8, 21], the members participating in the threshold signature protocol need to reconstruct the secret shares of the group members that are currently inaccessible to them. In this way both protocols essentially equate a temporarily inaccessible group member with a corrupt one, whose secrets might just as well be reconstructed. This is not a desirable feature for a proactive signature scheme, but it can be tolerated in one of the main applications of such schemes, which is a fault-tolerant on-line certification authority.

Unfortunately, such proactive signature schemes cannot be applied to groups whose members are not always active and/or are not always connected with all other members.

¹Such an execution of the proactive share update protocol is necessary to disable the revoked member from participating in future access control decisions of the group. However, to improve the efficiency (although at a cost of decreased resilience to simultaneous corruptions), the members can defer the proactive update protocol and execute it to simultaneously revoke up to t members.

²The difficulty in distributing RSA signatures is caused by the seeming necessity to perform computations on the secret shares modulo $\phi(N) = (p-1)(q-1)$, where $N = pq$ is the RSA modulus. Performing operations modulo $\phi(N)$ is difficult because this number must stay secret for the distributed RSA scheme to be secure. Note that that revealing $\phi(N)$ enables anyone to immediately compute the RSA private key $d = e^{-1} \pmod{\phi(N)}$ from the RSA public key (e, N) . In contrast, both the standard DSS signatures and BLS signatures of [4] are based on variants of the discrete logarithm problem where all the moduli used in computations involving secret shares can be made public.

This is the case in peer-to-peer on-line communities, in mobile ad-hoc networks, in sensor networks, and in many other applications. In such settings we need to enable isolated but large enough subgroups of members to operate without reconstructing everyone else's secrets, because such reconstruction would be insecure in the presence of just two corrupted players, each one a member of a different disconnected subgroup.

Insecurity of the Proactive RSA Scheme Proposed in the URSA Ad-Hoc Network Access Control Protocol.

In an effort to mitigate the above problem of the known proactive RSA signatures, Luo, et al. [17, 14, 13, 18, 15] proposed a new proactive RSA scheme, geared towards providing a security service, called "URSA", to mobile ad hoc networks (MANETs). The original description of this proactive RSA scheme and the URSA application can be found in [17]. Subsequently both the proactive RSA scheme and URSA were described in [14, 13, 18], and most recently in a journal version [15]. The URSA proactive RSA scheme can be applicable to MANETs because it avoids the need to access all shares during the threshold signature protocol. This is because it relies solely on Shamir's *polynomial* secret sharing scheme [24], as opposed to resorting to an additional layer of *additive* secret sharing, as is done by the two most efficient provably secure proactive RSA schemes [8, 21] discussed above. The core of the URSA proactive RSA scheme is the so-called *t*-bounded *offsetting algorithm* which is used to reconstruct the RSA signature $m^d \pmod{N}$ from $t+1$ signature shares produced individually by the $t+1$ members participating in the signing protocol.

The first problem with this scheme was pointed out in [19]. Namely, contrary to what the authors of the proposal claimed, their scheme does not provide robustness in signature generation in the presence of t malicious members. Simply speaking, the robustness mechanisms proposed by the authors are faulty because they require certain verification equations to hold even though they in fact do not hold, because the equations involve computation in two different groups (see [19] for more details). Hence, the set of t malicious members can prevent the honest members from efficiently creating a valid signature. However, this robustness problem in the t -bounded proactive RSA scheme can be solved, for example if the scheme is amended by special-purpose zero-knowledge proof protocols for proving equality of discrete logarithms in two different groups, e.g. the proofs of Camenish and Michels [5]. Such proof protocols are not very fast, but their expense can be tolerated because they would need to be executed only in the (rare) case of a corrupted member providing an incorrect partial signature to other members.³

Since the t -robustness of this scheme can be ensured by the above modifications, there remains a question if the (modified) URSA proactive RSA scheme is *secure* against a coalition of corrupt t members whose goal is not to prevent members from issuing signatures but to learn the secret-shared RSA signature key and to be thus able to forge signatures on the group's behalf. The question is interesting because the proposed URSA proactive RSA scheme, amended

³This will be a rare occurrence because a malicious member behaving in such a manner would be detected by the honest players, and therefore would be subsequently revoked from the group.

as described above, would provide efficiency and functionality advantages over the best known provably secure proactive RSA schemes [7, 8, 21]. However, the answer turns out to be negative.⁴

Our Contribution: Explicit Attack on the Proactive RSA Scheme in the URSA Protocol. We demonstrate the insecurity of the URSA proactive RSA signature scheme by constructing an explicit attack in which the admissible group of t corrupted members colludes in the proactive protocol in such a way so that they reconstruct the whole RSA secret key d after a realistic number of runs of the proactive update protocol and the threshold signature protocol.

Our attack exploits the fact that the t -bounded offsetting threshold RSA signature protocol, which is employed in the URSA proactive RSA scheme, leaks certain seemingly innocuous information about the secret signature key. The information that the adversary learns about the secret key in a run of the signature protocol depends on the current sharing of the secret key and on which group of members participates in the protocol. While it is not clear how dangerous this released information is for a single secret sharing, in a *proactive* signature scheme the secret sharing is refreshed with every proactive update, and therefore the released information about the secret key can be different in each update interval. It turns out that the corrupted members can influence an execution of the *update* protocol in such a way that the executions of the *signature* protocol during the subsequent update interval will release information which is both new and correlated with the information the adversary has gained so far. Thus our attack can be seen as a simple search algorithm, where the information learned in a signature protocol tells the adversary which branch to pick next, and the proactive update protocol allows the adversary to pick that branch.

The attack poses a realistic threat. For example, for the threshold size $t = 7$ and for the RSA public key of $e = 65537$ (utilized in the implementation of URSA [14]) and 1024-bit RSA modulus N , our attack needs 163 executions of the proactive update protocol and 1148 runs of the signature protocol to succeed. The attack succeeds assuming that throughout these 163 update periods, the t corrupted players belong to the so-called “update group” of players which play an active role in the proactive update protocol.

However, the vulnerability of the URSA proactive RSA protocol is stronger than what is immediately implied by the above attack. First, our attack does not make use of all the information leaked in the signature protocol, thus it is quite possible that another attack, which does utilize all the available information succeeds in recovering the private RSA key even faster and/or succeeds in recovering the key even if, say, only a smaller subset of the special “update group” of players is corrupted. Moreover, even if the attack we describe is slowed down, for example by slowing down the rate of the proactive updates, it recovers $512 + (r - \log_{t+1}(e)) * \log_2(t + 1)$ most significant bits of d after $r > \log_{t+1}(e)$ update rounds, which gives $512 + 3(r - 5)$ MSB bits of d for the above $e = 65537$ and $t = 7$. Therefore our attack raises doubts about the security of the URSA proactive RSA scheme even for smaller number of rounds. It is hard to say

much about the security of this protocol, for example after $r = 34$ rounds, because while it is not currently known how to recover the whole RSA key knowing 600 of the MSBs of d , we also do not know any arguments that RSA remains secure with this side information about d revealed, and it would be rather surprising if such arguments existed.

Positive Results in a Related Work. Recently, two authors of this paper examined a related question of whether the URSA proactive RSA signature scheme can be fixed, and at what cost, to make a *provably secure* signature scheme [12]. It turns out that if polynomial secret-sharing in the URSA scheme is replaced with additive sharing in the first layer (with a second layer of polynomial sharing, as in the provably secure RSA schemes of Frankel et al. [8] and Tal Rabin [21]), the resulting scheme is actually provably secure, and it is factor of two or four more efficient for the signing players than the most efficient previously known scheme of Tal Rabin *as analyzed in the original Tal Rabin’s paper* [21]. However, in the same paper the authors also improved the security analysis of the [21] scheme by the same factors, and thus the two schemes appear now to be almost equivalent from efficiency point of view, although the new scheme is arguably simpler and more elegant [12].

However, it is important to note that because of the additive sharing, while the new scheme is applicable to on-line certification authorities or on-line time-stamping services, it is *not* applicable to group access control for groups like MANETs, on-line peer-to-peer groups, sensor nets, etc. Thus the efficient provably secure proactive RSA signature which would avoid additive sharing and be applicable in such contexts, remains an open problem.

Paper Organization. The rest of this paper is organized as follows: Section 2 summarizes the notation and introduces some definitions. Section 3 describes the URSA proactive RSA signature scheme. We then present an attack on this scheme in Section 4. Finally, we elaborate on the implications of our attack for the security of the URSA scheme in Section 5.

2. NOTATION AND SETTING

An Adversary: We assume the worst case where all corrupted members collude, and in fact all corruptions are simply scheduled and controlled by a single entity, called an “adversary”. An adversary against a proactive signature scheme, and thus also against a MANET group access control mechanism like URSA which utilizes a proactive signature scheme, is able to compromise any set of at most t members in the group in every *update round* (see below). After compromising a member, the adversary learns its corresponding secret share and can force this member to behave arbitrarily in the protocol.

Share Update “Round”, or “Interval”: The time period between any two consecutive share update procedures. In the URSA design, the share update protocols occur periodically, for example twice a day. The reason for period execution of this procedure is to re-randomize the secret sharing in such a way that all the currently non-revoked players receive new random shares of the same RSA secret key d . (We describe this protocol in detail in Section 3.3.)⁵

⁵We note that if the update protocol is triggered only peri-

⁴In particular, although Luo, et al. claim that their scheme is provably secure, the security proofs that appear in [17] are incorrect.

We will assume that in every share update round, there is a significant amount of signature requests. Our attack requires just t such signature requests per update round to recover the key at a maximum speed.

Table 1: Notation

M_i	group member with unique index i
ss_i	secret share of M_i
t	“adversarial threshold”, i.e.
	the number of tolerated corruptions
n	total number of members M_i
(N, e)	RSA public key
TD	trusted dealer
\mathcal{A}	the adversary
$ x $	number of bits in a binary representation of x
$lg(x)$	$\log_2(x)$

3. THE PROACTIVE RSA SIGNATURE SCHEME IN URSA

In this section we describe the proactive RSA signature scheme of [17, 14, 13, 18, 15] used in the URSA ad hoc network access control protocol. We will refer to this scheme as an “URSA proactive RSA signature scheme”. The description includes the system set-up, the “ t -bounded offsetting” threshold signature generation protocol, and the proactive share update protocol.

3.1 The Setup Procedure

A trusted dealer TD is involved in a one-time setup to bootstrap the system. The dealer is not required hereafter and in fact is assumed to vanish from the scene, or, equivalently, to erase his memory. TD generates the standard RSA private/public key pair, i.e. it picks two random primes p and q , sets $N = pq$, sets (e, N) as a public key where $\gcd(e, N) = 1$, and as a private key it sets a number $d < N$ s.t. $ed = 1 \pmod{\phi(N)}$.

Once the standard RSA key pair is chosen, TD secret-shares the RSA secret key d using a slight modification of Shamir secret sharing [24]. Namely, TD selects a random polynomial $f(z)$ over \mathbb{Z}_N of degree t , such that the group secret is $f(0) = d \pmod{N}$. Next, TD gives to each member M_i , for $i = 1, \dots, n$, a secret share $ss_i = f(i) \pmod{N}$. Notice that the secret d is shared over a composite modulus N as opposed to a prime modulus as in the original scheme of Shamir, but our attack does not depend on what modulus is used in the secret sharing.

3.2 The Threshold Signature Protocol

The goal of the threshold RSA signature protocol is to generate in a distributed manner an RSA signature $s = m^d$ odically as in the URSA design, this has a consequence that only up to t members can be meaningfully revoked from the group within an update interval, because a player is fully revoked only when his share of the secret key d becomes invalid, and that happens only when all the remaining non-revoked players perform a share update protocol. Alternatively, as suggested by [23], an update protocol can also be triggered *reactively*, as an immediate response in the case of revocation of an unusually large number of players.

(mod N) under the secret-shared key d . The URSA threshold RSA signature protocol consists of two phases: First each participating member creates its *partial signature* on the intended message and sends it to the signature recipient, and then the recipient locally reconstructs the RSA signature from these partial signatures.

Partial Signature Generation: Let G denote the set of identifiers of the $t + 1$ members in the group who participate in the threshold signature protocol. Using polynomial interpolation we can write the secret key d as

$$d = \sum_{j \in G} ss_j l_j^{(G)} \pmod{N}$$

where $l_j^{(G)} = \prod_{i \in G, i \neq j} \frac{(j-i)}{j-i}$ (mod N) Notice that $N = pq$ has only two very large factors, and therefore all the elements $(j - i)$ for $i, j \in G$ will have inverses modulo N . Each member M_j , for $j \in G$, outputs his partial signature $s_j^{(G)}$ on m as

$$s_j^{(G)} = m^{d_j^{(G)}} \pmod{N}, \text{ where } d_j^{(G)} = ss_j l_j^{(G)} \pmod{N} \quad (1)$$

Signature Reconstruction: On receiving $t + 1$ partial signatures $s_j^{(G)}$ from the $t + 1$ group members M_j in G , the signature recipient reconstructs the RSA signature s using the “ t -bounded-offsetting” algorithm which works as follows. Since $\sum_{j \in G} d_j^{(G)} = d \pmod{N}$ and $0 \leq d_j^{(G)} \leq N - 1$ for all j 's, therefore

$$d = \sum_{j \in G} d_j^{(G)} - \alpha^{(G)} N \pmod{N} \quad (\text{over the integers}) \quad (2)$$

for some integer $\alpha^{(G)} \in [0, t]$. Equation (2) implies that

$$s = m^d = \left(\prod_{j \in G} s_j^{(G)} \right) m^{-\alpha^{(G)} N} \pmod{N}$$

for some integer $\alpha^{(G)} \in [0, t]$. Since there can be at most $t + 1$ possible values of $\alpha^{(G)}$, the signature recipient can recover $s = m^d \pmod{N}$ by trying each of the $t + 1$ possible values $Y_\alpha = Y(m^{-N})^\alpha \pmod{N}$ for $Y = \prod_{j \in G} s_j^{(G)}$ and $\alpha = 0, \dots, t$, and returning $s = Y_\alpha$ if $(Y_\alpha)^e = m \pmod{N}$. The most significant cost factor in this procedure in an exponentiation $m^{-N} \pmod{N}$, and therefore the computational cost of the URSA threshold RSA signature protocol for each of the signers and for the recipient is about one full (1024 bit) exponentiation modulo N .

Remark: It is important to note that the above t -bounded offsetting threshold RSA signature algorithm **reveals the value of $\alpha^{(G)}$** , which, as will be described in section 4, leaks some information about the secret-shared private key d to an adversary who corrupts t of the players participating in group G . This information leakage in fact exposes the whole proactive RSA scheme to a key-recovery attack.

3.3 The Proactive Share Update Protocol

The goal of the proactive share update protocol is to randomize the secret sharing of the private RSA key d held by the group members. This protocol was first proposed for both proactive secret sharing and for proactive cryptosystems and signature schemes by Herzberg et al. [11, 10]. The URSA proactive share update protocol is a variant of this protocol which is more efficient, especially in settings where

some players can be inactive or temporarily disconnected from others, as in MANETs.

The “Classic” Share Update Protocol. The proactive share update protocol of [11, 10] proceeds as follows (when sharing is done modulo N): Every member M_j chooses a random *partial update polynomial* $\delta_j(z)$ over \mathbb{Z}_N of degree t with the constant term being zero. The sum of these partial update polynomials defines the *update polynomial* $\delta(z) = \sum_{j=1, n} \delta_j(z) \pmod{N}$. Note that $\delta(0) = 0$. For each pair of members (M_j, M_i) , player M_j gives a share $\delta_j(i)$ of his partial update polynomial to M_i . Each member M_i then computes his new secret share (to be used in the threshold signature protocol in the subsequent update interval) as

$$ss_i = ss'_i + \sum_{j=1}^n \delta_j(i) = ss'_i + \delta(i) \pmod{N}$$

where ss'_i is M_i 's existing share, i.e. a share this player used in the previous interval. All the information pertaining to this protocol except of the new share ss_i is then erased. Note that if $ss'_i = f'(i) \pmod{N}$ for all i then the new secret-sharing polynomial $f(z)$ is defined as $f(z) = f'(z) + \delta(z) \pmod{N}$, and it is therefore a t -degree polynomial s.t. $f(0) = f'(0) = d \pmod{N}$. The scheme can be made robust against malicious participants using verifiable secret sharing techniques (see [11] for details).

However, it is important to notice that the new secret-sharing polynomial $f(z)$ is not necessarily a *random* t -degree polynomial s.t. $f(0) = d \pmod{N}$. This is because a corrupt player M_j can distribute its update polynomial $\delta_j(z)$ only after all the other corrupt players M_i see their shares of all the other update polynomials. In this way, the corrupt players can control the new secret-sharing polynomial $f(z)$ to some degree, by controlling the shares of $f(z)$ held by the corrupted players. In provably-secure proactive schemes that employ this proactive share update protocol, like the proactive DSS or BLS signatures [9, 2], this adversarial ability does not pose any harm. However, as we will see in section 4, this control ability means trouble for the URSA proactive RSA scheme since the information about shared secret d leaked in the threshold signature protocol depends precisely on the shares held by the corrupted players.

The URSA Two-Stage Modification of this Protocol. The URSA proactive RSA scheme utilizes a simple modification of the above share update protocol which improves the protocol's efficiency. This modification can indeed be used to speed up all proactive cryptosystems that use the [11] protocol, as the above mentioned schemes of [9, 2]. The URSA proactive share update protocol consists of two stages. The protocol relies on an existence of a designated group of players Ω , which we will call an “update group”, consisting of t group members.⁶ The first stage of the protocol proceeds exactly like the above protocol of Herzberg et al., except that only the players Ω participate in it. In other words, players $M_j \in \Omega$ create their update polynomials $\delta_j(z)$, send their shares $\delta_j(i)$ to other members $M_i \in \Omega$, and thus the players in Ω can be said to hold in secret-shared form the update

polynomial $\delta(z)$ defined as $\delta(z) = \sum_{j \in \Omega} \delta_j(z) \pmod{N}$. In the second stage of the URSA proactive share update protocol, the members of the update group Ω provide shares of this update polynomial $\delta(i)$ to *all* remaining (and non-revoked) group members $M_i \notin \Omega$. In this way all group members M_i will get their update share and can compute the new share $f(i) = f'(i) + \delta(i) \pmod{N}$ as before.

The URSA papers describe two protocols for how these $\delta(i)$ update shares are transferred from the Ω players to their final destinations. Even though the details of the second version of this protocol are a little unclear, these details do not affect the attack we describe in this paper. For completeness, we sketch the two variants as follows: In the first version, each M_i gets its $\delta(i)$ share by communicating with each of the players $M_i \in \Omega$ directly. The players in Ω jointly reconstruct the $\delta(i)$ value by first sharing masking random values among themselves. In the second version, either the $\delta(i)$ update shares or the whole $\delta(z)$ polynomial (this version is not very clear to us), appear to be distributed to the rest of the group members encrypted under the group public key (e, N) . Presumably, each member M_i reconstructs his share of this update polynomial $\delta(i)$ by contacting her $t + 1$ neighbors who either decrypt her encrypted share or decrypt the whole polynomial $\delta(z)$ and evaluate it at point i at the same time. It is not clear how this second version can be implemented securely, but the first version is standard and we see no vulnerabilities in it.

Choosing the Ω Update Group. We need to note here that the attack on the URSA proactive RSA scheme presented in this paper depends crucially on all of the t players in the above update group Ω to be the corrupted players. It is therefore important for the practical feasibility of the attack how this Ω group is decided. From the initial reply of the URSA authors to the attack presented in this paper [16], it appears that the details of how the Ω group is decided are not set in stone in the design of the URSA scheme. This is not surprising since the idea of modifying the Herzberg et al. protocol by delegating the update work to a smaller set of players Ω was introduced for the reasons of efficiency, not security.

The fixes proposed in [16], e.g. choosing the Ω group as the t players with smallest IDs, seem problematic: First, such players would need to be identified in some distributed protocol, and second, the resulting protocol would now be still under the attack, but only if the t players with lowest IDs are corrupted. Moreover, the resulting protocol would then need all these Ω players to be always present and connected, which goes against the URSA philosophy of providing group access control in environments where some players can become inactive and disconnected. Possibly, the tweak that would slow the attack we describe in this paper the most would choose the Ω group differently in every round. This tweak, however, has similar problems: It is not clear how to make sure that the Ω membership rotates without the global knowledge of the current membership list, which would reduce the applicability of the resulting protocol, and (2) the resulting protocol would again need the scheduled players to be up and connected at the right times. In our understanding, the attractive idea of the original philosophy of URSA was that the Ω group can be formed by *any* t players, and moreover that several such groups can be independently created in disconnected fragments of the network.

⁶In some URSA descriptions it seems that the Ω group needs to have $t + 1$ and not t members. We believe that t members is enough, and that the issue does not have a significant bearing on anything considered in this paper.

Unfortunately, while the ability for any group of t currently active and connected players to form the Ω update group would make the URSA scheme most reliable and attractive, because of the security vulnerability of this scheme which we describe in this paper, such freedom in choosing Ω would also lead to the fastest key-recovery attack on the resulting scheme. While it is still possible that there exists a smart tweak of the Ω -choosing process which would both slow down in practice the attack we describe here, and would not impact the applicability of the URSA scheme to the “on/off presence, on/off communication links” setting it targets, we believe that what is really needed is a replacement of the URSA proactive RSA scheme with a *provably secure* proactive signature scheme.

4. AN ATTACK ON THE URSA PROACTIVE RSA SCHEME

In this section we present an efficient key-recovery attack on the URSA proactive RSA signature scheme summarized in the previous section. The roots of this attack lie in the *t-bounded offsetting algorithm* which is the core of the URSA threshold signature protocol.

4.1 Overview of the Attack

The goal of the adversary in our attack is to recover the secret-shared private RSA key d . Consider an adversary \mathcal{A} who compromises t members. The full attack holds as long as these t members form the “update group” Ω (see section 3.3 above), and it holds regardless of what indices these players hold. However, for the sake of simplicity in the exposition, we will assume that the adversary corrupts members M_1, \dots, M_t throughout the lifetime of the scheme, and that these players also always form the Ω update group. We note, however, that our attack does *not* depend on the ability of the adversary to corrupt a *different* set of members every update period. This means that for example, as long as the Ω group is allowed not to change between the updates, the adversary can recover the private RSA key quite quickly if only he corrupts that subset and otherwise follows the protocol so that to avoid detection and revocation of these corrupted members from the group.

Information Leakage in the Signature Protocol. Assume that \mathcal{A} participates in the threshold signature protocol on some message in which the set of participating members G (see section 3.2 above) is made of all the corrupted members M_1, \dots, M_t and a single honest member M_p , for some $p \in [t+1, \dots, 2t]$. (Here too, the attack works for other players M_p , but we fix the above t values of p to simplify the presentation.) Let G_p represent the set of identifiers $\{1, \dots, t, p\}$ corresponding to the members participating in this run of the threshold signature protocol. By equation (2), the secret key d satisfies the following equation for some integer $\alpha^{(G_p)} \in [0, t]$:

$$d = \sum_{j \in G_p, j \neq p} d_j^{(G_p)} + d_p^{(G_p)} - \alpha^{(G_p)} N \quad (\text{over the integers})$$

Let us denote $S_p = \sum_{j \in G_p, j \neq p} d_j^{(G_p)}$ (over integers) and $D_p = S_p \pmod{N}$. Note that since \mathcal{A} knows ss_1, ss_2, \dots, ss_t , he can compute S_p and D_p .

By employing the reconstruction using the *t-bounded offsetting algorithm*, \mathcal{A} learns the value of $\alpha^{(G_p)}$ corresponding

to this signing group G_p . Now, note that from values $\alpha^{(G_p)}$ and S_p , the adversary also learns whether the shared secret $d \in [0, \dots, N-1]$ is less than or greater than D_p . This is because $S_p < \alpha^{(G_p)} N$ if and only if $d < D_p$; and therefore $S_p \geq \alpha^{(G_p)} N$ if and only if $d \geq D_p$.

Utilizing the Information Leakage to Recover the Key. At first sight, the information of whether the secret RSA key d is left or right of some value D_p in the $[0, \dots, N]$ range seems to provide only information on the few most significant bits of d . However, recall that over the lifespan of the system, the members update their secret shares by performing the proactive share update procedure. As we will see below, it turns out that during this procedure, as long as the “update group” Ω is formed by the corrupted players $\{M_1, \dots, M_t\}$, the adversary can choose the values of his new shares ss_1, ss_2, \dots, ss_t , which gives him complete freedom in specifying the resulting values D_p , for $p = t+1, \dots, 2t$, to be any values that he wants (we describe this process in subsections 4.2 and 4.3 below). Since in any subsequent run of the threshold signature protocol involving members M_1, \dots, M_t, M_p the adversary learns whether the secret d lies to the left or to the right of the corresponding value D_p (for $p = t+1, \dots, 2t$), the adversary can learn most about d if the chosen values D_{t+1}, \dots, D_{2t} divide the range $[0, N-1]$ into $t+1$ equally spaced intervals $\{[0, D_{t+1}-1], [D_{t+1}, D_{t+2}-1], \dots, [D_{2t}, N-1]\}$.

In this case, \mathcal{A} learns from t instances of the threshold signature protocol, for t different values $p = t+1, \dots, 2t$, whether d lies to the left or to the right of each of these D_p 's. Consequently \mathcal{A} shrinks the search interval for the secret d from $[0, N-1]$ to some interval $[D_p, D_{p+1}-1]$ which is smaller than the original interval by the factor of $t+1$. If the adversary repeats this attack recursively, then with every share update protocol his search range narrows by the factor of $t+1$. This is equivalent to saying that in every update interval, the adversary learns the $\lg(t+1)$ new most significant bits (MSBs) of the secret d . Therefore, this search procedure will end and the secret key d will be completely recovered after $\lceil \frac{|N|}{\lg(t+1)} \rceil$ share update rounds. We refer to this search procedure as a “ $(t+1)$ -ary search”.

Example with $t=1$: Consider a simple example with $t=1$. Assume that the adversary \mathcal{A} compromises member M_1 . Assume also that M_1 collaborates with member M_2 in a threshold signature protocol. The signing group is therefore $G_2 = \{1, 2\}$, which yields the following equation:

$$d = d_1^{(G_2)} + d_2^{(G_2)} - \alpha^{(G_2)} N$$

Here $S_2 = D_2 = d_1^{(G_2)}$. \mathcal{A} now employs the *t-bounded offsetting algorithm* and learns the value of $\alpha^{(G_2)}$. If $\alpha^{(G_2)} = 0$, \mathcal{A} learns that $d \geq d_1^{(G_2)} \pmod{N}$; otherwise if $\alpha^{(G_2)} = 1$, he learns that $d < d_1^{(G_2)} \pmod{N}$. Assuming that in the share update procedure \mathcal{A} can pick his secret share ss_1 so that the resulting $D_2 = d_1^{(G_2)}$ is whatever \mathcal{A} wants, \mathcal{A} can set $D_2 = \lceil \frac{N-1}{2} \rceil$. Then, with every share update round, \mathcal{A} halves the search interval, and thus he performs a *binary search* which recovers the secret d completely in $\lg(N) = \lceil \frac{|N|}{\lg(2)} \rceil$ rounds.

Attack Speed-ups: Since in many cases the RSA secret key d can be efficiently recovered once some number of the

most significant bits are recovered, the number of rounds in the attack can be further reduced. Moreover, for the commonly used small values of e , like $e = 3, 17$, or 65537 , the attack can be sped-up by the factor of two because the first few MSB bits of d enable \mathcal{A} to efficiently compute the first half of the MSBs of d . We describe such speed-up mechanisms in subsection 4.4 below.

4.2 Optimal Choice of New Secret Shares

In each share update protocol the adversary's goal is to set the t values D_{t+1}, \dots, D_{2t} which will hold in the subsequent update period in a manner described above. In order to do that, the adversary first solves for the optimal new secret shares ss_1, ss_2, \dots, ss_t which would result in values D_{t+1}, \dots, D_{2t} he desires, using the following system of t modular linear equations:

$$ss_1 l_1^{(G_{t+1})} + ss_2 l_2^{(G_{t+1})} + \dots + ss_t l_t^{(G_{t+1})} = D_{t+1} \pmod{N}$$

$$ss_1 l_1^{(G_{t+2})} + ss_2 l_2^{(G_{t+2})} + \dots + ss_t l_t^{(G_{t+2})} = D_{t+2} \pmod{N}$$

.....

.....

$$ss_1 l_1^{(G_{2t})} + ss_2 l_2^{(G_{2t})} + \dots + ss_t l_t^{(G_{2t})} = D_{2t} \pmod{N}$$

These equations are linearly independent if the following matrix \mathbb{L} is invertible:

$$\mathbb{L} = \begin{pmatrix} l_1^{(G_{t+1})} & l_2^{(G_{t+1})} & \dots & l_t^{(G_{t+1})} \\ l_1^{(G_{t+2})} & l_2^{(G_{t+2})} & \dots & l_t^{(G_{t+2})} \\ \dots & \dots & \dots & \dots \\ l_1^{(G_{2t})} & l_2^{(G_{2t})} & \dots & l_t^{(G_{2t})} \end{pmatrix} = (-1)^t \prod_{i=1}^t \left(\frac{\prod_{j=1, j \neq i}^t (-j)}{\prod_{j=1, j \neq i}^t (i-j)} \right) \begin{pmatrix} \frac{1}{t} & \frac{1}{t-1} & \dots & 1 \\ \frac{1}{t+1} & \frac{1}{t} & \dots & \frac{1}{2} \\ \dots & \dots & \dots & \dots \\ \frac{1}{2t-1} & \frac{1}{2t-2} & \dots & \frac{1}{t} \end{pmatrix}$$

Thus, by inverting the above matrix the adversary can compute the optimal secret share values ss_1, \dots, ss_t he needs in the next update round, in order to achieve his optimal values S_{t+1}, \dots, S_{2t} in the signature protocols performed in that update round. The adversary is now left with the task of forcing the proactive update protocol to actually arrive at these optimal secret shares ss_1, \dots, ss_t for the corrupted members M_1, \dots, M_t .

4.3 Adversarial Behavior in the Update

We show that as long as the adversary controls the players in the update group Ω , the adversary can easily influence the proactive share update protocol to arrive at the optimal shares ss_1, \dots, ss_n he computed above. In the case of larger Ω groups, the attack succeeds as long as the adversary corrupts t members in Ω , and as long as some of these members can "speak last" in the first phase of the URSA share update protocol (see section 3.3).

Let us describe the attack assuming the most general case that $|\Omega| \geq t$. Let $B \subseteq \Omega$ denotes the subset of t members corrupted by the adversary, and let $M_b \in B$ be the corrupted member who "speaks last" in the first phase of

the share update protocol. Since there is no established sequence or order in which the members in Ω take part in the secret share update procedure, the adversary can wait until each member in Ω except of M_b distributes their shares of the random update polynomials $\delta_j, j \in \Omega \setminus \{M_b\}$, before distributing the shares of his polynomial $\delta_b(z)$ as the last one. (If the order is somehow fixed, although it's not clear how it could be without heavy performance penalty for the protocol, the adversary would still win assuming that he corrupts the player who is entitled to speak last.) Recall that the update polynomial is equal to

$$\delta(z) = \sum_{j \in \Omega \setminus \{M_b\}} \delta_j(z) + \delta_b(z) \pmod{N}$$

and that the new shares of each members are computed as $ss_i = ss'_i + \delta(i)$ where ss'_i is the current share of M_i .

To fix the resulting shares of the corrupted members to come out as the optimal values ss_1, \dots, ss_t specified above, member M_b chooses his partial update polynomial $\delta_b(z)$ in such a way that the resulting update polynomial $\delta(z)$ satisfies $\delta(i) = ss_i - ss'_i \pmod{N}$ for $i = 1, \dots, t$. To do that, M_b sets values $\delta_b(i)$ for $i = 1, \dots, t$ as

$$\delta_b(i) = ss_i - ss'_i - \sum_{j \in \Omega \setminus \{M_b\}} \delta_j(i) \pmod{N}$$

The M_b player then interpolates these values to recover the $\delta_b(z)$ update polynomial he should use.

Importantly, note that this adversarial behavior is indistinguishable to outside observers from prescribed behavior an honest player exhibits in the protocol. The attack succeeds if M_b picks his partial update protocol in the above way instead of the proscribed way of picking this polynomial at random, but the difference cannot be observed by the honest players, and thus this attack would be undetected.

4.4 Speeding-up the Attack

By following the above attack procedure, in every update round \mathcal{A} learns new $lg(t+1)$ most significant bits (MSBs) of d . Assuming that \mathcal{A} needs to discover all the $|N|$ -bits of the RSA secret key d , \mathcal{A} needs $\lceil \frac{|N|}{lg(t+1)} \rceil$ update rounds, and t signature protocol instances within each update interval as described above, to complete the attack. However, there are several ways in which \mathcal{A} can speed up this search. First, we can assume that at least the last 40-bits of the secret d can be obtained by a brute-force search once all the other bits are found, because the candidate d can be efficiently tested given the public key (e, N) . This reduces the number of rounds in the attack to $\lceil \frac{1}{lg(t+1)} (|N| - 40) \rceil$. Secondly, we can speed up this search by making a simple observation about half of the MSBs of d for small e values, and by utilizing several known results regarding the security of the RSA cryptosystem under partial key exposure [3, 1]. Below we explain the speed up for small e 's, and we list the other applicable results and explain how they speed up our search algorithm. The graph in Figure (1) summarizes this discussion by showing the number of rounds required in the attack w.r.t the range of the public exponent e for 1024-bit RSA modulus N , taking as an example a threshold value $t = 7$.

THEOREM 1. *Given only the first $\log_2(e)$ MSBs of d , the first half of MSBs of d can be efficiently computed.*

PROOF. Note that $ed = 1 \pmod{\phi(N)}$ implies that $d = 1/e(1 + k\phi(N))$ for some integer $k = 1, \dots, e - 1$. Therefore,




Figure 1: Currently required # of proactive update rounds to recover d for a given value of $\log_2(e)$, assuming $|N| = 1024, t = 7$.

since $N - \phi(N) < \sqrt{N}$, it follows that $0 \leq \hat{d}_k - d < \sqrt{N}$ for $\hat{d}_k = \lfloor 1/e(1 + kN) \rfloor$ for one of the $e - 1$ choices of k . Note that the \hat{d}_k values can be publicly computed, and note that $\hat{d}_{k+1} - \hat{d}_k \approx N/e$ for every k , and therefore that the $\log_2(e)$ MSBs of d determine the appropriate k (and \hat{d}_k) value, and therefore, since $|\hat{d}_k - d| < \sqrt{N}$ for that k , they also determine the 512 MSBs of d . \square

The import of the above observation for our attack is very simple: Since the above choices of the $e - 1$ values $\hat{d}_1, \dots, \hat{d}_{e-1}$, are neatly spread in the $[0, N - 1]$ interval in distances of N/e apart from each other, in our attack based on the $(t + 1)$ -ary search, the adversary can identify the appropriate \hat{d}_k (and k) value, and thus recover the first $|N|/2$ MSBs of d by the above theorem, after just $\lceil \frac{\lg(e)}{\lg(t+1)} \rceil$ rounds of the share update protocol.

Therefore, for small e 's, the attack requires only

$$r > \frac{1}{\lg(t+1)} \left(\lg(e) + \frac{|N|}{2} - 40 \right) \quad (3)$$

rounds of share updates for the adversary to learn the whole secret d . This means that the current implementation of URSA which uses the well-known value $e = 65537$ [14] can be attacked in just 163 update rounds for a modest threshold of $t = 7$.

For larger values of e , the results of [3, 1] on the RSA key security with partial key exposure imply the following speed-ups in our RSA key recovery attack:

THEOREM 2. [3] *If e is a prime in the range $[2^m, 2^{m+1}]$, with $\frac{|N|}{4} \leq m \leq \frac{|N|}{2}$, then given m MSBs of d , there is a polynomial time algorithm to compute d .*

THEOREM 3. [3] *If e is in range $[2^m, 2^{m+1}]$ and is a product of at most r primes, with $\frac{|N|}{4} \leq m \leq \frac{|N|}{2}$, then given the factorization of e and m MSBs of d , there is a polynomial time algorithm to compute all of d .*

If e meets either of the above criteria, the number of rounds r required in our attack reduces to within the range

$$\left\lceil \frac{|N|}{4\lg(t+1)} \right\rceil \leq r \leq \left\lceil \frac{|N|}{2\lg(t+1)} \right\rceil$$

THEOREM 4. [1] *If e is in range $[N^{0.5}, N^{0.725}]$, then the number of MSBs needed to completely recover d is given by $\frac{|N|}{8}(3 + 2\alpha + \sqrt{36\alpha^2 + 12\alpha - 15})$ where $\alpha = \log_N(e)$.*

If e meets the above criteria, the number of rounds required to recover the secret key d is given by $\left\lceil \frac{|N|}{8\lg(t+1)}(3 + 2\alpha + \sqrt{36\alpha^2 + 12\alpha - 15}) \right\rceil$

5. DISCUSSION: EFFICIENCY OF THE ATTACK AND INSECURITY OF THE URSA SCHEME

While the above results for general e values are interesting, in practice people want to use RSA with small e values, like $e = 3, 17$, or 65537 (all of which are prime numbers of the form $2^i + 1$ for a small value of i , which makes the exponentiation $s^e \pmod{N}$ involved in the RSA signature verification take only $i + 1$ modular multiplications), and for these values our attack holds if (1) t members of the update group Ω are corrupted and one of the corrupted player speaks last, (2) if in every update interval some t chosen honest players M_p are coaxed into participating in a signature protocol (note that the adversary's attack does not depend on what message is used in this protocol!), and (3) if the system lasts for $r = 163$ update rounds, which given the twice a day rate of updates gives only two months.

We think that these are quite reasonable assumptions. There's certainly nothing in the adversarial model of a proactive signature scheme and of the URSA group access control scheme, which would disallow the adversary from satisfying each of the above criteria.

As we discussed in the last subsection of section 3.3, the attack can be probably slowed down in practice if some modifications are employed in the process of choosing the "update group" Ω , to make it harder for the adversary to corrupt this group, or, equivalently, to make it harder for the adversary to have whatever players he does corrupt be chosen as the Ω group. However, summing up the discussion from that section, it is not clear how to modify the protocol in order to make the adversary's success significantly harder, and at the same time not to severely limit the applicability of the URSA scheme.

Perhaps more importantly, the attack we exhibit shows that even if the adversary does not satisfy all the above criteria, the adversary still learns meaningful information about the RSA private key d , which makes the security of the resulting system doubtful. For example, as we discussed in the introduction, if the adversary successfully participates in just $r = 34$ instead of $r = 163$ update rounds (for 1024-bit N , $e = 65537$ and $t = 7$), the adversary will learn 600 most significant bits of d . Given the steady progress in the ability to recover the full d from partial knowledge [3, 1], there is little hope that such partial information can be shown not to weaken the RSA system. In other words, all bets are off.

Finally, our attack uses only very specific part of the information leaked in the URSA threshold RSA signature protocol. We showed that if the adversary corrupts t of the signing $t + 1$ players, and if the remaining $(t + 1)$ -st player M_p can be known beforehand, the leaked information is equivalent to whether the shared secret d lies to the left or to the right of some value D_p which can be computed and in fact controlled by the adversary. However, information about d leaks also if (1) other players M_p , $p \notin [t + 1, \dots, 2t]$ partici-

pate in the threshold signature protocol, and if (2) less than t corrupted players participate in this protocol. The information revealed about d in this cases is more complicated than the $d \ll D_p$ information we used in our attack, but it is nevertheless easy to define as well, and it can very well be used to either speed up the attack we propose even more, or to extend it to adversaries who (1) corrupt less than t of the players participating in the threshold signature protocol, or (2) corrupt less than t players in the update group Ω , or (3) fail to predict which honest players will participate in the threshold signature protocol.

6. REFERENCES

- [1] J. Blomer and A. May. New Partial Key Exposure Attacks on RSA. In D. Boneh, editor, *CRYPTO '03*, number 2729 in LNCS, pages 27–43. IACR, 2003.
- [2] A. Boldyreva. Efficient threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In *Proceedings of International Workshop on Practice and Theory in Public Key Cryptography*, volume 2567 of LNCS, pages 31–46, 2003.
- [3] D. Boneh, G. Durfee, and Y. Frankel. An attack on RSA given a small fraction of the private key bits. In *ASIACRYPT'98*, number 1514 in LNCS, pages 25–34, 1998.
- [4] D. Boneh, B. Lynn, and H. Shacham. Short Signatures from the Weil Pairing. In C. Boyd, editor, *ASIACRYPT'01*, number 2248 in LNCS, pages 514–532. IACR, 2001.
- [5] J. Camenisch and M. Michels. Separability and efficiency for generic group signature schemes. In *In Advances in Cryptology - CRYPTO '99*, volume 1666 of LNCS, pages 106–121, 1999.
- [6] Y. Desmedt and Y. Frankel. Threshold cryptosystems. In G. Brassard, editor, *CRYPTO '89*, number 435 in LNCS, pages 307–315. IACR, 1990.
- [7] Y. Frankel, P. Gemmell, P. D. MacKenzie, and M. Yung. Optimal-Resilience Proactive Public-Key Cryptosystems. In *Foundations of Computer Science FOCIS'97*, pages 384–393, 1997.
- [8] Y. Frankel, P. Gemmell, P. D. MacKenzie, and M. Yung. Proactive RSA. In *Proc. of Crypto'97*, pages 440–454, 1997.
- [9] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Robust Threshold DSS Signature. In M. Abadi, editor, *Information and Computation*, vol. 164 (1), pages 54–84. 2001.
- [10] A. Herzberg, M. Jakobsson, S. Jarecki, H. Krawczyk, and M. Yung. Proactive public key and signature systems. In *ACM Conference on Computers and Communication Security*, 1997.
- [11] A. Herzberg, S. Jarecki, H. Krawczyk, and M. Yung. Proactive Secret Sharing, Or How To Cope With Perpetual Leakage. In D. Coppersmith, editor, *CRYPTO '95*, number 963 in LNCS, pages 339–352. IACR, 1995.
- [12] S. Jarecki and N. Saxena. Further Simplifications in Proactive RSA Signature Schemes. In submission. Draft available from the authors and on <http://eprint.iacr.org>. August 2004.
- [13] J. Kong, H. Luo, K. Xu, D. L. Gu, M. Gerla, and S. Lu. Adaptive Security for Multi-level Ad-hoc Networks. In *Journal of Wireless Communications and Mobile Computing (WCMC)*, volume 2, pages 533–547, 2002.
- [14] J. Kong, P. Zerfos, H. Luo, S. Lu, and L. Zhang. Providing Robust and Ubiquitous Security Support for MANET. In *IEEE 9th International Conference on Network Protocols (ICNP)*, 2001.
- [15] H. Luo, J. Kong, P. Zerfos, S. Lu, and L. Zhang. URSA: Ubiquitous and Robust Access Control for Mobile Ad Hoc Networks, available on-line at <http://www.cs.ucla.edu/wing/publication/publication.html>. In *IEEE/ACM Transactions on Networking (ToN)*, to appear, Oct 2004.
- [16] Songwu Lu. Comments on Recent Advances in Cryptanalysis of URSA. A draft communicated to the authors by email by Songwu Lu, on August 16th, 2004.
- [17] H. Luo and S. Lu. Ubiquitous and Robust Authentication Services for Ad Hoc Wireless Networks. Technical Report TR-200030, Dept. of Computer Science, UCLA, 2000.
- [18] H. Luo, P. Zerfos, J. Kong, S. Lu, and L. Zhang. Self-securing Ad Hoc Wireless Networks. In *Seventh IEEE Symposium on Computers and Communications (ISCC '02)*, 2002.
- [19] M. Narasimha, G. Tsudik, and J. H. Yi. On the Utility of Distributed Cryptography in P2P and MANETs: The Case of Membership Control. In *IEEE 11th International Conference on Network Protocol (ICNP)*, pages 336–345, November 2003.
- [20] R. Ostrovsky and M. Yung. How to withstand mobile virus attacks. In *10th ACM Symp. on the Princ. of Distr. Comp.*, pages 51–61, 1991.
- [21] T. Rabin. A Simplified Approach to Threshold and Proactive RSA. In H. Krawczyk, editor, *CRYPTO '98*, number 1462 in LNCS, pages 89 – 104. IACR, 1998.
- [22] N. Saxena, G. Tsudik, and J. H. Yi. Admission Control in Peer-to-Peer: Design and Performance Evaluation. In *ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, pages 104–114, October 2003.
- [23] N. Saxena, G. Tsudik, and J. H. Yi. Identity-based Access Control for Ad Hoc Groups. In *Submission*, September 2004.
- [24] A. Shamir. How to Share a Secret. *Commun. ACM*, 22(11):612–613, Nov. 1979.
- [25] L. Zhou and Z. J. Haas. Securing Ad Hoc Networks. *IEEE Network Magazine*, 13(6):24–30, 1999.