

Bounded Vector Signatures and their Applications

Lei Wei
University of North Carolina
Chapel Hill, NC, USA
lwei@cs.unc.edu

Scott E. Coull
RedJack, LLC.
Silver Spring, MD, USA
scott.coull@redjack.com

Michael K. Reiter
University of North Carolina
Chapel Hill, NC, USA
reiter@cs.unc.edu

ABSTRACT

Although malleability is undesirable in traditional digital signatures, schemes with limited malleability properties enable interesting functionalities that may be impossible to obtain otherwise (*e.g.*, homomorphic signatures). In this paper, we introduce a new malleable signature scheme called *bounded vector signatures*. The proposed scheme allows a user to sign a multi-dimensional vector of values, along with a description of the context within which the vector should be interpreted. The scheme includes a unique malleability property, which we refer to as the *stretch property*, that allows the components of the signed vector to be increased up to a pre-defined limit without access to the signing key. Decreasing these values, however, remains computationally infeasible. We prove the security of our construction under the strong RSA and decisional Diffie-Hellman assumptions in the random oracle model. Finally, we underscore the utility of bounded vector signatures by discussing their use in distributed systems security applications.

Categories and Subject Descriptors

E.3 [Data]: Data Encryption; K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms

Security

Keywords

Algebraic Signatures, Malleable Signatures, Vectors

1. INTRODUCTION

Typically, digital signature schemes must remain secure against existential forgeries to achieve their purpose of verifying the authenticity of data. Therefore, it would seem as though any malleability allowed by a signature scheme would

be strictly harmful. However, several malleable signature schemes, including homomorphic signatures, have been proposed to achieve functionalities by allowing arbitrary parties to perform a limited set of operations to alter the signed data in logical ways. Transitive signatures [32, 6, 33, 23], for instance, allow a user with signatures on graph edges (x, y) and (y, z) to produce a third signature on the transitive closure, (x, z) , without access to the signing key.

Here, we propose a new type of malleable signature, called *bounded vector signatures*. Intuitively, our scheme allows a user to sign a vector of natural numbers, and includes a unique malleability property that allows components of these signed vectors to be increased, or *stretched*, up to a pre-defined bound without access to the signing key. Furthermore, each vector is associated with an arbitrary string, which we refer to as the *context*, that describes the way to interpret the values within the signed vector. To capture the security of our scheme, we define *shrink unforgeability*, which captures the inability of a computationally bounded adversary to decrease the value of any of the signed vector components.

The malleability property of our bounded vector signature scheme is particularly useful in multiparty settings, which are the primary focus of this paper. In such a setting, n signers may produce partial signatures on vectors, and any $t \leq n$ of those partial signatures can be efficiently combined to form a full signature on a vector representing all signers. In a typical, non-malleable signature scheme, these partially signed vectors would have to be exactly the same for the combiner to create a full signature. The malleability property of our bounded vector signatures, however, allows the signers to sign *different* vectors. In particular, the combiner receives t partial signatures on different vectors, applies the stretch operation to those signed vectors until all of them are the same, and then combines them as before.

This form of malleability leads to very efficient solutions to many problems that arise in the area of distributed systems security. In these distributed systems problems, information must be securely aggregated from multiple (potentially untrusted) parties while simultaneously limiting communications overhead and interaction among the parties. To that end, we show how our bounded vector signatures can be used to sign a number of interesting data structures commonly used in distributed environments, such as sets and intervals, with just a *small constant number* of elements (*e.g.*, at most four) of the underlying multiplicative group. Moreover, for each of these data structures we describe how the malleability of our bounded vector signatures enables us to

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.

ASIACCS '11, March 22–24, 2011, Hong Kong, China.
Copyright 2011 ACM 978-1-4503-0564-8/11/03 ...\$10.00.

non-interactively compute a full signature on exactly the intersection or union of given partially signed sets or intervals. We also discuss the application of these methods to distributed systems problems, including distributed cache reconciliation in content delivery networks [19, 9, 35], fault tolerant computation in sensor networks [31], and group key management [34, 42].

Related Work Malleable signatures, which can include homomorphic signature schemes and are sometimes referred to as algebraic signatures, were first discussed by Rivest during a series of talks at Cambridge University. The concept was later formalized by Johnson *et al.* [26], who described homomorphic signature schemes for redacted documents and set operations. Since then, new homomorphic signature schemes have been introduced, including transitive signature schemes on graphs [23, 6, 32, 33] and signatures for network coding [8, 22, 1]. Additionally, Kiltz *et al.* [28] introduced a malleable signature scheme that allows a user to append a message to a signature without requiring a signature on the appended message itself. Our bounded vector signatures, like the append-only signatures of Kiltz *et al.*, do not require a specific homomorphic property of the signature and instead rely upon a limited form of malleability. We note that the functionality allowed by our vector signatures is orthogonal to that provided by Kiltz *et al.*

The limited malleability allowed by our bounded vector signatures is achieved through an RSA-like signature construction, where the components of the vector being signed are encoded in the exponent of the signature, rather than the base. The first such construction was proposed by De Jonge and Chaum [16] in a study of several RSA signature variations. Gennaro *et al.* [21], and Hohenberger and Waters [25] have since used the construction, along with division-intractable and chameleon hashes, to build hash-and-sign signatures without random oracles. Perhaps the most closely related construction to our own is the dynamic accumulator scheme of Camenisch and Lysyanskaya [10], where identities are encoded as random primes in the exponent of the accumulator. None of these constructions, however, were developed with the goal of malleability in the form we pursue here, nor were they developed for use in multiparty settings typical of distributed systems problems.

2. PRELIMINARIES

To begin, we introduce bounded vector signatures, and define a new notion of unforgeability, called shrink unforgeability. Throughout the paper, we denote a d -dimensional vector of natural numbers as $\mathbf{v} = \langle v_1, \dots, v_d \rangle$, where the k^{th} component is denoted as $\mathbf{v}[k]$. Scalar values are denoted by non-bolded variables. Furthermore, we use the symbol $\overset{\$}{\leftarrow}$ to indicate sampling uniformly at random from a set of values, and the symbol \leftarrow to indicate the output of an algorithm.

2.1 Bounded Vector Signatures

A bounded vector signature scheme is a malleable signature on d -dimensional vectors wherein the components of the signed vector may be altered, within a predetermined bound, without access to the associated signing key. Specifically, we consider a signature on a pair (\mathbf{v}, c) , where $\mathbf{v} \in \mathbb{N}^d$ is a d -dimensional vector in the natural numbers and c is a context (*i.e.*, a text description) within which the vector is

interpreted. As an example, \mathbf{v} may represent distances in d independent dimensions and c may describe the starting point of the vector within that space. Furthermore, bounded vector signatures have a unique property that allows arbitrary parties to increase the values of any component of the signed vector up to a predefined limit without access to the signing key. For the remainder of this paper, we refer to this property as the *stretch property*.

Since we are most interested in distributed applications, we present a multiparty version of our scheme. The multiparty bounded vector signature scheme follows the typical threshold signature paradigm wherein n signers, which we denote as P_1, \dots, P_n , may produce partial signatures on vector-context pairs. Any subset of t distinct partial signatures can then be combined to produce a full signature that can be verified using a single verification key. Both the partial and full signatures in the multiparty version of our bounded vector signature scheme inherit the stretch property. Therefore, it is possible to combine partial signatures on different vectors by stretching them until they are identical, as long as their contexts are the same. In particular, the resultant vector in the full signature will be equal to (at least) the component-wise maximum from among all t vector signatures. For ease of exposition, we present the formal description of our scheme and its instantiation in a multiparty setting, and note that the single signer case may be trivially instantiated by setting $n = t = 1$.

More formally, a *bounded vector signature scheme* is defined by a tuple of algorithms $\mathcal{BVS} = (\text{KeyGen}, \text{Sign}, \text{Stretch}, \text{Combine}, \text{Verify})$ run by n signers P_1, \dots, P_n . The input-output specifications of these algorithms are as follows:

KeyGen($1^\kappa, t, n, \hat{\mathbf{v}}$): The key generation function takes as input a security parameter 1^κ , the threshold number t of partial signatures necessary to form a final signature, the number of signers n , and a d -dimensional vector $\hat{\mathbf{v}}$ containing the maximum value (*i.e.*, bound) for each dimension. The algorithm outputs signing keys sk_1, \dots, sk_n for each of the signers P_1, \dots, P_n , and a single verification key pk used to verify the final signatures.

Sign($sk_i, (\mathbf{v}, c)$): The signing algorithm takes in a signing key sk_i , along with a vector-context pair (\mathbf{v}, c) , where $0 \leq \mathbf{v}[k] \leq \hat{\mathbf{v}}[k]$ for each dimension k . The context describes how the verifier should interpret the signed value, such as the starting point of the vector. The signing algorithm outputs a partial signature on vector \mathbf{v} associated with context c for signer P_i , which we denote as σ_i .

Stretch($(\sigma, \mathbf{v}, c), k, a$): The stretch algorithm takes a partial or full signature σ on vector-context pair (\mathbf{v}, c) and an amount a to increment the value of the k^{th} component. The output of the algorithm is a new (partial or full) signature σ' on (\mathbf{v}', c) , where $\mathbf{v}'[k] = \min(\mathbf{v}[k] + a, \hat{\mathbf{v}}[k])$.

Combine($(\sigma_{i_1}, \mathbf{v}_{i_1}, c_{i_1}), \dots, (\sigma_{i_t}, \mathbf{v}_{i_t}, c_{i_t})$): The combine algorithm takes as input a set of t partial signatures on vector context pairs $(\mathbf{v}_{i_j}, c_{i_j})$ from distinct signers. These signatures must be associated with the same context c (*i.e.*, $c_{i_j} = c$), but may sign t different vectors. The algorithm outputs a full signature σ on the component-wise maximum values among the input signatures and their associated context (*i.e.*, (\mathbf{v}', c)), where $\mathbf{v}'[k] = \max(\mathbf{v}_{i_1}[k], \dots, \mathbf{v}_{i_t}[k])$ for all dimensions k .

$\text{Verify}(pk, \sigma, (\mathbf{v}', c'))$: The verify algorithm takes in the verification key pk , a full signature σ , and a vector-context pair (\mathbf{v}', c') that is to be verified as being signed. If the signature is valid for the vector-context pair, then the algorithm outputs 1. Otherwise, the algorithm outputs 0.

2.2 Security for Bounded Vector Signatures

Due to the malleable nature of our proposed signature scheme, the traditional definition of existential unforgeability under chosen message attack does not appropriately capture its security. Instead, we propose the notion of *shrink unforgeability* under chosen message attack. This definition captures the ability of untrusted parties to *increase*, but not decrease, the values of any component of the signed vector without access to the signing key. In the multiparty setting, this restriction translates to a limitation on the values within each component of the fully signed vector output from the Combine procedure. Specifically, the adversary should not be able to produce a full signature on a vector where the value in the k^{th} dimension is less than the t^{th} smallest value in that dimension from among all partial signatures given as input to the Combine algorithm. The intuition here is simply that the adversary would have to *decrease* at least one value in order to create a full signature on a vector containing values less than the t^{th} smallest in each dimension, which is disallowed by shrink unforgeability. This security notion is formalized as follows.

DEFINITION 1 (SHRINK UNFORGEABILITY). *Let \mathcal{BVS} be a bounded vector signature scheme, and let \mathcal{A} be an algorithm with oracle access to n signing oracles representing the signers P_1, \dots, P_n . Each query (\mathbf{v}, c) to oracle Oracle_{sk_i} returns $\text{Sign}(sk_i, (\mathbf{v}, c))$. For any context c and dimension k , let $\mathbf{v}_1^c[k], \dots, \mathbf{v}_n^c[k]$ denote the minimum values queried by \mathcal{A} with context c to the n oracles, listed in nondecreasing order. If \mathcal{A} made queries on context c to only $s < n$ oracles, then we let $\mathbf{v}_{s+1}^c[k] = \dots = \mathbf{v}_n^c[k] = \hat{\mathbf{v}}[k] + 1$. The experiment for shrink unforgeability under chosen message attack is then defined as follows:*

Experiment $\text{Exp}_{\mathcal{BVS}}^{\text{SU-CMA}}(\mathcal{A})$:

$(pk, sk_1, \dots, sk_n) \leftarrow \text{KeyGen}(1^\kappa, t, n, \hat{\mathbf{v}})$

$(\mathbf{v}', c', \sigma') \leftarrow \mathcal{A}^{\text{Oracle}_{sk_1}, \dots, \text{Oracle}_{sk_n}}(pk)$

Return 1 if $\text{Verify}(pk, \sigma', (\mathbf{v}', c')) = 1$ and either:

- (1) \mathcal{A} queried $(*, c')$ to less than t signing oracles, or
- (2) \mathcal{A} queried $(*, c')$ to at least t signing oracles and for some dimension k , $\mathbf{v}'[k] < \mathbf{v}_t^c[k]$.

Let $\text{Adv}_{\mathcal{BVS}}^{\text{SU-CMA}}(\mathcal{A}) = \Pr[\text{Exp}_{\mathcal{BVS}}^{\text{SU-CMA}}(\mathcal{A}) = 1]$, and define $\text{Adv}_{\mathcal{BVS}}^{\text{SU-CMA}}(T) = \max_{\mathcal{A}} \text{Adv}_{\mathcal{BVS}}^{\text{SU-CMA}}(\mathcal{A})$ where the maximum is taken over all adversaries \mathcal{A} running in time at most T .

We briefly explain the adversary's actions in this experiment. An adversary is provided with n partial signature oracles representing n signers. He can make partial signature queries on any (\mathbf{v}, c) pairs he chooses in any order and run the Combine algorithm to get full signatures from the partial signatures he received. In the end, he successfully breaks the shrink unforgeability if he outputs a valid signature on a pair (\mathbf{v}', c') and either of the two following cases occurred. The first case is when a forger \mathcal{A} makes signature queries of the form $(*, c')$ to less than t signing oracles. In

this case, he should not be able to generate full signatures on any pair of the form $(*, c')$, since he did not collect enough partial signatures. The second case is when \mathcal{A} made signature queries of the form $(*, c')$ to at least t signing oracles and there exists a dimension k such that $\mathbf{v}'[k] < \mathbf{v}_t^c[k]$. That is, $\mathbf{v}_t^c[k]$ is the smallest value in dimension k on which a signature should be possible for context c' .

2.3 Strong RSA Assumption

Here, we briefly recall the definition of the strong RSA assumption, which was first proposed by Barić and Pfitzmann [5]. In this paper we are primarily interested in the following variant of the strong RSA problem where the input is chosen as a random quadratic residue. It has been shown that this variant is not any easier than the standard strong RSA problem [13].

Let N be a positive integer that is the product of two distinct, random κ -bit primes p, q , let QR_N denote the quadratic residues in \mathbb{Z}_N^* , and let y be a random element of QR_N . Informally, given a random (N, y) , the Strong RSA assumption is that it is hard to compute a pair (x, e) such that $e > 1$ and $x^e \equiv y \pmod{N}$. Specifically we denote the advantage of an adversary \mathcal{B} in solving the strong RSA problem as

$$\text{Adv}_{\text{QR-SRSA}}(\mathcal{B}) = \Pr[e > 1 \text{ and } x^e \equiv y \pmod{N} : \\ p, q \xleftarrow{\$} \text{Primes} \cap \{0, 1\}^\kappa, N \leftarrow pq, \\ y \xleftarrow{\$} QR_N, (x, e) \leftarrow \mathcal{B}(N, y)]$$

and we define $\text{Adv}_{\text{QR-SRSA}}(T) = \max_{\mathcal{B}} \text{Adv}_{\text{QR-SRSA}}(\mathcal{B})$ where the maximum is taken over all adversaries \mathcal{B} executing in time at most T .

2.4 Decision Diffie-Hellman Assumption

Our proof of security also makes use of the decisional Diffie-Hellman (DDH) assumption in the quadratic residue group QR_N when N is the multiplication of two strong primes. That is, it is difficult for an adversary to distinguish a triple (g^a, g^b, g^{ab}) from (g^a, g^b, g^c) for a, b , and c chosen randomly and for a generator g of QR_N . In particular, we use a variant of the problem which says that DDH problem is hard even when the factorization of N is known. It has been shown that the DDH problem over QR_N does not depend on the hardness of factoring [27]. Formally, the DDH advantage of an adversary \mathcal{D} is defined as follows, where **SafePrimes** are the primes p of the form $p = 2p' + 1$ where p' is itself a prime:

$$\text{Adv}_{\text{DDH}}(\mathcal{D}) =$$

$$|\Pr[\mathcal{D}(p, q, g, g^a, g^b, g^{ab}) = 1 :$$

$$p, q \xleftarrow{\$} \text{SafePrimes} \cap \{0, 1\}^\kappa, N \leftarrow pq, g \xleftarrow{\$} QR_N,$$

$$a, b \xleftarrow{\$} \mathbb{Z}_{p'q'} \text{ where } p' = (p-1)/2, q' = (q-1)/2] -$$

$$\Pr[\mathcal{D}(p, q, g, g^a, g^b, g^c) = 1 :$$

$$p, q \xleftarrow{\$} \text{SafePrimes} \cap \{0, 1\}^\kappa, N \leftarrow pq, g \xleftarrow{\$} QR_N,$$

$$a, b, c \xleftarrow{\$} \mathbb{Z}_{p'q'} \text{ where } p' = (p-1)/2, q' = (q-1)/2] |$$

Let $\text{Adv}_{\text{DDH}}(T) = \max_{\mathcal{D}} \text{Adv}_{\text{DDH}}(\mathcal{D})$ where the maximum is taken over all adversaries \mathcal{D} executing in time at most T .

Below, we refer to a triple (g^a, g^b, g^{ab}) as a *valid* DDH triple, and (g^a, g^b, g^c) as an *invalid* DDH triple.

3. OUR CONSTRUCTION

In this section, we present a bounded vector signature construction in the multiparty setting using the threshold RSA signature scheme of Shoup [38] as the basis of our construction. The novelty of our construction lies in the observation that values being signed can be moved from the base of the signature to the exponent. By doing so, we can encode each value in such a way that arbitrary parties can increase this value by exponentiating the signature, but cannot decrease the value without being able to take roots in the RSA group. In particular, we encode the signature on the value in a vector dimension as the distance from that value to the maximum value for that dimension, and then embed that distance as a root in the RSA group. Thus, by exponentiating the signature by an exponent corresponding to this dimension, the distance in that dimension can be decreased (and so the value in that dimension is increased). The context portion of the signature acts as the base of the RSA signature and provides the meaning behind the values in the exponent. The context also ensures that vectors in different contexts cannot be combined.

3.1 Construction Description

Here we discuss the operation of each algorithm in our construction and provide the technical details in Figure 1.

The signature scheme is initialized by running the **KeyGen** algorithm, which generates the public verification key and the secret signing keys for each of the n signers. For the multiparty setting, we consider a trusted dealer¹ who generates two safe primes $p = 2p' + 1, q = 2q' + 1$, and a RSA modulus $N = pq$. Each dimension k in the vector space being signed is associated with a maximum value $\hat{v}[k]$ that is given as input to the algorithm and a distinct prime e_k chosen by the dealer. The dealer also computes the signing key $sk = \prod_{k=1}^d e_k^{-\hat{v}[k]+1} \bmod m$, where $m = p'q'$. The dealer then splits sk into n shares sk_1, \dots, sk_n using standard polynomial secret sharing techniques [37]. Essentially, this signing key encodes the maximum possible distance for each dimension in the index.

The **Sign** algorithm allows a signer with signing key sk_i to produce a partial signature on a d -dimensional vector. The signer first creates a context c , which is a text description of the way the verifier should interpret the signed vector. Practically speaking, the format of the context is application specific, but its general purpose in our construction is to ensure that vectors from different contexts cannot be combined into a full signature. The output of the **Sign** algorithm is a partial signature $\sigma_i = \mathcal{H}(c)^{n!sk_i \prod_{k=1}^d e_k^{\hat{v}[k]}} \bmod N$, where $\mathcal{H}(\cdot)$ is a hash into the group QR_N . The inclusion of $n!$ in the exponent is taken from the work of Shoup [38] to ensure that the Lagrange interpolation performed during the **Combine** algorithm produces integers.

Given a partial or full signature on a vector \mathbf{v} , an arbitrary party (without access to any signing key) can run the **Stretch** algorithm to increase the value of any component of the vector up to the pre-determined maximum set by the vector $\hat{\mathbf{v}}$ in the **KeyGen** algorithm. To do so, the party exponentiates the signature by the prime associated with the desired dimension to increment the value in that dimension (*i.e.*, decreasing the distance from the maximum). The

¹Existing techniques for removing the assumption of a trusted dealer may be applied to our construction (*e.g.*, [15]).

output of the algorithm is a new partial or full signature $\sigma' = (\sigma)^{e_k^a} \bmod N$, where k is the dimension being incremented and a is the amount by which it is increased.

The **Combine** algorithm takes t partial signatures on potentially different vectors from distinct signers and produces a full signature from them. In order to properly combine partial signatures on different vectors, the algorithm begins by applying the **Stretch** operation on each dimension of each partial signature to ensure all partial signatures provided as input sign the same vector. The resultant vector is actually the component-wise maximum among the t signed vectors given as input. Once all partial signatures embed the same vector, the **Combine** algorithm computes (modified) Lagrange coefficients for each partial signature and interpolates the signing key sk in the exponent. Thus, the output of the algorithm is a full signature $\sigma = \mathcal{H}(c)^{sk \prod_{k=1}^d e_k^{\hat{v}[k]}} = \mathcal{H}(c)^{\prod_{k=1}^d e_k^{-\hat{v}[k]-\mathbf{v}[k]+1}} \bmod N$ on the vector \mathbf{v} . Notice that in addition to requiring all vectors to be stretched to be identical, the context of these vectors must also be the same for the **Combine** algorithm to successfully produce a full signature. It is also important to note that this **Combine** procedure, after stretching, is exactly that of Shoup [38], and so interested readers should refer to that work for detailed technical explanation of the procedure.

Finally, the **Verify** algorithm takes as input a full signature and a vector-context pair, and outputs 1 if the full signature embeds the given vector-context pair. The verification procedure simply checks that $\mathcal{H}(c) = \sigma^{\prod_{k=1}^d e_k^{\hat{v}[k]-\mathbf{v}[k]+1}} \bmod N$, which only occurs if the signature encodes the given vector components as the appropriate distance from the predetermined maximum for each dimension.

We prove that this construction is shrink unforgeable under the strong RSA and DDH assumption in the random oracle model. Since we model $\mathcal{H}(\cdot)$ as a random oracle, we additionally quantify the number of queries q_h made to $\mathcal{H}(\cdot)$ in the adversary advantage. More formally, $\mathbf{Adv}_{BVS}^{\text{SU-CMA}}(T, q_h) = \max_{\mathcal{A}} \mathbf{Adv}_{BVS}^{\text{SU-CMA}}(\mathcal{A})$ where the maximum is taken over all adversaries \mathcal{A} running time at most T and making at most q_h queries to $\mathcal{H}(\cdot)$.

THEOREM 1. *Consider the above bounded vector signature construction with d dimensions and each vector dimension restricted to the range $[0, \hat{v}]$. Then, ignoring terms negligible in κ ,*

$$\begin{aligned} \mathbf{Adv}_{BVS}^{\text{SU-CMA}}(T, q_h) &\leq 2 \cdot q_h \cdot d \cdot \hat{v} \cdot \left(\frac{n}{t-1}\right)^{2(t-1)} \cdot \mathbf{Adv}_{\text{QR-SRSA}}(T') \\ &\quad + d \cdot \hat{v} \cdot \left(\frac{n}{t-1}\right)^{2(t-1)} \cdot (t-1) \cdot \mathbf{Adv}_{\text{DDH}}(T'') \end{aligned}$$

where $T' = T'' = T + O(t(d\hat{v} \log(n+d \log n) + \log(n!)))$ group operations.

The proof of the theorem can be found in our accompanying technical report [41].

3.2 Efficiency

With respect to space efficiency, the signing keys consist of a group element in \mathbb{Z}_m , while the public key is made up of an RSA modulus and d prime exponents. Both the partial and full signatures in our construction are represented by a single group element in QR_N .

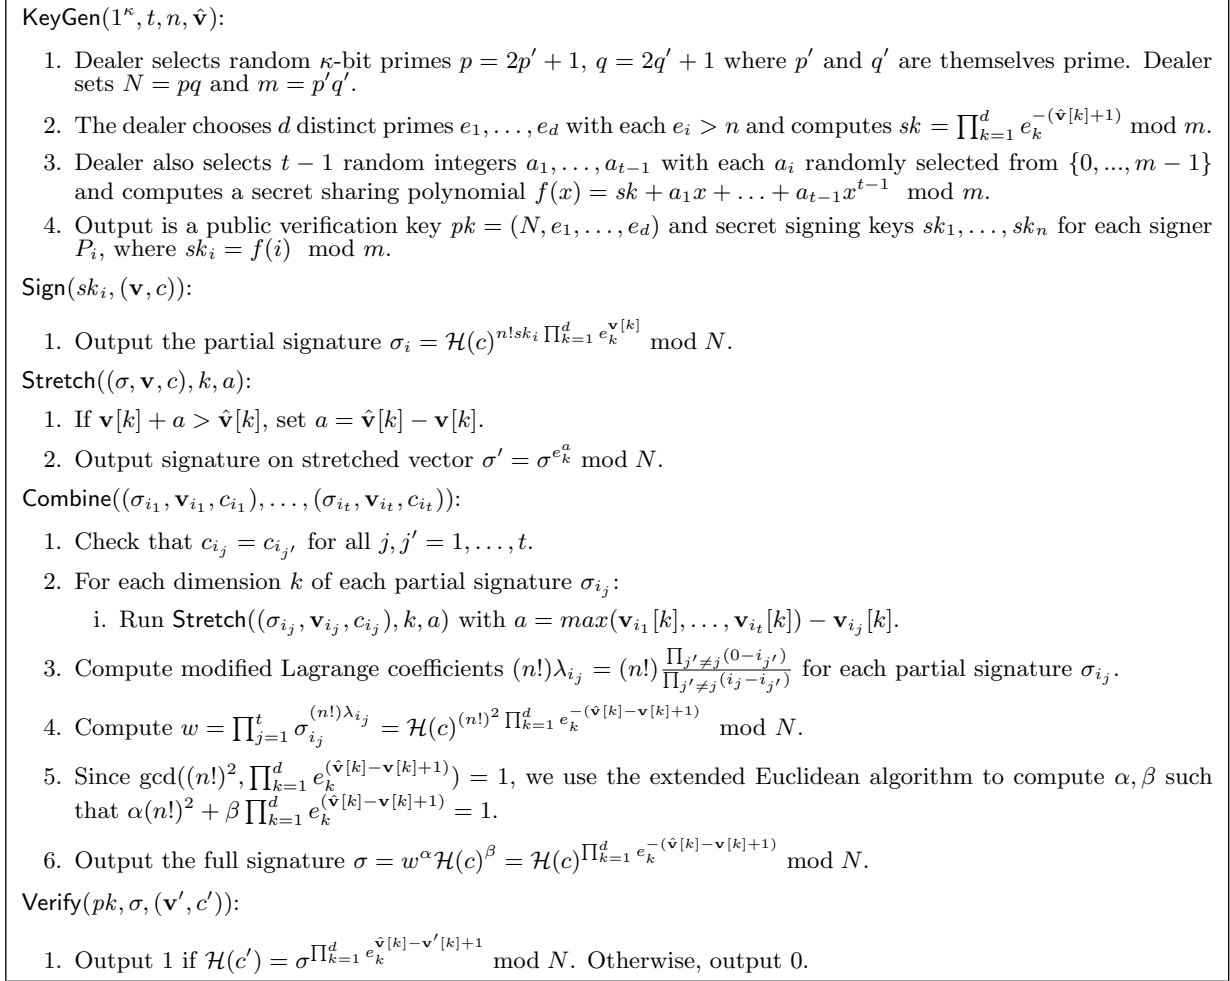


Figure 1: Algorithms for our bounded vector signature construction.

The **Sign** and **Verify** algorithms require at most $\sum_{k=1}^d \hat{\mathbf{v}}[k]$ modular exponentiations in the group QR_N . Likewise, the **Stretch** algorithm requires a number of modular exponentiations that is linear in the range of the dimension being stretched. Finally, the **Combine** algorithm requires t modular exponentiations in QR_N to perform the Lagrange interpolation, and two additional exponentiations to derive the full signature using the coefficients from the extended Euclidean algorithm.

4. EXTENSIONS

The basic bounded vector signature construction given in the previous section may be extended in a variety of ways to achieve additional functionalities. These extensions include the ability to dynamically limit the use of the **Stretch** operation and to verify partial signatures via robustness techniques from the threshold signature literature. Here, we briefly describe each of these extensions in turn.

4.1 Dynamically Limiting Malleability

The malleability properties of the basic bounded vector signature construction allow anyone who possesses a partial or full signature to produce valid signatures on any vector with components greater than those in the originally

signed vector, up to the maximum values specified during the **KeyGen** procedure. In certain instances, however, it is desirable to give each signer a degree of control over the use of the **Stretch** operation and, consequently, the signatures that other parties may produce. This can be accomplished by each signer taking steps to limit the degree to which its own partial signature can be stretched, and then leveraging the threshold t to limit what full signatures can then be produced using them.

Specifically, a signer i can create one bounded vector (partial) signature on its intended vector \mathbf{v}_i , and another bounded vector (partial) signature on the vector \mathbf{v}'_i , where $\hat{\mathbf{v}}[k] - \mathbf{v}'_i[k] \geq \mathbf{v}_i[k]$ for all dimensions k . This second signature can be viewed as bounding the extent to which the original vector can be stretched to only those vectors that can be simultaneously produced by both signatures. When $\mathbf{v}'_i[k] = \hat{\mathbf{v}}[k] - \mathbf{v}_i[k]$ for all dimensions k , all malleability has been removed from the partial signature.

Although the procedure above limits the malleability of partial signatures, its effect on the full signatures produced by **Combine** depends on the threshold t and its relation to the number of signers n . One particularly interesting parameterization allows us to ensure that a full signature (consisting of two bounded vector signatures) can be produced on only the

r^{th} smallest value in each dimension k (*i.e.*, value at rank r) from among the n values $\mathbf{v}_i[k]$, $i \in \{1 \dots n\}$. To implement this, we set the threshold for the first bounded vector signature instantiation, to which signer i contributes a partial signature for \mathbf{v}_i , to $t_1 = r$. In the second bounded vector instantiation, each signer contributes a partial signature for \mathbf{v}'_i where $\mathbf{v}'_i[k] = \hat{\mathbf{v}}[k] - \mathbf{v}_i[k]$, and we set the threshold to $t_2 = n - r + 1$. The first partial signatures can be stretched to increase the values $\mathbf{v}_i[k]$ of the $r - 1$ smallest such values, while the second partial signatures can be stretched to decrease the values $\hat{\mathbf{v}}[k] - \mathbf{v}'_i[k]$ of the $n - r$ largest such values. The respective full signatures only overlap at the rank- r values in each dimension k , and consequently cannot be altered via the **Stretch** algorithm. When $r = 1$, this is a non-malleable signature on the component-wise minimums, and when $r = n$, this is a non-malleable signature on the component-wise maximums, for example. In Section 5, we leverage this capability to perform intersection and union of signed data structures.

4.2 Adding Robustness

The notion of *robustness* in threshold signature schemes captures the inability of an adversary to produce incorrectly formatted partial signatures, which may cause the **Combine** algorithm to fail. Since the core of our bounded vector signature construction is based on the threshold RSA scheme of Shoup [38], we can make use of the robustness techniques presented therein. Broadly speaking, the robustness property can be achieved by requiring the signer to prove, in zero knowledge, that (i) the key used to sign the partial signature is the same as the one provided by the dealer, and (ii) that the partial signature is properly formed according to the algorithm description.

Specifically, the dealer chooses a random element $g \in QR_N$. Note that with overwhelming probability g is a generator of QR_N . He then computes and publishes the value $G_i = g^{s_i k_i}$ for each signer P_i . The partial signature σ_i is then accompanied by a zero-knowledge proof that $\log_{\mathcal{H}(c)}(\sigma_i^2) = \log_g(G_i^{2(n-1)^2} \prod_{k=1}^d e_k^{v_k^{[k]}})$, which can be accomplished using standard zero-knowledge techniques for proving the equality of discrete logarithms [12, 11]. Furthermore, this proof can be made non-interactive by using the Fiat-Shamir heuristic [20]. Due to space considerations and the similarity of our techniques with those of Shoup, we forgo a proof of the simulatability of this zero-knowledge proof.

5. APPLICATIONS

Bounded vector signatures can be used to efficiently represent signatures on a variety of data structures while still allowing certain operations on those data structures. These efficient, malleable signatures are particularly important in distributed systems applications, where disparate parties must securely share information while limiting communications overhead. Here, we describe two such data structures, namely sets and intervals. For each data structure, we describe how to represent it using our bounded vector signatures, specify operations on those signed data structures enabled by the **Stretch** operation, and suggest applications to problems in the area of distributed systems security.

The use of bounded vector signatures in the distributed applications described below provides a number of benefits. These include the compact representation of large groups of

items (*e.g.*, sets and intervals) with a constant number of group elements in QR_N , the ability to represent aggregated information from all members of a group as a single signature, non-interaction of the **Combine** and **Stretch** protocols, and the ability to locally update the full signature on the aggregated information. Furthermore, the threshold used in the **Combine** algorithm acts as a voting mechanism for the aggregated information, and may be of use in scenarios where Byzantine fault tolerance is required. Unless otherwise noted in this section, the bounded vector signatures are parameterized with $t = n$.

5.1 Sets and Multisets

The set data structure allows for the storage of unordered collections of unique items or values. In a multiset, this concept is extended to allow multiple copies of items (*i.e.*, associating a count with each item). Sets (*resp.*, multisets) are represented in our bounded vector signature scheme by associating each item in the universe of possible items with a dimension in the signed vector and limiting the maximum value of the dimension to one (*resp.*, the maximum per-item count). In essence, the vector representing a set reflects the presence or absence of each of the d items in the universe, while a vector for a multiset dictates the count of each item.

One downside of this approach for signing sets and multisets is that it requires a number of dimensions equal to the size of the universe of items represented, which may be quite large. In those cases, we may instead use our bounded vector signatures to encode the set in a *Bloom filter* [7]. Rather than associating each component of the vector with a specific item, we instead use the outputs of hash functions applied to the item to determine which dimensions to set to one. To test membership, we check that all dimensions to which the item hashes are set to one. In short, Bloom filters allow for a compact representation of sets by trading off a tunable probability of false positives when testing set membership (*i.e.*, an item is reported as a member even though it is not). Like simple sets, our bounded vector signatures represent a Bloom filter using a d -dimensional vector with maximum values for each dimension set to one, only in this case the number of dimensions may be much smaller than the number of items in the universe and is related to the desired probability of false positives.

Operations on Signed Sets.

The stretch property of the underlying bounded vector signatures allows us to perform the union operation among multiple set signatures simply by performing the **Combine** operation. As discussed in the previous section, **Combine** effectively takes the maximum value in each dimension from among the t partial signatures to form a final signature. As it turns out, this exactly defines the union among the signed sets. Furthermore, the security of our bounded vector signature scheme ensures that while items may be added to the (partially or fully) signed set, it is computationally infeasible for an adversary to remove an item. If additions to the fully signed set should not be permitted, then the rank-order technique described in Section 4.1 can be used to preclude this. The union operation and security guarantee also hold for multisets and Bloom filters, as well.

The intersection operation is achieved by having each signer produce a signature on a vector where each component encodes the distance between the signed value and the max-

imum allowable value in that dimension (*i.e.*, the complement for sets or the number of copies of items not in the multiset). As mentioned in Section 4.1, the **Stretch** operation on these signatures *decreases* the value being signed by increasing the distance from the maximum. Therefore, when **Combine** is run on these signatures, the largest value that can be signed is, in effect, the minimum value in each dimension. The security of bounded vector signatures ensures that the resultant signatures can have items removed from the set (resp., multiset, Bloom filter), but not added. Again, the technique of Section 4.1 can be used to preclude removals from the fully signed intersection, if so desired.

Distributed Systems Applications.

Signatures on sets, multisets, and Bloom filters can be used in a number of distributed data sharing applications, including peer-to-peer systems [14, 30, 24, 36], distributed caches [19, 9, 35], blacklisting services [17, 40, 39], and network flow monitoring systems [18]. In each of these applications, a collection of entities must provide users with an aggregate view of the data that they maintain as a group. These systems cannot assume interactivity, must provide an efficient method of verification to the user, and must be able to perform updates on the aggregated data without requiring additional interaction.

As a concrete example, consider the problem of spam blacklisting services, such as Spamhaus [40], SORBS [39], and DNSBL [17], which aggregate information about IP addresses sending spam e-mail. These blacklists must aggregate the IP addresses from several sources and provide users with a method of verifying that the given list does indeed represent that aggregate information (*i.e.*, names have not been spuriously added or removed, as the case may be). The blacklists can be efficiently represented in a Bloom filter, signed by our bounded vector signature method, and combined in a completely offline manner by the blacklist providers to give users a single signature to verify the aggregated blacklist. In the time between blacklist aggregation events, the individual sources can continue adding or removing IP addresses from the aggregated blacklist through the use of the **Stretch** operation without requiring additional interactions, if such operations are allowed by the signature instantiation. By contrast, traditional digital signatures would force the user to verify the blacklists from each of the sources independently and require the blacklist sources to continually generate new signatures on the published blacklists.

5.2 Intervals

Bounded vector signatures can also be used to represent contiguous intervals of values, in addition to discrete sets of items. There are potentially two ways in which to encode an interval of values within a bounded vector signature scheme. In the first, the signer creates 2-dimensional vectors, where the components of the vector encode the left and right endpoints of the signed interval, respectively. In the second, the signer may create two 1-dimensional vectors with each vector encoding one of the endpoints. The primary trade-off between these two methods is that in the first case the range is efficiently encoded in a single group element, whereas the second case allows for two unique parameterizations of the bounded vector signatures for the vectors (*e.g.*, two different thresholds t). As we will show later in this section, there are potential applications that are specific to both ap-

proaches. The concept of signing an interval can also be expanded to an arbitrary number of dimensions by simply increasing the dimensionality of the signed vector, or number of signatures in the case where each endpoint is a single vector. Specifically, if we want to encode a d -dimensional box (*i.e.*, the Cartesian product of d intervals), we create a $2d$ -dimensional vector to encode the left and right endpoints for the allowable range in each dimension.

Operations on Signed Intervals.

The use of bounded vector signatures provides the signer with a number of options for encoding the endpoint values of the signed interval, which in turn enable different functionalities. For example, if the signer were to encode the left-most point in the interval as the value itself and the right-most point as the distance to the maximum allowable value in the bounded vector scheme (*i.e.*, the vector $\langle \mathbf{v}[1], \hat{\mathbf{v}}[2] - \mathbf{v}[2] \rangle$ with the restriction $\mathbf{v}[1] \leq \hat{\mathbf{v}}[2] - \mathbf{v}[2]$), then the resultant signed range could only be *shrunk* and never expanded — this is equivalent to the dynamic malleability limitation described in Section 4.1. On the other hand, if the signer encoded the left-most point as the distance from the maximum and the right-most as the value itself (*i.e.*, $\langle \hat{\mathbf{v}}[1] - \mathbf{v}[1], \mathbf{v}[2] \rangle$ with $\hat{\mathbf{v}}[1] - \mathbf{v}[1] \leq \mathbf{v}[2]$), then the signed range could be *expanded* but never shrunk. Notice that if we set both endpoints to be encoded relative to the same value (*e.g.*, $\langle \mathbf{v}[1], \mathbf{v}[2] \rangle$ or $\langle \hat{\mathbf{v}}[1] - \mathbf{v}[1], \hat{\mathbf{v}}[2] - \mathbf{v}[2] \rangle$), we can slide and expand the interval in one direction only.

The range encodings above, along with the **Stretch** operation on the underlying bounded vector signatures, allow a combiner to efficiently compute the intersection and union of signed intervals. To compute intersections, each signer encodes the endpoints of its signed interval such that the range can only be shrunk. When the **Combine** algorithm is run on these partial signatures, the combiner will be forced to shrink the intervals to equal the smallest overlapping range and so the output will be a full signature on the intersection. For the union operation, the signers use an encoding that allows expansion of the signed interval and the combiner will expand all intervals such that the full signature will contain the left and right-most points from among all partially signed intervals. We note that this is slightly different from the typical definition of union among intervals, since in our system the union operation is allowed even if the intervals do not overlap. The security of the bounded vector signatures ensures that the full signature output by the combine procedure cannot be altered to sign a value not found within the intersection or union of the individual partially signed intervals.

Distributed Systems Applications.

Like signatures on sets and multisets, our bounded vector signature representations of intervals can be used to accomplish data sharing tasks within distributed systems for values that have natural ordering properties. One example is a time interval. For example, an approach to the management of public-key certificates (*e.g.*, [29]) employs an *online revocation authority ORA to countersign*, for short intervals, long-term certificates created by a more trusted offline certificate authority *CA*. In a traditional implementation of this idea, both the signature by *CA* and the countersignature by *ORA* would be implemented using separate signatures, resulting in two signature verifications per use of the certificate. Our

bounded vector signatures support the consolidation of these two signatures into one: the *CA* provides to *ORA* a signed certificate with a large validity period, which *ORA* can parcel out in small portions by “sliding” the ends of the validity interval accordingly (rather than separately countersigning each short validity interval).

A second area for using bounded vector signatures to represent intervals is sensor applications. Marzullo [31], for example, considered a setting in which multiple sensors produce an interval in which it senses a value (e.g., temperature, time) to lie, and where a sensor is *correct* if the interval it returns contains the actual value (and is *faulty* otherwise). His algorithm uses intervals from n sensors to produce the tightest interval possible in which the actual value lies, provided that fewer than f sensors are faulty. A direct use of our bounded vector signatures would permit each sensor to sign its interval and subsequently combine these signed intervals using this fault-tolerant intersection algorithm to produce a signature on the tightest interval possible in which the actual value lies. Others can be disallowed from tightening the interval further using the rank-order techniques of Section 4.1.

A third interesting application of signatures on intervals is their use to represent a broad class of access-control structures used in key management tasks, like hierarchical access controls [2, 4] and group key management [34, 42]. In these settings, attributes that determine access can often be encoded within a d -dimensional space. For instance, in some location-based access control schemes [3, 43], the dimensions encode the location of the entity in 3-dimensional space, plus an extra dimension for time. A service provider can sign subranges describing the allowable access attributes for the entities in the distributed system and provide them with their respective subranges as a non-interactive access token. These entities can then derive a “key”² for any subspace of their allowable attributes using the *Stretch* operation on the signature given to them. Such a scheme can derive several benefits from our bounded vector signatures, such as security against key recovery [2], which roughly states that it is infeasible for the adversary to derive a key for an access attribute for which she does not have a token.

6. CONCLUSION

In this paper, we described a new malleable signature scheme for signing vectors of natural numbers, which we call bounded vector signatures. The primary contribution of our scheme is a malleability property that allows arbitrary parties to *increase* the value embedded in any component of the signed vector without access to the signing key, while making it computationally infeasible to *decrease* values. In the multiparty setting, this malleability property allows each signer to sign a potentially different vector of values and yet still produce a valid full signature representing the component-wise maximum of any t vectors. We described an efficient construction for our scheme and proved it secure under the strong RSA and decisional Diffie-Hellman assumptions in the random oracle model. Finally, we showed that bounded vector signatures could be used to sign interesting data structures, such as sets and intervals, and that the malleabil-

²Depending on the scheme in use, additional key derivation steps may be necessary. For simplicity we just denote the signature as a “key”.

ity of those signatures enabled efficient implementations of standard operations on the signed structures. Moreover, we illustrated the unique benefits of these malleable signatures in providing compact and non-interactive solutions to several problems in the area of distributed systems security.

Acknowledgements

This work was supported in part by NSF grants 0756998, 0910483, and by a grant from Intel. Support was also provided by NSF grant 0937060 that was awarded to the Computing Research Association for the CIFellows Project.

7. REFERENCES

- [1] S. Agrawal, D. Boneh, X. Boyen, and D. Mandell-Freeman. Preventing pollution attacks in multi-source network coding. Cryptology ePrint Archive, Report 2010/183, 2010. <http://eprint.iacr.org/>.
- [2] M.J. Atallah, M. Blanton, N. Fazio, and K.B. Frikken. Dynamic and efficient key management for access hierarchies. *ACM Transactions on Information and System Security*, 12(3):1–43, 2009.
- [3] M.J. Atallah, M. Blanton, and K.B. Frikken. Efficient techniques for realizing geo-spatial access control. In *2nd ACM Symposium on Information, Computer and Communications Security*, pages 82–92, 2007.
- [4] G. Ateniese, A. De Santis, A.L. Ferrara, and B. Masucci. Provably-secure time-bound hierarchical key assignment schemes. In *13th ACM Conference on Computer and Communications Security*, pages 288–297, 2006.
- [5] N. Baric and B. Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *Advances in Cryptology—Eurocrypt ’97*, pages 480–494, 1997.
- [6] M. Bellare and Neven G. Transitive signatures: new schemes and proofs. *IEEE Transactions on Information Theory*, 51(6):2133–2151, 2005.
- [7] B.H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [8] D. Boneh, D. Freeman, J. Katz, and B. Waters. Signing a linear subspace: Signature schemes for network coding. In *12th International Conference on Practice and Theory in Public Key Cryptography*, pages 68–87, 2009.
- [9] J.W. Byers, J. Considine, M. Mitzenmacher, and S. Rost. Informed content delivery across adaptive overlay networks. *IEEE/ACM Transactions on Networking*, 12(5):767–780, 2004.
- [10] J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 101–120, 2002.
- [11] J. Camenisch and M. Michels. Separability and efficiency for generic group signature schemes. In *Advances in Cryptology—Crypto ’99*, volume 1666, pages 413–430, 1999.
- [12] D. Chaum and T. P. Pedersen. Wallet databases with observers. In *Advances in Cryptology—Crypto ’92*, pages 89–105, 1992.

- [13] R. Cramer and V. Shoup. Signature schemes based on the strong RSA assumption. *ACM Transactions on Information and System Security*, 3(3):161–185, 2000.
- [14] F.M. Cuenca-Acuna, C. Peery, R.P. Martin, and T.D. Nguyen. Planetp: Using gossiping to build content addressable peer-to-peer information sharing communities. In *12th IEEE Symposium on High Performance Distributed Computing*, pages 236–246, 2003.
- [15] I. Damgård and M. Koprowski. Practical threshold RSA signatures without a trusted dealer. In *Advances in Cryptology—Eurocrypt ’01*, pages 152–165, 2001.
- [16] W. de Jonge and D. Chaum. Some variations on RSA signatures and their security. In *Advances in Cryptology—Crypto ’86*, pages 49–59, 1986.
- [17] DNS-based blacklist, 2010. <http://www.dnsbl.info>.
- [18] C. Estan and G. Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems*, 21(3):270–313, 2003.
- [19] L. Fan, P. Cao, J. Almeida, and A.Z. Broder. Summary cache: a scalable wide-area web cache sharing protocol. *IEEE/ACM Transactions on Networking*, 8(3):281–293, 2000.
- [20] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology—Crypto ’86*, volume 263 of LNCS, pages 186–194, 1986.
- [21] R. Gennaro, S. Halevi, and T. Rabin. Secure hash-and-sign signatures without the random oracle. In *Advances in Cryptology—Eurocrypt ’99*, pages 123–139, 1999.
- [22] R. Gennaro, J. Katz, H. Krawczyk, and T. Rabin. Secure network coding over the integers. *Cryptology ePrint Archive*, Report 2009/569, 2010. <http://eprint.iacr.org/>.
- [23] A. Hevia and D. Micciancio. The provable security of graph-based one-time signatures and extensions to algebraic signature schemes. In *Advances in Cryptology—Asiacrypt ’02*, pages 379–396, 2002.
- [24] T.D. Hodes, S.E. Czerwinski, B.Y. Zhao, A.D. Joseph, and R.H. Katz. An architecture for secure wide-area service discovery. *Wireless Networks*, 8(2/3):213–230, 2002.
- [25] S. Hohenberger and B. Waters. Short and stateless signatures from the RSA assumption. In *Advances in Cryptology—Crypto ’09*, pages 654–670, 2009.
- [26] R. Johnson, D. Molnar, D. Song, and D. Wagner. Homomorphic signature schemes. In *The Cryptographer’s Track at the RSA Conference on Topics in Cryptology*, pages 244–262, 2002.
- [27] A. Kiayias and M. Yung. Group signatures: provable security, efficient constructions and anonymity from trapdoor-holders. *Cryptology ePrint Archive*, Report 2004/076, 2004. <http://eprint.iacr.org/>.
- [28] E. Kiltz, A. Mityagin, S. Panjwani, and B. Raghavan. Append-only signatures. In *32th International Colloquium on Automata, Languages and Programming*, 2005.
- [29] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, November 1992.
- [30] J. Ledlie, J.M. Taylor, L. Serban, and M. Seltzer. Self-organization in peer-to-peer systems. In *10th ACM SIGOPS Workshop*, pages 125–132, 2002.
- [31] K. Marzullo. Tolerating failures of continuous-valued sensors. *ACM Transactions on Computer Systems*, 8(4):284–304, 1990.
- [32] S. Micali and R. L. Rivest. Transitive signature schemes. In *The Cryptographer’s Track at the RSA Conference on Topics in Cryptology*, pages 236–243, 2002.
- [33] G. Neven. A simple transitive signature scheme for directed trees. *Theoretical Computer Science*, 396(1-3):277 – 282, 2008.
- [34] A. Perrig, D. Song, and D. Tygar. ELK: a new protocol for efficient large-group key distribution. In *2001 IEEE Symposium on Security and Privacy*, pages 247–262, 2002.
- [35] P. Reynolds and A. Vahdat. Efficient peer-to-peer keyword searching. In *2003 ACM/IFIP/USENIX International Conference on Middleware*, pages 21–40, 2003.
- [36] S.C. Rhea and J. Kubiatowicz. Probabilistic location and routing. In *21st Joint Conference of the IEEE Computer and Communication Societies*, pages 1248–1257, 2002.
- [37] A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- [38] V. Shoup. Practical threshold signatures. In *Advances in Cryptology—Eurocrypt ’00*, pages 207–220, 2000.
- [39] Spam and open relay blocking system, 2010. <http://www.sorbs.net>.
- [40] The spamhaus project, 2010. <http://www.spamhaus.org>.
- [41] L. Wei, S. E. Coull, and M. K. Reiter. Bounded vector signatures and their applications. Technical Report 10-018, Department of Computer Science, University of North Carolina at Chapel Hill, December 2010.
- [42] C. K. Wong, M. Gouda, and S. S. Lam. Secure group communications using key graphs. *IEEE/ACM Transactions on Networking*, 8(1):16–30, 2000.
- [43] H. Yuan and M.J. Atallah. Efficient and secure distribution of massive geo-spatial data. In *17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 440–443, 2009.