

Private Revocation Test using Oblivious Membership Evaluation Protocol

Hiroaki Kikuchi

Tokai University

Dept. of Information Media Technology,

1117 Kitakaname, Hiratsuka, Kanagawa, 259-1292, Japan

kikn@tokai.ac.jp

Abstract This paper presents a cryptographic protocol for the authenticated dictionary, namely, an untrusted directory provides a verifiable answer to a membership query for a given element. In our protocol, a user is able to retrieve whether or not a target element belongs to a database that the directory has without revealing which element he/she wishes to know against the untrusted directory. Our protocol requires linear exponentiations to the number of elements in the database, but achieves a constant size communication complexity between a user and a directory. The privacy of query is assured under the Φ -hiding assumption introduced by Cachin.

1 Introduction

1.1 The PKI Issue

Certificate revocation is a current topic of interest in public-key infrastructure (PKI). Traditionally, a list of revoked certificates (CRL) has been used to represent the periodic distribution of revoked information. To improve the bandwidth consumption of the entire CRL transmission, some mail agents have begun supporting an online protocol for providing users the status of a target certificate alone, instead of the full CRL. The Online Certificate Status Protocol (OCSP)[6] is the standard protocol now in common use. There have been several attempts to improve the efficiency and security of CRLs. Kocher proposed a hash-tree based revocation protocol known as CRT[8], Micali presented a linear linking scheme with $O(1)$ communication cost (CRS)[7], and Naor and Nissim formalized the problem as an authenticated dictionary [9] in which a B-tree is used to balance the tree while the tree itself is skewed while updating the database.

1.2 Privacy Issues

As these online protocols are now in widespread use, a new privacy issue has arisen. The OCSP method uses the following steps. Each time a digitally signed mail is received, then the mail agent picks up a certificate from the mail and automatically sends a query to check if the certificate is revoked to a server specified in the certificate. Hence, the server, known as the *CRL distribution point*, acquires the significant statistics of the PKI – who sends a message to whom, how often, and, even worse, a digital signature, which is often used when we send significant messages whose privacy we wish to preserve the most.

1.3 Privacy Information Retrieval

To overcome the privacy issues of revoked certificates, the private information retrieval (PIR) method is a suitable technique for a user to be able to retrieve a target data item from a database while hiding the identity of the target item from the server. The notion of a PIR was introduced by Chor, Goldreich, Kushievitz, and Sudan [4], and has already improved retrieval in terms of its communication and computation costs. One of the recent results by Beimel, Ishai, and Malkin [5] archives, for a given constant, $k \geq 2$, and the number of items in a database n , a k -server protocol with $O(n^{1/(2k-1)})$ communication, and $O(n/\log^{2k-2} n)$ computations at the server. The servers, however, are considered as untrustworthy parties in the PKI model because servers must be online and, thus, have greater chance of being compromised by an intruder. Therefore, the behavior on the server side is not guaranteed to be correct. In addition, the average user may have poor computational power and narrow bandwidth, with even just one server. Thus, a single server protocol making the cost at user side as small as possible is preferable for solving the CRL distribution problem.

1.4 Our Contribution

In this paper, we present a simple solution to the problem. Given an element, $x \in X$, a user performs a membership test if x is in a subset $L = \{x_1, \dots, x_n\} \subset X$, requesting a query for a single non-trusted server who manages L steps without revealing x to the server. Our proposed protocol achieves a single server PIR with an optimal communication cost of $O(1)$ between a server and a user, and an optimum computation cost of $O(1)$ at the user side. To prevent the server from answering an improper response, a verification protocol that authorizes the answer from an authority is also provided.

2 Preliminaries

2.1 The PKI Model and Requirements

We have three types of parties: The *source* S or *Certification Authority* (CA), which is a trusted party that certifies the list of revoked certificates, *directories* D is non-trusted party who maintains the list and answers the questions that a target certificate is still active, instead of CA, and *users*, U , who wish to keep in touch with the current status of the certificates via the non-trusted D .

The CA is a source of information of revoked certificates and has the replication of the information distributed among directories. The directories of D s work as carriers of the revoked information and are thus not responsible for the integrity of the database provided from the source. In terms of security, the directories have no secret information inside so that, even if one of directories is compromised, any rebuild of PKI is not necessary. The directories have a powerful computational power, e.g., the state-of-the-art CPUs, a secure coprocessor, and broad-bandwidth connections to each party. Since the directories are widely distributed over the network, we assume the risk that some of directories might perform an analysis of the access log from the end users using the data mining techniques. A user U communicates with one of the directories and checks if a target certificate is revoked or not and examines the integrity of responses from the directory server. We assume that some of the users may have limited computational power and a poor link of limited bandwidth. (In particular, this can happen when the user is mobile and with a PDA).

Oblivious Membership Evaluation:

Let X be the universal set of identities of certificate (64-bit serial numbers are often used in actual services), and $L = \{x_1, \dots, x_n\}$ be a subset of X . The S gets L distributed among directories D . Given an element $x \in X$, U performs a membership query to D whether or not, $x \in L$ without revealing x to D .

The requirements of oblivious membership evaluation should satisfy are as follows:

1. **Privacy of query.** From a membership query of $x \in L$, D learns no information about x .
2. **Authenticity of source.** From the response from D , U verifies that the result of membership query is authorized by S and that D follows the steps properly.
3. **Efficiency.** The sizes of both query and answer should be independent of the number of PKI users, to which the size of CRL n seems to be proportional, and we want the sizes to be as small as possible. The computational costs at users should be also minimized.

2.2 Dynamic Accumulator

The RSA accumulator is proposed by Benaloh and de Mare[10], where a set of values are accumulated into a single object for which a witness that a given

value was incorporated into it is provided. Camenish and Lysyanskaya improves the RSA accumulator so that dynamic operations of insertion and deletion are feasible with independent cost from the number of values[12]. Goodrich, Tamassia, and Hasic show the pre-computations of witness reduces the computation overhead at the directory with the cost of communication consumption[11].

Informally, the RSA accumulator works as follows. The source picks strong primes p and q and publishes $N = pq$. Let L be a set of primes $\{x_1, \dots, x_n\}$, representing identities (of the revoked certificates in PKI). The source then computes *accumulator*

$$A = a^{x_1 x_2 \dots x_n} \pmod{N},$$

where a is a public constant that is relatively prime to N and publishes A together with digital signature $\sigma_S(A)$ on A . To prove an element $x_i \in L$, the directory computes *witness*

$$A_i = a^{x_1 \dots x_{i-1} x_{i+1} \dots x_n} \pmod{N}.$$

The user verifies witness by $A_i^{x_i} \pmod{N} \stackrel{?}{=} A$. Under the strong RSA assumption[12], the directory, which does not have the knowledge of factorization of N , is able to compute the witness A_i only when x_i belongs to L .

2.3 Φ -Hiding Assumption

Cachin present an efficient secure auction protocol that an oblivious party blindly compares two inputs bit-by-bit under the the ϕ -hiding assumption (Φ HA)[13]. Informally, the Φ HA states that it is computationally infeasible to decide whether a given prime divides $\phi(N)$, where m is a composite number of unknown factorization.

We say modulus m *hides* a prime p if N is a composite number $p'q'$ such that $p' = 2pp_1 + 1$ and $q' = 2q_1 + 1$ with primes p_1, q_1 . Note that N hides p if and only if $p|\phi(N)$. The Φ HA states that, for a randomly chosen $N \in Z_N^*$ and primes p_0, p_1 such that N hides p_0 but does not hide p_1 , the (N, p_0) and (N, p_1) is computationally indistinguishable.

An integer x is a p -th *residue* modulo m if there exists an α such that $\alpha^p = x \pmod{m}$. Let $R_N(p)$ denote a set of all p -th residues in Z_N^* . Then, note that only the party that knows the factorization of N and thus $\phi(N)$ is able to test if any given integer is a p -th residue by

$$a^{\phi(N)/p} \equiv 1 \pmod{m},$$

which holds if a is a p -th residue modulo m .

2.4 Proof of Conjunctive Knowledge

Cramer, Cramer, Damgard, and Schoenmakers presents an efficient zero-knowledge proof of conjunctive propositions [1]. By $PK\{(\alpha) : y_1 = g_1^\alpha \wedge y_2 = g_2^\alpha\}$, we denote a proof of knowledge of discrete logarithms of elements y_1 and y_2 to the

bases g_1 and g_2 . Selecting random numbers r_1 and $r_2 \in Z_q$, a prover sends $t_1 = g_1^{r_1}$ and $t_2 = g_2^{r_2}$ to a verifier, who then sends back a random challenge $c \in \{0, 1\}^k$. The prover shows $s = r - cx \pmod{q}$, which should satisfy both $g_1^s y_1^c = t_1$ and $g_2^s y_2^c = t_2$.

3 Oblivious Membership Evaluation

3.1 Overview

Our construction is based on Φ HA in order for users to blindly query a membership to a directory that has the list L . A user generates a modulus m that hides a prime x specifying the identity of a given certificate, and then sends a query consisting of non x -th residue c . The directory D then raises c to the power of all primes in S modulo m and sends the answer back to U , who then performs an x -th residue test using secret knowledge of factorization of m . In addition, we need a verification protocol to prevent a dishonest directory from cheating users. The witness in RSA accumulator cannot be applied here because the directory does not know which element is to be tested. Instead, we employ a zero-knowledge proof technique to show that the directory has raised a base to the power exactly the same exponents to that used by accumulator A .

3.2 Accumulator Setup

We begin with a set up protocol in which a source S notifies to the directories the list of currently revoked certificates.

1. The S picks strong primes P and Q and publishes $N = PQ$. For the list of revoked certificates $L = \{x_1, x_2, \dots, x_n\}$, where x_i are small primes corresponding identities of revoked certificates, S computes accumulation

$$A = a^{x_1 x_2 \dots x_n} \pmod{N},$$

where a is a public constant that is relatively prime to N and publishes L, A, a together with a digital signature $\sigma_S(A, a, t)$, where t is the current time interval.

2. On receiving the list L and accumulator A periodically, every directory D updates the current (at a time t) database by L after it verifies the digital signature and accumulator $a^{x_1 x_2 \dots x_n} = A \pmod{N}$.

3.3 Membership Test

Given a certificate to be examined, a user performs the following membership test protocol with one of the directories.

1. Given a target certificate specified by prime x , U chooses strong primes p and q such that $m = pq$ hides prime x . U picks an integer c that is not

p -th residue modulo m . U sends a query of the form (c, m) to one of the directories, D .

- Then, D computes an answer

$$z = c^{x_1 x_2 \cdots x_n} \pmod{m}$$

and responds to U the answer z together with the accumulator A and digital signature $\sigma_S(A, a, t)$.

- Finally, U locally performs the membership test

$$z^{\phi(m)/x} \pmod{m} = \begin{cases} 1 & \text{if } x \in L, \\ 1^{1/x} & \text{otherwise.} \end{cases}$$

Note that answer z becomes the x -th residue when there is an element in L that is equal to the target x .

3.4 Authenticity of the Source

To prevent a dishonest D from cheating users with improperly computed z , we require D to provide the proof of accumulating every element L into A by the form

$$PK\{\beta : a^\beta = A \wedge c^\beta = z\},$$

where private information β is ℓ defined by $\ell = x_1 x_2 \cdots x_n$ (note that this is not a modular multiplication), for which both $z = c^\ell$ and $A = a^\ell$ are satisfied. In other words, ℓ is a witness for which accumulator A is consistent with the answer z . Since D does not know the factorization of N nor m , we need the modified version of the proof of conjunctive knowledge mentioned in Section 2.4.

- The D randomly picks r that is properly large (but is less than N and m) and computes

$$T = a^r \pmod{N}, \quad V = c^r \pmod{m}.$$

For T and V , D applies a secure cryptographic hash function H with properly large range to obtain a challenge $d = H(T||V)$, and computes (not modular arithmetic)

$$s = r + d\ell$$

and sends the proof (T, V, s) to U .

- Then, U computes $d = H(T||V)$ and verifies that

$$\begin{aligned} a^s / A^d &\stackrel{?}{=} T \pmod{N}, \\ c^s / z^d &\stackrel{?}{=} V \pmod{m}. \end{aligned}$$

4 Evaluation

4.1 Security

Under the assumption of a secure digital signature scheme used by the source, the accumulator A at the time t is unable to be forged. Consider a dishonest directory that is trying to manipulate z to z' so that the membership test will fail for z' when x is in L . To convince users that the answer was correctly computed, the directory has to predict s that satisfies the above-mentioned equations for proof of knowledge. The probability of passing the test is negligibly small.

4.2 Privacy

If a malicious directory is able to determine which prime is hidden by a given m and c , it can immediately distinguish two composite numbers m_0 and m_1 that hide distinct primes, which contradicts the Φ -hiding assumption. Therefore, D is not able to learn the target x under the Φ HA. Moreover, D does not even know the result of the membership test at all.

4.3 Efficiency

The proposed scheme has the following performance:

- a size of query (c) sent from user to directory is $|m|$;
- a size of answer (z) sent from directory to user is $|m| + |N| + |\sigma|$ (without proof of knowledge);
- a size of proof (T, V, s) is $O(n)$ (since the magnitude of ℓ is linear to n);
- a number of modular exponentiations at the user is 1;
- a number of modular exponentiations at the directory is n .

Without the knowledge of $\phi(m)$, the size of ℓ increases with the number of elements in L ; thus, the verification at the last step in the scheme requires $O(n|m|)$ modular multiplications, which is impractically heavy when n is too large.

One more inefficiency we should address is the key generation cost to the user, who should always pick a new modulus that hides the given prime.

5 Conclusions

We have presented a protocol for oblivious membership evaluation using the Φ -hiding assumption. The proposed protocol is efficient in terms of directory-and-user communication with $O(1)$, preserves the privacy of a query as to which certificate is to be examined, and provides verification steps that result in the membership query being correctly computed. Future studies include an efficient

verification independent of n and an improvement of n -size modular exponentiations at the directory.

References

- [1] R. Cramer, I. Damgard, and B. Schoenmakers, "Proofs of partial knowledge and simplified design of witness hiding protocols," in Proc. of *CRYPTO '94*, pp.174-187, 1994.
- [2] J. Camenisch, and M. Michels, "Proving in zero-knowledge that a number is the product of two safe primes," in Proc. of *EUROCRYPT '99*, pp. 107-122, 1999.
- [3] R. Cramer, R. Gennaro, and B. Schoenmakers, "A secure and optimally efficient multi-authority election scheme," *EUROCRYPT 1997*.
- [4] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan, "Private information retrieval," in Proc. 36th IEEE Symposium on Foundations of Computer Science (FOCS), 1995.
- [5] Amos Beimel, Yuval Ishai, and Tal Malkin, "Reducing the servers computation in private information retrieval: pir with preprocessing," in Proc. CRYPTO '00, LNCS vol. 1880, pp. 550, 2000.
- [6] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams, "X.509 Internet public key infrastructure online certificate status protocol - OCSP," Internet RFC 2560, 1999.
- [7] S. Micali, "Efficient certificate revocation", Technical Report TM-542b, MIT Laboratory for Computer Science, 1996.
- [8] P. Kocher, "On certificate revocation and validation," in Proc. of Financial Cryptography'98, *Springer LNCS 1465*, pp. 172-177, 1998.
- [9] M. Naor, and K. Nissim, "Certificate revocation and certificate update," in Proc. of Seventh USENIX Security Symposium, pp. 217-228, 1998.
- [10] J. Benaloh, and M. de Mare, "One-way accumulators: A decentralized alternative to digital signatures," in Proc. of *EUROCRYPT*, LNCS vol. 839, Springer, pp. 216-233, 1994.
- [11] M. Goodrich, R. Tamassia, and J. Hasic, "An efficient dynamic and distributed cryptographic accumulator," in Proc. of Information Security Conference (ISC 2002), LNCS Vol. 2433, Springer, pp. 372-388, 2002.
- [12] J. Camenisch, and A. Lysyanskaya, "Dynamic accumulators and application to efficient revocation of anonymous credentials," in Proc. *CRYPTO 2002*, 2002.

- [13] C. Cachin, "Efficient private bidding and auctions with an oblivious third party," ACM Conference on Computer and Communications Security (CCS), pp. 120-127, 1999.