Collusion Resistant Broadcast Encryption With Short Ciphertexts and Private Keys

Dan Boneh[1*], Craig Gentry[2], and Brent Waters[1]

[1] Stanford University, {dabo,bwaters}@cs.stanford.edu
[2] DoCoMo USA Labs, cgentry@docomolabs-usa.com

Abstract. We describe two new public key broadcast encryption systems for stateless receivers. Both systems are fully secure against any number of colluders. In our first construction both ciphertexts and private keys are of constant size (only two group elements), for any subset of receivers. The public key size in this system is linear in the total number of receivers. Our second system is a generalization of the first that provides a tradeoff between ciphertext size and public key size. For example, we achieve a collusion resistant broadcast system for n users where both ciphertexts and public keys are of size $O(\sqrt{n})$ for any subset of receivers. We discuss several applications of these systems.

1 Introduction

In a broadcast encryption scheme [FN93] a broadcaster encrypts a message for some subset S of users who are listening on a broadcast channel. Any user in S can use his private key to decrypt the broadcast. However, even if all users outside of S collude they can obtain no information about the contents of the broadcast. Such systems are said to be collusion resistant. The broadcaster can encrypt to any subset S of his choice. We use n to denote the total number of users.

Broadcast encryption has several applications including access control in encrypted file systems, satellite TV subscription services, and DVD content protection. As we will see in Section 4 we distinguish between two types of applications:

- Applications where we broadcast to large sets, namely sets of size $n - r$ for $r \ll n$. The best systems [NNL01,HS02,GST04] achieve a broadcast message containing $O(r)$ ciphertexts where each user's private key is of size $O(\log n)$.
- Applications where we broadcast to small sets, namely sets of size t for $t \ll n$. Until now, the best known solution was trivial, namely encrypt the broadcast message under each recipient's key. This broadcast message contains t ciphertexts and each user's private key is of size $O(1)$.

In this paper we construct fully collusion secure broadcast encryption systems with short ciphertexts and private keys for *arbitrary* receiver sets. Our

* Supported by NSF.

constructions use groups with an efficiently computable bilinear map. Our first construction provides a system in which both the broadcast message and user private keys are of constant size (a precise statement is given in the next section). No matter what the receiver set is, our broadcast ciphertext contains only two group elements. Each user's private key is just a single group element. Thus, when broadcasting to small sets our system generates far shorter ciphertexts than the trivial solution discussed above. However, the public key size in this system is linear in the number of recipients. This is not a large problem in applications such as encrypted file systems where the receivers have access to a large shared storage medium in which the public key can be stored. For other applications, such as content protection, we need to minimize both public key and ciphertext size.

Our second system is a generalization of the first that enables us to tradeoff public key size for ciphertext size. One interesting parametrization of our scheme gives a system where both the public key and the ciphertext are of size $O(\sqrt{n})$. This means that we can attach the public key to the encrypted broadcast and still achieve ciphertext size of $O(\sqrt{n})$. Consequently, we obtain a fully collusion secure broadcast encryption scheme with $O(\sqrt{n})$ ciphertext size (for any subset of users) where the users have a constant size private key.

In Section 2 we define our security model and the complexity assumption we use. In Section 3 we describe our systems and prove their semantic security. In Section 4 we discuss in detail several applications for these systems. Finally, in Section 5 we describe how to make our systems chosen-ciphertext secure.

1.1 Related Work

Fiat and Naor [FN93] were the first to formally explore broadcast encryption. They presented a solution for n users that is secure against a collusion of t users and has ciphertext size of $O(t \log^2 t \log n)$.

Naor et al. [NNL01] presented a fully collusion secure broadcast encryption system that is efficient for broadcasting to all but a small set of revoked users. Their scheme is useful for content protection where broadcasts will be sent to all but a small set of receivers whose keys have been compromised. Their scheme can be used to encrypt to $n - r$ users with a header size of $O(r)$ elements and private keys of size $O(\log^2 n)$. Further improvements [HS02,GST04] reduce the private key size to $O(\log n)$. Dodis and Fazio [DF02] extend the NNL (subtree difference) method into a public key broadcast system for a small size public key.

Other broadcast encryption methods for large sets include Naor and Pinkas [NP00] and Dodis and Fazio [DF03] as well as [AMM99,TT01]. For some fixed t all these systems can revoke any $r < t$ users where ciphertexts are always of size $O(t)$ and private keys are constant size. By running $\log n$ of these systems in parallel, where the revocation bound of the i'th system is $t_i = 2^i$ (as in [YJCK04]), one obtains a broadcast encryption system with the same parameters as [GST04]. Private key size is $O(\log n)$ and, when revoking r users, ciphertext size is proportional to $2^{\lceil \log_2 r \rceil} = O(r)$. This simple extension to the

Naor and Pinkas system gives a broadcast system with similar parameters as the latest NNL derivative.

Wallner et al. [WHA97] and Wong [WGL98] independently discovered the logical-key-hierarchy scheme (LKH) for multicast group key management. Using these methods receivers maintain state and remain connected to receive key-update messages. The parameters of these schemes are improved in later work [CGI+99,CMN99,SM03]. Our broadcast system also gives a group key management method with short key update messages.

The security of our broadcast encryption relies on computational assumptions. Several other works [Sti97,ST98,SW98,GSY99,GSW00] explore broadcast encryption and tracing from an information theoretic perspective.

Boneh and Silverberg [BS03] show that n-linear maps give the ultimate fully collusion secure scheme with constant public key, private key, and ciphertext size. However, there are currently no known implementations of cryptographically useful n-linear maps for $n > 2$. Our results show that we can come fairly close using bilinear maps alone.

2 Preliminaries

We begin by formally defining public-key broadcast encryption systems. For simplicity we define broadcast encryption as a key encapsulation mechanism. We then state the complexity assumption needed for our proof of security.

2.1 Broadcast Encryption Systems

A broadcast encryption system is made up of three randomized algorithms:

Setup(n) Takes as input the number of receivers n. It outputs n private keys $d_1, \ldots, d_n$ and a public key PK.

Encrypt(S, PK) Takes as input a subset $S \subseteq \{1, \ldots, n\}$, and a public key PK. It outputs a pair (Hdr, K) where Hdr is called the header and $K \in \mathcal{K}$ is a message encryption key chosen from a finite key set $\mathcal{K}$. We will often refer to Hdr as the broadcast ciphertext.

Let M be a message to be broadcast that should be decipherable precisely by the receivers in S. Let C_M be the encryption of M under the symmetric key K. The broadcast consists of (S, Hdr, C_M). The pair (S, Hdr) is often called the full header and C_M is often called the broadcast body.

Decrypt$(S, i, d_i, \mathrm{Hdr}, PK)$ Takes as input a subset $S \subseteq \{1, \ldots, n\}$, a user id $i \in \{1, \ldots, n\}$ and the private key d_i for user i, a header Hdr, and the public key PK. If $i \in S$, then the algorithm outputs a message encryption key $K \in \mathcal{K}$. Intuitively, user i can then use K to decrypt the broadcast body C_M and obtain the message body M.

As usual, we require that the system be correct, namely that for all subsets $S \subseteq \{1, \dots, n\}$ and all $i \in S$,

if $(PK, (d_1, \dots, d_n)) \xleftarrow{R} \mathit{Setup}(n)$ and $(\mathrm{Hdr}, K) \xleftarrow{R} \mathit{Encrypt}(S, PK)$
then $\mathit{Decrypt}(S, i, d_i, \mathrm{Hdr}, PK) = K$.

We define chosen ciphertext security of a broadcast encryption system against a static adversary. Security is defined using the following game between an attack algorithm $\mathcal{A}$ and a challenger. Both the challenger and $\mathcal{A}$ are given n, the total number of users, as input.

Init. Algorithm $\mathcal{A}$ begins by outputting a set $S^* \subseteq \{1, \dots, n\}$ of receivers that it wants to attack.

Setup. The challenger runs $\mathit{Setup}(n)$ to obtain a public key PK and private keys $d_1, \dots, d_n$. It gives $\mathcal{A}$ the public key PK and all private keys d_j for which $j \notin S^*$.

Query phase 1. Algorithm $\mathcal{A}$ issues decryption queries $q_1, \dots, q_m$ adaptively where a decryption query consists of (u, S, Hdr) where $S \subseteq S^*$ and $u \in S$. The challenger responds with $\mathit{Decrypt}(S, u, d_u, \mathrm{Hdr}, PK)$.

Challenge. The challenger runs algorithm *Encrypt* to obtain $(\mathrm{Hdr}^*, K) \xleftarrow{R} \mathit{Encrypt}(S, PK)$ where $K \in \mathcal{K}$. Next, the challenger picks a random $b \in \{0, 1\}$. It sets $K_b = K$ and picks a random $K_{1-b} \in \mathcal{K}$. It then gives $(\mathrm{Hdr}^*, K_0, K_1)$ to algorithm $\mathcal{A}$.

Query phase 2. Algorithm $\mathcal{A}$ adaptively issues more decryption queries $q_{m+1}, \dots, q_{q_D}$ where $q_i = (u, S, \mathrm{Hdr})$ with $S \subseteq S^*$ and $u \in S$. The only constraint is that $\mathrm{Hdr} \neq \mathrm{Hdr}^*$. The challenger responds as in phase 1.

Guess. Algorithm $\mathcal{A}$ outputs its guess $b' \in \{0, 1\}$ for b and wins the game if $b = b'$.

Let $\mathsf{AdvBr}_{\mathcal{A},n}$ denote the probability that $\mathcal{A}$ wins the game when the challenger and $\mathcal{A}$ are given n as input.

Definition 1. *We say that a broadcast encryption system is (t, ϵ, n, q_D) CCA secure if for all t-time algorithms $\mathcal{A}$ that make a total of q_D decryption queries, we have that* $|\mathsf{AdvBr}_{\mathcal{A},n} - \frac{1}{2}| < \epsilon$.

The game above models an attack where all users not in the set S^* collude to try and expose a broadcast intended for users in S^* only. The set S^* is chosen by the adversary. Note that the adversary is non-adaptive; it chooses S^*, and obtains the keys for users outside of S^*, before it even sees the public key PK. An adaptive adversary could request user keys adaptively. We only prove security of our system in the non-adaptive settings described above. It is an open problem to build a broadcast encryption system with the performance of our system which is secure against adaptive adversaries. We note that similar formal definitions for broadcast encryption security were given in [BS03,DF03].

As usual, we define semantic security for a broadcast encryption scheme by preventing the attacker from issuing decryption queries.

Definition 2. *We say that a broadcast encryption system is (t, ϵ, n) semantically secure if it is $(t, \epsilon, n, 0)$ CCA secure.*

In Section 3 we first construct *semantically secure* systems with constant ciphertext and private key size. We come back to chosen ciphertext security in Section 5.

2.2 Bilinear Maps

We briefly review the necessary facts about bilinear maps and bilinear map groups. We use the following standard notation [Jou00,JN03,BF01]:

1. $\mathbb{G}$ and $\mathbb{G}_1$ are two (multiplicative) cyclic groups of prime order p;
2. g is a generator of $\mathbb{G}$.
3. $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_1$ is a bilinear map.

Let $\mathbb{G}$ and $\mathbb{G}_1$ be two groups as above. A bilinear map is a map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_1$ with the following properties:

1. For all $u, v \in \mathbb{G}$ and $a, b \in \mathbb{Z}$, we have $e(u^a, v^b) = e(u, v)^{ab}$, and
2. The map is not degenerate, i.e., $e(g, g) \neq 1$.

We say that $\mathbb{G}$ is a bilinear group if the group action in $\mathbb{G}$ can be computed efficiently and there exists a group $\mathbb{G}_1$ and an efficiently computable bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_1$ as above. Note that $e(,)$ is symmetric since $e(g^a, g^b) = e(g, g)^{ab} = e(g^b, g^a)$.

2.3 Complexity Assumptions

Security of our system is based on a complexity assumption called the bilinear Diffie-Hellman Exponent assumption (BDHE). This assumption was previously introduced in [BBG05].

Let $\mathbb{G}$ be a bilinear group of prime order p. The ℓ-BDHE problem in $\mathbb{G}$ is stated as follows: given a vector of $2\ell + 1$ elements

$$\left(h, g, g^{\alpha}, g^{(\alpha^2)}, \ldots, g^{(\alpha^\ell)}, g^{(\alpha^{\ell+2})}, \ldots, g^{(\alpha^{2\ell})}\right) \in \mathbb{G}^{2\ell+1}$$

as input, output $e(g, h)^{(\alpha^{\ell+1})} \in \mathbb{G}_1$. Note that the input vector is missing the term $g^{(\alpha^{\ell+1})}$ so that the bilinear map seems to be of little help in computing the required $e(g, h)^{(\alpha^{\ell+1})}$.

As shorthand, once g and α are specified, we use g_i to denote $g_i = g^{(\alpha^i)} \in \mathbb{G}$. An algorithm $\mathcal{A}$ has advantage ϵ in solving ℓ-BDHE in $\mathbb{G}$ if

$$\Pr\left[\mathcal{A}\left(h, g, g_1, \ldots, g_\ell, g_{\ell+2}, \ldots, g_{2\ell}\right) = e(g_{\ell+1}, h)\right] \geq \epsilon$$

where the probability is over the random choice of generator g in $\mathbb{G}$, the random choice of h in $\mathbb{G}$, the random choice of α in $\mathbb{Z}_p$, and the random bits used by $\mathcal{A}$.

The decisional version of the ℓ-BDHE problem in $\mathbb{G}$ is defined analogously. Let $\boldsymbol{y}_{g,\alpha,\ell} = (g_1, \ldots, g_\ell, g_{\ell+2}, \ldots, g_{2\ell})$. An algorithm $\mathcal{B}$ that outputs $b \in \{0,1\}$ has advantage ϵ in solving decision ℓ-BDHE in $\mathbb{G}$ if

$$\left| \Pr\left[\mathcal{B}(g, h, \boldsymbol{y}_{g,\alpha,\ell}, e(g_{\ell+1}, h)) = 0\right] - \Pr\left[\mathcal{B}(g, h, \boldsymbol{y}_{g,\alpha,\ell}, T) = 0\right] \right| \geq \epsilon$$

where the probability is over the random choice of generators g, h in $\mathbb{G}$, the random choice of α in $\mathbb{Z}_p$, the random choice of $T \in \mathbb{G}_1$, and the random bits consumed by $\mathcal{B}$. We refer to the distribution on the left as $\mathcal{P}_{BDHE}$ and the distribution on the right as $\mathcal{R}_{BDHE}$.

Definition 3. *We say that the (decision) (t, ϵ, ℓ)-BDHE assumption holds in $\mathbb{G}$ if no t-time algorithm has advantage at least ϵ in solving the (decision) ℓ-BDHE problem in $\mathbb{G}$.*

Occasionally we drop the t and ϵ and refer to the (decision) ℓ-BDHE in $\mathbb{G}$. We note that the ℓ-BDHE assumption is a natural extension of the bilinear-DHI assumption previously used in [BB04,DY05]. Furthermore, Boneh et al. [BBG05] show that the ℓ-BDHE assumption holds in generic bilinear groups [Sho97].

3 Construction

We are now ready to present our broadcast encryption system. We first present a special case system where ciphertexts and private keys are always constant size. The public key grows linearly with the number of users. We then present a generalization that allows us to balance the public key size and the ciphertext size. Private keys are still constant size. We prove security of this general system.

3.1 A Special Case

We begin by describing a broadcast encryption system for n users where the ciphertexts and private keys are constant size. The public key grows linearly in the number of users.

Setup(n)**:** Let $\mathbb{G}$ be a bilinear group of prime order p. The algorithm first picks a random generator $g \in \mathbb{G}$ and a random $\alpha \in \mathbb{Z}_p$. It computes $g_i = g^{(\alpha^i)} \in \mathbb{G}$ for $i = 1, 2, \ldots, n, n+2, \ldots, 2n$. Next, it picks a random $\gamma \in \mathbb{Z}_p$ and sets $v = g^\gamma \in \mathbb{G}$. The public key is:

$$PK = (g, g_1, \ldots, g_n, g_{n+2}, \ldots, g_{2n}, v) \in \mathbb{G}^{2n+1}$$

The private key for user $i \in \{1, \ldots, n\}$ is set as: $d_i = g_i^\gamma \in \mathbb{G}$.
Note that $d_i = v^{(\alpha^i)}$. The algorithm outputs the public key PK and the n private keys $d_1, \ldots, d_n$.

Encrypt(S, PK)**:** Pick a random t in $\mathbb{Z}_p$ and set $K = e(g_{n+1}, g)^t \in \mathbb{G}_1$.
The value $e(g_{n+1}, g)$ can be computed as $e(g_n, g_1)$. Next, set

$$\text{Hdr} = \left(g^t, \;\; (v \cdot \prod_{j \in S} g_{n+1-j})^t \right) \in \mathbb{G}^2$$

and output the pair (Hdr, K).

Decrypt$(S, i, d_i, \text{Hdr}, PK)$**:** Let $\text{Hdr} = (C_0, C_1)$ and recall that $d_i \in \mathbb{G}$. Then, output

$$K = e(g_i, C_1) \;\; / \;\; e(d_i \cdot \prod_{\substack{j \in S \\ j \neq i}} g_{n+1-j+i}, \;\; C_0)$$

Note that a private key is only one group element in $\mathbb{G}$ and the ciphertext, Hdr, is only two group elements. Furthermore, since $e(g_{n+1}, g)$ can be precomputed, encryption requires no pairings. Nevertheless, the system is able to broadcast to any subset of users and is fully collusion resistant. We prove security in Section 3.3 where we discuss a more general system.

We verify that the system is correct — i.e., that the decryption algorithm works correctly — by observing that, for any $i \in S$, the quotient of the terms

$$e(g_i, C_1) = e(g,g)^{\alpha^i \cdot t(\gamma + \sum_{j \in S} \alpha^{n+1-j})} = e(g,g)^{t(\gamma\alpha^i + \sum_{j \in S} \alpha^{n+1-j+i})} \quad \text{and}$$

$$e(C_0, \;\; d_i \cdot \prod_{\substack{j \in S \\ j \neq i}} g_{n+1-j+i}) = e(g,g)^{t \cdot (\gamma\alpha^i + \sum_{j \neq i}^{j \in S} \alpha^{n+1-j+i})}$$

is $K = e(g_{n+1}, g)^t = e(g,g)^{t\alpha^{n+1}}$, as required.

Efficient implementation. For any large number of receivers, decryption time will be dominated by the $|S| - 2$ group operations needed to compute $\prod_{j \neq i}^{j \in S} g_{n+1-j+i}$. However, we observe that if the receiver had previously computed the value $w = \prod_{j \neq i}^{j \in S'} g_{n+1-j+i}$ for some receiver set S' that is similar to S then, the receiver can compute $\prod_{j \neq i}^{j \in S} g_{n+1-j+i}$ with just δ group operations using the cached value w, where δ is the size of the set difference between S and S'.

This observation is especially useful when the broadcast system is intended to broadcast to large sets, i.e. sets of size $n - r$ for $r \ll n$. The private key d_i could include the value $\prod_{j \neq i}^{j \in [1,n]} g_{n+1-j+i} \in \mathbb{G}$ which would enable the receiver to decrypt using only r group operations. Furthermore, user i would only need r elements from the public key PK.

We note that computation time for encryption will similarly be dominated by the $|S| - 1$ group operations to compute $\prod_{j \in S} g_{n+1-j}^t$ and similar performance optimizations (e.g. precomputing $\prod_{j=1}^{n} g_{n+1-j}$) apply. We also note that in the secret key settings (where the encryptor is allowed to keep secret information) the encryptor need only store (g, v, α) as opposed to storing the entire public key vector.

3.2 A General Construction

Next, we present a more general broadcast encryption system. The idea is to run A parallel instances of the system in the previous section where each instance can broadcast to at most $B < n$ users. As a result we can handle as many as $n = AB$ users. However, we substantially improve performance by sharing information between the A instances. In particular, all instances will share the same public key values $g, g_1, \ldots, g_B, g_{B+2}, \ldots, g_{2B}$.

This generalized system enables us to tradeoff the public key size for ciphertext size. Setting $B = n$ gives the system of the previous section. However, setting $B = \lfloor\sqrt{n}\rfloor$ gives a system where both public key and ciphertext size are about $\sqrt{n}$ elements. Note that either way, the private key is always just one group element.

For fixed positive integer B, the B-broadcast encryption system works as follows:

Setup$_B(n)$**:** The algorithm will set up $A = \lceil \frac{n}{B} \rceil$ instances of the scheme. Let $\mathbb{G}$ be a bilinear group of prime order p. The algorithm first picks a random generator $g \in \mathbb{G}$ and a random $\alpha \in \mathbb{Z}_p$. It computes $g_i = g^{(\alpha^i)} \in \mathbb{G}$ for $i = 1, 2, \ldots, B, B+2, \ldots, 2B$. Next, it picks random $\gamma_1, \ldots, \gamma_A \in \mathbb{Z}_p$ and sets $v_1 = g^{\gamma_1}, \ldots, v_A = g^{\gamma_A} \in \mathbb{G}$. The public key is:

$$PK = (g, g_1, \ldots, g_B, g_{B+2}, \ldots, g_{2B}, v_1, \ldots, v_A) \in \mathbb{G}^{2B+A}$$

The private key for user $i \in \{1, \ldots, n\}$ is defined as follows: write i as $i = (a-1)B + b$ for some $1 \le a \le A$ and $1 \le b \le B$ (i.e. $a = \lceil i/B \rceil$ and $b = i \bmod B$). Set the private key for user i as:

$$d_i = g_b^{\gamma_a} \in \mathbb{G} \qquad \text{(note that} \quad d_i = v_a^{(\alpha^b)} \text{)}$$

The algorithm outputs the public key PK and the n private keys $d_1, \ldots, d_n$.

Encrypt(S, PK)**:** For each $\ell = 1, \ldots, A$ define the subsets $\hat{S}_\ell$ and S_ℓ as

$$\hat{S}_\ell = S \cap \{(\ell-1)B+1, \ldots, \ell B\}, \qquad S_\ell = \{x - \ell B + B \mid x \in \hat{S}_\ell\} \subseteq \{1, \ldots, B\}$$

In other words, $\hat{S}_\ell$ contains all users in S that fall in the ℓ'th interval of length B and S_ℓ contains the indices of those users relative to the beginning of the interval. Pick a random t in $\mathbb{Z}_p$ and set $K = e(g_{B+1}, g)^t \in \mathbb{G}_1$. Set

$$\text{Hdr} = \left(g^t,\ (v_1 \cdot \prod_{j \in S_1} g_{B+1-j})^t,\ \ldots,\ (v_A \cdot \prod_{j \in S_A} g_{B+1-j})^t \right) \in \mathbb{G}^{A+1}$$

Output the pair (Hdr, K). Note that Hdr contains $A+1$ elements.

Decrypt$(S, i, d_i, \text{Hdr}, PK)$**:** Let $\text{Hdr} = (C_0, C_1, \ldots, C_A)$ and recall that $d_i \in \mathbb{G}$. Write i as $i = (a-1)B + b$ for some $1 \le a \le A$ and $1 \le b \le B$. Then

$$K = e(g_b, C_a) \ / \ e(d_i \cdot \prod_{\substack{j \in S_a \\ j \neq b}} g_{B+1-j+b},\ C_0)$$

Verifying that the decryption algorithm works correctly is analogous to the calculation in the previous section. We note that when $B = n$ then $A = 1$ and we obtain the system of the previous section.

Efficiency A user's private key size again will only consist of one group element. The ciphertext consists of $A+1$ group elements and the public key is $2B+A$ elements. Our choice of B depends on the application. As we will see, in some cases we want $B=n$ to obtain the smallest possible ciphertext. In other cases we want $B=\sqrt{n}$ to minimize the concatenation of the ciphertext and public key.

The decryption time for user $i=(a-1)B+b$ will be dominated by $|S_a|-2<B$ group operations. Similar caching techniques to those described in the end of Section 3.1 can be used to improve performance.

3.3 Security

We now prove the semantic security of the general system of Section 3.2.

Theorem 1. *Let $\mathbb{G}$ be a bilinear group of prime order p. For any positive integers B,n $(n>B)$ our B-broadcast encryption system is (t,ϵ,n) semantically secure assuming the decision (t,ϵ,B)-BDHE assumption holds in $\mathbb{G}$.*

Proof. Suppose there exists a t-time adversary, $\mathcal{A}$, such that $\mathsf{AdvBr}_{\mathcal{A},n}>\epsilon$ for a system parameterized with a given B. We build an algorithm, $\mathcal{B}$, that has advantage ϵ in solving the decision B-BDHE problem in $\mathbb{G}$. Algorithm $\mathcal{B}$ takes as input a random decision B-BDHE challenge $(g,h,\boldsymbol{y}_{g,\alpha,B},Z)$, where $\boldsymbol{y}_{g,\alpha,B}=(g_1,\ldots,g_B,g_{B+2},\ldots,g_{2B})$ and Z is either $e(g_{B+1},h)$ or a random element of $\mathbb{G}_1$ (recall that $g_i=g^{(\alpha^i)}$ for all i). Algorithm $\mathcal{B}$ proceeds as follows.

Init Algorithm $\mathcal{B}$ runs $\mathcal{A}$ and receives the set S of users that $\mathcal{A}$ wishes to be challenged on.

Setup $\mathcal{B}$ needs to generate a public key PK and private keys d_i for $i\notin S$. The crux of the proof is in the choice of $v_1,\ldots,v_A$. Algorithm $\mathcal{B}$ chooses random $u_i\in\mathbb{Z}_p$ for $1\le i\le A$. We again define the subsets $\hat{S}_i$ and S_i as

$$\hat{S}_i=S\cap\{(i-1)B+1,\ldots,iB\}\quad\text{and}\quad S_i=\{x-iB+B \mid x\in\hat{S}_i\}\subseteq\{1,\ldots,B\}$$

For $i=1,\ldots,A$ algorithm $\mathcal{B}$ sets $v_i=g^{u_i}\left(\prod_{j\in S_i}g_{B+1-j}\right)^{-1}$. It gives $\mathcal{A}$ the public key

$$PK=(g,g_1,\ldots,g_B,g_{B+2},\ldots,g_{2B},v_1,\ldots,v_A)\in\mathbb{G}^{2B+A}$$

Note that since g,α and the u_i values are chosen uniformly at random, this public key has an identical distribution to that in the actual construction.

Next, the adversary needs all private keys that are not in the target set S. For all $i\notin S$ we write i as $i=(a-1)B+b$ for some $1\le a\le A$ and $1\le b\le B$. Algorithm $\mathcal{B}$ computes d_i as

$$d_i=g_b^{u_i}\cdot\prod_{j\in S_a}(g_{B+1-j+b})^{-1}$$

Indeed, we have that $d_i = (g^{u_i} \prod_{j \in S_a} (g_{B+1-j})^{-1})^{(\alpha^b)} = v_a^{(\alpha^b)}$ as required. The main point is that since $i \notin S$ we know that $b \notin S_a$ and hence the product defining d_i does not include the term g_{B+1}. It follows that algorithm $\mathcal{B}$ has all the necessary values to compute d_i.

Challenge To generate the challenge, $\mathcal{B}$ computes Hdr as $(h, h^{u_1}, \ldots, h^{u_A})$. It then randomly chooses a bit $b \in \{0,1\}$ and sets $K_b = Z$ and picks a random K_{1-b} in $\mathbb{G}_1$. It gives (Hdr, K_0, K_1) as the challenge to $\mathcal{A}$.

We claim that when $Z = e(g_{B+1}, h)$ (i.e. the input to $\mathcal{B}$ is a B-BDHE tuple sampled from $\mathcal{P}_{BDHE}$) then (Hdr, K_0, K_1) is a valid challenge to $\mathcal{A}$ as in a real attack. To see this, write $h = g^t$ for some (unknown) $t \in \mathbb{Z}_p$. Then, for all $i = 1, \ldots, A$ we have

$$h^{u_i} = (g^{u_i})^t = (g^{u_i} (\prod_{j \in S_i} g_{B+1-j})^{-1} (\prod_{j \in S_i} g_{B+1-j}))^t = (v_i \prod_{j \in S_i} g_{B+1-j})^t$$

Therefore, by definition, $(h, h^{u_1}, \ldots, h^{u_A})$ is a valid encryption of key $e(g_{B+1}, g)^t$. Furthermore, $e(g_{B+1}, g)^t = e(g_{B+1}, h) = Z = K_b$ and hence (Hdr, K_0, K_1) is a valid challenge to $\mathcal{A}$.

On the other hand, when Z is random in $\mathbb{G}_1$ (i.e. the input to $\mathcal{B}$ is sampled from $\mathcal{R}_{BDHE}$) then K_0, K_1 are just random independent elements of $\mathbb{G}_1$.

Guess The adversary, $\mathcal{A}$ outputs a guess b' of b. If $b' = b$ the algorithm $\mathcal{B}$ outputs 0 (indicating that $Z = e(g_{B+1}, h)$). Otherwise, it outputs 1 (indicating that Z is random in $\mathbb{G}_1$).

We see that $\Pr[\mathcal{B}(g, h, \boldsymbol{y}_{g,\alpha,B}, Z) = 0] = \frac{1}{2}$ if $(g, h, \boldsymbol{y}_{g,\alpha,B}, Z)$ is sampled from $\mathcal{R}_{BDHE}$. If $(g, h, \boldsymbol{y}_{g,\alpha,B}, Z)$ is sampled from $\mathcal{P}_{BDHE}$ then $|\Pr[\mathcal{B}(g, h, \boldsymbol{y}_{g,\alpha,B}, Z) = 0] - \frac{1}{2}| = \mathsf{AdvBr}_{\mathcal{A},n} \geq \epsilon$. It follows that $\mathcal{B}$ has advantage at least ϵ in solving decision B-BDHE in $\mathbb{G}$. This concludes the proof of Theorem 1. $\square$

Note that the proof of Theorem 1 does not use the random oracle model. The system can be proved secure using the weaker *computational* B-BDHE assumption (as opposed to decision B-BDHE), using the random oracle model. In that case the advantage of the simulator is at least ϵ/q, where q is the maximum number of random oracle queries made by the adversary.

4 Applications

We describe how our system can be used for a number of specific applications. The first application, file sharing in encrypted file systems, is an example of broadcasts to small sets. The second application, encrypted email for large mailing lists, shows that the majority of the public key can be shared by many broadcast systems so that the public key for a new broadcast system is constant size. The third application, DVD content protection, is an example of broadcasts to large sets.

4.1 File Sharing in Encrypted File Systems

Encrypted file systems let users encrypt files on disk. For example, Windows EFS encrypts the file contents using a file encryption key K_F and then places an encryption of K_F in the file header. If n users have access to the file, EFS encrypts K_F under the public keys of all n users and places the resulting n ciphertexts in the file header. Related designs can be found in the SiRiUS [GSMB03] and Plutus [KRS+03] file systems.

Abstractly, access control in an encrypted file system can be viewed as a broadcast encryption problem. The file system is the broadcast channel and the key K_F is broadcast (via the file header) to the subset of users that can access file F. Many encrypted file systems implement the straightforward broadcast system where the number of ciphertexts in the file header grows linearly in the number of users that can access the file. As a result, there is often a hard limit on the number of users that can access a file. For example, the following quote is from Microsoft's knowledge base:

> "EFS has a limit of 256KB in the file header for the EFS metadata. This limits the number of individual entries for file sharing that may be added. On average, a maximum of 800 individual users may be added to an encrypted file."

A natural question is whether we can implement file sharing in an encrypted file system without resorting to large headers. Remarkably, there is no known combinatorial solution that performs better than the straightforward solution used in EFS. The broadcast system of Section 3.1 performs far better and provides a system with the following parameters:

- The public key (whose size is linear in n) is stored on the file system. Even for a large organization of 100,000 users this file is only 4MB long (using a standard security parameter where each group element is 20 bytes).
- Each user is given a private key that contains only one group element.
- Each file header contains $([S], C)$ where $[S]$ is a description of the set S of users who can access F and C is a fixed size ciphertext consisting of only two group elements.

Since S tends to be small relative to n, its shortest description is simply an enumeration of the users in S. Assuming 32-bit user ID's, a description of a set S of size r takes $4r$ bytes. Hence, the file header grows with the size of S, but only at a rate of 4 bytes per user. In EFS the header grows by one public key ciphertext per user. For comparison, we can accommodate sharing among 800 users using a header of size $4 \times 800 + 40 = 3240$ bytes which is far less than EFS's header size. Even if EFS were using short ElGamal ciphertexts on elliptic curves, headers would grow by 44 bytes per user which would result in headers that are 11 times bigger than headers in our system.

Our system has a few more properties that make it especially useful for cryptographic access control. We describe these next.

1. Incremental sharing. Suppose a file header contains a ciphertext $C = (C_0, C_1)$ which is the encryption of K_F for a certain set S of users. Let $C_0 = g^t$

and $C_1 = (v \cdot \prod_{j \in S} g_{n+1-j})^t$. Suppose the file owner wishes to add access rights for some user $u \in \{1, \ldots, n\}$. This is easy to do given t. Simply set $C_1 \leftarrow C_1 \cdot g_{n+1-u}^t$. Similarly, to revoke access rights for user u set $C_1 \leftarrow C_1 / g_{n+1-u}^t$.

This incremental sharing mechanism requires the file owner to remember the random value $t \in \mathbb{Z}_p$ for every file. Alternatively, the file owner can embed a short nonce T_F in every file header and derive the value t for that file by setting $t \leftarrow PRF_k(T_F)$ where k is a secret key known only to the file owner. Hence, changing access permissions can be done efficiently with the file owner only having to remember a single key k. Note that when access rights to a file F change it is sometimes desirable to re-encrypt the file using a new key K_F^{new}. Modifying the existing header to encrypt a new K_F^{new} for the updated access list is just as easy.

2. Incremental growth of the number of users. In many cases a broadcast encryption system must be able to handle the incremental addition of new users. It is desirable to have a system that does not a-priori restrict the total number of users it can handle. Our system supports this by slowly expanding the public key as the number of users in the system grows. To do so, at system initialization the setup algorithm picks a large value of n (say $n = 2^{64}$) that is much larger than the maximum number of users that will ever use the system. At any one time if there are j users in the system the public key will be $g_{n-j+1}, \ldots, g_n, g_{n+2}, \ldots, g_{n+j}$. Whenever a new user joins the system we simply add two more elements to the public key. Note that user i must also be given g_i as part of the private key and everything else remains the same.

4.2 Sending Encrypted Email to Mailing Lists

One interesting aspect of our broadcast encryption system is that the public values $\boldsymbol{y}_{g,\alpha,n} = (g_1, \ldots g_n, g_{n+2}, \ldots, g_{2n})$ can be shared among many broadcast systems and α can be erased. Suppose this $\boldsymbol{y}_{g,\alpha,n}$ is distributed globally to a large group of users (for example, imagine $\boldsymbol{y}_{g,\alpha,n}$ is pre-installed on every computer). Then creating a new broadcast system is done by simply choosing a random $\gamma \in \mathbb{Z}_p$, setting $v = g^\gamma$, and assigning private keys as $d_i = g_i^\gamma$. Since all broadcast systems use the same pre-distributed $\boldsymbol{y}_{g,\alpha,n}$, the actual public key for this new broadcast system is just one element, v. Theorem 1 shows that using the same $\boldsymbol{y}_{g,\alpha,n}$ for many broadcast systems is secure.

We illustrate this property with an example of secure mailing lists. Sending out encrypted email to all members of a mailing list is an example of a broadcast encryption system. Suppose the global public vector $\boldsymbol{y}_{g,\alpha,n}$ is shipped with the operating system and installed on every computer. We arbitrarily set $n = 50,000$ in which case the size of $\boldsymbol{y}_{g,\alpha,n}$ is about 2MB.

For every secure mailing list, the administrator creates a separate broadcast encryption system. Thus, every mailing list is identified by its public key $v = g^\gamma$. We assume the maximum number of members in a mailing list is less than $n = 50,000$ (larger lists can be partitioned into multiple smaller lists). Each time a new user is added onto the list, the user is assigned a previously unused index $i \in \{1, \ldots, n\}$ and given the secret key $d_i = g_i^\gamma$. The broadcast set S is updated to include i and all mailing list members are notified of the change in S. Similarly, if a user with index j is removed from the list, then j is removed from the set S and all members are notified. We note that any member, i, does not need to actually store S. Instead, member i need only store the value $\prod_{j \in S j \neq i} g_{n+1-j+i}$ needed for decryption. The member updates this value every time a membership update message is sent. To send email to a mailing list with public key v the server simply does a broadcast encryption to the current set of members S using $(\boldsymbol{y}_{g,\alpha,n}, v)$ as the public key.

In this mail system, the header of an email message sent to the list is constant size. Similarly, membership update messages are constant size. A mailing list member only needs to store two group elements for each list he belongs to (although we have to keep in mind the cost of storing $\boldsymbol{y}_{g,\alpha,n}$ which is amortized over all mailing lists). It is interesting to compare our solution to one using an LKH scheme [WHA97,WGL98]. In LKH email messages are encrypted under a group key. Using this type of a system each update message contains $O(\log(m))$ ciphertexts, and private keys are of size $O(\log(m))$ (per system), where m is the current group size. In our system, update messages and private user storage are much smaller.

4.3 Content Protection

Broadcast encryption applies naturally to protecting DVD content, where the goal is to revoke compromised DVD players. Recall that the public key in our system is needed for decryption and hence it must be embedded in the header of every DVD disk. Consequently, we are interested in minimizing the total length of the header and public key, namely minimize $|\mathrm{Hdr}| + |PK|$.

Let n be the total number of DVD players (e.g. $n = 2^{32}$) and let r be the number of revoked players. Let $\ell_{\mathrm{id}} = \log_2 n$ (e.g. $\ell_{\mathrm{id}} = 32$) and let $\bar{k}$ be the size of a group element (e.g. $\bar{k} = 160$ bits). Then using our $\sqrt{n}$-broadcast system ($B = \sqrt{n}$) we can broadcast to sets of size $n - r$ using the following parameters:

$$\text{priv-key-size} = 4\bar{k}, \quad \text{and} \quad |\mathrm{Hdr}| + |PK| + |S| = 4\bar{k}\lceil\sqrt{n}\rceil + r\ell_{\mathrm{id}}$$

In comparison, the NNL system [NNL01] and its derivatives [HS02,GST04] can broadcast to sets of size $n - r$ using:

$$\text{priv-key-size} = O(k \log n), \quad \text{and} \quad \text{header-size} = O((k + \ell_{\mathrm{id}}) \cdot r)$$

where k is the length of a symmetric key (e.g. $k = 128$ bits). Note that the broadcast header grows by $O(k + \ell_{\mathrm{id}})$ bits per revoked player. With our system the broadcast header only grows by ℓ_{id} bits per revoked player.

Example. Plugging in real numbers we obtain the following. When $n = 2^{32}, \bar{k} =$ 20 bytes, and $\ell_{\mathrm{id}} = 4$ bytes, header size in our system is 5.12MB and each revocation adds 4 bytes to the header. In NNL-like systems, using $k = 128$-bit symmetric keys, each revocation adds about 40 bytes to the header, but there is no upfront 5MB fixed cost.

The best system is obtained by combining NNL with our system (using NNL when $r < \sqrt{n}$ and our system when $r > \sqrt{n}$). Thus, as long as things are stable, DVD disk distributors use NNL. In case of a disaster where, say, a DVD player manufacturer loses a large number of player keys, DVD disk distributors can switch to our system where the header size grows slowly beyond $O(\sqrt{n})$.

5 Chosen Ciphertext Secure Broadcast Encryption

We show how to extend the system of Section 3.1 to obtain chosen ciphertext security. The basic idea is to compose the system with the IBE system of [BB04] and then apply the ideas of [CHK04]. The resulting system is chosen ciphertext secure without using random oracles.

We need a signature scheme (*SigKeyGen*, *Sign*, *Verify*). We also need a collision resistant hash function that maps verification keys to $\mathbb{Z}_p$. Alternatively, we can simply assume (as we do below) that signature verification keys are encoded as elements of $\mathbb{Z}_p$. This greatly simplifies the notation.

As we will see, security of the CCA-secure broadcast system for n users is based on the $(n+1)$-BDHE assumption (as opposed to the n-BDHE assumption for the system of Section 3.1). Hence, to keep the notation consistent with Section 3.1 we will describe the CCA-secure system for $n-1$ users so that security will depend on the n-BDHE assumption as before. The system works as follows:

Setup$(n-1)$**:** Public key PK is generated as in Section 3.1. The private key for user $i \in \{1, \ldots, n-1\}$ is set as: $d_i = g_i^\gamma \in \mathbb{G}$. The algorithm outputs the public key PK and the $n-1$ private keys $d_1, \ldots, d_{n-1}$.

Encrypt(S, PK)**:** Run the *SigKeyGen* algorithm to obtain a signature signing key K_{SIG} and a verification key V_{SIG}. Recall that for simplicity we assume $V_{\text{SIG}} \in \mathbb{Z}_p$. Next, pick a random t in $\mathbb{Z}_p$ and set $K = e(g_{n+1}, g)^t \in \mathbb{G}_1$. Set

$$C = \left(g^t,\ \big(v \cdot g_1^{V_{\text{SIG}}} \cdot \prod_{j \in S} g_{n+1-j}\big)^t\right) \in \mathbb{G}^2, \quad \text{Hdr} = \big(C, \mathit{Sign}(C, K_{\text{SIG}}), V_{\text{SIG}}\big)$$

and output the pair (Hdr, K). Note that the only change to the ciphertext from Section 3.1 is the term $g_1^{V_{\text{SIG}}}$ and the additional signature data.

Decrypt$(S, i, d_i, \text{Hdr}, PK)$**:** Let $\text{Hdr} = \big((C_0, C_1),\ \sigma,\ V_{\text{SIG}}\big)$.

1. Verify that σ is a valid signature of (C_0, C_1) under the key V_{SIG}. If invalid, output '?'.
2. Otherwise, pick a random $w \in \mathbb{Z}_p$ and compute

$$\hat{d}_0 = \Big(d_i \cdot g_{i+1}^{V_{\text{SIG}}} \cdot \prod_{\substack{j \in S \\ j \neq i}} g_{n+1-j+i}\Big) \cdot \Big(v \cdot g_1^{V_{\text{SIG}}} \cdot \prod_{j \in S} g_{n+1-j}\Big)^w \quad \text{and} \quad \hat{d}_1 = g_i g^w$$

3. Output $\quad K = e(\hat{d}_1, C_1)/e(\hat{d}_0, C_0)$.

Correctness can be shown with a similar calculation to the one in Section 3.1. Note that private key size and ciphertext size are unchanged.

Unlike the system of Section 3.1, decryption requires a randomization value $w \in \mathbb{Z}_p$. This randomization ensures that for any $i \in S$ the pair $(\hat{d}_0, \hat{d}_1)$ is chosen from the following distribution

$$\left(g_{n+1}^{-1} \cdot \left(v \cdot g_1^{V_{\mathrm{SIG}}} \cdot \prod_{j \in S} g_{n+1-j} \right)^r , \quad g^r \right)$$

where r is uniform in $\mathbb{Z}_p$. Note that this distribution is independent of i implying that all members of S generate a sample from the same distribution. Although this randomization slows down decryption by a factor of two, it is necessary for the proof of security.

We briefly recall that a signature scheme (*SigKeyGen*, *Sign*, *Verify*) is (t, ϵ, q_S) strongly existentially unforgeable if no t-time adversary who makes at most q_S signature queries is able to produce some new (message,signature) pair with probability at least ϵ. A complete definition is given in, e.g., [CHK04]. The following theorem proves chosen ciphertext security.

Theorem 2. *Let $\mathbb{G}$ be a bilinear group of prime order p. For any positive integer n, the broadcast encryption system above is $(t, \epsilon_1 + \epsilon_2, n-1, q_D)$ CCA-secure assuming the decision (t, ϵ_1, n)-BDHE assumption holds in $\mathbb{G}$ and the signature scheme is $(t, \epsilon_2, 1)$ strongly existentially unforgeable.*

We give the proof of Theorem 2 in the full version of the paper [BGW05]. The proof does not use the random oracle model implying that the system is chosen-ciphertext secure in the standard model.

We also note that instead of the signature-based method of [CHK04] we could have used the more efficient MAC-based method of [BK05]. We chose to present the construction using the signature method to simplify the proof. The MAC-based method would also work.

6 Conclusions and Open Problems

We presented a fully collusion resistant broadcast encryption scheme with constant size ciphertexts and private keys for arbitrary receiver sets. In Section 5 we built a chosen-ciphertext secure broadcast system with the same parameters. A generalization of our basic scheme gave us a tradeoff between public key size and ciphertext size. With the appropriate parametrization we achieve a broadcast encryption scheme with $O(\sqrt{n})$ ciphertext and public key size. We discussed several applications such as encrypted file systems and content protection.

We leave as an open problem the question of building a public-key broadcast encryption system with the same parameters as ours which is secure against adaptive adversaries. We note that any non-adaptive scheme that is (t, ϵ, n) secure is also $(t, \epsilon/2^n, n)$ secure against adaptive adversaries. However, in practice this reduction is only meaningful for small values of n.

Another problem is to build a tracing traitors system [CFN94] with the same parameters as our system. Ideally, one could combine the two systems to obtain an efficient trace-and-revoke system. Finally, it is interesting to explore alternative systems with similar performance that can be proved secure under a weaker assumption.

References

[AMM99] J. Anzai, N. Matsuzaki, and T. Matsumoto. A quick key distribution scheme with entity revocation. In *Proc. of Asiacrypt '99*, pages 333–347. Springer-Verlag, 1999.

[BB04] D. Boneh and X. Boyen. Efficient selective-ID identity based encryption without random oracles. In *Proc. of Eurocrypt 2004*, volume 3027 of *LNCS*, pages 223–238. Springer-Verlag, 2004.

[BBG05] D. Boneh, X. Boyen, and E. Goh. Hierarchical identity based encryption with constant size ciphertext. In *Proc. of Eurocrypt 2005*. Springer-Verlag, 2005.

[BF01] D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In *Proceedings of Crypto 2001*, volume 2139 of *LNCS*, pages 213–29. Springer-Verlag, 2001.

[BGW05] D. Boneh, C. Gentry, and B. Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. Cryptology ePrint Archive, Report 2005/018, 2005. Full version of current paper.

[BK05] D. Boneh and J. Katz. Improved efficiency for CCA-secure cryptosystems built using identity based encryption. In *Proc. of RSA-CT 2005*, volume 3376 of *LNCS*, pages 87–103. Springer-Verlag, 2005.

[BS03] D. Boneh and A. Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324:71–90, 2003.

[CFN94] B. Chor, A. Fiat, and M. Naor. Tracing traitors. In *Proc. of Crypto 1994*, volume 839 of *LNCS*, pages 257–270. Springer-Verlag, 1994.

[CGI+99] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. In *Proc. IEEE INFOCOM 1999*, volume 2, pages 708–716. IEEE, 1999.

[CHK04] R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. In *Proc. of Eurocrypt 2004*, volume 3027 of *LNCS*, pages 207–222. Springer-Verlag, 2004.

[CMN99] R. Canetti, T. Malkin, and K. Nissim. Efficient communication-storage tradeoffs for multicast encryption. In *Proc. of Eurocrypt 1999*, pages 459–474. Springer-Verlag, 1999.

[DF02] Y. Dodis and N. Fazio. Public key broadcast encryption for stateless receivers. In *Proc. of DRM 2002*, volume 2696 of *LNCS*, pages 61–80. Springer-Verlag, 2002.

[DF03] Y. Dodis and N. Fazio. Public key broadcast encryption secure against adaptive chosen ciphertext attack. In *Proc. of PKC 2003*, pages 100–115. Springer-Verlag, 2003.

[DY05] Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In *Proc. of PKC 2005*, pages 416–431. Springer-Verlag, 2005.

[FN93] A. Fiat and M. Naor. Broadcast encryption. In *Proc. of Crypto 1993*, volume 773 of *LNCS*, pages 480–491. Springer-Verlag, 1993.

[GSMB03] E. Goh, H. Shacham, N. Modadugu, and D. Boneh. Sirius: Securing remote untrusted storage. In *Proc. of NDSS 2003*, pages 131–145, 2003.

[GST04] M.T. Goodrich, J.Z. Sun, and R. Tamassia. Efficient tree-based revocation in groups of low-state devices. In *Proc. of Crypto 2004*, volume 3152 of *LNCS*, pages 511–527. Springer-Verlag, 2004.

[GSW00] J. Garay, J. Staddon, and A. Wool. Long-lived broadcast encryption. In *Proc. of Crypto 2000*, volume 1880 of *LNCS*, pages 333–352. Springer-Verlag, 2000.

[GSY99] E. Gafni, J. Staddon, and Y.L. Yin. Efficient methods for integrating traceability and broadcast encryption. In *Proc. of Crypto 1999*, volume 1666 of *LNCS*, pages 372–387. Springer-Verlag, 1999.

[HS02] D. Halevy and A. Shamir. The lsd broadcast encryption scheme. In *Proc. of Crypto 2002*, volume 2442 of *LNCS*, pages 47–60. Springer-Verlag, 2002.

[JN03] A. Joux and K. Nguyen. Separating decision Diffie-Hellman from Diffie-Hellman in cryptographic groups. *J. of Cryptology*, 16(4):239–247, 2003. Early version in Cryptology ePrint Archive, Report 2001/003.

[Jou00] A. Joux. A one round protocol for tripartite Diffie-Hellman. In *Proc. of ANTS IV*, volume 1838 of *LNCS*, pages 385–94. Springer-Verlag, 2000.

[KRS+03] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, and K. Fu. Plutus: Scalable secure file sharing on untrusted storage. In *Proc. of USENIX Conf. on File and Storage Technologies (FAST)*, 2003.

[NNL01] D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In *Proc. of Crypto 2001*, volume 2139 of *LNCS*, pages 41–62. Springer-Verlag, 2001.

[NP00] M. Naor and B. Pinkas. Efficient trace and revoke schemes. In *Proc. of Financial cryptography 2000*, volume 1962 of *LNCS*, pages 1–20. Springer-Verlag, 2000.

[Sho97] V. Shoup. Lower bounds for discrete logarithms and related problems. In *Proc. of Eurocrypt 1997*, volume 1233 of *LNCS*, pages 256–266. Springer-Verlag, 1997.

[SM03] A.T. Sherman and D.A. McGrew. Key establishment in large dynamic groups using one-way function trees. *IEEE Trans. Softw. Eng.*, 29(5):444–458, 2003.

[ST98] D.R. Stinson and T.V. Trung. Some new results on key distribution patterns and broadcast encryption. *Des. Codes Cryptography*, 14(3):261–279, 1998.

[Sti97] D.R. Stinson. On some methods for unconditionally secure key distribution and broadcast encryption. *Des. Codes Cryptography*, 12(3):215–243, 1997.

[SW98] D.R. Stinson and R. Wei. Combinatorial properties and constructions of traceability schemes and frameproof codes. *SIAM J. Discret. Math.*, 11(1):41–53, 1998.

[TT01] W. Tzeng and Z. Tzeng. A public-key traitor tracing scheme with revocation using dynamic shares. In *Proc. of PKC 2001*, pages 207–224. Springer-Verlag, 2001.

[WGL98] C.K. Wong, M. Gouda, and S. Lam. Secure group communications using key graphs. In *Proc. of SIGCOMM 1998*, 1998.

[WHA97] D.M. Wallner, E.J. Harder, and R.C. Agee. Key management for multicast: Issues and architectures. IETF draft wallner-key, 1997.

[YJCK04] E. Yoo, N. Jho, J. Cheon, and M. Kim. Efficient broadcast encryption using multiple interpolation methods. In *Proc. of ICISC 2004*, LNCS. Springer-Verlag, 2004.