

Correction of Adversarial Errors in Networks

Sidharth Jaggi*, Michael Langberg†, Tracey Ho,* and Michelle Effros*

*Department of Electrical Engineering,
California Institute of Technology,
Pasadena, CA 91125, USA,

Email: {jaggi,effros}@caltech.edu, trace@mit.edu

† Department of Computer Science,
California Institute of Technology,
Pasadena, CA 91125, USA,

Email: mikel@caltech.edu

Abstract—We design codes to transmit information over a network, some subset of which is controlled by a malicious adversary. The computationally unbounded, hidden adversary knows the message to be transmitted, and can observe and change information over the part of the network he controls. The network nodes do not share resources such as shared randomness or a private key. We first consider a unicast problem in a network with $|\mathcal{E}|$ parallel, unit-capacity, directed edges. The rate-region has two parts. If the adversary controls a fraction $p < 0.5$ of the $|\mathcal{E}|$ edges, the maximal throughput equals $(1-p)|\mathcal{E}|$. We describe low-complexity codes that achieve this rate-region. We then extend these results to investigate more general multicast problems in directed, acyclic networks.

I. INTRODUCTION

Consider the following point-to-point adversarial channel coding problem. The network \mathcal{G} consists of $|\mathcal{E}|$ parallel, directed, binary-input, binary-output edges $\mathcal{E} = \{e_1, e_2, \dots, e_{|\mathcal{E}|}\}$ between the source s with encoder Xavier and the sink t with decoder Yvonne. At time i , source s generates R Bernoulli- $(1/2)$ random bits $\{X_{i,j}\}_{j=1}^R$. Encoder Xavier wishes to describe these bits across the network. Xavier also has access to a fair coin, which he can use to generate as many bits as he wants. Xavier and Yvonne share no private key or common randomness. Xavier wishes to transmit all the information generated by s to Yvonne, who wishes to decode the received message with asymptotically negligible probability of error. Xavier and Yvonne agree on low-complexity encoding and decoding schemes in advance. The encoding and decoding schemes are also known to the computationally unbounded adversary Zorba. The adversary Zorba knows the message generated by s but not the outcomes of Xavier's coin flips. Zorba can also see and control the transmissions on $\mathcal{Z} \subseteq \mathcal{E}$, where \mathcal{Z} has size M ; Zorba cannot observe or change transmissions on $\mathcal{E} - \mathcal{Z}$. Zorba wishes to minimize the rate R at which Yvonne can reconstruct the information from s with asymptotically negligible probability of error. Zorba's interference patterns on the links he controls can be based only on the knowledge he already possesses (code-design, source message, and causal knowledge of symbols transmitted on links he controls). We note, however, that the rate region is identical even if Zorba has non-causal knowledge of the information transmitted on links he controls. In this paper, for simplicity, we concentrate on *causal* adversaries and leave the details when considering non-causal adversaries to [1].

Previous work [2] exhibits a low-complexity algorithm for each receiver to detect an adversarial attack with high

probability as long as there is at least one packet in the network whose contents the adversary cannot infer.

We obtain an intriguing two-part rate-region for the corresponding error-correction problem. We construct low-complexity block codes, which for large enough block-length n asymptotically achieve the capacity $C_{Adv}(M, |\mathcal{E}|) = (|\mathcal{E}| - M)1(M/|\mathcal{E}| < 0.5)$ of this channel model. Indicator function $1(\cdot)$ is one when its argument is true and zero otherwise. Viewing the ratio $M/|\mathcal{E}|$ as the noise parameter of this adversarial channel, the capacity of the channel for the regime $M/|\mathcal{E}| < 0.5$ equals $|\mathcal{E}|(1 - M/|\mathcal{E}|)$. That is, it equals the capacity of $|\mathcal{E}|$ parallel binary erasure channels (BECs) with erasure probability $M/|\mathcal{E}|$. This result is striking since the location of all erasures is explicitly known to the decoder of an erasure channel whereas \mathcal{Z} is unknown to Yvonne. Indeed, our code construction relies on BEC channel codes. The construction also employs parity information, which enables Yvonne to estimate with high reliability the subset $\mathcal{J}_Z \subseteq \mathcal{Z}$ that Zorba corrupts. Yvonne then decodes the messages on $\mathcal{E} - \mathcal{J}_Z$. Conversely, we show that no matter which code Xavier uses, if he transmits at a rate higher than $C_{Adv}(M, |\mathcal{E}|)$, then there exists a strategy by which Zorba can force Yvonne's probability of decoding error to be bounded away from 0.

Section II we present our results for the case where the network consists of parallel edges. These set the stage for the more interesting multicast model of Section III. In general non-trivial coding needs to be performed at internal nodes in order to achieve the multicast capacity [3]. This makes error-correction harder than in the parallel link case, since in principle the information injected into the network by an adversary controlling even a single link can contaminate all of the information reaching any receiver.

In Section IV treats generalizations. These include allowing small amounts of feedback, which increases the rate-region to $(|\mathcal{E}| - M)$, and knowledge at the sinks of the adversary's location, which enlarges the rate-region to $(|\mathcal{E}| - M)$. In contrast, knowledge by the source of the adversary's location leaves the region unchanged. We also show a separation between channel and network coding for this problem. That is, if the links in the network in addition to possible adversarial interference also have random noise, then network coding to combat the adversary's actions can be overlaid on top of link-by-link channel coding. We also provide an algorithm for detecting which edges need to be removed from the network so as to eliminate the contamination from the information being

injected by adversaries. Lastly, we consider the case where the adversary does not know the message at the source, and show that the maximal rate at which secret information can be embedded in an information-theoretically secure manner into the message being transmitted equals $(1-2p)1(M/|\mathcal{E}| < 0.5)$.

II. UNICAST MODEL

We start with results for the parallel-edge unicast model.

We block both source bits and Xavier's random coin flips into m -dimensional vectors that we treat as elements of the finite field \mathbb{F}_q , where $q = 2^m$. The source input vector is $X = (X(1), X(2), \dots, X(n))^T$, where each $X(i)$ vector comprises R elements of \mathbb{F}_q . Thus $X(1), \dots, X(n)$ represents the R source bits from the first mn units of time. The m -vector of random coin outcomes is $\rho \in \mathbb{F}_q$.

A code \mathcal{C} is defined by its encoder $\{f_e\}_{e \in \mathcal{E}}$ and decoder h . For each $e \in \mathcal{E}$, $f_e : (\mathbb{F}_q)^{nR} \times \mathbb{F}_q \rightarrow (\mathbb{F}_q)^n$ maps a source vector X and random symbol ρ to the length- n vector $Y(e) = Y(e, i) = f_e(X, \rho)$ transmitted across edge e . We use $Y = f(X, \rho) = (f_e(X, \rho))_{e \in \mathcal{E}}$ to denote the full channel input and $\hat{Y} = (\hat{Y}(e))_{e \in \mathcal{E}}$ to describe the full channel output. In particular, we use the length- $|\mathcal{E}|$ vector $Y(i)$ to denote the channel input at time i and $\hat{Y}(i)$ to denote the channel output at time i . A decoder $h : (\mathbb{F}_q)^{n|\mathcal{E}|} \rightarrow (\mathbb{F}_q)^{nR}$ maps the collection \hat{Y} of received channel outputs to a reconstruction $\hat{X} = h(\hat{Y})$ of source message X .

Xavier and Yvonne together choose a code $\mathcal{C} = ((f_e)_{e \in \mathcal{E}}, h)$. This code choice is fixed and known to Zorba, who also has full knowledge of the source message X to be transmitted. Zorba uses this information to choose the jamming function g used to corrupt the channel input Y to give channel output \hat{Y} . In designing his jamming function, Zorba first chooses a set \mathcal{Z} of edges to control. The size of \mathcal{Z} cannot exceed M , the jamming dimension. For each $e \in \mathcal{E} - \mathcal{Z}$, $\hat{Y}(e) = Y(e)$. For each $e \in \mathcal{Z}$, Zorba uses jamming functions $g_e : (\mathbb{F}_q)^{nR} \times (\mathbb{F}_q)^{nM} \rightarrow (\mathbb{F}_q)^n$ to produce $\hat{Y}(e) = g_e(X, (Y(e))_{e \in \mathcal{Z}})$; thus the corrupted information on any edge $e \in \mathcal{Z}$ can rely on both the source message X and causally on the channel inputs $Y(e)$ on all edges $e \in \mathcal{Z}$. For notational simplicity we henceforth write $\hat{Y} = g(X, Y)$ to denote the full collection of channel outputs.

The error probability is defined as $P_e^{(n)} = \Pr[h(g(X, (Y(e))_{e \in \mathcal{Z}}))] \neq X$. Rate R is achievable for the channel g for jamming dimension M if for any $\epsilon > 0$ and n sufficiently large there exists a blocklength- n code \mathcal{C} with $P_e^{(n)} < \epsilon$ for every jamming function g in the family of jamming functions described above. The capacity $C_{Adv}(M, |\mathcal{E}|)$ equals the maximal achievable rate over all g .

We now state and prove our main result for unicast channels.

Theorem 1:

$$C_{Adv}(M, |\mathcal{E}|) = (|\mathcal{E}| - M)1(M/|\mathcal{E}| < 0.5)$$

Further, for any n and any $m = \omega(\log(n|\mathcal{E}|))$ there exists block-length n codes with $R = (1 - (|\mathcal{E}| + 1)/n)C_{Adv}(M, |\mathcal{E}|)$, $P_e^{(n)} < n|\mathcal{E}|2^{-m}$, and complexity of design and encoding and decoding implementation equal to $\mathcal{O}((nm|\mathcal{E}|)^2)$.

Proof: Upper Bounds: The bound $R \leq |\mathcal{E}| - M$ is immediate since Zorba can set $\hat{Y}(e) = 0^n$ for all $e \in \mathcal{Z}$, thereby giving rate zero on all edges controlled by Zorba.

If $M \geq |\mathcal{E}|/2$, $R = 0$ since Zorba can use the following strategy to make decoding with $P_e^{(n)} < 1/2$ impossible. Zorba selects an arbitrary jamming subset \mathcal{J}_Z of size $|\mathcal{E}| - M$ of \mathcal{Z} . Then, for arbitrary $X' \neq X$ and ρ , Zorba sets $\hat{Y}(e) = f_e(X', \rho)$ for each $e \in \mathcal{J}_Z$ and $\hat{Y}(e) = 0^n$ for $e \in \mathcal{Z} - \mathcal{J}_Z$. Yvonne does not know \mathcal{Z} and is therefore unable to decide which of X and X' to decode to, leading to a probability of error of at least $1/2$.

Lower Bound: We first sketch the achievability argument and then give a precise code construction. Assume $M/|\mathcal{E}| < 1/2$ and $R = |\mathcal{E}| - M$. In the first $n - |\mathcal{E}| - 1$ symbols on each $e \in \mathcal{E}$ Xavier transmits X using an erasure code. Xavier uses the remaining $\mathcal{E} + 1$ symbols to send a header containing ρ and $D = (D(e))_{e \in \mathcal{E}}$. The vector D is a hash of the vectors $(Y(i))_{i=1}^{n-|\mathcal{E}|-1}$ with ρ . Yvonne decodes by looking for consistency among the received channel outputs. Since Zorba controls fewer than half of the edges, Yvonne can determine (ρ, D) using a majority rule. She then recomputes the hash using ρ and the received transmissions. Since Zorba does not know ρ a priori, any changes he makes on $(y(e))_{e \in \mathcal{Z}}$ will with high probability be inconsistent with the hash values. This enables Yvonne to determine which edges have been corrupted. She then uses $\hat{Y}(e)$ from $e \notin \mathcal{J}_Z$ to reconstruct \hat{X} , via the erasure code.

We now describe our coding scheme in detail. For any n and $m = \omega(\log(n|\mathcal{E}|))$, fix $R = \lfloor (1 - (|\mathcal{E}| + 1)/n)(|\mathcal{E}| - M) \rfloor$ and design the functions f_e using the following procedure.

Let \mathcal{L} be any $(n - |\mathcal{E}| - 1)|\mathcal{E}| \times nR$ Vandermonde matrix over \mathbb{F}_q (such a matrix exists since $q = 2^m$ and $m = \omega(\log(n|\mathcal{E}|))$) [4]. For the i th edge $e_i \in \mathcal{E}$, the matrix $\mathcal{L}(e_i)$, known a priori to Xavier, Yvonne, and Zorba, is defined to be the $(n - |\mathcal{E}| - 1) \times nR$ matrix consisting of row $[(n - |\mathcal{E}| - 1)(i - 1) + 1]$ through $[(n - |\mathcal{E}| - 1)i]$ of \mathcal{L} . For all $e \in \mathcal{E}$ we define $T(e)$, U and D as

$$\begin{aligned} T(e) &= (\mathcal{L}(e)X)^T, \\ U &= (1, \rho, \dots, \rho^{n-|\mathcal{E}|-1}) \text{ and} \\ D &= U[T(e_1) \dots T(e_{|\mathcal{E}|})]. \end{aligned}$$

and set $Y(e) = [T(e), D, \rho]$. Thus for each $e \in \mathcal{E}$, the first $n - |\mathcal{E}| - 1$ symbols in $Y(e)$ are the erasure-coded message symbols, the next $|\mathcal{E}|$ symbols are the hash function output, and the last symbol is the hash-function's key ρ .

Yvonne's decoding scheme h is as follows. Let $\hat{Y}(e) = (\hat{T}(e), \hat{D}, \hat{\rho})$ denote the channel output on $e \in \mathcal{E}$. As described above, Yvonne first determines the correct value of the header (D, ρ) by choosing the value that appears on the majority of the links. She then checks, for the i th edge $e_i \in \mathcal{E}$, whether $D(e_i)$ equals the i th symbol in $U(\hat{T}(e_1) \dots \hat{T}(e_{|\mathcal{E}|}))$. She calls the set of edges for which this is true \mathcal{E}_D .

In the second stage of decoding Yvonne constructs \mathcal{L}_D , a $|\mathcal{E}_D|(n - |\mathcal{E}| - 1) \times nR$ matrix created by concatenating the matrices in $\{\mathcal{L}(e)\}_{e \in \mathcal{E}_D}$. Since \mathcal{L} is a Vandermonde matrix,

so is \mathcal{L}_D . Yvonne obtains \hat{X} by inverting the matrix equation $\hat{Y}_D = \mathcal{L}_D X$, where \hat{Y}_D is the dimension $|\mathcal{E}_D|(n - |\mathcal{E}| - 1)$ vector obtained by the ordered concatenation of $\hat{Y}(e)$, $e \in \mathcal{E}_D$. There is a decoding error only if $\mathcal{E}_J \cap \mathcal{E}_D \neq \emptyset$, where $\mathcal{E}_J \subseteq \mathcal{Z}$ is the set of edges for which $\hat{T}(e) \neq T(e)$. We next bound the probability of this event.

It suffices to prove that with probability $1 - n|\mathcal{E}|2^{-m}$, for no $e_i \in \mathcal{E}_J$ is $D(e_i) = U\hat{T}(e_i)$ true. By definition of \mathcal{E}_J , $T(e) \neq \hat{T}(e)$. Thus $U(T(e))^T = U(\hat{T}(e))^T$, i.e., $U(T(e) - \hat{T}(e))^T = 0$, only if ρ is a root of the degree $n - |\mathcal{E}| - 1$ polynomial $U(T(e) - \hat{T}(e))^T$. Zorba does not know the value of ρ , and the polynomial contains at most $n - |\mathcal{E}| - 1$ roots in the field of size 2^m . Therefore $e \notin \mathcal{E}_D$ are inconsistent with probability at least $1 - (n - |\mathcal{E}| - 1)/2^m$. Since there are fewer than $|\mathcal{E}|/2$ edges in \mathcal{E}_J , the total probability that $\mathcal{E}_J \cap \mathcal{E}_D \neq \emptyset$ is at most $(n - |\mathcal{E}| - 1)|\mathcal{E}|/2^{m+1} < n|\mathcal{E}|2^{-m}$.

Lastly, it can be verified that the complexity of encoder f_e at each edge e is determined by the complexity of computing the vectors $T(e)$ over a field of size q , and that the complexity of decoder h is determined by the complexity of inverting a Vandermonde matrix of dimension nR over the same finite field [4]. \square

Note 1: Most Maximum Distance Separable codes [5] can be used in place of \mathcal{L} . We choose Vandermonde matrices due to their low design and implementation complexity.

Note 2: Even if Zorba's jamming functions g_e are allowed to violate causality, a result with identical rate-regions and codes with similar parameters to those in Theorem 1 still holds. A proof using a variant of Verifiable Secret Sharing ([6],[7]) can be found in [1].

III. MULTICAST MODEL

We now examine the problem of multicasting information on more complex networks with a hidden adversary. We assume that $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed acyclic network with unit-capacity directed edges. For a node $v \in \mathcal{V}$, $\Gamma_O(v)$ denotes set of the edges outgoing from v and $\Gamma_I(v)$ denotes the set of edges entering v . An edge originating at vertex v and terminating at vertex v' is said to have *tail* v (denoted by $v = v_t(e)$) and *head* v' (denoted by $v' = v_h(e)$). The encoder Xavier at the source node s uses the network \mathcal{G} to transmit the source's information $X = (X(1), X(2), \dots, X(n))^T$ as defined in Section II to a set of decoders, $\{Y_{\text{vonne}}_1, \dots, Y_{\text{vonne}}_{|T|}\}$, located respectively at the sink nodes $T = \{t_1, \dots, t_{|T|}\}$. Xavier uses $M|T|$ random m -vectors, denoted $\rho = (\rho_k(i))_{i \in \{1, \dots, M\}, k \in \{1, \dots, |T|\}}$.

A (v, v', S) -cut between any $v, v' \in \mathcal{V}$ is a partition of \mathcal{V} into S and $\mathcal{V} - S$ such that $v \in S$ and $v' \in \mathcal{V} - S$. The *cut edge-set* $\mathcal{E}(v, v', S)$ comprises $e \in \mathcal{E}$ such that $v_t(e) \in S$ and $v_h(e) \in \mathcal{V} - S$. The *min-cut capacity* [3] is defined as $C_{Mul} = \min_{k \in \{1, \dots, |T|\}} \min_S |\mathcal{E}(s, t_k, S)|$, and any such minimizing $\mathcal{E}(s, t_k, S)$ -cut is called a *min-cut*.

A *network code* \mathcal{C}_N is defined by its source encoder, internal encoders, and decoders at receiver nodes.

The source encoder comprises a collection of functions $\{f_e\}_{e \in \Gamma_O(s)}$. For each $e \in \Gamma_O(s)$, $f_e : (\mathbb{F}_q)^{nR} \times (\mathbb{F}_q)^{M|T|} \rightarrow$

$(\mathbb{F}_q)^n$ maps X and a set of random symbols $\rho_k(i)$ to the length n vector $Y(e)$ transmitted across edge e . We denote by $Y(e, i)$ the i th symbol input to $Y(e)$ for each $e \in \mathcal{E}$, and denote by $\hat{Y}(e, i)$ the i th symbol output on edge e .

The internal encoders for all edges $e \notin \Gamma_O(s)$ are functions $f_e : (\mathbb{F}_q)^{n|\Gamma_I(v_t(e))|} \rightarrow (\mathbb{F}_q)^n$ which map messages $Y(e')$ on all links e' incoming to $v_t(e)$ to the vector $Y(e)$ transmitted across edge e .

For each $k \in \{1, \dots, |T|\}$, decoder $h_k : (\mathbb{F}_q)^{n|\Gamma_I(t_k)|} \rightarrow (\mathbb{F}_q)^{nR}$ maps the collection $\hat{Y}_k = (\hat{Y}(e, i))_{e \in \Gamma_I(t_k), i \in \{1, \dots, n\}}$ of received channel outputs to a reconstruction \hat{X}_k of source X . In particular, denote $\hat{Y}_k(i) = (\hat{Y}(e, i))_{e \in \Gamma_I(t_k)}$ and $\hat{Y}_k(e) = (\hat{Y}(e, i))_{i \in \{1, \dots, n\}}$.

Xavier and the Yvonnees together choose a code $\mathcal{C} = ((f_e)_{e \in \mathcal{E}}, (h_k)_{k \in \{1, \dots, |T|\}})$, and inform each e of f_e . This code choice is fixed and known to Zorba, who also has full knowledge of the source message X to be transmitted. Zorba uses this information to choose the set $\mathcal{Z} \subseteq \mathcal{E}$ of edges to control. The size of \mathcal{Z} cannot exceed the jamming dimension M . We note that adversarial control of a vertex $v \in \mathcal{V}$ is equivalent to adversarial control of all edges in $\Gamma_O(v)$, and therefore we only consider the case where Zorba controls edges.

For each $e \in \mathcal{E}$ we use $\hat{Y}(e)$ to describe the channel output of link e received by node $v_h(e)$. For each $e \notin \mathcal{Z}$ $\hat{Y}(e) = Y(e)$. For each $e \in \mathcal{Z}$, Zorba uses jamming function $g_e : (\mathbb{F}_q)^{nR} \times (\mathbb{F}_q)^{nM} \rightarrow (\mathbb{F}_q)^n$ to produce corrupted output $\hat{Y}(e) = g_e(X, (Y(e))_{e \in \mathcal{Z}})$. As in the unicast cast, the g_e s are required to be causal in $(Y(e))_{e \in \mathcal{Z}}$, and defer the discussion of the corresponding theorem for the case with non-causal jamming functions to [1].

The error probability is defined as $P_e^{(n)} = \Pr[\exists k \text{ such that } h_k(g(X, (Y(e))_{e \in \mathcal{Z}})) \neq X]$. Rate R is achievable for jamming dimension M if for any $\epsilon > 0$ and n sufficiently large there exists a blocklength- n code \mathcal{C} with $P_e^{(n)} < \epsilon$ for every jamming function g in the family of jamming functions described above. The capacity $C_{Adv, Mul}(M, |\mathcal{E}|)$ of the given adversarial channel model equals the maximal achievable rate.

Our key observations and ideas are as follows.

For each $v \in \mathcal{V}$ the encoding functions $\{f_e\}_{e \in \Gamma_O(v)}$ perform approximately n rounds of a robust algebraic network code ([8], [9]). In the i th round the input to this algebraic network code is $X(i)$. After these rounds of transmitting the actual information, s transmits to each receiver in succession $M(R + 1)$ symbols of header information using C_{Mul} edge-disjoint paths.

We model the effect of the jamming functions as follows. Let \mathcal{G}_Z be the graph obtained by attaching a new unit-rate source node $s_{e,z}$ to the midpoint of e for each $e \in \mathcal{Z}$. The message $X(e, i)$ generated over the i th time interval by $s_{e,z}$ may be an arbitrary function of X and $(Y(e', i'))_{e' \in \mathcal{Z}, i' \leq i}$. For each $e \in \mathcal{Z}$ the link output $\hat{Y}(e)$ equals $Y(e) + X(e, i)$. Denote by $X_Z(i)$ the length- M vector $(X(e, i))_{e \in \mathcal{Z}}$.

Since the set \mathcal{Z} is fixed and \mathcal{C} is linear, for each $k \in$

$\{1, \dots, |T|\}$

$$\hat{Y}_k(i) = T_k X(i) + T_{Z,k} X_Z(i) \quad (1)$$

for some fixed linear transforms T_k and $T_{Z,k}$. We define the *interference at t_k* as $\delta_k(i) = T_{Z,k} X_Z(i)$. The linear span of $\{X(i)\}_{i \in \{1, \dots, n-M|T|(C_{Mul}+1)\}}$ is a vector-space (denoted V_X) of dimension at most R . Denote by $T_k V_X$ the linear span of $\{T_k X(i)\}_{i \in \{1, \dots, n-M|T|(C_{Mul}+1)\}}$. The linear span of $\{X_Z(i)\}_{i \in \{1, \dots, n-M|T|(C_{Mul}+1)\}}$ is a vector-space (denoted by V_Z) of dimension at most M . Denote by $T_{Z,k} V_Z$ the linear span over i of $\{\delta_k(i)\}_{i \in \{1, \dots, n-M|T|(C_{Mul}+1)\}}$. By a direct corollary of [10, Theorem 4], with high probability over code design $X(i)$ is retrievable from $T_k X(i)$, and $V_X \cap V_Z$ equals only the zero vector (indeed, this is the property that implies the existence of robust network codes). This implies that if $T_{Z,k} V_Z$ is known to Yvonne _{k} , then $X(i)$ is recoverable for all i . (In contrast to the unicast case, she cannot here infer the set \mathcal{E}_j . She can, however, cancel out the interference effect. Theorem 5 in Section IV shows a scheme for detecting the set of edges which need to be cut to isolate \mathcal{Z} from the network.) To ascertain $T_{Z,k} V_Z$, we use a scheme similar to the one developed in Section II.

Theorem 2:

$$C_{Adv, Mul}(M, C_{Mul}) = (C_{Mul} - M)1(M/C_{Mul} < 0.5)$$

Further, for any n and any $m = \omega(\log(n|\mathcal{E}||T|))$, there exist block-length n codes with $R = (1 - M|T|(C_{Mul} + 1)/n)C_{Adv, Mul}(M, C_{Mul})$, $P_e^{(n)} < n|\mathcal{E}|2^{-m(C_{Mul}-R)}$, complexity of design and encoding $O(nm)$, and of decoding equal to $\mathcal{O}(nmC_{Mul})^3$.

Proof: Upper bound: The bound $R \leq C_{Mul} - M$ follows since Zorba can choose \mathcal{Z} to be in a cutset, and set $\hat{Y}(e) = 0^n$ for all $e \in \mathcal{Z}$.

If $M > C_{Mul}/2$ edges, $R = 0$ by the following argument. Zorba chooses \mathcal{Z} to be a subset of some min-cut $\mathcal{E}(s, t, S)$, and an arbitrary *multicast jamming subset* $\mathcal{J}_Z \subset \mathcal{Z}$ of size $C_{Mul} - M$. For an arbitrary $X' \neq X$ and an arbitrary ρ' Zorba mimics the network code \mathcal{C} and for each $e \in \mathcal{J}_Z$ sets $\hat{Y}(e)$ to what the message would have been on \mathcal{J}_Z if s had input (X', ρ') , and $\hat{Y}(e) = 0^n \forall e \in \mathcal{Z} - \mathcal{J}_Z$. As in Theorem 1, Yvonne _{k} is unable to decide which X and X' to decode to.

Lower bound: We present a coding strategy using ideas from the proof of Theorem 1.

Let $R = (1 - M|T|(C_{Mul} + 1)/n)(C_{Mul} - M)$ and $m = \Theta(\log(n|T||\mathcal{E}|))$. We show the existence of codes that achieve R with $P_e^{(n)} < 2^{-\Omega(m(C_{Mul}-R))}$. There are two encoding steps.

First, Xavier uses a robust network code $n - M|T|(C_{Mul} + 1)$ times to multicast information to each t_k . The input to \mathcal{C} during the i th use is $X(i)$.

In the second step, the paths $\{P_{i,k}\}_{i \in \{1, \dots, C_{Mul}\}, k \in T}$ (where $\{P_{i,k}\}_{i \in \{1, \dots, C_{Mul}\}}$ comprise C_{Mul} edge-disjoint paths from s to t_k) are used to transmit identical copies of the header information. This header information consists of M blocks,

each of length $R + 1$, for each receiver. Since there are $|T|$ receivers, this process is at most $(R + 1)M|T|$ channel uses over \mathbb{F}_q . The header information sent to t_k is $(D_k(j), \rho_k(j))_{j=1}^M$. That is, each of the M blocks of length $R + 1$ in the header to t_k contains the random symbol $\rho_k(i)$, and the length- R hash-vector $D_k(i)$. Each hash-vector $D_k(j)$ is a distinct linear combination of the $X(i)$ s, defined as

$$D_k(j) = \sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} X(i).$$

For any k Zorba may control edges in less than half of $\{P_{i,k}\}_{i=1}^{C_{Mul}}$, hence the header information each sink receives on more than half the paths is identical. At each t_k Yvonne _{k} retrieves $(D_k(j), \rho_k(j))_{j=1}^M$ by a majority decision.

Decoding by Yvonne _{k} proceeds as follows. For all $j \in \{1, \dots, M\}$ she computes the vectors $T_k D_k(j)$ and the vectors $\sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} \hat{Y}_k(i)$. Using (1) we have

$$\begin{aligned} & \sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} \hat{Y}_k(i) \\ &= T_k D_k(j) + T_{Z,k} \left(\sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} X_Z(i) \right) \\ &= T_k D_k(j) + \sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} \delta_k(i). \end{aligned}$$

Hence Yvonne can retrieve M length- R vectors in $T_{Z,k} V_Z$, namely $\sum_{i=1}^{n-(R+1)M|T|} \rho_k(j)^{i-1} \delta_Z(i)$, denoted respectively by $A_k(j)$. We now prove that with high probability, for each $k \in \{1, \dots, |T|\}$, $\{A_k(j)\}_{j \in \{1, \dots, M\}}$ forms a basis for $T_{Z,k} V_Z$.

We denote by $[\Delta]$ the matrix which has $\delta_k(i)$ s as row vectors. Let $U_k(i) = (\rho_k(i)^{j-1})_{j=1}^{n-(R+1)M|T|}$. We denote by $[U]$ the matrix which has $U_k(i)$ as row vectors. We note that since Zorba controls at most M links, $rank([\Delta])$ is at most M . We choose $[\Delta']$ to be any set of $rank([\Delta])$ linearly independent columns of $[\Delta]$. Suppose that $\{A_k(i)\}_{i=1}^{C_{Mul}}$ does not form a basis for $T_{Z,k} V_Z$. This means that for some linear combination $c_k = (c_k(1) \dots c_k(rank([\Delta'])))$ the length- M column vector $[U][\Delta']c_k$ equals the zero vector, though the column vector $[\Delta']c_k$ is non-zero (since by definition $[\Delta']$ has full column rank). Thus the adversary would have to choose the matrix $[\Delta]$ so that the M polynomials which are the elements of the column vector $[U][\Delta']c_k$ are all zero. By a similar argument as in Theorem 1, the probability that this happens is $(n/q)^M$. Taking the union bound over all receivers, the total probability of error equals $|T|(n/q)^M$. \square

IV. VARIATIONS ON THE THEME

We now analyze various related models.

1. Suppose, in addition to the conditions described in Section III, Xavier (but none of the Yvonne)s knows \mathcal{Z} , we denote the capacity by $C_{\mathcal{Z} \rightarrow X}(M, C_{Mul})$. Alternatively, if all of the Yvonne)s (but not Xavier) know \mathcal{Z} , we denote the capacity by

$C_{Z \rightarrow Y}(M, C_{Mul})$. The following shows knowledge of Z at $t_k s$ is more useful than knowledge of Z at s .

Theorem 3:

$$C_{Z \rightarrow X}(M, C_{Mul}) = (C_{Mul} - M)1(M/C_{Mul} < 0.5)$$

$$C_{Z \rightarrow Y}(M, C_{Mul}) = (C_{Mul} - M)$$

Sketch of Proof: Both $C_{Z \rightarrow X}(M, C_{Mul})$ and $C_{Z \rightarrow Y}(M, C_{Mul})$ must be at least as large as $C_{Adv, Mul}(M, C_{Mul})$ since Xavier and Yvonne_k can still follow the same strategy as when neither of them knew Z . For all values of M , if Yvonne_k does not know Z , then Zorba can still follow the same strategy as in the upper bound of Theorem 1, and therefore $C_{Z \rightarrow X}(M, C_{Mul}) = C_{Adv, Mul}(M, C_{Mul})$. However, if Xavier uses f_e as in Theorem 2 and Yvonne_k knows Z , she can with high probability infer $T_{k, Z}$ and cancel the effect of $X_Z(i)$ to decode. Hence $C_{Z \rightarrow Y}(M, C_{Mul}) = C_{Mul} - M$ for all values of M . \square

2. Suppose each $e \in \mathcal{E}$ is noisy and has channel capacity $C_{Noise} < 1$, and Zorba controls some M edges. We denote the overall capacity of this channel by $C_{Adv, Noise}(M, C_{Mul})$.

Theorem 4:

$$C_{Adv, Noise}(M, C_{Mul}) = C_{Noise}C_{Adv, Mul}(M, C_{Mul}).$$

Sketch of Proof: Xavier first uses a channel code to make each e noiseless and then uses the code of Theorem 1. No higher rate is achievable since Zorba can use the same strategy as in the upper bound in Theorem 1. \square

3. Suppose Yvonne_k wishes to find a set of links L which, when removed from the network, would neutralize the effect of Zorba without diminishing the multicast capacity. Let $T_{L, k}$ be the linear dependence between $X(i)$ and $\hat{Y}_k(i)$ on removing L from \mathcal{G} .

Theorem 5: A set of edges $L \subset \mathcal{E}$ such that the network code on the graph $(\mathcal{V}, \mathcal{E})$ induces a network code with the same achievable rate on the graph $(\mathcal{V}, \mathcal{E} - L)$ such that $\hat{Y}_k(i) = T_{L, k}X(i)$ for all k exists and can be determined by each Yvonne_k.

Sketch of Proof: The same codes as in Theorem 2 are used. Each Yvonne_k first determines $T_{L, k}V_Z$, and then sequentially considers all subsets of size M of \mathcal{E} , and sees if any of them induces the transform $T_{L, k}$. She chooses the first such set and calls it L . Due to random code design, with high probability such a choice suffices. \square

4. Suppose we allow secret feedback from each Yvonne_k to Xavier. We denote the capacity of this channel by $C_{Feedback}(M, C_{Mul})$.

Theorem 6:

$$C_{Feedback}(M, C_{Mul}) = C_{Mul} - M.$$

Sketch of Proof: We use essentially the same codes as in Theorem 2. Each Yvonne_k transmits a secret key (not known to Zorba) to Xavier. Instead of transmitting just the header, as in Theorem 2, Xavier signs the header with the secret key using an information-theoretic authentication scheme (e.g., [11]). This enables each Yvonne_k to receive an uncorrupted header even if only a single path from Xavier is uncorrupted. \square

5. Finally, suppose Zorba is unaware of X , and embedded within X is a message X_s that we wish to keep secret from Zorba.

Theorem 7:

$$C_{Adv, Secret}(M, C_{Mul}) = (C_{Mul} - 2M)1(M/C_{Mul} < 0.5).$$

Sketch of Proof: Every set U of links of size $C_{Mul} - M$ in every min-cut must contain enough information to be able to decode X correctly. Therefore the number of links Zorba does not observe in any set U is of size at most $C_{Mul} - 2M$, which proves our upper bound. We use the same codes from Theorem 2 to prove achievability. The null-space of the linear transform which takes X to $\{Y(e)\}_{e \in \mathcal{Z}}$ is of dimension $n(|\mathcal{E}| - 2M)$, which is asymptotically the same as the dimension of the vector space in which the secret message X_s sits. Thus with high probability over network code design X_s will be secret from Zorba. \square

ACKNOWLEDGMENTS

The authors were partially supported by NSF Grants CCR-0325324, CCR-0325673 and CCF-0346991 and a grant from the Lee Center for Advanced Networking at Caltech. We thank Yuval Ishai for pointing out the connection between our results and the known results on VSS.

REFERENCES

- [1] S. Jaggi, M. Langberg, T. Ho, M. Effros, B. Leong, R. Koetter, M. Médard, and D. Karger. Error detection, correction and secrecy: Network coding in the presence of adversaries. *In preparation*, 2005.
- [2] T. C. Ho, B. Leong, R. Koetter, M. Médard, M. Effros, and D. R. Karger. Byzantine modification detection in multicast networks using randomized network coding. *In International Symposium on Information Theory*, 2004.
- [3] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.
- [4] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.
- [5] R. J. McEliece. *The Theory of Information and Coding*, volume 3 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, Reading, Mass., 1977.
- [6] R. Cramer, I. Damgård, and S. Fehr. On the cost of reconstructing a secret, or vss with optimal reconstruction phase. *In Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, pages 503–523, 2001.
- [7] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. *In Proceedings of the Eighth Annual ACM Symposium on Principles of Distributed Computing: Edmonton, Alberta, Canada*, pages 73–85, 1989.
- [8] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 11(5):782–795, October 2003.
- [9] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. submitted to the IEEE Transactions on Information Theory, 2003.
- [10] S. Jaggi, P. A. Chou, and K. Jain. Low complexity algebraic multicast network codes. *In IEEE International Symposium on Information Theory (ISIT)*, page 368, Yokohama, July 2003.
- [11] P. Gemmell and M. Naor. Codes for interactive authentication. *In Proceedings of CRYPTO 93: Lecture Notes in Computer Science*, 773:355–367, 1993.