

Broadcast Secret-Sharing, Bounds and Applications^{*}

Ivan Damgård, Kasper Green Larsen, and Sophia Yakoubov

Aarhus University, {ivan, larsen, sophia.yakoubov}@cs.au.dk

Abstract. Consider a sender \mathcal{S} and a group of n recipients. \mathcal{S} holds a secret message \mathbf{m} of length l bits and the goal is to allow \mathcal{S} to create a secret sharing of \mathbf{m} with privacy threshold t among the recipients, by broadcasting a single message \mathbf{c} to the recipients. Our goal is to do this with information theoretic security in a model with a simple form of correlated randomness. Namely, for each subset A of recipients of size q , \mathcal{S} may share a secret random bit string with all recipients in A . We call this *Broadcast Secret-Sharing (BSS)* with parameters l , n , t and q .

Our main question is: how large must \mathbf{c} be, as a function of the parameters? We show that $\frac{n-t}{q}l$ is a lower bound, and we show an upper bound of $(\frac{n(t+1)}{q+t} - t)l$, matching the lower bound whenever $t = 0$, or when $q = 1$ or $n - t$.

When $q = n - t$, the size of \mathbf{c} is exactly l which is clearly minimal. The protocol demonstrating the upper bound in this case requires \mathcal{S} to share a key with *every* subset of size $n - t$. We show that this overhead cannot be avoided when \mathbf{c} has minimal size.

We also show that if access is additionally given to an idealized PRG, the lower bound on ciphertext size becomes $\frac{n-t}{q}\lambda + l - \text{negl}(\lambda)$ (where λ is the length of the input to the PRG). The upper bound becomes $(\frac{n(t+1)}{q+t} - t)\lambda + l$.

BSS can be applied directly to secret-key threshold encryption. We can also consider a setting where the correlated randomness is generated using computationally secure and non-interactive key exchange, where we assume that each recipient has an (independently generated) public key for this purpose. In this model, any protocol for non-interactive secret sharing becomes an *ad hoc threshold encryption (ATE) scheme*, which is a threshold encryption scheme with no trusted setup beyond a PKI. Our upper bounds imply new ATE schemes, and our lower bound becomes a lower bound on the ciphertext size in any ATE scheme that uses a key exchange functionality and no other cryptographic primitives.

^{*} This research was supported by: the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO); the NSF MACS project.

Table of Contents

Broadcast Secret-Sharing, Bounds and Applications	1
<i>Ivan Damgård, Kasper Green Larsen, and Sophia Yakoubov</i>	
1 Introduction	1
1.1 Applications	3
(Secret-Key) Threshold Encryption	3
Secure Multiparty Computation	3
1.2 Implementing Shared Keys	4
Using NIKE to get (Public-Key) Ad-Hoc Threshold Encryption .	4
Using Quantum Agreement	5
1.3 Related Work	5
Threshold Secret-Key Cryptosystems	5
Threshold Public-Key Cryptosystems	6
Pseudorandom Secret-Sharing	7
1.4 Open Problems	7
2 Definitions	7
2.1 BSS Syntax	8
2.2 BSS Security	8
3 Lower Bounds for Broadcast Secret Sharing	9
3.1 Warm-Up: BSS with $t = 0$ and $q = 1$	9
3.2 BSS with $t = 0$	10
3.3 Final BSS Lower Bound	14
4 Upper Bound on Ciphertext Size	16
4.1 Building Block: Pseudorandom Secret Sharing	16
4.2 Lower Bounding the Correlated Randomness When $H(C) = H(M)$	17
4.3 The Upper Bound	17
5 Bounds Additionally Assuming an Idealized PRG	18
5.1 Lower Bound on Ciphertext Size	19
5.2 Upper Bound	20
6 Application: Ad hoc Threshold Encryption	21
6.1 NIKE Definitions	21
6.2 ATE Definitions	21
6.3 ATE from NIKE and BSS	22
6.4 From ATE and NIKE to BSS	22

1 Introduction

In this paper, we consider the following scenario: We have a sender \mathcal{S} and a group of n recipients. \mathcal{S} holds a secret message m of length l bits, and the goal is to allow \mathcal{S} to create a secret sharing of m with privacy threshold t among the recipients. This should be done by broadcasting a single message c to the recipients, followed by local computation by the recipients.

Our goal is to do this with information theoretic security, and since this is clearly impossible in the plain model, we consider a model with correlated randomness. In doing so, one should be careful not to assume something “too strong” so the problem becomes trivial¹. We therefore choose the arguably simplest and easiest to implement form of correlated randomness where \mathcal{S} shares random strings with one or more of the recipients. More precisely, for each subset \mathcal{A} of recipients of size q , \mathcal{S} may share a secret random bit string $\mathfrak{s}_{\mathcal{A}}$ with all recipients in \mathcal{A} . Note that this particular form of correlated randomness is useful for applications with computational security because it can be extended by only local computation using a PRF, see Section 1.1 for details.

For any q , we also allow \mathcal{S} to share a secret with any subset smaller than q ². This means that, for larger q , we have stronger forms of correlated randomness.

We consider protocols where \mathcal{S} computes \mathfrak{c} from \mathfrak{m} and all the shared secrets ($\mathfrak{s}_{\mathcal{A}}$'s). Then \mathfrak{c} is broadcast, and each recipient computes his share of \mathfrak{m} from \mathfrak{c} and the shared secrets he holds. Security means that \mathfrak{c} and the information held by up to t recipients contain no information on \mathfrak{m} , but \mathfrak{c} and the information held by any $t + 1$ recipients determine \mathfrak{m} .

We call the notion we just sketched *Broadcast Secret-Sharing (BSS)*, with parameters l , n , t and q . In the following, we will sometimes refer to \mathfrak{c} as the *ciphertext* and the correlated randomness as *shared keys*, which is motivated by the fact that any broadcast secret sharing scheme can be used as is for a secret key threshold encryption scheme. More on this interpretation below.

Our main question is: how large must \mathfrak{c} be, as a function of the parameters? And, as a secondary question, how much secret correlated data do we need? To the best of our knowledge, these questions, as well the notion of broadcast secret-sharing, have not been considered before.

Let $l_{\mathfrak{c}}$ be the length of \mathfrak{c} . It is easy to see that

$$l \leq l_{\mathfrak{c}} \leq n \cdot l.$$

Namely, \mathfrak{c} must always carry enough information to transmit \mathfrak{m} to the receivers — and on the other hand, \mathcal{S} can always solve the problem by sharing a one-time pad key with each receiver, then making a standard secret sharing of \mathfrak{m} and letting \mathfrak{c} consist of the one-time pad encryptions of each of the shares.

In this paper, we show the much stronger conditions

$$\frac{n-t}{q}l \leq l_{\mathfrak{c}} \leq \left(\frac{n(t+1)}{q+t} - t\right)l.$$

Note that our upper bound matches the lower bound whenever $t = 0$ or when $q = 1$ or $n - t$. Note also that when $q = n - t$, the size of \mathfrak{c} is exactly l which is

¹ For instance, we could ask that \mathcal{S} has a random secret r of the same length as \mathfrak{m} and the recipients have shares of r in some linear secret-sharing scheme. Now, \mathcal{S} can broadcast $\mathfrak{m} - r$ which is clearly of minimal size, and the recipients adjust their shares accordingly.

² The motivation is that, for virtually any way to implement the shared randomness, \mathcal{S} could always share with $q' < q$ parties by imagining $q - q'$ virtual parties and emulate these herself.

minimal, so $q = n - t$ is the largest value it makes sense to consider. The protocol demonstrating the upper bound in this case requires \mathcal{S} to share a key with *every* subset of size $n - t$. We show that this (possibly exponential) overhead cannot be avoided when \mathbf{c} has minimal size.

Finally, we also show that if access is additionally given to an idealized PRG, the lower bound on ciphertext size becomes $\frac{n-t}{q}\lambda + l - \text{negl}(\lambda)$ (where λ is the length of the input to the PRG). The upper bound becomes $(\frac{n(t+1)}{q+t} - t)\lambda + l$. Namely, the sender chooses a PRG-seed, shares it among the receivers using the best available *BSS* and one-time pad encrypts the message using the output from the PRG.

1.1 Applications

We believe broadcast secret-sharing is interesting in its own right, and we describe below a couple of applications that make use of a BSS-scheme “out of the box”. As further motivation, we also consider in the following subsection two different ways to provide the correlated randomness, leading to other applications.

(Secret-Key) Threshold Encryption The first application is to secret-key threshold encryption, where a sender sends a ciphertext to set of receivers such that only large enough subsets can decrypt. The main difference between broadcast secret sharing and secret-key threshold encryption is that, in secret-key threshold encryption, it is important that the shared keys be *reusable*. We can easily achieve this by interpreting each key shared between \mathcal{S} and a (subset of) receiver(s) as a key for a pseudorandom function (PRF) ϕ . To encrypt, \mathcal{S} chooses a random nonce r , and for each shared key K , computes $\phi_K(r)$. Note that these PRF values form a (pseudorandom) set of values that can be used as fresh correlated randomness for the broadcast secret-sharing scheme we use. \mathcal{S} now uses this scheme to share her message \mathbf{m} among the receivers, resulting in a ciphertext \mathbf{c} , and sends the pair (r, \mathbf{c}) . Decryption can clearly be done by any subset consisting of at least $t+1$ receivers, and no smaller subset learns anything, which follows easily from security of the PRF and the underlying BSS-scheme. Note that decryption requires minimal interaction: each receiver just has to send his share to the others.

Note also that this application works exactly for the simple form of correlated randomness we use, where \mathcal{S} knows some keys, and each receiver knows a subset of them. Had we allowed a more complicated correlation, the receivers could not have generated new (pseudorandom) correlations of the same form simply by applying the PRF locally.

Secure Multiparty Computation A second application of BSS is to use it to non-interactively supply input to a secret-sharing based multiparty computation protocol, where the shared keys can be generated in an earlier setup phase.

Given an ideal functionality for distributing keys, we get information theoretic security if the shared keys are used once. But if we are happy with computational security, we can use a PRF as explained in the previous subsection to extend the key material and support any number of inputs. Note that this will not work when using the well-known method of “pre-cooking” a Shamir secret sharing of a random value known to the sender. Note also that our construction generates Shamir secret-sharings and so is compatible with standard MPC protocols.

1.2 Implementing Shared Keys

Broadcast secret sharing assumes keys shared between the sender and (subsets of) the receiver(s). To discuss the use of BSS in practice, we must also consider the distribution of these keys. We suggest two approaches: non-interactive key exchange (NIKE), and quantum key agreement.

Using NIKE to get (Public-Key) Ad-Hoc Threshold Encryption In this subsection, we discuss a way to generate the shared keys on the fly, via computationally secure and non-interactive key exchange. Here, we assume that each recipient has an (independently generated) public key and secret key for this purpose.

In this model, any protocol for BSS (including our upper bounds) implies a (public-key) *ad hoc threshold encryption (ATE) scheme*, which is a threshold encryption scheme with no trusted setup beyond a PKI. Namely, the sender creates a ciphertext that includes the information required for the key exchange as well as the c created for broadcast secret-sharing of the message m . To decrypt, at least $t + 1$ recipients will first compute the shared randomness using the key exchange, then use this to compute their shares, and finally exchange the shares to reconstruct m . In the related work section below, we give more background on ATE and its relation to standard threshold encryption.

Note that for $q = 1$ the non-interactive key exchange can be done very efficiently based on the DDH assumption: if each receiver i has a public key of form g^{x_i} in some appropriate group, then \mathcal{S} just needs to include a single element g^r for random r in the ciphertext, then the shared key will be of form $g^{x_i r}$ for receiver i . A similar solution for $q = 2$ can be designed using pairing friendly groups. Thus, for these cases, our upper bounds become (essentially) upper bounds on the ciphertext size of the corresponding ATE-scheme. In particular, the ATE-scheme that follows from this and our construction for $q = 2$ has smaller ciphertext size than the best previous scheme of Daza *et. al* [DHMR08]. For instance, when $t = 1$, that scheme has ciphertext size $(n - 1)l$ while we can obtain $(\frac{2n}{3} - 1)l$.

Less efficient non-interactive key exchange solutions also exist for larger values of q . They can be constructed from multilinear maps, indistinguishability obfuscation [BZ14], universal samplers [HJK⁺16,GPSZ17] (which can be built from indistinguishability obfuscation or functional encryption), or encryption combiners satisfying perfect independence [MZ17] (which can be built from universal samplers).

On the other hand, in this setting, our lower bound becomes a lower bound on the ciphertext size in any ATE scheme that uses an ideal functionality for key exchange (and perhaps for PRG), and no other cryptographic primitives. We formalize the demand that no other cryptographic primitives are used by requiring that the scheme is information theoretically secure when using the ideal functionalities.

We stress that these lower bounds hold for ATE-schemes that have access to the cryptographic primitives only via the ideal functionalities they implement. This is more restrictive than if black-box access were given to the corresponding algorithms; one might say that we allow the protocol to use them “only as intended”. However, to the best of our knowledge, no general lower bound was known for ATE before.

Using Quantum Agreement The correlated randomness needed for BSS can also be provided in a setting where the sender shares entangled quantum states with each of the receivers. As is well known, if sender and receiver share a pair of particles that are in the so-called EPR state, then measuring each particle results in the same random bit being obtained by both parties. Moreover, as long as the state really is the pure EPR state, no third party has any information on the randomness obtained. Thus this setting gives us exactly what we want for $q = 1$, with perfect security assuming perfect ability to prepare states and measure them. The same is true if one assumes that sender and receiver has executed a secure quantum key exchange protocol at some earlier time.

The case of $q > 1$ also has a quantum implementation, namely if we assume that the sender shares multipartite entangled states with subsets of receivers. In a multipartite entangled state, each involved party holds a particle, and the global state of the particles can be designed to be fully entangled so that local measurements return the same random result for all parties.

1.3 Related Work

Threshold Secret-Key Cryptosystems There is not much work on secret-key (symmetric) cryptosystems where the decryption and/or the encryption process can be distributed among a number of parties. A formal study of this was done by Agrawal et al. [AMMR18], in which formal security definitions and constructions were given for the case where both encryption and decryption is distributed. Our construction is in a different model where only the decryption is distributed. This allows us to offer new tradeoffs for constructions using only secret-key primitives and no public-key techniques, which is usually the more efficient case. The one construction from [AMMR18] using only secret-key primitives (a PRF) is very similar to our solution where $q = n - t$. It has minimal ciphertext size l but requires $\binom{n}{t}$ keys, potentially leading to exponential in n overhead. At the other extreme, we have the trivial solution where $q = 1$ and the sender secret-shares the message and sends a share to each receiver, leading to ciphertext size nl and a total of n keys. However, the construction leading to

our upper bound implies a spectrum of options “in between”, namely we can get ciphertext size $(\frac{n(t+1)}{q+t} - t)l$ using $\frac{n}{q+t} \binom{q+t}{t}$ keys.

Threshold Public-Key Cryptosystems The concept of public-key threshold encryption is very well known. It goes back at least to Desmedt *et. al* [DDFY94], and has since then been studied in a very long line of research. For this type of scheme, the key generation outputs a public key pk and a set of secret keys $\text{sk}_1, \dots, \text{sk}_n$ which are generated with respect to a threshold value t , where $0 \leq t < n$. Informally, the important security properties are that given any set of at least $t+1$ secret keys, one can decrypt a ciphertext encrypted under pk , while the encryption remains secure even given any set of t secret keys. For efficiency, ciphertexts should have size independent of n .

Requiring a single trusted execution of key generation can be very limiting, particularly in a system where parties may join at any point, or where senders want to dynamically choose subsets of the parties to be the recipients of a particular message. *Dynamic* threshold public-key encryption, introduced by Delerablée and Pointcheval [DP08], has a reduced setup requirement where the sender can pick the set of n recipients at encryption time; however, each recipient’s secret key must be derived from a common master secret key, so a trusted authority is still necessary. *Ad hoc threshold encryption* (ATE), first introduced by Daza *et. al* [DHMR08] as *threshold broadcast encryption*³ (motivated by its applicability to mobile ad hoc networks), requires no trusted setup beyond the absolute minimum — a PKI.

ATE considers a universe of users, where each user i has a public key pk_i and corresponding secret key sk_i , and where all key pairs are independently generated. A sender can select a set \mathcal{R} of n users and a threshold value t at the time at which he decides to send a message m . He can then construct a ciphertext $\mathbf{c} = E_{\text{pk}_{\mathcal{R}}, t}(m)$, where $\text{pk}_{\mathcal{R}}$ is the set of public keys belonging to parties in \mathcal{R} . ATE requires properties similar to those of standard threshold encryption: namely, that any $t+1$ parties in \mathcal{R} can decrypt, while the encryption remains semantically secure even given the secret keys of any t parties in \mathcal{R} .

Clearly, ATE has a number of attractive properties that standard threshold encryption lacks: no trusted authority, and the ability to decide on the set of receivers and the threshold on the fly. On the other hand, it is not clear that an ATE ciphertext can be as small as a standard one. The best known solution is from Daza *et. al* [DHMR08]. They show how to get ciphertext size linear in $n-t$. This solution is in our model discussed earlier (though it was not presented this way). Namely, it combines a BSS-scheme with non-interactive key exchange, where $q = 1$. In fact, their BSS scheme is a special case of our upper bound.

In this context, our lower bound shows that the ATE scheme of Daza *et. al* has optimal ciphertext size in the class of ATE schemes that use non-interactive

³ One should note that ATE for $t = 0$ is very similar to broadcast encryption: each party can decrypt on his own. However, in broadcast encryption, centralized key generation is usually allowed (or at least key generation is coordinated between receivers). This is exactly what is not allowed in ATE.

key exchange with $q = 1$ and no other cryptographic tools (but as mentioned above, it can be improved using $q = 2$). To the best of our knowledge, our bound is the first lower bound obtained for ATE schemes.

Reyzin *et. al* [RSY18] show that using indistinguishability obfuscation, as well as few standard primitives, it is possible to get ciphertext size independent of n . There are several reasons, however, why this is not a very satisfactory answer. For one thing, the construction requires that also senders have public and secret keys, which is not usually assumed for ATE (this is also one reason why that construction does not contradict our lower bound). Moreover, obfuscation requires strong assumptions; and with current state of the art techniques, it comes at the price of a huge loss of efficiency in practice.

Pseudorandom Secret-Sharing In [CDI05], Cramer *et. al* show that, in a model where sufficiently many independent random values are generated and each player is given an appropriate subset of these, the players can locally convert this information to a random Shamir secret-sharing (with a fixed threshold that depends on the set-up). This model is somewhat similar to ours. The crucial difference, however, is that we have a distinguished player - the sender - who knows all the values and can send a single message to the others. This allows us to create secret-sharings with any threshold, and while we do make use of their technique in our construction, we need additional new ideas to do so.

1.4 Open Problems

There is a very rich space of problems to explore. The most obvious open question is of course to close the gap between the upper and the lower bound on ciphertext size. Another problem is to understand how large the correlated randomness must be. Can the lower bound for minimal ciphertext size be generalized, or is there a way to get polynomial size randomness when the ciphertext is (close to) minimal size?

2 Definitions

In this section, we give the syntax and security definitions for broadcast secret sharing (BSS).

We consider the following random variables:

- $S_{\mathcal{A}}$, the random variable shared by the sender with the q parties in the set \mathcal{A} ,
- the message M , and
- the ciphertext C .

For ease of notation, we also let U be the random variable giving all the secrets $S_{\mathcal{A}}$ shared by the sender with any subset of receivers, U_i be the random variable giving all the secrets held by party \mathcal{P}_i (that is, $U_i = \{S_{\mathcal{A}}\}_{i \in \mathcal{A}}$), and $U_{\mathcal{A}}$ be the random variable giving the union of all the secrets held by parties in \mathcal{A} .

We use uppercase variables — S, U, M, C — to refer to distributions, and lowercase variables — s, u, m, c — to refer to concrete values.

2.1 BSS Syntax

We assume that any BSS scheme comes with a specification of finite sets from where the random variables are to be chosen. Hence, when we say in the following “any distribution of M ”, for instance, this means any distribution over the specified set of outcomes.

A BSS scheme with parameters (l, n, t, q) consists of two algorithms, described below.

$E_{u_{\mathcal{R}}}(\mathbf{m}) \rightarrow \mathbf{c}$ is a secret sharing algorithm (which we also sometimes dub *encryption*) that uses a set of keys $u_{\mathcal{R}} = \{u_i\}_{i \in \mathcal{R}}$ belonging to the parties in the size- n set \mathcal{R} of intended recipients (where each u_i consists of all secrets known to sets \mathcal{A} where $i \in \mathcal{A}$) to transform a length- l message \mathbf{m} into a secret sharing (or *ciphertext*) \mathbf{c} .

$D_{u_{\mathcal{A}}}(\mathbf{c}) \rightarrow \mathbf{m}$ is a reconstruction (or *decryption*) algorithm that uses keys $u_{\mathcal{A}} = \{u_i\}_{i \in \mathcal{A}}$ belonging to a subset \mathcal{A} of the intended recipient set \mathcal{R} (where $|\mathcal{A}| > t$) to recover the message \mathbf{m} from the sharing / ciphertext \mathbf{c} .

2.2 BSS Security

Informally, a BSS scheme is secure if any t parties in the designated set of receivers \mathcal{R} can learn nothing about a message from a ciphertext, but any $t + 1$ parties in \mathcal{R} can recover the message. More precisely:

Definition 1 (BSS Perfect Security). *A BSS scheme (E, D) is perfectly secure with threshold t if for any set of receivers \mathcal{R} of size n , for $C = E_{u_{\mathcal{R}}}(M)$, the following two properties hold for any distribution of M :*

Security *For any $\mathcal{A} \subset \mathcal{R}$ of size at most t , we have $H(M|C, U_{\mathcal{A}}) = H(M)$.*

Correctness *For any $\mathcal{A} \subset \mathcal{R}$ of size greater than t , we have $H(M|C, U_{\mathcal{A}}) = 0$. Furthermore, $M = D_{u_{\mathcal{R}}}(C)$.*

We can define statistical security similarly, where we assume that the distribution of the variables may also depend on a security parameter λ , but we always assume that the parameters l, n, t are polynomial in λ .

Definition 2 (BSS Statistical Security). *A BSS scheme (E, D) is statistically secure with threshold t if for any set of receivers \mathcal{R} of size n , for $C = E_{u_{\mathcal{R}}}(M)$, the following two properties hold for any distribution of msg :*

Security *For any $\mathcal{A} \subset \mathcal{R}$ of size at most t , we have $H(M|C, U_{\mathcal{A}}) \geq H(M) - \text{negl}(\lambda)$.*

Correctness *For any $\mathcal{A} \subset \mathcal{R}$ of size greater than t , we have $H(M|C, U_{\mathcal{A}}) \leq \text{negl}(\lambda)$. Furthermore, $M = D_{u_{\mathcal{R}}}(C)$ with overwhelming probability.*

Finally we define a different type of security that we will need later for technical reasons. It is designed for a situation where $t = 0$, so C alone reveals nothing about the message. Moreover, each player on her own can learn l' bits of the message, but not necessarily the entire message.

Definition 3 (BSS l' -Security). *A BSS scheme (E, D) is l' -secure if for any set of receivers \mathcal{R} of size n , for $C = E_{\mathcal{U}_{\mathcal{R}}}(\mathbf{M})$, the following two properties hold for any distribution of \mathbf{M} and some $l' \leq H(\mathbf{M})$:*

Security $H(\mathbf{M}|C) \geq H(\mathbf{M}) - \text{negl}(\lambda)$.

Correctness For any receiver \mathcal{P}_i we have $H(\mathbf{M}|C, \mathbf{U}_i) \leq H(\mathbf{M}) - l' + \text{negl}(\lambda)$.

Clearly, if a BSS-scheme is l' -secure for $l' = H(\mathbf{M})$, it is statistically secure in the case where $t = 0$.

3 Lower Bounds for Broadcast Secret Sharing

In this section, we prove a lower bound for BSS schemes with statistical security. Throughout the proofs, we consider sending a uniform random message \mathbf{M} of l bits. We then prove that the corresponding ciphertext of a BSS scheme must (roughly) satisfy $H(\mathbf{C}) \geq nH(\mathbf{M})/q = nl/q$. Since the entropy of a random variable giving a bit string is a lower bound on its expected length (Shannon's source coding theorem), this also lower bounds the length of the ciphertext. We prove the lower bound in steps, starting with the warm-up case $t = 0, q = 1$ and then extending it to arbitrary q and finally also to arbitrary t .

3.1 Warm-Up: BSS with $t = 0$ and $q = 1$

We start with a lower bound proof in the simple setup with threshold $t = 0$ and shared keys among $q = 1$ recipients. We let the message \mathbf{M} be a uniform random bit string of length l (hence $H(\mathbf{M}) = l$). We prove the following lower bound, where $\text{negl}(\lambda)$ may be replaced by 0 for perfect security:

Theorem 1. *For any BSS scheme with statistical security, n recipients, threshold $t = 0$ and sharing of keys with $q = 1$ recipients, we must have:*

$$H(\mathbf{C}) \geq n(l - \text{negl}(\lambda)).$$

To prove the lower bound, let \mathbf{S}_i for $i = 1, \dots, n$ denote the shared key received by the i 'th recipient (for $q = 1$, only i receives that random key). The high level idea in our proof is to argue that \mathbf{C} must contain a lot of information about the randomness \mathbf{S}_i for every index i . Since the shared keys are independent, this implies a lower bound on the entropy of \mathbf{C} . More formally, consider the mutual information $I(\mathbf{S}_i; \mathbf{C} \mid \mathbf{M}, \mathbf{S}_1, \dots, \mathbf{S}_{i-1})$. We will show:

Lemma 1. *For all recipients i , it holds that $I(\mathbf{C}; \mathbf{S}_i \mid \mathbf{M}, \mathbf{S}_1, \dots, \mathbf{S}_{i-1}) \geq l - \text{negl}(\lambda)$.*

Before proving Lemma 1, let us see how we use it to prove Theorem 1. Using non-negativity of entropy and the chain rule of mutual information, we have

$$\begin{aligned}
H(C) &\geq H(C | M) \\
&\geq H(C | M) - H(C | M, S_1, \dots, S_n) \\
&= I(C; S_1, \dots, S_n | M) \\
&= \sum_{i=1}^n I(C; S_i | M, S_1, \dots, S_{i-1}) \\
&\geq n(l - \text{negl}(\lambda)).
\end{aligned}$$

This completes the proof of Theorem 1. Thus what remains is to prove Lemma 1.

Proof (of Lemma 1). The basic idea in the proof of Lemma 1 is that C and S_i together reveal M , thus collectively they must have $l - \text{negl}(\lambda)$ bits of information about M . Since S_1, \dots, S_i alone have no information about M , those $l - \text{negl}(\lambda)$ bits must be accounted for in $I(C; S_i | M, S_1, \dots, S_{i-1})$. We prove that formally in the following. By definition, the mutual information in Lemma 1 equals:

$$\begin{aligned}
I(C; S_i | M, S_1, \dots, S_{i-1}) &= \\
H(S_i | M, S_1, \dots, S_{i-1}) - H(S_i | C, M, S_1, \dots, S_{i-1}).
\end{aligned}$$

The message M and all the shared keys are independent, hence $H(S_i | M, S_1, \dots, S_{i-1}) = H(S_i)$. Since entropy may only increase by dropping variables we condition on, we also conclude $H(S_i | C, M, S_1, \dots, S_{i-1}) \leq H(S_i | C, M)$. Using the definition of mutual information, we thus have:

$$\begin{aligned}
I(S_i; C | M, S_1, \dots, S_{i-1}) &\geq H(S_i) - H(S_i | C, M) \\
&= I(S_i; C, M) \\
&= H(C, M) - H(C, M | S_i).
\end{aligned}$$

Since the ciphertext C contains no information about M alone (up to $\text{negl}(\lambda)$), we have $H(C, M) = H(C) + H(M | C) \geq H(C) + H(M) - \text{negl}(\lambda)$. By the chain rule of entropy, we have $H(C, M | S_i) = H(C | S_i) + H(M | C, S_i) \leq H(C) + H(M | C, S_i)$. But $H(M | C, S_i) \leq \text{negl}(\lambda)$ since recipient i can recover M from C and S_i . We therefore have:

$$\begin{aligned}
I(S_i; C | M, S_1, \dots, S_{i-1}) &\geq H(C) + H(M) - \text{negl}(\lambda) - (H(C) + \text{negl}(\lambda)) \\
&= H(M) - \text{negl}(\lambda) \\
&= l - \text{negl}(\lambda).
\end{aligned}$$

■

3.2 BSS with $t = 0$

In this section, we generalize the lower bound from Section 3.1 to $q \geq 1$ (still assuming $t = 0$ and that the message M is a uniform random l bit string):

Theorem 2. *For any BSS scheme with statistical security, n recipients, security threshold $t = 0$ and sharing of keys with q recipients, we must have:*

$$H(\mathbf{C}) \geq n(l - \text{negl}(\lambda))/q.$$

To show this, we will show a stronger statement that will be useful for other purposes in the following:

Theorem 3. *For any l' -secure BSS scheme with n recipients, and sharing of keys with q recipients, we must have:*

$$H(\mathbf{C}) \geq n(l' - \text{negl}(\lambda))/q.$$

Clearly, this result implies Theorem 2: when \mathbf{M} is uniform and $H(\mathbf{M}) = l$, the assumption in Theorem 2 is equivalent to requiring l -security.

The basic idea in the proof for $q = 1$ was to argue that the ciphertext \mathbf{C} contained a lot of information about each \mathbf{S}_i . Formally, Lemma 1 showed that $I(\mathbf{C}; \mathbf{S}_i \mid \mathbf{M}, \mathbf{S}_1, \dots, \mathbf{S}_{i-1}) \geq l - \text{negl}(\lambda)$. In the following, we discuss the obstacles we face when generalizing the proof to $q \geq 1$ and show how we overcome them.

First, in order to prove Lemma 1, we used the fact that \mathbf{S}_i together with \mathbf{C} revealed \mathbf{M} to conclude that $I(\mathbf{C}; \mathbf{S}_i \mid \mathbf{M}, \mathbf{S}_1, \dots, \mathbf{S}_{i-1}) \geq l - \text{negl}(\lambda)$. Considering instead l' -security this statement would be $I(\mathbf{C}; \mathbf{S}_i \mid \mathbf{M}, \mathbf{S}_1, \dots, \mathbf{S}_{i-1}) \geq l' - \text{negl}(\lambda)$ and it could be proved in exactly the same way for $q = 1$.

However, since a recipient may now use all his shared keys to recover \mathbf{M} , we define a random variable \mathbf{U}_i for each recipient i : We let \mathbf{U}_i denote all shared keys held by recipient i ($\mathbf{U}_i = \{\mathbf{S}_{\mathcal{A}}\}_{i \in \mathcal{A}}$). Intuitively, the analog of Lemma 1 would state that $I(\mathbf{C}; \mathbf{U}_i \mid \mathbf{M}, \mathbf{U}_1, \dots, \mathbf{U}_{i-1}) \geq l' - \text{negl}(\lambda)$.

With this definition of \mathbf{U}_i we again have that \mathbf{U}_i and \mathbf{C} together reveal l' bits of \mathbf{M} . Unfortunately, the sets of shared keys held by different recipients are not disjoint. This means that \mathbf{U}_i may depend on $\mathbf{U}_1, \dots, \mathbf{U}_{i-1}$ and thus the lower bound on the mutual information is not necessarily true.

Our key idea for addressing the above issue is to further partition \mathbf{U}_i into subset $\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,q}$ where $\mathbf{U}_{i,k}$ contains all shared keys $\mathbf{S}_{\mathcal{A}}$ for which i is the k 'th smallest index in \mathcal{A} . Note that with this definition $\mathbf{U}_{i,k}$ and $\mathbf{U}_{j,k}$ with $i \neq j$ are disjoint sets of shared keys (only one index can be the k 'th smallest in a set \mathcal{A}) and thus are independent. The same holds for $\mathbf{U}_{i,j}$ and $\mathbf{U}_{i,k}$ with $j \neq k$ (i cannot both be the j 'th and k 'th smallest index in \mathcal{A}). Finally, we also define $F_{i,k}$ to denote the set of all shared keys $\mathbf{S}_{\mathcal{A}}$ in which i is the largest index in \mathcal{A} and $|\mathcal{A}| < k$. Our generalization of Lemma 1 then becomes:

Lemma 2. *There is an index $k \in \{1, \dots, q\}$ such that*

$$\sum_{i=1}^n I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) \geq n(l' - \text{negl}(\lambda))/q.$$

Before proving Lemma 2, let us see that it implies Theorem 2. We have:

$$\begin{aligned}
H(\mathbf{C}) &\geq H(\mathbf{C} \mid \mathbf{M}) \\
&\geq H(\mathbf{C} \mid \mathbf{M}) - H(\mathbf{C} \mid \mathbf{M}, \mathbf{U}_{1,k}, F_{1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) \\
&= I(\mathbf{C}; \mathbf{U}_{1,k}, F_{1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k} \mid \mathbf{M}) \\
&= \sum_{i=1}^n I(\mathbf{C}; \mathbf{U}_{i,k}, F_{i,k} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) \\
&\geq n(l' - \text{negl}(\lambda))/q.
\end{aligned}$$

What remains is thus to prove Lemma 2. The key step in doing so is to replace each mutual information in the sum by a term that only depends on the sets $\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,q}$ seen by the i 'th recipient. The rewriting is quite non-trivial and crucially relies on the fact that we applied the chain rule in reverse order of indices such that we condition on $\mathbf{U}_{j,k}, F_{j,k}$ for indices $j > i$. The rewriting we make uses the following:

Lemma 3. *For every recipient i and every index $k \in \{1, \dots, q\}$ we have*

$$I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) \geq I(\mathbf{U}_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}).$$

Let us first use Lemma 3 to prove Lemma 2.

Proof (of Lemma 2). Consider summing over all recipients and all choices of k , applying Lemma 3 on each term:

$$\begin{aligned}
\sum_{k=1}^q \sum_{i=1}^n I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) &\geq \\
\sum_{k=1}^q \sum_{i=1}^n I(\mathbf{U}_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}) &= \\
\sum_{i=1}^n \sum_{k=1}^q I(\mathbf{U}_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}) &= \\
\sum_{i=1}^n I(\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,q}; \mathbf{C} \mid \mathbf{M}) &= \\
\sum_{i=1}^n I(\mathbf{U}_i; \mathbf{C} \mid \mathbf{M}). &
\end{aligned}$$

Since \mathbf{U}_i and \mathbf{M} are independent, we have $I(\mathbf{U}_i; \mathbf{C} \mid \mathbf{M}) = H(\mathbf{U}_i \mid \mathbf{M}) - H(\mathbf{U}_i \mid \mathbf{C}, \mathbf{M}) = H(\mathbf{U}_i) - H(\mathbf{U}_i \mid \mathbf{C}, \mathbf{M}) = I(\mathbf{U}_i; \mathbf{C}, \mathbf{M}) = H(\mathbf{C}, \mathbf{M}) - H(\mathbf{C}, \mathbf{M} \mid \mathbf{U}_i)$. Since \mathbf{M} cannot be recovered from \mathbf{C} , we have

$$H(\mathbf{C}, \mathbf{M}) = H(\mathbf{C}) + H(\mathbf{M} \mid \mathbf{C}) \geq H(\mathbf{C}) + H(\mathbf{M}) - \text{negl}(\lambda).$$

By the chain rule, $H(\mathbf{C}, \mathbf{M} \mid \mathbf{U}_i) = H(\mathbf{C} \mid \mathbf{U}_i) + H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_i) \leq H(\mathbf{C}) + H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_i)$. But, by l' -security, l' bits of \mathbf{M} are determined from \mathbf{C} and \mathbf{U}_i , more precisely

$$H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_i) \leq H(\mathbf{M}) - l' + \text{negl}(\lambda).$$

We have thus shown $I(\mathbf{U}_i; \mathbf{C} \mid \mathbf{M}) \geq H(\mathbf{C}) + H(\mathbf{M}) - \text{negl}(\lambda) - (H(\mathbf{C}) + H(\mathbf{M}) - l' + \text{negl}(\lambda)) = l' - \text{negl}(\lambda)$. We therefore have:

$$\begin{aligned} \sum_{k=1}^q \sum_{i=1}^n I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) &\geq \\ &\sum_{i=1}^n l' - \text{negl}(\lambda) = \\ &n(l' - \text{negl}(\lambda)). \end{aligned}$$

Averaging over all choices of k completes the proof of Lemma 2. ■

To finish, we thus need to prove Lemma 3:

Proof (of Lemma 3). We need to show that for all recipients i and every index k , it holds that

$$I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) \geq I(\mathbf{U}_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}).$$

The main observation needed in the proof is the fact every shared key in $\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k}$ also appears in $\mathbf{U}_{i,k}, F_{i,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}$. More formally, we start by observing that:

$$\begin{aligned} I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) &\geq \\ I(\mathbf{U}_{i,k}; \mathbf{C} \mid \mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) &= \\ H(\mathbf{U}_{i,k} \mid \mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, \dots, F_{n,k}) - H(\mathbf{U}_{i,k} \mid \mathbf{C}, \mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, \dots, F_{n,k}). \end{aligned}$$

Notice that the set of shared keys $\mathbf{U}_{i,k}$ is disjoint from the sets $\mathbf{U}_{j,k}$ with $j \neq i$. This holds since for any set of receivers \mathcal{A} , only one receiver can be the k 'th smallest. Moreover, $\mathbf{U}_{i,k}$ is also disjoint from $F_{j,k}$ for all j . This is true since $F_{j,k}$ contains only shared keys for sets of receivers with cardinality less than k . This means that $\mathbf{U}_{i,k}$ is independent of $\mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}$ and thus we have

$$H(\mathbf{U}_{i,k} \mid \mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, \dots, F_{n,k}) = H(\mathbf{U}_{i,k}).$$

We therefore have:

$$\begin{aligned} I(\mathbf{U}_{i,k} F_{i,k}; \mathbf{C} \mid \mathbf{M}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}) &\geq \\ H(\mathbf{U}_{i,k}) - H(\mathbf{U}_{i,k} \mid \mathbf{C}, \mathbf{M}, F_{i,k}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}). \end{aligned}$$

Since entropy may only increase by removing variables that we condition on, we remove all shared keys from $F_{i,k}, \mathbf{U}_{i+1,k}, F_{i+1,k}, \dots, \mathbf{U}_{n,k}, F_{n,k}$ which do not appear in $\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}$. We claim that we are left with precisely the full set of shared keys appearing in $\mathbf{U}_{i,1}, \dots, \mathbf{U}_{i,k-1}$. To see this, consider a shared key $\mathbf{S}_{\mathcal{A}}$ appearing in $\mathbf{U}_{i,j}$ for some $j < k$. Assume first that i is the largest index in the set \mathcal{A} . Then the cardinality of \mathcal{A} is $j < k$ and we have $\mathbf{S}_{\mathcal{A}} \in F_{i,k}$ by definition of $F_{i,k}$. Next, assume that the cardinality of \mathcal{A} is less than k , but i is not the largest index in \mathcal{A} . Let $i' > i$ be the largest index. Then by definition, we have

$S_{\mathcal{A}} \in F_{i',k}$. Finally, assume that the cardinality of \mathcal{A} is at least k . Let $i' > i$ be the k 'th smallest index in \mathcal{A} , then $S_{\mathcal{A}} \in U_{i',k}$. In all cases, we have that $S_{\mathcal{A}}$ is in one of $F_{i,k}, U_{i+1,k}, F_{i+1,k}, \dots, U_{n,k}, F_{n,k}$ and we conclude that we are left with $U_{i,1}, \dots, U_{i,k-1}$. We therefore have:

$$\begin{aligned} I(U_{i,k} F_{i,k}; C \mid M, U_{i+1,k}, F_{i+1,k}, \dots, U_{n,k}, F_{n,k}) &\geq \\ H(U_{i,k}) - H(U_{i,k} \mid C, M, U_{i,1}, \dots, U_{i,k-1}). \end{aligned}$$

Conditioning on a random variable may only decrease entropy, we can therefore bound the above by:

$$\begin{aligned} I(U_{i,k} F_{i,k}; C \mid M, U_{i+1,k}, F_{i+1,k}, \dots, U_{n,k}, F_{n,k}) &\geq \\ H(U_{i,k} \mid M, U_{i,1}, \dots, U_{i,k-1}) - H(U_{i,k} \mid C, M, U_{i,1}, \dots, U_{i,k-1}) &= \\ I(U_{i,k}; C \mid M, U_{i,1}, \dots, U_{i,k-1}). \end{aligned}$$

This concludes the proof of Lemma 3 and thus also of Theorem 3. ■

3.3 Final BSS Lower Bound

In this section, we finally extend the lower bound in Theorem 2 to the general case of $t \geq 0$ and $q \geq 1$. Our final result is the following:

Theorem 4. *For any BSS scheme with statistical security, n recipients, security threshold t and sharing of keys with q recipients, we must have:*

$$H(C) \geq (n - t)(l - \text{negl}(\lambda))/q.$$

The proof follows via a reduction from the case with $t = 0$ (Theorem 2). The basic idea is to show that any BSS scheme for arbitrary threshold $t \geq 0$ can be converted into a scheme for $t = 0$ and $n - t$ receivers. This is done by treating the first t receivers as dummy receivers for which all shared keys are public information. This way, we get a BSS scheme with $t = 0$ for the remaining receivers $t + 1, \dots, n$.

In detail, consider all shared keys U_1, \dots, U_t held by the first t parties in a BSS scheme with threshold t . Consider any concrete instantiation u_1, \dots, u_t of the random variables and let E_{u_1, \dots, u_t} denote the event that $U_i = u_i$ for $i = 1, \dots, t$. We will prove that for most instantiations of $U_1 = u_1, \dots, U_t = u_t$, conditioned on E_{u_1, \dots, u_t} , the BSS statistical security definitions hold for the remaining $n - t$ receivers with threshold $t = 0$. Formally, we require that:

Security We have $H(M \mid C, E_{u_1, \dots, u_t}) \geq H(M) - \text{negl}(\lambda)$.

Correctness For any receiver i with $i \in \{t + 1, \dots, n\}$, we have

$$H(M \mid C, U_i, E_{u_1, \dots, u_t}) \leq \text{negl}(\lambda).$$

Call u_1, \dots, u_t *typical* if they satisfies the above Security and Correctness. If u_1, \dots, u_t are typical, then we have a BSS scheme with threshold $t = 0$ for the

remaining $n - t$ receivers $t + 1, \dots, n$ if we hard code $\mathbf{U}_1 = \mathbf{u}_1, \dots, \mathbf{U}_t = \mathbf{u}_t$ and let those be shared knowledge. Therefore, by Theorem 2, it must be the case for typical $\mathbf{u}_1, \dots, \mathbf{u}_t$, that

$$H(\mathbf{C} \mid E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) \geq \frac{n-t}{q}(1 - \text{negl}(\lambda)).$$

We will show:

Lemma 4. $\mathbf{U}_1, \dots, \mathbf{U}_t$ are typical with probability at least $1 - \text{negl}(\lambda)$.

Before we prove Lemma 4, we use the lemma to finish the proof of Theorem 4. We see that

$$\begin{aligned} H(\mathbf{C}) &\geq H(\mathbf{C} \mid \mathbf{U}_1, \dots, \mathbf{U}_t) \\ &= \sum_{\mathbf{u}_1, \dots, \mathbf{u}_t} H(\mathbf{C} \mid E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) \Pr[E_{\mathbf{u}_1, \dots, \mathbf{u}_t}] \\ &\geq \sum_{\mathbf{u}_1, \dots, \mathbf{u}_t: \mathbf{u}_1, \dots, \mathbf{u}_t \text{ are typical}} H(\mathbf{C} \mid E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) \Pr[E_{\mathbf{u}_1, \dots, \mathbf{u}_t}] \\ &\geq \frac{n-t}{q}(1 - \text{negl}(\lambda)) \Pr[\mathbf{U}_1, \dots, \mathbf{U}_t \text{ are typical}] \\ &= \frac{n-t}{q}(1 - \text{negl}(\lambda)). \end{aligned}$$

What remains is thus to prove Lemma 4.

Proof (of Lemma 4). Let $X(\mathbf{u}_1, \dots, \mathbf{u}_t)$ take the value $H(\mathbf{M}) - H(\mathbf{M} \mid \mathbf{C}, E_{\mathbf{u}_1, \dots, \mathbf{u}_t})$. Observe that since \mathbf{M} is independent of $\mathbf{U}_1, \dots, \mathbf{U}_t$, we have $H(\mathbf{M}) = H(\mathbf{M} \mid E_{\mathbf{u}_1, \dots, \mathbf{u}_t})$ and thus $X(\mathbf{u}_1, \dots, \mathbf{u}_t) = H(\mathbf{M} \mid E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) - H(\mathbf{M} \mid \mathbf{C}, E_{\mathbf{u}_1, \dots, \mathbf{u}_t})$. Conditioning on \mathbf{C} may only decrease entropy, hence X is non-negative for all $\mathbf{u}_1, \dots, \mathbf{u}_t$. It follows by Markov's inequality that

$$\Pr \left[X(\mathbf{U}_1, \dots, \mathbf{U}_t) > \sqrt{\mathbb{E}[X(\mathbf{U}_1, \dots, \mathbf{U}_t)]} \right] < \sqrt{\mathbb{E}[X(\mathbf{U}_1, \dots, \mathbf{U}_t)]}.$$

Now recall from the security requirements of a BSS scheme with threshold t that:

$$\begin{aligned} H(\mathbf{M}) - \text{negl}(\lambda) &\leq H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_1, \dots, \mathbf{U}_t) \\ &= \sum_{\mathbf{u}_1, \dots, \mathbf{u}_t} H(\mathbf{M} \mid \mathbf{C}, E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) \Pr[E_{\mathbf{u}_1, \dots, \mathbf{u}_t}], \end{aligned}$$

which implies

$$\begin{aligned} \mathbb{E}[X(\mathbf{U}_1, \dots, \mathbf{U}_t)] &= H(\mathbf{M}) - \sum_{\mathbf{u}_1, \dots, \mathbf{u}_t} H(\mathbf{M} \mid \mathbf{C}, E_{\mathbf{u}_1, \dots, \mathbf{u}_t}) \Pr[E_{\mathbf{u}_1, \dots, \mathbf{u}_t}] \\ &\leq \text{negl}(\lambda). \end{aligned}$$

Thus by Markov's, we have $\Pr \left[X(\mathbf{U}_1, \dots, \mathbf{U}_t) > \text{negl}(\lambda) \right] < \text{negl}(\lambda)$.

Next, for any receiver $i > t$, define $Y_i(\mathbf{u}_1, \dots, \mathbf{u}_t)$ to take the value $H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_i, E_{u_1, \dots, u_t})$. Since entropy is always non-negative, so is Y_i . By definition of conditional entropy, we have $\mathbb{E}[Y_i(\mathbf{U}_1, \dots, \mathbf{U}_t)] = H(\mathbf{M} \mid \mathbf{C}, \mathbf{U}_i, \mathbf{U}_1, \dots, \mathbf{U}_t)$. Thus from Markov's we again have $\Pr[Y_i(\mathbf{U}_1, \dots, \mathbf{U}_t) > \text{negl}(\lambda)] < \text{negl}(\lambda)$. It finally follows by a union bound that with probability at least $1 - (n - t + 1)\text{negl}(\lambda) = 1 - \text{negl}(\lambda)$, we simultaneously have $X(\mathbf{U}_1, \dots, \mathbf{U}_t) < \text{negl}(\lambda)$ and $Y_i(\mathbf{U}_1, \dots, \mathbf{U}_t) < \text{negl}(\lambda)$ for all $i = t + 1, \dots, n$. That is, $\mathbf{U}_1, \dots, \mathbf{U}_t$ are typical with probability at least $1 - \text{negl}(\lambda)$. ■

4 Upper Bound on Ciphertext Size

In this section, we explore constructions of broadcast secret-sharing.

4.1 Building Block: Pseudorandom Secret Sharing

Our results in this section leverage *pseudorandom secret sharing*, which is a technique for the local (that is, non-interactive) conversion of a replicated secret sharing to a Shamir secret sharing.

A *replicated secret sharing* for the $(t + 1)$ -out-of- n threshold access structure proceeds as follows. First, the dealer splits the secret \mathbf{M} into $\binom{n}{t}$ additive secret shares, where each share $r_{\mathcal{A}}$ corresponds to a different maximally unqualified set \mathcal{A} of size t . Then, the complement of each set \mathcal{A} (that is, the $n - t$ parties that are *not* in \mathcal{A}) are all given $r_{\mathcal{A}}$. It is then clear that any maximally unqualified set \mathcal{A} is only missing knowledge of one share $r_{\mathcal{A}}$, which any additional party holds.

Pseudorandom secret sharing [CDI05] locally converts such a replicated secret sharing into a Shamir secret sharing (a degree- t polynomial f with $f(0) = \mathbf{M}$ as the secret, and $f(i) = s_i$ as party i 's share for $i \in [1, \dots, n]$). Pseudorandom secret sharing proceeds as follows: let $f_{\mathcal{A}}$ be the degree- t polynomial such that $f_{\mathcal{A}}(0) = 1$, and $f_{\mathcal{A}}(i) = 0$ for all $i \in [n] \setminus \mathcal{A}$. Each player \mathcal{P}_i can then compute their Shamir share as

$$s_i = \sum_{\mathcal{A} \subseteq [n]: |\mathcal{A}|=n-t, i \in \mathcal{A}} r_{\mathcal{A}} f_{\mathcal{A}}(i).$$

We stress that, despite the name, pseudorandom secret-sharing as presented here provides perfect information theoretic security. The name comes from an application of the technique that uses pseudorandom functions.

Cramer, Damgård and Ishai [CDI05] also prove a lower bound, stated in Theorem 5.

Theorem 5 (From [CDI05]). *Fewer than $\binom{n}{t}$ independent random values shared among various subsets of parties cannot be locally converted into a $(t + 1)$ -out-of- n threshold secret sharing.*

4.2 Lower Bounding the Correlated Randomness When $H(\mathbf{C}) = H(\mathbf{M})$

Theorem 6. *For any perfectly secure BSS scheme with threshold $t = \theta(n)$, if $H(\mathbf{C}) = H(\mathbf{M})$, then correlated randomness of exponential size is necessary.*

Proof. If $H(\mathbf{C}) = H(\mathbf{M})$, then for any distribution of keys, there is exactly one ciphertext that corresponds to any given message. Therefore, choosing a ciphertext at random (without considering the correlated randomness) will always give a valid ciphertext that corresponds to some message, no matter which value the randomness takes. Choosing the randomness and ciphertext simultaneously independently at random thus produces a random $(t + 1)$ -out-of- n secret sharing (where the ciphertext is simply an additional random value given to all parties). So, the exponential lower bound by Cramer *et. al* [CDI05] (Theorem 5) on amount of independent randomness that can be converted into a $(t + 1)$ -out-of- n secret sharing applies. ■

4.3 The Upper Bound

Construction 1 below achieves optimal ciphertext size whenever $t = 0$, or when $q = 0$ or when q is the maximal relevant value $n - t$. We this result this by leveraging the techniques of replicated or pseudorandom secret sharing. The price we pays is that the overhead in terms of size of correlated randomness is sometimes exponential (that is, the sender and each of the receivers must use an exponential number of shared random values). Whether this happens depends on the parameter values.

Construction 1 *Let $n' = q + t$. We partition the recipients into $\frac{n}{n'}$ subsets of size $n' = q + t$. (We assume for simplicity that $n' = q + t$ divides n .) An arbitrary but fixed one of these subsets is chosen and named B . This is done publicly once and for all. We also assign once and for all a unique point in a suitable finite field to each recipient.*

Consider now any of the above subsets A . We set up the correlated randomness such that the sender \mathcal{S} shares a random value with any subset of A , of size $n' - t = q$. These values form a random replicated secret-sharing among the players in A and hence, using the technique from [CDI05], \mathcal{S} can share a random polynomial f_A of degree at most t with the participants in A , using only the correlated randomness. Concretely, \mathcal{S} knows f_A and each player in A knows a point on f_A .

The ciphertext consists of $\mathbf{m} + f_B(0)$ and $f_B - f_A$ for every subset $A \neq B$.

Each recipient locally computes from the correlated randomness $f_A(i)$ where A is the subset she is in and i is her assigned point in the field. Then she computes $f_B(i) = f_A(i) + (f_B - f_A)(i)$. To reconstruct, any subset of size at least $t + 1$ can interpolate f_B and compute $\mathbf{m} = (\mathbf{m} + f_B(0)) - f_B(0)$.

The security of this construction follows trivially from the security of replicated secret sharing: each f_A is uniformly random of degree at most t and so

$f_B - f_A$ contains no information on \mathbf{m} , even given $\mathbf{m} + f_B(0)$. Since each polynomial $f_B - f_A$ can be specified using $t + 1$ coefficients, the ciphertext size is

$$((t + 1)(n/(q + t) - 1) + 1)l = (n(t + 1)/(q + t) - t)l.$$

The size of the shared keys (correlated randomness) is $n/(q + t) \cdot \binom{q+t}{t}$ field elements. This can be as much as $\binom{n}{t}$ and so may be exponential in n . But as we showed above, at least when $q = n - t$, this overhead cannot be avoided.

5 Bounds Additionally Assuming an Idealized PRG

In this section, we add to our BSS model an idealized pseudorandom generator (PRG); an idealized functionality that takes in a random length- λ seed, and outputs a longer random value. (As long as the output is at least one bit longer than the input, we can bootstrap the PRG to give arbitrarily long outputs. In our case, the output length that most often makes sense is l , the length of the message.) Our BSS algorithms are augmented with oracle access to the idealized PRG.

We make some assumptions on how the BSS protocol may use the idealized PRG:

Definition 4. *An admissible BSS-protocol satisfies the following:*

- *For any subset of receivers, any PRG-seed chosen by the sender can either be computed using what that subset of receivers knows, or has full entropy (possibly up to a negligible loss).*
- *During the sharing phase, the sender chooses all seeds that are input to PRG uniformly, independently of anything else.*
- *The idealized PRG is not called with any shared keys as input.*

In the following we will only consider admissible BSS constructions. The motivation for this is as follows:

- We want to make sure that an admissible protocol can be turned into a construction in the real world by replacing the idealized PRG by a real PRG construction. Now, if a seed has (essentially) full entropy in the view of the adversary, then (and only then) can we use the standard security of a real PRG to conclude that the output is pseudorandom. Seeds for which the adversary has partial information are not useful in this sense, and we may as well give the adversary full information on that seed for free.

This is why we assume that in the view of a subset of receivers, any seed that the sender chose can either be computed or has (essentially) full entropy. However, for a seed to be potentially useful it must have full entropy in the first place, which is why we assume that the sender chooses all seeds uniformly, independently of anything else.

- We assume that the idealized PRG is not called using shared keys as input for simplicity, because this does not cost us any generality: calls to the PRG using shared keys as input is equivalent to asking for longer shared keys. In both cases, the result is a greater amount of correlated randomness.

Finally, we will assume that privacy only needs to hold given ability to call the PRG a polynomial number of times. The reason for this is that otherwise protocols that actually make use of the PRG could not ensure that the message is hidden from a non-qualified subset of receivers. As an example, suppose the sender secret-shares a seed s and includes in the ciphertext a one-time pad encryption $m \oplus PRG(s)$. A completely unbounded adversary can call the PRG on all inputs and now the only uncertainty she has is which seed the sender used. Then, if m is longer than s , it cannot have full entropy.

To be able to talk about the information a set of receivers can get from the oracle, we abuse notation and let $PRG(C, U_A)$ denote the random variable that is obtained by calling the PRG on inputs that are selected by an unbounded randomized algorithm that gets C, U_A as input. The algorithm only returns a polynomial number of outputs. For simplicity of notation, we suppress the algorithm and the random coins it uses.

Definition 5 (BSS Statistical Security with PRG). *A BSS scheme (E, D) is statistically secure with threshold t with respect to a random oracle PRG if for any set of receivers \mathcal{R} of size n , for $C = E_{U_{\mathcal{R}}}^{PRG}(M)$, the following two properties hold for any distribution of M :*

Security *For any $\mathcal{A} \subset \mathcal{R}$ of size at most t , we have $H(M|C, U_{\mathcal{A}}, PRG(C, U_{\mathcal{A}})) \geq H(M) - \text{negl}(\lambda)$.*

Correctness *For any $\mathcal{A} \subset \mathcal{R}$ of size greater than t , $H(M|C, U_{\mathcal{A}}, PRG(C, U_{\mathcal{A}})) \leq \text{negl}(\lambda)$. Furthermore, $M = D_{U_{\mathcal{R}}}^{PRG}(C)$ with overwhelming probability.*

5.1 Lower Bound on Ciphertext Size

Theorem 7. *Consider any BSS scheme that is statistically secure with threshold t with respect to PRG (which takes inputs of size λ) and shares messages of length $l \geq \lambda$. If the scheme is admissible it holds that*

$$H(C) \geq \frac{n-t}{q} \lambda + l - \delta(\lambda)$$

for a negligible function $\delta(\lambda)$.

To show the above theorem, consider first a scheme that satisfies the assumption with threshold $t = 0$, so then the only unqualified set of receivers is the empty set. Since the scheme is admissible, there is a (possibly empty) set of seeds \mathcal{S} that were chosen by the sender, but where each seed in \mathcal{S} has full entropy given the ciphertext C , and all other seeds are determined by C .

We claim that we can transform this scheme into a new one (for a different distribution of messages) that is l' -secure (Definition 3) with $l' = \lambda$. In particular,

this will be a scheme where the PRG is not available. Recall that in such a scheme a qualified subset of receivers can determine at least l' bits of the message.

To this end, we define the message M' in the new scheme to be the original M concatenated with the seeds in \mathcal{S} . Reconstruction in the new scheme by a qualified set \mathcal{A} works as follows: If at least one seed $s \in \mathcal{S}$ is determined by $C, U_{\mathcal{A}}$, then return s . Otherwise, by admissibility, all seeds in \mathcal{S} have full entropy given $C, U_{\mathcal{A}}$. Consider the random variable $PRG(C, U_{\mathcal{A}})$ that would have been used for reconstruction in the original scheme. Notice that since this variable is formed by calling the PRG a polynomial number of times, the inputs used will overlap with \mathcal{S} with only negligible probability. Therefore unless this overlap event happens, access to the PRG can be perfectly simulated without calling the PRG, simply by choosing fresh randomness to play the role of the PRG's output. Hence, we can return M with overwhelming probability without calling the PRG, so $H(M|C, U_{\mathcal{A}})$ is negligible, even without access to the PRG.

Since $l \geq \lambda$, we have shown that given $C, U_{\mathcal{A}}$ for a qualified set \mathcal{A} , the entropy of M' drops by at least l' bits (up to a negligible amount), and this is the correctness property of Definition 3.

The security property of Definition 3 follows immediately from admissibility and from the security property of Definition 5: given only C , all seeds in \mathcal{S} have full entropy and $H(M|C, U_{\mathcal{A}}, PRG(C, U_{\mathcal{A}}))$ can only increase if we take away the PRG and therefore do not condition on $PRG(C, U_{\mathcal{A}})$.

We can now apply Theorem 3 and since we did not change the distribution of C , we conclude:

Lemma 5. *For any BSS-scheme satisfying Definition 5 with $t = 0$, we have:*

$$H(C) \geq n(\lambda - \delta(\lambda))/q.$$

Proof (of Theorem 7). Given any BSS-scheme satisfying Definition 5, we can construct from this a new scheme for $n' = n - t$ receivers and threshold 0 (but the same ciphertext distribution). This is done by fixing the shared keys of the first t players and making them public, exactly as in the proof of Theorem 4, so we will not repeat the details here. We then apply the above lemma, and conclude that $H(C) \geq (n - t)(\lambda - \delta(\lambda))/q$. We finally obtain Theorem 7 by also noting that C must carry enough information to determine the message, so we can add l to the lower bound.

■

5.2 Upper Bound

Construction 2 describes how, using an idealized PRG in addition to shared keys, we can achieve

$$H(C) = (n(t + 1)/(q + t) - t)\lambda + l.$$

Construction 2 *The sender chooses a random PRG seed, uses the scheme from Construction 1 to share this seed among the receivers, and uses the PRG output on this seed to one-time-pad-encrypt the message.*

Ciphertext size and reconstruction follows trivially from Construction 1. As for security, it follows from security of Construction 1 that an unqualified set \mathcal{A} of receivers has no information on the seed chosen by the sender. Hence the event that the (polynomial number of) inputs to the PRG chosen by \mathcal{A} include the sender's seed has negligible probability. Unless this event happens, the message has full entropy, so the security property follows.

6 Application: Ad hoc Threshold Encryption

We can use any (l, n, t, q) BSS scheme together with any non-interactive key exchange (NIKE) scheme for $q + 1$ parties to get (l, n, t) ad hoc threshold encryption (ATE). Informally, the message sender uses the NIKE scheme to set up the correlated randomness for BSS non-interactively. She simply generates a fresh NIKE key pair, uses the secret key to derive shared secrets with every size- q subset of receivers, uses those shared secrets to run BSS, and sends the NIKE public key along with the resulting ciphertext to enable the recipients to derive the same shared secrets.

We sketch the definitions of NIKE and ATE below, and formalize how ATE can be instantiated from NIKE and BSS.

6.1 NIKE Definitions

A non-interactive key exchange (NIKE) scheme consists of two algorithms:

$KG(1^\lambda) \rightarrow (\mathbf{pk}, \mathbf{sk})$ is a randomized key generation algorithm that takes in the security parameter λ and returns a public-private key pair.

$KA(\mathbf{sk}_i, \mathbf{pk}_{\mathcal{A}}) \rightarrow \mathbf{s}$ is a key agreement algorithm that takes in one secret key and q public keys $\mathbf{pk}_{\mathcal{A}} = \{\mathbf{pk}_j\}_{j \in \mathcal{A}}$ and returns a shared secret.

Informally, a NIKE scheme for q parties is *correct* as long as, for any $i \in \mathcal{A}$ (where $|\mathcal{A}| = q + 1$), $\mathbf{s}_{\mathcal{A}} \leftarrow KA(\mathbf{sk}_i, \{\mathbf{pk}_j\}_{j \in \mathcal{A}, j \neq i})$ gives the same value. It is *secure* as long as, given $\{\mathbf{pk}_i\}_{i \in \mathcal{A}}$ (but none of the associated secret keys \mathbf{sk}_i), $\mathbf{s}_{\mathcal{A}}$ is computationally indistinguishable from random.

6.2 ATE Definitions

An ad hoc threshold encryption (ATE) scheme consists of three algorithms:

$KG(1^\lambda) \rightarrow (\mathbf{pk}, \mathbf{sk})$ is a randomized key generation algorithm that takes in the security parameter λ and returns a public-private key pair.

$E_{\mathbf{pk}_{\mathcal{R}}}(\mathbf{m}) \rightarrow \mathbf{c}$ is an encryption algorithm that encrypts a message \mathbf{m} to a set of public keys $\mathbf{pk}_{\mathcal{R}} = \{\mathbf{pk}_i\}_{i \in \mathcal{R}}$ belonging to the parties in the intended recipient set \mathcal{R} in such a way that any size- $(t + 1)$ subset of the recipient set should jointly be able to decrypt.

$D_{\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{A}}}(\mathbf{c}) \rightarrow \mathbf{m}$ is a decryption algorithm that uses secret keys $\text{sk}_{\mathcal{A}} = \{\text{sk}_i\}_{i \in \mathcal{A}}$ belonging to a subset \mathcal{A} of the intended recipient set \mathcal{R} (where $|\mathcal{A}| > t$) to decrypt the ciphertext \mathbf{c} and recover the message \mathbf{m} .

Informally, an (l, n, t) ATE scheme is *correct* if $D(E(\mathbf{M})) = \mathbf{M}$ (where D and E are run with the appropriate keys). It is *secure* if, for any two messages \mathbf{m}_0 and \mathbf{m}_1 of the same length l , $\mathbf{c}_0 = E_{\text{pk}_{\mathcal{R}}}(\mathbf{M}_0)$ and $\mathbf{c}_1 = E_{\text{pk}_{\mathcal{R}}}(\mathbf{M}_1)$ are computationally indistinguishable even given t or fewer of the secret keys sk_i , $i \in \mathcal{A}$.

6.3 ATE from NIKE and BSS

We can build an ATE scheme from a NIKE scheme and a BSS scheme as follows:

$KG(1^\lambda) \rightarrow (\text{pk}, \text{sk})$:

1. Return $(\text{pk}, \text{sk}) \leftarrow \text{NIKE.KG}(1^\lambda)$.

$E_{\text{pk}_{\mathcal{R}}}(\mathbf{m})$:

1. Run $(\text{pk}, \text{sk}) \leftarrow \text{NIKE.KG}(1^\lambda)$.
2. For every size- q subset $\mathcal{A} \subseteq \mathcal{R}$, run $\mathbf{s}_{\mathcal{A}} \leftarrow \text{NIKE.KA}(\text{sk}, \text{pk}_{\mathcal{A}})$.
3. Run $\text{BSS.c} \leftarrow \text{BSS.E}_{\mathbf{u}_{\mathcal{R}}}(\mathbf{m})$.
4. Return $(\text{pk}, \text{BSS.c})$.

$D_{\text{pk}_{\mathcal{R}}, \text{sk}_{\mathcal{A}}}(\mathbf{c} = (\text{pk}, \text{BSS.c}))$:

1. For every party $i \in \mathcal{A}$, for every size- q subset \mathcal{A}' such that $i \in \mathcal{A}'$, run

$$\mathbf{s}_{\mathcal{A}'} \leftarrow \text{NIKE.KA}(\text{sk}_i, \{\text{pk}\} \cup \{\text{pk}_j\}_{j \in \mathcal{A}', j \neq i}).$$

2. Recall that $\mathbf{u}_{\mathcal{A}}$ denotes $\{\mathbf{s}_{\mathcal{A}'}\}_{\mathcal{A}' \cup \mathcal{A} \neq \emptyset}$. Return $\mathbf{m} \leftarrow \text{BSS.D}_{\mathbf{u}_{\mathcal{A}}}(\text{BSS.c})$.

The size of a ciphertext in this scheme will be equal to the size of the corresponding BSS ciphertext plus the size of a NIKE public key.

6.4 From ATE and NIKE to BSS

Assume we have an ATE-scheme whose algorithms use an ideal NIKE functionality. We also assume that the ATE scheme is statistically secure when using the ideal NIKE functionality, that is, ciphertexts of different messages are statistically indistinguishable, and the message has full entropy in the view of a non-qualified set of receivers (up to a negligible amount).

From this, we can obtain a BSS scheme: the keys returned from the NIKE functionality become the correlated randomness, the encryption algorithm becomes the sharing algorithm, and the view of each receiver (including the ciphertext) is a share. Reconstruction is done by emulating the decryption protocol.

It therefore follows that our lower bound for BSS ciphertext size is also a lower bound for ciphertext size in any ATE scheme of the type we assumed.

References

- AMMR18. Shashank Agrawal, Payman Mohassel, Pratyay Mukherjee, and Peter Rindal. Dis: distributed symmetric-key encryption. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1993–2010, 2018.
- BZ14. Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, August 2014.
- CDI05. Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 342–362. Springer, Heidelberg, February 2005.
- DDFY94. Alfredo De Santis, Yvo Desmedt, Yair Frankel, and Moti Yung. How to share a function securely. In *26th ACM STOC*, pages 522–533. ACM Press, May 1994.
- DHMR08. Vanesa Daza, Javier Herranz, Paz Morillo, and Carla Ràfols. Ad-hoc threshold broadcast encryption with shorter ciphertexts. *Electron. Notes Theor. Comput. Sci.*, 192(2):3–15, May 2008.
- DP08. Cécile Delerablée and David Pointcheval. Dynamic threshold public-key encryption. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 317–334. Springer, Heidelberg, August 2008.
- GPSZ17. Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. Breaking the sub-exponential barrier in obfustopia. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 156–181. Springer, Heidelberg, April / May 2017.
- HJK⁺16. Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. How to generate and use universal samplers. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 715–744. Springer, Heidelberg, December 2016.
- MZ17. Fermi Ma and Mark Zhandry. Encryptor combiners: A unified approach to multiparty NIKE, (H)IBE, and broadcast encryption. Cryptology ePrint Archive, Report 2017/152, 2017. <http://eprint.iacr.org/2017/152>.
- RSY18. Leonid Reyzin, Adam Smith, and Sophia Yakubov. Turning hate into love: Homomorphic ad hoc threshold encryption for scalable mpc. Cryptology ePrint Archive, Report 2018/997, 2018. <https://eprint.iacr.org/2018/997>.