A Broadcast Encryption Scheme with Free-Riders but Unconditional Security

Andre Adelsbach and Ulrich Greveler*

Horst Görtz Institute for IT Security
Ruhr University Bochum
Germany
(andre.adelsbach,ulrich.greveler)@nds.rub.de

This is an extended version of a conference paper presented at
First International Conference on Digital Rights Management:
Technology, Issues, Challenges and Systems
Sydney, Australia, 31 October - 2 November 2005

Abstract. We propose two schemes for efficient broadcast key establishment that enables a sender to communicate to any subset of the user-base by allowing a small ratio of *free-riders*. The schemes do not require stateful receivers and one scheme is unconditionally secure. The free-riders are unable to learn from the past whether they might become free-riders for a certain transmission again.
We present a new trade-off facet for broadcast encryption, namely the number (or ratio) of free-riders vs. the number of messages to be sent or the number of keys stored by each user.

1 Introduction

A number of applications need solutions to the problem of transmitting data to a group of receivers in a way that only the correct subset of all possible receivers can decrypt the data: Pay-TV, Digital Rights Management (DRM) controlled media, audio streaming, real-time business data, multicast communication are current examples. The subset of receivers can change for every transmission (e.g. pay-per-view) so an efficient scheme for a quick establishment of a secure channel to the new subset is desirable.

In the literature one can find very efficient revocation schemes which are suitable for a small set $\mathcal{R}$ with $|\mathcal{R}| \ll |\mathcal{N}|$ of revoked receivers (e.g. pirate receivers or traitors) compared to a huge number of total users $\mathcal{N}$ so that the broadcast communication can only be decrypted by the users $\mathcal{N} - \mathcal{R}$. The most efficient

* This author was supported by the European Commission through IST-2002-507932 ECRYPT.

known schemes [1–4] require a message header of length $O(|\mathcal{R}|)$ and user's individual private key size of $O(\log(|\mathcal{N}|))$. However, these revocation schemes are not intended nor suitable for a general subset case, e.g., cases where $|R| \approx \frac{1}{2}|\mathcal{N}|$.

The trivial scheme to address exactly all the users in the target set $\mathcal{N} - \mathcal{R}$ is to send the message encrypted individually for each user yielding a total number of $|\mathcal{N} - \mathcal{R}|$ messages to be sent via broadcast and only $O(1)$ keys to be stored by a user. This straightforward scheme is still the best known approach for the case where $|\mathcal{N} - \mathcal{R}| \ll |\mathcal{N}|$.

In this paper we will consider the case where *arbitrary* subsets are addressed by the sender. Assuming that any subset of $\mathcal{N}$ is chosen with equal probability for a transmission an average number of $\frac{1}{2}|\mathcal{N}|$ messages needs to be sent via the broadcast channel for every transmission if the trivial scheme is used. In order to reduce this number it is possible to assign keys to certain or all subsets of $\mathcal{N}$ and make these keys known only to the members of the subset. But even in the best (and not realistic) case where each user is provided with a key for all $2^{|\mathcal{N}|-1}$ subsets it belongs to, the numbers of bits needed to encode the subset key identifier is approximately $|\mathcal{N}|$ so any scheme which addresses the exact subsets would need to send $O(|\mathcal{N}|)$ message bits. Apart from that lower bound, a trade-off between the number of keys stored by each user and the number of messages to be sent to establish a transmission session key needs to be considered. The number of colluders (users outside the target group cooperating to break the scheme) the system can tolerate is another major parameter. Finally, we are interested in the level of security (existence of one-way functions, number-theoretic or information-theoretic security) we can establish.

1.1 Relaxed Requirements, New Trade-Offs

In order to set up schemes that are more efficient than sending $|\mathcal{N}| - |\mathcal{R}|$ messages we are relaxing the requirement that only the users in the target group $\mathcal{T} := \mathcal{N} - \mathcal{R}$ can decrypt the message by allowing a certain (small) number of users in $\mathcal{R}$ to decrypt the transmission as long as every user in $\mathcal{T}$ can receive the transmission. In this case new requirements on a relaxed scheme are to be considered: The number of users who can receive a transmission they have not subscribed to, i.e., the number of *free-riders*, shall be minimized and—following economic, game-theoretic requirements (see e.g. [5])—a user shall not gain any information whether she might be a free-rider for a future transmission by examining the past transmissions.

For example in a pay-TV scenario, we want to avoid a situation where two users $u_1, u_2 \in \mathcal{N}$ are put in one subscription set so that each time user u_1 subscribes to a transmission the user u_2 becomes a free-rider. The user u_2 might learn that he often becomes a free-rider for a certain kind of transmission preferred by u_1 (e.g., *Tarantino* movies) and will stop subscribing for these transmissions to avoid *unnecessary* payment.

The main area of trade-off parameters to be considered in this relaxed notion of broadcast encryption is the number (or ratio) of free-riders versus the message header length versus the user key size. Other major requirements on a scheme

are collusion resiliency (i.e., the number of non-subscribers that may collude without being able to access the secured transmission) and underlying security assumptions (e.g., unconditional security versus computational security).

1.2 Related Work

The notion of broadcast encryption was introduced by Fiat and Naor in [6]. Their work described several methods making it possible to remove users from the target group by setting the requirement that only t users may collude where $t \leq k$ (k-*resiliency*). One method achieves a message header of size $O(k^3 \log |\mathcal{N}|)$ and a user key storage of $O(k|\mathcal{N}| \log |\mathcal{N}|)$ with unconditional security. The method is improved to user key storage of $O(k \log^2 |\mathcal{N}|)$ by assuming the existence of one-way-functions and to user key storage $O(k \log |\mathcal{N}|)$ by assuming hardness of root extraction modulo a composite.

Naor, Naor and Lotspiech [1] presented their *complete subtree* method that is secure under any number of colluders ($|\mathcal{R}|$-resiliency) and requires a header length of $|\mathcal{R}| \log(|\mathcal{N}|/|\mathcal{R}|)$ and $O(\log |\mathcal{N}|)$ keys per user. An improved version, the *subtree difference* method, requires header length $2|\mathcal{R}| - 1$ and $O(\log^2 |\mathcal{N}|)$ keys per user. Both methods are very efficient in the $|\mathcal{R}| \ll |\mathcal{N}|$ case and use PRNGs to assign keys in a tree structure.

Halevy and Shamir [2] presented a modified subset difference method with $O(\log^{1+\epsilon}(|\mathcal{N}|)$ key storage and $O(|\mathcal{R}|/\epsilon)$ message header size where ϵ can be chosen ($\epsilon = 1/2$ is a natural choice).

Dodis and Fazio [7] extended all three schemes, i.e., CS, SD and LSD, to the public key setting.

Boneh and Silverberg [8] showed that by using n-linear maps a collusion secure scheme with a fixed size public key and message header length can be achieved; Boneh and Waters [9] improved this by limiting a modified scheme to bilinear maps. Both schemes do not provide information-theoretic security.

Luby and Staddon [10] considered the information theoretic case and give general lower bounds for revocation schemes. Applying these bounds to the general case (i. e., not assuming $|\mathcal{R}| \ll |\mathcal{N}|$) shows that broadcast schemes with unconditional security are never efficient in the sense that either the message header length is large or the user key size is large.

1.3 Summary of Results and Outline

In the following section we will propose two new broadcast encryption schemes operating in a pseudo-probabilistic way. Both schemes realize their efficiency by accepting an adjustable ratio of free-riders. The first scheme is unconditionally secure, but puts certain undesirable constraints on the abilities of attackers; the improved scheme is information-theoretically secure and lacks these constraints. We will give a calculation of the parameter trade-offs of our schemes and discuss the collusion resiliency.

2 The Biased Sub-set Scheme[1]

2.1 Notations, Definitions and Basic Idea

Let $\mathcal{N}$ be the set of all users of a broadcast scheme and $\mathcal{T}$ be the set of users which shall receive a certain transmission[2].

Each user $u \in \mathcal{N}$ is provided with a fixed set of secret keys K_u which are assigned to him before receiving any transmission. Each user owns at least one individual key $k_u^{indv} \in K_u$ only known to him and the sender; the other keys might also be shared between several users, which is not known to the users sharing a key. During a transmission any user might receive further one-time usage keys (session keys, key encryption keys) which are not re-used and do not need to be stored after the transmission (thus we have a *stateless receiver*).

The basic idea of our scheme is to transmit the session key for a certain transmission bit-wise in a probabilistic way to all users in $\mathcal{N}$ so that the users in $\mathcal{T}$ receive on average more key bits than the users in $\mathcal{N} - \mathcal{T}$, thus only a small fraction of the users in $\mathcal{N} - \mathcal{T}$ is able to decrypt the transmission. Most users in $\mathcal{T}$ are provided with enough bits of the session key to derive the full key after exhaustive search. For the great majority of the users in $\mathcal{N} - \mathcal{T}$ it is infeasible to derive the session key in due time (e.g., before the transmission starts or before the transmitted data becomes outdated).

We choose a security parameter s and the generated session key k_S consists of $|k_S| = s$ bits. This key is valid for one transmission only. For the users in $\mathcal{T}$ a minimum of $d < s$ bits is needed to derive k_S (d is chosen according to the computation power of the users). We assume potential attackers could be more powerful than the ordinary users, so they only need $d' \leq d < s$ bits to derive k_S in due time. The goal of the scheme is then that at a protocol step, the great majority of users in $\mathcal{T}$ has received more than d bits when at the same time only a small minority of users in $\mathcal{N} - \mathcal{T}$ has received d' or more bits (see Figure 1).

Our scheme works in two phases: First a number of messages each carrying one key bit of k_S is broadcasted (each message can only be decrypted by a different subset of $\mathcal{N}$ provided with the right subset key) so that a certain number of the users $\mathcal{T}' \subset \mathcal{T}$ has received at least d bits (e.g., targeting $\frac{|T'|}{|T|} > 0.95$). In the second phase each user in $\mathcal{T} - \mathcal{T}'$ is provided individually with the full session key using the keys k_u^{indv} for all $u \in \mathcal{T} - \mathcal{T}'$.

Remark 1. Our approach broadcasts a secret by gradually broadcasting parts (bits or shares) of the overall secret to certain subsets so that any party having enough bits (or shares) can compute or recover the complete secret, e.g., the session key of a pay-TV broadcast transmission. The gradual transmission of secrets has been previously applied in the context of fair exchange [11]. In this

[1] This work is subject of German patent DE 1020 0404 2094 B3 (issued 2005).

[2] We will not use the notion of a set $\mathcal{R}$ of revoked users in this paper further as we address the problem of broadcasting to arbitrary subsets, so the set $\mathcal{N} - T$ does not refer to a small set of revoked users but to a set of users that have not subscribed to a certain transmission but might subscribe again to future transmissions.

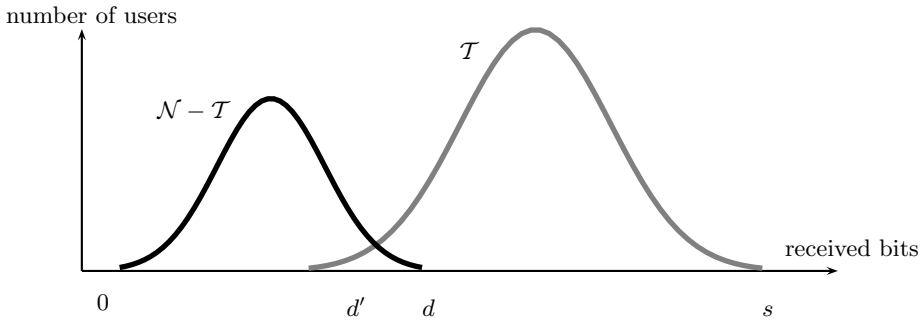

Fig. 1. Distribution of received key bits

context there is an additional "verifiability" requirement, as released parts of the secret have to be verifiable, such that a cheating party cannot gain valid parts of an honest party's secret, while sending random bits to this party. In the broadcast encryption setting the verifiability requirement can be neglected as the sender is trusted in this classical model and it is only a one-way release of secrets. This gradual probabilistic broadcast of secrets represents, to the best of our knowledge, a new probabilistic approach to broadcast encryption, which may foster further advance in broadcast encryption.

2.2 Setting Up the Scheme

The sender selects NK subsets $N'_1 \dots N'_{NK} \subset \mathcal{N}$, which are chosen uniformly from the set of all subsets of $\mathcal{N}$ with $\frac{1}{2}|\mathcal{N}|$ elements, so $|N'_i| = \frac{1}{2}|\mathcal{N}|$ for all $i = 1, \dots, NK$. For each subset N'_i a *subset key* $k_{N'_i}$ is generated and let $k_{N'_i} \in K_u \Leftrightarrow u \in N'_i$. So each user knows the key assigned to each subset she belongs to, but she does not know any other subset keys, so she stores $\approx \frac{1}{2}NK$ subset keys in total (note that key storage could be reduced heavily if a PRNG-based algorithm is used to generate keys before usage, but then unconditional security is not achievable).

We simplify the notion here as for implementation we do intend to split the users $\mathcal{N}$ in equal-sized batches of users and choose the subsets and the keys for each batch individually so we can parameterize the batch size, reduce the number of necessary subsets and are able to add new users to the broadcast scheme batch-wise after the scheme is in broadcasting operation. As the scheme is then set up and run for each batch individually we can still assume $\mathcal{N}$ to be the set of users, although $|\mathcal{N}|$ might be a rather small fixed-length number (e.g. ten thousand) compared to the possible millions of users of the full broadcasting group.

2.3 Broadcasting

For a transmission a set $\mathcal{T} \subset \mathcal{N}$ of valid subscribers is given and a session key k_S is generated for this transmission.

In order to broadcast k_S to the set $\mathcal{T} \subset \mathcal{N}$ of users, the sender sorts the subsets N'_i so that we can assume for the sorted subsets $N_i := N'_{\pi(i)}$ that

$$N_i \geq_{\mathcal{T}} N_{i+1} \qquad \forall i : 1 \leq i < NK \tag{1}$$

where for arbitrary subsets N_a, N_b we define

$$N_a \geq_{\mathcal{T}} N_b :\Leftrightarrow |N_a \cap \mathcal{T}| \geq |N_b \cap \mathcal{T}| \tag{2}$$

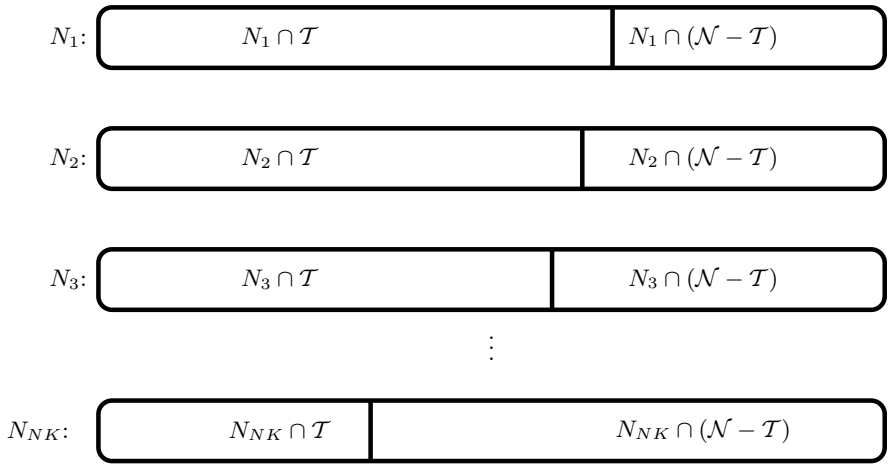

Fig. 2. Biased subsets

using a suitable permutation π for sorting the subsets. Loosely speaking, we let N_1 be the subset containing the highest number of subscribed users, N_2 is next with the second greatest bias towards the number of subscribers, so for small indices i, the bias of the subsets N_i towards $\mathcal{T}$ is high, see Figure 2 for illustration.

The scheme consists of two phases.

Phase 1 The session key k_S will be transmitted bit-wise: Let $k_{S,1}$, $k_{S,2}$, ..., $k_{S,s}$ denote the session key bits. First bit $k_{S,1}$ is sent to subset N_1 using subset key k_{N_1}, then $k_{S,2}$ is sent to N_2 etc. until all bits are sent. Let $\mathcal{T}'_j \subset \mathcal{T}$ be the set of users in $\mathcal{T}$ which have received at least d session key bits after $k_{S,j}$ is broadcasted.

Phase 2 For each user in $u \in \mathcal{T} - \mathcal{T}'_s$ we provide the full session key by using her unique secret key k_u^{indv} to encrypt an individual message for her and send it via the broadcast channel.

We are now interested in the number of messages sent in the two phases and the number of free-riders who can decrypt the session key although being a user in $\mathcal{N}-\mathcal{T}$ because they received d' key bits or more.

Theorem 1. *The number of free-riders given as a ratio of all users in $\mathcal{N}-\mathcal{T}$ can be approximated by $FR_{rat} = \Phi_{\mathcal{N}-\mathcal{T}}(d')$ where $\Phi_{\mathcal{N}-\mathcal{T}}$ is the distribution function of the normal distribution $N((1-\overline{t_s})s, \sqrt{s\overline{t_s}(1-\overline{t_s})^2})$ and where $\overline{t_s}$ is the average key-bit information received by the subscribed users per key-bit transmission: $\overline{t_s}$ can be approximated by $\overline{t_s} = \Phi^{-1}(1-\frac{s}{NK})$ where Φ^{-1} is the quantile function of the Gauss distribution $N(\frac{|\mathcal{T}|}{2}, \sqrt{\frac{|\mathcal{T}|}{2}\frac{|\mathcal{T}|}{|\mathcal{N}|}(1-\frac{|\mathcal{T}|}{|\mathcal{N}|})})$.*

The ratio of users receiving at least d bits in phase 1 is $SUC_{rat} = \Phi_{\mathcal{T}}(d)$ where $\Phi_{\mathcal{T}}$ is the distribution function of the the normal probability distribution $N(t_s s, \sqrt{s\overline{t_s}^2(1-\overline{t_s})})$.

The number of messages to be sent in phase 2, which is the the number of users not having received at least d bits in phase 1, can be given by $(1-SUC_{rat})|T| = (1-\Phi_{\mathcal{T}}(d))|T|$.

Proof. We first calculate the bias of the sub-sets $N_1 \dots N_j$. As the subsets were chosen uniformly we can approximate the binomial distribution of the values $t'_i := |N'_i \cap \mathcal{T}|$ for unsorted subsets N'_i to be Gaussian i.e. $\forall i = 1 \dots NK : t'_i \sim N(\frac{|\mathcal{T}|}{2}, \sqrt{\frac{|T|}{2}\frac{|\mathcal{T}|}{|\mathcal{N}|}(1-\frac{|\mathcal{T}|}{|\mathcal{N}|})})$. After sorting the NK subsets we have the most biased s values $t_i := |N_i \cap \mathcal{T}| > \frac{|\mathcal{T}|}{2}$ for $i = 1 \dots s$ with average values

$$\overline{t_i} = \Phi^{-1}(1-\frac{i}{NK}) \tag{3}$$

where Φ^{-1} be the quantile function of the Gaussian probability distribution $N(\frac{|\mathcal{T}|}{2}, \sqrt{\frac{|\mathcal{T}|}{2}\frac{|\mathcal{T}|}{|\mathcal{N}|}(1-\frac{|\mathcal{T}|}{|\mathcal{N}|})})$. As we assume the number $s \ll NK$ we approximate the value $t_i \approx \overline{t_s}$ for all $i < s$, so the bias' of all the subsets used for transmitting the key-bits are estimated to be equal to the bias of the last used sub-set in step s (note that, the scheme is more efficient than approximated here as the other biases are higher).

Using these approximations we can now calculate that each user in $\mathcal{T}$ has received every key-bit with probability $\overline{t_s} > \frac{1}{2}$, thus he has received $\overline{t_s}s$ key-bits on average and the number of key-bits received by each user is (by approximation) Gaussian distributed with parameters $N(\overline{t_s}s, \sqrt{s\overline{t_s}^2(1-\overline{t_s})})$. For the users is $\mathcal{N}-\mathcal{T}$ we have the probability $1-\overline{t_s} < \frac{1}{2}$, thus the distribution $N((1-\overline{t_s})s, \sqrt{s\overline{t_s}(1-\overline{t_s})^2})$.

Note, that on average a subscribed user has received $2s(\overline{t_s}-\frac{1}{2})$ more key-bits than non-subscribed users. □

Trade-off corollary result of the theorem: The scheme can be parameterized with the values s, d, d', NK (and with $|\mathcal{N}|$ and $|\mathcal{T}|$). From these values we can calculate (by approximation) the free-rider ratio FR_{rat} and success-ratio SUC_{rat}, hence the number of messages: $(1-SUC_{rat})|T| + s$.

2.4 Batches of Users

As mentioned before we intend to divide the user set into batches of a certain size which we still denote $|\mathcal{N}|$ to avoid unnecessary notations and run the scheme for each batch serially; let the number of batches that make up the real user-base be denoted by m so our total number of users is $|\mathcal{N}|m$. We now face the problem of selecting the parameters during the set-up phase: batch size $|\mathcal{N}|$, number of subsets NK – and for each transmission the parameters s, d and d'. There is obviously a tradeoff: For a smaller batch size, we have better biases and need less subsets (and less keys to be stored by the users), but we need to run the whole scheme more often and increase the transmission length.

Finally we have to identify the user's and attacker's computation power in order to select for a key size s the partial-key size values d and d'. In the next section we will improve our scheme so that this is not necessary anymore.

2.5 Improvements Based on Secret Sharing

The scheme introduced in the previous section uses a bit-wise distribution of the session key k_S. The scheme's security stems from the statistical certainty that unauthorized users receive on the average fewer bits (d' bits) of the session key than authorized users (d bits).

This basic scheme has some shortcomings, which are summarized below:

- The non-authorized users who are not free-riders do receive partial information as they receive a certain amount of key-bits.
- Authorized users have to perform an exhaustive search for up to $s - d$ bits of the session key. This could be costly.
- Each bit that an unauthorized user does not receive doubles his computational expense required for computing the full session key k_S. However, this still requires a rather large spread $d - d'$ between the number of bits received by authorized users in $\mathcal{T}$ and those received by unauthorized users in $\mathcal{N} \setminus \mathcal{T}$. Furthermore, estimating the computational power of adversaries is difficult, since exhaustive key-search can be easily parallelized and media content is sufficiently popular to attract many users in participating in a parallelized search for session keys.

Improved Biased Subset Scheme Distributing the secret session key in a bit-wise manner can be seen as a naive way of sharing the secret and distributing its shares to certain sub-sets, which cover the set of authorized users. We will see in the sequel that we can overcome these disadvantages by applying a cryptographic secret-sharing scheme.

We will use the notion of a (k, n) *secret sharing scheme* consisting of two algorithms: `Share` and `Reconstruct`. Given a secret s the sharing algorithm `Share`(s) outputs n shares $s_1, \ldots, s_n$. Given shares $s_{i_1}, \ldots, s_{i_k}$, the reconstruction algorithm `Reconstruct`$(s_{i_1}, \ldots, s_{i_k})$ outputs the original shared secret s

so given any k of the n shares, the original secret s can be reconstructed, but knowledge of less than k shares does not reveal any information.

For our construction one of the first proposed schemes (Shamir's scheme: [12]) can be used. This scheme shares a secret $s \in F$ (e.g., $F = Z_p$ with $p > n$) by choosing a random polynomial pol of degree $k-1$ and with constant term s (i.e., $pol(0) = s$) over F. The shares are defined as $s_i := (i, s(i))$, $i = 1, \ldots, n$, i.e., each share is a point of the polynomial. Given k different shares, the polynomial pol (and consequently the secret $s = pol(0)$) can be efficiently and uniquely reconstructed by performing a Lagrange interpolation.

The idea of applying secret sharing to overcome the limitations of the basic scheme is quite simple: The improvement is to replace the *bit-wise* broadcasting by broadcasting shares of the secret to the subsets instead.

In **Phase 1** of the improved scheme the sender applies a (d, s)-secret-sharing scheme to the key k_S, which results in the shares $s_1, \ldots, s_s$. Instead of encrypting and broadcasting single bits of the key k_S, the sender encrypts the share s_1 with sub-set key k_{N_1} and broadcasts the encrypted share (which can only be decrypted by members of N_1). Afterwards s_2 is sent to N_2, etc.

Given at least d shares a receiver can apply `Reconstruct` to efficiently reconstruct the secret k_S. Therefore, instead of performing an exhaustive search for the missing key bits, a receiver only performs a Lagrange interpolation to compute the complete session key. Moreover any unauthorized user in $N - T$ is unable to gain any information about the session key as long as he receives less than d shares. This is a significant improvement over the basic scheme, where an attacker could use extra time or extra computation power to derive more key-bits than ordinary users: the threshold value d in the improved scheme is a hard threshold and provides unconditional security.

2.6 Resiliency of the Schemes

The proposed schemes are highly vulnerable to colluders being able to combine their respective set of subset keys as these users would receive more key-bits or shares than any other user. In the case of two users sharing their subset key pool they would increase their portion of known subset keys from 0.5 to 0.75 each. This is higher than a reasonable bias being achieved by sorted subsets, thus the two users would become free-riders for all transmissions. Hence, the scheme does not offer any resiliency for colluders being a member of the same user batch. However, users from different batches can not gain anything from collusion as the scheme is run serially for each batch and different key encryption keys would be used. Therefore, the partial key information can not be combined at all.

The resiliency is therefore 1 from a worst-case point of view or dependent on the number of batches from an average-case point of view. It can easily be seen that the birthday paradox could be applied here if every user was assigned to a certain user batch uniformly chosen. So the resiliency of our proposed schemes can be approximated by the square root of the number of batches.

Fiat and Naor [6] describe a general applicable method to convert a scheme with low resiliency (1-resilient) to one with high resiliency by randomly grouping

users in small *random* sets (batches) and applying 1-resilient broadcast encryption in parallel to broadcast shares of the broadcasted secret. This construction could also be used to achieve higher resiliency for our scheme.

3 Further Improvements in Practice

3.1 Getting Rid of Free-Riders

It is possible to avoid the existence of free-riders if the biased-subset schemes are connected with a revocation scheme (e. g., with Naor et al.'s subset-difference scheme.

We take advantage of the following observations

- The set of free-riders is known by the sender. The set could be determined before the protocol is started in order to take reasonable precautions.
- Different boradcast encryption schemes can be combined, i. e. they can be run consecutively in a way that the first scheme distributes a pre-session key that is used for encrypting all the communication of the second scheme, thus, only the privileged users from the first scheme are able to take part in the session key distribution of the second scheme and the result is a set intersection operation of the privileged user sets. The same result can be achieved if two independent pre-session keys are distributed with the two schemes and the session key results from the exclusive-or operation of both pre-session keys.
- The biased-subset schemes introduce with respect to suitable parametrization a *small* amount of free-riders while the revocation schemes are designed for the exclusion of a *small* amount of revoked users so there is a straightforward concept to let the different schemes complement one another.

Taking this into account we can combine a revocation scheme with our scheme by first identifying the free-riders of the biased subset scheme and then run the revocation scheme **before** the biased subset scheme in order to distribute the pre-session key that excludes the set of anticipated free-riders from the protocol communication. As both protocols are unidirectional there is no real impact in the decision which scheme is run first because a receiver could always (if capable) record the broadcast and choose by itself which transmission is parsed first.

Note that the extension of the biased subset scheme with a revocation scheme does not support unconditional security for the combined scheme because the revocation schemes only offer computational security, but the construction is still useful for practical implementations, especially when the sender wants to avoid free-riders only for few transmissions.

In the annex we provide some sample data for a combined scheme consisting of the improved biased subset scheme and Naor at al.'s SD scheme.

3.2 Re-Using Establishment Keys for Stateful Receivers

If the receivers are not stateless, the agreed key for a certain transmission can be learnt and re-used as a subset key for a future protocol run. In practice it is likely that the target set of one transmission is very *similar* to the target set of a related transmission, so if the key is used as a future subset key the bias towards the target set will often be much higher than that of a normal sorted subset. Hence, the scheme will become more efficient for future runs when the receivers have stored their transmission session keys for each time the receiver was in the privileged set of users.

4 Conclusion

We proposed schemes for efficient broadcast key establishment that offer a trade-off between the ratio of free-riders and other parameters (overall key size or message header size). The schemes do not require stateful receivers and the second one is unconditionally secure (disregarding the existence of free-riders). Free-riders can also be prevented if revocation schemes are used together with our proposed schemes.

5 Evaluation (Sample Data)

In this section we evaluate our probabilistic broadcast encryption scheme by comparing its performance with that of existing broadcast encryption schemes. Comparison will be mainly in terms of communication overhead (i.e., broadcast message length), storage (number of keys and public storage) per user, as well as computational complexity per user. The latter will be measured in terms of generic operations, i.e., we will count the number of XOR-, PRNG, multiplication-, addition- and exponentiation operations.

Table 1. Performance results: our 1-resilient scheme compared to the 1-resilient schemes of Fiat and Naor [6] and revocation schemes of Naor, Naor and Lotspiech [1]. $|\mathcal{N}| = 10000$, $|\mathcal{T}| = 5000$, size of keys and shares is 64 bits

BE scheme	keys p. user	shares s	bits (header)	#ops p. user	FR_{rat}
Fiat&Naor [6]	10001	NA	10000	5001	0
Fiat&Naor [6] (OWF)	14	NA	10000	≈ 24974	0
Fiat&Naor [6] (Root)	10000 PKs	NA	10000	4999	0
Our scheme	1,000,000	1300	99200	650	0.05
Our scheme (Sect. 3.1)	1,000,014	1300	99200+138301	664	0
Trivial 1	2^{10000}	NA	64	1	0
Trivial 2	1	NA	320000	1	0
CS Revocation [1]	14	NA	320000	1	0
SD Revocation [1]	196	NA	639936	14 + 1	0

Furthermore, we focus our comparison on *1-resilient schemes* and rather *small numbers of users* n (in the order of 10000). These restrictions make comparisons between the different schemes possible: 1-resilient schemes are the basic building blocks for constructing k-resilient schemes by clever batching of users and serving each batch by an independent 1-resilient scheme. Note that this comparison is *unfair* to schemes being infinity-resilient; these are the trivial schemes and the revocation schemes (CS, SD).

Table 1 shows several sample data values for the proposed share-wise key distribution scheme. In all cases half of the user base is in the privileged set $\mathcal{T}$ (i.e., $|\mathcal{T}| = |\mathcal{N}|/2$), while the other half is not (generic case), and resilience is fixed as $k = 1$.

We summarize the performance for different user batch sizes $|\mathcal{N}|$. The free-rider ratios FR_{rat} are a parameter so that different number of shares (the phase 1 messages of our scheme) s and total number of messages are calculated from that parameter. The values are approximated average numbers.

References

1. Naor, D., Naor, M., Lotspiech, J.B.: Revocation and tracing schemes for stateless receivers. In: CRYPTO '01: Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology, Springer-Verlag (2001) 41–62
2. Halevy, D., Shamir, A.: The LSD broadcast encryption scheme. In Yung, M., ed.: CRYPTO '02. Volume 2442 of Lecture Notes in Computer Science., Springer (2002) 47–60
3. Goodrich, M.T., Sun, J.Z., Tamassia, R.: Efficient tree-based revocation in groups of low-state devices. In Franklin, M.K., ed.: CRYPTO '04. Volume 3152 of Lecture Notes in Computer Science., Springer (2004) 511–527
4. Jho, N.S., Hwang, J.Y., Cheon, J.H., Kim, M.H., Lee, D.H., Yoo, E.S.: One-way chain based broadcast encryption scheme. In: EUROCRYPT '05. (2005) to appear.
5. Shavitt, Y., Winkler, P., Wool, A.: On the economics of multicasting. Netnomics **6**(1) (2004) 1–20
6. Fiat, A., Naor, M.: Broadcast encryption. In: CRYPTO '93. (1993) 480–491
7. Dodis, Y., Fazio, N.: Public-key broadcast encryption for stateless receivers. In Feigenbaum, J., ed.: ACM Workshop in Digital Rights Management—DRM 2002. Volume 2696 of Lecture Notes in Computer Science., Springer-Verlag (2003) 61–80
8. Boneh, D., Silverberg, A.: Applications of multilinear forms to cryptography. Contemporary Mathematics **324** (2003) 71–90
9. Boneh, D., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. Cryptology ePrint Archive, Report 2005/018 (2005) http://eprint.iacr.org/.
10. Luby, M., Staddon, J.: Combinatorial bounds for broadcast encryption. In: EUROCRYPT '98. (1998) 512–526
11. Damgård, I.B.: Practical and provably secure release of a secret and exchange of signatures. Journal of Cryptology **8**(4) (1995) 201–222
12. Shamir, A.: How to share a secret. Commun. ACM **22**(11) (1979) 612–613