

Coding constructions for blacklisting problems without computational assumptions

Ravi Kumar¹, Sridhar Rajagopalan¹, and Amit Sahai²

¹ IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120.

{ravi, sridhar}@almaden.ibm.com

² Laboratory for Computer Science, M.I.T., Cambridge, MA 02139. This work was done while the author was visiting IBM Almaden Research Center. Also supported in part by DOD/NDSEG fellowship and in part by DARPA grant DABT63-96-C-0018.

amits@theory.lcs.mit.edu

Abstract. We consider the *broadcast exclusion* problem: how to transmit a message over a broadcast channel shared by $N = 2^n$ users so that *all but some* specified coalition of k excluded users can understand the contents of the message. Using error-correcting codes, and avoiding any computational assumptions in our constructions, we construct natural schemes that completely avoid any dependence on n in the transmission overhead.

Specifically, we construct: (i) (for illustrative purposes,) a randomized scheme where the server's storage is exponential (in n), but the transmission overhead is $O(k)$, and each user's storage is $O(kn)$; (ii) a scheme based on polynomials where the transmission overhead is $O(kn)$ and each user's storage is $O(kn)$; and (iii) a scheme using algebraic-geometric codes where the transmission overhead is $O(k^2)$ and each user is required to store $O(kn)$ keys. In the process of proving these results, we show how to construct very good cover-free set systems and combinatorial designs based on algebraic-geometric codes, which may be of independent interest and application.

Our approach also naturally extends to solve the problem in the case where the broadcast channel may introduce errors or lose information.

Keywords: blacklisting, broadcast encryption, copyrights protection, error-correcting codes

1 Introduction

Consider the problem of secure communication over a broadcast channel. Such channels, for example radio, cable TV, or the Internet MBONE, are designed to transmit messages to large groups of users. Over time, however, the broadcaster may not always desire that *all* users receive its messages. Thus one can consider the problem of transmitting a message only to some small specified subset of the users, and indeed this has been considered in the literature [1]. For the case of a broadcast channel shared by numerous users, we argue that an equally, if not more, interesting problem is the complement of the above problem: how to transmit a message over a broadcast channel shared by an exponential number $N = 2^n$ of users so that *all but* some specified small coalition of k excluded users can decipher the message, *even if these excluded users collude with each other in an arbitrary manner*. This is what we call the *broadcast exclusion problem* (also known as the *blacklisting problem*). We consider this problem as a communication problem in the most general setting possible, in particular without making any computational assumptions. We use the underlying tool of encryption *only as a black box*, and thus do not rely on any particular assumptions about how the encryption is implemented. Our focus is to design efficient solutions to this problem which minimize the communication overhead, while keeping the storage required by users and broadcasters feasible, i.e., polynomial in n and k . In this work, we use error-correcting codes and cover-free set systems to yield natural constructions to solve this problem, without the need for any computational assumptions. Our main result is an efficient scheme which achieves a communication blowup that depends only on the number of excluded users k . Such a result was not known previously even given computational assumptions. Our main technical tool in establishing these results is a new construction of cover-free set systems using algebraic-geometric codes.

Our results. Arguably, the two most important parameters in a broadcast exclusion scheme are the transmission overhead (i.e., the blowup in communication) and the storage requirement for each user. In optimizing these, however, it is important to ensure that the broadcaster's storage and computation requirements also remain feasible. Our main result states that it is possible to completely avoid any dependence on n in the transmission overhead, as long as the message is of some minimal length. We provide a constructive scheme using algebraic-geometric codes where the transmission overhead is $O(k^2)$ regardless of N , and each user is required to have storage $O(kn)$. We also present a scheme based on polynomials where the transmission overhead is $O(kn)$ and each user has storage $O(kn)$. Finally, we also present a scheme that depends on a randomized construction and thus requires the broadcaster to maintain a large database, but which has transmission overhead only $O(k)$ regardless of n and in which each user has storage $O(kn)$. All of our schemes naturally extend to provide efficient solutions to the broadcast exclusion problem where the broadcast channel may be noisy. In the process of proving these results, we show how to construct very good cover-free set systems and combinatorial designs which may be of independent interest.

We note that while it is standard when discussing communication problems to state results that are asymptotic with regard to the message length, we take care to show that our results hold even when messages are of moderate length (for example the length of

what one might expect an encryption key to be in the future). And indeed, our results are interesting even when the minimum length requirement is not met.

Background. Recently, there has been growing interest and research in providing security for transmissions over broadcast mediums. The general setting for such broadcast security problems is where one has access to pairs of black-box encryption and decryption devices, which can be used to send encrypted messages over the broadcast channel and decrypt them. In the most general computational setting, for binary broadcast channels these devices could function by means of preset one-time pads, while for analog broadcast channels (such as radio or cable), these devices could encrypt based on non-computational means taking advantage of the analog nature of the communication, or other properties of the channel. In a computationally limited world, these devices could function using private-key encryption, or using public-key encryption.¹ In this framework, to solve broadcast security problems like the one we consider, the objective would be to distribute decryption devices to users such that the following property would hold: given some set of excluded users, one could determine a small set of encryption devices, such that none of the excluded users would have the corresponding decryption devices, while each non-excluded user would have at least one such decryption device. One could then broadcast the message using these encryption devices, and only the designated recipients would be able to decrypt it.

The two major works on blacklisting problems that have appeared in the literature have been the relatively early work of Fiat and Naor [13] and the more recent work of Wallner *et. al.* [27], which have each led to several derivative works. In the framework laid out above, both these works take a similar approach: They take a universe of an exponential (in n) number of decryption devices, and distribute them according to some scheme to the users. As one might imagine, when given such a large number of decryption devices, it is quite simple to design distribution schemes which achieve the goals stated above. Indeed, [27] uses a simple balanced binary tree based scheme. By using such a large number of decryption keys, both [13] and [27] are in fact able to achieve stronger properties, for example the scheme of [27] can allow the number of excluded users k to be arbitrary, rather than fixed in advance. To deal with this structural problem of needing so many decryption keys, for the special case of private-key based broadcast, [13], and to some extent work following [27] by Canetti *et. al.* [4], make clever use of pseudo-random generators to efficiently access many pseudo-random private keys using only a few random ones.

To achieve constructions with feasible storage that do not require computational assumptions, however, these methods seem unfruitful. Thus, we take a completely different approach. Our approach yields results which compare quite favorably even against results that make computational assumptions, and also extend naturally to cover the case where the broadcast channel may introduce errors or lose information.

Our approach. The approach we take can be interpreted from the point of view of error-correcting codes. Fundamentally, we view the broadcaster's role as creating a noisy

¹ When we say user or server "storage," we abuse terminology to mean either the physical storage of encryption or decryption devices, or in the case that these are implemented digitally, to mean the storage of the information that allows them to operate.

channel—one that is very noisy indeed for the excluded users, but not so bad for the non-excluded users. Thus, by first passing the message through an appropriate error-correcting code, we are able to guarantee that non-excluded users will have enough information to reconstruct the message, while excluded users will gain no information. Interestingly, even the task of *creating* the noisy channel can be solved by using an error-correcting code with certain properties. This use of a two layer coding scheme, where the *outer code* is the one used to encode the message, and the *inner code* is used to create the noisy channel, is critical to realizing the gains in our results.

The inner code implements our noisy channel by building an interesting family of sets, which we call a (k, α) -cover-free family. This is a natural generalization of a family introduced by Erdős, Frankl, and Füredi [12]: A (k, α) -cover-free family $\mathcal{S} = \{S_1, S_2, \dots, S_N\}$ is one where the union of any k sets covers less than an α fraction of the elements in any other. A suitably chosen random construction can be used to design a (k, α) -cover-free family. This, however, has the disadvantage that it requires the explicit storage (space $N = 2^n$) of the set system by the broadcaster. To avoid this, a careful choice of codes turns out to be rather critical. We first demonstrate our construction based on Reed-Solomon codes. We then show that algebraic-geometric codes have specific properties that make them very suitable for our purpose, yielding a construction which allows us to achieve a transmission overhead that is independent of the number of users $N = 2^n$. The set systems we can construct using algebraic-geometric codes have quite strong properties, and thus may be of independent interest as well.

It is also clear from our approach that straightforward modifications to our construction (in particular, the outer code) yield blacklisting protocols that are robust against noisy broadcast channels.

A new application: Public broadcast groups. We now present a new application of broadcast security protocols which falls into our general framework. Recent and proposed improvements in the quality and bandwidth of networks (see for instance IEEE 1394 [14] or the Firewire website [15]) have led to a growing interest in building and using broadcast channels to provide content, but have also sparked growing concern that digital media, such as movies, can be easily duplicated and pirated. The proposed solution to this problem is for the broadcaster to maintain a *broadcast group*: a group of users who share a common key which can be used to encrypt all broadcasted messages to the group. The problem arises when some set of users need to leave the group—perhaps because they were caught misusing the content provided by the broadcaster. Now, the broadcaster needs to establish a new common key for all the users in the group *except* the set of users leaving the group. This is precisely the problem we address. Note that offending users or coalitions of users need only be blacklisted once, since all future communication (including key updates) occurs using a new common key unknown to the ousted users. Therefore it suffices to design a blacklisting protocol that can deal with a small number of excluded users at a time. This problem has motivated much study in broadcast security.

Now consider the same scenario, except where there may be many broadcasters utilizing the same broadcast channel to reach their broadcast groups. To avoid unnecessary duplication of effort, one can have a trusted facilitator set up users with pub-

lic key/private key based encryption and decryption mechanisms. Then by publishing these public keys, an arbitrary number of broadcasters could use the same underlying system to maintain their broadcast groups. Users would only have to remember one set of keys aside from a single group key for each broadcast group they belong to. We call a construction which implements this a *public broadcast group scheme*. Our results apply directly to this case, and allow for efficient public broadcast group schemes where the database of public keys is of feasible size. However, since the keys being used must come from some externally specified public key encryption scheme, pseudo-random keys need not suffice. Thus, applying the results of [13] or schemes descending from the scheme of [27] without computational assumptions would require an exponential number of public keys be stored for the purpose of maintaining public broadcast groups, whereas our results require only polynomially many public keys.

Prior work. The investigation of encrypted broadcasts to selected groups of users seems to have been initiated by Berkovits [1]. A number of works have followed, focusing mainly on the problem of broadcast security for the private-key case, under the assumption that one-way functions exist. Serious study was initiated in this area by [13] by defining the term *broadcast encryption*, a generalization of the blacklisting problem where the issue is how to send a message excluding an arbitrary subset S of N users so that no coalition $S' \subset S$ of at most k users can decrypt the message. In this definition, however, the computation time of the broadcaster is allowed to be proportional to the size of the excluded group. The result of Fiat and Naor is based on the idea of first constructing a scheme that works for excluding a single user (using exponentially many keys) and then using multi-layered hashing techniques to break up coalitions of excluded users into single exclusions from sub-groups. Finally by taking advantage of properties of pseudo-random generators, they are able to collapse the number of keys required. The constructions from [13], when applied to solve the broadcast exclusion problem, yield (assuming the existence of one-way functions) a scheme where the transmission overhead is $O(k^2 n \log^2 k)$, and each user needs to maintain $O(k n^2 \log k)$ keys.

A second major step in this area occurred when Wallner *et. al.* [27] proposed a hierarchical scheme for the broadcast exclusion problem where k , the number of users excluded, is allowed to be arbitrary. The scheme, however, requires an exponential number $O(N = 2^n)$ number of keys, but achieves a $O(kn)$ transmission overhead and requires only $O(n)$ keys for every user. Work of Canetti *et. al.* [4] show how to use pseudo-random generators to reduce by a constant factor the transmission overhead, but still require an exponential number of total keys. Canetti, Malkin, and Nissim [5] explore further trade-offs possible between the transmission overhead, user keys, and server keys possible by varying the scheme of [27]. Their constructions do not achieve feasible server storage, and indeed they give evidence that it is not possible to have communication overhead, user storage, and server storage all be polynomial in k and n if one requires security for arbitrary k .

A number of papers have also appeared following [13] attacking variants of the broadcast encryption and blacklisting problem under no computational assumptions. One variant that has been examined carefully is the notion of an unconditionally secure one-time broadcast encryption scheme, where the broadcasting can only be done once and all security guarantees must be unconditional. Blundo, Mattos, and Stinson [3]

are able to essentially give lower bounds and upper bounds of $O(1)$ for both the key size per user compared to the size of the message and the transmission overhead, but they require the message size be $O(N = 2^n)$. Without such restrictions but still in the one-time model, Stinson and van Trung [22] give a scheme with transmission overhead approximately worse than [13] but with user key size better than the unconditional [13]. Stinson and Wei [23] improve this further, yielding a scheme which still requires an exponential number of keys, but transmission overhead roughly $O(k^2 n)$. These papers on the one-time variant of blacklisting use some set systems like those we employ, but do so in conjunction with the techniques of [13] and in a way quite different from ours.

In related work, two methods—*watermarking* and *blacklisting*—are seen as the first line of defense against large-scale content piracy. The idea is for these two techniques to work in tandem. Each decoder chip will insert a watermark containing its identity into content that it decodes. Watermarking technologies have the following properties: (i) watermarks should be hard to replicate, erase, or modify, and (ii) the watermark reveals the identity of the owner of the copy. When this watermarked media is repeatedly leaked by someone, the identity of the offender is obtained from the watermark and blacklisted. For details on watermarking and similarly motivated techniques, see [8, 6, 11, 10].

We also note that the set system we construct from algebraic-geometric codes can be used to improve some aspects of the unconditional *Tracing Traitors* construction in [6].

Organization. Section 2 discusses preliminaries including cover-free families and background on codes. Section 3 gives an overview of our approach. Section 4 gives our randomized construction. Section 5 gives the construction based on polynomials and Section 6 gives the construction based on algebraic-geometric codes.

2 Preliminaries

2.1 Cover-free set systems

The notion of a *cover-free set system* is central to our approach. Let $[n]$ denote the set $\{1, 2, \dots, n\}$. Let N denote the number of users and B denote the set of excluded users.

We first provide a formal definition of cover-free set systems

Definition 1 ((k, α) -cover-free family). *Let U be a universe, where we write n for the size of U . A family of sets $\mathcal{S} = \{S_1, \dots, S_N\}$ where $S_i \subseteq U$ is a (k, α) -cover-free family if for all $S, S_1, \dots, S_k \in \mathcal{S}$,*

$$\left| S \setminus \bigcup_{i=1}^k S_i \right| \geq \alpha |S|.$$

A *k -cover-free family* is a (k, α) -cover-free family where α is unspecified, but strictly positive. In our applications, α will be a constant.

Constructing such a family with maximal N for a given universe size n was considered in [12]. In particular, the following theorem is established:

Theorem 1 ([12]). *Given any $N > 0$, there are k -cover-free set systems with N sets such that $|S_i| = O(kn)$ for all i , and $n = O(k^2n)$.*

However, Erdős *et. al.* establish this theorem non-constructively by showing that any maximal set system with sets of a size $n/4$ with the property that no two sets intersect at more than $n/(4k)$ points will have enough sets to make N large enough for Theorem 1 to hold.

We provide an alternative constructive, although probabilistic, proof of this theorem in Section 4.

It is interesting to note that the work of [6] in the context of tracing traitors also seems to involve the need to construct cover-free set systems.

2.2 Background on codes

Let \mathbb{F}_q be the finite field.

Definition 2 (Linear Code). *An $[n, k, d]_q$ -linear code is a k -dimensional subspace of \mathbb{F}_q^n where the minimum Hamming distance between elements of this subspace is d .*

The main codes we use in our constructions are Reed-Solomon and algebraic-geometric codes.

We first give the definition of Reed-Solomon codes. Let $F_{q,k}$ denote the set of polynomials on \mathbb{F}_q of degree less than k .

Definition 3 (Reed-Solomon Code). *Let $x_1, \dots, x_n \in \mathbb{F}_q$ be distinct and $k > 0$. The $(n, k)_q$ -Reed-Solomon code is given by the subspace $\{\langle f(x_1), \dots, f(x_n) \rangle \mid f \in F_{q,k}\}$.*

The fact that two distinct degree k polynomials can agree on at most k points can be restated as:

Property 1. The $(n, k)_q$ -Reed-Solomon code is an $[n, k, n - k + 1]_q$ -linear code.

For many applications, Reed-Solomon codes are quite restrictive since they require $n \leq q$. This motivates the use of algebraic-geometric codes (or *AG codes*), which are based on algebraic curves over finite fields [17].

We now give a definition of AG codes, also known as geometric Goppa codes (see [21, 20]). Let K/\mathbb{F}_q be an algebraic function field of one variable with field of constants \mathbb{F}_q and genus g . We denote the places of degree one of K by $\mathbb{P}(K)$. Suppose $|\mathbb{P}(K)| \geq n + 1$. For $\alpha < n$ and Q a place of degree one, let $L(\alpha Q)$ denote the set of all $f \in K$ which have no poles except one at Q of order at most α .

Definition 4 (AG Code). *Let $Q, P_1, \dots, P_n \in \mathbb{P}(K)$ be distinct. The AG code $C(\alpha Q; P_1, \dots, P_n)$ is given by the subspace $\{\langle f(P_1), \dots, f(P_n) \rangle \mid f \in L(\alpha Q)\}$.*

By the Riemann-Roch Theorem,

Property 2. The AG code $C(\alpha Q; P_1, \dots, P_n)$ is an $[n, k, d]_q$ -linear code for $d \geq n - \alpha$ and $k \geq \alpha - g + 1$.

Reed-Solomon codes are straightforward to construct. There are a few AG codes in existence today (see [26, 18, 16]). We choose for sake of clarity (as did [20]) to focus on the family of codes given by the explicit and fairly simply Garcia-Stichtenoth function fields [16]. More details are given in Section 6.

2.3 Terminology

Although we consider our problem in a very general setting, for ease of understanding we will often talk of “keys” for encryption and decryption as if we were in the private-key setting. We use “key” to simply tie together a pair of encryption and decryption devices. Then when we talk of either encrypting or decrypting using a particular key, we mean to use the corresponding encryption or decryption device. When we talk of “giving a key to user” or a “user’s key,” we are referring to the decryption mechanism corresponding to that key.

2.4 Warmup: A solution when $|B| = 1$

A simple solution to the broadcast exclusion problem exists when the excluded set is a singleton, i.e., $B = \{i\}$. Let $n = n$. We choose a set S of $2n$ keys and give a different subset S_i of the keys S , $|S_i| = n$ to each user i . Since $N < \binom{2n}{n}$, this is possible. To exclude user i , we transmit the message encrypted n times: once using each key in $S \setminus S_i$. Since each user other than i has at least one key in $S \setminus S_i$, the message can be decoded by the user. Clearly, i cannot decode the message. The storage overhead is n and the transmission overhead is n .

A central result in [13] is that for the case of private-key encryption, the transmission overhead can be reduced to $O(1)$ for this special case. This is done via a clever use of one-way functions which allows all but the excluded user to reconstruct the session key for the transmission.

3 The overall construction

In this section, we describe our basic construction for the broadcast exclusion problem. Suppose we want that up to k people can be excluded at any given time. The basic idea is to have a set $K = \{k_1, k_2, \dots, k_m\}$ of encryption/decryption keys, of which each user x gets assigned some subset S_x of u out of the m keys. When some message M is to be broadcast, it is “digested” using an error-correcting code into m smaller but redundant pieces (M_1, M_2, \dots, M_m) , which are then encrypted according to the keys K to produce $(E_{k_1}(M_1), E_{k_2}(M_2), \dots, E_{k_m}(M_m))$. The pieces corresponding to keys belonging to users who have been excluded are then discarded, and the remaining encrypted pieces are broadcast to all users. By decrypting the pieces corresponding to the keys that each valid user has, the user reconstructs the original message M .

For this scheme to work, we must have two properties: (i) After discarding the pieces of the message that an excluded user could intercept, we must ensure that enough pieces are left for each non-excluded user. We must try to achieve this while trying to maximize the number of possible users N we can support (i.e., the number of subsets we can make) in terms of the number of pieces m into which we split the message. This is what we call the *inner code* problem. (ii) We must make sure that given the number of pieces of the message guaranteed by the inner code, each non-excluded user can reconstruct the original message. On the other hand, we do not want to waste resources, so the pieces should be as small as possible. This is what we call the *outer code* problem.

The issues. In our construction, we will try to optimize several quantities. First, we want to minimize the blowup in the communication complexity of broadcast in terms of the number of excluded users and the total number of users. We recall the best known construction has a blowup that is related to *both* the number of total users (which we think of as quite large—e.g., the number of cable TV subscribers) as well as the number of excluded users (which we think of as much smaller—e.g., the number of criminal TV pirating organizations). Ideally, we would like the blowup in the communication to be related only to the number of excluded users.

We also consider the problem of storage. For the user, we want to minimize the number of keys that each user must store. We would also like to avoid having the broadcaster store a large database with information for each user. Rather, we would like the broadcaster to be able to efficiently generate the relevant information about a user given only his identity. Although we do give one randomized scheme that requires a large database, we focus mainly on the problem where storage requirements of the broadcaster are not to depend polynomially on the number of users.

The methods. We will focus on the inner code and the outer code separately. First, we observe that the outer coding problem can be solved trivially by a simple repetition code—i.e., the message M is simply repeated so $M = M_1 = M_2 = \dots = M_m$.

The inner code. Now, if we can simply solve the inner coding problem so that at least one of the message pieces can be decrypted by each good user after all the message pieces corresponding to the k excluded users have been removed, then our construction will work. Constructing such set systems is precisely the problem Erdős *et al.* considered [12], i.e., finding families of finite sets such that no set is covered by the union of k others. In Section 4, we describe a randomized construction that matches the lower bound of [12]. In Sections 5 and 6, we give constructions based on error correcting codes. When building an inner code, not just any “good” error correcting code will do. The fact that we want sets to have very small intersections corresponds to the underlying code having good minimum distance at very low rates. Polynomial codes have this property. To optimize transmission overhead in conjunction with the outer code, however, it will turn out that we want the sets in our set system to be fairly large. This corresponds to underlying field of the code being small. Polynomial codes fail this property but this is precisely what algebraic-geometric codes were designed for. These intuitions are made formal in Sections 5 and 6.

The outer code. The purpose of the outer code is to eliminate waste in the broadcast transmission, i.e., to allow users to fully utilize the (limited) information they are able to decrypt during a broadcast. For this reason, the outer code should be a good low-rate erasure code. In this exposition, we will focus on polynomial-based codes, which have the advantage of being *perfect* erasure codes in that no information is wasted. We could also employ other erasure codes with smaller alphabet sizes to minimize how long the message being transmitted has to be, but we omit the details of these issues in this extended abstract.

We define an $[n, k, m]_q$ (constructive) erasure code to be a (polynomial-time) function $C : \mathbb{F}_q^k \mapsto \mathbb{F}_q^m$ such that there exists a (polynomial-time) function $D : \mathbb{F}_q^m \mapsto \mathbb{F}_q^k$,

where $\bar{\mathbb{F}} = \mathbb{F} \cup \{\perp\}$, such that: For all $v \in \mathbb{F}_q^k$, if $u \in \bar{\mathbb{F}}_q^n$ is such that u agrees with $C(v)$ on at least m places, and is \perp elsewhere, then $D(u) = v$. In other words, one can erase up to $n - m$ positions of the codeword and still reconstruct the original message. Clearly, the best one can hope for is to have $m = k$. Such an erasure code is called a *perfect code*.

Polynomials. Here we recall that polynomials allow one to construct perfect $[n, k, k]_q$ erasure codes for any $k \leq n \leq q$. This is called the *Reed-Solomon code*. Given a vector $v = (v_0, v_1, \dots, v_{k-1}) \in \mathbb{F}_q^k$, we construct the polynomial $p_v(x) = v_0 + v_1x + \dots + v_{k-1}x^{k-1}$. Let e_1, e_2, \dots, e_n be some canonically chosen elements of \mathbb{F}_q . Then $C(v) = (p_v(e_1), p_v(e_2), \dots, p_v(e_n))$. In the decoding scenario, we are given k pairs of the form $(e_i, p_v(e_i))$. Since we know $\deg(p_v) < k$, simple polynomial interpolation over \mathbb{F}_q will efficiently allow us to reconstruct the coefficients of p_v , yielding the original vector v .

Communication. It follows from the definition that if C is an $[n, k, m]_q$ erasure code, then the length of $C(v)$ is $n \log q$ bits, whereas the length of the original message v is $k \log q$ bits. Hence the communication blowup is n/k .

Fault tolerance. We see immediately that by changing the parameters of the outer code, we can easily make all our construction tolerant to faults in the broadcast channel. If the channel can drop packets, then simply by using the erasure correction properties we have already discussed, robustness can be achieved. If the channel also corrupts data, one can use a code with both erasure and error-correcting capabilities. This is possible for example with Reed-Solomon codes [2]. We omit the details of how these parameters should be chosen to achieve different levels of robustness in this extended abstract.

4 A randomized construction

We now describe a randomized construction of a broadcast exclusion scheme that works with high probability.

The inner code. We begin with a randomized construction that with high probability yields a set system with the desired properties. This construction matches the lower bound given in [12]. Let the universe $U = [m]$. We wish to construct a (k, α) -cover-free family \mathcal{S} . For sake of clarity, we give our construction only for $\alpha = 1/2$ (which we will use later), although the construction can easily be seen to work for any constant α . Also, for sake of clarity, we give present an analysis that is not tight in the calculation of some of the constants, but captures the essence of what we can achieve in this randomized construction given below:

Partition U into $m/2k$ sets $U_1, U_2, \dots, U_{m/2k}$ each of size $2k$. Pick each of N sets as follows: Pick $x_i \in U_i$ uniformly for each i , and let $S = \{x_1, \dots, x_{m/2k}\}$.

We now show the following theorem:

Theorem 2. *There exists a universal constant c such that for any $t > 0$, if*

$$N = \exp\left(\frac{cm}{k(k+1)} - \frac{t}{k+1}\right),$$

then the probability that the above set system is not $(k, 1/2)$ -cover free is at most $\exp(-t)$.

Proof. Without loss of generality, let $k \geq 2$. Fix sets $S, S_1, S_2, \dots, S_k \in \mathcal{S}$. For any $x \in U_j$ for some j , we have

$$\Pr\left[x \notin \bigcup S_i\right] = \left(1 - \frac{1}{2k}\right)^k \geq \frac{9}{16}.$$

Hence,

$$\mathbb{E}\left[\left|S \cap \left(\bigcup S_i\right)\right|\right] = \left(1 - \left(1 - \frac{1}{2k}\right)^k\right) \cdot |S| \leq \frac{7}{16} \cdot |S|.$$

By Chernoff bounds [7],

$$\Pr\left[\left|S \cap \left(\bigcup S_i\right)\right| > \frac{|S|}{2}\right] \leq \exp\left(2 \cdot \left(\frac{1}{16}\right)^2 \cdot |S|\right) = \exp\left(\frac{m}{256k}\right).$$

By the union bound, the probability that this occurs for any choice of S, S_1, S_2, \dots, S_k —i.e., that \mathcal{S} does not satisfy the condition given in the theorem—is at most

$$N^{k+1} \cdot \exp\left(\frac{m}{256k}\right).$$

Letting N as in the statement of the theorem, we see that this probability is at most $\exp(-t)$. \square

It is immediate that if we are willing to tolerate the broadcaster having a large database to store the keys assigned to each user, this randomized construction, when used in conjunction with a trivial repetition outer code, yields a broadcast exclusion scheme with a blowup in communication of $O(k^2 n)$.

The outer code. We will now see how to use an outer code to reduce this communication blowup to $O(k)$. The inner code construction above shows that to support N users, we need to have $m = O(k^2 n)$ total keys and total messages sent, with each user receiving $m/2k = O(kn)$ keys. Moreover, the construction guarantees that with high probability, even with k users blacklisted, every good user will be able to decrypt at least $m/4k$ of the messages.

Hence, we may use a $[m, m/4k, m/4k]_q$ constructive erasure code as the outer code, such as the polynomial construction given earlier. Here, we pick the field size $q \geq m$. Thus, if we assume the message is at least $(m/4k) \cdot \log q = \Omega(kn \log k \log n)$ bits long (which is reasonable even if the message is a single strong cryptographic key), the communication blowup is reduced to simply $4k$. Note that this bound on the message size can be even improved further using more sophisticated erasure codes. The details are omitted from this extended abstract.

5 Construction based on polynomials

We now give a deterministic construction based on polynomials.

The inner code. Let m be such that $q = \sqrt{m}$ is a prime power and let $\mathbb{F}_q = \{u_1, \dots, u_q\}$ be an enumeration of the field. Let $F_{q,d}$ be the set of polynomials on \mathbb{F}_q on degree at most d . Let the universe $U = \mathbb{F}_q^2$, $|U| = m$ and let $d = q/2k$. Let the set system $\mathcal{S} = \{S_f \mid f \in F_{q,d+1}\}$, where $S_f = \{\langle u_1, f(u_1) \rangle, \dots, \langle u_q, f(u_q) \rangle\} \subset U$. We denote the size of \mathcal{S} by N .

We now show the following theorem:

Theorem 3. *The set system \mathcal{S} is $(k, 1/2)$ -cover free with*

$$N = \exp\left(\frac{\sqrt{m} \log m}{4k}\right).$$

Proof. Notice that the construction of \mathcal{S} corresponds exactly to an $(q, d+1)_q$ -Reed-Solomon code. Since two distinct degree d polynomials can agree on at most d points, we see that $|S_i \cap S_j| \leq d$ for any distinct $S_i, S_j \in \mathcal{S}$. So, for any $S, S_1, \dots, S_k \in \mathcal{S}$,

$$\left|S \setminus \bigcup_{i=1}^k S_i\right| \geq q - k \left(\frac{q}{2k}\right) = \frac{q}{2},$$

as desired. Since $|F_{q,d}| = q^{d+1}$, the bound on N follows. \square

This scheme, when combined with a trivial repetition outer code, yields a deterministic broadcast exclusion scheme with a blowup in communication of $O(k^2 \log^2 N)$.

As Theorem 3 shows, the value of N is sub-optimal. In Section 6, we show how to improve this construction using the more powerful AG codes.

The outer code. As in Section 4, we can use an outer code to significantly reduce the communication blowup. The inner code construction above shows that to support N users, we need to have $m = O(k^2 \log^2 N)$ total keys and total messages sent, with each user receiving $q = O(kn)$ keys. Moreover, the construction guarantees that even with k users blacklisted, every good user will be able to decrypt at least $q/2$ of the messages.

Let $q' \geq q^2$. We can use an $[m, q/2, q/2]_{q'}$ constructive erasure code as the outer code, such as the polynomial construction given earlier. Thus, if we assume the message is at least $q \log q = \Omega(k \log kn \log n)$ bits long, the communication blowup is just $2m/q = 2q$ which is $O(kn)$.

6 Construction based on AG codes

As we can see, the fundamental limitation of the polynomial-based scheme given in Section 5 was that the size of the sets had to be much smaller than the total number of possible elements. This was because the number of points where we could evaluate the polynomials was bounded by the field size. It is precisely this problem that algebraic-geometric (Goppa) Codes were designed to address. We will show that this property

allows one to construct very good set systems using them as well. There are a few Goppa codes in existence today that would address our problems [26, 18, 16], we chose for sake of clarity (following [20]) to focus on the family of codes given by the explicit and fairly simply Garcia-Stichtenoth function fields [16].

Garcia and Stichtenoth [16] give an explicit construction of function fields F_n extending \mathbb{F}_{q^2} . The main theorem of the paper is that F_n has at least $s(n) + 1 \stackrel{\text{def}}{=} q^{n-1}(q^2 - 1) + 1$ places of degree one and genus at most $g(n) \stackrel{\text{def}}{=} q^{n-1}(q + 1)$. Thus, the asymptotic ratio of the number of places of degree one to the genus is $q - 1$, attaining the Drinfeld-Vladut upper bound. Indeed, such function fields were known to exist by the celebrated Tsfasman-Vladut-Zink theorem [26], and Manin and Vladut [18] showed that the corresponding Goppa codes are constructible in polynomial time. The function fields underlying the work of [26], however, are far from being explicit. The main contribution of [16]— can be seen as giving a simple construction of function fields attaining the Drinfeld-Vladut bound. Indeed, we can describe the construction here: Let $F_1 = \mathbb{F}_{q^2}(x_1)$, the rational function field over one variable x_1 . For $n \geq 2$, let $F_n = F_{n-1}(z_n)$, where z_n satisfies the equation $z_n^q + z_n = x_{n-1}^{q+1}$ and for $n \geq 3$, $x_{n-1} \stackrel{\text{def}}{=} z_{n-1}/x_{n-2}$.

By Property 2, for every $n \geq 3$ and $\alpha \leq s(n)$ the Garcia-Stichtenoth construction yields an $[s(n), k, d]_{q^2}$ -linear code $C = C(\alpha Q; P_1, \dots, P_s(n))$, where $d \geq s(n) - \alpha$, $k \geq \alpha - g(n) + 1$ and $Q, P_1, \dots, P_s(n) \in \mathbb{P}(F_n)$, i.e., are distinct places of degree one in F_n .

The inner code. Let q be chosen so that $k = \lfloor q/6 \rfloor$, and let $n \geq 3$. Let $Q, P_1, \dots, P_{s(n)} \in \mathbb{P}(F_n)$ be distinct places of degree one. Let the universe $U = \{P_1, \dots, P_{s(n)}\} \times \mathbb{F}_{q^2}$, and we denote the size of U by $m = s(n)q^2$. Finally, let $\alpha = g(n) + m/q^3$. Then we define the set system $\mathcal{S} = \{S_f \mid f \in L(\alpha Q)\}$, where $S_f = \{\langle P_1, f(P_1) \rangle, \dots, \langle P_{s(n)}, f(P_{s(n)}) \rangle\} \subset U$. We denote the size of \mathcal{S} by N .

Theorem 4. *The set system \mathcal{S} is $(k, 1/2)$ -cover free with*

$$N = \exp\left(O\left(\frac{m \log k}{k^3}\right)\right).$$

Proof. The construction of \mathcal{S} above obviously corresponds to the AG code $C = C(\alpha Q; P_1, \dots, P_{s(n)})$. Thus, because C is a code with minimum distance at least $s(n) - m/q^3 - g(n)$, it follows that $|S_i \cap S_j| \leq m/q^3 + g(n)$ for any distinct $S_i, S_j \in \mathcal{S}$. So, for any $S, S_1, \dots, S_k \in \mathcal{S}$,

$$\begin{aligned} \left|S \setminus \bigcup_{i=1}^k S_i\right| &\geq s(n) - k \left(\frac{m}{q^3} + g(n)\right) \\ &= q^{n+1} - q^{n-1} - \frac{2q^{n+1} - q^n - q^{n-1}}{6} \\ &\geq \frac{q^{n+1} - q^{n-1}}{2} + \frac{q^{n+1} - q^n}{6} \\ &\geq \frac{s(n)}{2} \end{aligned}$$

as desired. Since $N = |L(\alpha Q)| = (q^2)^{\alpha-g(n)} = (q^2)^{m/q^3}$, the bound on N follows. \square

This scheme, combined with the trivial repetition outer code, yields a deterministic broadcast exclusion scheme with a blowup in communication of $O(k^3 n)$.

The outer code. We will now show how to use an appropriate outer code to reduce the communication blowup to just $O(k^2)$, eliminating dependence on N and almost matching the randomized construction of Section 4. The inner code construction above shows that to support N users, we need to have $m = O(k^3 n)$ total keys and total messages sent, with each user receiving $m/q^2 = O(kn)$ keys. Moreover, the construction guarantees that even with k users blacklisted, every good user will be able to decrypt at least $m/2q^2$ of the messages.

Let $q' \geq m$ be a prime power. We can use an $[m, m/2q^2, m/2q^2]_{q'}$ constructive erasure code as the outer code, such as the polynomial construction given earlier. Thus, if we assume the message is at least $(m/q^2) \cdot \log q' = \Omega(k \log k n \log n)$ bits long, the communication blowup is just $2q^2$ which is $O(k^2)$, independent of the number of users N .

References

1. S. Berkovits. How to broadcast a secret. *Proc. of EUROCRYPT*, Springer LNCS 547:535–541, 1991.
2. E. R. Berlekamp. Bounded distance + 1 soft decision Reed-Solomon coding. *IEEE Trans. on Information Theory*, 42:704–720, 1996.
3. C. Blundo, L. F. Mattos, and D. R. Stinson. Generalized Beimel-Chor schemes for broadcast encryption and interactive key distribution. *Theoretical Computer Science*, 200(1-2):313–334, 1998.
4. R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. *Proc. IEEE INFOCOM'99*.
5. R. Canetti, T. Malkin, and K. Nissim. Efficient communication-storage tradeoffs for multicast encryption. *Proc. EUROCRYPT 99*, to appear.
6. B. Chor, A. Fiat, and M. Naor. Tracing traitors. *Proc. CRYPTO*, pp. 257–270, 1994.
7. H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–509, 1952.
8. I. J. Cox, J. Kilian, T. Leighton, and T. Shamoon. A secure, robust watermark for multimedia. *Workshop on Information Hiding*. Newton Institute, University of Cambridge, 1996.
9. I. J. Cox and M. L. Miller. A review of watermarking and the importance of perceptual modeling. *Proc. Electronic Imaging*, 1997.
10. C. Dwork. Copyright? Protection? *The Mathematics of Coding, Extraction, and Distribution: IMA Volumes in Mathematics and its Applications*, 107, 1996.
11. C. Dwork, J. Lotspiech, and M. Naor. DigitalSignets: Self-enforcing protection of digital information. *Proc. 28th ACM Symposium on Theory of Computing*, pp. 489–498, 1996.
12. P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51:79–89, 1985.
13. A. Fiat and M. Naor. Broadcast encryption. *Proc. CRYPTO*, pp. 480–491, 1993.
14. IEEE. *1394 Specifications*. Available from customer.service@ieee.org.

15. <http://firewire.org/>.
16. A. Garcia and H. Stichtenoth. A tower of Artin-Schreier extensions of function fields attaining the Drinfeld-Vladut bound. *Invent. Math.*, 121:211–222, 1995.
17. V. D. Goppa. Codes on algebraic curves. *Soviet Math. Dokl.*, 24:170–172, 1981.
18. Y. I. Manin and S. G. Vladut. Linear codes and modular curves. *J. Soviet Math.*, 30:2611–2643, 1985.
19. S. Mitra. Iolus: A framework for scalable secure multicasting. *Proc. ACM SIGCOMM Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication*, pp. 277–288, 1997.
20. M. A. Shokrollahi and H. Wasserman. Decoding algebraic-geometric codes beyond the error-correction bound. *Proc. 30th ACM Symposium on Theory of Computing*, pp. 241–248, 1998.
21. H. Stichtenoth. *Algebraic Function Fields and Codes*. Universitext, Springer-Verlag, 1993.
22. D. R. Stinson and T. van Trung. Some new results on key distribution patterns and broadcast encryption. *Designs, Codes, and Cryptography*, to appear.
23. D. R. Stinson and R. Wei. An application of ramp schemes to broadcast encryption. Manuscript.
24. D. R. Stinson and R. Wei. Key preassigned traceability schemes for broadcast encryption. Manuscript.
25. L. Trevisan. Constructions of near-optimal extractors using pseudorandom generators. *31st ACM Symposium on Theory of Computing*, pp. 141–148, 1999.
26. M. A. Tsfasman, S. G. Vladut, and Th. Zink. Modular curves, Shimura curves, and Goppa codes better than the Varshamov-Gilbert bound. *Math. Nachrichten*, 109:21–28, 1982.
27. D. M. Wallner, E. J. Harder, and R. C. Agee. Key management for multicast: Issues and architectures. <ftp://ietf.org/internet-drafts/draft-wallner-key-arch-01.txt>