

Multicast Authentication in Fully Adversarial Networks

ANNA LYSYANSKAYA ROBERTO TAMASSIA NIKOS TRIANDOPOULOS
anna@cs.brown.edu rt@cs.brown.edu nikos@cs.brown.edu

Department of Computer Science
Brown University
Providence, RI, 02912-1910

An early version of the paper [22] was presented in 2004 IEEE Symposium on Security and Privacy.

Abstract

We consider a general version of an important and well studied data authentication problem, the multicast authentication problem, where – in this general version – we wish to authenticate a packet stream transmitted over a *fully adversarial network*. By fully adversarial network, we mean that the network is controlled by an adversary who may perform any possible type of attack: he may drop or modify chosen packets, rearrange the order of the packets in any way, and inject new, random or of specific structure, packets into the stream. In contrast, prior work on the multicast authentication problem has either focused on a less powerful adversarial model for the network, where random rather than adversarially-selected packets may be dropped or altered or no additional packets may be injected into the stream, or has examined considerably less general settings, where specific assumptions about the network limit its adversarial behavior significantly.

In this paper, we provide the first formal definition of multicast authentication in a network controlled by an adversary as above and define the notion of correctness and security with which any authentication scheme should comply. We model the ability of the network to modify a stream of n packets transmitted by the sender with two parameters: the *survival rate* α , $0 < \alpha \leq 1$, denoting the minimum fraction of the packets that are guaranteed to reach any particular receiver unmodified (i.e., at least αn packets in the received stream will be valid), and the *flood rate* β , $\beta \geq 1$, indicating the maximum factor by which the size of the received stream at any particular receiver may exceed the size of the transmitted stream (i.e., at most βn packets will be in the received stream). Such a network, which we call an (α, β) -network, forms a very general version for the multicast authentication problem. Both of the rates express two intrinsic characteristics of the problem, which avoid some degenerated and extreme cases, rather than some limitations of the adversary or weakness of our model.

We describe the first efficient authentication scheme for an (α, β) -network. This scheme gives almost the same security guarantees as if each packet were individually signed, but requires only one signature operation for the entire stream and adds to each transmitted packet only a small amount of authentication information, proportional to β/α^2 . We use a novel combination of error-correcting codes with standard cryptographic primitives (hashing and digital signatures). We prove the security and correctness of our scheme and analyze its performance in terms of computational effort at the sender and receiver and communication overhead. We also discuss specific design and implementation choices and we compare our scheme with previously proposed approaches.

1 Introduction

The authentication of multicast transmissions is one central problem in network security. Data transmission in a multicast setting involves a sender – the source of the data – sending data to a large set of receivers, where data is transmitted as a stream of packets over an underlying network. Considering an honest receiver, the authentication problem that arises is the verification of the received stream being authentic, that is, being (part of) the original stream sent by the data source. A large and still growing set of Internet’s applications that are based on multicast data transmissions justify the importance of this problem. Distributed data management and peer-to-peer systems, digital broadcasts, public subscription systems usually include multicast of information. Also numerous multicast group applications involve large scale dissemination of high volumes of data from one source to many users. Since data is critical for the reliability – and the existence, in fact – of the target application, data authenticity is a necessity, not a feature of quality of service.

Technically, the problem of multicast authentication is a challenging one and has attracted a great amount of interest during the last years. One of the distinguishing properties of the multicast data transmission is certainly the fact that for the majority of the applications packet losses are tolerated and, consequently, data transmission need not be reliable. Indeed, IP multicast is implemented with a best-effort delivery mechanism over the UDP transport protocol and packets can be lost due to failures. Thus, in principle, the stream that reaches a receiver may differ from the transmitted one. As a result, any authentication scheme for multicast streams should verify as many as possible of the received packets without assuming the availability of the entire original stream. In addition, it should resist against any type of attack by an adversary, even when the adversary completely controls the underlying network. Although what happens in practice is that packets get lost because of errors, what authentication dictates is protection against an adversarial network behavior. Indeed, the main characteristic of the problem is by definition the existence of an entity acting maliciously in between the sender and an honest receiver. In practice, the role of the adversary may play any of the parts of the underlying network, such as ISPs, routers or malicious users. Finally, an authentication scheme should be efficient, scalable, lightweight with regard to the final application and most importantly as general as possible with regard to adversarial network behavior. That is, it is essential that the adversary is modelled as an entity of great power and that no assumptions about the network that limit adversary’s acting exist.

Therefore, in the *multicast authentication problem*, we wish to authenticate a packet stream transmitted over a network that may adversarially drop packets, arbitrarily rearrange the order of the packets, and inject new packets into the stream. As explained, the authentication mechanism should be as general as possible and should not depend on any specific assumptions about the underlying network. Although this problem has been extensively studied, no formal definition has been given for it for this general version. Prior work on the subject has focused on a network model where either all the received packets are valid (authentic) or packets are lost according to some predefined random patterns (e.g., [13, 23, 29, 34]) or no packet injections occur (e.g., [24, 25]) or relatively strong to meet conditions are assumed about the network behavior (e.g., [28, 29]). Thus, most of the previously proposed schemes rely on less general network models, tolerate only *erroneous* network behavior and are not resilient against an *adversarial* behavior of the network.

Of course, if each packet were signed by the sender, then the only damage the adversarial network could inflict is packet loss, as the receiver would simply reject packets whose signature is not verified. However, this simple “sign-all” solution is undesirable because of the repeated use by the sender of the critical and computationally expensive sign primitive for each transmitted packet and the heavy communication overhead caused by the addition of a signature to each packet.

Additionally, this solution suffers by a simple denial-of-service attack at the receiver; one signature verification must be performed for each received packet, valid or not.

In this paper, we formally define a general model for multicast authentication where an adversary can perform various attacks on the transmitted streams. In this model, two parameters of the network, the *survival rate* and the *flood rate*, characterize the power of the adversary. Using this network model, we formally define the authentication problem for multicast transmissions and describe the notions of correctness and security that any authentication scheme should respect. We describe an efficient authentication scheme for this model that gives almost the same security guarantees as if each packet were individually signed, but requires only one signature operation for the entire stream and adds only a constant size authentication overhead per packet. Our technique uses a novel combination of Reed-Solomon error-correcting codes with standard cryptographic primitives, such as collision-resistant hashing and digital signatures. The use of error correcting codes for multicast authentication was inspired by the previous use of erasure codes, as in [24, 25], for this problem.

In the rest of this section, we introduce our model, summarize our contributions and present previous work on multicast authentication. The organization of the rest of the paper is as follows. The cryptographic primitives and error-correcting codes used in this paper are reviewed in Section 2. In Section 3, we describe in detail our adversarial network model and multicast authentication framework. Section 4 describes the construction of our multicast authentication scheme and gives proofs of correctness and security. In Section 5 we analyze the performance of our scheme and compare it with various other proposed schemes in terms of security assumptions, underlying network model, resilience to packet loss and injection, computational effort at the sender and receiver, and communication overhead. We conclude in Section 6.

1.1 Model and Contributions

We consider the problem of authenticating a stream of packets transmitted over a *fully adversarial* network. Namely, the network is controlled by an adversary who can destroy packets of her choice, arbitrarily rearrange the order of the packets, and inject new, arbitrarily constructed, packets. We limit the power of the adversary to modify a stream of n packets transmitted by the sender by introducing two parameters of the network, the *survival rate* α , $0 < \alpha \leq 1$, and the *flood rate* β , $\beta \geq 1$, which are assumed to be constants. A network with these two parameters, which we call an (α, β) -network, guarantees that despite the presence of the adversary, at least αn packets in the received stream are valid and the received stream contains at most βn packets.

The model is formally described in Section 3. For now, we note that these rates of the network are only used to model the adversary's action; not to limit it. An authentication scheme should operate correctly and securely for any values of α and β , even non constants. We briefly justify the introduction of the survival and flood rates – and in essence the choice of them to be constants – with the following observations. If too many packets are dropped or corrupted by the adversary, then the main problem is the *loss of data*, as the small number of valid packets received may be useless even if authenticated. On the other hand, if the adversary can inject a very large number of packets, then we have a *denial-of-service attack*. In both cases, as α gets smaller or β gets larger, the authentication problem degenerates to data loss and denial-of-service attack. No need for authentication really exists in both of these extreme cases.

The contributions of our work can be summarized as follows:

- We provide a formal definition of multicast authentication over an (α, β) -network, where arbitrary packets are lost, injected, and rearranged, subject to a given survival rate α and

flood rate β . We also give the requirements for an authentication scheme to be *correct* and *secure*.

- We present the first efficient and scalable multicast authentication scheme for an (α, β) -network. Our scheme is based on digital signatures, cryptographic hash functions and Reed-Solomon error-correcting codes. This last feature of our scheme provides a new interesting connection between coding theory and security.
- We prove the correctness and security of our scheme, analyze its performance in terms of various cost parameters, discuss design and implementation choices, and compare it with previous approaches. In particular, we show that our scheme adds to each transmitted packet only a small amount of authentication information, proportional to β/α^2 , and that all the valid packets received are recognized, while all the invalid packets are rejected.

The only prior approaches that provide security in our adversarial model is (i) the inefficient “sign every packet” solution, which consists of either signing each packet individually or using a Merkle hash tree [36] by Wong and Lam and (ii) a recently proposed scheme [17] by Karlof *et al.* that uses signature dispersal and a Merkle hash tree. The trivial solution of signing each packet individually is not viable due to heavy computational operations at both the sender and the receiver, but also because secret-key operations are expensive in terms of the security architecture as well. Our scheme, by amortizing one signature per a stream of size n , suffers from no such problem. On the other hand, the Merkle-tree-based authentication schemes [17, 36] have the drawback that the per-packet communication overhead grows logarithmically with the number of packets sent. Indeed, each packet of a stream of size n carries authentication information of size $O(\log n)$. In contrast, our scheme achieves per-packet communication overhead independent of n and, thus, more efficiency and scalability.

1.2 Prior and Related Work

Previous work on multicast authentication considers both unconditionally secure and computationally secure authentication. Approaches based on the information theoretic model (see, e.g., [8, 33]) tend to be less practical. In the rest of this section, we overview approaches that use computationally secure authentication. We do so, by appropriately categorizing previous work according to the underlying authentication technique used.

MAC-Based Approaches. Various approaches use message authentication codes (MACs) and secret-key cryptography. The trivial solution here is having the multicast group members (i.e., all the receivers) sharing a secret key and including a MAC into every packet sent, but this scheme is not secure, as any user can spoof packets. In another MAC-based trivial solution, each receiver has her own secret key and the sender possesses all such keys. To authenticate a stream, the sender adds to each packet a MAC for every receiver. This approach is not scalable because of the high communication cost.

Canetti *et al.* [4] describe a MAC-based scheme that is secure with high probability against any coalition of w corrupted users and where $O(w)$ MACs are appended to each packet. This scheme is not fully scalable due to its communication overhead. Perrig *et al.* [28, 29] present another MAC-based scheme, TESLA, where a MAC is appended to every packet and the key of the MAC is provided in some subsequent packet. To tolerate packet losses, the keys are generated by means of a hash chain. This approach has low communication overhead. However, it requires time synchronization between the parties. Two MAC-based schemes that make explicit use of the

topology of a multicast tree are proposed in [37] by Xu and Sandhu. Both schemes are similar in concept to [28] and take denial-of-service and access control into consideration (namely, a corrupted packet is filtered out as soon as possible in the multicast tree and only legitimate group subscribers can authenticate the multicast packets). Both schemes assume the existence of secure and trusted routers at the nodes of the tree. In addition, the first scheme uses clock synchronization, whereas the second scheme relies on the existence of secure channels between the source and each of the receivers.

Boneh et al. [1] generalize MACs to a multicast setting by defining a new primitive for multicast authentication called Multicast MAC (MMAC). A MMAC is a triplet of algorithms ($key-gen, mac-gen, mac-ver$) where: $key-gen$ produces secret key sk for the sender and secret keys rs_1, \dots, rs_n for the receivers; $mac-gen(M, sk)$ computes an authentication tag τ for message M and $mac-ver(M, \tau, rk_i)$ returns a “yes” or “no”, checking whether τ is an authentication tag of message M . A MMAC must satisfy certain correctness and security constraints. In their work, Boneh *et al.* study the existence of efficient MMACs, i.e., MMACs with short tag τ . They show that any MMAC scheme can be transformed into a digital signature scheme of almost the same efficiency. Thus, any multicast authentication scheme not relying on additional assumptions on the network (such as synchronization, trusted routers, or secure channels) may as well use a signature scheme, which brings us to *signature amortization*. The result is also extended in the case where the adversary possesses a limited number of the receivers’ keys and a lower bound on the length of the authentication tag is derived. The construction by Canetti *et al.* in [4] meets this bound, whereas TESLA by Perrig *et al.* in [28, 29], using time synchronization and one digital signature for bootstrapping, is not an exact MMAC scheme.

In view of the previous result, research efforts also focused on building faster signature schemes for signing every packet separately. Work on this direction includes (i) the use of one-time digital signatures by Gennaro and Rohatgi in [10] for on-line data stream multicast transmission but only in reliable in terms of packet delivery communication channels, (ii) the use of k -time digital signatures by Rohatgi in [32] to speed up the signing rate with relatively short signature sizes and (iii) Perrig’s BiBa one-time broadcast protocol in [27].

Signature Amortization. Many approaches use the technique of signature amortization, where a single digital signature is used for the authentication of multiple packets. A first scheme that uses signature amortization over a hash chain appears by Gennaro and Rohatgi in [10]. Each packet p_i is augmented with authentication information a_i , which is recursively defined as the hash of $p_{i+1} \circ a_{i+1}$ (\circ denotes concatenation). Also, the augmented first packet $p_1 \circ a_1$ is digitally signed. This scheme has constant authentication overhead per packet but does not tolerate packets losses. In [36], a Merkle hash tree is used by Wong and Lam to amortize a signature over n packets. Namely, a hash tree is built on top of the hashes of the packets and the root hash value is digitally signed. Each packet is augmented with authentication information that consists of the signed root hash and the hashes of the siblings of the nodes on the path between the root and the leaf associated with the packet. The scheme tolerates packet losses but has logarithmic communication overhead per packet. In contrast, our approach, which also uses signature amortization, has *constant* per-packet communication overhead.

Graph-Based Authentication. Graph-based authentication [13, 23, 29, 34] generalizes the idea of amortizing a signature over a hash chain in such a way as to tolerate packet losses. A single-sink directed acyclic graph (DAG) G is defined, where each vertex corresponds to a packet. A directed edge from packet p_i to packet p_j indicates that the authentication information a_j of packet p_j

includes the hash of $p_i \circ a_i$. Also, the augmented packet $p_1 \circ a_1$ of the sink of the DAG is digitally signed. The validation of packets proceeds backward along the edges of the graph. Namely, if packet p_j has been validated and edge (p_i, p_j) exists in G , then the validity of packet p_i can be determined using the authentication information a_j of p_j . Graph-based authentication schemes offer probabilistic security guarantees provided packet losses occur randomly (i.e., they are not adversarially selected). In particular, they require that the signature packet will reach the receiver intact. Two packet loss patterns have been studied: the uniform model, where each packet is lost with a fixed probability and independently of other packets being lost, and the bursty model, where a packet is lost with a fixed probability and then a given number of successive packets are also lost.

In [29], Perrig *et al.* choose G to be an augmented-chain graph, consisting of a path plus additional edges that connect vertices at various distances. Golle and Modadugu in [13] propose the use of another augmented-chain graph which is designed specifically to tolerate bursty packet losses. Random graphs and a new scheme that is resilient to multiple bursty losses are studied by Minner and Staddon in [23]. Finally, in [34], expander graphs are used by Song *et al.* The efficiency of graph-based authentication schemes is analyzed in [5] by Chan and experimentally studied in [7] by Cucinotta *et al.*

Erasure Codes. Park *et al.* [25] and Pannetrat and Molva [24] employ the use of erasure codes (e.g., [20, 21, 30]) for multicast authentication to tolerate adversarially-chosen packet losses and disperse one signature over a group of packets. In particular, information sufficient – if reconstructed at the receiver – to authenticate the received packets is encoded using an erasure code so that delivery of a constant fraction of packets guarantee the successful decoding of this information. The constructions are efficient in terms of communication cost and similar in principle. The two schemes only differ in that in [24], encoding is performed twice to reduce the size of the authentication information. The idea here is that a significant portion of the encoded information reaches the receiver for free through the valid packets. Both schemes are, however, vulnerable to a very simple attack: a single injected packet can compromise the correctness of the decoding procedure at the receiver.

In [26], Park *et al.* identify a special case of the problem, where packets are altered, as a denial-of-service attack. They suggest the use of distributed fingerprints [18] (which, in turn, are based on error correcting encoding) to tolerate a small number of symbol modifications (through packet alterations) of the erasure-encoded information. However, the proposed scheme lacks efficiency, since distributed fingerprints are used on top of the erasure encoding and, more importantly, collapses for a specific packet-injection type of attack, where more than one packets (symbols) claim to be a specific packet (symbol). Note that our adversarial model includes this type of attack: many injected packets may pretend to be a specific original packet.

Recently, Karlof *et al.* [17] propose a solution to the packet injection problem that the erasure-based schemes above have. They refer to this problem as *packet pollution* and they introduce the notion of *distillation codes*, codes that tolerate both erasures and pollution of symbols. They first realize a distillation code by using an erasure code and an one-way accumulator. Symbols are erasure encoded as usual, but also appended by a witness of set inclusion in the set of pre-encoded symbols. At the receiver, the idea is to partition the received symbols into groups of symbols so that received symbols of the same group belong in the same group of transmitted (not necessarily authentic) pre-encoded packets. Partitioning is feasible relying on the set inclusion properties of the one-way accumulator and, consequently, each group can be decoded and then examined to be authentic. Built on the scheme [26] by Park *et al.* and using Merkle’s hash tree as an one-way accumulator, this approach leads to an multicast authentication scheme, where both packet

injections and packet erasures are tolerated. However, this scheme achieves *not constant* per-packet communication overhead, since each packet carries information of size $O(\log n)$ as witness (because of the use of Merkle tree as an one-way accumulator), and suffers from the cumbersome decoding procedure with respect to the processing of packets at the receiver: the partition operation adds a considerable amount of computation and hashing and a considerable number of erasure decodings.

Recently Gunter *et al.* [14] use erasure codes to tolerate packet injections but only random packet losses. In particular, in a less general network model, the adversary shares the transmission channel with the sender and is allowed to inject packets up to a certain transmission rate. The scheme involves the use of three different streams: (i) the data stream where packets only contain data, (ii) the hash and parity stream, which consists respectively of hash packets and encoded packets and (iii) the signature stream. Signatures are verified selectively and injected packets are filtered out by processing all received packets in an exhausting way. The scheme achieves only probabilistic guarantees about the authentication of received packets. Obviously, no adversarially chosen packets are tolerated.

Finally, in a recent related work [19], Krohn *et al.* use erasure encoding techniques in combination with *homomorphic hashing* for the on-the-fly verification of erasure-encoded blocks. Applications involve peer-to-peer content distribution but also multicast transfers. However, their scheme relies on the following strong assumption: the receiver knows in advance the cryptographic hashes of the transmitted shares.

2 Preliminaries

In this section, we introduce some notation and define the cryptographic and coding primitives that we use in our construction.

2.1 Notation

Let A be an algorithm. By $A(\cdot)$ we denote that A has one input (resp., by $A(\cdot, \dots, \cdot)$ we denote that A has several inputs). By $y \leftarrow A(x)$, we denote that y was obtained by running A on input x . If A is deterministic, then this y is unique; if A is probabilistic, then y is a random variable. If S is a finite set, then $y \leftarrow S$ denotes that y was chosen from S uniformly at random. By $y \in A(x)$ we mean that the probability that y is output by $A(x)$ is positive.

By $A^O(\cdot)$, we denote an algorithm that makes queries to an oracle O . I.e., this algorithm (Turing machine) will have an additional (read/write-once) query tape, on which it will write its queries in binary; once it is done writing a query, it inserts a special symbol “#”. By external means, once the symbol “#” appears on the query tape, oracle O is invoked and its answer appears on the query tape adjacent to the “#” symbol. By $Q = Q(A^O(x)) \leftarrow A^O(x)$ we denote the contents of the query tape once A terminates, with oracle O and input x . By $(q, a) \in Q$ we denote the event that q was a query issued by A , and a was the answer received from oracle O .

Let b be a boolean function. By $(y \leftarrow A(x) : b(y))$, we denote the event that $b(y)$ is TRUE after y was generated by running A on input x . The statement $\Pr[\{x_i \leftarrow A_i(y_i)\}_{1 \leq i \leq n} : b(x_n)] = \alpha$ means that the probability that $b(x_n)$ is TRUE after the value x_n was obtained by running algorithms A_1, \dots, A_n on inputs y_1, \dots, y_n , is α , where the probability is over the random choices of the probabilistic algorithms involved.

2.2 Cryptographic Primitives

The following definition is due to Goldwasser, Micali, and Rivest [12], and has become the standard definition of security for signature schemes. Schemes that satisfy it are also known as signature schemes secure against *adaptive chosen-message attack*.

Definition 1 (Signature scheme). Probabilistic polynomial-time algorithms $(G(\cdot), \text{Sign}_{(\cdot)}(\cdot), \text{Verify}_{(\cdot)}(\cdot, \cdot))$, where G is the key generation algorithm, Sign is the signature algorithm, and Verify the verification algorithm, constitute a digital signature scheme for a family (indexed by the public key PK) of message spaces $\mathcal{M}_{(\cdot)}$ if the following two hold:

Correctness If a message m is in the message space for a given public key PK , and SK is the corresponding secret key, then the output of $\text{Sign}_{SK}(m)$ will always be accepted by the verification algorithm Verify_{PK} . More formally, for all values m and k :

$$\Pr[(PK, SK) \leftarrow G(1^k); \sigma \leftarrow \text{Sign}_{SK}(m) : m \leftarrow \mathcal{M}_{PK} \wedge \neg \text{Verify}_{PK}(m, \sigma)] = 0.$$

Security Even if an adversary has oracle access to the signing algorithm that provides signatures on messages of the adversary's choice, the adversary cannot create a valid signature on a message not explicitly queried. More formally, for all families of probabilistic polynomial-time oracle Turing machines $\{A_k^{(\cdot)}\}$, there exists a negligible function¹ $\nu(k)$ such that

$$\Pr[(PK, SK) \leftarrow G(1^k); (Q, m, \sigma) \leftarrow A_k^{\text{Sign}_{SK}(\cdot)}(1^k) : \text{Verify}_{PK}(m, \sigma) = 1 \wedge \neg(\exists \sigma' \mid (m, \sigma') \in Q)] = \nu(k).$$

For completeness, we give a standard definition of a family of collision-resistant hash functions.

Definition 2 (Collision-resistant Hash Function). Let \mathcal{H} be a probabilistic polynomial-time algorithm that, on input 1^k , outputs an algorithm $H : \{0, 1\}^* \mapsto \{0, 1\}^k$. Then \mathcal{H} defines a family of collision-resistant hash functions if:

Efficiency For all $H \in \mathcal{H}(1^k)$, for all $x \in \{0, 1\}^*$, it takes polynomial time in $k + |x|$ to compute $H(x)$.

Collision-resistance For all families of probabilistic polynomial-time Turing machines $\{A_k\}$, there exists a negligible function $\nu(k)$ such that

$$\Pr[H \leftarrow \mathcal{H}(1^k); (x_1, x_2) \leftarrow A_k(H) : x_1 \neq x_2 \wedge H(x_1) = H(x_2)] = \nu(k).$$

2.3 Error Correcting Codes

Error correcting codes allow recovering a message that is transmitted over a noisy channel. Let $q \geq 2$ be the size of alphabet $[q] = \{1, 2, \dots, q\}$. An *error correcting code* $[n, k]_q$, $k < n$, is a function² $C : [q]^k \rightarrow [q]^n$ that takes as input a k -length message x over $[q]$ and outputs a longer n -length codeword $C(x)$ over the same alphabet. That is, an error correcting code processes k characters (symbols) in $[q]$ and adds redundancy to form n characters. If only redundancy is added by a code C such that the k first symbols of $C(x)$ form word x , then the code is called *systematic*.

¹A function $\nu : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every positive polynomial $p(\cdot)$ and for sufficiently large k , $\nu(k) < \frac{1}{p(k)}$.

²There are more than one possible definitions. Here we choose to define an error correcting code as a function.

The processing and the added redundancy help correcting up to e errors of the codeword $C(x)$; that is, given a received word $y \in [q]^n$ such that y and $C(x)$ differ in at most e characters of the alphabet, one can unambiguously correct (decode) y to $C(x)$. The value e depends on the exact choice of the code C . In particular, for *unambiguous* decoding, e is always bounded from above by $d/2$, i.e., $e < d/2$, where d is the *diameter* of the code, defined as follows. The *distance* of two codewords of C is the number of positions where their symbols differ and the distance of code C is the minimum distance of two codewords of C , measured over all pairs of codewords.

List-decoding allows to correct a number of errors that is beyond the bound above. Handling even more errors comes at a price, though. In fact, one can *ambiguously* correct beyond this bound. That is, given that $C(x)$ is received as word $y \in [q]^n$ and $C(x)$ and y differ in at most e positions with $e \geq d/2$, list-decoding provides a list of candidate initial messages in $[q]^k$, such that x belongs in this list. In general, as e grows, the size of the output list grows as well and, typically, we put an upper bound on e so that list-decoding is efficient (i.e., is performed in polynomial time) and the list size is reasonable to work with.

Reed-Solomon codes [31] are a family of codes that are based on properties of univariate polynomials over finite fields. An $[n, k+1]_q$ Reed-Solomon code, $k < n \leq q$, has diameter $d = n - k$ and, as long as the number of errors e is at most $(d - 1)/2 = (n - k - 1)/2$, one can unambiguously decode in quadratic time [35]. If $e > (n - k - 1)/2$, then list-decoding is considered to be feasible [11] (i.e., it can be performed in polynomial time) as long as

$$e \leq n - \sqrt{nk}.$$

We next define the $[n, k+1]_q$ Reed-Solomon code, where n , k and q are positive integers and parameters of the code, with $k < n \leq q$,³ and list decoding is considered. We present here a slightly modified definition of Reed-Solomon codes than the one commonly used in the literature, so that, by definition, Reed-Solomon codes considered in this paper are systematic. This property is not necessary for the correctness of our scheme, but offers an extra desired property in our construction, discussed in Section 5 and requires no extra computational effort at the encoder.

Definition 3 (Reed-Solomon code). A $[n, k+1]_q$ Reed-Solomon code consists of the following components:

Alphabet The alphabet is a finite field \mathbb{F}_q of size $q \geq n$, with q being a prime power.

Encoder The code is a function $C : \mathbb{F}_q^{k+1} \rightarrow \mathbb{F}_q^n$, where $n > k$. In more detail, the encoder:

1. takes as input the parameters n and k , and $k+1$ points (i, y_i) , $i \in \mathbb{F}_q$, $y_i \in \mathbb{F}_q$, $1 \leq i \leq k+1$,
2. finds a unique univariate polynomial $p \in \mathbb{F}_q[x]$ over elements of \mathbb{F}_q and of degree at most k , such that $p(i) = y_i$, $1 \leq i \leq k+1$, and
3. outputs points $(i, p(i))$, for $1 \leq i \leq n$. The ratio $\frac{k}{n} < 1$ is called the *rate* of the code.⁴

We have that C is a systematic code.

Decoder Let $0 < \epsilon < 1$ be a parameter that controls the performance of the decoder. The decoder takes as input parameters n and k , the maximum number of errors e that may occur, and n

³For simplicity, in this definition we require that $k < n$ rather than $k+1 < n$, allowing thus the extreme case, where, for $k = n - 1$, the code adds no redundancy at all.

⁴For convenience, we note that in fact the rate of an $[n, k+1]_q$ code is $\frac{k+1}{n}$, but for simplicity we adopt the ratio $\frac{k}{n}$ as being the rate.

points (x_i, y_i) , $1 \leq i \leq n$, and list-decodes, i.e., it outputs a list of all univariate polynomials $p \in \mathbb{F}_q[x]$ of degree at most k such that $y_i \neq p(x_i)$ for less than e values of i , $1 \leq i \leq n$.

For a $[n, k + 1]_q$ Reed-Solomon code, we use a decoder that runs in polynomial time and is due to Guruswami and Sudan [15, 16]. The maximum number of errors that can be tolerated by the list-decoding procedure matches the theoretical bound of $n - \sqrt{kn}$. However, the decoding algorithm is prohibitively expensive (e.g., $O(n^{12})$) when the exact bound is met. By reducing this upper bound of the maximum number of errors to

$$n - \sqrt{(1 + \epsilon)kn}, \quad \epsilon \in (0, 1),$$

a quadratic decoding algorithm exists [16]. We refer to this decoder as GS-Decoder. For this decoder, parameter ϵ controls how far away from the bound of $n - \sqrt{kn}$ list-decoding operates (in the worst case). We will use the following result.

Theorem 1 (Guruswami-Sudan). *Consider a $[n, k + 1]_q$ Reed-Solomon code. For any $\epsilon \in (0, 1)$, given n points with at most $e = n - \sqrt{(1 + \epsilon)kn}$ errors, GS-Decoder outputs a list of size $O(\epsilon^{-1} \sqrt{n/k})$ in $O(n^2 \epsilon^{-5} \log^2 q \log^{O(1)} \log q)$ time, performing $O(n^2 \epsilon^{-5} \log q)$ field operations.*

We view ϵ as a parameter of the GS-Decoder, $0 < \epsilon < 1$, and use $\text{GSDecode}_\epsilon(n, k, e, \{(x_i, y_i) | 1 \leq i \leq n\})$ to denote that GS-Decoder runs with parameter ϵ having as input parameters the integers n, k, e and the points (x_i, y_i) , $1 \leq i \leq n$.

In practice, codes with constant expansion are used, where $k = \rho n$ for some constant $\rho < 1$, with ρ being the rate of the code in use. In particular, our construction makes use of a $[n, \rho n + 1]_q$ Reed-Solomon code, with $\rho < 1$, over some large alphabet of size $q = 2^c$ for some constant c . From Theorem 1 we have the following, where by $\tilde{O}(\cdot)$ we denote that some logarithmic factors are omitted.

Corollary 2. *For any $[n, \rho n + 1]_q$ Reed-Solomon code, for any constants $\epsilon \in (0, 1)$ and $\rho < 1$ such that $\sqrt{(1 + \epsilon)\rho} \leq 1$, on an input with at most $e = (1 - \sqrt{(1 + \epsilon)\rho})n$ errors, GS-Decoder outputs a list of $O(1)$ size in $\tilde{O}(n^2)$ time, performing $\tilde{O}(n^2)$ field operations.*

Proof. It follows directly from Theorem 1 for $k = \rho n$, $\rho < 1$ and using the $\tilde{O}(\cdot)$ notation to hide constants and logarithmic on q factors. We need $\sqrt{(1 + \epsilon)\rho} \leq 1$ such that $e \geq 0$. \square

GS-Decoder is based on an algorithm that solves the *polynomial reconstruction* problem: given k, t , and n points $\{(x_i, y_i), 1 \leq i \leq n\}$, where $x_i, y_i \in \mathbb{F}_q$, find a list that contains all univariate polynomials $p \in \mathbb{F}_q[x]$ of degree at most k such that $y_i = p(x_i)$ for at least t values of i , $1 \leq i \leq n$. Parameter t is usually referred as *agreement*. Polynomial reconstruction and Reed-Solomon list-decoding are equivalent problems [15]. In particular, the following corollary is equivalent to Corollary 2.

Corollary 3. *For any constants $\epsilon \in (0, 1)$ and $\rho < 1$ such that $\sqrt{(1 + \epsilon)\rho} \leq 1$, polynomial reconstruction on input $\rho n, t$, and n points in $\mathbb{F}_q \times \mathbb{F}_q$ can be solved in $\tilde{O}(n^2)$ time, provided $t \geq \sqrt{(1 + \epsilon)\rho n}$, where $\tilde{O}(n^2)$ field operations are performed and the output list has $O(1)$ size.*

Proof. It follows immediately from Corollary 2 by considering the equivalence between the two problems. The upper bound e of the number of errors guarantees agreement t of at least $n - e = \sqrt{(1 + \epsilon)\rho n}$. Also, we need $\sqrt{(1 + \epsilon)\rho} \leq 1$ in order to be consistent with the requirement that $t \leq n$. \square

3 Network Model and Multicast Authentication Framework

Considering data transmission in a multicast setting, a *sender*, the source of the data, transmits a data stream over an underlying “best-effort” network. Data packets are received by a large set of *receivers*. Without loss of generality, we focus our attention to one such receiver, such that the two honest parties of the authentication protocol are the sender and the receiver. No guarantees about the delivery of the packets exist in general. Furthermore, the network is an *adversary* of great, yet not unlimited, power. In our model, packets may be *adversarially lost, altered, delayed, or injected*. However, this adversary is not given complete freedom — if it were, then no messages (data) would ever get delivered, and so our task would be hopeless.

Conventionally and without loss of generality, we consider data streams consisting of n packets, that is, at the sender, the data for transmission is arranged in groups of size n . Each group of n packets is identified by a – unique per distinct group – *group identification tag* GID . That is, packets of a group are marked with the corresponding GID . Note that the existence of tag GID adds no new assumption about the transmitted stream. It corresponds to a means for the packets to be grouped together, which in practice can be provided by any network-layer transmission protocol in use. In our framework, the GID is used as an abstract quantity of constant size; in practice, it is a string of some small constant size, e.g., the size of a hash value (20 bytes for SHA-1).

3.1 The (α, β) -Network Model

We model the network as an *adversarial* entity, i.e., an entity that can *simultaneously* inflict *any possible type* of attack to the transmitted data stream. The repertoire of attacks consists of packet *losses, injections, alterations* and *rearrangements*. These modifications of the data stream are adversarially chosen so that the adversary can cause the loss of any selected packets. The ability to tolerate packet losses has been widely considered an important property of multicast authentication schemes [13, 23, 25, 29, 34, 36]. However, only a few previous schemes [17, 25, 36] tolerate *adversarial losses*, i.e., the capability by the adversary to choose which packets are dropped and which survive. The adversarial loss model is the strongest, and also the most realistic one, since it makes the least assumptions on how the traffic is routed.

Also, the adversary can inject packets of random or malicious structure into the stream. This type of network failure has not been studied as widely in the context of multicast authentication. In contrast, we develop robust techniques for dealing with it. In their recent paper [17], Karloff *et al.*, in studying the *pollution attack* by an adversary when an erasure-code based authentication scheme is in use, they in essence consider packets injections as well.

Finally the adversary can arbitrarily modify, delay or rearrange packets. Note that changing a packet corresponds to destroying (losing) it and injecting a new packet.

An adversarial network modelled with the above capabilities in terms of how the adversary is acting is what we call a *fully adversarial network*.

Definition 4 (Fully adversarial network). A *fully adversarial network* is a network that is used for the transmission of a data stream and is *controlled* by an adversary. In particular, the adversary can:

- cause packets of *her choice* to be lost;
- inject packets (either *random* ones or with a *specific malicious structure*); and
- *arbitrarily* alter, delay or rearrange packets.

It is realistic to assume that even if an adversary controls part of the network, there are still some honest routers and at least a fraction of the data packets goes through them. Thus, we expect some *reliability* from the network. Namely, the network will faithfully deliver at least a constant fraction, α , of all the packets of a given stream. This assumption is also justified by the fact that if fewer than a constant fraction of the packets survive, then it is unlikely that meaningful information can be extracted from the surviving packets. Also, in modelling the ability of the adversary to maliciously inject invalid packets, we take the following into consideration: if the adversary injects packets at too high a rate, this will result in a denial-of-service attack. In this case, the receiver’s primary concern is unlikely to be authentication. Thus, we assume that authentication is useful when the stream is expanded by no more than a constant factor β through adversarial packet injections.

Our two assumptions about the power of the adversary to modify a stream of n packets transmitted by the sender are expressed by two *parameters* of the adversarial network: the *survival rate* α and the *flood rate* β . In this paper, both rates are considered to be constants. In a network with survival rate α , $0 < \alpha \leq 1$, if a stream of n packets is sent, at least αn packets of this stream will arrive at the receiver intact. In a network with flood rate β , $\beta \geq 1$, if a stream of n packets is sent, then the received stream will have at most βn packets.

Definition 5 (Network parameters). Consider an adversarial network through which a stream of n packets is transmitted by the sender. Consider any receiver of the data stream. With respect to this receiver:

- The *survival rate* α , $0 < \alpha \leq 1$, is the minimum fraction of the packets that are guaranteed to reach the receiver unmodified. I.e., at least αn packets in the received stream are valid.
- The *flood rate* β , $\beta \geq 1$, is the maximum factor by which the size of the stream that reaches the receiver may exceed the size of the transmitted stream. I.e., at most βn packets are in the received stream.

We claim that the survival and flood rates of a network are reasonable limitations that better model the adversary’s ability to modify the transmitted stream. In particular, as an assumption, the network’s reliability expressed by means of α and β , it does not affect the generality and strength of our model, but rather only denotes some intrinsic characteristics of the authentication problem itself. Indeed, the extreme cases where too few packets survive or too many packets are injected are both degenerate cases of the multicast authentication problem: for no need, at the receiver, for authentication really exists when $\alpha \rightarrow 0$ or $\beta \rightarrow \infty$. At the same time, our scheme presented in Section 4 is *parameterized* by the two rates α and β ; that is, it operates for *any values* of these two network parameters.

A network with the above characteristics in terms of adversarial behavior and reliability is what we call an (α, β) -*network* and is the basis for our multicast authentication framework. Although our discussion focuses on one particular receiver, for generality and completeness, in the following definition we require that the network provides the same amount of reliability (expressed by rates α and β) to any of the receivers⁵.

Definition 6 ((α, β) -network). An (α, β) -*network* is a fully adversarial network with survival rate α and flood rate β with respect to any receiver.

3.2 Authentication Framework

We describe a new multicast authentication framework that is based on the (α, β) -network model. Our definition of a *multicast authentication scheme* essentially mimics the classical definition of

⁵This does not, of course, mean that the same data packets, i.e., the same stream, reaches all the receivers.

security for signatures [12]. This is not surprising since in [1] it is shown that the two problems are equivalent. A signature scheme consists of key generation, signature, and verification algorithms (see Definition 1). Similarly, we have key generation, authentication, and decoding algorithms, specified below. The key generation algorithm is run in advance to produce the private and public keys used by the involved parties, the sender and the receiver. The other two algorithms, authenticator *Auth* and decoder *Decode*, are executed by the sender and the receiver respectively. The sender runs *Auth* to process data packets and create the authenticated packets. The receiver runs *Decode* to decode the received packets and recognize the valid ones.

Key generation The key generation algorithm *KeyGen* is a probabilistic polynomial-time algorithm that takes as input the security parameter 1^k , and outputs the key pair (PK, SK) . We write $(PK, SK) \leftarrow \text{KeyGen}(1^k)$. We assume that the sender knows both the public key PK and the secret key SK and that the receiver knows the public key PK .

Authenticator The authenticator algorithm *Auth* takes as input:

- (SK, PK) : the secret key and the public key.
- GID : the group identification tag of the data stream.
- n : the size of the data stream, i.e., the number of packets that need to be authenticated.
- α : the survival rate that determines what fraction of the packets are guaranteed to reach the receiver intact, $0 < \alpha \leq 1$.
- β : the flood rate that determines an upper bound of the packets that reach the receiver; namely, the received stream consists of at most βn packets that claim to belong to a given GID , $\beta \geq 1$.
- $DP = \{p_1, \dots, p_n\}$: the *data packets*, i.e., the data stream that needs to be authenticated.

The output of the authenticator algorithm is the set AP of *authenticated packets*, with $AP = \{a_1, \dots, a_n\}$. We write: $AP \leftarrow \text{Auth}(SK, PK, GID, n, \alpha, \beta, DP)$.

Decoder The decoder algorithm *Decode* takes as input:

- PK : the public key.
- GID : the group identification tag of the data stream.
- n : the number of the original data packets.
- α : the survival rate.
- β : the flood rate.
- $RP = \{r_1, \dots, r_m\}$: the received packets.

The decoder either *rejects* the input (when less than αn of the received packets are valid, or more than $(\beta - \alpha)n$ packets are injected by the adversary⁶), or produces the *output packets* $OP = \{p'_1, \dots, p'_n\}$. Some of these packets may be empty — an empty output packet is denoted by \emptyset , and corresponds to the event that the decoder did not receive the corresponding authenticated packet. We write: $\{OP, \text{reject}\} \leftarrow \text{Decode}(PK, GID, n, \alpha, \beta, RP)$.

⁶Note that both, of course, cannot be true for an (α, β) -network.

A signature scheme has two requirements: correctness and security. We have similar requirements for a multicast authentication scheme.

A multicast authentication scheme is (α, β) -correct if, whenever at least αn correct authenticated packets are received among βn total packets, all and only the valid received packets will be decoded correctly, i.e., the corresponding data packets will be among the output packets.

A multicast authentication scheme is *secure* if, even if the adversary is allowed to query the authenticator on any number of chosen inputs, the adversary cannot make the decoder output a non-authenticated set of packets.

Definition 7 (Multicast Authentication Scheme). Probabilistic polynomial-time algorithms (KeyGen, Auth, Decode) constitute an (α, β) -correct and secure multicast authentication scheme if no probabilistic polynomial-time adversary \mathcal{A} can win non-negligibly often in the following game:

1. A key pair is generated:

$$(PK, SK) \leftarrow \text{KeyGen}(1^k).$$

2. The adversary \mathcal{A} is given:

- The public key PK as input.
- Oracle access to the authenticator, i.e., for $1 \leq i \leq \text{poly}(k)$, where $\text{poly}(\cdot)$ is a polynomial, the adversary can specify the values $(GID_i, n_i, \alpha_i, \beta_i, DP_i)$ and obtain $AP_i \leftarrow \text{Auth}(SK, PK, GID_i, n_i, \alpha_i, \beta_i, DP_i)$. However, the adversary cannot issue more than one query with the same group identification tag. That is to say, for all $i \neq j$, $GID_i \neq GID_j$.

3. At the end, \mathcal{A} outputs a group identification tag, GID , the values n , α and β , and a set of packets, RP .

The adversary *wins* the game if one of the following violations occurs:

Violation of the (α, β) -correctness property: The adversary did manage to construct RP in such a way that even though it contains $\alpha_i n_i$ packets of some authenticated packet set AP_i for group identification tag $GID_i = GID$, the decoder still failed at identifying all the correct packets. Namely, the adversary wins if *all* of the following hold:

- For some i , the adversary's query i contained $GID_i = GID$, $n_i = n$, and $\alpha_i = \alpha$. Let $DP_i = \{p_1, \dots, p_n\} = DP$ be the data packets associated with that query, and let $AP_i = \{a_1, \dots, a_n\} = AP$ be the response of the authenticator.
- At least αn of the authenticated packets (a_1, \dots, a_n) are included in the received packets RP , i.e., $|RP \cap AP| \geq \alpha n$.
- The number of received packets is at most βn , i.e., $|RP| \leq \beta n$.
- For some $1 \leq j \leq n$, p_j is the j 'th packet in the original set of data packets DP , such that the corresponding authenticated packet a_j was received, i.e., $a_j \in RP \cap AP$, and yet was not decoded correctly. Namely, let $(p'_1, \dots, p'_n) \leftarrow \text{Decode}(PK, GID, n, \alpha, \beta, RP)$. For p_j it holds that $p_j \neq p'_j$.

Violation of the security property: The adversary did manage to construct RP in such a way that the decoder will output packets $OP = \{p'_1, \dots, p'_n\}$ that were never authenticated by the authenticator algorithm for the group identification tag GID . More precisely, the adversary wins if *one* of the following happens:

- The authenticator was never queried with group identification tag GID and the size n , and yet the decoder algorithm does not reject. That is, $\text{Decode}(PK, GID, n, \alpha, \beta, RP) \neq \text{reject}$ but $\text{Decode}(PK, GID, n, \alpha, \beta, RP) = OP$.
- The authenticator was queried with the group identification tag GID , the values n , α and β , and the data packets $DP = \{p_1, \dots, p_n\}$. However, the decoder algorithm does not reject and some output packet $p'_j \neq \emptyset$ is different from the corresponding data packet p_j , where $OP = \{p'_1, \dots, p'_n\}$.

4 Construction

In this section we describe a multicast authentication scheme ($\text{KeyGen}, \text{Auth}, \text{Decode}$) that meets the definitions of the previous section. In the sequel, we denote with $0 < \epsilon < 1$ the *tolerance parameter* of the decoder, which yields a trade-off between the error-tolerance ability of the decoder and its performance. Both the authenticator and the decoder know and use the value of this parameter. Recall that this parameter expresses how far away from the ultimate bound for efficient list decodability the encoding is performed. The higher the ϵ the further away from the bound the encoding is performed, thus, the higher the communication overhead is, but the faster (up to constant factors) the list decoder operates. Similarly, values of ϵ that are closer to zero reduce the communication overhead at the cost of increasing (by constant factors) the decoding time. We will talk in detail about this trade-off when we give efficiency analysis. By \circ , we denote concatenation and by \emptyset we appropriately denote either a packet that is empty or the empty string. We also often omit the floor and ceiling notation in order to avoid notational overload.

4.1 Key Generation

We assume that a signature scheme $(G(\cdot), \text{Sign}_{(\cdot)}(\cdot), \text{Verify}_{(\cdot)}(\cdot, \cdot))$ and a family \mathcal{H} of collision-resistant hash functions are given (see Definitions 1 and 2). If $(PK_s, SK_s) \leftarrow G(1^k)$ and $H \leftarrow \mathcal{H}(1^k)$, we set $PK = (PK_s, H)$ and $SK = SK_s$.

4.2 Authenticator Auth

Input Secret key SK , public key PK , group identification tag GID , data stream size n , parameters α and β of the network and data packets $DP = \{p_1, \dots, p_n\}$.

Output Authenticated packets $AP = \{a_1, \dots, a_n\}$.

Algorithm

1. For $1 \leq i \leq n$, compute the hash value $h_i = H(p_i)$. The concatenation of all the hash values, together with the value GID , is digitally signed:

$$\sigma \leftarrow \text{Sign}_{SK}(GID \circ h_1 \circ \dots \circ h_n).$$

The string $S = h_1 \circ \dots \circ h_n \circ \sigma$ is called the *authentication information*. We want to guarantee that, even if only an α fraction of the packets survive, and a large number of packets $(\beta - \alpha)n$ are injected, the receiver still gets all the authentication information. To that end, we encode S using a $[n, \rho n + 1]_q$ Reed-Solomon code in a manner that is tolerant to packet loss and insertion, as in the following steps.

2. Let the rate of the code be

$$\rho = \frac{\alpha^2}{(1 + \epsilon)\beta}.$$

Recall that α and β are the survival and flood rates of the network, respectively, whereas ϵ is the tolerance parameter of the decoder. Observe that since $\alpha \leq 1$, $\beta \geq 1$ and $0 < \epsilon < 1$, we have $\rho < 1$.

3. Split S into $\rho n + 1$ substrings of size $\lceil \frac{|S|}{\rho n + 1} \rceil$, where each substring is viewed as a value of \mathbb{F}_q , with $q = 2^{\lceil \frac{|S|}{\rho n + 1} \rceil}$ ⁷. If S is not an exact multiple of $\rho n + 1$, pad S with ℓ 0's, such that $|S \circ 0^\ell| \bmod \rho n + 1 \equiv 0$.
4. Treat the resulting set of $\rho n + 1$ field elements as an input to the Reed-Solomon encoder (see Definition 3). Compute the corresponding codeword $C(S)$ using an $[n, \rho n + 1]_q$ Reed-Solomon code^{8,9}. $C(S)$ consists of n elements of \mathbb{F}_q , denoted as (s_1, \dots, s_n) .
5. Let $AP = \{a_1, \dots, a_n\}$, where for $1 \leq i \leq n$, we have $a_i = GID \circ i \circ p_i \circ s_i$.

4.3 Decoder Decode

Our decoder Decode uses a modification of the GS-Decoder (see Definition 3 and Theorem 1) as a subroutine. The standard GS-Decoder expects to receive, as input, n pairs (x_i, y_i) , and outputs a list L of all the polynomials of degree at most k such that every $p \in L$ has the property that for at least $\sqrt{(1 + \epsilon)kn}$ of the i 's, $p(x_i) = y_i$. We write $L \leftarrow \text{GSDecode}_\epsilon(n, k, \sqrt{(1 + \epsilon)kn}, \{(x_i, y_i) | 1 \leq i \leq n\})$. The modified decoder is specified by parameters that are slightly different: it takes as input up to βn points (x_i, y_i) and finds a list of candidate inputs (set of points) that can be encoded by polynomials of degree at most ρn (with $\rho = \frac{\alpha^2}{(1 + \epsilon)\beta}$) such that each polynomial agrees with at least αn of the input points (see Corollary 3).

What is important is that the modified GS-Decoder operates even in the presence of an *adversarially chosen* set of (x_i, y_i) pairs. In other words, the modified decoder corresponds to the alphabet and encoder of the Reed-Solomon code described in Definition 3 (the encoder is used in the Auth algorithm above), but now, the set of points (x_i, y_i) that constitute the input of the decoder is in principle different than the set of *valid* output points of the encoder; this corresponds to the various attacks by the adversary. For instance, this set may not include some of the original points, may be larger since some new points are added and may even contain points that are *vertically aligned*, i.e., some of the x_i 's are not distinct. The modified decoder is obtained by adapting the original GS-Decoder, as follows.

Modified GS-Decoder $\text{MGSDecode}_\epsilon$

Input n , α , β , and m points (x_i, y_i) , $1 \leq i \leq m$.

Output List of all computed candidates $\{c_1, \dots, c_\ell\}$ or reject.

Algorithm

⁷We see that q is in fact a function of n , α and β , thus it does not need be transmitted to the receiver. Observe that $|S|$ is a function of n .

⁸We assume that the value of ϵ is known also to the encoder; thus, in fact $C(S) = C_\epsilon(S)$.

⁹Here, we implicitly assume that the size of the problem n along with the parameters of the network α and β are such that $\rho n + 1 < n$, or equivalently, $\frac{\alpha^2}{(1 + \epsilon)\beta} + \frac{1}{n} < 1$, so that the $[n, \rho n + 1]_q$ Reed-Solomon code does not degenerate. This technical requirement is easily satisfied as n and β get larger and α gets smaller.

1. If $m > \beta n$, reject.
2. Else, if there are fewer than αn distinct values of x_i , reject.
3. Else, run the GS-Decoder, that is, let $L \leftarrow \text{GSDecode}_\epsilon(m, \rho n, \alpha n, \{(x_i, y_i) | 1 \leq i \leq m\})$, where $\rho = \frac{\alpha^2}{(1+\epsilon)\beta}$. If L is empty, reject.
4. Process $L = \{Q_1(x), \dots, Q_\ell(x)\}$ as follows: for each $Q(i) \in L$, evaluate $Q(i)$ for $1 \leq i \leq \rho n + 1$ and let the string $Q_j(1) \circ Q_j(2) \circ \dots \circ Q_j(\rho n + 1)$ be a candidate c_i .

For this decoder, operating on input points subject to constraints that are in accordance with our (α, β) -network, we have the following.

Lemma 1. *When at least αn out of its at most βn input points are valid, $\text{MGSDecoder}_\epsilon$ does not reject, runs in time $\tilde{O}(n^2)$, where $\tilde{O}(n^2)$ field operations are involved, and outputs the constant size list of all candidate inputs that are consistent with αn of the received points.*

Proof. $\text{MGSDecoder}_\epsilon$ does not reject in Steps 1 and 2 of the algorithm. All claims follow considering Step 3, Theorem 1, Corollary 3 and the fact that GS-Decoder operates even when the x_i 's are not distinct (see Guruswami and Sudan [16]). Corollary 3 holds, since, if $m = \gamma n$, $\alpha \leq \gamma \leq \beta$, then $\rho n = \frac{\rho}{\gamma} m$, thus, we consider the polynomial reconstruction problem on inputs $\frac{\rho}{\gamma} m$, t , m points and for $\rho = \frac{\alpha^2}{(1+\epsilon)\beta}$ we have that

$$t \geq \alpha n \geq \sqrt{\frac{\gamma}{\beta}} \alpha n = \sqrt{(1+\epsilon)\rho n m} = \sqrt{(1+\epsilon)\frac{\rho}{\gamma} m},$$

as needed. Of course, $\tilde{O}(m^2) = \tilde{O}(n^2)$ for $m = \gamma n$ and γ a constant. Finally, we have that $\sqrt{(1+\epsilon)\frac{\rho}{\gamma}} = \sqrt{\frac{\alpha^2}{\beta\gamma}} \leq 1$ as required. Note that the correct answer (candidate input) is guaranteed to be contained in the output list, since it is a polynomial of degree at most ρn that is consistent with αn points. Thus, $\text{MGSDecoder}_\epsilon$ does not reject in Steps 3 either. \square

Now, we are ready to describe our decoder. The idea is simple: we resist every attack by the adversary by treating injected, altered and lost packets as errors, essentially, in a *crypto-enhanced* list-decoder. What matters is only the achieved *agreement*, i.e., the valid packets.

Decoder Decode_ϵ

Input Public key PK , group identification tag GID , n , parameters α and β and received packets $RP = \{r_1, \dots, r_m\}$.

Output $OP = \{p'_1, \dots, p'_n\}$ or reject.

Algorithm

1. View packets in RP as $r_i = GID_i \circ j_i \circ p_i \circ s_i$.
2. Discard all non-conforming packets, i.e., all packets for which $GID_i \neq GID$ or packets with $j_i \notin [1..n]$. Let $(r_1, \dots, r_{m'})$ be the remaining packets in RP . Each of them is viewed as $r_i = GID \circ j_i \circ p_i \circ s_i$, such that $j_i \in [1..n]$.
3. If $m' < \alpha n$ or $m' > \beta n$, then reject.
4. For $1 \leq i \leq m'$, set $(x_i, y_i) = (j_i, s_i)$.

5. Run algorithm $\text{MGSDecoder}_\epsilon$ with input parameters n, α, β and the m' points (x_i, y_i) , $1 \leq i \leq m'$. If $\text{MGSDecoder}_\epsilon$ rejects, reject; otherwise, obtain the candidate codewords $\{c_1, \dots, c_\ell\}$.
6. For $1 \leq i \leq n$, set $h_i = \emptyset$. Let $j = 1$. While $j \leq \ell$:
 - Parse the codeword c_j as string $h_1^j \circ \dots \circ h_n^j \circ \sigma$.
 - If $\text{Verify}_{PK_s}(GID \circ h_1^j \circ \dots \circ h_n^j, \sigma) = 1$, then set $h_i = h_i^j$ for $1 \leq i \leq n$ and break out of the loop; otherwise, increment j .
7. If $(h_1, \dots, h_n) = (\emptyset, \dots, \emptyset)$, reject. Else, compute the output packets OP as follows:
 - Initialize $OP = \{p'_1, \dots, p'_n\}$: for each $1 \leq i \leq n$, set $p'_i = \emptyset$.
 - For $1 \leq i \leq m'$:
 - view r_i as $r_i = GID \circ j \circ p_j \circ s_j$, such that $j \in [1..n]$.
 - if $H(p_j) = h_j$, set $p'_j = p_j$.
8. Let $OP = \{p'_1, \dots, p'_n\}$.

For this decoder we have the following.

Lemma 2. *When operating on a stream of packets encoded by authenticator Auth and transmitted through an (α, β) -network, algorithm Decode does not reject.*

Proof. From the properties of the (α, β) -network, the algorithm does not reject in Steps 3 and 5. Since at least αn packets are valid, from Lemma 1 we have that the correct codeword (authentication information) is among the candidates of the output list of Step 5. Thus, the corresponding signature verification in Step 6 excludes the rejection in Step 7. \square

We postpone the analysis of the running time of the algorithms (KeyGen, Auth, Decode) of our scheme until the next section.

4.4 Correctness and Security Proofs

Let us show that the authentication scheme (KeyGen, Auth, Decode) described in subsections 4.1, 4.2 and 4.3 satisfies Definition 7. Suppose that we have an adversary \mathcal{A} who manages to break the (α, β) -correctness or security of our scheme with (non-negligible) probability $\pi(k)$. Then one of the following is true:

- With probability (at least) $\pi(k)/2$, the adversary \mathcal{A} violates the (α, β) correctness property.
- With probability (at least) $\pi(k)/2$, the adversary \mathcal{A} violates the security property.

Let us show that a non-negligible probability of either event contradicts the security properties of the underlying signature scheme and hash function.

Claim 1. *If a polynomial-time adversary \mathcal{A} violates the (α, β) -correctness property of our scheme, then the underlying signature scheme is not secure, or the underlying hash function is not collision-resistant.*

Proof. Let us prove the claim by exhibiting a reduction which transforms an attack that violates the correctness of our scheme, into an attack on the underlying signature scheme.

Reduction. The input to the reduction is the public key PK_s of the signature scheme. Our reduction is also given oracle access to the corresponding signer $\text{Sign}_{SK_s} (= \text{Sign}_{SK})$. The reduction sets up the public key $PK = (PK_s, H)$. Our reduction does not know the corresponding secret key. Our reduction invokes the adversary \mathcal{A} on input PK . It now needs to be able to answer the adversary's queries to the authenticator Auth . In order to respond to a query $(GID_i, n_i, \alpha_i, \beta_i, DP_i)$, run the algorithm Auth with the following modification: in Step 1, at the beginning of the algorithm Auth , instead of computing the signature σ_i , obtain it by querying the signature oracle Sign_{SK} . Everything else is carried out as prescribed by the algorithm Auth .

It is clear that the view of the adversary in this reduction will be identical to the view that the adversary obtains in real life. Therefore, with the same probability as in real life, the adversary violates the correctness property. Namely, it outputs values GID, n, α, β and the set of received packets RP , such that all of the following hold:

1. $GID = GID_i, n = n_i, \alpha = \alpha_i$, and $\beta = \beta_i$ for some i . Let $DP_i = \{p_1, \dots, p_n\}$ be the data packets associated with that query, and let AP be the response that we gave to the adversary. In particular, let σ_i be the signature associated with this query, that is, $\sigma_i \leftarrow \text{Sign}_{SK}(GID \circ H(p_1) \circ \dots \circ H(p_n))$.
2. $|RP \cap AP| \geq \alpha n$ and $|RP| \leq \beta n$.
3. For some j such that $r_j \in RP, p_j \neq p'_j$, where $(p'_1, \dots, p'_n) \leftarrow \text{Decode}(PK, GID, n, \alpha, \beta, RP)$ and r_j is a packet that corresponds to packet $p_j \in DP_i$. (Note that from 2 and Lemma 2, it follows that algorithm Decode does not reject.)

Case 1. Suppose that $p'_j \neq \emptyset$. From 3, we get that either $H(p_j) \neq H(p'_j)$, or it is easy to find a collision to the hash function. By definition of Decode , if $r_j \in RP$ and $p'_j \neq \emptyset$, then, in Step 6, the algorithm Decode processes a candidate $c = h_1 \circ \dots \circ h_n \circ \sigma$ such that $\text{Verify}_{PK_s}(GID \circ h_1 \circ \dots \circ h_n, \sigma) = 1$. We must argue that our signature oracle was never queried on input $(GID \circ h_1 \circ \dots \circ h_n)$. Note that the only time it was queried with this GID , it was when we obtained σ_i on input $(GID \circ H(p_1) \circ \dots \circ H(p_n))$. Moreover, in Step 7, Decode includes p'_j into OP if and only if $H(p'_j) = h_j$. Therefore, $h_j \neq H(p_j)$, and so our signature oracle was never queried with $(GID \circ h_1 \circ \dots \circ h_n)$, and yet our adversary has caused us to compute a signature σ such that $\text{Verify}_{PK_s}(GID \circ h_1 \circ \dots \circ h_n, \sigma) = 1$. Thus, the underlying signature scheme is insecure.

Case 2. So, suppose that $p'_j = \emptyset$. From 1 and 2, we know that αn of the original authenticated packets were received, among the total of βn packets. Then, by the properties of $\text{MGSDecoder}_\epsilon$ (Lemma 1), Step 5 of the algorithm Decode includes the candidate value $c = H(p_1) \circ \dots \circ H(p_n) \circ \sigma_i$. Then, by construction, it cannot be the case that in Step 7, $(h_1, \dots, h_n) = (\emptyset, \dots, \emptyset)$. If $(h_1, \dots, h_n) = (H(p_1), \dots, H(p_n))$, then by construction of Decode , if (as is the case according to 3) $r_j \in RP$, then $p'_j \neq \emptyset$, because p'_j is set to p_j when the packet r_j is considered in Step 7. Therefore, $(h_1, \dots, h_n) \neq (H(p_1), \dots, H(p_n))$, and yet $\text{Verify}_{PK_s}(GID \circ h_1 \circ \dots \circ h_n, \sigma) = 1$. But the only query with GID that we ever issued to the signer was for the message $(GID \circ H(p_1) \circ \dots \circ H(p_n)) \neq (GID \circ h_1 \circ \dots \circ h_n)$. Thus σ is a successful forgery. \square

Claim 2. *If a polynomial-time adversary \mathcal{A} violates the security property of our scheme, then the underlying signature scheme is not secure, or the underlying hash function is not collision-resistant.*

Proof. Consider setting up the reduction in exactly the same way as in the proof of Claim 1. Again, the adversary's view in the reduction is the same as in real life. So, just as often as in real life, the adversary will violate the security property of our scheme, namely, one of the following will hold:

1. The authenticator was never queried with group identification tag GID and size n , and yet the decoder algorithm does not reject. I.e., $\text{reject} \neq \text{Decode}(PK, GID, n, \alpha, \beta, RP) = OP$.
2. The authenticator was queried with the group identification tag GID , with the values n , α and β , and data packets $DP = \{p_1, \dots, p_n\}$. However, the decoder algorithm does not reject and some output packet $p'_j \neq \emptyset$ is different from the corresponding data packet p_j , where $OP = \{p'_1, \dots, p'_n\}$.

Suppose 1 holds. Then, from the description of the decoder, we know that the only way that it will produce some non-empty set of output packets is if, in Step 6, it sees a string c and a signature σ such that $\text{Verify}_{PK_s}(GID \circ c, \sigma) = 1$. Since the signature oracle was never queried for this GID and n , σ is a successful forgery.

So, suppose that 2 holds. Then this is exactly the same situation as Case 1 of the proof of Claim 1, and we obtain either a successful forgery or a hash function collision in the same manner. \square

4.5 Authenticated Error Correcting Code

Our authentication scheme consists of probabilistic algorithms (KeyGen , Auth , Decode) described in subsections 4.1, 4.2 and 4.3. We note that essentially our scheme uses an *authenticated* Reed-Solomon error correcting code. Using this term we refer to the fact that, by allowing the use of cryptographic primitives in the data that is encoded, list-decoding succeeds in operating (even) in the presence of *adversarial behavior* of the “transmission channel” (a network in this case). In general, depending on this adversarial behavior, our authenticated decoder either rejects or *correctly* reconstructs the original codeword (i.e., authentication information) that was sent. For an (α, β) -network, Lemma 2 and Claim 1 guarantee that the correct reconstruction is produced. On the other hand, Claim 2 guarantees that the authenticated Reed-Solomon error correcting code is *secure*. For this reason, we refer to our multicast authentication scheme as AuthECC .

Definition 8. AuthECC is the multicast authentication scheme consisting of the triplet of probabilistic algorithms (KeyGen , Auth , Decode) described above, which, in turn, are based on an authenticated Reed-Solomon error-correcting code.

We have, thus, proved the following result.

Theorem 4. *Multicast authentication scheme AuthECC is an (α, β) -correct and secure multicast authentication scheme for any (α, β) -network.*

5 Analysis

We now analyze our scheme in terms of the various cost parameters. Recall that:

- α is the survival rate of the network, where $0 < \alpha \leq 1$;
- β is the flood rate of the network, where $\beta \geq 1$;
- ϵ is the tolerance parameter of the list-decoder, where $0 < \epsilon < 1$; and
- ρ is the rate of the encoder, where $\rho = \frac{\alpha^2}{(1+\epsilon)\beta}$ and $\rho < 1$.

We start by discussing the complexity of our authentication scheme in terms of computational and communication costs and introduced delay, we then examine how our scheme can be tuned and extended and finally we close the section by comparing it with various previous proposed schemes. In the sequel, by h we denote the size of a hash value and by s the size of a digital signature.

Computational Cost. The sender and the receiver execute algorithms Auth and Decode, respectively. Both algorithms involve field operations (additions and multiplications) over finite field \mathbb{F}_q of size $q = 2^{\lceil \frac{nh+s}{\rho n+1} \rceil} \simeq 2^{\frac{h}{\rho}}$. Both operations take $O\left(\frac{h}{\rho} \log^{O(1)} \frac{h}{\rho}\right)$ time [15]. Setting $N = \frac{h}{\rho}$, both operations take $O(N \log^{O(1)} N)$ time. Note that N is independent of n .

Authenticator: The cost to encode n packets is as follows. First, n hashes are computed and one signature operation is performed over the hashes. Then, a Reed-Solomon code is applied on the authentication information, which consists of a polynomial interpolation on $\rho n + 1$ positions and a polynomial evaluation in $n - \rho n - 1$ positions. These tasks require $O(n \log n)$ field operations or $O(n \log n N \log^{O(1)} N)$ time, since both polynomial evaluation and interpolation for polynomials of degree at most n can be solved using $O(n \log n)$ field operations (thus, Reed-Solomon encoding requires a quasi-linear $O(n \log n)$ number of field operations). Observe that the use of a systematic Reed-Solomon code adds no extra computational cost.

Decoder: From Theorem 1 and Lemma 1, we have that $O(\beta^2 n^2 N) = \tilde{O}(n^2)$ field operations are required and thus $O(\beta^2 n^2 N^2 \log^{O(1)} N) = \tilde{O}(n^2)$ time is needed for the decoder to run. Also, for each of $O(1)$ candidate polynomials, we perform a polynomial evaluation at $\rho n + 1$ positions, thus $O(n \log n)$ field operations and so $O(n \log n N \log^{O(1)} N)$ time, and one signature verification. In total, we have $O(n^2 N^2 \log^{O(1)} N) = \tilde{O}(n^2)$ processing time and $O(1)$ signature verifications. Finally, $O(n)$ hash values are computed.

Communication Cost. The size of the authentication information is $\frac{n}{\rho n+1}(s + hn)$. That is, we have *constant communication overhead* per packet

$$\frac{s + hn}{\rho n + 1} < \frac{h}{\rho} + \frac{s}{\rho n} = \frac{h}{\rho} + o(1).$$

We see that $1/\rho$ hash values are included in each packet, with $\rho = \frac{\alpha^2}{(1+\epsilon)\beta} < 1$. The larger the value of ρ the smaller the authentication overhead.

Delay. As *delay*, we count the number of packets that the authenticator or decoder algorithm has to buffer. Of course, by definition, any authentication scheme according to our model needs to process n packets. However, delay is a cost parameter that is useful even in our model, since it captures the ability of the authenticator or the decoder to process packets in an *on line* fashion. In our scheme the sender processes n packets and the receiver processes βn packets in the worst case. However, the receiver can invoke an decoding procedure only after $\rho n + 1$ or αn packets have been received.

In particular, the receiver can try to compute the authentication information exactly after $\rho n + 1$ packets are received: the used code is systematic and the first $\rho n + 1$ symbols of $E(S)$ equal S (where S is the authentication information). Of course, we need no packet loss to occur among these packets. If the correct polynomial is computed (and verified) from the first $\rho n + 1$ packets, the authentication information is computed without any decoding overhead. Similarly, the receiver can try to compute the authentication information after αn packets are received: this time the decoder runs completely, but computation is a less expensive, and if the correct authentication information is computed, no attack is in process and the delay is αn . Otherwise, if no polynomial can be verified, the receiver is under attack and βn delay is required in the worst case. In other words, our scheme can distinguish between the less expensive *detection* of an attack by an adversary from the more expensive *verification* of the valid received packets. We believe that this feature is desirable, for less computational effort is spent when no adversary acts.

We can summarize the performance of our scheme as follows.

Theorem 5. *For any (α, β) -network, multicast authentication scheme AuthECC achieves the following performance in authenticating n packets.*

- *At the sender, n packets are processed and one signature operation, n hash computations and $O(n \log n)$ field operations are performed.*
- *At the receiver, at most βn packets are processed and $O(1)$ signature verifications, βn hash computations and $\tilde{O}(n^2)$ field operations are performed.*
- *Communication overhead is constant per packet, proportional to $\frac{\beta}{\alpha^2}$.*

5.1 Tuning and Extensions

Given specific values of the survival rate α and flood rate β of the network, the parameter ρ , which controls the communication overhead, can be tuned by the tolerance parameter ϵ . This gives one degree of freedom in implementing the exact encoding-decoding procedures. Namely, bandwidth consumption can be decreased at the cost of increasing by a constant factor the time complexity and vice versa. A realistic deployment of our scheme can consider α and β as an additional degree of freedom: early packet streams (groups of packets) are encoded for bigger values of α and smaller values of β . Depending on the observed network's behavior, the network parameters can be later adjusted to a new desired level of security. Table 1 shows the communication overhead per packet for specific values of α , β and ϵ .

α	β	ϵ	$1/\rho$	cost c (bytes)	α	β	ϵ	$1/\rho$	cost c (bytes)
0.33	1.5	0.1	15.15	303	0.5	1	0.01	4.04	81
0.5	1.5	0.1	6.6	132	0.5	2	0.01	8.08	162
0.75	1.5	0.1	2.93	59	0.5	3	0.01	12.12	243
0.33	1.5	0.5	20.66	414	0.5	1	0.1	4.4	88
0.5	1.5	0.5	9	180	0.5	2	0.1	8.8	176
0.75	1.5	0.5	4	80	0.5	3	0.1	13.2	264

Table 1: Communication cost c per packet for various values of the survival rate α , flood rate β and tolerance parameter ϵ . We assume the use of the SHA-1 hashing algorithm, that is, $h = 20$ bytes. The communication cost should be compared with the size s of the signature in use (e.g., an RSA signature with $s = 256$ bytes). Recall that $\rho = \frac{\alpha^2}{(1+\epsilon)\beta}$ is the rate of the code in use and that $c = \frac{h}{\rho} = \frac{\beta(1+\epsilon)}{\alpha^2}h$.

Independently of the choice of parameters, our scheme can be further modified in two ways, achieving different trade-offs between communication cost and computational efficiency. First, we can decrease the communication overhead, by applying the technique of [24]. The idea is that, since (at least) αn packets are guaranteed to be received intact, a significant portion of the authentication information is obtained by the decoder for free and without decoding: the (at least) αn hash values of the valid packets. Thus, less authentication information can be used and less redundancy is added to packets. To implement this idea, one has to encode the n hash values appropriately and, thus, Reed-Solomon codes are applied twice. Interestingly, as opposed to the case of erasure codes [24], in our case where Reed-Solomon error correcting codes are used, the decrease of the communication overhead occurs only for appropriate ranges of values for the network parameters α and β .

In particular, let $\{X, X'\} \leftarrow C[n, k+1]_q(X)$ denote the application of systematic Reed-Solomon code $[n, k+1]_q$ on word X , where X' is the added redundancy. Also let $H = h_1 \circ \dots \circ h_n$ be the hash values of the n packets. We get the modified scheme by encoding

$$\{H, H'\} \leftarrow C[\gamma n, n+1]_{q_1}(H)$$

and then

$$\{A, A'\} \leftarrow C[n, \rho n + 1]_{q_2}(A),$$

where

$$A = H' \circ \text{Sign}_{SK}(H), \quad \gamma = 1 - \alpha + \sqrt{(1 + \epsilon)\beta}, \quad q_1 = 2^h, \quad \rho = \frac{\alpha^2}{(1 + \epsilon)\beta}, \quad q_2 = 2^{\lceil \frac{|A|}{\rho n + 1} \rceil}.$$

As in our basic scheme, $A \circ A'$ is split in n equal shares A_i and packet p_i corresponds to authenticated packet $a_i = GID \circ i \circ p_i \circ A_i$. At the decoder, by the network reliability (at least αn packets will be valid) it is guaranteed that a constant size list of candidate strings for $H' \circ \text{Sign}_{SK}(H)$ is produced; also, list-decoding is transformed to unambiguous decoding by verifying a constant number of signatures. Furthermore, the receiver is always capable to list-decode the packet hashes H . If in total δn packets reach the receiver, $\delta \leq \beta$, then Corollary 3 holds, since, $t \geq \alpha n + (\gamma - 1)n \geq \sqrt{(1 + \epsilon)\delta n}$. The per packet communication overhead of this scheme is $\frac{(\gamma - 1)h}{\rho}$. When $\gamma < 2$, with this scheme we save in communication overhead. That is, for network parameters α and β in appropriate ranges so that $\beta < \frac{(\alpha + 1)^2}{1 + \epsilon}$ we can decrease the communication cost by the constant factor γ at the cost of increasing the computational cost by roughly a factor of 2, since two applications of Reed-Solomon codes are required.

Also, by decreasing the field size, we can reduce the cost of performing field operations. For instance, we could split the authentication information into $\gamma \rho n + 1$ substrings of size ℓ , $\gamma > 1$ (e.g., $\gamma = 10$), consider each substring as a field element in \mathbb{F}_q , with $q = 2^\ell$, encode with a $[\gamma n, \gamma \rho n + 1]_q$ Reed-Solomon code, and split the augmented authentication information into n pieces (each of γ field elements). In this way, the communication cost stays the same, but field operations become faster. The number of field operations at the encoder or decoder is increased by only a constant factor. Depending on the hardware architecture, this modification may be useful. A drawback here is that one injected packet by the adversary is now affecting the decoding algorithm by a factor γ .

5.2 Comparison with other schemes

We compare our schemes against various classes of proposed multicast authentication schemes.

Sign-All and Merkle Tree Schemes. The sign-all and Merkle-tree [36] authentication schemes are resilient to fully adversarial networks. The sign-all scheme involves one signature (resp. verification) operation per packet and a communication overhead that is equal to the signature size. Depending on the specific signature scheme in use, the parameters of our scheme or the architecture, both communication and computational costs of our scheme are comparable to the corresponding costs of the sign-all scheme.

Very short signature schemes have recently been proposed [3]. While the length of a signature can be as low as 160 bits, the security of this signature scheme is only proven in the random oracle model, and only under a strong assumption (Diffie-Hellman assumption in gap-DH groups, see Boneh and Franklin [2] for more on these groups). Signing every packet with this short signature, therefore, has a communication advantage over our construction, but loses in provable security. On

the other hand, signing every packet with a provably secure signature, such as the the Cramer-Shoup [6] signature or its modification due to Fischlin [9], will add about 500 bytes to each packet — which is more than what we have for reasonable α and β .

Additionally, signing every packet is undesirable in practice. Indeed, by signing every packet separately we lose both in efficiency and in architecture design since the secret key operations are computationally expensive and require extra need of security. Invoking a signature operation involves fetching the private key and temporarily storing it in the main memory of the system. When secret key operations are performed at high rates, the secret key resides almost exclusively in the memory of the system increasing the danger of the key being compromised to other running processes in the system. Special-purpose hardware can be used to overcome this problem, but of course at a higher cost. In terms of secure architecture design costs, and also for provable security or efficiency reasons, the sign-all approach is inferior to ours.

Finally, since one signature verification must be performed for each received packet, valid or not, the sign-all solution suffers by the following denial-of-service attack at the receiver: by injecting invalid packets an adversary can increase the computation resources spent at the receiver for signature verifications. In our scheme, where signature dispersal is used, no such attack is possible.

On the other hand, the Merkle-tree scheme [36] has better time complexity than our scheme. For a group of packets of size n , only $2n$ hash computations and one signature computation (resp. verification) are performed at the sender (resp. receiver). However, the Merkle-tree scheme has communication cost that grows with the number of packets, thus, this scheme is not scalable. Our scheme is efficient in terms of communication cost: packets have constant authentication overhead.

Another drawback for the Merkle-tree scheme operating in a fully adversarial network is the *signature flooding* attack, identified and described in [17]: in a way similar to the denial-of-service attack against the “sign-all” scheme, injected packets cause a signature verification at the receiver for the Merkle-tree scheme as well. Of course, we have to note that at the receiver, by appropriately caching hash values, we can significantly resist against the signature flooding attack: once the first valid packet is verified, its (authenticated) hashes are stored and subsequent packets need only be verified with respect to the hashes they carry and not with respect to the signature they carry. Because of that, injected packets that are received afterwards do not cause signature verifications. Although this kind of attack is thus terminated after the first correct signature verification, still, in a fully adversarial network with packets rearrangements, injected packets will precede the valid ones. In our (α, β) -network model, the Merkle-tree scheme [36] needs $(\beta - \alpha)n$ signature verifications. Instead, our scheme performs only a constant number of signature verifications at the receiver.

Graph-Based Schemes. These schemes [13, 29, 23, 34] assume the reliable receipt of a signature packet. However, a fully adversarial network will capture the signature packet and invalidate the scheme. Even if the signature packet is assumed to arrive intact, any efficient scheme in terms of communication overhead (i.e., with constant overhead for packet) will have the undesirable property that $O(1)$ critical packets can be adversarially chosen to disconnect from the authentication chain the signature node (packet). In the piggybacking scheme in [23], this number of critical packets can be $O(n)$ at the expense of a communication overhead of $O(n)$ per packet. Our scheme does not have these drawbacks since the signature is dispersed among all the packets. As opposed to graph-based authentication where the authentication of a packet crucially depends on other packets (with packets closer to the signature packet being more important), our scheme is symmetric in this context: all packets share the authentication information.

Erasure-Code Schemes. The first two erasure-code based proposed schemes [24, 25] make use of erasure codes to tolerate packet losses, up to a constant fraction. However, no packet injections are tolerated: a single injected packet suffices to fail the decoding procedure. For networks where packets get only lost, they perform slightly better than our scheme in terms of communication cost and time complexity. This is due to the fact that erasure codes are more efficient than error-correcting codes in terms of time complexity and space requirement. Moreover, erasure codes can tolerate more erasures than the theoretical limit $d/2$ for error correcting codes (d is the diameter of the code). In our authentication scheme, tolerating injected packets comes at this small price of having *slightly* worse performance than erasure-code schemes.

In [17], Karlof *et al.* using distillation codes address the vulnerability to packet injections that any scheme based on erasure-codes has, but their proposed scheme has high communication overhead and is thus less scalable, because a Merkle hash tree is used to “filter out” the injected packets (and thus the communication cost is $O(\log n)$). For such a logarithmic communication overhead, the scheme by Wong and Lam [36] may be actually preferable since it has both lower time complexity and better resiliency to adversarial network behavior. In terms of computational effort at the sender and receiver this scheme is similar to our scheme except from the following two points regarding the computational effort at the receiver. In [17], partitioning the packets into groups introduces an extra computational overhead and the total number of hashing values computed is by a logarithmic factor larger. In our scheme, the constants involved in the quadratic decoding process are higher than in the scheme by Karlof *et al.*

Finally, in [14], the shared channel model that is used does not tolerate adversarially chosen packet losses.

Other Schemes. TESLA [29, 28] and the scheme by Xu and Sandhu [37] have very different assumptions from our model. They are both based on MACs and on strong time-synchronization requirements about the nodes of the networks that do not fit our model. For instance, in [37], the routers of the networks are considered trusted entities.

Tables 2 and 3 summarize the above discussion, where selected schemes are compared with our scheme. In particular, Table 2 compares our scheme with the sign-all solution and various selected schemes that are not (α, β) -correct and secure. We consider two graph-based authentication schemes, one of constant degree (expander construction [34]) and one of $O(n)$ degree (piggybacking scheme with parameterized performance [23], where we assume a constant number of classes), and one erasure scheme (optimized in terms of communication scheme [24]). Table 3 compares our scheme with the only two (α, β) -correct and secure previous approaches, namely the schemes by Wong and Lam [36] and Karlof *et al.* [17].

6 Conclusion

In this paper, we propose a new general framework for the multicast authentication problem, where the network is controlled by an adversary that has great, yet not unlimited, power in modifying the transmitted stream. Our model is realistic in terms of adversarial model and the security assumptions. The limitations on the adversary’s power, characterized by the survival and flood rates, exclude from consideration only degenerate cases, where the authentication problem actually disappears.

Our work establishes a new direction in multicast authentication by going beyond erroneous networks and addressing fully adversarial networks. Our authentication scheme is efficient, lightweight and practical. It is as secure as the “sign-all” solution, but more efficient in both computational

	Sign-all	GB [34]	GB [23]	Erasure [24]	Our
Delay (Sender)	1	n	n	n	n
Computation (Sender)					
Sign	n	1	1	1	1
hash	—	$O(n)$	$O(n^2)$	n	n
field op	—	—	—	$O(n \log n)$	$O(n \log n)$
Communication	sn	$O(hn)$	$O(hn^2)$	$\frac{1-\alpha}{\alpha}hn$	$\frac{\beta(1+\epsilon)}{\alpha^2}hn$
Delay (Receiver)	1	n	n	n	βn
Computation (Receiver)					
Verify	n	1	1	1	$O(1)$
hash	—	$O(n)$	$O(n^2)$	n	βn
field op	—	—	—	$O(n^2)$	$\tilde{O}(n^2)$
Secret key protection	—	•	•	•	•
Resiliency					
Chosen packet loss	•	—	•	•	•
Chosen packet injection	•	•	•	—	•
Signature dispersal	•	—	—	•	•

Table 2: Comparison of selected multicast authentication approaches, no (α, β) -correct and secure, with respect to various aspects of efficiency, security and resiliency. By Sign, we denote a signature operation, Verify denotes a signature verification, hash denotes the total hashing cost, where we consider that the complexity of hashing a string is a linear function of the string size. Also, we use the following notation: n is the number of packets in the data stream, s is the signature size and h is the hash size. Both the communication overhead and the computational costs refer to n packets.

effort and communication overhead. Its constant communication overhead makes it scalable and preferable to the approaches by Wong and Lam [36] and Karlof *et al.* [17]. When compared with this Merkle-tree based scheme, the $O(n^2)$ time complexity of our scheme is a shortcoming. However, it is possible that in practice this may not be a serious concern. Additionally, our scheme can be tuned by the network parameters α and β and distinguishes between the less expensive detection of an attack by the adversary and the more expensive task of verification.

Open problems to address in future work are as follows. First, we would like to investigate the practical performance of our multicast authentication approach by implementing it and conducting an experimental study. Also, a natural question to explore is whether other classes of error correcting codes can be employed in our framework. But even for our proposed scheme, one open question is whether the decoding procedure can be simplified and the whether the time complexity can be improved.

Moreover, in this paper we showed a connection between coding theory and cryptography. In particular, we employed cryptographic primitives to unambiguously list-decode an error correcting code. It would be interesting to study whether there are other connections between the two areas. Finally, we would like to explore the use of our technique in other authentication problems.

	Merkle [36]	Distillation Code [17]	Our
Delay (Sender)	n	n	n
Computation (Sender)			
Sign	1	1	1
hash	$2n$	$3n$	n
field op	—	$O(n \log n)$	$O(n \log n)$
Communication	$(s + h \log n)n$	$(\frac{1}{\alpha} + \log n)hn$	$\frac{\beta(1+\epsilon)}{\alpha^2}hn$
Delay (Receiver)	1	βn	βn
Computation (Receiver)			
Verify	$(\beta - \alpha)n$	$\frac{\beta}{\alpha}$	$O(1)$
hash	$2n$	βn	βn
field op	—	$O(n^2)$	$\tilde{O}(n^2)$
Secret key protection	•	•	•
Resiliency			
Chosen packet loss	•	•	•
Chosen packet injection	•	•	•
Signature dispersal	•	•	•

Table 3: Comparison of the three (α, β) -correct and secure multicast authentication schemes with respect to various aspects of efficiency, security and resiliency. Again, both the communication overhead and the computational costs refer to a group of n packets.

Acknowledgments

We would like to thank Philip Klein for useful discussions and also Christoph Schuba and the anonymous referees for their detailed comments.

This work was supported in part by the National Science Foundation under grants CCR-0311510 and IIS-0324846 (Information Technology Research program).

References

- [1] D. Boneh, G. Durfee, and M. Franklin. Lower bounds for multicast message authentication. In B. Pfitzmann, editor, *Advances in Cryptology — EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 437–452. Springer Verlag, 2001.
- [2] D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003.
- [3] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In *Advances in Cryptology, ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 514–532. Springer-Verlag, 2001.
- [4] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. In *IEEE INFOCOM*, pages 708–716, 1999.
- [5] A. C.-F. Chan. A graph-theoretical analysis of multicast authentication. In *Proc. 23rd International Conference on Distributed Computing Systems – ICDCS*, pages 155–162, 2003.

- [6] R. Cramer and V. Shoup. Signature schemes based on the strong RSA assumption. *ACM Transactions on Information and System Security*, 3(3):161–185, 2000.
- [7] T. Cucinotta, G. Cecchetti, and G. Ferraro. Adopting redundancy techniques for multicast stream authentication. In *Proc. 9th IEEE International Workshop on Future Trends of Distributed Computing Systems (FTDCS 2003)*, pages 189–201, 2003.
- [8] Y. Desmedt, Y. Frankel, and M. Yung. Multi-receiver/multi-sender network security: Efficient authenticated multicast/feedback. In *IEEE Conference on Computer Communications — INFOCOM '92*, pages 2045–2054. IEEE-Press, 1992.
- [9] M. Fischlin. The Cramer-Shoup strong-RSA signature scheme revisited. In Y. Desmedt, editor, *Public Key Cryptography - PKC 2003*, volume 2567 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.
- [10] R. Gennaro and P. Rohatgi. How to sign digital streams. In B. Kaliski, editor, *Advances in Cryptology — CRYPTO '97*, volume 1294 of *Lecture Notes in Computer Science*, pages 180–197. Springer Verlag, 1997.
- [11] O. Goldreich, R. Rubinfeld, and M. Sudan. Learning polynomials with queries: The highly noisy case. *SIAM Journal on Discrete Mathematics*, 13(4):535–570, November 2000.
- [12] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, Apr. 1988.
- [13] P. Golle and N. Modadugu. Authenticating streamed data in the presence of random packet loss. In *Network and Distributed System Security Symposium — NDSS '01*, pages 13–22, 2001.
- [14] C. Gunter, S. Khanna, K. Tan, and S. Venkatesh. DoS protection for reliably authenticated broadcast. In *Eleventh Annual Network and Distributed Systems Security Symposium — NDSS '04*, 2004.
- [15] V. Guruswami. *List Decoding of Error-correcting Codes*. PhD thesis, Massachusetts Institute of Technology, Boston, MA, 2001.
- [16] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometric codes. In *IEEE Transactions on Information Theory*, pages 45:1757–1767, 1999.
- [17] C. Karlof, N. Sastry, Y. Li, A. Perrig, and J. Tygar. Distillation codes & applications to DoS resistant multicast authentication. In *Eleventh Annual Network and Distributed Systems Security Symposium — NDSS '04*, 2004.
- [18] H. Krawczyk. Distributed fingerprints and secure information dispersal. In *13th ACM Symposium on Principles of Distributed Computing*, pages 207–218. ACM, 1993.
- [19] M. Krohn, M. Freedman, and D. Mazieres. On-the-fly verification of rateless erasure codes for efficient content distribution. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 226–240, May 2004.
- [20] M. Luby. LT codes. In *43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '02)*, 2002.
- [21] M. G. Luby, M. Mitzenmacher, M. A. Shokrollahi, and D. A. Spielman. Efficient erasure correcting codes. *IEEE Transactions on Information Theory*, 47(2):569–584, February 2001.

- [22] A. Lysyanskaya, R. Tamassia, and N. Triandopoulos. Multicast authentication in fully adversarial networks. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 241–255, May 2004.
- [23] S. Miner and J. Staddon. Graph-based authentication of digital streams. In *IEEE Symposium on Security and Privacy*, pages 232–246, 2001.
- [24] A. Pannetrat and R. Molva. Efficient multicast packet authentication. In *Proc. Network and Distributed System Security Symposium — NDSS '03*, 2003.
- [25] J. M. Park, E. K. P. Chong, and H. J. Siegel. Efficient multicast packet authentication using signature amortization. In *IEEE Symposium on Security and Privacy*, pages 227–240, 2002.
- [26] J. M. Park, E. K. P. Chong, and H. J. Siegel. Efficient multicast packet authentication using erasure codes. *ACM Transactions on Information and System Security*, pages 6(2):258–285, May 2003.
- [27] A. Perrig. The BiBa one-time signature and broadcast authentication protocol. In *Eighth ACM Conference on Computer and Communication Security*, pages 28–37, November 2001.
- [28] A. Perrig, R. Canetti, D. Song, and J. Tygar. Efficient and secure source authentication for multicast. In *Proceedings of Network and Distributed System Security Symposium – NDSS '01*, 2001.
- [29] A. Perrig, R. Canetti, J. Tygar, and D. Song. Efficient authentication and signing of multicast stream over lossy channels. In *IEEE Symposium on Security and Privacy*, pages 56–73, 2000.
- [30] M. O. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of ACM*, 36(2):335–348, 1989.
- [31] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *SIAM Journal of Applied Mathematics*, 8(2):300–304, 1960.
- [32] P. Rohatgi. A compact and fast hybrid signature scheme for multicast packet authentication. In *Proc. 6th ACM Conference on Computer and Communications Security*, pages 93–100. ACM, 1999.
- [33] G. J. Simmons. Authentication theory/coding theory. In *Proceedings of the Conference on Advances in Cryptology (CRYPTO'84, Santa Barbara, CA)*, LNCS 196, Springer-Verlag, pages 411–431, 1984.
- [34] D. Song, D. Zuckerman, and J. D. Tygar. Expander graphs for digital stream authentication and robust overlay networks. In *IEEE Symposium on Security and Privacy*, pages 258–27, 2002.
- [35] L. Welch and E. Berlekamp. Error correction of algebraic block codes. U.S. Patent Number 4,633,470, issued December 1986.
- [36] C. K. Wong and S. S. Lam. Digital signatures for flows and multicasts. *IEEE/ACM Transactions on Networking*, 7(4):502–513, August 1999.
- [37] S. Xu and R. Sandhu. Authenticated multicast immune to denial-of-service attack. In *Proc. ACM Symposium on Applied Computing*, pages 196–200, Madrid, Spain, March 2002.