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

† Department of Computer Science, California Institute of Technology, Pasadena, CA 91125, USA, Email: mikel@caltech.edu

‡ Coordinated Science Laboratory, University of Illinois, Urbana, IL, 61801 USA, Email: trace@mit.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 at the source 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 unicasting over a network with $|\mathcal{E}|$ parallel, unit-capacity, directed edges for which we show a two-part rate region. If the adversary controls a fraction $p < 0.5$ of the $|\mathcal{E}|$ edges, the maximum throughput equals $p|\mathcal{E}|$, and we exhibit low complexity codes that achieve this rate. Conversely, if the adversary controls at least half of the $|\mathcal{E}|$ edges, the maximum throughput equals 0. We extend these results to the case of multicasting over a directed acyclic network and exhibit a similar two part rate-region. We also consider cases when the source or the sink has knowledge of which edges are controlled by the adversary, the links are noisy or have different capacities, feedback from the sink to the source, and a model where the adversary does not know the message to be transmitted from the source to the sink.**

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 $\mathcal{E} = \{e_1, e_2, \ldots, e_{|\mathcal{E}|}\}$ between the source s with encoder Xavier and the sink t with decoder Yvonne. At time j, source s generates R Bernoulli-$(1/2)$ random bits $\{X_{i,j}\}_{i=1}^{R}$. 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 it with asymptotically negligible probability of error. A pre-determined encoding/decoding scheme is agreed upon between Xavier and Yvonne, who can perform low-complexity computations over some finite field $\mathbb{F}_{2^m}$ (m is a design parameter). The encoding/decoding scheme is also known to the computationally unbounded adversary Zorba. Zorba knows the message generated by s (but not the outcomes of Xavier's coin flips), and can see and control the transmissions on $\mathcal{Z} \subset \mathcal{E}$, where $\mathcal{Z}$ has size M, but 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.

We obtain an intriguing two-part rate-region for this 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. Here $1(.)$ is the indicator function. 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 is the same as the capacity if Xavier used a BEC code (Binary Erasure Channel code [14]) and $\mathcal{Z}$ were known to Yvonne. Indeed, our code construction enables Yvonne to estimate with high reliability the subset $\mathcal{J}_Z \subset \mathcal{Z}$ that Zorba uses to corrupt data, and then decode the messages on $\mathcal{E} - \mathcal{J}_Z$. Conversely, we show that for any encoding/decoding scheme if Xavier transmits at any 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.

We extend our results to the problem of multicasting over arbitrary directed acyclic graphs, parts of which may be controlled by an adversary. Previous work [7] exhibits a low-complexity algorithm for each receiver to detect, with high probability an adversarial attack as long as there is at least one packet in the network whose contents the adversary cannot infer. We consider the corresponding error-correction problem and give a tight characterization of the achievable rates at which each sink can reconstruct, with high reliability, all the information from s. We also provide a corresponding encoding/decoding scheme.

Previous work on adversarial channels considers models where Zorba can change an arbitrary fraction of the n-bit channel input [6], for which the tightest known bounds are by [12]. Other works consider cases when the source and sink share a short private key [11] or common randomness [2], or when Zorba is computationally bounded [3], [13].

In Section II we present our results for the case when the network consists of parallel edges. In Section III we consider variants of this model. In Section IV we present our main result, on multicasting over an arbitrary directed acyclic graph, some fraction of which is controlled by an adversary.

II. UNICAST MODEL

In the description that follows, we block bits into m-dimensional vectors that we treat as elements of the finite field $\mathbb{F}_{2^m}$. The coding algorithm takes R length-n blocks of these finite field elements as its input. Let $x = (x(1), x(2), \ldots, x(R))^T$, where $x(i) = (x(i,1), x(i,2), \ldots, x(i,n))$ and $x(i,k) = (X_{i,(k-1)m+1}, \ldots, X_{i,km}) \in \mathbb{F}_{2^m}$. Similarly, the bits from Xavier's random coin are blocked into m-vectors, which are likewise treated as elements of $\mathbb{F}_{2^m}$. Xavier uses $|\mathcal{E}|$ such vectors, here denoted by $r = (r(e))_{e\in\mathcal{E}} \in (\mathbb{F}_{2^m})^{|\mathcal{E}|}$.

A code $\mathcal{C}$ is defined by its encoder and decoder. An encoder comprises a collection of functions $\{f_e\}_{e\in\mathcal{E}}$. For each $e \in \mathcal{E}$, $f_e : (\mathbb{F}_{2^m})^{nR} \times \mathbb{F}_{2^m} \rightarrow (\mathbb{F}_{2^m})^n$ maps a source vector x and random symbol $r(e)$ to the vector $y(e) = f_e(x, r(e))$ transmitted across edge e. We use $y = f(x, r)$ to denote the full channel input.

A decoder $h : (\mathbb{F}_{2^m})^{n|\mathcal{E}|} \rightarrow (\mathbb{F}_{2^m})^{nR}$ maps the collection $(\hat{y}(e))_{e\in\mathcal{E}}$ of received channel outputs to a reconstruction $\hat{x} = h(\hat{y})$ of source 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 x to be transmitted. Zorba uses this information to choose the set $\mathcal{Z}$ of edges to control and the jamming functions $\{g_e\}_{e\in\mathcal{Z}}$ (defined below) to apply on these edges. The size of $\mathcal{Z}$ is less than or equal to M. We call M the jamming dimension.

We use $\hat{y} = (\hat{y}(e))_{e\in\mathcal{E}}$ to describe the channel output received by Yvonne. For each $e \in \mathcal{E} - \mathcal{Z}$, $\hat{y}(e) = y(e)$. For each $e \in \mathcal{Z}$, Zorba uses jamming function $g_e : (\mathbb{F}_{2^m})^{nR} \times (\mathbb{F}_{2^m})^{nM} \rightarrow (\mathbb{F}_{2^m})^n$ to produce $\hat{y}(e) = g_e(x, (y(e))_{e\in\mathcal{Z}})$.

The error probability is defined as $P_e^{(n)} = \mathbf{Pr}[h(g(f(x, r)))] \neq x)$. 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. The capacity $C_{Adv}(M, |\mathcal{E}|)$ of the given adversarial channel model equals the maximal achievable rate.

Theorem 1:

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

Further, for any n and any $\Delta_n = \omega(\log(n|\mathcal{E}|))$, there exist block-length n codes with $R = (1 - \Delta_n/n)(|\mathcal{E}| - M)$, $P_e^{(n)} < 2^{-\Omega(\Delta_n)}$, and complexity of design and implementation $\mathcal{O}((n\Delta_n|\mathcal{E}|)^2)$ for both f_e and h.

Proof: **Upper Bound:** If $M/|\mathcal{E}| < 1/2$, $R > |\mathcal{E}| - M$ is not achievable since for all $e \in \mathcal{Z}$ Zorba can set $\hat{y}_e = 0^n$ (i.e., regardless of x Zorba can transmit the all-zero vector on all edges in $\mathcal{Z}$).

If $M \geq |\mathcal{E}|/2$, $R > 0$ is not achievable since Zorba may use the following strategy. Zorba selects an arbitrary *jamming subset* $\mathcal{J}_Z$ of size $|\mathcal{E}| - M$ of $\mathcal{Z}$. Then, for an arbitrary $x' \neq x$ and an arbitrary r', Zorba sets $\hat{y}(e) = f_e(x', r'(e))$ for $e \in \mathcal{J}_Z$ 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: For the regime $M/|\mathcal{E}| < 1/2$ we design codes achieving $R = |\mathcal{E}| - M$. The main underlying idea is that each $y(e)$ contains a small secret header with $r(e)$ and $|\mathcal{E}|$ symbols $D(e, e')$, where for each $e' \in \mathcal{E}$, $D(e, e')$ is a hash of $r(e)$ with $y(e')$. Roughly speaking, Yvonne decodes by checking the headers of all pairs $(y(e), y(e'))$ for mutual consistency. Since Zorba does not know $(r(e))_{e\in\mathcal{E}-\mathcal{Z}}$, the headers in $y(e)_{e\in\mathcal{E}-\mathcal{Z}}$ enable Yvonne to detect with high probability any changes made on the edges in $\mathcal{Z}$. Yvonne then uses $y(e)$ only from $e \notin \mathcal{Z}$ to reconstruct $\hat{x}$.

We now describe our coding scheme in detail. Xavier and Yvonne first decide on Δ_n and n so as to trade off (using the equations in Theorem 1) the probability of error, rate and complexity of design and implementation according to their own preferences, and fix $R = \lfloor(1 - \Delta_n/n)(|\mathcal{E}| - M)\rfloor$. They then choose $m = \Theta(\Delta_n)$ 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* [5] over $\mathbb{F}_{2^m}$ (since $m = \omega(\log(n|\mathcal{E}|))$ such a matrix exists). For each $e_i \in \mathcal{E}$, the matrix $\mathcal{L}(e_i)$, known *a priori* to Xavier, Yvonne and Zorba, is defined as the $(n - |\mathcal{E}| - 1) \times nR$ matrix consisting of the $[(n - |\mathcal{E}| - 1)(i - 1) + 1]$th to the $[(n-|\mathcal{E}|-1)i]$th rows of $\mathcal{L}$. For each $e \in \mathcal{E}$ we denote by $T(e)$ the matrix product $(\mathcal{L}(e)x)^T$. For all $e, e' \in \mathcal{E}$ we denote by $D(e, e')$ the dot-product $u(e)(T(e'))^T$, where $u(e)$ is defined as $(1, r(e), \ldots, r(e)^{n-|\mathcal{E}|-1})$. For each $e \in \mathcal{E}$, $y(e)$ is then set equal to $(T(e), r(e), (D(e, e'))_{e\in\mathcal{E}})$. Thus for each $e \in \mathcal{E}$, the first $n-|\mathcal{E}|-1$ symbols of $y(e)$ comprise of $T(e) = (\mathcal{L}(e)x)^T$, the next symbol equals $r(e)$, and the remaining $|\mathcal{E}|$ symbols are the dot-products $D(e, e') = u(e)(T(e'))^T$.

Yvonne's decoding scheme h is as follows. Denote $\hat{y}(e)$ by $\left(\hat{T}(e), \hat{r}(e), (\hat{D}(e, e'))_{e'\in\mathcal{E}}\right)$ (the first $n - |\mathcal{E}| - 1$ symbols comprise $\hat{T}(e)$, the next symbol $\hat{r}(e)$, and the remaining symbols respectively $\hat{D}(e, e')$). Let $\mathcal{E}_J$ be the set of edges e in $\mathcal{E}$ for which $\hat{T}(e) \neq T(e)$. In the first stage of decoding, Yvonne constructs $\mathcal{E}_D$, a *large* set of edges which with high probability does not contain any edges in $\mathcal{E}_J$. $\mathcal{E}_D$ will then be used to reconstruct the value of the transmitted information x (we will need $\mathcal{E}_D$ to be of size at least $|\mathcal{E}| - M$). To do this Yvonne builds a *consistency graph* $\mathcal{G}_C$. For each $e \in \mathcal{E}$ there is a vertex $v(e) \in \mathcal{G}_C$, and two vertices $v(e), v(e') \in \mathcal{G}_C$ are connected if e and e' are *consistent*. Two edges e and e' are said to be consistent if $\hat{y}(e)$ and $\hat{y}(e')$ satisfy $\hat{D}(e, e') = \hat{u}(e)(\hat{T}(e'))^T$ and $\hat{D}(e', e) = \hat{u}(e')(\hat{T}(e))^T$, where Yvonne computes $\hat{u}(e)$ as $(1, \hat{r}(e), \ldots, \hat{r}(e)^{n-|\mathcal{E}|-1})$. Self-loops in $\mathcal{G}_C$ are allowed. The vertex-set of $\mathcal{G}_C$ consists of the union of three disjoint subsets $V_{\mathcal{E}-\mathcal{Z}}$, $\mathcal{E}_J$ and $V_{\mathcal{Z}-J}$, where vertices in $V_{\mathcal{E}-\mathcal{Z}}$ correspond to edges in $\mathcal{E} - \mathcal{Z}$, vertices in V_J correspond to edges in $\mathcal{E}_J$, and vertices in $V_{\mathcal{Z}-J}$ correspond to the remaining edges $\mathcal{Z} - \mathcal{E}_J$. Lemma 2 below proves that with high probability $\mathcal{G}_C$ has a "nice" structure. The *degree of a vertex* v is defined as the number of vertices v' it is connected to; a self-loop contributes 1 to the degree of a vertex.

Lemma 2: The graph $\mathcal{G}_C$ satisfies: (a) each vertex in $V_{\mathcal{E}-\mathcal{Z}}$ has degree at least $|\mathcal{E}| - M$, and (b) with probability $1 - 2^{-\Omega(\Delta_n)}$ no vertex in V_J has degree larger than M.

Proof: It suffices to prove that (a) $|V_{\mathcal{E}-\mathcal{Z}}| = |\mathcal{E}| - M$ and any two vertices in $V_{\mathcal{E}-\mathcal{Z}}$ are connected, and (b) with probability $1-2^{-\Omega(\Delta_n)}$ there are no edges in $\mathcal{G}_C$ connecting vertices from $V_{\mathcal{E}-\mathcal{Z}}$ and V_J.

To prove (a), note that for any $e, e' \in \mathcal{E} - \mathcal{Z}$, $\hat{y}(e) = y(e)$ and $\hat{y}(e') = y(e')$, hence e and e' are consistent and thus $\mathcal{G}_C$ contains an edge between $v(e)$ and $v(e')$. Moreover, by our definitions $|V_{\mathcal{E}-\mathcal{Z}}| = |\mathcal{E}| - M$.

To prove (b), consider any $e \in \mathcal{E} - \mathcal{Z}$, $e' \in \mathcal{E}_J$ and their corresponding vertices $v(e)$ and $v(e')$ in $\mathcal{G}_C$. By definition of $\mathcal{E}_J$ and $\mathcal{E} - \mathcal{Z}$ it follows that $T(e') \neq \hat{T}(e')$, $\hat{u}(e) = u(e)$, and $\hat{D}(e, e') = D(e, e')$. Thus for e and e' to be consistent $u(e)(T(e'))^T = u(e)(\hat{T}(e'))^T$, i.e., $u(e)(T(e') - \hat{T}(e'))^T = 0$, which happens only if $r(e)$ is a root of $u(e)(T(e') - \hat{T}(e'))^T$ viewed as a degree $n - |\mathcal{E}| - 1$ polynomial. Zorba does not know the value of $r(e)$, and the field is of size 2^m and contains at most $n - |\mathcal{E}| - 1$ roots, therefore with probability at least $1 - (n - |\mathcal{E}| - 1))/2^m$, e and e' will be inconsistent. Since there are at most $|\mathcal{E}|^2/4$ such $(e, e') \in (\mathcal{E} - \mathcal{Z}) \times \mathcal{E}_J$, using $m = \Theta(\Delta_n)$ the total probability of there being an edge between $V_{\mathcal{E}-\mathcal{Z}}$ and V_J is $(n - |\mathcal{E}| - 1))|\mathcal{E}|^2/2^{m+2} = 2^{-\Omega(\Delta_n)}$. □

Using Lemma 2, Yvonne constructs the set $\mathcal{E}_D$ (of $e \in \mathcal{E}$ with corresponding vertices in $\mathcal{G}_C$ of degree at least $|\mathcal{E}| - M$). Let $\mathcal{L}_D$ be the $|\mathcal{E}_D|(n - |\mathcal{E}| - 1) \times nR$ matrix whose row vectors are exactly the rows of the matrices in $\{\mathcal{L}(e)\}_{e \in \mathcal{E}_D}$. Since $\mathcal{L}$ is a Vandermonde matrix, so is $\mathcal{L}_D$. Thus Yvonne obtains $\hat{x}$ by inverting the matrix equation $\hat{y}_D = \mathcal{L}_D x$, where $\hat{y}_D$ is the $|\mathcal{E}_D|(n - |\mathcal{E}| - 1)$ length vector obtained by the ordered concatenation of $y(e)$, $e \in \mathcal{E}_D$. There is a decoding error only if $\mathcal{Z} \cap \mathcal{E}_D \neq \phi$, and therefore by Lemma 2 $P_e^{(n)} = 2^{-\Omega(\Delta_n)}$. Lastly, it can be verified that each f_e's complexity is determined by the complexity of computing the vectors $T(e)$ over a field of size $\mathcal{O}(\Delta_n)$, and that of h by the complexity of inverting a Vandermonde matrix of dimension nR over the same finite field ([5]). □

III. Variations on the theme

We now analyze various related models.

1. Suppose, in addition to the conditions described in Section II, Xavier (but not Yvonne) knows $\mathcal{Z}$, we denote the capacity by $C_{Adv,\mathcal{Z}\to X}(M, |\mathcal{E}|)$. If Yvonne (but not Xavier) knows $\mathcal{Z}$, we denote the capacity by $C_{Adv,\mathcal{Z}\to Y}(M, |\mathcal{E}|)$. The following shows that Yvonne knowing $\mathcal{Z}$ is more powerful than Xavier knowing $\mathcal{Z}$.

Theorem 3:

$$\begin{aligned} C_{Adv,\mathcal{Z}\to X}(M, |\mathcal{E}|) &= (|\mathcal{E}| - M).1(M/|\mathcal{E}| < 0.5) \\ C_{Adv,\mathcal{Z}\to Y}(M, |\mathcal{E}|) &= |\mathcal{E}| - M. \end{aligned}$$

Sketch of Proof: Both $C_{Adv,\mathcal{Z}\to X}(M, |\mathcal{E}|)$ and $C_{Adv,\mathcal{Z}\to Y}(M, |\mathcal{E}|)$ must be at least as large as $C_{Adv}(M, |\mathcal{E}|)$, since Xavier or Yvonne can still follow the same strategy as when neither of them knew $\mathcal{Z}$. For all values of M, if Yvonne does not know $\mathcal{Z}$, then Zorba can still follow the same strategy as in the upper bound of Theorem 1, and therefore $C_{Adv,\mathcal{Z}\to X}(M, |\mathcal{E}|) = C_{Adv}(M, |\mathcal{E}|)$. However, if Yvonne knows $\mathcal{Z}$, Xavier uses f_e as in Theorem 1, and Yvonne uses only $(\hat{y}(e))_{\mathcal{E}-\mathcal{Z}}$ to decode, we have that $C_{Adv,\mathcal{Z}\to Y}(M, |\mathcal{E}|) = |\mathcal{E}| - M$ for all values of M. □

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 $C_{Adv,Noise}(M, |\mathcal{E}|)$.

Theorem 4:

$$C_{Adv,Noise}(M, |\mathcal{E}|) = C_{Noise} \cdot C_{Adv}(M, |\mathcal{E}|).$$

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. □

3. Suppose now that each e has a different capacity $C(e)$. Define $C_E = \sum_{e\in\mathcal{E}} C(e)$, and $C_Z = \sum_{e\in\mathcal{Z}} C(e)$. We denote the capacity of this channel by $C_{Adv,Var}(C_Z, C_E)$.

Theorem 5:

$$C_{Adv,Var}(C_Z, C_E) = (C_E - C_Z).1(C_Z/C_E < 0.5)$$

Sketch of Proof: Xavier calculates u, the greatest common divisor of C_i, $\forall i \in \mathcal{E}$. He then divides each edge into equal capacity u edges, and uses the codes devised in Theorem 1. The proof from Theorem 1 then carries over. □.

4. Finally, suppose we allow secret feedback from Yvonne to Xavier. We denote capacity of this channel by $C_{Adv,Feedback}(M, |\mathcal{E}|)$.

Theorem 6:

$$C_{Adv,Feedback}(M, |\mathcal{E}|) = M - |\mathcal{E}|.$$

Sketch of Proof: Yvonne transmits a secret key (not known to Zorba) to Xavier. Xavier encodes x using the Vandermonde matrices described in Theorem 1, then *signs* each message using the secret key and an information theoretic authentication scheme (*e.g.* [4]), and sends the signed messages over the communication edges. Yvonne decodes using only the $\hat{y}(e)$ which pass the authentication test. □.

IV. 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 general directed acyclic network with unit-capacity directed edges. The directed acyclic nature induces the partial orders $\geq_V$ on $\mathcal{V}$, with $s \geq_V v$ for all $v \in \mathcal{V}$; and $\geq_E$ on $\mathcal{E}$, with $\Gamma_O(s) \geq_E e$ for all $e \in \mathcal{E}$. Here, for a node $v \in \mathcal{G}$, $\Gamma_O(v)$ denotes the out going edges of v (we will also use $\Gamma_I(v)$ to denote the 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 is located at the source node s and wishes to use the edges in $\mathcal{E}$ to transmit the source's information x as defined in Section II to a set of decoders, $\{\text{Yvonne}_1, \ldots, \text{Yvonne}_{|T|}\}$, located respectively at the sink nodes $T = \{t_1, \ldots, t_{|T|}\}$. Xavier uses $|\Gamma_O(s)||T|$ random m-vectors, denoted as $r = (r(e,k))_{e\in\Gamma_O(s), k\in\{1,\ldots,|T|\}}$. A (v, v', S)-*cut* between any $v, v' \in V$ is a partition of V into S and $V - S$ such that $v \in S$ and $v' \in V - S$. The *cut edge-set* $\mathcal{E}(v, v', S)$ comprises of $e \in \mathcal{E}$ such that $v_t(e) \in S$ and $v_h(e) \in V - S$. The *min-cut capacity* ([1]) is defined as $C_{Mul} = \min_{k\in\{1,\ldots,|T|\}} \min_S |\mathcal{E}(s, t_k, S)|$ with any such minimizing $\mathcal{E}(s, t_k, S_{MC})$ 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 $e \in \Gamma_O(s)$, $f_e : (\mathbb{F}_{2^m})^{nR} \times (\mathbb{F}_{2^m})^{|T|} \to (\mathbb{F}_{2^m})^n$ maps a source vector x and a set of random symbols

$r(e,k)$, $k \in \{1,\ldots,|T|\}$ to the vector $y(e)$ transmitted across edge e. We use $y = f(x,r)$ to denote the full channel input.

The internal encoders for all edge $e \notin \Gamma_O(s)$ is a function $f_e : (\mathbb{F}_{2^m})^{n|\Gamma_I(v_t(e))|} \to (\mathbb{F}_{2^m})^n$ which maps messages on links e' such that $v_h(e') = v_t(e)$ to the vector $y(e) = F_e(x,r)$ transmitted across edge e. Here F_e is recursively defined in terms of $f_{e'}$ for all $e' \geq_E e$.

For each $k \in \{1,\ldots,|T|\}$ decoder $h_k : (\mathbb{F}_{2^m})^{n|\Gamma_I(t_k)|} \to (\mathbb{F}_{2^m})^{nR}$ maps the collection $(\hat{y}(e))_{e\in\Gamma_I(t_k)}$ of received channel outputs to a reconstruction $\hat{x}_k$ of source x.

The encoder and decoders together choose a code $\mathcal{C} = ((f_e)_{e\in\mathcal{E}}, (h_k)_{k\in\{1,\ldots|T|\}})$. This code choice is fixed and known to Zorba, who also has full knowledge of the source x to be transmitted. Zorba uses this information to choose the set $\mathcal{Z}$ of edges to control and the jamming functions $\{g_e\}_{e\in\mathcal{Z}}$ to apply on these edges. The size of $\mathcal{Z}$ is at most M.

For each $e \in \mathcal{E}$ we use $\hat{y}(e)$ to describe the channel output of link e received by node $v_h(e)$, and $\hat{y}_k = (\hat{y}(e))_{e\in\Gamma_I(t_k)}$ to describe the channel output received by Yvonne$_k$. For each $e \in \mathcal{E} - \mathcal{Z}$, $\hat{y}(e) = y(e)$. For each $e \in \mathcal{Z}$, Zorba uses jamming function $g_e : (\mathbb{F}_{2^m})^{nR} \times (\mathbb{F}_{2^m})^{nM} \to (\mathbb{F}_{2^m})^n$ to produce $\hat{y}(e) = g_e(x, (y(e))_{e\in\mathcal{Z}})$.

The error probability is defined as $P_e^{(n),|T|} = \mathbf{Pr}[\exists k, h_k(g(f(x,r)))] \neq x)$. Rate R_{Mul} is achievable for jamming dimension M if for any $\epsilon > 0$ and n sufficiently large there exists a blocklength-n code $\mathcal{C}_N$ with $P_e^{(n),|T|} < \epsilon$ for every jamming function g. The capacity $C_{Adv,Mul}(M,|\mathcal{E}|)$ of the given adversarial channel model equals the maximal achievable rate.

Our key observation is as follows. Let $\mathcal{G}_Z$ be the graph obtained by removing $e \in \mathcal{Z}$ from $\mathcal{G}$ and attaching a new unit-rate source node $s_{e,z}$ to each $v_h(e)$. If $\mathcal{C}_N$ has the properties of a *robust network code* (*e.g.* [10], [9]), and if $\mathcal{E}_Z$ is known to the decoders, the information from both s and $s_Z = \{s_{e,Z}\}_{e\in\mathcal{Z}}$ injected into $\mathcal{G}$ by Zorba can be completely reconstructed by each decoder. To ascertain $\mathcal{Z}$ we again use the consistency ideas developed in Section II.

Theorem 7:

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

Proof: **Upper bound:** If $M > C_{Mul}/2$ edges, $R_{Mult} > 0$ is not achievable since Zorba may use the following strategy. Let $\mathcal{E}_{MC}(s,t_k,S)$ be a min-cut. Zorba chooses $\mathcal{Z} \subset \mathcal{E}(s,t,S_{MC})$, and an arbitrary *multicast jamming subset* $\mathcal{J}_{MC,Z} \subset \mathcal{Z}$ of size $C_{Mul} - M$. For an arbitrary $x' \neq x$ and an arbitrary r' Zorba sets $\hat{y}(e) = F_e(x',r') \forall e \in \mathcal{J}_{MC,Z}$, i.e., the message which would have been on $\mathcal{J}_{MC,Z}$ if s had input (x',r'), and $\hat{y}(e) = 0^n \forall e \in \mathcal{Z} - \mathcal{J}_{MC,Z}$. As in Theorem 1, Yvonne$_k$ is unable to decide which x and x' to decode to.

If $M/C_{Mul} < 1/2$, $R_{Mul} > C_{Mul} - M$ is not achievable since for all $e \in \mathcal{Z} \subset \mathcal{E}(s,t_k,S_{MC})$ Zorba can set $\hat{y}(e) = 0^n$.

Lower bound: We now present a coding strategy which achieves rate $(C_{Mul} - M).1(M/C_{Mul} < 0.5)$ using ideas from the proof of Theorem 1. We set $R_{Mul} = (1 - \Delta_n/n)(C_{Mul} - M)$ (where $\Delta_n = \omega(\log(nC_{Mul}))$) and $m = \Theta(\Delta_n)$, and show the existence of codes which achieve R_{Mul} with $P_e^{(n),|T|} < 2^{-\Omega(\Delta_n)}|T|$. There are three encoding steps.

First, as in Theorem 1, let $\mathcal{L}_{Mul}$ be any $(n - (C_{Mul} - 1)|T|)C_{Mul} \times nR$ Vandermonde matrix over $\mathbb{F}_{2^m}$ known *a priori* to Xavier, Yvonne and Zorba; $\{\mathcal{L}_{Mul}(i)\}$ for $i \in \{1,\ldots,C_{Mul}\}$ be $(n - (C_{Mul} - 1)|T|) \times nR$ sub-matrices of $\mathcal{L}_{Mul}$ with distinct rows; and $T(i)$ the matrix product $(\mathcal{L}_{Mul}(i)x)^T$. Now, we generate $|T|$ sets of consistency information. For each sink t_k, let $\{P_{i,k}\}_{i=1}^{C_{Mul}}$ be C_{Mul} predetermined edge-disjoint paths from s to t_k. Denote by $r(i,k)$ the symbol $r(e,k)$ such that $e \in \Gamma_O(s)$, $e \in P_{i,k}$. For each $t_k \in t$, we denote by $D(i,i',k)$ the dot-product $u(i,k)(T(i'))^T$ where $u(i)$ is defined as $(1, r(i,k), \ldots, r(i,k)^{n-C_{Mul}-1})$.

In the second step, the paths $P_{i,k}$ are used to transmit the *consistency information* $\left(r(i,k), (D(i,i',k))_{i'=1}^{C_{Mul}}\right)$ to t_k. This step requires $m(C_{Mul} + 1)$ coding instants per t_k, for a total of $m(C_{Mul} + 1)|T|$ coding instants. For any k Zorba may control e in (less than half of) $\{P_{i,k}\}_{i=1}^{C_{Mul}}$, hence each Yvonne$_k$ at t_k receives (partially corrupted) consistency information denoted by $\left(\hat{r}(i,k), (\hat{D}(i,i',k))_{i'=1}^{C_{Mul}}\right)_{i=1}^{C_{Mul}}$.

In the third step, Xavier uses the robust network code described below.

Lemma 8: For any $\mathcal{G}$, any subset of edges $\mathcal{Z}$ and corresponding $\mathcal{G}_Z$, rate C_{Mul} on $\mathcal{G}_Z$ is achievable with a single code $\mathcal{C}_R$. Moreover, the design and implementation complexity of $\mathcal{C}_R$ is $\mathcal{O}(\log(|\mathcal{E}|))$ and $\mathcal{C}_R$ can be constructed randomly with probability of error $2^{-\Omega(n\Delta_n + \log(|\mathcal{E}|.|T|))}$.

Sketch of proof: The graph $\mathcal{G}_Z$ has min-cut capacity at least C_{Mul} between $s \cup s_Z$ and each t_k. Therefore ([1]) rate C_{Mul} is achievable between $s \cup s_Z$ and each receiver. We use random code design to construct such a $\mathcal{C}_R$. Namely, each edge e takes random linear combinations over $\mathbb{F}_{2^m}$ of the $\hat{y}(e')$s incident on $v_t(e)$ to produce $y(e)$. Theorem 4 of [8] guarantees the required performance. □

Since the rate at which adversarial source nodes s_Z generate information is at most M, a rate of $C_{Mul} - M$ is achievable between s and any receiver t_k if $\mathcal{Z}$ was known to each decoder.

Decoding by Yvonne$_k$ proceeds as follows. For each possible $\mathcal{E}' \subset \mathcal{E}$ of size at most M (there are at most $|\mathcal{E}|^M$ of these), Yvonne$_k$ decodes to $\hat{x}_{\mathcal{E}'}$ by guessing $\mathcal{G}_Z = \mathcal{G}_{\mathcal{E}'}$ and using the code $\mathcal{C}_R$. More specifically, the code $\mathcal{C}_R$ is used on the symbols of $(\hat{y}(e))_{e\in\Gamma_I(t_k)}$ corresponding to the third step of the encoding process. If $\mathcal{E}'$ indeed equals $\mathcal{E}_Z$, Lemma 8 proves that $\hat{x}_{\mathcal{E}'} = x$ (w.h.p. over the choice of $\mathcal{C}_R$), but if $\mathcal{E}' \neq \mathcal{Z}$ then $\hat{x}_{\mathcal{E}'}$ may not equal x. We now provide a consistency check similar to the one in Theorem 1, which w.h.p. distinguishes between the cases $\hat{x}_{\mathcal{E}'} = x$ and $\hat{x}_{\mathcal{E}'} \neq x$. This will conclude the proof of our theorem.

As noted earlier, for each k at most M of the received consistency information packets $\left(\hat{r}(i,k), (\hat{D}(i,i',k))_{i'=1}^{C_{Mul}}\right)_{i=1}^{C_{Mul}}$ may have been viewed or corrupted by Zorba. After decoding to $\hat{x}_{\mathcal{E}'}$ Yvonne$_k$ computes $\{\hat{T}(i,k) = (\mathcal{L}_{Mul}(i)x_{\mathcal{E}'})^T\}_{i=1}^{C_{Mul}}$ (where $\mathcal{L}_{Mul}(i)$ are the matrices used by Xavier in the first step of encoding). The information $\left(\hat{T}(i,k), \hat{r}(e,k), (\hat{D}(e,e',k))_{i'=1}^{C_{Mul}}\right)_{i=1}^{C_{Mul}}$ is used to construct a consistency graph $\mathcal{G}_{Mul,C}(k)$ with C_{Mul} vertices $\{v(i)\}_{i\in\{1,\ldots,C_{Mul}\}}$ and an edge between $v(i)$ and $v(i')$ iff i and i' are consistent, i.e., $\hat{D}(i,i',k) = \hat{u}(i,k)(\hat{T}(i',k))^T$ and $\hat{D}(i',i,k) = \hat{u}(i',k)(\hat{T}(i,k))^T$. Here $\hat{u}(i,k) = (1, \hat{r}(i,k), \ldots, (\hat{r}(i,k))^{n-C_{Mul}-1})$ that Yvonne$_k$ computes given $\hat{r}(i,k)$. We now show that if Yvonne$_k$ declares successful decoding when $\mathcal{G}_{Mul,C}(k)$ has at least $C_{Mul}-M$ vertices each of degree at least $C_{Mul}-M$ then w.h.p. $\hat{x}_{\mathcal{E}'} = x$.

Lemma 9: With probability greater than $1 - |\mathcal{E}|^M 2^{-\Delta_n}$ $\mathcal{G}_{Mul,C}(k)$ has at least $C_{Mul}-M$ vertices each of degree at least $C_{Mul}-M$ iff $\hat{x}_{\mathcal{E}'} = x$.

Proof: To prove one direction note that for all k, $\hat{T}(i,k) = T(i,k)$ for all $i = 1,\ldots,C_{Mul}$; and the consistency information satisfies $\left(\hat{r}(i,k), (\hat{D}(i,i',k))_{i'=1}^{C_{Mul}}\right) = \left(r(i,k), (D(i,i',k))_{i'=1}^{C_{Mul}}\right)$ for at least $C_{Mul}-M$ indices i, which are therefore pairwise consistent.

To prove the other direction assume $\hat{x}_{\mathcal{E}'} \neq x$. Define the following four subsets of vertices of $\mathcal{G}_{Mul,C}(k)$ whose union equals the entire vertex-set of $\mathcal{G}_{Mul,C}(k)$.

$$A = \left\{v(i)|\hat{T}(i,k) = T(i,k) \text{ and } P_{i,k}\cap\mathcal{Z} = \phi\right\}.$$
$$B = \left\{v(i)|\hat{T}(i,k) = T(i,k) \text{ and } P_{i,k}\cap\mathcal{Z} \neq \phi\right\}.$$
$$C = \left\{v(i)|\hat{T}(i,k) \neq T(i,k) \text{ and } P_{i,k}\cap\mathcal{Z} = \phi\right\}.$$
$$D = \left\{v(i)|\hat{T}(i,k) \neq T(i,k) \text{ and } P_{i,k}\cap\mathcal{Z} \neq \phi\right\}.$$

We now use three observations to prove our assertion. First we note that $|A\cup B| < C_{Mul}-M$. If not, let $\mathcal{L}_{Mul,A\cup B}$ be the matrix whose rows are exactly the rows of the matrices $\{\mathcal{L}_{Mul}(i)\}_{v(i)\in A\cup B}$. By construction $\mathcal{L}_{Mul,A\cup B}$ has full rank and therefore one can obtain x by using the inverse transformation. This contradicts the assumption that $x_{\mathcal{E}'} \neq x$. Secondly, we note that $|B\cup D| \leq M$, as for all $v(i)\in B\cup D$, the information on path $P_{i,k}$ was viewed by Zorba, which by assumption can occur on at most M indices i. Finally, we notice that w.h.p. over random $r(i,k)$ there are no edges in $\mathcal{G}_{Mul,C}(k)$ between vertices $v(i)$ for which Zorba does not control $P_{i,k}$, and vertices $v(i')$ for which $\hat{T}(i',k) \neq T(i',k)$. This is because for such (i,i'), as analyzed in Lemma 2, $\hat{D}(i,i',k) = \hat{u}(i,k)(\hat{T}(i',k))^T$ with probability at most $(n-C_{Mul}-1))/2^m$.

We now analyze the degree of vertices $\mathcal{G}_{Mul,C}(k)$. By our third observation $v(i)$ in C cannot be connected to any vertex $v(i')\in A$, any vertex $v(i')\in C$ (including self-loops), and any vertex $v(i')\in D$. Thus the degree of vertices in C is at most $|B|$ which is at most M by our second observation, which in turn is strictly less than $C_{Mul}-M$.

For $v(i)\in D$, by our third observation $v(i)\in D$ cannot be connected to any $v(i')\in A$ or any $v(i')\in C$. Thus the degree of vertices in D is at most $|B\cup D|$ which is at most $M < C_{Mul}-M$ as above.

Finally, by our first observation $|A\cup B| \leq C_{Mul}-M-1$, and thus there are at most $C_{Mul}-M-1$ vertices in $\mathcal{G}_{Mul,C}(k)$ with degree at least $C_{Mul}-M$. □ □

Note: A similar proof can also be obtained by extending the Byzantine-error detection scheme in [7].

V. Discussion

We consider data transmission over a network containing a hidden adversary who can observe and corrupt data on network resources he controls. We first consider networks consisting of parallel edges and derive an intriguing two-part rate region. We show that if the adversary controls less than half of the edges, his presence can be detected and corrected for. We design low-complexity codes that achieve this performance. Conversely, if the adversary controls more than half the edges, we show that no positive throughput is possible.

We then consider multicasting over a directed acyclic network and show a similar two-part rate region. We give a tight characterization of this rate region in terms of the number of edges that the adversary controls.

Lastly, we consider extensions to the results on the parallel edge case, by considering models where either the source or the sink have knowledge of which edges are controlled by the adversary, noisy and variable capacity edges and feedback from the sink to the source.

Acknowledgments

The authors were partially supported by NSF Grants CCR-0220039 and CCR-0325673 and a grant from the Lee Center for Advanced Networking at Caltech.

References

[1] 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.

[2] I. Csiszár and P. Narayan. Common randomness and secret key generation with a helper. *IEEE Transactions on Information Theory*, 46(2):344–366, March 2000.

[3] Y. Z. Ding, P. Gopalan, and R. J. Lipton. Error correction against computationally bounded adverseries. In *DIMACS Workshop on Codes and Complexity*, December 2001.

[4] P. Gemmell and M. Naor. Codes for interactive authentication. *In Proceedings of CRYPTO 93': Lecture Notes in Computer Science*, 773:355–367, 1993.

[5] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.

[6] R. W. Hamming. Error detecting and error correcting codes. *Bell System Technical Journal*, 29:147–160, 1950.

[7] 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.

[8] 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.

[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] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 11(5):782–795, October 2003.

[11] M. Langberg. Private codes or succint random codes that are (almost) perfect. In *45th Annual Symposium on the Foundations of Computer Science*, 2004.

[12] R. J. McEliece, E. R. Rodemich, L. Rumsey, and L. R. Welch. New upper bounds on the rate of a code via the delsarte-mcwilliams inequalities. *IEEE Transactions on Information Theory*, 23:157–166, 1977.

[13] S. Micali, C. Peikert, M. Sudan, and D. Wilson. Optimal error correction against computationally bounded noise. In preparation., 2004.

[14] H. Pfister, I. Sason, and R. Urbanke. Capacity-achieving ensembles for the binary erasure channel with bounded complexity. In *IEEE International Symposium on Information Theory (ISIT)*, page 207, Chicago, July 2004.