

Trust-Preserving Set Operations

Ruggero Morselli Bobby Bhattacharjee Jonathan Katz Pete Keleher
 Department of Computer Science, University of Maryland, College Park, Maryland, USA
 {ruggero, bobby, jkatz, keleher}@cs.umd.edu

Abstract— We describe a method of performing trust-preserving set operations by untrusted parties. Our motivation for this is the problem of securely reusing content-based search results in peer-to-peer networks. We model search results and indexes as data sets. Such sets have value for answering a new query only if they are *trusted*. In the absence of any system-wide security mechanism, a data set is trusted by a node a only if it was generated by some node trusted by a .

Our main contributions are a formal definition of the problem, and an efficient scheme that solves this problem by allowing untrusted peers to perform set operations on trusted data sets, and to produce unforgeable proofs of correctness. This is accomplished by requiring trusted nodes to sign appropriately-defined *digests* of generated sets; each such digest consists of an RSA accumulator and a Bloom filter. The scheme is general, and can be applied to other applications as well. We give an analysis that demonstrates the low overhead of the scheme and we include experimental data which confirm the analysis.

Index Terms—System design, simulations.

I. INTRODUCTION

This paper describes a method of performing trust-preserving computations on sets. The party performing the computation does not need to be trusted, but the result is a set which is trusted to the same extent as the original input. These properties allow the scheme to be used in distributed environments with many untrusted hosts. Our motivation for addressing this issue is the problem of securely reusing content-based search results in peer-to-peer (P2P) networks. However, the result is more general than result caching, and can potentially be used in other areas.

For purposes of this work, we define a P2P network as a large, distributed, and decentralized structure that allows autonomous *peers* to cooperatively provide access to a potentially large set of data. A number of approaches to providing lookup services on such data have been proposed. Distributed hash tables (DHTs) (Chord [1], CAN [2], Tapestry [3], Pastry [4] etc.), for example, use cryptographic hashes to provide near-random association of objects to sites that “publish” the object to the rest of the system. Objects are “looked up” by using the hash of

the object name to route to the corresponding peer, usually in $O(\log n)$ or so overlay hops. DHTs provide excellent balance of routing load because paths to a given node vary widely based on the path’s origin, and because related objects (even with only slightly different names) are randomly distributed through the system.

This random distribution is the strength of the approach, but it also destroys object locality. A set of objects chosen by common attributes or characteristics must necessarily include members mapped to peers throughout the system. This distribution, together with the sheer scale and diversity of the source data, makes it impractical to consider approaches to providing content searches that (even periodically) flood the network in order to satisfy queries. Similarly, flooding the network should not be considered a practical primitive when creating indexes for satisfying queries.

Instead, searches must rely on incrementally created and maintained indexes. Each “index” (or *attribute index*) records nodes that have a particular attribute, or attribute value. Given the nature of P2P applications, the indexes should be as distributed and as decentralized as the underlying system. However, a straightforward use of distributed indexes is still potentially quite expensive. A simple conjunction of two indexes will usually require the contents of at least one of the indexes to be sent across the network. These indexes may be large, as their size potentially grows linearly with the number of system peers.

The costs of satisfying queries can be reduced by caching query results in the system. Locality in query streams and object attributes can then be exploited by using prior results to help satisfy subsequent queries. The key feature is that only conjunctions of two or more attributes are stored. Such results are typically much smaller than the original attribute indexes, and can reduce the cost of satisfying multi-item queries by several orders of magnitude. Using such views efficiently is non-trivial, but has been addressed in other work [5].

Our motivation is the problem of making such cached results *secure*. Our main contribution is to show an efficient scheme that allows an untrusted peer to perform set operations on trusted data sets, and to supply an unforge-

able proof to the client of the correctness of the result. The approach we take is to reduce the trustworthiness of the derived value (the cached result) to a function of the trustworthiness of the sources (the original attribute values). The latter are more easily secured, as they are long-lived and fairly heavyweight. Cached results, however, are light-weight, ephemeral, and hence may be stored at arbitrary (possibly untrusted) nodes.

We briefly introduce an example of the problem we are trying to solve, along with some terminology. We model each index or cached query as a set of elements. The operations we wish to perform are standard set operations, such as union, difference, and intersection.

Our work may be viewed as extending the concept of *authenticated data structures* [6], [7], [8], [9], [10]. Briefly, these allow an untrusted host to answer queries about a *single* trusted data set in a trustworthy way (e.g., given a set S generated by a trusted source, they allow an untrusted directory d to answer queries of the form “is x in S ?”). Our work focuses on the more challenging (and more broadly-applicable) case of answering queries about *multiple* data sets in an efficient way (e.g., the intersection problem can be cast as a query “give me every x in $S_1 \cap S_2$ ”).

Two trusted *source* nodes, s_1 and s_2 , each store an index in the form of a set S_1 and S_2 (stored by s_1 and s_2 , respectively, cf. Fig. 1). An untrusted *directory* d caches the result of some set operation on S_1 and S_2 ; for this example, we assume it stores their intersection. The following also assumes that the computation of $S_1 \cap S_2$ is performed by d , although this need not necessarily be the case.

The problem that we analyze in this paper is how to construct a scheme that allows an arbitrary *client* c to verify that d did not falsify the result of the query. If such a verification can be performed, c can obtain the cached result from d and use it during evaluation of a new query that includes the same expression.

We assume that s_1 and s_2 each possess a private-public key pair for a signature scheme, and that c knows these public keys. We do not necessarily require the existence of a public key infrastructure; the public key will be tied to the concept of identity in the system. The fact that c knows and trusts s_1 and s_2 implies that c also knows their public keys.

The idea we exploit is the following: when s_1 sends a copy of S_1 to d , it also appends a *digest* of S_1 , which contains authenticity information, together with a signature produced using s_1 ’s signing key. Source s_2 supplies similar information. In the figure, s_1 sends to d the digest, D_1 , of S_1 , and the digest signature, σ_1 , on D_1 .

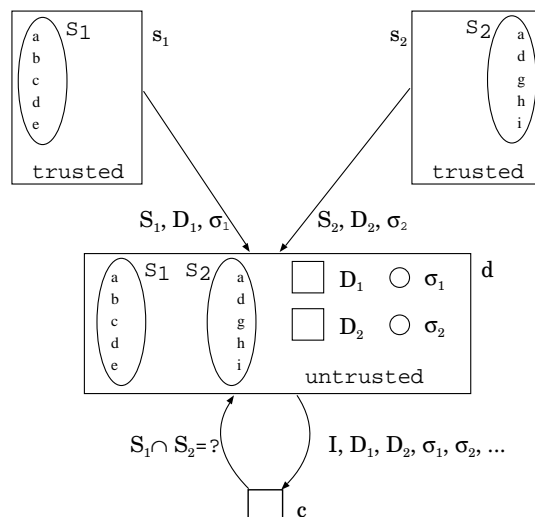


Fig. 1. The problem: an untrusted node computes intersection of two sets obtained from two trusted nodes. D_1 and D_2 are the digests of S_1 and S_2 respectively. σ_1 (resp. σ_2) is s_1 ’s (resp. s_2 ’s) signature on D_1 (resp. D_2). I is the intersection of S_1 and S_2 .

When d is queried for the intersection of the two sets, it provides the intersection I , the two digests, and other information derived during the intersection evaluation. Together, this information should be sufficient to prove the result accurate (at least at the time it was created).

The dangers are that a malicious directory could insert extraneous elements in the reply, (but still include all the elements of the intersection; we will call this an *insertion attack*), or it could fail to include elements that really are in the intersection (without inserting; we will call this a *deletion attack*). More generally, it could both insert and delete elements.

As a trivial solution, the two sources could sign the respective sets. The directory could then provide both of the original sets in response to the query, together with the signatures. (This is equivalent to choosing $D_1 = S_1$ and $D_2 = S_2$.) This is secure, because the client would be able to verify the authenticity of the two original sets and compute the intersection himself. Unfortunately, this solution is very expensive, because the two original sets might be much bigger than their intersection, so the communication overhead for the response may be excessive.

Our detailed study shows that, of the three basic set operations, intersection is indeed the most challenging to compute in provably correct fashion. We therefore first solve the intersection problem, and then show how to construct solutions for the case of union and set difference.

The rest of the paper is organized as follows. In Section II we give a formal definition for an intersection scheme that solves the sample intersection problem described above. In Section III, we summarize two exist-

ing techniques that we will use in our construction. In Section IV, we introduce a first scheme to solve the intersection problem and we discuss its limitations. In Section V, we give a secure and efficient intersection scheme that solves the problems of the previous one. In Section V-C, we show how to make the scheme composable, that is to say, how to securely compute intersections of sets that are themselves intersections performed by other untrusted nodes. In Section VI we show how the scheme can be easily extended to include the other two sets operations, union and difference. In Section VI-C we discuss implementation details. Finally, in Section VII, we give experimental results showing the efficiency of our approach and in Section VIII we summarize and conclude.

II. DEFINING AN INTERSECTION SCHEME

We define an intersection scheme formally in order to make it usable as a cryptographic primitive, like signature schemes or encryption schemes. As mentioned above, the secure computation of intersection is the most challenging of the secure set operations; it is easy to extend this definition to include union and difference.

Definition 1—Intersection scheme: An *intersection scheme* is a triple of efficient algorithms (Dgst , CV , Vrfy) (called, respectively, *digest*, *check value* and *verify* algorithm) with the following properties:

- Dgst takes as input a set S of elements in some universe \mathcal{U} and outputs a bit string $D \leftarrow \text{Dgst}(S)$ that we will call a *digest* of S .
- CV takes as input two sets S_1, S_2 and two digests D_1, D_2 and outputs a bit string $C \leftarrow \text{CV}(S_1, S_2, D_1, D_2)$, that we will call a *check value* for that quadruple of inputs.
- Vrfy takes as input one set I , two digests D_1 and D_2 and a check value C ; it produces a boolean value $b = \text{Vrfy}(D_1, D_2, I, C)$.
- The following correctness property holds. For all S_1, S_2, I if:

$$\begin{aligned} D_1 &\leftarrow \text{Dgst}(S_1) \wedge D_2 \leftarrow \text{Dgst}(S_2) \\ \wedge C &\leftarrow \text{CV}(S_1, S_2, D_1, D_2) \wedge I = S_1 \cap S_2 \end{aligned}$$

then $\text{Vrfy}(D_1, D_2, I, C) = 1$.

- The following security property holds. It is computationally infeasible for an adversary, on input S_1, S_2, D_1, D_2 to find a set I' and a value C' , such that $I' \neq I$ and $\text{Vrfy}(D_1, D_2, I', C') = 1$.

In words an intersection scheme is used as follows.

In the setting of the introduction, s_1 computes $D_1 \leftarrow \text{Dgst}(S_1)$ and a signature σ_1 on D_1 (computed using its

private key); s_2 computes $D_2 \leftarrow \text{Dgst}(S_2)$ and σ_2 analogously. d has a copy of D_1 and D_2 , together with signatures σ_1, σ_2 plus a copy of S_1 and S_2 (Fig. 1). When c queries d , then a honest d computes $I = S_1 \cap S_2$ and $C \leftarrow \text{CV}(S_1, S_2, D_1, D_2)$; finally d sends to c the values $(I, D_1, \sigma_1, D_2, \sigma_2, C)$.

Node c , in response to its query, receives $(I', D_1, \sigma_1, D_2, \sigma_2, C')$. If d is malicious, then I' and C' might be different from I and C . The signature scheme prevents the attacker d from lying on the remaining values. The client first checks that the signatures are valid and then it runs $\text{Vrfy}(D_1, D_2, I', C')$; this will output 1, if the answer to the query is correct, by the definition above; we would like this to return 0 if $I' \neq I$. Indeed the security property states that it is hard for the directory to find an incorrect response that passes the verification test.

III. BACKGROUND

In this section we briefly summarize two existing techniques that we will use in the secure intersection scheme. The first is the RSA accumulator, which allows an untrusted directory to securely prove membership of elements in a set. The second is the (counting) Bloom filter.

A. The RSA accumulator

Previous work [11], [6], [8] has shown how to use *cryptographic accumulators* to solve a related problem. A *source entity* s generates some set S . A copy of the set S is stored in the *untrusted directory entity* d . A *client entity* c queries the directory asking whether an element x_i belongs to the set S or not. If the directory replies with an affirmative answer, it must be able to prove that x_i is actually in the set.

To allow for this, s computes the value of a cryptographic accumulator $\text{Acc}(S)$ from the set S , signs it and sends a copy of the signature to d . When d wants to prove that $x_i \in S$, it computes a value w_i , called the *witness* of x_i , from the inputs S and x_i . Then it shows $\text{Acc}(S)$, w_i and the signature on $\text{Acc}(S)$. The client, after verifying the signature, can verify from $\text{Acc}(S)$, w_i, x_i that the answer is correct.

The work cited above pertains to the RSA accumulator, which has the important property that its size does not depend on the size of the set S .

Let $S = \{x_1, \dots, x_n\}$ be the set stored at a source node. Each x_i is represented by u bits. The source chooses a random RSA modulus N and a random $a \in \mathbb{Z}_N^*$, then it makes those public information¹.

¹ \mathbb{Z}_N^* is the set of integers between 1 and $N - 1$, that are relatively prime with N

Suppose there exists an efficient algorithm R that, on input of a u -bit element x produces an odd integer $e = R(x)$, called a *representative* of x . We require that the algorithm R implements a *division intractable function* [12]. That is to say, it is infeasible for an adversary to find elements x_1, \dots, x_n, x' such that $R(x')$ divides the product $R(x_1) \cdots R(x_n)$.

In order to compute the accumulator of S , compute representatives e_1, \dots, e_n of x_1, \dots, x_n , then output:

$$\text{Acc}(S) = a^{e_1 e_2 \cdots e_n} \pmod{N}. \quad (1)$$

In order to compute the witness w_i of x_i in S , output:

$$w_i = a^{e_1 e_2 \cdots e_{i-1} e_{i+1} \cdots e_n} \pmod{N}.$$

In order to verify the correctness of an answer $\text{Acc}(S)$, w_i , x_i , verify that:

$$w_i^{e_i} = \text{Acc}(S) \pmod{N}.$$

It can be proved ([6], [12]) that, under *the strong RSA assumption*, it is infeasible for an adversary to “fool” a client into believing that some x is in S when in fact it is not; i.e., it is infeasible for an adversary to find (x, w) such that $x \notin S$ but $w^{R(x)} = \text{Acc}(S) \pmod{N}$.

B. Counting Bloom filters

We assume that the reader has some basic familiarity with Bloom filters [13]. Suppose we use a Bloom filter as a digest for the intersection scheme. Also suppose that the directory answers the intersection query with a “claimed” intersection I' and the two Bloom filters B_1, B_2 for the two original sets S_1, S_2 . The client can verify that each element of I' appears to be in the two original sets, using the Bloom filters. This does not prevent the attacker from returning a subset of the real intersection, so Bloom filters are not the whole solution. Also, the Bloom filter would need to be prohibitively large in order to ensure that an adversary would not be able to find a false positive in either of the two original sets.

Our protocol employs a generalization of Bloom filters, called *counting Bloom filters* [14], for a different purpose. A counting Bloom filter is a data structure characterized by two parameters k and m . It requires k independent hash functions, which we denote h_1, \dots, h_k , with range in $\{1, \dots, m\}$. In particular, given a set S , the counting Bloom filter of S (denoted by $\text{Bl}(S)$) is a vector of m non-negative integers (counters). It can be incrementally constructed, starting with a vector of all zeros and then, for each element x in S , by incrementing the counters with indices $h_1(x), \dots, h_k(x)$. We denote by $\text{Bl}(S)_j$ the j -th counter of the filter; it represents the number of elements that hash to that index.

	0	1	
cat	1	1	cat
	0	0	
cat, dog	2	1	cat
	0	2	whale, chicken
dog	1	2	whale, monkey
	0	1	chicken
	$\text{Bl}(S_1)$	$\text{Bl}(S_2)$	
	$S_1 = \{\text{cat}, \text{dog}\}$	$S_2 = \{\text{cat}, \text{chicken}, \text{monkey}, \text{whale}\}$	

Fig. 2. Counting Bloom filters of two sets

IV. A FIRST ATTEMPT

In this section we show an intersection scheme based on counting Bloom filters. This scheme, as we will discuss, is neither secure nor efficient, but it will be useful in understanding the main idea of the correct secure intersection scheme.

Consider the two sets S_1, S_2 in Fig. 2, together with the corresponding counting Bloom filters (with $k = 2$ hash functions and $m = 7$ counters). The label on the side of each Bloom filter counter is the list of items in the set that hash to the index of that counter; e.g. in $\text{Bl}(S_1)$, the element “dog” appears in the label of the fourth and sixth counters, meaning that, for example, $h_1(\text{“dog”}) = 4$ and $h_2(\text{“dog”}) = 6$.

The client receives a signed copy of these two Bloom filters, together with the supposed intersection I' (for a malicious directory, this could be different from the real intersection I). To verify correctness, the client computes the Bloom filter $\widehat{\text{Bl}}$, obtained as the element-by-element minimum of $\text{Bl}(S_1)$ and $\text{Bl}(S_2)$, and the Bloom filter $\text{Bl}(I')$ of the returned intersection (Fig. 3). The condition:

$$\text{Bl}(I')_j \leq \widehat{\text{Bl}}_j \quad \forall j = 1, \dots, m. \quad (2)$$

holds for the correct intersection, because each element in I also belongs to both S_1 and S_2 ; so the client can check this condition using I' instead of I .

In the given example (but not in general), any insertion attack (that is to say, any attack that would make I' a superset of I) would be detected, because it would make at least one of the counters of the Bloom filter of I' (Fig. 3 on the right) greater than the corresponding counter of $\widehat{\text{Bl}}$ (same figure, on the left).

This test does not prevent an attacker from performing a deletion attack, i.e. returning an I' that is a subset of I (in the example, the directory could claim that the intersection is empty), because the only effect of removing

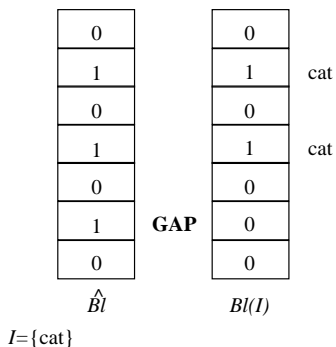


Fig. 3. The elementwise minimum, the Bloom filter of the intersection and the gap for the example in Fig. 2

elements is to decrease some of the counters. Ideally, if $\widehat{\text{Bl}}$ and $\text{Bl}(I')$ were equal, then we would know for certain that a deletion attack was not performed (otherwise, without the attack the Bloom filter of the intersection would have some counters strictly greater than the corresponding in $\widehat{\text{Bl}}$, which is a contradiction). The problem is that, even for the legitimate intersection I , some counters of the Bloom filter could be strictly less than the corresponding counter in $\widehat{\text{Bl}}$. In our example, the counter with index 6 of $\text{Bl}(I)$ is 0 while the corresponding in $\widehat{\text{Bl}}$ is 1; this is due to the fact that there are at least one element (“dog”) which are only in the first set and at least one element (“whale” and “monkey”) which are only in the second set, all of which hash to the same index 6. We will say that index 6 is a *gap*, meaning that there is a gap between the two counters in question.

Definition 2—Gap: An index j is called a *gap* if $\text{Bl}(I)_j$ is strictly less than $\widehat{\text{Bl}}_j$.

A deletion attack would create more gaps, as perceived by the client. Therefore we *require the directory to justify each gap*. In our example, if the directory returns the correct intersection, it can justify the gap 6 by including in the answer all the elements in both sets that map to index 6 (“dog” for the first set, “whale” and “monkey” for the second set); we will say that those are the *check elements* of the response. If the directory tries to return an empty intersection, it must justify three gaps (2, 4 and 6), for example by finding two strings that hash to indexes 2 and 4 and adding them to the check elements. The attack is restricted but still feasible.

More formally, the untrusted directory d will include in the answer to the intersection query a check value C that consists of a pair of sets (C_1, C_2) , that are respectively a set of check elements in $S_1 \setminus I$ and $S_2 \setminus I$. Upon receiving the answer $(I', \text{Bl}(S_1), \sigma_1, \text{Bl}(S_2), \sigma_2, C'_1, C'_2)$, the client c computes the Bloom filter $\widehat{\text{Bl}}$ obtained as the element-by-element minimum of $\text{Bl}(S_1)$ and $\text{Bl}(S_2)$. Here we denoted with I', C'_1, C'_2 the values that d has returned for

the intersection and the check elements; if d is malicious, those can be different from the I, C_1, C_2 an honest node would return.

The client c will check that condition (2) holds and also will check that for each $j = 1, \dots, m$ such that $\text{Bl}(I)_j < \widehat{\text{Bl}}_j$:

$$\text{Bl}(I'_j) + \text{Bl}(C'_1)_j = \text{Bl}(S_1)_j \quad (3)$$

$$\text{Bl}(I'_j) + \text{Bl}(C'_2)_j = \text{Bl}(S_2)_j \quad (4)$$

and will reject if any of these is not satisfied. Condition (3) states that for each gap j the check element set C_1 must contain some elements $\{y_1, \dots, y_s\}$ from $S_1 \setminus I$ that collectively hash $\text{Bl}(S_1 \setminus I)_j$ times to index j . That is to say:

$$\text{Bl}(S_1 \setminus I)_j = \text{Bl}(\{y_1, \dots, y_s\})_j.$$

Analogously (4) states that C_2 must contain some elements from $S_2 \setminus I$ that hash $\text{Bl}(S_2 \setminus I)_j$ times to index j .

In the next subsection we state some results that estimate the number of check elements needed by the scheme. In Section IV-B we use this result to estimate the level of security and efficiency of the scheme and in Section V, we will show how to make an efficient scheme out of this.

A. Number of check elements

Definition 3—Load of the filter: The *load* of a counting Bloom filter $\text{Bl}(S)$ of a set S of n elements, with k hash functions and m counters, is the expected value l of each counter in the filter, i.e.

$$l = \frac{kn}{m}. \quad (5)$$

Consider the counting Bloom filter intersection scheme, where the two sets S_1 and S_2 have respectively n_1 and n_2 elements each and their intersection has $|I| = q$ elements. Let n be the maximum set size allowed by the scheme ($n_1, n_2 \leq n$). Let m be the number of counters and k the number of hash functions. Let $l_1 = \frac{kn_1}{m}$, $l_2 = \frac{kn_2}{m}$ be the loads of $\text{Bl}(S_1)$, $\text{Bl}(S_2)$ and let $l = \frac{kn}{m}$ be the *maximum load* of the scheme.

We assume that if we randomly generate n distinct objects x_1, \dots, x_n from the universe according to any distribution of interest, the values $\{h_i(x_a) : i = 1, \dots, k \wedge a = 1, \dots, n\}$ are independent random variables uniformly distributed in $\{1, \dots, m\}$.

Theorem 1: The expected number of check elements for both sets satisfies:

$$E[|C_1|] \leq m(1 - e^{-l_2})l_1 \leq ml_1l_2 \leq ml^2 \quad (6)$$

$$E[|C_2|] \leq m(1 - e^{-l_1})l_2 \leq ml_1l_2 \leq ml^2. \quad (7)$$

A proof can be found in the appendix.

B. Security and efficiency considerations

The intersection scheme described above suffers from several kinds of attacks. For example, the attacker could insert an element in the intersection that is not in any of the original sets, as long as all the indices it maps to are gaps; it can be shown that making the probability of this attack negligible can be done only at the cost of Bloom filters with a prohibitively large number of counters (for a probability of 2^{-50} or less, we need $m \geq 37n$).

Independently of the security considerations, the load of the Bloom filters must be small enough to keep the number of check elements small. This implies that m must be much bigger than n , making the size of the Bloom filters themselves bigger than the size of encoding the original sets. We address this problem next.

In order to improve the efficiency of the counting Bloom Filter scheme, we use compressed counting Bloom filters [15]. We can reduce the size of the counting Bloom filters by applying a compression algorithm. This is still not sufficient to make the intersection scheme secure and efficient, but is the first step towards the solution to the problem. In this section we introduce some notation and some results on the compressed counting Bloom filters that we will need in Section V.

Let S be a set of n elements and consider its counting Bloom Filter $\text{Bl}(S)$ with m counters and k hash functions. We can apply a data compressor to the filter, obtaining a compressed filter of some size z .

Claim 1: An upper bound to the size of the compressed counting Bloom filter, assuming optimal compression, is given by

$$z = mH(l) \quad (8)$$

where $H(l)$ is the entropy of a Poisson distribution with mean l .

We believe the bound is tight. Proof available in [16].

V. A CRYPTOGRAPHICALLY SECURE INTERSECTION SCHEME

In this section we describe our full solution to the intersection problem. In Section V-B, we analyze the scheme to determine the optimal values of the parameters and the resulting overhead.

The scheme works as follows:

- Source s_i , to produce the digest of a set S_i , generates:
 - the compressed counting Bloom filter $\text{Bl}(S_i)$;
 - the RSA accumulator $\text{Acc}(S_i)$;
 - a signature σ_i on both.
- When a directory receives a query, it returns:
 - the intersection I ;

- the Bloom filters $\text{Bl}(S_1), \text{Bl}(S_2)$;
- the RSA accumulators $\text{Acc}(S_1), \text{Acc}(S_2)$;
- the check element sets C_1, C_2 (computed exactly as described in Section IV);
- two RSA accumulator witnesses w_1, w_2 . The two witnesses prove that $(I \cup C_1) \subset S_1$ and $(I \cup C_2) \subset S_2$ respectively.
- the two signatures σ_1, σ_2 .
- The client verifies an answer by:
 - checking that the signatures are correct;
 - checking that the witnesses are correct;
 - checking that, for each gap, there are enough check elements (3, 4).

The RSA accumulators prevent any attack that inserts elements in the intersection, because the adversary cannot prove that any such element belongs to both sets (Section III-A). Also the malicious directory cannot use something that is not in S_1 (S_2) as a check element for that set. As a consequence, the following holds [16]:

Claim 2: No attacker can return an incorrect intersection, unless it can break the security of the RSA accumulators or of the signature scheme.

We illustrate this scheme through the example in Fig. 2 and 3. When the client asks for the intersection, the directory returns $I = \{\text{“cat”}\}$, the check element sets ($\{\text{“dog”}\}, \{\text{“whale”}, \text{“monkey”}\}$) and a proof (via RSA accumulators) that “cat”, “dog” are in S_1 and another proof that “cat”, “whale”, “monkey” are in S_2 .

If the directory maliciously attempts to return an empty intersection, two extra gaps (2 and 4) are created for the Bloom filters. In order to justify gap 2, the directory has to find a check element for the first set that maps to index 2. Such a check element must belong to S_1 ; otherwise it won't be possible to find a witness for the RSA accumulator. The only option, therefore, is to use “cat”. The directory also has to find a check element for S_2 . Again, the only option is “cat”. Therefore the client will detect the attack by noticing that the same check element is returned for both sets.

A. Hashing the check elements

We here show how to reduce the size of the check elements, using hashing, and we introduce the parameter u (the size of an element).

In general the set elements may belong to an arbitrary domain and, therefore, be arbitrary length bit strings. In the intersection scheme we propose, the directory must include in the reply to the query several check elements, each of which can be significantly large. The client does not really care about what those check elements are, as long as they exist.

With this in mind, we choose a hash function (for example, SHA-1), then we map each element in the universe to its hash and we run the intersection scheme, not on the original elements, but on their hashes. When a Bloom filter or an RSA accumulator is computed for set S , it is actually computed for the set of the hashes of the elements of S . The directory replies to the client with the *real elements* of the intersection, but it provides only the hashes of the individual check elements. The hash function must be collision resistant, because, if the attacker could find two elements X_1, X_2 that map to the same value, it could substitute X_1 with X_2 in any set that contains X_1 .

If we call u the length of the output of the hash function, then we can think of the scheme as operating on elements of u bits. The collision resistance property requires u to be large enough; on the other hand the cost of the check elements is proportional to u . We believe that using $u = 160$ (SHA-1 output size) be a reasonable compromise; we will use this value in numerical examples and experiments, but our analysis holds for any value of u .

B. Choosing the parameters of the Bloom filters

Let us consider the overhead of the protocol in the size of the response to the query. Recall that the response, in addition to the intersection, contains the following parts:

$$(D_1, D_2, CV, \sigma_1, \sigma_2) = (\text{Acc}(S_1), \text{Bl}(S_1), \text{Acc}(S_2), \text{Bl}(S_2), C_1, C_2, w_1, w_2, \sigma_1, \sigma_2)$$

The elements $\text{Acc}(S_1), \text{Acc}(S_2), w_1, w_2$ are 4 RSA values. If, for example, we employ 1024 bit RSA, then the total overhead for those values is 512 bytes. The two signatures σ_1, σ_2 have a size that depends on the signature scheme considered. For example, they could be 1024 bits each for RSA signatures, bringing the total of these 6 components to 768 bytes.

From the analysis in Section IV-B, the size of the compressed Bloom filter is $z = mH(l)$. The encoded size of the check element set C_1 (or C_2) is u bits for each element; if we use the upper bound of theorem 1 as an estimate for the number of check elements, $|C_1| = ml^2 = knl$. Therefore the total overhead of the protocol (in the size of the response) is given by:

$$Ov = \frac{z}{n} + \frac{u|C_1|}{n} = k\left(\frac{H(l)}{l} + ul\right) = kf(l) \quad (9)$$

The overhead of the protocol is evaluated as a fraction of the total number of elements $2n$ in the two original sets.

The optimal parameters are $k = 1$ and the choice l_{opt} of l that minimizes $f(l)$. Note that, for $k = 1$, the counting

Bloom filter is actually a hash table, where each bucket is replaced by the number of elements in the bucket.

For $u = 160$, plotting f numerically shows that it has a minimum approximately at $l_{opt} = 0.01$, for which it takes a value $f(l_{opt}) \approx 9.7$. This means that the overhead of the protocol is, *in the worst case*, approximately of 9.7 bits per element in the original set. In practice, a more careful analysis (12) shows that the overhead decreases significantly whenever the size of the intersection is non-negligible compared to the two original sets or when the sizes of S_1 and S_2 are quite different from each other; e.g. it's below 2 bits per element if $n_1/n_2 = 1000$. Compare this with the overhead of the trivial protocol, in which the directory returns to the client the entire sets S_1 and S_2 (signed by the sources), which incurs in a overhead of $u = 160$ bits per element.

For example, if the two original sets contain both $n = 100000$ elements, the intersection contains 100 elements (so we are in the worst-case scenario), in our scheme the directory needs to send 2 Kbytes for the intersection and about 250 Kbytes for the accumulators and the Bloom filters; it saves a factor of 16 compared to the trivial scheme that would instead require 4 Mbytes and it still offer an arbitrarily high level of security (for example 1024 or 2048 bit strong RSA security). In a better scenario, for $|S_1| = 100000, |S_2| = 100$ and $|I| = 10$, our scheme requires 200 bytes for the intersection and about 27.5 Kbytes of overhead, compared to the 2 Mbytes of the trivial scheme.

For the optimal configuration, the Bloom filter uses $\frac{m}{n} = 100$ counters for each element in the original set. If the cost of that many counters, in terms of storage of the uncompressed Bloom filter and of processing time, is considered excessively high, then a suboptimal and larger value of the load can be used. For example, for $l = 0.03$, we can get away with $m = 33n$ counters and the overhead increases only to $f(l) = 11.3$ bits per element (plus the constant overhead due to the accumulators and the signatures).

C. Composability

In this section we show that the scheme is composable; i.e. an untrusted directory can perform a trust-preserving intersection of two sets, even if one or both of the two sets was obtained by another untrusted directory, as a different trust-preserving intersection.

Consider the following example, involving four source nodes and three untrusted directories (Fig. 4). The source nodes generate sets S_1, S_2, S_3, S_4 respectively. The first directory has a copy of S_1 and S_2 , together with *certificates* $Cert_1, Cert_2$ for those two sets, respectively; the

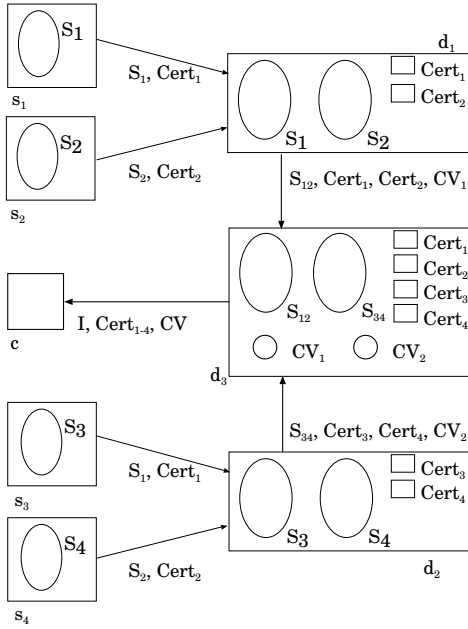


Fig. 4. Four trusted nodes (s_1 through s_4) and three untrusted directories (d_1 through d_3) involved in a two-level intersection computation.

certificate for a set consists, as described in the previous sections, in the digest of the set (a counting Bloom filter plus an RSA accumulator) and a signature on the digest by the source. The second directory, analogously, has a copy of S_3 and S_4 and the corresponding certificates. The third directory (d_3) queries the first directory (d_1) for the intersection $S_{12} = S_1 \cap S_2$. The answer will contain a certificate for the set S_{12} , consisting of the certificates for the two base sets and of a check value CV_1 , (which in turn consists of the two check element sets and the two witnesses). Analogously, d_3 obtains $S_{34} = S_3 \cap S_4$ from d_2 , also with a certificate. When a client (c) queries the third directory for $I = S_{12} \cap S_{34}$, the composability property of the scheme means that the directory can construct a certificate proving that I was computed correctly, using as inputs the sets S_{12} and S_{34} and their certificates, without seeing any of the original sets and without having to send the whole S_{12} and S_{34} to the client.

We now illustrate the algorithm to produce such certificate for the second-level intersection, through the example in Fig. 4, with the specific sets and Bloom filters as in Fig. 5.

The directory d_3 knows, among other things, the check values CV_1 and CV_2 . We have $CV_1 = (C_1, C_2, w_1, w_2)$, where $C_1 = \{\text{“dog”}\}$ and $C_2 = \{\text{“whale”, “monkey”}\}$ are the check elements for $S_1 \cap S_2$, and w_1, w_2 are the corresponding witnesses (e.g. w_1 witnesses that $\{\text{“cat”, “mouse”, “horse”, “dog”}\} \subset S_1$). Analogously $CV_2 = (C_3, C_4, w_3, w_4)$, where C_3 and C_4 are empty.

The directory d_3 computes the Bloom filters $\text{Bl}(S_{12})$

and $\text{Bl}(S_{34})$ from the two sets, then it computes the check elements for the intersection, using the same algorithm described in Section V for intersecting two basic sets; let C_{12} and C_{34} be the two check element sets thus obtained; in this example $C_{12} = \{\text{“mouse”}\}$ and $C_{34} = \{\text{“sheep”}\}$. The check value returned by d_3 is given by:

$$CV = (C_1, C_2, C_3, C_4; C_{12}, C_{34}; w'_1, w'_2, w'_3, w'_4). \quad (10)$$

Here w'_1 is a witness that $I \cup C_1 \cup C_{12} \subset S_1$. The directory can compute this, because it knows w_1 (which proves $S_{12} \cup C_1 \subset S_1$) and by construction $I \cup C_{12} \subset S_{12}$. The remaining w'_i are analogously defined.

Note that the client can perform verification and that the scheme inherits the security guarantee of claim 2.

We claim that the performance of the scheme does not degrade. Equation (10) shows that the size of a certificate for the set I is the same as the sum of the sizes of the certificates for the two intermediate intersections, plus the second-level check element sets C_{12} and C_{34} ; those do not represent a problem, because the worst-case overhead happens exactly in the case when their size is negligible. We will now justify this, with reference to C_{12} . If S_{12} has a negligible size compared to S_1 and S_2 , then the load of its Bloom filter will be very low and therefore C_{12} will have almost no elements. In cases where S_{12} is progressively larger compared to the two original sets, the size of C_{12} increases, *but the size of both C_1 and C_2 decreases accordingly, making the total size of the response from d_3 to c smaller.*

VI. FROM SECURE INTERSECTION TO SECURE FULL SET ALGEBRA

In Sections II through V-C, we have given a solution to the problem of secure intersection. In particular we have shown a recursive scheme that, from two sets, each accompanied by a suitable certificate, can generate a certificate for their intersection. In this section, we describe how to construct a suitable certificate for their union and their difference. With this extension, our proposed scheme allows untrusted directories to perform arbitrarily complex set operations on data from trusted nodes, and to prove the correctness of the result.

The described scheme satisfies the following important property:

Given a certificate for the set S , it is always possible to construct $\text{Bl}(S)$, without knowing S (11)

Note that, in the case of the intersection, this is true: if $S = S_1 \cap S_2$, the certificate of S includes either the Bloom

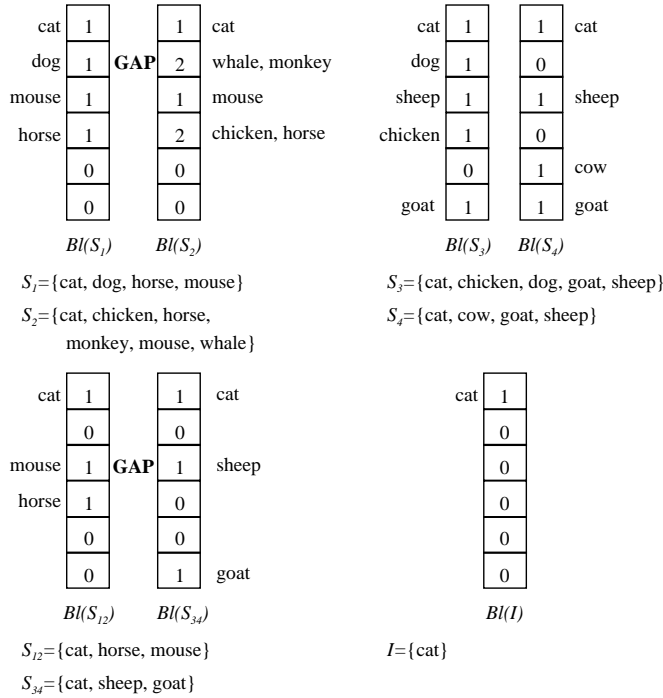


Fig. 5. An example of the scheme of an intersection of intersections. The figure shows the four original sets, the two intermediate intersections and the final intersection, together with the corresponding Bloom filters

filter of S_1 and S_2 or enough information to reconstruct them; it also includes the check element sets C_1, C_2 of the intersection. Therefore the Bloom filter of S can be obtained as the elementwise minimum of the two vectors $Bl(S_1) - Bl(C_1)$ and $Bl(S_2) - Bl(C_2)$.

A. Performing set difference

We now show how to perform set difference, through a modified version of the example described in Fig. 4 and 5. Suppose that the query that c asks of d is, instead, $S_{12} \setminus S_{34} = (S_1 \cap S_2) \setminus (S_3 \cap S_4)$. The answer to the query is: $A = \{\text{“horse”, “mouse”}\}$.

To prove the correctness of the answer, the directory needs to compute the check element sets C_{12} and C_{34} for the set difference.

- The check element set for the first set consists of all the elements of the *intersection*: $C_{12} = S_{12} \cap S_{34}$. In this case $C_{12} = \{\text{“cat”}\}$. Note that the client, during verification, will reconstruct $Bl(S_{12})$ and $Bl(S_{34})$. It will notice that the counter at index 1 is equal to 1 for both Bloom filters (Fig. 5). By showing “cat”, which maps to index 1, and proving that “cat” is in the intersection, the directory assures the client that it did not delete an element that happened to map to that index from the difference.

- The check element set C_{34} for the second set consists of all the elements of the opposite set difference ($S_{34} \setminus S_{12}$) that map to a gap. In this case $C_{34} = \{\text{“sheep”}\}$. By showing this element, we prove to the client that “mouse” is indeed in the set difference and not in the intersection. Note that no such proof is required for “horse”, since it does not map to a gap.

Once the check sets are known, d_3 can answer the query from c in the same form as in Fig. 4, with a check value of the form (10), but this time w'_1 is a witness that $C_1 \cup C_{12} \cup A \subset S_1$ and analogously w'_2 , while w'_3 is a witness that $C_3 \cup C_{12} \cup C_{34} \subset S_3$ and analogously w'_4 .

B. Performing set union

Suppose the directory is queried for $S = S_1 \cup S_2$; the directory has a certificate for both S_1 and S_2 : it can simply provide the two sets along with their certificates.

The catch is that, if the result S of the union is used as operand in a subsequent operations, then Property (11) does not hold anymore. To fix this, we add the following rule: the certificate for $S_1 \cup S_2$ includes the check element sets C_1 and C_2 computed exactly as in the intersection case, in addition to the certificates of S_1 and S_2 .

Additionally, when S is later used as operand in an intersection or in a difference, each element of S that appears in the final result, or in a check element set, should carry the information about whether it belongs to S_1 only, or to S_2 only, or to both S_1 and S_2 .

C. Standard Bloom filter sizes

When a source computes the digest of its set S , it does not know the size of the sets S will be intersected with. As a practical matter, then, we introduce a set of *standard set sizes* $\{\hat{n}_1, \hat{n}_2, \dots\}$ and the corresponding *standard Bloom filter sizes* $\{\hat{m}_1, \hat{m}_2, \dots\}$, with the property that $\hat{m}_i = \hat{n}_i/l$, and l is the desired value of the load (for example, the optimal $l = 0.01$). For example, we suggest the following standard set sizes $\{\hat{n}_1 = 10^3, \hat{n}_2 = 10^5, \hat{n}_3 = 10^7\}$ and a maximum set size of $4 \cdot 10^7$.

When the source generates the digest for S , it generates one different Bloom filter for each standard size. In the suggested configuration, it would generate three Bloom filters with respectively $10^5, 10^7, 10^9$ counters. The source will then sign each of these Bloom filters individually and provide all of the signatures to the directory.

When the directory computes the intersection of two sets, it will choose the standard Bloom filter size that minimizes the communication overhead with the client. For example, if the two sets contain $n = 8 \cdot 10^3$ elements

each, then the choice is between a Bloom filter with 10^5 counters, which corresponds to a load of $l = 0.08$ and therefore to an overhead of $f(l) \approx 18$ or a filter with 10^7 counters, which corresponds to a load of $l = 8 \cdot 10^{-4}$ and therefore to an overhead of $f(l) \approx 12$; this second choice is obviously better.

We can extend the analysis of the overhead per element, (Section V-B) to handle the case when the size of the two original sets are different. We obtain, for $k = 1$:

$$\begin{aligned} Ov &= \frac{u(|C_1| + |C_2|) + z_1 + z_2}{n_1 + n_2} = \\ &= \frac{2ul_1l_2 + H(l_1) + H(l_2)}{l_1 + l_2}. \end{aligned} \quad (12)$$

Interestingly, we note that if we are intersecting a set of size $n_1 = 10^5$, with a set of size $n_2 = 100$ (therefore much smaller), then choosing the standard Bloom filter size of 10^7 , leads to $Ov = 8.1$ bits per element, while the filter size of 10^5 , leads to $Ov = 2.2$ bits per element, which is much better. The maximum overhead of the scheme (found using numerical unconstrained optimization in Matlab) happens for $|S_1| = |S_2| \approx 4.2 \cdot \hat{n}_i$ (for $\hat{n}_i = 10^3, 10^5$) with a value of $Ov = 12.7$ bits per element.

VII. EXPERIMENTAL RESULTS

We implemented a prototype of the scheme algorithms (digest, check value and verification) and we used them to test the performance of the scheme. For simplicity, we tested separately the RSA accumulator computation and the remaining components of the scheme. Our experiments concentrate mostly on the intersection of two base sets, because we believe intersection to be the most challenging of the operations.

a) Cost of RSA accumulator: We implemented a C program that takes as input an RSA modulus N , a base a , and a set S , and outputs the corresponding RSA accumulator. We use the PARI library [17] for big number computation.

Observe that all the principals in this scheme need this computation: the source has to compute the accumulator of the whole original set S ; the directory has to compute two accumulators (witnesses), for the sets $S_1 \setminus (I \cup C_1)$ and $S_2 \setminus (I \cup C_2)$; the client has to compute also an accumulator operation, for sets $I \cup C_1$ and $I \cup C_2$ (using the witnesses as a base). Analogously to what is done for RSA signatures, note that the source can store the prime factors p_1, p_2 of the modulus N , taking care to keep them secret. Instead of using (1), it can, then, compute:

$$\text{Acc}(S) \bmod p_i = a^{e_1 e_2 \dots e_n \bmod (p_i - 1)} \bmod p_i$$

TABLE I
TIME REQUIRED TO COMPUTE AN RSA ACCUMULATOR

Set Size	Source time (sec)	Directory/Client time (sec)
100	0.021s	0.79
1000	0.033s	7.8
10000	0.17s	79
100000	1.5s	780

for $i = 1, 2$, and reconstruct $\text{Acc}(S)$ from $\text{Acc}(S) \bmod p_1$ and $\text{Acc}(S) \bmod p_2$ using Chinese remaindering.

We computed the accumulators for sets of different sizes. We employed a 512-bit RSA modulus and 962-bit representatives, generated with SHA-1 [16]. We ran the experiments on a 3 GHz Intel P4 with 2 GBytes RAM. Table I shows the effective CPU time required for the computation of an accumulator of one set as a function of the set size. For each set size, we repeated the experiment for 10 randomly generated sets of strings representing small integers; times were averaged. There are two columns: the first one for the source (which can employ the optimization mentioned above), the second for the directory and the client (which cannot).

As expected, the cost grows linearly with the set size. Note that the computation for the source is very efficient. Also note that, in practice, the costs at the client would be much lower because the client performs its computations on sets that are usually much smaller than the source sets.

b) Cost of Bloom filters and check elements: We created an unoptimized Perl implementation of the three algorithms that comprise the intersection scheme (digest, check value and verification) and used it to measure the communication overhead (encoded size of Bloom filters and check elements) and the computation overhead (CPU time to perform the algorithm). Our Perl program calls `arith_coder` [18] to perform Bloom filter compression.

We ran the set of experiments on the host described in the previous paragraph. Each experiment runs the four steps of the scheme (two digests, one intersection with check value computation, one verification) on a pair of random sets S_1, S_2 , of sizes n_1, n_2 , with an intersection of size q . The experiment is run ten times on different pairs of sets and the results are averaged. For each experiment the number of counters m of the Bloom filters is chosen in order to obtain a value l_1 of the load of $\text{Bl}(S_1)$ according to Table II; this is the value that minimizes (12), for the given value of n_2/n_1 . In all runs, the size of the intersection is $q = 0.01n_2$.

Table III shows the values of the running time for the

TABLE II

OPTIMAL BLOOM FILTER LOAD FOR DIFFERENT RATIOS OF THE SIZES OF THE TWO SOURCE SETS

n_2/n_1	l_1
1	0.010
0.1	0.065
0.01	0.51
0.001	2.9

TABLE III

EXECUTION TIMES OF THE SCHEME ALGORITHMS (ACCUMULATOR EXCLUDED) IN SECONDS

n_1	n_2/n_1	dig. S_1	dig. S_2	inters.	verif.
1000	1.000	0.449	0.446	0.137	0.413
	0.100	0.169	0.151	0.112	0.168
	0.010	0.126	0.105	0.125	0.126
	0.001	0.122	0.099	0.107	0.111
10000	1.000	3.400	3.630	0.681	2.720
	0.100	0.881	0.651	0.421	0.671
	0.010	0.389	0.167	0.346	0.275
	0.001	0.304	0.098	0.312	0.159
100000	1.000	34.10	35.20	6.520	24.10
	0.100	7.570	4.940	3.470	5.180
	0.010	3.300	0.736	3.130	1.900
	0.001	2.430	0.201	2.710	0.705

TABLE IV

COMMUNICATION OVERHEAD (ACCUMULATOR EXCLUDED)

n_1	n_2/n_1	Absolute overhead (bytes)	Relative overhead (bits)
1000	1.000	3431	13.7
	0.100	1725	12.5
	0.010	1042	8.26
	0.001	478	3.82
10000	1.000	27533	11.0
	0.100	12450	9.06
	0.010	7025	5.56
	0.001	3158	2.52
100000	1.000	269572	10.8
	0.100	118924	8.65
	0.010	64699	5.12
	0.001	27334	2.18

TABLE V

OVERHEAD OF COMPLEX QUERIES

Query	Absolute overhead (bytes)	Saving factor
$S_1 \setminus S_2$	4457	13.5
$(S_5 \cap S_6) \setminus (S_1 \cup S_2)$	209633	109.0
$(S_1 \cap S_4) \cap S_3$	17880	17.9
$(S_5 \setminus S_4) \cap S_2$	105062	11.8
$(S_3 \cup (S_5 \cap S_2)) \setminus S_5$	122927	9.3

computation, *not including the time to compute the RSA accumulators*, of (column 3 to 6): the digest of S_1 , the digest of S_2 , the operations performed by the directory (intersection and check elements) and verification.

In the third column of Table IV, we report the absolute overhead of the scheme (also accumulators excluded), i.e. the total number of bytes required to encode the two compressed Bloom filters of the source sets and the hashes of the check elements ($u = 160$). In the fourth column, we show the relative overhead $Over$ as defined in Section V-B (absolute overhead divided by $n_1 + n_2$) expressed in bits per element.

Note that if the bigger set is large enough ($n_1 \geq 10000$), the measured overhead is essentially what is yielded by the analysis (12). The slightly larger overhead for $n_1 = 1000$ is due to the non-optimality of the compression algorithm, which is more evident for small input sizes.

We also ran experiments (not reported in the tables) for different values of q and we noticed that this parameter has little influence on the computation times and the scheme overhead.

c) Overhead of complex queries: Finally, we ran a simple experiment to test the complete scheme, including the extensions in Sections V-C through VI-C. We generated six random sets (S_1, \dots, S_6), of sizes resp. 1000, 2000, 5000, 10000, 50000, 100000, from a universe of 500000 elements (small integers); then we executed our scheme for five sample queries, using our Perl implementation. Results are shown in Table V: the second column shows the measured absolute overhead of our scheme, while the third column shows the ratio between the absolute overhead of the trivial scheme and the value in the second column. We note that, in all cases, our scheme offers significant savings.

VIII. CONCLUSIONS

In this paper we have formally defined the notion of *secure set operations*. Secure set operations allow any principal to perform a set operation on a pair of trusted sets, and to provide a proof of the result's validity. We then show an efficient construction of a set operation scheme that, recursively, allows secure operations on certified results. We demonstrate, through analysis and experiments,

that our scheme produces certificates that are a *factor of 9 to 100* smaller than the trivial scheme, in which a signed copy of all the original sets is used as a certificate. To the best of our knowledge, no other scheme that solves this problem is known. A more exhaustive description of our algorithms can be found in [16].

We believe that this scheme has an important application in the context of efficient searches in P2P systems. Clients in such systems can profit from reusing the results of previous queries, which are cached at untrusted peers. With a secure set operation scheme, a client can retrieve such a result from an untrusted peer, and use the corresponding certificate to verify that the data was not polluted during the computation. Therefore, if the client trusts the sources, then it can also trust the retrieved data.

Finally, we note that our results are not specific to result caching. The results hold for any type of set, and therefore may have other applications.

REFERENCES

- [1] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan, “Chord: A scalable peer-to-peer lookup service for internet applications,” in *Proceedings of the ACM SIGCOMM '01 Conference*, San Diego, California, August 2001.
- [2] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Shenker, “A scalable content addressable network,” in *Proceedings of ACM SIGCOMM 2001*, 2001.
- [3] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, “Tapestry: An infrastructure for fault-tolerant wide-area location and routing,” Tech. Rep. UCB/CSD-01-1141, UC Berkeley, Apr. 2001.
- [4] Antony Rowstron and Peter Druschel, “Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems,” in *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, 2001, pp. 329–350.
- [5] Bobby Bhattacharjee, Sudarshan Chawathe, Vijay Gopalakrishnan, Pete Keleher, and Bujor Silaghi, “Efficient peer-to-peer searches using result-caching,” in *The 2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*, February 2003.
- [6] Niko Baric and Birgit Pfitzmann, “Collision-free accumulators and fail-stop signature schemes without trees,” in *EUROCRYPT*, 1997, pp. 480–494.
- [7] P. Devanbu, M. Gertz, C. Martel, and S. Stubblebine, “Authentic third-party data publication,” in *IFIP Conf. Database Security*, 2000.
- [8] M. Goodrich, A. Schwerin, and R. Tamassia, “An efficient dynamic and distributed cryptographic accumulator,” Tech. Rep., Johns Hopkins Information Security Institute, 2000.
- [9] M.T. Goodrich and R. Tamassia, “Efficient authenticated dictionaries with skip lists and commutative hashing,” Tech. Rep., Johns Hopkins Information Security Institute, 2000.
- [10] M.T. Goodrich, R. Tamassia, N. Triandopoulos, and R. Cohen, “Authenticated data structures for graph and geometric searching,” in *CT-RSA*, 2003.
- [11] Jan Camenisch and Anna Lysyanskaya, “Dynamic accumulators and application to efficient revocation of anonymous credentials,” in *CRYPTO*, 2002.
- [12] Rosario Gennaro, Shai Halevi, and Tal Rabin, “Secure hash-and-sign signatures without the random oracle,” in *EUROCRYPT*, 1999, pp. 123–139.
- [13] Burton H. Bloom, “Space/time trade-offs in hash coding with allowable errors,” *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [14] Li Fan, Pei Cao, Jussara Almeida, and Andrei Z. Broder, “Summary cache: a scalable wide-area web cache sharing protocol,” *IEEE/ACM Transactions on Networking (TON)*, vol. 8, no. 3, pp. 281–293, 2000.
- [15] M. Mitzenmacher, “Compressed bloom filters,” in *20th Annual ACM Symposium on Principles of Distributed Computing*, 2001, pp. 144–150, To appear in *IEEE/ACM Trans. on Networking*.
- [16] Ruggero Morselli, Samrat Bhattacharjee, Jonathan Katz, and Pete Keleher, “Trust preserving set operations,” Tech. Rep. CS-TR-4499, University of Maryland, College Park, July 2003.
- [17] The PARI Group, Bordeaux, *PARI/GP, Version 2.1.5*, 2000, available from <http://www.parigp-home.de/>.
- [18] John Carpinelli et al., “arith_coder: Word, character, integer, and bit based compression using arithmetic coding,” http://www.cs.mu.oz.au/~alistair/arith_coder/, 1999.

APPENDIX

PROOF OF THEOREM 1

Fix a counter index j . What is the probability that it contains a gap? Assume that the sets S_1 and S_2 are randomly generated sets of n_1, n_2 elements respectively and that $|I| = q$. If we call n the maximum set size that our scheme supports, then $n_1, n_2 \leq n$.

$$\begin{aligned} \Pr[\text{gap at } j] &= \Pr[\text{Bl}(S_1 \setminus I)_j > 0 \wedge \text{Bl}(S_2 \setminus I)_j > 0] = \\ &= \Pr[\text{Bl}(S_1 \setminus I)_j > 0] \cdot \Pr[\text{Bl}(S_2 \setminus I)_j > 0] \end{aligned}$$

where we used the fact that the elements of $S_1 \setminus I$ and $S_2 \setminus I$ are distinct, therefore the two events are independent. Remembering the derivation in [14]

$$\begin{aligned} \Pr[\text{Bl}(S_1 \setminus I)_j > 0] &= [1 - (1 - \frac{1}{m})^{k(n_1 - q)}] \\ &\simeq [1 - e^{-\frac{k(n_1 - q)}{m}}] \leq [1 - e^{-\frac{kn_1}{m}}] \leq [1 - e^{-\frac{kn}{m}}]. \end{aligned}$$

Any gap j requires $\text{Bl}(S_1 \setminus I)_j$ check elements for S_1 and $\text{Bl}(S_2 \setminus I)_j$ check elements for S_2 . In the worst case, no check element will be useful to cover more than one gap. Therefore, the expected number of check elements in C_1 needed for an index j is given by:

$$\begin{aligned} \Pr[\text{gap in } j] \mathbb{E}[\text{Bl}(S_1 \setminus I)_j | \text{gap at index } j] &= \\ \Pr[\text{Bl}(S_1 \setminus I)_j > 0] \Pr[\text{Bl}(S_2 \setminus I)_j > 0] & \\ \mathbb{E}[\text{Bl}(S_1 \setminus I)_j | \text{Bl}(S_1 \setminus I)_j > 0] &= \\ \Pr[\text{Bl}(S_1 \setminus I)_j > 0] \Pr[\text{Bl}(S_2 \setminus I)_j > 0] & \\ \frac{\mathbb{E}[\text{Bl}(S_1 \setminus I)_j]}{\Pr[\text{Bl}(S_1 \setminus I)_j > 0]} &= \\ \Pr[\text{Bl}(S_2 \setminus I)_j > 0] \mathbb{E}[\text{Bl}(S_1 \setminus I)_j] &\leq (1 - e^{-l_2}) l_1 \\ &\leq l_1 l_2. \end{aligned}$$

The claim follows from linearity of expectation.