

Byzantine Modification Detection in Multicast Networks with Random Network Coding

Tracey Ho, Ben Leong, Ralf Koetter, Muriel Médard, Michelle Effros, and David R. Karger

Abstract

This paper gives an information theoretic approach for detecting Byzantine modifications in networks employing random linear network coding. Each exogenous source packet is augmented with a flexible number of hash symbols that are obtained as a polynomial function of the data symbols. This approach depends only on the adversary not knowing the random code coefficients of all other packets received by the sink nodes when designing its adversarial packets. We show how the detection probability varies with the overhead (ratio of hash to data symbols), coding field size and the amount of information unknown to the adversary about the random code that reaches a sink node. This approach can be used in conjunction with higher overhead schemes that are activated only upon detection of a Byzantine node.

I. INTRODUCTION

We consider the problem of information theoretic detection of Byzantine, i.e. arbitrary, modifications of transmitted data in a network coding setting.

Interest in network coding has grown following demonstrations of its various advantages: in network capacity [1], robustness to non-ergodic network failures [2] and ergodic packet erasures [3], [4], and distributed network operation [5]. Multicast in overlay and ad hoc networks is a promising application. Since packets are forwarded by end hosts to other end hosts, such networks are susceptible to Byzantine errors introduced by compromised end hosts.

We show that Byzantine modification detection capability can be added to a multicast scheme based on random linear block network coding [5], [6], with modest additional computational and communication overhead, by incorporating a simple polynomial hash/check value in each packet. With this approach, a sink node can detect Byzantine modifications with high probability, as long as these modifications have not been designed with knowledge of the random coding combinations present in all other packets obtained at the sink: the only essential condition is the adversary's incomplete knowledge of the random network code seen by the sink. No other assumptions are made regarding the topology of the network or the adversary's power to corrupt or inject packets. The adversary can know the entire message as well as portions of the random network code, and can have the same (or greater) transmission capacity compared to the source. This approach works even in the extreme case where every packet received by a sink has been corrupted by being coded together with an independent adversarial packet. This new adversarial model may be useful for application scenarios in which conventional assumptions of an upper bound on adversarial transmission capacity are less appropriate. For instance, in some peer-to-peer or wireless ad hoc settings we may not know how many adversarial nodes might join the network, while it may be more likely that there will be some transmissions that are not received by the adversarial nodes. In such cases, our approach can provide a useful alternative to existing methods.

Our approach provides much flexibility in trading off between the detection probability, the proportion of redundancy, the coding field size and the amount of information about the random code that is not observed by the adversary. This approach can be used for low overhead monitoring during normal conditions when no adversary is known to be present, in conjunction with more complex, higher overhead techniques which are activated upon detection of a Byzantine error, such as adding more redundancy for error correction.

A preliminary version of this work with less general assumptions appeared in [7]. The security model is substantially generalized and strengthened in this work.

A. Background and related work

The problem of secure network communications in the presence of Byzantine adversaries has been studied extensively, e.g. [8], [9], [10], [11]. A survey of information theoretic research in this area is given in [12]. Two important issues are secrecy and authenticity¹; this work concerns the latter. Like one-time pads [13], our approach relies on the generation of random values unknown to the adversary, though the one-time pad provides secrecy and not authenticity.

In the network coding context, the problem of ensuring secrecy in the presence of a wire tap adversary has been considered in [14], [15], [16]. The problem of correcting adversarial errors, which is complementary to our work, has been studied in [17], [18], [19], [20], [21].

Tracey Ho and Michelle Effros are with the California Institute of Technology, Pasadena, CA 91125, e-mail: {tho,effros}@caltech.edu

Ben Leong is with the National University of Singapore, e-mail: benleong@comp.nus.edu.sg

Ralf Koetter is with the Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801, e-mail: koetter@csl.uiuc.edu

Muriel Médard is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, e-mail: medard@mit.edu

David R. Karger is with the Computer Science and Artificial Intelligence Laboratory, MIT, e-mail karger@csail.mit.edu

¹These are independent attributes of a cryptographic system [13].

Adversarial models in existing works on information theoretic authenticity techniques commonly assume some upper bound on the number of adversarial transmissions, which leads to a requirement on the amount of redundant network capacity. For the problem of adversarial error correction or resilient communication, the number of links/transmissions controlled by the adversary must necessarily be limited with respect to the number of links/transmissions in a minimum source-sink cut or the amount of redundancy transmitted by the source. For instance, in the resilient communication problem of Dolev et al. [9], the source and sink are connected by n wires, and their model requires that no more than $(n - 1)/2$ wires are disrupted by an adversary for resilient communication to be possible. In the network coding error correction problems of [17], [20], [21], the rate of redundant information that the source needs to transmit is between one and two times the maximum rate of information that can be injected by the adversary, depending on the specific adversarial model.

The above techniques can also be considered in the context of error detection. For example, in one phase of the secret sharing based algorithm in [9], the source communicates a degree τ polynomial $f(x) \in \mathbb{F}_q(x)$ by sending $f(i)$ on the i th wire. If the adversary controls at most $n - \tau$ wires, any errors it introduces can be detected. In general, for approaches based on error-correcting codes such as in [17], the number of adversarial errors that can be detected is given by the difference between the source-sink minimum cut and the source information rate.

Such approaches have a threshold nature in that they do not offer graceful performance degradation when the number of adversarial transmissions exceeds the assumed upper bound. Their efficiency is also sensitive to overestimates of adversarial transmission capacity, which determines the amount of redundancy required.

The adversarial model considered in this paper is slightly different. Instead of assuming a limit on the number of adversarial errors, our only assumption is on the incompleteness of the adversary's knowledge of the random code (the adversary can know the entire source message). In this case, the overhead (proportion of redundant information transmitted by the source) is no longer a function of the estimated upper bound on the number of adversarial errors. Instead, it is a design parameter which, as we will show, can be flexibly traded off against detection probability and coding field size. Unlike approaches based on secret sharing and its variants where the required proportional overhead is a function of the adversarial strength, in our approach, for any nonzero proportional overhead and any adversarial strength short of full knowledge or control of network transmissions, the detection probability can be made arbitrarily high by increasing the field size. The former has the advantage of deterministic guarantees, while our approach has the advantage of greater flexibility with additional performance parameters that can be traded off against one another.

The use of our error detection technique for low-overhead monitoring under normal conditions when no adversary is known to be present, in conjunction with a more complex technique activated upon detection of an adversary, has a parallel in works such as [22] and [23]. These works optimize for normal conditions by using less complex message authentication codes and signed digests respectively during normal operation, resorting to more complex recovery mechanisms only upon detection of a fault.

B. Notation

In this work we denote matrices with bold uppercase letters and vectors with bold lowercase letters. All vectors are row vectors unless indicated otherwise with a subscript T . We denote by $[\mathbf{x}, \mathbf{y}]$ the concatenation of two row vectors \mathbf{x} and \mathbf{y} .

II. MODEL AND PROBLEM FORMULATION

Consider random linear block network coding [5], [6], [24] of a block of r exogenous packets which originate at a source node and are multicast to one or more sink nodes. We assume that the network coding subgraph is given by some separate mechanism, the details of which we are not concerned with². An adversary observes some subset of packets transmitted in the network, and can corrupt, insert or delete one or more packets, or corrupt some subset of nodes. The only assumption we make is that the adversary's observations are limited such that when designing the adversarial packets, the adversary does not know the random coding combinations present in all other packets obtained at the sinks. This assumption is made precise using the notion of secret packets which we define below. The source and sinks do not share any keys or common information.

Each packet p in the network is represented by a row vector \mathbf{w}_p of $d + c + r$ symbols from a finite field \mathbb{F}_q , where the first d entries are data symbols, the next c are redundant hash symbols and the last r form the packet's (*global*) *coefficient vector* \mathbf{t}_p . The field size is 2 to the power of the symbol length in bits. The hash symbols in each exogenous packet are given by a function $\psi_d : \mathbb{F}_q^d \rightarrow \mathbb{F}_q^c$ of the data symbols. The coefficient vector of the i th exogenous packet is the unit vector with a single nonzero entry in the i th position. The coefficient vectors are used for decoding at the sinks as explained below.

Each packet transmitted by the source node is an independent random linear combination of the r exogenous packets, and each packet transmitted by a non-source node is an independent random linear combination of packets received at that node. The coefficients of these linear combinations are chosen with the uniform distribution from the finite field \mathbb{F}_q , and the same linear operation is applied to each symbol in a packet. For instance, if packet p_3 is formed as a random linear combination of packets p_1 and p_2 , then $\mathbf{w}_{p_3} = \gamma_{1,3}\mathbf{w}_{p_1} + \gamma_{2,3}\mathbf{w}_{p_2}$ where $\gamma_{1,3}$ and $\gamma_{2,3}$ are random scalar coding coefficients distributed uniformly over \mathbb{F}_q .

²The network coding subgraph defines the times at which packets are or can be transmitted on each network link (see, e.g. [25])

Let row vector $\mathbf{m}_i \in \mathbb{F}_q^{(c+d)}$ represent the concatenation of the data and hash symbols for the i th exogenous packet, and let \mathbf{M} be the matrix whose i th row is \mathbf{m}_i . A packet p is *genuine* if its data/hash symbols are consistent with its coefficient vector, i.e. $\mathbf{w}_p = [\mathbf{t}_p \mathbf{M}, \mathbf{t}_p]$. The exogenous packets are genuine, and any packet formed as a linear combination of genuine packets is also genuine. *Adversarial packets*, i.e. packets transmitted by the adversary, may contain arbitrary coefficient vector and data/hash values. An adversarial packet p can be represented in general by $[\mathbf{t}_p \mathbf{M} + \mathbf{v}_p, \mathbf{t}_p]$, where \mathbf{v}_p is an arbitrary vector \mathbb{F}_q^{c+d} . If \mathbf{v}_p is nonzero, p (and linear combinations of p with genuine packets) are non-genuine.

A set \mathcal{S} of packets can be represented as a block matrix $[\mathbf{T}_\mathcal{S} \mathbf{M} + \mathbf{V}_\mathcal{S} | \mathbf{T}_\mathcal{S}]$ whose i th row is \mathbf{w}_{p_i} where p_i is the i th packet of the set. A sink node t attempts to decode when it has collected a *decoding set* consisting of r linearly independent packets (i.e. packets whose coefficient vectors are linearly independent). For a decoding set \mathcal{D} , the decoding process is equivalent to pre-multiplying the matrix $[\mathbf{T}_\mathcal{D} \mathbf{M} + \mathbf{V}_\mathcal{D} | \mathbf{T}_\mathcal{D}]$ with $\mathbf{T}_\mathcal{D}^{-1}$. This gives $[\mathbf{M} + \mathbf{T}_\mathcal{D}^{-1} \mathbf{V}_\mathcal{D} | \mathbf{I}]$, i.e. the receiver decodes to $\mathbf{M} + \tilde{\mathbf{M}}$, where

$$\tilde{\mathbf{M}} = \mathbf{T}_\mathcal{D}^{-1} \mathbf{V}_\mathcal{D} \quad (1)$$

gives the disparity between the decoded packets and the original packets. If at least one packet in a decoding set is non-genuine, $\mathbf{V}_\mathcal{D} \neq \mathbf{0}$, and the decoded packets will differ from the original packets. A decoded packet is *inconsistent* if its data and hash values do not match, i.e. applying the function ψ_d to its data values does not yield its hash values. If one or more decoded packets are inconsistent, the sink declares an error.

The coefficient vector of a packet transmitted by the source is uniformly distributed over \mathbb{F}_q^r ; if a packet whose coefficient vector has this uniform distribution is linearly combined with other packets, the resulting packet's coefficient vector has the same uniform distribution. We are concerned with the distribution of decoding outcomes conditioned on the adversary's information, i.e. the adversary's observed and transmitted packets, and its information on independencies/dependencies among packets. Note that in this setup, scaling a packet by some scalar element of \mathbb{F}_q does not change the distribution of decoding outcomes.

For given \mathbf{M} , the value of a packet p is specified by the row vector $\mathbf{u}_p = [\mathbf{t}_p, \mathbf{v}_p]$. We call a packet p *secret* if, conditioned on the value of \mathbf{v}_p and the adversary's information, its coefficient vector \mathbf{t}_p is uniformly distributed over $\mathbb{F}_q^r \setminus W$ for some (possibly empty) subspace or affine space $W \subset \mathbb{F}_q^r$.³ Intuitively, secret packets include genuine packets whose coefficient vectors are unknown (in the above sense) to the adversary, as well as packets formed as linear combinations involving at least one secret packet. A set \mathcal{S} of secret packets is *secrecy-independent* if each of the packets remains secret when the adversary is allowed to observe the other packets in the set; otherwise it is *secrecy-dependent*. Secrecy-dependencies arise from the network transmission topology, for instance, if a packet p is formed as a linear combination of a set \mathcal{S} of secret packets (possibly with other non-secret packets), then $\mathcal{S} \cup \{p\}$ is secrecy-dependent.

To illustrate these definitions, suppose that the adversary knows that a sink's decoding set contains an adversarial packet p_1 as well as a packet p_4 formed as some linear combination $k_2 \mathbf{w}_{p_2} + k_3 \mathbf{w}_{p_3}$ of an adversarial packet p_2 with a genuine packet p_3 , so the adversary knows $\mathbf{t}_{p_1}, \mathbf{t}_{p_2}, \mathbf{v}_{p_1}, \mathbf{v}_{p_2}$ and $\mathbf{v}_{p_3} = \mathbf{0}$. Since a decoding set consists of packets with linearly independent coefficient vectors, the adversary knows that \mathbf{t}_{p_1} and \mathbf{t}_{p_3} are linearly independent. Suppose also that the adversary does not observe the contents of any packets dependent on p_3 . Thus, the distribution of \mathbf{t}_{p_4} , conditioned on the adversary's information and any potential value $k_2 \mathbf{v}_{p_2}$ for \mathbf{v}_{p_4} , is uniform over $\mathbb{F}_q^r \setminus \{k \mathbf{t}_{p_1} : k \in \mathbb{F}_q\}$. Also, packets p_3 and p_4 are secrecy-dependent.

Consider a decoding set \mathcal{D} containing one or more secret packets. Choosing an appropriate packet ordering, we can express $[\mathbf{T}_\mathcal{D} | \mathbf{V}_\mathcal{D}]$ in the form

$$[\mathbf{T}_\mathcal{D} | \mathbf{V}_\mathcal{D}] = \left[\begin{array}{c|c} \mathbf{A} + \mathbf{B}_1 & \mathbf{V}_1 \\ \hline \mathbf{C}\mathbf{A} + \mathbf{B}_2 & \mathbf{V}_2 \\ \hline \mathbf{B}_3 & \mathbf{V}_3 \end{array} \right], \quad (2)$$

where for any given values of $\mathbf{B}_i \in \mathbb{F}_q^{s_i \times r}$, $\mathbf{V}_i \in \mathbb{F}_q^{s_i \times (d+c)}$, $i = 1, 2, 3$, and $\mathbf{C} \in \mathbb{F}_q^{s_2 \times s_1}$, the matrix $\mathbf{A} \in \mathbb{F}_q^{s_1 \times r}$ has a conditional distribution that is uniform over all values for which $\mathbf{T}_\mathcal{D}$ is nonsingular. The first $s_1 + s_2$ rows correspond to secret packets, and the first s_1 rows correspond to a set of secrecy-independent packets. $s_2 = 0$ if there are no secrecy-dependencies among the secret packets in \mathcal{D} .

This notion of secret packets provides the most general characterization of the conditions under which the scheme succeeds. For a given network topology, a requirement on the number of secrecy-independent secret packets received at the sink can be translated into constraints on the subsets of links/packets the adversary can observe and/or modify. For instance, if information is sent on n parallel paths from a source to a sink node, then the number of secrecy-independent secret packets is the number of linearly independent packets received on paths that are not observed or controlled by the adversary.

Note that we allow each packet of the decoding set to be corrupted with an independent adversarial packet, as long as at least one of the packets has been formed as a linear combination with some secret packet.

³This definition of a secret packet is conservative as it does not distinguish between packets with a nonuniform conditional distribution and packets that are fully known to the adversary. Taking this distinction into account would make the analysis more complicated but would in some cases give a better bound on detection probability.

III. MAIN RESULTS

In the following theorem, we consider decoding from a set of packets that contains some non-genuine packet, which causes the decoded packets to differ from the original exogenous packets. The first part of the theorem gives a lower bound on the number of equally likely potential values of the decoded packets– the adversary cannot narrow down the set of possible outcomes beyond this regardless of how it designs its adversarial packets. The second part provides, for a simple polynomial hash function, an upper bound on the proportion of potential decoding outcomes that can have consistent data and hash values, in terms of $k = \lceil \frac{d}{c} \rceil$, the ceiling of the ratio of the number of data symbols to hash symbols. Larger values for k correspond to lower overheads but lower probability of detecting an adversarial modification. This tradeoff is a design parameter for the network.

Theorem 1: Consider a decoding set \mathcal{D} containing a secrecy-independent subset of s_1 secret (possibly non-genuine) packets, and suppose the decoding set contains at least one non-genuine packet.

a) The adversary cannot determine which of a set of at least $(q-1)^{s_1}$ equally likely values of the decoded packets will be obtained at the sink. In particular, there will be at least s_1 packets such that, for each of these, the adversary cannot determine which of a set of at least $q-1$ equally likely values will be obtained.

b) Let $\psi : \mathbb{F}_q^k \rightarrow \mathbb{F}_q$ be the function mapping (x_1, \dots, x_k) , $x_i \in \mathbb{F}_q$, to

$$\psi(x_1, \dots, x_k) = x_1^2 + \dots + x_k^{k+1}, \quad (3)$$

where $k = \lceil \frac{d}{c} \rceil$. Suppose the function ψ_d mapping the data symbols x_1, \dots, x_d to the hash symbols y_1, \dots, y_c in an exogenous packet is defined by

$$\begin{aligned} y_i &= \psi(x_{(i-1)k+1}, \dots, x_{ik}) \quad \forall i = 1, \dots, c-1 \\ y_c &= \psi(x_{(c-1)k+1}, \dots, x_d). \end{aligned}$$

Then the probability of not detecting an error is at most $\left(\frac{k+1}{q}\right)^{s_1}$.

Corollary 1: Let the hash function ψ_d be defined as in Theorem 1b. Suppose a sink obtains more than r packets, including a secrecy-independent set of s secret packets, and at least one non-genuine packet. If the sink decodes using two or more decoding sets whose union includes all its received packets, then the probability of not detecting an error is at most $\left(\frac{k+1}{q}\right)^s$.

Example: With 2% overhead ($k = 50$), symbol length=7 bits, $s = 5$, the detection probability is at least 98.9%; with 1% overhead ($k = 100$), symbol length=8 bits, $s = 5$, the detection probability is at least 99.0%.

IV. DEVELOPMENT, PROOFS AND ANCILLARY RESULTS

A. Vulnerable scenarios

Before analyzing the scenario described in the previous sections, we first point out when this approach fails to detect adversarial modifications.

First, the sink needs some way of knowing if the source stops transmitting, otherwise the assumption of no shared secret information results in the adversary being indistinguishable from the source. One possibility is that the source either transmits at a known rate or is inactive, and that the sink knows at what rates it should be receiving information on various subsets of incoming links when the source is active. If the adversary is unable to reproduce those information rates, e.g. because it does not control the same part of the network as the source, then the sink knows when the source is inactive.

Second, if the adversary knows that the genuine packets received at a sink have coefficient vectors that lie in some w -dimensional subspace $W \subset \mathbb{F}_q^r$, the following strategy allows it to control the decoding outcome and so ensure that the decoded packets have consistent data and hash values.

The adversary ensures that the sink receives w genuine packets with linearly independent coefficient vectors in W , by supplying additional such packets if necessary. The adversary also supplies the sink with $r-w$ non-genuine packets whose coefficient vectors $\mathbf{t}_1, \dots, \mathbf{t}_{r-w}$ are not in W . Let $\mathbf{t}_{r-w+1}, \dots, \mathbf{t}_r$ be a set of basis vectors for W , and let \mathbf{T} be the matrix whose i th row is \mathbf{t}_i . Then the coefficient vectors of the r packets can be represented by the rows of the matrix

$$\left[\begin{array}{c|c} \mathbf{I} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{K} \end{array} \right] \mathbf{T},$$

where \mathbf{K} is a nonsingular matrix in $\mathbb{F}_q^{w \times w}$. From (5), we have

$$\begin{aligned} \left[\begin{array}{c|c} \mathbf{I} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{K} \end{array} \right] \mathbf{T} \tilde{\mathbf{M}} &= \left[\begin{array}{c} \tilde{\mathbf{V}} \\ \mathbf{0} \end{array} \right] \\ \tilde{\mathbf{M}} &= \mathbf{T}^{-1} \left[\begin{array}{c|c} \mathbf{I} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{K}^{-1} \end{array} \right] \left[\begin{array}{c} \tilde{\mathbf{V}} \\ \mathbf{0} \end{array} \right] \\ &= \mathbf{T}^{-1} \left[\begin{array}{c} \tilde{\mathbf{V}} \\ \mathbf{0} \end{array} \right]. \end{aligned}$$

Since the adversary knows \mathbf{T} and controls $\tilde{\mathbf{V}}$, it can determine $\tilde{\mathbf{M}}$.

B. Byzantine modification detection

We next consider the scenario described in Section II, where the adversary designs its packets without knowing the contents of one or more secret packets the receiver will use for decoding, and prove the results of Section III.

We first establish two results that are used in the proof of Theorem 1. Consider the hash function defined in (3). We call a vector $(x_1, \dots, x_{k+1}) \in \mathbb{F}_q^{k+1}$ *consistent* if $x_{k+1} = \psi(x_1, \dots, x_k)$.

Lemma 1: At most $k + 1$ out of the q vectors in a set

$$\{\mathbf{u} + \gamma\mathbf{v} : \gamma \in \mathbb{F}_q\},$$

where $\mathbf{u} = (u_1, \dots, u_{k+1})$ is a fixed vector in \mathbb{F}_q^{k+1} and $\mathbf{v} = (v_1, \dots, v_{k+1})$ is a fixed nonzero vector in \mathbb{F}_q^{k+1} , can be consistent.

Proof: Suppose some vector $\mathbf{u} + \gamma\mathbf{v}$ is consistent, i.e.

$$u_{k+1} + \gamma v_{k+1} = (u_1 + \gamma v_1)^2 + \dots + (u_k + \gamma v_k)^{k+1}. \quad (4)$$

Note that for any fixed value of \mathbf{u} and any fixed nonzero value of \mathbf{v} , (4) is a polynomial equation in γ of degree equal to $1 + \tilde{k}$, where $\tilde{k} \in [1, k]$ is the highest index for which the corresponding $v_{k'}$ is nonzero, i.e. $v_{\tilde{k}} \neq 0, v_{k'} = 0 \forall k' > \tilde{k}$. By the fundamental theorem of algebra, this equation can have at most $1 + \tilde{k} \leq 1 + k$ roots. Thus, the property can be satisfied for at most $1 + k$ values of γ . ■

Corollary 2: Let \mathbf{u} be a fixed row vector in \mathbb{F}_q^n and \mathbf{Y} a fixed nonzero matrix in $\mathbb{F}_q^{n \times (k+1)}$. If row vector \mathbf{g} is distributed uniformly over \mathbb{F}_q^n , then the vector $\mathbf{u} + \mathbf{g}\mathbf{Y}$ is consistent with probability at most $\frac{k+1}{q}$.

Proof: Suppose the i th row of \mathbf{Y} , denoted \mathbf{y}_i , is nonzero. We can partition the set of possible values for \mathbf{g} such that each partition consists of all vectors that differ only in the i th entry g_i . For each partition, the corresponding set of values of $\mathbf{u} + \mathbf{g}\mathbf{Y}$ is of the form $\{\mathbf{u}' + g_i\mathbf{y}_i : g_i \in \mathbb{F}_q\}$. The result follows from Lemma 1 and the fact that g_i is uniformly distributed over \mathbb{F}_q . ■

Proof of Theorem 1: We condition on any given values of $\mathbf{B}_i, \mathbf{V}_i, i = 1, 2, 3$, and \mathbf{C} in (2). Writing $\mathbf{A}' = \mathbf{A} + \mathbf{B}_1$, $\mathbf{T}_{\mathcal{D}}$ becomes

$$\begin{bmatrix} \mathbf{A}' \\ \mathbf{C}(\mathbf{A}' - \mathbf{B}_1) + \mathbf{B}_2 \\ \mathbf{B}_3 \end{bmatrix}.$$

From (1), we have

$$\begin{aligned} \begin{bmatrix} \mathbf{A}' \\ \mathbf{C}(\mathbf{A}' - \mathbf{B}_1) + \mathbf{B}_2 \\ \mathbf{B}_3 \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{V}_2 \\ \mathbf{V}_3 \end{bmatrix} \\ \begin{bmatrix} \mathbf{A}' \\ -\mathbf{C}\mathbf{B}_1 + \mathbf{B}_2 \\ \mathbf{B}_3 \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{V}_2 - \mathbf{C}\mathbf{V}_1 \\ \mathbf{V}_3 \end{bmatrix} \end{aligned}$$

which we can simplify to

$$\begin{bmatrix} \mathbf{A}' \\ \mathbf{B}' \end{bmatrix} \tilde{\mathbf{M}} = \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{V}'_2 \end{bmatrix} \quad (5)$$

by writing

$$\mathbf{B}' = \begin{bmatrix} -\mathbf{C}\mathbf{B}_1 + \mathbf{B}_2 \\ \mathbf{B}_3 \end{bmatrix}, \quad \mathbf{V}'_2 = \begin{bmatrix} \mathbf{V}_2 - \mathbf{C}\mathbf{V}_1 \\ \mathbf{V}_3 \end{bmatrix}.$$

Since the determinant of a matrix is not changed by adding a multiple of one row to another row, and $\begin{bmatrix} \mathbf{A}' \\ \mathbf{B}' \end{bmatrix}$ is obtained from $\mathbf{T}_{\mathcal{D}}$ by a sequence of such operations, we have

$$\begin{bmatrix} \mathbf{A}' \\ \mathbf{B}' \end{bmatrix} \text{ is nonsingular} \Leftrightarrow \mathbf{T}_{\mathcal{D}} \text{ is nonsingular.}$$

Thus, matrix $\mathbf{A}' \in \mathbb{F}_q^{s_1 \times r}$ has a conditional distribution that is uniform over the set \mathcal{A} of values for which $\begin{bmatrix} \mathbf{A}' \\ \mathbf{B}' \end{bmatrix}$ is nonsingular.

The condition that the decoding set contains at least one non-genuine packet corresponds to the condition $\mathbf{V}_{\mathcal{D}} \neq \mathbf{0}$. We consider two cases. In each case we show that we can partition the set \mathcal{A} such that at most a fraction $\left(\frac{k+1}{q}\right)^{s_1}$ of values in each partition give decoding outcomes $\mathbf{M} + \tilde{\mathbf{M}}$ with consistent data and hash values. The result then follows since the conditional distribution of values within each partition is uniform.

Case 1: $\mathbf{V}'_2 \neq \mathbf{0}$. Let \mathbf{v}_i be some nonzero row of \mathbf{V}'_2 , and \mathbf{b}_i the corresponding row of \mathbf{B}' . Then $\mathbf{b}_i \tilde{\mathbf{M}} = \mathbf{v}_i$.

We first partition \mathcal{A} into cosets

$$\mathcal{A}_n = \{\mathbf{A}_n + \mathbf{r}^T \mathbf{b}_i : \mathbf{r} \in \mathbb{F}_q^{s_1}\}, n = 1, 2, \dots, \chi$$

where

$$\chi = \frac{|\mathcal{A}|}{q^{s_1}}.$$

This can be done by the following procedure. Any element of \mathcal{A} can be chosen as \mathbf{A}_1 . Matrices $\mathbf{A}_2, \mathbf{A}_3, \dots, \mathbf{A}_\chi$ are chosen sequentially; for each $m = 2, \dots, \chi$, \mathbf{A}_m is chosen to be any element of \mathcal{A} not in the cosets $\mathcal{A}_n, n < m$. Note that this forms a partition of \mathcal{A} , since the presence of some element c in two sets \mathcal{A}_n and $\mathcal{A}_m, n < m$, implies that \mathbf{A}_m is also in \mathcal{A}_n , which is a contradiction. It is also clear that each coset has size $|\{\mathbf{r} : \mathbf{r} \in \mathbb{F}_q^{s_1}\}| = q^{s_1}$.

For each such coset \mathcal{A}_n , the corresponding values of $\tilde{\mathbf{M}}$ satisfy, from (5),

$$\begin{aligned} \begin{bmatrix} \mathbf{A}_n + \mathbf{r}^T \mathbf{b}_i \\ \mathbf{B}' \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{V}'_2 \end{bmatrix} \\ \begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{V}_1 - \mathbf{r}^T \mathbf{v}_i \\ \mathbf{V}'_2 \end{bmatrix} \\ \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{V}_1 - \mathbf{r}^T \mathbf{v}_i \\ \mathbf{V}'_2 \end{bmatrix}, \end{aligned}$$

where $\mathbf{r} \in \mathbb{F}_q^{s_1}$. Let \mathbf{U} be the submatrix consisting of the first s_1 columns of $\begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix}^{-1}$. Since \mathbf{U} is full rank, we can find a set $\mathcal{J} \subset \{1, \dots, r\}$ of s_1 indexes that correspond to linearly independent rows of \mathbf{U} . Let $[\mathbf{U}_1 \mid \mathbf{U}_2]$ be the $s_1 \times r$ submatrix of $\begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix}^{-1}$ consisting of rows with indexes in \mathcal{J} . Consider the corresponding rows of $\mathbf{M} + \tilde{\mathbf{M}}$, which can be expressed in the form

$$\mathbf{M}_{\mathcal{J}} + \mathbf{U}_1 \mathbf{V}_1 - \mathbf{U}_1 \mathbf{r}^T \mathbf{v}_i + \mathbf{U}_2 \mathbf{V}'_2 \quad (6)$$

where $\mathbf{M}_{\mathcal{J}}$ is the submatrix of \mathbf{M} consisting of rows corresponding to set \mathcal{J} . Since \mathbf{U}_1 is non-singular by the choice of \mathcal{J} , $\mathbf{U}_1 \mathbf{r}^T$ takes potentially any value in $\mathbb{F}_q^{s_1}$. Thus, the set of potential values for each row of (6), for any given value of $\mathbf{M}_{\mathcal{J}}, \mathbf{A}_n, \mathbf{B}', \mathbf{V}_1, \mathbf{V}'_2, \mathbf{v}_i$ and the other rows, is of the form $\{\mathbf{u} + \gamma \mathbf{v}_i : \gamma \in \mathbb{F}_q\}$ where \mathbf{u} is a function of $\mathbf{M}_{\mathcal{J}}, \mathbf{A}_n, \mathbf{B}', \mathbf{V}_1, \mathbf{V}'_2$. Applying Lemma 1 yields the result for this case.

Case 2: $\mathbf{V}'_2 = \mathbf{0}$, i.e. $\mathbf{V}_2 - \mathbf{C}\mathbf{V}_1 = \mathbf{V}_3 = \mathbf{0}$. Then $\mathbf{V}_1 \neq \mathbf{0}$, since otherwise $\mathbf{V}_1 = \mathbf{V}_2 = \mathbf{0}$ and $\mathbf{V}_D = \mathbf{0}$ which would contradict the assumption that there is at least one non-genuine packet.

We partition \mathcal{A} such that each partition consists of all matrices in \mathcal{A} that have the same row space:

$$\mathcal{A}_n = \{\mathbf{R}\mathbf{A}_n : \mathbf{R} \in \mathbb{F}_q^{s_1 \times s_1}, \det(\mathbf{R}) \neq 0\}, n = 1, 2, \dots, \chi$$

where

$$|\mathcal{A}_n| = \prod_{i=0}^{s_1-1} (q^{s_1} - q^i), \chi = \frac{|\mathcal{A}|}{|\mathcal{A}_n|}.$$

This can be done by choosing any element of \mathcal{A} as \mathbf{A}_1 , and choosing $\mathbf{A}_n, n = 2, \dots, \chi$ sequentially such that \mathbf{A}_n is any element of \mathcal{A} not in $\mathcal{A}_m, m < n$.

For each $\mathcal{A}_n, n = 1, \dots, \chi$, the corresponding values of $\tilde{\mathbf{M}}$ satisfy, from (5),

$$\begin{aligned} \begin{bmatrix} \mathbf{R}\mathbf{A}_n \\ \mathbf{B}' \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{V}_1 \\ \mathbf{0} \end{bmatrix} \\ \begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix} \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{R}^{-1} \mathbf{V}_1 \\ \mathbf{0} \end{bmatrix} \\ \tilde{\mathbf{M}} &= \begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{R}^{-1} \mathbf{V}_1 \\ \mathbf{0} \end{bmatrix}. \end{aligned}$$

Let \mathbf{U} be the submatrix consisting of the first s_1 columns of $\begin{bmatrix} \mathbf{A}_n \\ \mathbf{B}' \end{bmatrix}^{-1}$. We can find an ordered set $\mathcal{J} = \{i_1, \dots, i_{s_1} : i_1 < \dots < i_{s_1}\} \subset \{1, \dots, r\}$ of s_1 indexes that correspond to linearly independent rows of \mathbf{U} . Let $\mathbf{U}_{\mathcal{J}}$ and $\mathbf{M}_{\mathcal{J}}$ be the submatrices of \mathbf{U} and \mathbf{M} respectively consisting of the s_1 rows corresponding to \mathcal{J} . Then $\mathbf{U}_{\mathcal{J}}$ is nonsingular, and the value of the matrix representation of the corresponding decoded packets is uniformly distributed over the set

$$\{\mathbf{M}_{\mathcal{J}} + \mathbf{R}' \mathbf{V}_1 : \mathbf{R}' \in \mathbb{F}_q^{s_1 \times s_1}, \det(\mathbf{R}') \neq 0\}. \quad (7)$$

Let ν be the rank of \mathbf{V}_1 . Consider a set of ν linearly independent rows of \mathbf{V}_1 . Denote by \mathcal{I} the corresponding set of row indexes, and denote by $\mathbf{V}_{\mathcal{I}}$ the submatrix of \mathbf{V}_1 consisting of those rows. We can write

$$\mathbf{V}_1 = \mathbf{L}\mathbf{V}_{\mathcal{I}}$$

where $\mathbf{L} \in \mathbb{F}_q^{s_1 \times \nu}$ has full rank ν . We define $\mathbf{R}_{\mathcal{I}} = \mathbf{R}'\mathbf{L}$, noting that

$$\mathbf{R}_{\mathcal{I}}\mathbf{V}_{\mathcal{I}} = \mathbf{R}'\mathbf{L}\mathbf{V}_{\mathcal{I}} = \mathbf{R}'\mathbf{V}_1$$

and that $\mathbf{R}_{\mathcal{I}}$ is uniformly distributed over all matrices in $\mathbb{F}_q^{s_1 \times \nu}$ that have full rank ν . Thus, (7) becomes

$$\{\mathbf{M}_{\mathcal{J}} + \mathbf{R}_{\mathcal{I}}\mathbf{V}_{\mathcal{I}} : \mathbf{R}_{\mathcal{I}} \in \mathbb{F}_q^{s_1 \times \nu}, \text{rank}(\mathbf{R}_{\mathcal{I}}) = \nu\}. \quad (8)$$

Denote by $\mathbf{r}_1, \dots, \mathbf{r}_{s_1}$ the rows of $\mathbf{R}_{\mathcal{I}}$, and by \mathbf{R}_n the submatrix of $\mathbf{R}_{\mathcal{I}}$ consisting of its first n rows. We consider the rows sequentially, starting with the first row \mathbf{r}_1 . For $n = 1, \dots, s_1$, we will show that conditioned on any given value of \mathbf{R}_{n-1} , the probability that the i_n th decoded packet $\mathbf{M}_{i_n} + \mathbf{r}_n\mathbf{V}_{\mathcal{I}}$ is consistent is at most $\frac{k+1}{q}$.

Case A: \mathbf{R}_{n-1} has zero rank. This is the case if $n = 1$, or if $n > 1$ and $\mathbf{R}_{n-1} = \mathbf{0}$.

Suppose we remove the restriction $\text{rank}(\mathbf{R}_{\mathcal{I}}) = \nu$, so that \mathbf{r}_n is uniformly distributed over \mathbb{F}_q^ν . By Corollary 2, $\mathbf{m}_{i_n} + \mathbf{r}_n\mathbf{V}_{\mathcal{I}}$ would have consistent data and hash values with probability at most $\frac{k+1}{q}$. With the restriction $\text{rank}(\mathbf{R}_{\mathcal{I}}) = \nu$, the probability of \mathbf{r}_n being equal to $\mathbf{0}$ is lowered. Since the corresponding decoded packet $\mathbf{m}_{i_n} + \mathbf{r}_n\mathbf{V}_{\mathcal{I}}$ is consistent for $\mathbf{r}_n = \mathbf{0}$, the probability that it is consistent is less than $\left(\frac{k+1}{q}\right)$.

Case B: $n > 1$ and \mathbf{R}_{n-1} has nonzero rank.

Conditioned on \mathbf{r}_n being in the row space of \mathbf{R}_{n-1} , $\mathbf{r}_n = \mathbf{g}\mathbf{R}_{n-1}$ where \mathbf{g} is uniformly distributed over \mathbb{F}_q^{n-1} . Since $\mathbf{V}_{\mathcal{I}}$ has linearly independent rows, $\mathbf{R}_{n-1}\mathbf{V}_{\mathcal{I}} \neq \mathbf{0}$, and by Corollary 2, the corresponding decoded packet

$$\mathbf{m}_{i_n} + \mathbf{r}_n\mathbf{V}_{\mathcal{I}} = \mathbf{m}_{i_n} + \mathbf{g}\mathbf{R}_{n-1}\mathbf{V}_{\mathcal{I}}$$

is consistent with probability at most $\frac{k+1}{q}$.

Conditioned on \mathbf{r}_n not being in the row space of \mathbf{R}_{n-1} , we can partition the set of possible values for \mathbf{r}_n into cosets

$$\{\mathbf{r} + \mathbf{g}\mathbf{R}_{n-1} : \mathbf{g} \in \mathbb{F}_q^{n-1}\}$$

where \mathbf{r} is not in the row space of \mathbf{R}_{n-1} ; the corresponding values of the i_n th decoded packet are given by

$$\{\mathbf{m}_{i_n} + \mathbf{r}\mathbf{V}_{\mathcal{I}} + \mathbf{g}\mathbf{R}_{n-1}\mathbf{V}_{\mathcal{I}} : \mathbf{g} \in \mathbb{F}_q^{n-1}\}.$$

Noting as before that $\mathbf{R}_{n-1}\mathbf{V}_{\mathcal{I}} \neq \mathbf{0}$ and applying Corollary 2, the i_n th decoded packet is consistent with probability at most $\frac{k+1}{q}$. ■

Proof of Corollary 1: Suppose two or more different sets of packets are used for decoding. If not all of them contain at least one non-genuine packet, the decoded values obtained from different decoding sets will differ: sets containing only genuine packets will be decoded to \mathbf{M} , while sets containing one or more non-genuine packets will not. This will indicate an error.

Otherwise, suppose all the decoding sets contain at least one non-genuine packet. Let S denote the set of s secrecy-independent packets. Consider the decoding sets in turn, denoting by s'_i the number of unmodified packets from S in the i th decoding set that are not in any set $j < i$. Conditioned on any fixed values of packets in sets $j < i$, there remain s'_i secrecy-independent packets in the i th decoding set, and we have from Theorem 1 that at most a fraction $\left(\frac{k+1}{q}\right)^{s'_i}$ of decoding outcomes for set i have consistent data and hash values. Thus, the overall fraction of consistent decoding outcomes is at most $\left(\frac{k+1}{q}\right)^{\sum_i s'_i} = \left(\frac{k+1}{q}\right)^s$. ■

V. CONCLUSION

This paper has described an information theoretic approach for detecting Byzantine modifications in networks employing random network coding. Byzantine modification detection capability is added by augmenting each packet with a small, flexible number of hash symbols; this overhead can be traded off against the detection probability and symbol length. The hash symbols can be obtained as a simple polynomial function of the data symbols. The only necessary condition is that the adversarial packets are not all designed with knowledge of the random code coefficients of all other packets received by the sink nodes. This approach can be used in conjunction with higher overhead schemes that are activated only upon detection of a Byzantine node.

REFERENCES

- [1] R. Ahlswede, N. Cai, S.-Y.R. Li, and R.W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
- [2] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, October 2003.
- [3] A. F. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of wireless erasure networks," *IEEE Transactions on Information Theory*, vol. 52, pp. 789–804, 2006.
- [4] D. S. Lun, M. Mardar, and M. Effros, "On coding for reliable communication over packet networks," in *In Proc. 42nd Annual Allerton Conference on Communication, Control, and Computing, Invited paper*, September–October 2004.
- [5] T. Ho, R. Koetter, M. Médard, D. R. Karger, and M. Effros, "The benefits of coding over routing in a randomized setting," in *Proc. IEEE International Symposium on Information Theory*, 2003.
- [6] P. A. Chou, Y. Wu, and K. Jain, "Practical network coding," in *Proc. 41st Annual Allerton Conference on Communication, Control, and Computing*, 2003.
- [7] Tracey Ho, Ben Leong, R. Koetter, M. Médard, M. Effros, and D. R. Karger, "Byzantine modification detection for multicast networks using randomized network coding," in *Proc. 2004 IEEE Int. Symp. Information Theory (ISIT 2004)*, Chicago, IL, June/July 2004, p. 144.
- [8] R. Perlman, *Network Layer Protocols with Byzantine Robustness*, Ph.D. thesis, Massachusetts Institute of Technology, October 1988.
- [9] D. Dolev, C. Dwork, O. Waarts, and M. Yung, "Perfectly secure message transmission," *Journal of the ACM*, vol. 40, no. 1, pp. 17–47, January 1993.
- [10] D. Malkhi and M. Reiter, "A high-throughput secure reliable multicast protocol," *Journal of Computer Security*, vol. 5, pp. 113–127, 1997.
- [11] P. Papadimitratos and Z.J. Haas, "Secure routing for mobile ad hoc networks," January 2002.
- [12] Y. Desmedt, "Unconditionally private and reliable communication in an untrusted network," in *IEEE Information Theory Workshop on Theory and Practice in Information-Theoretic Security*, October 2005, pp. 38–41.
- [13] J. L. Massey, "Contemporary cryptography: An introduction," *Contemporary Cryptology: The Science of Information Integrity*, pp. 1–39, 1991.
- [14] N. Cai and Raymond W. Yeung, "Secure network coding," in *Proc. 2002 IEEE Int. Symp. Information Theory (ISIT 2002)*, Lausanne, Switzerland, June/July 2002, p. 323.
- [15] J. Feldman, T. Malkin, C. Stein, and R. A. Servedio, "On the capacity of secure network coding," in *Proc. 42nd Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sept./Oct. 2004.
- [16] K. Bhattad and K. R. Nayayanan, "Weakly secure network coding," in *Proc. WINMEE, RAWNET and NETCOD 2005 Workshops*, Riva del Garda, Italy, Apr. 2005.
- [17] Raymond W. Yeung and N. Cai, "Network error correction, part I: Basic concepts and upper bounds," *Commun. Inf. Syst.*, vol. 6, no. 1, pp. 19–36, 2006.
- [18] N. Cai and Raymond W. Yeung, "Network error correction, part II: Lower bounds," *Commun. Inf. Syst.*, vol. 6, no. 1, pp. 37–54, 2006.
- [19] K.K. Chi and X.M. Wang, "Analysis of network error correction based on network coding," *IEE Proc. -Communications*, vol. 152(4), 2005.
- [20] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Médard, "Resilient network coding in the presence of byzantine adversaries," in *Proc. IEEE INFOCOM 2007*, Anchorage, AK, May 2007.
- [21] R. Koetter and F. R. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Transactions on Information Theory*, 2007, submitted.
- [22] M. Castro and B. Liskov, "Practical byzantine fault tolerance," in *OSDI: Symposium on Operating Systems Design and Implementation*. 1999, USENIX Association, Co-sponsored by IEEE TCOS and ACM SIGOPS.
- [23] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith, "The SecureRing protocols for securing group communication," in *Proc. 31st Annual Hawaii International Conference on System Sciences (HICSS)*. 1998, vol. 3, pp. 317–326, IEEE Computer Society Press.
- [24] T. Ho, M. Médard, R. Koetter, D.R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4413–4430, October 2006.
- [25] D. S. Lun, N. Ratnakar, M. Médard, R. Koetter, D. R. Karger, T. Ho, E. Ahmed, and F. Zhao, "Minimum-cost multicast over coded packet networks," *IEEE Transactions on Information Theory*, to appear, 2006.