# Resilient Network Coding in the Presence of Byzantine Adversaries 

S. Jaggi<br>jaggi@mit.edu

M. Langberg<br>mikel@caltech.edu

S. Katti<br>skatti@mit.edu

T. Ho<br>tho@caltech.edu

D. Katabi<br>dk@mit.edu

M. Médard<br>medard@mit.edu


#### Abstract

- Network coding substantially increases network throughput. But since it involves mixing of information inside the network, a single corrupted packet generated by a malicious node can end up contaminating all the information reaching a destination, preventing decoding.

This paper introduces the first distributed polynomial-time rateoptimal network codes that work in the presence of Byzantine nodes. We present algorithms that target adversaries with different attacking capabilities. When the adversary can eavesdrop on all links and jam $z_{O}$ links, our first algorithm achieves a rate of $C-2 z_{O}$, where $C$ is the network capacity. In contrast, when the adversary has limited snooping capabilities, we provide algorithms that achieve the higher rate of $C-z_{O}$.

Our algorithms attain the optimal rate given the strength of the adversary. They are information-theoretically secure. They operate in a distributed manner, assume no knowledge of the topology, and can be designed and implemented in polynomialtime. Furthermore, only the source and destination need to be modified; non-malicious nodes inside the network are oblivious to the presence of adversaries and implement a classical distributed network code. Finally, our algorithms work over wired and wireless networks.


## I. Introduction

Network coding allows the routers to mix the information content in packets before forwarding them. This mixing has been theoretically proven to maximize network throughput [1], [19], [13]. It can be done in a distributed manner with low complexity, and is robust to packet losses and network failures [8], [23]. Furthermore, recent implementations of network coding for wired and wireless environments demonstrate its practical benefits [16], [6].

But what if the network contains malicious nodes? A malicious node may pretend to forward packets from source to destination, while in reality it injects corrupted packets into the information flow. Since network coding makes the routers mix packets' content, a single corrupted packet can end up corrupting all the information reaching a destination. Unless this problem is solved, network coding may perform much worse than pure forwarding in the presence of adversaries.

The interplay of network coding and Byzantine adversaries has been examined by a few recent papers. Some detect the presence of an adversary [10], others correct the errors he injects into the codes under specific conditions [7], [12], [20], and a few bound the maximum achievable rate in such adverse environments [3], [29]. But attaining optimal rates using distributed and low-complexity codes is still an open problem.

This paper designs distributed polynomial-time rate-optimal network codes that combat Byzantine adversaries. We present three algorithms that target adversaries with different strengths. The adversary can inject $z_{O}$ packets per unit time, but his
listening power varies. When the adversary is omniscient, i.e., he observes transmissions on the entire network, our codes achieve the rate of $C-2 z_{O}$, with high probability. When the adversary's knowledge is limited, either because he eavesdrops only on a subset of the links or the source and destination have a low-rate secret-channel, our algorithms deliver the higher rate of $C-z_{O}$.

The intuition underlying all of our algorithms is that the aggregate packets from the adversarial nodes can be thought of as a second source. The information received at the destination is a linear transform of the source's and the adversary's information. Given enough linear combinations (enough coded packets), the destination can decode both sources. The question however is how does the destination distill out the source's information from the received mixture. To do so, the source's information has to satisfy certain constraints that the attacker's data cannot satisfy. This can be done by judiciously adding redundancy at the source. For example, the source may add redundancy to ensure that certain functions evaluate to zero on the original source's data, and thus can be used to distill the source's data from the adversary's. The challenge addressed herein is to design the redundancy that achieves the optimal rates.

This paper makes several contributions. The algorithms presented herein are the first distributed algorithms with polynomial-time complexity in design and implementation, yet are rate-optimal. In fact, since pure forwarding is a special case of network coding, being rate-optimal, our algorithms also achieve a higher rate than any approach that does not use network coding. They assume no knowledge of the topology and work in both wired and wireless networks. Furthermore, implementing our algorithms involves only a slight modification of the source and destination while the internal nodes can continue to use standard network coding.

## II. Illustrating Example

We illustrate the intuition underlying our approach using the toy example in Fig. 1. Calvin wants to prevent the flow of information from Alice to Bob, or at least minimize it. All links have a capacity of one packet per unit time. Further, Calvin connects to the three routers through an intermediate node. The intermediate node just relays all the packets Calvin sends him to the three routers. The network capacity, $C$, is by definition the min-cut from Alice to Bob. It is equal to 3 packets per unit time. The min-cut from Calvin to the destination is $z_{O}=1$ packet per unit time. Hence, the maximum rate from Alice to Bob in this scenario is bounded by $C-z_{O}=2$ packets per unit time as proven in [12].


Fig. 1-A simple example. Alice transmits to Bob. Calvin injects corrupted packets into their communication. The grey nodes in the middle perform network coding.

We express each packet as a vector of $n$ bytes, where $n$ is a sufficiently large number. The routers create random linear combinations of the packets they receive. Hence, every unit of time Bob receives the packets:

$$
\begin{equation*}
\tilde{\mathbf{y}}_{\mathbf{i}}=\alpha_{i} \tilde{\mathbf{x}}_{\mathbf{i}}+\beta_{i} \tilde{\mathbf{z}}, i \in\{1,2,3\} \tag{1}
\end{equation*}
$$

where $\tilde{\mathbf{x}}_{\mathrm{i}}$ 's are vectors representing the three packets Alice sent, $\tilde{\mathbf{z}}$ is the packet Calvin sent, $\alpha_{i}$ and $\beta_{i}$ are random coefficients.

In our example, the routers operate over bytes; the $i^{t h}$ byte in an outgoing packet is a linear combination of $i^{t h}$ bytes in the incoming packets. Thus, (1) also describes the relation between the individual bytes in $\tilde{\mathbf{y}}_{\mathrm{i}}$ 's and the corresponding bytes in $\tilde{\mathbf{x}}_{\mathbf{i}}$ 's and $\tilde{\mathbf{z}}$.

Since the routers mix the content of the packets, Alice cannot just sign her packets and have Bob discard all packets with incorrect signatures. To decode, Bob has to somehow distill the $\tilde{x}_{i}$ 's from the $\tilde{\mathbf{y}}_{\mathrm{i}}$ 's he receives.

As a first attempt at solving the problem, let us assume that Bob knows the topology, i.e., he knows that the packets he receives are produced using (1). Further, let us assume that he knows the random coefficients used by the routers to code the packets, i.e., he knows the values of $\alpha_{i}$ 's and $\beta_{i}$ 's. To decode, Bob has to solve (1). Since each packet contains $n$ bytes, the system in (1) represents $3 n$ equations, one equation per received byte. Bob has $3 n$ equations and $4 n$ unknowns ( $n$ unknown bytes per each packet $\tilde{\mathbf{z}}, \tilde{\mathbf{x}}_{1}, \tilde{\mathbf{x}}_{2}$ and $\tilde{\mathbf{x}}_{3}$ ). Hence, Bob cannot decode.

To address the above situation, Alice needs to add redundancy to her transmitted packets. After all, as noted above, for the particular example in Fig. 1, Alice's rate is bounded by 2 packets per unit time. Thus, Alice should send no more than 2 packets worth of information. She can use the third packet for added redundancy. Suppose Alice sets

$$
\begin{equation*}
\tilde{x}_{3}=\tilde{x}_{1}+\tilde{x}_{2} \tag{2}
\end{equation*}
$$

This coding strategy is public to both Bob and Calvin. Since each packet contains $n$ bytes, combining (2) with (1), Bob obtains a system of $4 n$ equations with $4 n$ unknowns, which he can solve to decode.

But in the general case, Bob knows nothing about the coefficients used by the routers, the topology, or the overall network transform. Said differently, the 6 coefficients corresponding to the $\alpha_{i}$ 's and the $\beta_{i}$ 's are usually unknown to Bob. Thus, given (1) and (2), Bob is faced with $4 n$ equations and $4 n+6$ unknowns, and thus cannot decode. The matter is further complicated by the non-linearity of (1), which involves the product of unknown terms $\alpha_{i} \tilde{\mathbf{x}}_{\mathbf{i}}$ and $\beta_{i} \tilde{\mathbf{z}}$.

The first idea we exploit in our solution is that while $\tilde{\mathbf{z}}$ is a whole unknown packet of $n$ bytes, each of the coefficients
$\beta_{i}$ is a single byte. Thus, instead of devoting a whole vector of $n$ bytes for added redundancy (as in (2)), Alice just needs to introduce 6 extra bytes of redundancy to compensate for the $\alpha_{i}$ 's and $\beta_{i}$ 's being unknown.

Alice imposes constraints on her data to help Bob to decode. For instance, a simple constraint could be that the first byte in each packet equals zero. This constraint provides Bob with 2 additional equations (recall that the first byte in $\tilde{\mathrm{x}}_{3}$ is forced to 0 due to (2), and hence the new constraint produces 2 additional equations rather than 3). Rewriting (1) for the first byte of each packet, we obtain:

$$
\begin{equation*}
y_{i, 1}=\alpha_{1} x_{i, 1}+\beta_{i} z_{1}=\beta_{i} z_{1},, i \in\{1,2,3\} \tag{3}
\end{equation*}
$$

where $y_{i, j}$ denotes the $j^{\text {th }}$ byte in the $i^{\text {th }}$ received packet. The above equations provide Bob with a scaled version of the $\beta_{i}$ 's, i.e., they are all multiplied by $z_{1}$.

Our second observation is that the scaled version of the $\beta_{i}$ 's suffices for Bob to decode $\tilde{\mathbf{x}}$. This can be seen by a simple algebraic manipulation of (1). Bob can rewrite the equations in (1) by multiplying and dividing the second term with $z_{1}$ and appending (2) to obtain

$$
\begin{equation*}
\tilde{\mathbf{y}}_{\mathbf{i}}=\alpha_{i} \tilde{\mathbf{x}}_{\mathbf{i}}+\left(\beta_{i} z_{1}\right)\left(\tilde{\mathbf{z}} / z_{1}\right), i \in\{1,2,3\} \tag{4}
\end{equation*}
$$

Notice that Bob already knows all three $\beta_{i} z_{1}$ terms from (3). The term $\left(\tilde{\mathbf{z}} / z_{1}\right)$ can be considered a single unknown because Bob does not care about estimating the exact value of $\tilde{\mathbf{z}}$.

To allow Bob to discover the $\alpha_{i}$ 's, Alice similarly adds 4 more bytes of redundancy by imposing constraints on the second and third bytes in her packets. For example, she chooses $x_{1,2}=$ $x_{2,2}=1$ and $x_{1,3}=-x_{2,3}=1$ (combined with (2), these constraints force $x_{3,2}=2$ and $x_{3,3}=0$ ). Substituting the values of $\left(\beta_{1} z_{1}\right),\left(\beta_{2} z_{1}\right)$ and $\left(\beta_{3} z_{1}\right)$ from (3) gives Bob the following equations.

$$
\begin{array}{cc}
y_{1,2}=\alpha_{1}+y_{1,1}\left(z_{2} / z_{1}\right), & y_{1,3}=\alpha_{1}+y_{1,1}\left(z_{3} / z_{1}\right) \\
y_{2,2}=\alpha_{2}+y_{2,1}\left(z_{2} / z_{1}\right), & y_{2,3}=-\alpha_{2}+y_{2,1}\left(z_{3} / z_{1}\right)  \tag{5}\\
y_{3,2}=2 \alpha_{3}+y_{3,1}\left(z_{2} / z_{1}\right), & y_{3,3}=y_{3,1}\left(z_{3} / z_{1}\right)
\end{array}
$$

Now Bob has 6 linear equations with the 5 unknowns $\alpha_{1}$, $\alpha_{2}, \alpha_{3}, z_{2} / z_{1}$ and $z_{3} / z_{1}$, and they can be solved to obtain the $\alpha_{i}$ 's. Hence we are essentially back to the situation where Bob knows the $\alpha_{i}$ 's and $\beta_{i}$ 's, and can solve for $\tilde{\mathbf{x}}_{i}$ 's.

One complication still remains. If Calvin knows the constraints on Alice's data, he will try to assign values to his bytes to prevent Bob from decoding. For example, if Calvin knows that the first byte of each of Alice's packets is zero, he too would set the first byte in his packet $z_{1}$ to zero, in which case Bob does not obtain any information about the $\beta_{i}$ 's from (4).

There are two ways out of this situation. Suppose Alice could communicate to Bob a small message that is secret from Calvin. In this case, she could compute a small number of hashes of her data, and transmit them to Bob. These hashes correspond to constraints on her data, which enables Bob to decode. If Alice cannot communicate secretly with Bob, she leverages the fact that Calvin can inject only one fake packet. Since Calvin's packet is $n$ bytes long, he can cancel out at most $n$ hashes. If Alice injects $n+1$ hashes, there must be at least one hash Calvin cannot cancel. This hash enables Bob to find the $\beta_{i}$ 's and decode. Notice, however, that the $n+1$ additional constraints imposed on the bytes in $\vec{x}_{1}$ and $\vec{x}_{2}$ mean that Alice
can only transmit at most $n-1$ bytes of data to Bob. For a large number of bytes $n$ in a packet, this rate is asymptotically optimal against an all-knowing adversary [3].

After giving some intuition on how our scheme works, the rest of this paper considers the general problem of network coding over completely unknown topology, in the presence of an adversary who has partial or full knowledge of the network and transmissions in it.

## III. Related Work

We start with a brief summary of network coding, followed by a survey of prior work on Byzantine adversaries in networks.

## A. Network Coding Background

Work on network coding started with a pioneering paper by Ahlswede et al. [1], which establishes the value of coding in the routers and provides theoretical bounds on the capacity of such networks. The combination of [21], [19], [13] shows that, for multicast traffic, linear codes achieve the maximum capacity bounds, and coding and decoding can be done in polynomial time. Additionally, Ho et al. show that the above is true even when the routers pick random coefficients [8]. Researchers have extended the above results to a variety of areas including wireless networks [23], [15], [16], energy [28], secrecy [2], content distribution [6], and distributed storage [14].

## B. Byzantine Adversaries in Networks

A Byzantine attacker is a malicious adversary hidden in a network, capable of eavesdropping and jamming communications. Prior research has examined these attacks in the presence of network coding and without it. In the absence of network coding, Dolev et al. [5] consider the problem of communicating over a known graph containing Byzantine adversaries. They show that for $k$ adversarial nodes, reliable communication is possible only if the graph has more than $2 k+1$ vertex connectivity. Subramaniam extends this result to unknown graphs [26]. Pelc et al. address the same problem in wireless networks by modeling malicious nodes as locally bounded Byzantine faults, i.e., nodes can overhear and jam packets only in their neighborhood [24].

The interplay of network coding and Byzantine adversaries was first examined in [10], which detects the existence of an adversary but does not provide an error-correction scheme. This has been followed by the work of Cai and Yeung [29], [3], who generalize standard bounds on error-correcting codes to networks, without providing any explicit algorithms for achieving these bounds. Our work presents a constructive design to achieve those bounds.

The problem of correcting errors in the presence of both network coding and Byzantine adversaries has been considered by a few prior proposals. Earlier work [20], [7] assumes a centralized trusted authority that provides hashes of the original packets to each node in the network. More recent work by Charles et al. [4] obviates the need for a trusted entity under the assumption that the majority of packets received by each node is uncorrupted. In contrast to the above two schemes which are cryptographically secure, in a previous work [12], we consider an information-theoretically rate-optimal solution to Byzantine attacks for wired networks, which however requires a centralized design. This paper builds on the above prior schemes to combine

| Scheme | Charles et.al. [4] | Jaggi et.al. [12] | Ours |
| :--- | :---: | :---: | :---: |
| Info. Theoretic Security | No | Yes | Yes |
| Distributed | Yes | No | Yes |
| Internal Node Complexity | High | Low | Low |
| Decoding Complexity | High | Exponential | Low |
| General Graphs | No | Yes | Yes |
| Universal | No | No | Yes |

## TABLE I-Comparison between the results in this paper and some prior papers.

their desirable traits; it provides a distributed solution that is information-theoretically rate optimal and can be designed and implemented in polynomial time. Furthermore, our algorithms have new features; they assume no knowledge of the topology, do not require any new functionality at internal nodes, and work for both wired and wireless networks. Recent work [17] has considered the same problem from a different perspective, their results and bounds are similar to ours. Table I highlights similarities and differences from prior work.

## IV. Model \& Definitions

We use a general model that encompasses both wired and wireless networks. To simplify notation, we consider only the problem of communicating from a single source to a single destination. But similar to most network coding algorithms, our techniques generalize to multicast traffic.

## A. Threat Model

There is a source, Alice, and a destination, Bob, who communicate over a wired or wireless network. There is also an attacker Calvin, hidden somewhere in the network. Calvin aims to prevent the transfer of information from Alice to Bob, or at least to minimize it. He can observe some of the transmissions, and can inject his own. When he injects his own packets, he pretends they are part of the information flow from Alice to Bob.

Calvin is quite strong. He is computationally unbounded. He knows the encoding and decoding schemes of Alice and Bob, and the network code implemented by the interior nodes. He also knows the exact network realization.

## B. Network and Code Model

This section describes the network model, the packet format, and how the network transforms the packets.
Network Model: The network is modeled as a hypergraph [22]. Each packet transmission corresponds to a hyperedge directed from the transmitting node to the set of observer nodes. The hypergraph model captures both wired and wireless networks. For wired networks, the hyperedge is a simple point-to-point link. For wireless, each such hyperedge is determined by instantaneous channel realizations (packets may be lost due to fading or collisions) and connects the transmitter to all nodes that hear the transmission. The hypergraph is unknown to Alice and Bob prior to transmission.
Source: Alice generates incompressible data that she wishes to deliver to Bob over the network. To do so, Alice encodes her data as dictated by the encoding algorithm (described in subsequent sections). She divides the encoded data into batches of $b$ packets. For clarity, we focus on the encoding and decoding of one batch.

A packet contains a sequence of $n$ symbols from the finite field $\mathbb{F}_{q}$. All arithmetic operations henceforth are done over symbols from $\mathbb{F}_{q}$. (See the treatment in [18]). Out of the $n$ symbols in Alice's packet, $\delta n$ symbols are redundancy added by the source.

Alice organizes the data in each batch into a matrix $X$ as shown in Fig. 2. We denote the $(i, j)^{t h}$ element in the matrix by $x(i, j)$. The $i^{t h}$ row in the matrix $X$ is just the $i^{t h}$ packet in the batch. Fig. 2 shows that similarly to standard network codes [8], some of the redundancy in the batch is devoted to sending the identity matrix, $I$. Also, as in [8], Alice takes random linear combinations of the rows of $X$ to generate her transmitted packets. As the packets traverse the network, the internal nodes apply a linear transform to the batch. The identity matrix receives the same linear transform. The destination discovers the linear relation between the packets it receives and those transmitted by inspecting how $I$ was transformed.

Adversary: Let the matrix $Z$ be the information Calvin injects into each batch. The size of this matrix is $z_{O} \times n$, where $z_{O}$ is the size of the min-cut from Calvin to the destination.

Destination: Analogously to how Alice generates $X$, the destination Bob organizes the received packets into a matrix $Y$. The $i^{\text {th }}$ received packet corresponds to the $i^{\text {th }}$ row of $Y$. Note that the number of received packets, and therefore the number of rows of $Y$, is a variable dependent on the network topology. The column rank of $Y$, however, is $b+z_{O}$. Bob attempts to reconstruct Alice's information, $X$, using the matrix of received packets $Y$.

## C. Definitions

We define the following concepts.

- The network capacity, denoted by $C$, is the time-average of the maximum number of packets that can be delivered from Alice to Bob, assuming no adversarial interference, i.e., the max flow. It can be also expressed as the min-cut from source to destination. (For the corresponding multicast case, $C$ is defined as the minimum of the min-cuts over all destinations.)
- The error probability is the probability that Bob's reconstruction of Alice's information is inaccurate.
- The rate, $R$, is the number of information bits in a batch amortized by the length of a packet in bits.
- The rate $R$ is said to be achievable if for any $\epsilon>0$, any $\delta>0$, and sufficiently large $n$, there exists a block-length- $n$ network code with a redundancy $\delta$ and a probability of error less than $\epsilon$.
- A code is said to be universal if the code design is independent of $z_{O}$.


## V. Network Transform

This section explains how Alice's packets get transformed as they travel through the network. It examines the effect the adversary has on the received packets, and Bob's decoding problem.

The network performs a classical distributed network code [8]. Specifically, each packet transmitted by an internal node is a random linear combination of its incoming packets.


Fig. 2-Alice, Bob and Calvin's information matrices.

| Variable | Definition |
| :--- | :--- |
| $b$ | Number of packets in a batch. |
| $z_{O}$ | Number of packets Calvin can inject. |
| $z_{I}$ | Number of packets Calvin can hear. |
| $n$ | Length of each packet. |
| $\delta$ | Fractional redundancy introduced by Alice. |
| $\hat{T}$ | Proxy of the transfer matrix $T$ representing the <br> network transform. |

## TABLE II-Terms used in the paper.

Thus, the effect of the network at the destination can be summarized as follows.

$$
\begin{equation*}
Y=T X+T_{Z \rightarrow Y} Z \tag{6}
\end{equation*}
$$

where $X$ is the batch of packets sent by Alice, $Z$ refers to the packets Calvin injects into Alice's batch, and $Y$ is the received batch. The variable $T$ refers to the linear transform from Alice to Bob, while $T_{Z \rightarrow Y}$ refers to the linear transform from Calvin to Bob.

As explained in §IV, a classical random network code's $X$ includes the identity matrix as part of each batch. The identity matrix sent by Alice incurs the same transform as the rest of the batch. Thus,

$$
\begin{equation*}
\hat{T}=T I+T_{Z \rightarrow Y} L \tag{7}
\end{equation*}
$$

where $\hat{T}$ and $L$ are the columns corresponding to $I$ 's location in $Y$ and $Z$ respectively, as shown in Fig. 2.

In standard network coding, there is no adversary, i.e., $Z=$ 0 and $L=0$, and thus $\hat{T}=T$. The destination receives a description of the network transform in $\hat{T}$ and can decode $X$ as $\hat{T}^{-1} Y$. In the presence of the adversary, however, the destination needs to solve (6) and (7) to extract the value of $X$.

By substituting $T$ from (7), (6) can be simplified to get

$$
\begin{align*}
Y & =\hat{T} X+T_{Z \rightarrow Y}(Z-L X)  \tag{8}\\
& =\hat{T} X+E \tag{9}
\end{align*}
$$

where $E$ is a $C \times n$ matrix that characterizes Calvin's interference. Note that the matrix $\hat{T}$, which Bob knows, acts as a proxy transfer matrix for $T$, which he doesn't know.

Note that in (6), all terms other than $Y$ are unknown. Further, it is non-linear due to the cross-product terms, $T X$ and $T_{Z \rightarrow Y} Z$. In contrast, (9) is linear in the unknowns $X$ and $E$. The rest of this work focuses on solving (9) under different assumptions on Calvin's strength.

## VI. Summary of Results

We have three main results. Each result corresponds to a distributed, rate-optimal, polynomial-time algorithm that defeats an adversary of a particular type. The optimality of these rates
has been proven by prior work [3], [29], [12]. Our work, however, provides a construction of distributed codes/algorithms that achieve optimal rates. In what follows, let $|\mathcal{T}|$ denote the number of receivers, and $|\mathcal{E}|$ denote the number of transmissions in the network.
(1) Shared Secret Model: This model assumes that Alice and Bob have a very low rate secret channel, the transmissions on which are unknown to Calvin. It considers the transmission of information via network coding in a network where Calvin can observe all transmissions, and can inject some corrupt packets.

Theorem 1: The Shared Secret algorithm achieves a rate of $C-z_{O}$ with code-complexity $\mathcal{O}\left(n C^{2}\right)$. This is the maximum achievable rate.
In §VII, we prove the above theorem by constructing an algorithm that achieves the bounds. Note that [7] proves a similar result for a more constrained model where Alice shares a very low rate secret channel with all nodes in the network, and the operations performed by internal nodes are computationally expensive. Further, their result guarantees cryptographic security, while we provide information-theoretic security.
(2) Omniscient Adversary Model: This model assumes an omniscient adversary, i.e., one from whom nothing is hidden. In particular, Alice and Bob have no shared secrets hidden from Calvin. It also assumes that the min-cut from the adversary to the destination, $z_{O}$, is less than $C / 2$. Prior work proves that without this condition, it is impossible for the source and the destination to reliably communicate without a secret channel [12]. In $\S$ VIII, we prove the following.

Theorem 2: The Omniscient Adversary algorithm achieves a rate of $C-2 z_{O}$ with code-complexity $\mathcal{O}\left((n C)^{3}\right)$. This is the maximum achievable rate.
(3) Limited Adversary Model: In this model, Calvin is limited in his eavesdropping power; he can observe at most $z_{I}$ transmitted packets. Exploiting this weakness of the adversary results in an algorithm that, like the Omniscient Adversary algorithm operates without a shared secret, but still achieves the higher rate possible via the Shared Secret algorithm. In particular, in §IX we prove the following.

Theorem 3: If $z_{I}<C-2 z_{O}$, the Limited Adversary algorithm achieves a rate of $C-z_{O}$ with code-complexity $\mathcal{O}\left(n C^{2}\right)$. This is the maximum achievable rate.

## VII. Shared Secret Model

In the Shared Secret model, Alice and Bob have use of a strong resource, namely a secret channel over which Alice can transmit a small amount of information to Bob that is secret from Calvin. Note that since the internal nodes mix corrupted and uncorrupted packets, Alice cannot just sign her packets and have Bob check the signature and throw away corrupted packets, in extreme cases this might lead to Bob not receiving any uncorrupted packets. Alice uses the secret channel to send a hash of her information $X$ to Bob, which Bob can use to distill the corrupted packets he receives, as explained below.
Shared Secret: Alice generates her secret message in two steps. She first chooses $C$ parity symbols uniformly at random from the field $\mathbb{F}_{q}$. The parity symbols are labeled $r_{d}$, for $d \in$ $\{1, \ldots, C\}$. Corresponding to the parity symbols, Alice's paritycheck matrix $P$ is defined as the $n \times C$ matrix whose $(i, j)^{t h}$
entry equals $\left(r_{j}\right)^{i}$, i.e., $r_{j}$ to the $i^{t h}$ power. The second part of Alice's secret message is the $b \times C$ hash matrix $H$, computed as the matrix product $X P$. We assume Alice communicates both the set of parity symbols and the hash matrix $H$ to Bob over the secret channel. The combination of these two creates the shared secret, denoted $\mathbb{S}$, between Alice and Bob. The size of $\mathbb{S}$ is $C(b+1)$ symbols, which is small in comparison to Alice's information $X$. (The size of $X$ is $b \times n$; it can be made arbitrarily large compared to the size of $\mathbb{S}$ by increasing the packet size $n$.)

Alice's Encoder: Alice implements the classical random network encoder described in §IV-B.
Bob's Decoder: Not only is $P$ used by Alice to generate $H$, but is also used by Bob in his decoding process. To be more precise, Bob computes $Y P-\hat{T} H$ using the messages he gets from the network and the secret channel. We call the outcome the syndrome matrix $S$.

By substituting the value of $H$ and using (9), we obtain

$$
\begin{equation*}
S=Y P-\hat{T} H=(Y-\hat{T} X) P=E P \tag{10}
\end{equation*}
$$

Thus, if no adversary was present, the packets would not be corrupted (i.e., $E=0$ ) and $S$ would be an all-zero matrix. As shown in $\S$ IV, $X$ then equals $\hat{T}^{-1} Y$. If Calvin injects corrupt packets, $S$ will be a non-zero matrix.

Claim 1: The rank of $E$ is at most $z_{O}$.
Claim 2: The columns of $S$ span the same vector-space as the columns of $E$ with probability at least $1-C n^{C} q^{-1}$.
Claim 1 follows from the definition of $E=T_{Z \rightarrow Y}(Z-L X)$.
Claim 2 is proved in the Appendix. Together, they imply that Calvin's interference, $E$, can be written as linear combinations of the columns of a $C \times z_{O}$ submatrix $S^{\prime}$ of $S$, i.e., $E=S^{\prime} A$, where $A$ is a $z_{O} \times n$ matrix. This enables Bob to rewrite (9) as the matrix product

$$
Y=\left[\begin{array}{ll}
\hat{T} & S^{\prime}
\end{array}\right]\left[\begin{array}{l}
X  \tag{11}\\
A
\end{array}\right]
$$

Bob does not care about $A$, but to obtain $X$, he must solve (11). Let $|\mathcal{T}|$ and $|\mathcal{E}|$ be the number of terminals and links in the underlying network.

Claim 3: The matrix [ $\left.\begin{array}{c}T \\ T_{Z \rightarrow Y}\end{array}\right]$, and thus the matrix $\left[\hat{T} S^{\prime}\right.$ ], has full column-rank with probability at least $1-|\mathcal{T}||\mathcal{E}| q^{-1}$.
Claim 3, proved in the Appendix, means that Bob can decode by simply inverting the $C \times C$ matrix $\left[\hat{T} S^{\prime}\right.$ ] and multiplying the result by $Y$. Thus, the shared secret algorithm achieves the rate of $C-z_{O}-b^{2} / n$. Here, the asymptotically negligible term $b^{2} / n$ corresponds to the overhead due to the identity matrix Alice appends to $X$. This rate is shown to be optimal by prior work [12]. The probability of error is at most the sums of the probabilities of error in Claims 2 and 3, i.e., $\left(n^{C} C+|\mathcal{T}||\mathcal{E}|\right) q^{-1}$. Of code design, encoding and decoding, both encoding and decoding require $\mathcal{O}\left(n C^{2}\right)$ steps. The costliest step for Alice is the computation of the hash matrix $H$, and for Bob is the computation of the syndrome matrix $S$.

The scheme presented above is universal, i.e., the parameters of the code do not depend on any knowledge about $z_{O}$, which in some sense functions as the "noise parameter" of the network. Alice therefore has flexibility in tailoring her batch size to the size of the data which she wishes to transmit and the packet size allowed by the network.

## VIII. Omniscient Adversary Model

What if we face an omniscient adversary, i.e., Calvin can observe everything, and there are no shared secrets between Alice and Bob? We design a network error-correcting code to defeat such a powerful adversary. Our algorithm achieves a rate of $R=C-2 z_{O}$, which is lower than in the Shared Secret model. This is a direct consequence of Calvin's increased strength. Recent bounds [3] on network error-correcting codes show that in fact $C-2 z_{O}$ is the maximum achievable rate for networks with an omniscient adversary.

Alice's Encoder: Alice encodes in two steps. To counter the adversary's interference, she first generates $X$ by adding redundancy to her information. She then encodes $X$ using the encoder defined in §IV-B.

Alice adds redundancy as follows. Her original information is a length- $\left(b n-\delta n-b^{2}\right)$ column vector $\tilde{\mathbf{U}}$. (Here the fractional redundancy $\delta$, is dependent on $z_{O}$, the number of packets Calvin may inject into the network.) Alice converts $\tilde{\mathbf{U}}$ into $\tilde{\mathbf{X}}$, a length$b n$ vector $\left(\begin{array}{lll}\tilde{\mathbf{U}} & \tilde{\mathbf{R}} & \tilde{\mathbf{I}}\end{array}\right)^{T}$, where $\tilde{\mathbf{I}}$ is just the column version of the $b \times b$ identity matrix. It is generated by stacking columns of the identity matrix one after the other. The second term, $\tilde{\mathbf{R}}$ represents the redundancy Alice adds. The redundancy vector $\tilde{\mathbf{R}}$ is a length $-\delta n$ column vector generated by solving the matrix equation for $\tilde{\mathbf{R}}$.

$$
D\left(\begin{array}{lll}
\tilde{\mathbf{U}} & \tilde{\mathbf{R}} & \tilde{\mathbf{I}}
\end{array}\right)^{T}=\mathbf{0}
$$

where $D$ is a $\delta n \times b n$ matrix defined as the redundancy matrix. $D$ is obtained by choosing each element as an independent and uniformly random symbol from the finite field $\mathbb{F}_{q}$. Due to the dependence of $D$ on $\delta$ and thus on $z_{O}$, the Omniscient Adversary algorithm is not universal. The redundancy matrix $D$ is known to all parties - Alice, Bob, and Calvin - and hence does not constitute a shared secret.

Alice then proceeds to the standard network encoding. She rearranges $\tilde{\mathbf{X}}$, a length- $b n$ vector, into the $b \times n$ matrix $X$. The $j^{\text {th }}$ column of $X$ consists of symbols from the $((j-1) b+1)^{\text {th }}$ through $(j b)^{t h}$ symbols of $\mathbf{X}$. From this point on, Alice's encoder implements the classical random network encoder described in §IV-B, to generate her transmitted packets.
Bob's Decoder: As shown in (9), Bob's received data is related to Alice and Calvin's transmitted data as $Y=\hat{T} X+E$. Bob's objective, as in $\S \mathrm{VII}$, is to distill out the effect of the error matrix $E$ and recover the vector $X$. He can then retrieve Alice's data by extracting the first $\left(b n-b^{2}-\delta n\right)$ symbols to obtain $\tilde{\mathbf{U}}$.

To decode, Bob performs the following steps, each of which corresponds to an elementary matrix operation.

- Determining Calvin's strength: Bob first determines the strength of the adversary $z_{O}$, which is the column rank of $T_{Z \rightarrow Y}$. Bob does not know $T_{Z \rightarrow Y}$, but since $T$ and $T_{Z \rightarrow Y}$ span disjoint vector spaces (Claim 3), the column rank of $Y$ is equal to the sum of the column ranks of $T$ and $T_{Z \rightarrow Y}$. Since the column rank of $T$ is simply the batch size $b$, Bob determines $z_{O}$ by subtracting $b$ from the column rank of the matrix $Y$.
- Discarding irrelevant information: Since the classical random network code is run without any central coordinating authority, the packets of information that Bob receives
may be highly redundant. Of the packets Bob receives, he selectively discards some so that the resulting matrix $Y$ has $b+z_{O}$ rows, and has full row rank. For him to consider more packets is useless, since at most $b+z_{O}$ packets of information have been injected into the network, $b$ from Alice and $z_{O}$ from Calvin. This operation has the additional benefit of reducing the complexity of linear operations that Bob needs to perform henceforth. This reduces the dimensions of the matrix $\hat{T}$, since Bob can discard the rows corresponding to the discarded packets.
- Estimating a "basis" for E: If Bob could directly estimate a basis for the column space of $E$, then he could simply decode as in the Shared Secret algorithm. However, there is no shared secret that enables him to discover a basis for the column space of $E$. So, he instead chooses a proxy error matrix $T^{\prime \prime}$ whose columns (which are, in general, linear combinations of columns of both $X$ and $E$ ) act as a proxy error basis for columns of $E$. This is analogous to step (9), where the matrix $\hat{T}$ acts as a proxy transfer matrix for the unknown matrix $T$.
The matrix $T^{\prime \prime}$ is obtained as follows. Bob selects $z_{O}$ columns from $Y$ such that these columns, together with the $b$ columns of $\hat{T}$, form a basis for the columns of $Y$. Without loss of generality, these columns correspond to the first $z_{O}$ columns of $Y$ (if not, Bob simply permutes the columns of $Y$ to make it so $)$. The $\left(b+z_{O}\right) \times z_{O}$ matrix corresponding to these first $z_{O}$ columns is denoted $T^{\prime \prime}$.
- Changing to proxy basis: Bob rewrites $Y$ in the basis corresponding to the columns of the $\left(b+z_{O}\right) \times\left(b+z_{O}\right)$ matrix $\left[T^{\prime \prime} \widehat{T}\right]$. Therefore $Y$ can now be written as

$$
Y=\left[T^{\prime \prime} \hat{T}\right]\left[\begin{array}{ccc}
I_{z_{O}} & F^{Z} & 0  \tag{12}\\
0 & F^{X} & I_{b}
\end{array}\right] .
$$

Here $\left[\begin{array}{l}F^{Z} \\ F^{X}\end{array}\right]$ is defined as the $\left(b+z_{O}\right) \times\left(n-\left(b+z_{O}\right)\right)$ matrix representation of the columns of $Y$ (other than those in $\left[T^{\prime \prime} \hat{T}\right]$ ) in the new basis, with $F^{Z}$ and $F^{X}$ defined as the sub-matrices of appropriate dimensions.
Bob splits $X$ as $X=\left[\begin{array}{lll}X_{1} & X_{2} & X_{3}\end{array}\right]$, where $X_{1}$ corresponds to the first $z_{O}$ columns of $X, X_{3}$ to the last $b$ columns of $X$, and $X_{2}$ to the remaining columns of $X$. We perform linear algebraic manipulations on (12), to reduce it to a form in which the variables in $X$ are related by a linear transform solely to quantities that are computable by Bob. Claim 4 summarizes the effect of these linear algebraic manipulations (proof in Appendix).

Claim 4: The matrix equation (12) is exactly equivalent to the matrix equation $\hat{T} X_{2}=\hat{T}\left(F^{X}+X_{1} F^{Z}\right)$.
To complete the proof of correctness of our algorithm, we need only the following claim, proved in the Appendix.

Claim 5: For $\delta n>n\left(z_{O}+\varepsilon\right)$, with probability greater than $1-q^{-n \varepsilon}$, the system of linear equations

$$
\begin{align*}
\hat{T} X_{2} & =\hat{T}\left(F^{X}+X_{1} F^{Z}\right)  \tag{13}\\
D \tilde{\mathbf{X}} & =0 \tag{14}
\end{align*}
$$

is solvable for $X$.
The final claim enables Bob to recover $X$, which contains Alice's information at asymptotic rate $R=C-2 z_{O}$. (There is an asymptotically negligible rate overhead equalling $b^{2} / n+\epsilon$.

The $b^{2} / n$ term corresponds, as before, to the identity matrix appended to $X$. The term $\epsilon$ takes any positive value, and the probability of error also depends on it.) The probability of error equals the sums of the probabilities of error in Claims 3 and 5, i.e., $|\mathcal{T} \| \mathcal{E}| q^{-1}+q^{-n \varepsilon}$. Of code design, encoding and decoding, the most computationally expensive is decoding. The costliest step involves inverting the linear transform corresponding to (13)-(14), which is of dimension $\mathcal{O}(n C)$.

## IX. Limited Adversary Model

We combine the strengths of the Shared Secret algorithm and the Omniscient Adversary algorithm, to achieve the higher rate of $C=C-z_{O}$, without needing a secret channel. The caveat is that Calvin's strength is more limited; the number of packets he can transmit, $z_{O}$, and the number he can eavesdrop on, $z_{I}$, satisfy the technical constraint

$$
\begin{equation*}
2 z_{O}+z_{I}<C \tag{15}
\end{equation*}
$$

We call such an adversary a Limited Adversary.
The main idea underlying our Limited Adversary algorithm is simple. Alice uses the Omniscient Adversary algorithm to transmit a "short" message to Bob at rate $C-2 z_{O}$. By (15), $z_{I}<C-2 z_{O}$, the rate $z_{I}$ at which Calvin eavesdrops is strictly less than Alice's rate of transmission $C-2 z_{O}$. Hence Calvin cannot decode Alice's message, but Bob can. This means Alice's message to Bob is secret from Calvin. Alice then builds upon this secret, using the Shared Secret algorithm to transmit the bulk of her message to Bob at the higher rate $C-z_{O}$.

Though the following algorithm requires Alice to know $z_{O}$ and $z_{I}$, we describe in §IX-A how to change the algorithm to make it independent of these parameters. The price we pay is a slight decrease in rate.
Alice's Encoder: Alice's encoder follows essentially the schema described above, except for a technicality - the information she transmits to Bob via the Omniscient Adversary algorithm is padded with some random symbols. This is for two reasons. Firstly, since the Omniscient Adversary algorithm has a probability of error that decays exponentially with the size of the input, it isn't guaranteed to perform well to transmit just a small message. Secondly, the randomness in the padded symbols also ensures strong information-theoretic secrecy of the small secret message, i.e., we can then show (in Claim 6) that Calvin's best estimate of any function of the secret information is no better than if he made random guesses.

Alice's information $X$ decomposes into two parts [ $X_{1} X_{2}$ ]. She uses the information she wishes to transmit to Bob, at rate $R=C-z_{O}-\Delta$, as input to the encoder of the Shared Secret algorithm, thereby generating the $b \times n(1-\Delta)$ sub-matrix $X_{1}$. Here $\Delta$ is a parameter that enables Alice to trade off between the the probability of error and rate-loss.

The second sub-matrix, $X_{2}$, which we call the secrecy matrix is analogous to the secret $\mathbb{S}$ used in the Secret Sharing algorithm described in $\S$ VII. The size of $X_{2}$ is $b \times \Delta n$. In fact, $X_{2}$ is an encoding of the secret $\mathbb{S}$ Alice generates in the Shared Secret algorithm. The $b(C+1)$ symbols corresponding to the parity symbols $\left\{r_{d}\right\}$ and the hash matrix $H$ are written in the form of a length- $b(C+1)$ column vector. This vector is appended with symbols chosen uniformly at random from $\mathbb{F}_{q}$ to result in the length- $\left(C-z_{O}-\delta\right) \Delta n$ vector $\tilde{\mathbf{U}}^{\prime}$. This vector $\tilde{\mathbf{U}}^{\prime}$ could
function as the input $\tilde{\mathbf{U}}$ to the Omniscient Adversary algorithm operated over a packet-size $\Delta n$, with a probability of decoding error that is exponentially small in $\Delta n$; however, we actually use a hash of $\tilde{\mathbf{U}}^{\prime}$ to generate the input $\tilde{\mathbf{U}}$ to the Omniscient Adversary algorithm. To be more precise, $\tilde{\mathbf{U}}=V \tilde{\mathbf{U}}^{\prime}$, where $V$ is any square MDS code generator matrix ${ }^{1}$ of dimension $\left(C-z_{O}-\delta\right) \Delta n$, known to all parties Alice, Bob, and Calvin. As we see later, hashing $\tilde{\mathbf{U}}^{\prime}$ with $V$ strengthen the secrecy of $\mathbb{S}$ (and enables the proof of Claim 6 below). Alice then uses the encoder for the Omniscient Adversary algorithm to generate $X_{2}$ from $\tilde{\mathbf{U}}$.

The two components of $X$, i.e., $X_{1}$ and $X_{2}$, respectively correspond to the information Alice wishes to transmit to Bob, and an implementation of the low rate secret channel. The fraction of the packet-size corresponding to $X_{2}$ is "small", i.e., $\Delta$. Finally, Alice implements the classical random encoder described in §IV-B.
Bob's Encoder: Bob arranges his received packets into the matrix $Y=\left[\begin{array}{ll}Y_{1} & Y_{2}\end{array}\right]$. The sub-matrices $Y_{1}$ and $Y_{2}$ are respectively the network transforms of $X_{1}$ and $X_{2}$.

Bob decodes in two steps. Bob first decodes $Y_{2}$ to obtain $\mathbb{S}$. He begins by using the Omniscient Adversary decoder to obtain the vector $\tilde{\mathbf{U}}$. He obtains $\tilde{\mathbf{U}}^{\prime}$ from $\tilde{\mathbf{U}}$, by multiplying by $V^{-1}$. He then extracts from $\tilde{\mathbf{U}}^{\prime}$ the $b(C+1)$ symbols corresponding to $\mathbb{S}$. The following claim, proved in the Appendix, ensures that $\mathbb{S}$ is indeed secret from Calvin.

Claim 6: The probability that Calvin guesses $\mathbb{S}$ correctly is at most $q^{-b(C+1)}$, i.e., $\mathbb{S}$ is information-theoretically secret from Calvin.
Thus Alice has now shared $\mathbb{S}$ with Bob. Bob uses $\mathbb{S}$ as the side information used by the decoder of the Shared Secret algorithm to decode $Y_{1}$. This enables him to recover $X_{1}$, which contains Alice's information at rate $R=C-z_{O}$. (There is an asymptotically negligible rate overhead equalling $b^{2} / n+\Delta$. The $b^{2} / n$ term corresponds, as before, to the identity matrix appended to $X$. The term $\Delta$ takes any positive value, and the probability of error also depends on it.) The probability of error equals the sums of the probabilities of error in Theorems 1 and 2. The errors in Theorem 1 are analyzed in Claims 3 and 2. Theorem 2 is used to generate codes of blocklength $\Delta n$. This probability of error is analyzed in Claim 5. Together, an upper bound on the probability of error is $\left(|\mathcal{T}||\mathcal{E}|+n^{C} C\right) q^{-1}+q^{-\Delta n \varepsilon}$. Since the Limited Adversary algorithm is essentially a concatenation of the Shared Secret algorithm with the Omniscient Adversary algorithm, the computational cost is the sum of the computational costs of the two (with $\Delta n$ replacing $n$ as the block-length for the Shared Secret algorithm). This quantity therefore equals $\mathcal{O}\left(n C^{2}+(\Delta n C)^{3}\right)$. Choosing $\Delta$ appropriately (say $\Delta=\left(C^{-\frac{1}{3}} n^{-\frac{2}{3}}\right)$ makes the second term vanish.

## A. Limited Adversary: Universal Codes

We now discuss how to convert the above algorithm to be independent of the network parameters $z_{O}$ and $z_{I}$. Alice's challenge is to design for all possible $z_{O}$ and $z_{I}$ pairs that satisfy the constraint (15). For any specific $z_{I}$, Alice needs to worry only about the largest $z_{O}$ that satisfies (15) because what works

[^0]|  | Adversarial <br> Strength | Rate | Complexity |
| :--- | :--- | :---: | :---: |
| Shared <br> Secret | $z_{O}<C$, | $C-z_{O}$ | $\mathcal{O}\left(n C^{2}\right)$ |
| Omniscient | $z_{I}=$ network |  |  |
| Limited | $z_{I}=$ network |  |  |$\quad \mathcal{z _ { I } + 2 z _ { O } < C}$| $C-z_{O}$ |
| :--- |

TABLE III-Comparison of our three algorithms
against an attacker with a particular traffic injection strength works against all weaker attackers. Note that $C, z_{O}$, and $z_{I}$ are all integers, and thus there are only $C-1$ such attackers. For each of these attackers, Alice designs a different secrecy matrix $X_{2}$ as described above. She appends these $C-1$ matrices to her information $X_{1}$ and sends the result as described in the above section.

To decode Bob needs to estimate which secrecy matrix to use, i.e., which one of them is secret from the attacker. For this he needs a good upper bound on $z_{O}$. But, just as in the omniscient adversary algorithm, he can obtain this by computing the column rank of $Y$, and subtracting $b$ from it. He then decodes using the secrecy matrix corresponding to $\left(z_{O}, C-1-2 z_{O}\right)$. This secrecy matrix suffices since $z_{I}$ can at most be $C-1-2 z_{O}$, which corresponds to Calvin's highest eavesdropping strength for this $z_{O}$.

## X. CONCLUSION

Random network codes are vulnerable to Byzantine adversaries. This work makes them secure. We provide algorithms ${ }^{2}$ which are information-theoretically secure and rate-optimal for different adversarial strengths as shown in Table I. When the adversary is omniscient, we show how to achieve a rate of $C-2 z_{O}$, where $z_{O}$ is the number of packets the adversary injects and $C$ is the network capacity. If the adversary cannot observe everything, our algorithms achieve a higher rate, $C-z_{O}$. Both rates are optimal. Further our algorithms are practical; they are distributed, have polynomial-time complexity and require no changes at the internal nodes.

## Acknowledgments

The authors would like to thank Air Force grant FA9550-06-1-0155, Caltech's Lee Center For Advanced Networking and Microsoft Research for support.

## REFERENCES

[1] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung. Network information flow. IEEE Transactions on Information Theory, 46(5):1204-1216, July 2000.
[2] N. Cai and R. W. Yeung. Secure network coding. In Proceedings of International Symposium in Information Theory and its Applications, Lausanne, Switzerland, June 2002.
[3] N. Cai and R. W. Yeung. Network error correction, part 2: Lower bounds. submitted to Communications in Information and Systems, 2006.
[4] D. Charles, K. Jain, and K. Lauter. Signatures for network coding. In Proceedings of the fortieth annual Conference on Information Sciences and Systems, Princeton, NJ, USA, 2006.
[5] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. Journal of the Association for Computing Machinery, 40(1):17-47, January 1993.
[6] C. Gkantsidis and P. Rodriguez. Network coding for large scale content distribution. In Proceedings of IEEE Conference on Computer Communications (INFOCOM), Miami, March 2005.
[7] C. Gkantsidis and P. Rodriguez. Cooperative security for network coding file distribution. In Proceedings of IEEE Conference on Computer Communications (INFOCOM), Barcelona, April 2006.
${ }^{2}$ A refinement of some of the algorithms in this work can be found in [11].
[8] T. Ho., R. Koetter, M. Médard, D. Karger, and M. Effros. The benefits of coding over routing in a randomized setting. In IEEE International Symposium on Information Theory (ISIT), page 442, Yokohama, July 2003.
[9] T. Ho, M. Médard, J. Shi, M. Effros, and D. Karger. On randomized network coding. In Proceedings of 41st Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, 2003.
[10] 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, Chicago, USA, June 2004.
[11] S. Jaggi and M. Langberg. Resilient network coding in the presence of eavesdropping byzantine adversaries. submitted to ISIT, 2007.
[12] S. Jaggi, M. Langberg, T. Ho, and M. Effros. Correction of adversarial errors in networks. In Proceedings of International Symposium in Information Theory and its Applications, Adelaide, Australia, 2005.
[13] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. IEEE Transactions on Information Theory, 51(6):1973-1982, June 2005.
[14] A. Jiang. Network coding for joint storage and transmission with minimum cost. In Proceedings of International Symposium in Information Theory and its Applications, Seattle, Washington, USA, July 2006.
[15] S. Katti, D. Katabi, W. Hu, H. S. Rahul, and M. Médard. The Importance of Being Opportunistic: Practical Network coding for Wireless Environments. In 43rd Annual Allerton Conference on Communication, Control, and Computing, Allerton, 2005.
[16] S. Katti, H. Rahul, D. Katabi, W. H. M. Médard, and J. Crowcroft. XORs in the Air: Practical Wireless Network Coding. In ACM SIGCOMM, Pisa, Italy, 2006.
[17] R. Koetter and F. Kschischang. Coding for errors and erasures in random network coding. Under Submission.
[18] R. Koetter and M. Médard. Beyond routing: An algebraic approach to network coding. In Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), volume 1, pages 122-130, 2002.
[19] R. Koetter and M. Médard. An algebraic approach to network coding. IEEE/ACM Transactions on Networking, 11(5):782-795, October 2003.
[20] M. N. Krohn, M. J. Freedman, and D. Mazires. On-the-fly verification of rateless erasure codes for efficient content distribution. In Proceedings of the IEEE Symposium on Security and Privacy, 2004, Oakland, California.
[21] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. IEEE Transactions on Information Theory, 49(2):371-381, 2003.
[22] D. Lun, M. Médard, T. Ho, and R. Koetter. Network coding with a cost criterion. In Proceedings of International Symposium in Information Theory and its Applications, October 2004.
[23] D. S. Lun, M. Médard, and R. Koetter. Efficient operation of wireless packet networks using network coding. In International Workshop on Convergent Technologies (IWCT), Oulu, Finland, 2005.
[24] A. Pelc and D. Peleg. Broadcasting with locally bounded byzantine faults. Information Processing Letters, 93(3):109-115, February 2005.
[25] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In Proceedings of the twenty-first annual ACM symposium on Theory of computing, pages 73-85, Seattle, Washington, United States, 1989.
[26] L. Subramanian. Decentralized Security Mechanisms for Routing Protocols. PhD thesis, University of California at Berkeley, Computer Science Division, Berkeley, CA 94720, 2005.
[27] H. J. Wertz. On the numerical inversion of a recurrent problem: The Vandermonde matrix. AC-10:492, Oct. 1965.
[28] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In IEEE Infocom, volume 2, pages 585-594. IEEE, 2000.
[29] R. W. Yeung and N. Cai. Network error correction, part 1: Basic concepts and upper bounds. submitted to Communications in Information and Systems, 2006.

## A. Proof of Claim 2

## Appendix

The idea behind the claim is as follows. The parity-check matrix $P$ is, by construction, a Vandermonde matrix [27], and therefore has full column rank. Further, since $P$ is hidden from Calvin, with high probability he cannot choose interference such that the matrix product $E P$ has a lower column rank than does $E$.

To prove this we use a generalization of an argument used in [12]. Let $S_{i, j}$ denote the $(i, j)$ element of $S=E P$. We note that for each $(i, j), S_{i, j}$ can be thought of as a polynomial in $r_{j}$ with coefficients from the $i^{t h}$ row of $E$. Since $S_{i, j}$ has degree at most $n$ in $r_{j}$, at most
$n$ values of $r_{j}$ satisfy the equation $S_{i, j}=c$, for any scalar $c \in \mathbb{F}_{q}$. Since Calvin does not know the values of the $r_{j} \mathrm{~s}$, the probability he can choose entries in $E$ to satisfy any such equation is at most $n q^{-1}$.

In particular, the probability that the first row of $S$ consists of the length- $C$ zero vector is at most $\left(n q^{-1}\right)^{C}$. For a particular choice of the first row of $S$, the probability that the second row is linearly dependent on the first row (i.e., is any scalar multiple of the first row) is at most $n^{C} / q^{C-1}$. Similarly, the probability that the third row is any of the $q^{2}$ possible linear combinations of the first two rows is at most $n^{C} / q^{C-2}$. Continuing thus, the probability that the $i^{t h}$ row of $S$ is linearly dependent on the previous $i-1$ is at most $n^{C} / q^{C-i+1}$. Taking the union bound over all $C$ events, the probability that $S$ is singular is at most $n^{C} \sum_{i=1}^{C} 1 / q^{C-i+1}$. Since the largest summand equals $n^{C} / q$, therefore the probability of the undesirable event is at most $C n^{C} q^{-1}$. Hence, with probability at least $1-C n^{C} q^{-1}, E$ and $S$ are related via an invertible transformation. Note that $q$ is a design parameter and can be chosen to be much larger than $C n^{C}$ to make the probability of error arbitrarily small.

## B. Proof of Claim 3

The proof of Claim 3 follows directly from [9]. Essentially, it is a consequence of the following facts. First, due to [9], with probability at least $\left(1-|\mathcal{T}| q^{-1}\right)^{|\mathcal{E}|}$ over network code design, $\left[T T_{Z \rightarrow Y}\right]$ has full column rank. Here $|\mathcal{T}|$ is the number of terminals in the multicast connection, and $|\mathcal{E}|$ is the number of (hyper) links in the underlying network. Secondly, the matrix $\left[\hat{T} S^{\prime}\right]$ can be obtained via an invertible transformation from the matrix [ $T T_{Z \rightarrow Y}$ ]. Lastly, for large enough $q$, the quantity $\left(1-|\mathcal{T}| q^{-1}\right)^{|\mathcal{E}|}$ is strictly greater than $1-|\mathcal{T}||\mathcal{E}| q^{-1}$. $\square$

## C. Proof of Claim 4

Rewriting the right-hand side of (12) and substituting for $Y$ from (8) results in

$$
\hat{T} X+T_{Z \rightarrow Y}(Z-L X)=\hat{T}\left[\begin{array}{lll}
0 & F^{X} & I_{b}
\end{array}\right]+T^{\prime \prime}\left[\begin{array}{lll}
I_{z_{O}} & F^{Z} & 0 \tag{16}
\end{array}\right]
$$

Since the columns of $T^{\prime \prime}$ are spanned by the columns of $\left[\hat{T} T_{Z \rightarrow Y}\right]$, therefore we may write $T^{\prime \prime}$ as $\hat{T} M_{1}+T_{Z \rightarrow Y} M_{2}$, where the matrices $M_{1}$ and $M_{2}$ represent the appropriate basis transformation. Thus (16) becomes

$$
\begin{align*}
& \hat{T} X+T_{Z \rightarrow Y}(Z-L X)= \\
& \quad \hat{T}\left(\left[\begin{array}{lll}
0 & F^{X} & I_{b}
\end{array}\right]\right)+\left(\hat{T} M_{1}+T_{Z \rightarrow Y} M_{2}\right)\left[\begin{array}{lll}
I_{z_{O}} & F^{Z} & 0
\end{array}\right] \tag{17}
\end{align*}
$$

Since the vector spaces spanned by the columns of $\hat{T}$ and $T_{Z \rightarrow Y}$ are disjoint (except in the zero vector), therefore we may compare the term multiplying the matrix $\hat{T}$ on both sides of 17 (we may also compare the term corresponding to $T_{Z \rightarrow Y}$, but this gives us nothing useful). This comparison gives us the equation

$$
\hat{T} X=\hat{T}\left[\begin{array}{lll}
0 & F^{X} & I_{b}
\end{array}\right]+\hat{T} M_{1}\left[\begin{array}{ll}
I_{z_{O}} & F^{Z} \tag{18}
\end{array}\right] .
$$

We split the matrix equation (16) into three parts, corresponding to the sub-matrices $X_{1}, X_{2}$ and $X_{3}$ of $X$. Thus (18) now splits into the three equations

$$
\begin{align*}
& \hat{T} X_{1}=\hat{T} M_{1} I_{z_{O}},  \tag{19}\\
& \hat{T} X_{2}=\hat{T} F^{X}+\hat{T} M_{1} F^{Z}, \text { and }  \tag{20}\\
& \hat{T} X_{3}=\hat{T} . \tag{21}
\end{align*}
$$

Equation (21) is trivial, since it only reiterates that $X_{3}$ equals columns of an identity matrix. Equation (19) allows us to estimate that $M_{1}$ equals $X_{1}$. We are finally left with (20), which by substituting for $M_{1}$ from (19) reduces to

$$
\begin{equation*}
\hat{T} X_{2}=\hat{T}\left(F^{X}+X_{1} F^{Z}\right) \tag{22}
\end{equation*}
$$

## D. Proof of Claim 5

For $i=1,2$, we denote by $\tilde{\mathbf{X}}_{i}$ the vector obtained by stacking the columns of $X_{i}$ one after the other. Let $D=\left[\begin{array}{ll}D_{1} & D_{2}\end{array}\right]$, where $D_{2}$ corresponds to the last $b^{2}$ columns of $D$ and $D_{1}$ corresponds to the
remaining columns of $D$. Define $\alpha=n-\left(b+z_{O}\right)$. Denote by $\overrightarrow{F^{X}}$ the vector formed by stacking columns of the matrix $F^{X}$ one after the other, and by $f_{i, j}$ the $(i, j)^{t h}$ entry of the matrix $F^{Z}$. The system of linear equations (13)-(14) can be written in matrix form as

$$
A\binom{\tilde{\mathbf{X}}_{1}}{\tilde{\mathbf{X}}_{2}}=\binom{\hat{T} \overrightarrow{F^{X}}}{-D_{2} \tilde{\mathbf{I}}}
$$

where $A$ is given by

$\left[\right.$| $-f_{1,1} \hat{T}$ | $-f_{2,1} \hat{T}$ | $\ldots$ | $-f_{z_{O}, 1} \hat{T}$ | $\hat{T}$ | 0 | $\ldots$ | $\ldots$ | 0 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $-f_{1,2} \hat{T}$ | $-f_{2,2} \hat{T}$ | $\ldots$ | $-f_{z_{O}, 2} \hat{T}$ | 0 | $\hat{T}$ | 0 | $\ldots$ | 0 |
| $-f_{1,3} \hat{T}$ | $-f_{2,3} \hat{T}$ | $\ldots$ | $-f_{z_{O}, 3} \hat{T}$ | $\vdots$ | 0 | $\hat{T}$ | 0 | 0 |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | 0 | $\ddots$ | 0 |
| $-f_{1, \alpha} \hat{T}$ | $-f_{2, \alpha} \hat{T}$ | $\ldots$ | $-f_{z_{O}, \alpha} \hat{T}$ | 0 | 0 | 0 | 0 | $\hat{T}$ |
| $D_{1}$ |  |  |  |  |  |  |  |  |$]$

This matrix $A$ is described by smaller dimensional matrices as entries. The matrix $\hat{T}$ has dimensions $\left(b+z_{O}\right) \times b$. The $j^{t h}$ row of matrices in the top portion of matrix $A$ describes an equation corresponding to the $j^{\text {th }}$ column of the matrix equation in Equation 13. The bottom portion of $A$ corresponds to Equation 14. Bob can recover the variables $X(i, j)$ if and only if the above matrix $A$ has full column rank. We now analyze $A$ to show that this is indeed the case (with high probability) for sufficiently large $\delta n$. Using Claim 3, we can assume that $\hat{T}$ has full column-rank, and therefore the last $\alpha b$ columns of the matrix (represented by the right side of $A$ ) have full column rank.

We now address the left columns of $A$. Consider performing column operations from right to left, to zero out the $\hat{T}$ s in the left side of the top rows of $A$ (that is, to zero out the upper left sub-matrix of $A$ ). $A$ has full column rank iff after this process the lower left submatrix of $A$ has full column rank. We show that this is the case with high probability over the random elements of $D$ (when $\delta n$ is chosen to be sufficiently large). Let $f_{i j}$ 's be the values appearing in the upper left sub-matrix of $A$. We show that for any (adversarial) choice of $f_{i j}$ 's, with high probability, the act of zeroing out the $\hat{T}$ 's yields a lower left sub-matrix of $A$ with full column rank. Then using the union bound on all possible values of $f_{i j}$ we obtain our assertion.

For any fixed values of $f_{i j}$, let $C(j)$, for $j=1$ to $b z_{O}$, denote the columns of the lower left sub-matrix of $A$ after zeroing out the $\hat{T}$ 's. For each $j$, the vector $C(j)$ is a linear combination of the (lower part of the) $j^{t h}$ column of $A$ with columns from the lower right submatrix of $A$. As the entries of $D_{1}$ are independent random variables uniformly distributed in $\mathbb{F}_{q}$, the columns $C(j)$ for $j=1, \ldots, b z_{O}$ consist of independent entries that are also uniformly distributed in $\mathbb{F}_{q}$. Standard analysis shows that the probability that the columns $C(j)$ are not independent is $q^{b z_{O}-\delta n}$. For the union bound we would like this probability to be at most $q^{-\alpha z_{O}-n \varepsilon}=q^{-\left(n-\left(b+z_{O}\right)\right) z_{O}-n \varepsilon}$. Thus, it suffices to take $\delta n=n\left(z_{O}+\varepsilon\right)$ for an error probability of at most $q^{-n \varepsilon}$. Recall that $b=C-z_{O}$

## E. Proof of Claim 6

The vector $\tilde{\mathbf{U}}$ was generated from $\tilde{\mathbf{U}}^{\prime}$ via an MDS code generator matrix (see Footnote 1), and a folklore result about network codes is that with high probability over random network code design the linear transform between Alice and Calvin also has the MDS property. Thus, for Calvin to infer even a single symbol of the length- $(C-$ $\left.z_{O}-\delta\right) n \Delta$ vector $\tilde{\mathbf{U}}^{\prime}$, he needs to have received at least $\left(C-z_{O}-\right.$ ס) $n \Delta$ linear combinations of the variables in the secrecy matrix $X_{2}$. Since Calvin can overhear $z_{I}$ packets, he has access to $z_{I} n \Delta$ equations that are linear in the unknown variables. The difference between the number of variables unknown to Calvin, and the number of equations Calvin has, is linear in $n \Delta$ - for large enough $n \Delta$, this difference is larger than $b(C+1)$, the length of the vector $\mathbb{S}$. By a direct extension of [25], Calvin's probability of guessing any function of $\mathbb{S}$ correctly is $q^{-b(C+1)}$.


[^0]:    ${ }^{1}$ Secret Sharing protocols [25] demonstrate that using MDS code generator matrices guarantees that to infer even a single symbol of $\tilde{\mathbf{U}}^{\prime}$ from $\tilde{\mathbf{U}}$ requires the entire vector $\tilde{\mathbf{U}}$.

