Network Codes Resilient to Jamming and Eavesdropping

Hongyi Yao, Danilo Silva, Sidharth Jaggi, and Michael Langberg

***Abstract*—We consider the problem of communicating information over a network secretly and reliably in the presence of a hidden adversary who can eavesdrop and inject malicious errors. We provide polynomial-time distributed network codes that are information-theoretically rate-optimal for this scenario, improving on the rates achievable in prior work by Ngai *et al.* Our main contribution shows that as long as the sum of the number of links the adversary can jam (denoted by Z_O) and the number of links he can eavesdrop on (denoted by Z_I) is less than the network capacity (denoted by C) (i.e., $Z_O + Z_I < C$), our codes can communicate (with vanishingly small error probability) a single bit correctly and without leaking any information to the adversary. We then use this scheme as a module to design codes that allow communication at the source rate of $C - Z_O$ when there are no security requirements, and codes that allow communication at the source rate of $C - Z_O - Z_I$ while keeping the communicated message provably secret from the adversary. Interior nodes are oblivious to the presence of adversaries and perform random linear network coding; only the source and destination need to be tweaked. We also prove that the rate-region obtained is information-theoretically optimal. In proving our results, we correct an error in prior work by a subset of the authors in this paper.**

***Index Terms*—Achievable rates, adversary, error control, network coding, secrecy.**

I. Introduction

A SOURCE Alice wishes to multicast information to a set of receivers over a network. As shown in an elegant sequence of papers, *network coding*, i.e., allowing internal nodes in a network to perform nontrivial arithmetic operations on incoming information to generate their outgoing information, in general strictly increases the achievable rate region. An information-theoretic proof of this was provided in [1], and further, the work of [2] demonstrated that linear codes sufficed to achieve such performance. The techniques in [3] demonstrated an explicit procedure to design such codes over finite fields. Efficient code constructions were provided in [4] (deterministic codes) and in [5] (distributed randomized codes). An extensive account of the theory and applications of network coding can be found in [6].

However, if a network with even a single receiver Bob contains a malicious adversary Calvin, there are at least two security challenges—Calvin might eavesdrop on private communications, or he might disrupt communications by injecting fake information into the network, or in general he might do both. In the network coding model, this second danger may be even more pronounced since all nodes, including honest ones, mix information. In this case, even a small number of fake packets injected by Calvin may end up corrupting *all* the information flowing in the network, causing decoding errors. In particular, Calvin may use his knowledge of the network topology so as to sit in the bottleneck of the network, and thereby jam communications in a network location where it might do the most damage. Also, Calvin's eavesdropping capabilities have negative security implications in two ways. First, Calvin is able to eavesdrop and thus infer something about Alice's secret message to Bob. Second, Calvin might also be able to use the eavesdropped information to carefully design his jamming pattern so as to make it hard for Bob to correctly decode Alice's message.

In this paper, we consider the *secrecy* and *error control* issues together. Namely, we design schemes that allow reliable network communications in the presence of an adversary that can both jam and eavesdrop, without leaking information to him. In particular, suppose the network's min-cut from Alice to Bob is C, and Calvin eavesdrops on Z_I links and corrupts Z_O links.[1] Our main contribution is in the demonstration of schemes that are distributed, computationally efficient to design and implement, and can be used to communicate a *single* bit secretly and without error. We then use this scheme as a tool to improve on prior work [7] and achieve a provably optimal communication rate of $C - Z_O$ when no secrecy constraints are posed and a rate of $C - Z_O - Z_I$ when communication is kept secret from Calvin. In particular, the overall rate of communication thus achievable in the presence of Calvin's *adversarial* jamming is the same as if the jamming behaved like *random noise*. A preliminary version of the results in this paper was presented in [8].

Manuscript received February 27, 2012; revised November 12, 2012 and August 26, 2013; accepted October 28, 2013; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor D. Goeckel. Date of publication February 03, 2014; date of current version December 15, 2014. The work of D. Silva was supported in part by the Brazilian National Council for Scientific and Technological Development (CNPq) under Grant 482131/2010-1. The work of S. Jaggi was supported in part by the University Grants Committee of the Hong Kong Special Administrative Region, China, under Project No. AoE/E-02/08, the Microsoft-CUHK Joint Laboratory for Human-centric Computing and Interface Technologies under a grant, and an SHIAE grant. The work of M. Langberg was supported in part by ISF Grant 480/08 while the author was at The Open University of Israel.

H. Yao was with the California Institute of Technology, Pasadena, CA 91125 USA. He is now with Oracle, Inc., Redwood City, CA 94065 USA (e-mail: yaohongyi03@gmail.com).

D. Silva is with the Department of Electrical Engineering, Federal University of Santa Catarina (UFSC), Florianópolis, SC 88040-900, Brazil (e-mail: danilo@eel.ufsc.br).

S. Jaggi is with The Chinese University of Hong Kong, Hong Kong (e-mail: jaggi@ie.cuhk.edu.hk).

M. Langberg is with the State University of New York at Buffalo, Buffalo, NY 14260 USA, and also with The Open University of Israel, Ra'anana 43107, Israel (e-mail: mikel@openu.ac.il).

Digital Object Identifier 10.1109/TNET.2013.2294254

[1]We consider a model where network links rather than nodes are eavesdropped and corrupted; eavesdropping on a node is equivalent to eavesdropping on links incoming to it, and corrupting a node is equivalent to corrupting the links outgoing from it.

A. Prior Work

Related problems have been considered in the past. Prior results may be classified in the following three categories. An extensive discussion on the field of security problems for networks performing network coding can be found in [6, Ch. 7].

1) *Secrecy*: For networks containing adversaries that only eavesdrop on some links (without jamming transmissions), the work of [9] provided a tight information-theoretic characterization of the *secrecy capacity*, i.e., the optimal rate achievable without leaking any of Alice's information to Calvin. Efficient schemes achieving this performance were proposed by [10]–[13]. Cryptographically (but not information-theoretically) secret schemes for this scenario were also considered in [14].
2) *Error-control*: For networks containing adversaries with unlimited eavesdropping capabilities and limited jamming capabilities, prior related work has focused primarily on the detection of Byzantine errors (i.e., a bounded number of arbitrary—rather than random—errors) [15], nonconstructive bounds on the achievable *zero-error* rates [16], [17], and network error-correcting codes [18] (which have high design complexity) and [7] and [19]–[21] (which have low design complexity). Results for this setting are also available under cryptographic assumptions [22], [23].
3) *Secrecy + Error-control*: The scenario closest to the one considered in this paper, with limitations on both Calvin's eavesdropping power Z_I and his jamming power Z_O, have been considered in [7], [21], and [24]–[26]. Here, there are two questions one could consider.
   a) What is the maximum rate at which one can communicate reliably (without caring about hiding one's data from Calvin)?
   b) What is the maximum rate at which one could communicate *both* secretly and reliably?

   For Model 3b, under the requirement of *zero* error probability, the maximum rate of secret and reliable communication is given by $C-2Z_O-Z_I$. Schemes achieving this rate have been proposed in [25], [26] (high design complexity schemes), and [13], [24], and [27] (low design complexity schemes). The optimality of such a rate has been shown in [26] for single-letter coding and in [27] for block coding.

   For Model 3a, if the requirement of zero error probability is relaxed to *vanishingly small* error probability, as considered here, then higher rates than $C-2Z_O-Z_I$ may be achieved. In particular, the work in [7] provided computationally efficient communication schemes at rate $C-Z_O$ as long as the restrictive requirement $C > 2Z_O + Z_I$ was satisfied. Work by a subset of the authors of this paper claimed in [21] to improve this requirement to $C > Z_O + Z_I$. As we demonstrate in Section VIII, the prior proof of the claim was incorrect, and Section II gives a correct proof of the claim.

In this paper, for Model 3a, we present a communication scheme of rate $C - Z_O$ as long as the (improved) requirement $C > Z_O + Z_I$ is satisfied. For Model 3b, we obtain the rate (that we prove to be optimal) of $C - Z_O - Z_I$. These two results together complete the story for a line of work begun in [7]. Our results are obtained by combining ideas in [7] with the secrecy/"data-hiding" scheme of [13].

To put our results in perspective (in particular to compare with the most relevant prior work [7]), consider the following two scenarios that are specific examples of our general result.

1) Suppose $Z_O = Z_I$ (a "realistic" scenario, corresponding to the adversary being able to eavesdrop on links it can jam) and denote these quantities both by Z. Prior work [7], for either Model 3a or Model 3b, would apply only in the regime where Z is less than a third of the min-cut C. In contrast, our current work demonstrates that communication is possible (in either 3a or 3b) as long as Z is less than half the min-cut C.
2) An even more extreme example is as follows. Suppose $Z_I = 0$ (also a "realistic" scenario, corresponding to a "blind" adversary—one who cannot base his jamming function on what is being actually transmitted. In a wireless setting, this may happen perhaps because he has only one antenna). In this scenario, our schemes achieve a positive rate for *any* $Z_O < C$, whereas prior work [7] would be restricted to the setting wherein Z_O is less than half the min-cut C.

II. Main Results

The main results of this work are Theorems 1–3. Let q be the size of the finite field over which the network code operates, and let n be the block-length (number of symbols over $\mathbb{F}_q$) of the packets transmitted over the network.

Theorem 1: If $C > Z_O + Z_I$, then Alice can communicate a single bit correctly to Bob (while keeping it secret from Calvin) using codes of computational complexity $O(\mathrm{poly}(C, \log_2 q))$ and error probability $O(q^{-C})$.

Combining the codes in Theorem 1 with the "shared-secret" codes in [7] gives us Theorems 2 and 3. Roughly, we say that a message at a certain rate is "robustly achievable" if there exists a communication scheme that allows Bob to decode Alice's message (with high probability). A more formal definition follows in Section III.

Theorem 2: No rate higher than $C - Z_O$ is robustly achievable. If $C > Z_O + Z_I$, then a rate of $C - Z_O$ is robustly achievable with codes of computational complexity $O(n.\mathrm{poly}(C, \log_2 q))$.

We say that a message at a certain rate is "secretly and robustly achievable" if there exists a communication scheme that keeps Alice's message information-theoretically secure from Calvin, and simultaneously allows Bob to decode Alice's message (with high probability). A more formal definition follows in Section III.

Theorem 3: No rate higher than $C - Z_O - Z_I$ is secretly and robustly achievable. A rate of $C - Z_O - Z_I$ is secretly and robustly achievable with codes of computational complexity $O(n.\mathrm{poly}(C, \log_2 q))$.

Note that achievability in Theorem 2 does not immediately imply the achievability in Theorem 3 since the concatenation of an (outer) secrecy scheme with an (inner) error control scheme may not necessarily be secure (see, e.g., [13]). As will become clear later, such an implication (in fact, an equivalence) will follow from the fact that the secrecy and error control schemes that we use both have a linear structure, which naturally ensures their compatibility.

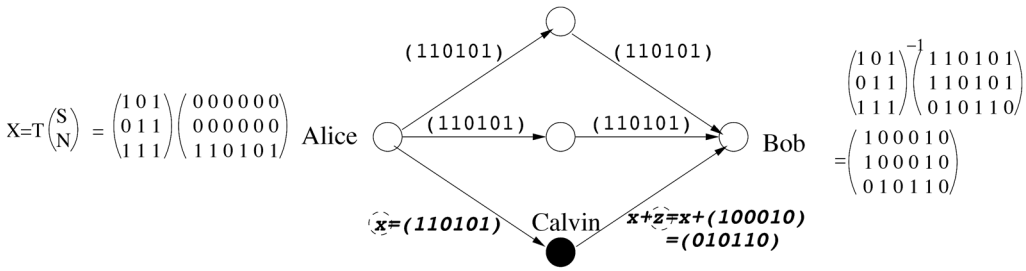

(a) Alice appends a rank−0 codeword S (corresponding to message 0) with a random N, mixes with T, and transmits rows of the resulting rank−1 matrix on her outgoing edges. (Arithmetic over the binary field)
Calvin observes x on his incoming link, and jams by adding z on his outgoing link.
Bob decodes by first inverting the effect of T on his received vectors, and noting that the rows corresponding to S are rank−deficient (rank−1)

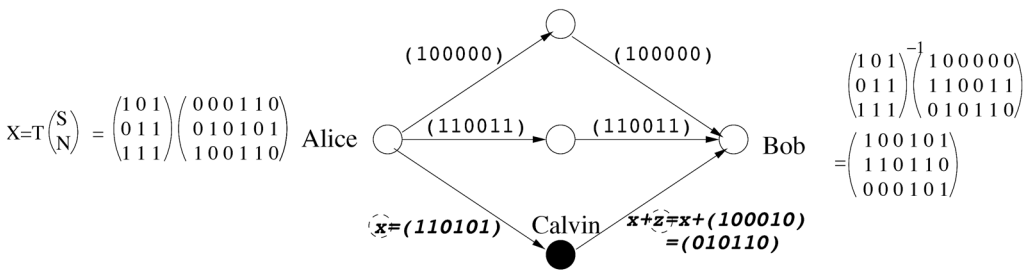

(b) Alice appends a rank−2 codeword S (corresponding to message 1) with a random N, mixes with T, and transmits rows of the resulting rank−3 matrix on her outgoing edges. (Arithmetic over the binary field)
Calvin observes x on his incoming link, and jams by adding z on his outgoing link.
Bob decodes by first inverting the effect of T on his received vectors, and noting that the rows corresponding to S are full−rank (rank−2)

Fig. 1. Toy example demonstrating how to share a single bit secretly and robustly: This example corresponds to the "Secret-sharing layer" referenced later in Fig. 3 and then in Section V. In this example, $C = 3$, and $Z_O = Z_I = 1$ (hence, $C > Z_O + Z_I$). The block-length n'' used in this example equals 6. Alice just wants to share a single bit with Bob secretly and reliably—if the bit equals 0, she uses the scheme in (a), else she uses the scheme in (b). Bob decodes by checking the rank of the received matrix. Note that we assume that Alice and Bob know the *value* of Z_O and Z_I in advance, though not Calvin's location—this is analogous to having a prior estimate of channel conditions.

A. Toy Example

We first begin with a toy example demonstrating the main ideas behind our central result, *viz.* the modular scheme that Alice and Bob use to communicate a single bit correctly, and without leaking information to Calvin. In our example network, Alice communicates with Bob over a network that comprises three parallel paths passing through distinct relays. Honest relays simply forward incoming data. One of the relays is controlled by the adversary Calvin, who therefore knows the transmission on the incoming link, and can corrupt the corresponding outgoing transmission arbitrarily. The identity of the node controlled by Calvin is unknown to Alice and Bob. In a slightly generalized version of this example, Calvin may "control" $Z = Z_I = Z_O$ relays (i.e., observes the packets incoming to Z relays and so $Z_I = Z$, and based on these observations corrupts the packets outgoing from these relays and so $Z_O = Z$), out of a total of C relays.

The scheme demonstrated in Fig. 1 is as follows. Let the block-length $n'' > C$ be a design parameter chosen by Alice and Bob so as to guarantee performance. The matrix T is known to all parties and is designed as an invertible matrix whose last Z_I columns are the transpose of a generator matrix of an MDS code. (The design of T is explained in Section IV-B.)

If Alice wishes to transmit a 0 to Bob, she transmits a "random low-rank codeword" over the parallel links. That is, she appends a $(C - Z_I) \times n''$ zero-matrix S to a random (hence full-rank with high probability) $Z_I \times n''$ matrix N. She "mixes" the rows of the resulting matrix by premultiplying it with T and transmits the resulting rows over her outgoing links [in Fig. 1(a), this corresponds to sending a single vector repeatedly over the different links].

Conversely, if Alice wishes to transmit a 1 to Bob, she transmits a "random high-rank codeword" over the parallel links. That is, she chooses S to be a random (hence full-rank with high probability) $(C - Z_I) \times n$ matrix and appends a random $Z_I \times n''$ matrix N. She again mixes this with T and transmits the resulting row-vectors over her outgoing links [in Fig. 1(b), this corresponds to sending three linearly independent vectors over the different links].

Note that Calvin knows the matrix T, as does Bob—it is part of code design. For this reason, T must be carefully designed, so that on observing any Z_I rows/packets, Calvin cannot distinguish between a zero bit-matrix and a high-rank bit-matrix.

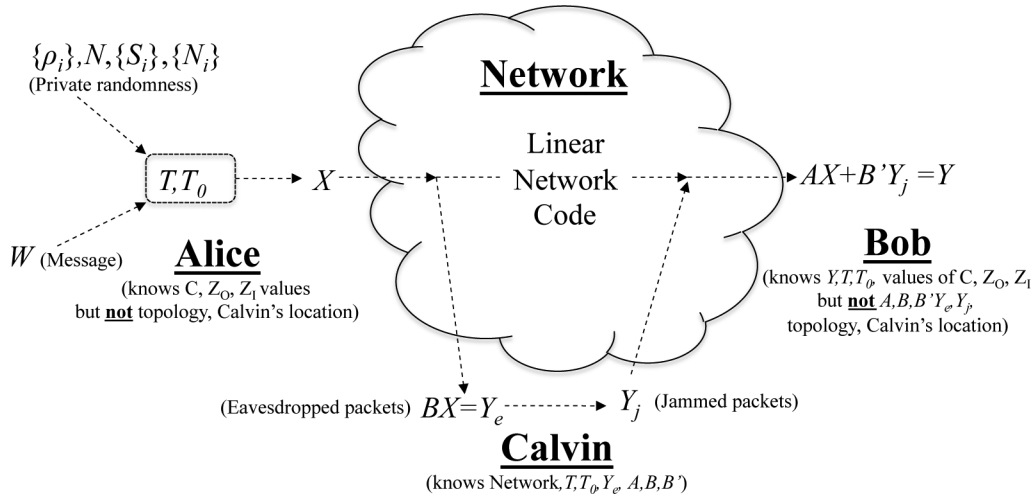

Fig. 2. System diagram: a pictorial representation of the system model described in Section III . Alice uses private randomness (known only to her) and "mixing" matrices T, T_0 (known to the other parties, Bob and Calvin) to encode her message W (of rate $R = C - Z_I - Z_O$) to X. (This encoding procedure is further detailed in Figs. 3 and 4, and in detail in Sections V and VI.) This X is transmitted over the "Network channel." Calvin overhears packets Y_e and, based on these, attempts to reconstruct W and also to "jam" transmission to Bob by injecting jamming packets Y_j. In addition to the packets Y_e he overhears, Calvin knows Alice's mixing matrices T, T_0, the network topology, and the network transforms A, B, and B' (respectively from Alice to Bob, Alice to Calvin, and Calvin to Bob)—Calvin's choice of Y_j may be a function of all of these.

(For instance, if the last column of T had a zero on the ith row, Calvin could easily distinguish between the two cases by observing the ith packet.)

Bob's decoding rule is straightforward. He first premultiplies the $C \times n''$ matrix with the inverse of T in an attempt to recover S and N. If the resulting row vectors corresponding to the locations of S are "rank-deficient" (of rank strictly less than $C - Z_I$), he decodes to a 0, otherwise he decodes to a 1.

To show that this scheme does not enable Calvin to estimate Alice's message bit, note that from his perspective the distribution over the messages he eavesdrops is identical regardless of Alice's message bit—in both cases, his observed packet is uniformly distributed over all length-n'' vectors.

To show that the scheme *also* enables Bob to decode Alice's 0 or 1 message correctly, with high probability, *regardless* of Calvin's adversarial jamming action, we proceed as follows. Note that Calvin has no information on what vectors are being transmitted on links not controlled by him. Hence, even though he can transmit arbitrary vectors on the links controlled by him, the probability that these vectors are linearly dependent on the other vectors (on links *not* controlled by him) is quite small (exponentially small in the block-length, and the block-length can be chosen to be large, to guarantee probability of success arbitrarily close to 1).[2]

Refining the ideas presented in this toy example to the general scenario requires several nontrivial extensions. Details of these extensions are in the following sections, but we summarize these here.

First, Alice and Bob need to design a *distributed* scheme that operates even when they are ignorant of network topology prior to communication. This requires that Alice's message bit remains secret from Calvin even if he receives random linear combinations of Alice's transmissions (rather than the specific vectors she injects into the network). It also requires that Calvin's injected jamming vectors be linearly independent of other randomly linearly combined vectors. Since the linear transforms applied by the network need not preserve the uniform probability distribution that Alice imposes on her transmitted vectors, a more delicate analysis is needed.

Second, Bob does not in general know the linear transform imposed by the network. To circumvent this problem, the "subspace metric" codes introduced by Kötter *et al.* [20] prove quite useful.

Lastly, we note that the ideas above can really only be used to transmit a "few" bits from Alice to Bob. This is because each use of the scheme requires Alice to send a somewhat bulky matrix simply to communicate a single bit to Bob, and if the scheme is repeated too many times, then the throughput of Alice's message goes down considerably. Fortunately, a "shared-secret" algorithm presented in [7] enables us to guarantee high-rate secret communication from Alice to Bob, as long as Alice can share even just a "few" bits secretly with Bob. We thus use our single-bit sharing scheme as a module for the shared-secret algorithm to obtain the desired result.

III. Network Model and Problem Statement

We use the general model proposed in [7] and pictorially represented in Fig. 2. To simplify notation, we consider only the problem of communicating from a single source to a single destination.[3]

A. Network Model

Alice communicates to Bob over a network with an attacker (adversary) Calvin hidden somewhere in it. Calvin aims to

[2]Note the following asymmetry: When Alice sends bit 0, Bob never makes an error; he makes an error (with small probability) if and only if bit 1 is sent and the received matrix is not full-rank. The reason for this asymmetry is as follows: If Alice's secret bit is 0, then the rank of the transmitted message is Z_I, and hence the maximum rank of the received message is $Z_I + Z_O < C$. However, in this case, by Bob's decoding rule, he (correctly) outputs a 0. On the other hand, if the secret bit is 1, it is possible (though "unlikely") that the packet injected by Calvin is able to lower the rank of the matrix Bob uses to decode.

[3]Similarly to many network coding algorithms, our techniques generalize to multicast problems.

disrupt the transfer of information from Alice to Bob and in the meantime to eavesdrop on the information Alice sends. He can observe some of the transmissions and can inject his own fake transmissions.

Calvin is computationally unbounded and knows the encoding and decoding schemes of Alice and Bob and the network code implemented by the interior nodes. He also knows the network topology, and he gets to choose which network links to eavesdrop on and which ones to corrupt.

The network is modeled as a directed and delay-free graph whose edges each have capacity equal to one symbol of a finite field of size q, $\mathbb{F}_q$, per unit time.[4] All computations are over $\mathbb{F}_q$. The *network capacity*, denoted by C, is the *min-cut from source to destination* .[5]

Each packet contains n symbols from $\mathbb{F}_q$. Alice's message is denoted $W \in \mathcal{S}$. To send this to Bob over the network, Alice encodes W into a matrix $X \in \mathbb{F}_q^{C\times n}$, possibly using a *stochastic encoder*.[6] The ith row in X is Alice's ith packet. As in [5], Alice and internal nodes take random linear combinations of their observed packets to generate their transmitted packets.

Analogously to how Alice generates X, Bob organizes received packets into a matrix Y. The ith received packet corresponds to the ith row of Y. The random linear network code used by Alice and all internal nodes induces a linear transform A from X to Y, such that $Y = AX$ when no error is induced by the adversary.[7] Thus, Y is a matrix in $\mathbb{F}_q^{C\times n}$, and $A \in \mathbb{F}_q^{C\times C}$. Hereafter, we assume that the matrix A is invertible, which happens with high probability if q is sufficiently large [5].

Calvin can eavesdrop on Z_I edges, and can inject (possibly fake) information at Z_O locations,[8] in the network. The matrix received by Bob is then $Y = AX + Z$, where Z corresponds to the information injected by Calvin as seen by Bob. Note that the limitation of Calvin's jamming capacity implies that $\mathtt{rank}(Z) \le Z_O$; in particular, Z can be thought of as $B'Y_j$, where Y_j correspond to the (at most Z_O) packets injected by Calvin, and $B' \in \mathbb{F}_q^{C\times Z_O}$ as the linear transform imposed by the network between Calvin and Bob. Similarly, Calvin's observation can be described as a matrix $Y_e = BX$, where $B \in \mathbb{F}_q^{Z_I\times C}$ is the linear transform undertaken by X as seen by Calvin. Both B and B' are known to Calvin *a priori*, but not to Alice or Bob (since Calvin is hidden). Neither Alice nor Bob are assumed to know the network transform, or indeed even the network topology prior to the commencement of communication, though Calvin is allowed to know these. However, Alice and Bob are both assumed to know the *values* of Z_O, Z_I, and C (or good upper bounds on these) since these are critical for them to decide the rates at which to transmit. This is analogous to the encoder/decoder pair needing to "estimate channel conditions" before deciding what rate/code to use.

[4]For ease of presentation, edges with nonunit capacities are not considered here (as in [7], they may be modeled via block coding and parallel edges).

[5]For the corresponding multicast case, C is defined as the minimum of the min-cuts over all destinations. It is well known that C also equals 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*.

[6]The random coin tosses made by Alice as part of her encoding scheme are not known to either Calvin or Bob.

[7]For the ease of notation, we assume Bob removes redundant incoming edges so that the number of edges reaching Bob equals the min-cut capacity C from Alice to Bob.

[8]We assume throughout that the information injected into the network by Calvin is *added* to the original information transmitted (here we consider addition over our field $\mathbb{F}_q$).

B. Problem Statement

Alice wishes to communicate with Bob with perfect secrecy and vanishingly small error probability. That is, Alice's scheme is *perfectly secret* if

$$I(\mathbf{W};\mathbf{Y}_e) = 0 \qquad \forall B \in \mathbb{F}_q^{Z_I\times C} \tag{1}$$

i.e., Calvin obtains no information about Alice's message regardless of which Z_I links he eavesdrops. The *error probability* is the probability (over all randomness introduced by Alice and Calvin) that Bob's reconstruction $\hat{W}$ of Alice's information W is inaccurate, i.e., $\Pr[\hat{W} \neq W]$. We consider the error probability of the worst-case scenario.[9] Namely, a scheme has error probability less than ϵ if $\Pr[\hat{W} \neq W] < \epsilon \quad \forall A, Z$, where A is assumed to be nonsingular, and $\mathtt{rank}(Z) \le Z_O$. The *rate* R of a scheme is the number of information bits of information Alice transmits to Bob, amortized by the size of a packet in bits, i.e., $R = \frac{1}{n}\log_q |\mathcal{S}|$. The rate R is said to be *secretly and robustly achievable* if for any $\epsilon > 0$, any $\delta > 0$, and sufficiently large n, there exists a perfectly secret block-length-n network code with rate at least $R - \delta$ and a probability of error less than ϵ. The rate R is said to be *robustly achievable* if for any $\epsilon > 0$, any $\delta > 0$, and sufficiently large n, there exists a block-length-n network code (which need not be perfectly secret) with rate at least $R - \delta$ and a probability of error less than ϵ.

IV. Auxiliary Tools

A. Mapping Between Finite Fields

We first define a mapping commonly used in the network error-correcting code literature (for example, see [13]) that maps between a vector space over a finite field and a corresponding extension field of the minute field. This mapping helps us translate between the field over which the internal nodes in the network perform network coding ($\mathbb{F}_q$) and the field over which the end-to-end codes operate ($\mathbb{F}_Q$).

Let $Q = q^C$, and let $\mathbb{F}_Q$ be an extension field of $\mathbb{F}_q$. Let $\phi : \mathbb{F}_Q \to \mathbb{F}_q^{1\times C}$ be a vector space isomorphism. In addition, let $\phi_{m,n} : \mathbb{F}_Q^{m\times n} \to \mathbb{F}_q^{m\times Cn}$ be a vector space isomorphism such that the ith row of $\phi_{m,n}(X)$ is given by $[\,\phi(X_{i,1}) \;\cdots\; \phi(X_{i,n})\,]$. In other words, we expand each element of $X \in \mathbb{F}_Q^{m\times n}$ as a length-C row vector over $\mathbb{F}_q$ (with the number of columns in matrix increasing accordingly). We will omit the subscript from $\phi_{m,n}$ when the dimensions of the argument are clear from the context.

Throughout, without loss of generality, we assume that n is divisible by C.

B. Secrecy/Data-Hiding Coding

For generality, let b (rather than C) denote the number of packets transmitted by Alice (so that the description below can apply to both cases $b = C$ and $b = C - Z_O$ as needed). Consider a special case of the problem where Calvin can eavesdrop on $Z_I < b$ packets, but *cannot* jam any packets. Below, we review a construction of a perfectly secret end-to-end scheme that

[9]Our interest is to design communication schemes that do not rely on the specific network topology or network code used.

asymptotically achieves the maximum possible rate (i.e., the secrecy capacity) $b - Z_I$ for this problem. The scheme, proposed in [13], is based on *Maximum Rank Distance* (MRD) codes. (For more details on MRD codes, see [13], [19], and [28].)

Recall that $Q = q^C$. Let $H \in \mathbb{F}_Q^{(b-Z_I)\times b}$ be the parity-check matrix of a $[b, Z_I]$ linear MRD code over $\mathbb{F}_Q$. Let $T \in \mathbb{F}_Q^{b\times b}$ be an invertible matrix chosen such that the first $b - Z_I$ rows of T^{-1} are equal to H. Equivalently, T should be chosen as any invertible matrix such that the last Z_I rows of T^T form a generator matrix of a $[b, Z_I]$ linear MRD code over $\mathbb{F}_Q$ [13, Proposition 9]. Such a matrix T can always be found since MRD codes exist for all parameters provided $b \leq C$ [13], [28].

Alice's encoding proceeds as follows. She first generates a random matrix $N \in \mathbb{F}_Q^{Z_I\times \bar{n}}$ uniformly and independently from any other variables. Then, she computes $\bar{X} = [\, I_b \quad \phi(M)\,]$, where $M = T\begin{bmatrix} W \\ N \end{bmatrix}$. Alice thus encodes a given message $W \in \mathbb{F}_Q^{(b-Z_I)\times \bar{n}}$, where $\bar{n} = n/C - 1$.

Assuming Bob receives $Y = AX = [\, A \quad A\phi(M)\,]$, Bob computes $X = A^{-1}Y$ to recover $M = \phi^{-1}(\phi(M))$. Then, Bob can directly obtain W since, by construction, $W = HM$.

Recall that Calvin's observation is given by $Y_e = BX$, where $B \in \mathbb{F}_q^{Z_I\times b}$. According to [13, Theorem 7], we have that $I(\mathbf{W}; \mathbf{Y}_e) = 0$ for all $\bar{B}$, and therefore (1) is satisfied. Thus, the scheme is indeed perfectly secret.

The decoding complexity is given by $O(\bar{n}b^2)$ operations over $\mathbb{F}_Q$, which can be done in $O(nC^4)$ operations over $\mathbb{F}_q$.

C. Error Control Under a Shared Secret Model

Consider now a second scenario, wherein Calvin can jam $Z_O < C$ packets and eavesdrop *any* number of packets he chooses. However, we posit the existence of a "low-rate side-channel," which Calvin cannot access, that enables Alice to transmit to Bob a "small" secret $\mathbb{S}$ (of size asymptotically negligible compared to Alice's message). We also drop the requirement of secret communication, i.e., all we require is that Bob can decode Alice's transmission correctly, with high probability over the transmissions on the side channel. Below, we review a coding scheme presented in [7] that can asymptotically achieve the maximum possible rate (the so-called *shared-secret capacity*) $C - Z_O$ for this case.

Let b denote $C - Z_O$. We first describe how Alice produces the secret bit string $\mathbb{S}$ based on a given message $\phi(M) \in \mathbb{F}_q^{b\times(n-b)}$. To begin with, she generates $\alpha = bC + 1$ symbols $\rho_1, \rho_2, \ldots, \rho_\alpha \in \mathbb{F}_q$ independently and uniformly at random. Let $P \in \mathbb{F}_q^{n\times\alpha}$ be the matrix given by $P_{(i,j)} = (\rho_j)^i$. Then, she computes a matrix $\mathbb{H} = \bar{X}P \in \mathbb{F}_q^{b\times\alpha}$, where $\bar{X} = [\, I_b \quad \phi(M)\,]$. The tuple $(\rho_1, \rho_2, \ldots, \rho_\alpha, \mathbb{H})$, consisting in total of $\alpha(b+1)$ symbols in $\mathbb{F}_q$, comprises the message "hash" that should be secretly transmitted to Bob. The bit representation of this tuple yields the string $\mathbb{S} \in \{0,1\}^k$, consisting of $k = \alpha(b+1)\log_2 q$ bits. Over the main channel, Alice transmits the $C \times n$ matrix $X = \begin{bmatrix} \bar{X} \\ 0 \end{bmatrix} = \begin{bmatrix} I_b & \phi(M) \\ 0 & 0 \end{bmatrix}$.

Assuming that $(\rho_1, \rho_2, \ldots, \rho_\alpha, \mathbb{H})$ is secretly and correctly received by Bob, let us proceed to the description of Bob's decoder. First, Bob reconstructs the matrix P. Bob obtains $Y = AX + Z$, where $Z \in \mathbb{F}_q^{C\times n}$ has rank at most Z_O. This can also be written as $Y = \tilde{A}\bar{X} + Z$, where $\tilde{A}$ consists of the first b columns of A. Let $\bar{Y}$ be the reduced row echelon form of Y. It is shown in [7] that, with probability at least $1 - O(1/q)$ for any fixed network, $\bar{X}$ can be written as $\bar{X} = U\bar{Y}$ for some $U \in \mathbb{F}_q^{b\times C}$. It is also shown in [7] that, with probability at least $1 - n^\alpha/q$, the system $U\bar{Y}P = \mathbb{H}$ has a unique solution in U. Bob solves this system to find U, computes $\bar{X} = U\bar{Y}$, and finally recovers $\phi(M)$.

Overall, the probability of error of the scheme is at most $n^\alpha/q + O(1/q) = O(n^{C^2}/q)$, while the decoding complexity is $O(nC^3)$ operations in $\mathbb{F}_q$.

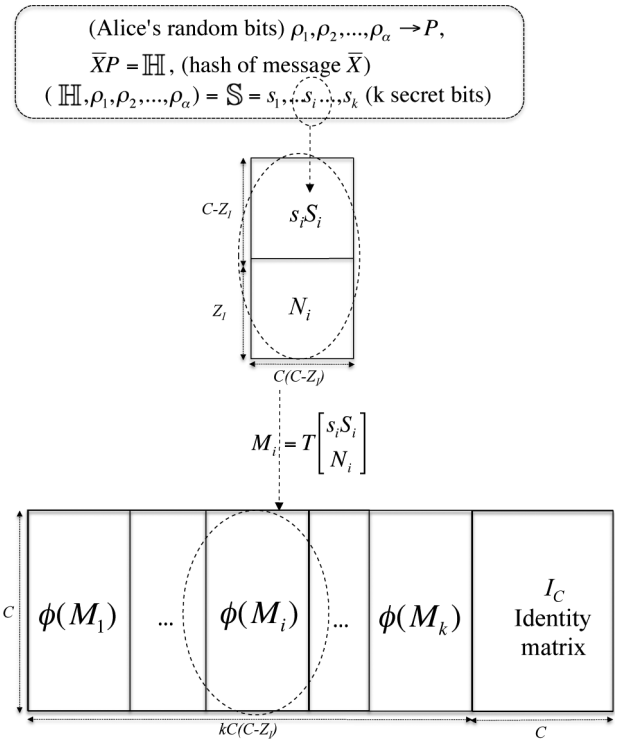

Fig. 3. Secret-sharing layer: (Described in detail in Section V.) Alice first generates a secret hash of her message W as follows. She chooses α symbols $\rho_1, \ldots, \rho_\alpha$ uniformly at random from $\mathbb{F}_q$, uses these symbols to generate a random parity-check matrix P, and uses P to generate a hash $\mathbb{H}$ of her secrecy-encoded/masked message $\bar{X}$ (generated as described in Fig. 4). The bits of $\mathbb{H}$ and $\rho_1, \ldots, \rho_\alpha$ together comprise her k secret bits $\{s_1, \ldots, s_k\} = \mathbb{S}$. She then uses these bits to encode bit-matrices s_iS_i—if s_i is a 0-bit, then s_iS_i is a zero-matrix, else it is a random full-rank matrix. As in the secrecy-encoding in Fig. 4, each of these bit-matrices is mixed with a random matrix N_i. The resulting mixed matrices (translated to $\mathbb{F}_q$ via $\phi(.)$), along with the "standard" identity matrix header used in, for instance, [5], comprise the secret-sharing layer.

V. Sending a Single Bit Secretly and Reliably

Let $C' = C - Z_I$. In this section, we show how Alice can transmit a secret bit s_i reliably to Bob when $C > Z_I + Z_O$. We assume that the block-length n'' for this single-bit scheme is $C(1 + C')$, as this is the smallest packet-length required for the scheme to work. Larger block-lengths can be handled by zero-padding the transmitted packets. A summary of the coding scheme (in which several secret bits are transmitted between Alice and Bob) is presented in Fig. 3.

Let $T \in \mathbb{F}_Q^{C\times C}$ and $H \in \mathbb{F}_Q^{C'\times C}$ be as given in Section IV-B with $b = C$.

A. Alice's Encoder

Initially, Alice chooses a matrix S_i uniformly at random from full-rank $\mathbb{F}_Q^{C'\times C'}$. If her secret bit s_i is 1, s_iS_i is nonzero; otherwise, if s_i is 0, $s_iS_i = 0$. Then, she sends s_iS_i to Bob using the data-hiding scheme described in Section IV-B. More precisely, she transmits $X = [\, I_C \quad \phi(M_i)\,]$, where $M_i = T\begin{bmatrix} s_iS_i \\ N_i \end{bmatrix}$ and $N_i \in \mathbb{F}_Q^{Z_I\times C'}$ is a uniformly random matrix chosen independently from S_i.

B. Bob's Decoder

For convenience, let $S = s_iS_i$ for the remainder of this section. Recall that Bob receives a matrix $Y = AX + Z$, where $A \in \mathbb{F}_q^{C\times C}$ is nonsingular and $Z \in \mathbb{F}_q^{C\times C(1+C')}$ has rank at most Z_O. Let $\bar{Y}$ denote the reduced row echelon form of Y. Consider first the case where $\bar{Y} = [\, I \quad \phi(r)\,]$, for some $r \in \mathbb{F}_Q^{C\times C'}$. It is possible to show that $Hr = S + E$, where $E \in \mathbb{F}_Q^{C'\times C'}$ is a matrix of rank at most Z_O. As will be shown later, with high probability, Hr is full-rank if and only if Alice's secret bit is 1. Thus, Bob can decode by computing the rank of Hr.

In general, however, $\bar{Y}$ may not have the form described above. Nevertheless, as shown in [19, Proposition 10] and [24, Ch. 5], it is possible to extract from $\bar{Y}$ some matrices $r \in \mathbb{F}_Q^{C\times C'}$, $\hat{L} \in \mathbb{F}_q^{C\times\mu}$ and $\hat{V} \in \mathbb{F}_Q^{\nu\times C'}$ such that

$$r = x + \hat{L}V^1 + L^2\hat{V} + L^3V^3 \tag{2}$$

for some $V^1 \in \mathbb{F}_Q^{\mu\times C'}$, $L^2 \in \mathbb{F}_q^{C\times\nu}$, $L^3 \in \mathbb{F}_q^{C\times f}$, and $V^3 \in \mathbb{F}_Q^{f\times C'}$. The matrices r, $\hat{L}$, and $\hat{V}$ can be obtained by converting $\bar{Y}$ to reduced row echelon form (see [24, Sec. 5.1.2]) and therefore are *known* to the decoder. The last three terms in (2) may be seen as generalized errors terms, as some of its factors ($\hat{L}$ and $\hat{V}$) are known. Note that a partially known error term is analogous to an erasure in classical coding theory (where the location of the error, but not its value, is known) and has the same effect of enabling the decoder to correct more errors than if such variables were unknown.

Additionally, it is shown in [24, Theorem 5.4] that μ, ν, and f [the inner dimensions of the three outer products in (2)] satisfy $\mu, \nu \le Z_O$ and $f \le Z_O - \max\{\mu, \nu\}$. Since $Z_O < C'$, it follows that

$$f < C' - \max\{\mu, \nu\}.$$

In possession of r, $\hat{L}$, and $\hat{V}$, Bob is now ready to decode the data-hiding layer that has been applied to x.

We have

$$\begin{aligned} Hr &= Hx + H\hat{L}V^1 + HL^2\hat{V} + HL^3V^3 \\ &= S + \hat{\Lambda}V^1 + \Lambda^2\hat{V} + \Lambda^3V^3 \end{aligned} \tag{3}$$

where $\hat{\Lambda} = H\hat{L}$, $\Lambda^2 = HL^2$ and $\Lambda^3 = HL^3$. Note that $\hat{\Lambda} \in \mathbb{F}_Q^{C'\times\mu}$ and $\hat{V} \in \mathbb{F}_Q^{\nu\times C'}$ are known.

Now, let $J \in \mathbb{F}_Q^{(C'-\mu)\times C'}$ and $K \in \mathbb{F}_Q^{C'\times(C'-\nu)}$ be full-rank matrices such that $J\hat{\Lambda} = 0$ and $\hat{V}K = 0$. Then, Bob can further simplify (3) by computing

$$JHrK = JSK + J\Lambda^3V^3K.$$

Note that $\mathtt{rank}(J\Lambda^3V^3K) \le f < C' - \max\{\mu, \nu\}$, while $C' - \max\{\mu, \nu\}$ is the maximum possible rank of $JHrK$.

Thus, Bob performs the following test. If $JHrK$ is full-rank, then Bob concludes that bit $s_i = 1$ was sent; otherwise, Bob concludes that bit $s_i = 0$ was sent.

With respect to complexity, computing $\bar{Y}$ takes $O(C^2n) = O(C^4)$ operations in $\mathbb{F}_q$. Computing $J, K, JHrK$, and the rank of $JHrK$ each take $O(C^3)$ operations in $\mathbb{F}_Q$, which amounts to $O(C^5)$ in $\mathbb{F}_q$. Thus, the overall decoding complexity is $O(C^5)$ operations in $\mathbb{F}_q$.

C. Probability of Error Analysis

When bit 0 is sent, Bob never makes an error; he makes an error if and only if bit 1 is sent and $JHrK$ is not full-rank. Recall that when bit 1 is sent, S is uniformly distributed over $\mathbb{F}_Q^{C'\times C'}$. Due to the data-hiding encoding, Calvin has no information about S, and therefore S is statistically independent from Λ^3V^3. It follows that $S' = S + \Lambda^3V^3$ is also uniformly distributed over $\mathbb{F}_Q^{C'\times C'}$. Thus, the probability of error when bit 1 is sent is equal to the probability that $JS'K \in \mathbb{F}_Q^{(C'-\mu)\times(C'-\nu)}$ is not full-rank for a uniform S'.

Lemma 4: If $S' \in \mathbb{F}_Q^{C'\times C'}$ is uniformly distributed then, for any $J \in \mathbb{F}_Q^{(C'-\mu)\times C'}$ and any $K \in \mathbb{F}_Q^{C'\times(C'-\nu)}$, the matrix $JS'K$ is full-rank with probability at least $1 - C'/Q$.

Proof: Without loss of generality, assume $\mu \ge \nu$. It suffices to prove the statement for $\mu = \nu$; If $\mu > \nu$, then removing $\mu - \nu$ columns from k cannot possibly increase the rank of $js'k$.

For any fixed J and K, consider the entries of S' as variables taking values in $\mathbb{F}_Q$. Then, each entry of $JS'K$ is a multivariate polynomial over $\mathbb{F}_Q$ with degree at most 1. It follows that $\det(JS'K)$ is a multivariate polynomial over $\mathbb{F}_Q$ with degree at most $C' - \mu \le C'$. Note that if $Q \le C'$, the statement follows trivially, so assume $Q > C'$. From [5, Lemma 4], we have that $P[\det(JS'K) = 0] \le C'/Q$. ∎

Thus, the probability of error of the scheme is upper-bounded by $C'/Q \le C/q^C$, which can be made arbitrarily small by choosing q sufficiently large. This proves Theorem 1.

VI. Achievability for Theorems 2 and 3

We start by addressing Theorem 2. The achievability of Theorem 3 will be shown in Section VI-D. The notation used in this section is summarized in Table I.

We describe a coding scheme that achieves rate $R = C - Z_O$ asymptotically in the packet length n. As before, assume that n is divisible by C, and let $n' = n/C - (1 + kC')$, where $k = (bC + 1)(b + 1)\log_2 q$ and $b = C - Z_O$.

Let $H \in \mathbb{F}_Q^{C'\times C}$ be the parity-check matrix of a $[C, Z_I]$ linear MRD code over $\mathbb{F}_Q$. Let $T \in \mathbb{F}_Q^{C\times C}$ be an invertible matrix such that the first $C - Z_I$ rows of T^{-1} are equal to H, as discussed in Section IV-B.

A. Alice's Encoder

First, given a message $W \in \mathbb{F}_Q^{R\times n'}$, Alice sets $M = W$ (this will be generalized in Section VI-D). Then, she generates a string $\mathbb{S} \in \{0,1\}^k$ of k bits according to the scheme described in Section VI-C. Next, for each ith bit s_i of $\mathbb{S}$, Alice produces a matrix $s_iS_i \in \mathbb{F}_Q^{C'\times C'}$ according to the scheme described in Section V. Then, for each $i = 1, \ldots, k$, she computes $M_i =$

TABLE I
SUMMARY OF COMMONLY USED NOTATION/PARAMETERS

Notation	Meaning
	Network notation/parameters
C	Min-cut of the network
Z_I	Eavesdropping rate
Z_O	Jamming rate
C'	$C - Z_I$
b	$C - Z_O$
A	Network transform from Alice to Bob
B	Network transform from Alice to Calvin
B'	Network transform from Calvin to Bob
	Alice's encoder
W	Alice's "payload" message
q	Network code field-size
n	Packet length (over $\mathbb{F}_q$)
$Q = q^C$	Extension field size
n'	Message packet-length (over $\mathbb{F}_Q$)
$\mathbb{S} = \{s_1, \ldots, s_k\}$	Alice's "small" secret message
T_0, T	Alice's "Mixing" matrices
α	$bC + 1$
k	$\alpha(b+1)\log_2 q$
ϕ	Mapping from $\mathbb{F}_Q$ to $(\mathbb{F}_q)^C$
n''	Packet length (over $\mathbb{F}_q$) for single-bit scheme

$T \begin{bmatrix} s_i S_i \\ N_i \end{bmatrix}$, where each $N_i \in \mathbb{F}_Q^{Z_I \times C'}$ is chosen uniformly at random and independently from any other variables. Finally, she produces a transmission matrix

$$X = \begin{bmatrix} I_C & \phi(M_1) & \phi(M_2) & \cdots & \phi(M_k) & \begin{bmatrix} \phi(M) \\ 0 \end{bmatrix} \end{bmatrix}.$$

B. Bob's Decoder

For each $i = 1, \ldots, k$, Bob extracts a submatrix Y_i from Y corresponding to the submatrix $[\, I_C \quad \phi(x_i)\,]$ from X (i.e., columns $1, \ldots, C, C + (i-1)C' + 1, \ldots, C + iC'$). He then applies on Y_i the decoder described in Section V to obtain each $s_i \in \mathbb{S}$.

Similarly, Bob extracts a submatrix Y_0 consisting of the first b and the last $n'C$ rows of Y. Note that $Y_0 = AX_0 + Z_0$, where $X_0 = \begin{bmatrix} I_b & \phi(M) \\ 0 & 0 \end{bmatrix} \in \mathbb{F}_q^{C \times (b+n'C)}$ and Z_0 has rank at most Z_O. Then, Bob applies the decoder described in Section IV-C to obtain $\phi(M)$.

Finally, Bob computes $W = M = \phi^{-1}(\phi(M))$.

C. Overall Analysis

1) Error Probability Analysis: By the union bound, the probability that Bob makes an error when decoding the k-bit secret $\mathbb{S}$ is at most $kC/q^C \le C^4(\log_2 q)/q^C = O(\frac{\log_2 q}{q^C})$. Given that the secret is decoded correctly, the probability that Bob makes an error when decoding the message is at most $O(n^{C^2}/q)$. Thus, the overall probability of error is at most $O(n^{C^2}/q)$.

2) Rate Analysis: The rate of the scheme is given by $Rn'C/n = R(1 - (1 + kC')C/n) \le R - RC^5(\log_2 q)/n$. Thus, the rate loss is $O(\frac{\log_2 q}{n})$.

3) Complexity Analysis: Decoding all the secret bits takes $O(kC^5) = O(C^8 \log_2 q)$ operations in $\mathbb{F}_q$, while the computational complexity of decoding the message is dominated by the secrecy/data-hiding decoding steps with $O(C^4 n)$ operations over $\mathbb{F}_q$.

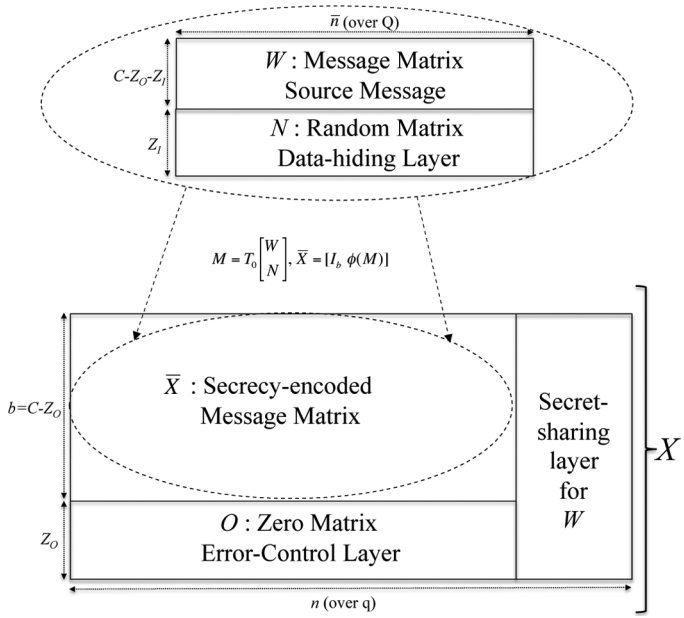

Fig. 4. Overall encoding: pictorial representation of the encoding (described in detail in Section VI). Alice generates her transmission X as follows. She first mixes the $C - Z_I - Z_O$ packets of her message (written as matrix W) with Z_I random packets (rows of the random matrix N) via the invertible matrix T_0 to obtain her secrecy-encoded/masked message $\bar{X}$ (for technical reasons, she also switches from the field $\mathbb{F}_Q$ to the field $\mathbb{F}_q$ via the isomorphism $\phi(.)$). This matrix $\bar{X}$ has the property that if Calvin observes any Z_I rows of it, or indeed any Z_I linear combinations of its rows, no information about W is leaked to him. To further protect her transmissions from the Z_O jamming packets Calvin may inject into the network, Alice adds redundancy by appending Z_O zero rows. Finally, she appends a "small secret-sharing layer" header, as described in Fig. 3.

Note: Both the rate loss and the error probability can be made asymptotically small by choosing q to grow faster than polynomially but slower than exponentially in n. For instance, we may choose $q = 2^{\lfloor \sqrt{n} \rfloor}$.

D. Achievability of Theorem 2

We now describe how the coding scheme described above can be adapted to achieve rate $R = C - Z_I - Z_O$ asymptotically in the packet length n without leaking any information to Calvin. The overall scheme is illustrated in Fig. 4. Essentially, Alice will hide her message from Calvin by applying the secrecy encoding scheme described in Section IV-B, and the rate will be reduced accordingly.

More precisely, let $H_0 \in \mathbb{F}_Q^{R \times b}$ be the parity-check matrix of a $[b, Z_I]$ linear MRD code over $\mathbb{F}_Q$, and let $T_0 \in \mathbb{F}_Q^{b \times b}$ be an invertible matrix such that the first R rows of T_0^{-1} are equal to H_0. Now, given a message $W \in \mathbb{F}_Q^{R \times n'}$, Alice computes $M = T_0 \begin{bmatrix} W \\ N \end{bmatrix}$, where $N \in \mathbb{F}_Q^{Z_I \times n'}$ is chosen independently and uniformly at random. The rest of the encoding is the same, as well as the decoding, except for one last step from Bob. Namely, after recovering M as described above, Bob computes $W = H_0 M$ to recover Alice's message. The secrecy of the message is guaranteed by this procedure as discussed in Section IV-B, while the remainder of the analysis of the scheme is the same as described in Section VI-C. Overall, Alice has to pay a price of Z_I packets of rate loss in order to guarantee secrecy from Calvin.

Note that, from an opposite angle, a scheme achieving Theorem 2 can be immediately derived from the above scheme if ones uses *both* W and N to encode the source's message at rate $C - Z_O$.

VII. Converse for Theorems 2 and 3

As before, we first address the converse for Theorem 3. The converse for Theorem 2 will follow. We start by presenting an attack that Calvin may use to force the achievable rate to at most $C - Z_O - Z_I$, thereby demonstrating that this is indeed an upper bound on the achievable rate. Let $\{e_1, e_2, \ldots, e_C\}$ be a set of edges that form a cut from Alice to Bob. Calvin jams the edges in $\{e_1, e_2, \ldots, e_{Z_O}\}$ by adding random errors on them. Furthermore, Calvin eavesdrops on edges in $\{e_{Z_O+1}, e_{Z_O+2}, \ldots, e_{Z_O+Z_I}\}$. Let $\mathbf{W}$ be the random variable denoting Alice's information. Let $\mathbf{Y}_j$, $\mathbf{Y}_e$, and $\mathbf{Y}_u$ be the random variables denoting the packets carried by the jammed edges $\{e_1, e_2, \ldots, e_{Z_O}\}$, eavesdropped edges $\{e_{Z_O+1}, e_{Z_O+2}, \ldots, e_{Z_O+Z_I}\}$, and untouched edges $\{e_{Z_O+Z_I+1}, e_{Z_O+Z_I+2}, \ldots, e_C\}$, respectively. Let $\mathbf{Y}$ be the random variable denoting the packets received by Bob. Then

$$
\begin{aligned}
nR = H(\mathbf{W}) &= H(\mathbf{W}|\mathbf{Y}) + I(\mathbf{W};\mathbf{Y}) && (4)\\
&\le 1 + \epsilon nR + I(\mathbf{W};\mathbf{Y}) && (5)\\
&\le 1 + \epsilon nR + I(\mathbf{W};\mathbf{Y}_j, \mathbf{Y}_e, \mathbf{Y}_u) && (6)\\
&\le 1 + \epsilon nR + I(\mathbf{W};\mathbf{Y}_e, \mathbf{Y}_u) && (7)\\
&= 1 + \epsilon nR + I(\mathbf{W};\mathbf{Y}_e) + I(\mathbf{W};\mathbf{Y}_u|\mathbf{Y}_e) && (8)\\
&= 1 + \epsilon nR + I(\mathbf{W};\mathbf{Y}_u|\mathbf{Y}_e) && (9)\\
&\le 1 + \epsilon nR + H(\mathbf{Y}_u) && (10)\\
&\le n\left[(C - Z_I - Z_O) + \epsilon R + \frac{1}{n}\right]. && (11)
\end{aligned}
$$

Here, ϵ refers to the probability of error. Equation (4) follows from the fact that Alice's message is uniformly distributed over $\mathbf{W}$, (5) from Fano's inequality, (6) from the data processing inequality, (7) since in the worst case Calvin adds random noise on the edges he jams and so $\mathbf{Y}_j$ is independent of $(\mathbf{W}, \mathbf{Y}_e, \mathbf{Y}_u)$, (8) by the chain rule for mutual information, (9) from the fact that information-theoretic secrecy is required and so $I(\mathbf{W};\mathbf{Y}_e) = 0$, (10) by the fact that conditioning reduces entropy and the definition of mutual information, and finally (11) by the fact that there are at most $C - Z_I - Z_O$ links corresponding to the random variable $\mathbf{Y}_u$ and the alphabet-size upper bound on entropy. Requiring $\epsilon \to 0$ as $n \to \infty$ gives the required result.

A. Converse for Theorem 2

The converse for Theorem 2 follows directly from observing that if Calvin may jam Z_O links in the min-cut of the network, no more than $C - Z_O$ rate can be robustly achievable.

VIII. Errata for [21]

We briefly reprise the scheme of [21] before demonstrating the flaw in the proof. In what follows, all operations are over $\mathbb{F}_q$.

In the scheme of [21], there exist two hash matrices D_0 and D_1 that are chosen independently and uniformly at random $C^2(C - Z_O) \times C^2$ *Vandermonde matrices* , i.e., each column of D_0 and D_1 is of the form $\mathbf{h}(u) = [u, u^2, \ldots, u^{C^2(C-Z_O)}]^T$, where the generator u is chosen independently and uniformly at random from $\mathbb{F}_q$. Both D_0 and D_1 are publicly known to all parties, including Bob and Calvin.

Alice's Encoder: Alice first chooses a random length-$(C^2(C - Z_O) - C^2)$ row vector $\mathbf{u}$. Let $I \in \{0, 1\}$ be the secret bit that Alice wishes to send to Bob. Alice then constructs the length-1 $\times$ C^2 row vector $\mathbf{r}$ such that $[\mathbf{u}, \mathbf{r}]D_I = 0$. Note that such $\mathbf{r}$ exists since the last C^2 rows of D_I form an invertible matrix. Finally, the vector $[\mathbf{u}, \mathbf{r}]$ is rearranged into a $(C - Z_O) \times C^2$ matrix that is sent through the network via random linear network coding.

Bob's Decoder: After receiving the $C \times C^2$ matrix Y, for each $I \in \{0, 1\}$, Bob checks whether there exists $C - Z_O$ length-C vectors $\{\mathbf{x}_i, i \in [1, C - Z_O]\}$ such that $[\mathbf{x}_1 Y, \mathbf{x}_2 Y, \ldots, \mathbf{x}_{C-Z_O} Y]D_I = 0$. If so, Bob decodes the secret bit as I. The idea is that if I is Alice's bit, such $\{\mathbf{x}_i, i \in [1, C - Z_O]\}$ exists for D_I with high probability [7].

Calvin's Successful Attack: When Calvin corrupts $Z_O \ge C - Z_O$ edges, Calvin could mimic Alice's behavior when she wishes to transmit a particular bit, say 1. As a result, Bob would always find length-C row vectors $\{\mathbf{x}_i, i \in [1, C - Z_O]\}$ such that $[\mathbf{x}_1 Y, \mathbf{x}_2 Y, \ldots, \mathbf{x}_{C-Z_O} Y]D_1 = 0$. In this case, Bob cannot determine whether the bit 1 is from Alice or from Calvin.

Even if Calvin can only inject $Z_O < C - Z_O$ errors, if $Z_O + Z_I \ge C - Z_O$, there is another successful attack for Calvin. To see that, without loss of generality, let $Z_O + Z_I = C - Z_O$. Since Calvin can eavesdrop on Z_I packets $\{\mathbf{y}_i, i \in [1, Z_I]\}$, he can carefully choose his Z_O injected error packets $\{\mathbf{z}_i, i \in [1, Z_O]\}$ so that $[\mathbf{y}_1, \ldots, \mathbf{y}_{Z_I}, \mathbf{z}_1, \ldots, \mathbf{z}_{Z_O}]D_1 = 0$. In this case, Bob also always decodes its bit as 1. Thus, the scheme in [21] only works for the case where $C > 2Z_O + Z_I$, which does not improve the result in [7].

Why Our Scheme Works: In our scheme in Section V, instead of distinguishing the bit by the hash matrices, Alice hides her secret in the rank of the bit matrix she transmits. In particular, there is a rank gap $C - Z_I$ between the bit matrix for bit 0 and the one for bit 1. Thus, as long as $C - Z_I > Z_O$, Calvin cannot mimic Alice any more since he can only inject Z_O errors. As a result, Bob can determine Alice's bit by examining the rank of the matrix he decodes.

IX. Conclusion

In this paper, we considered the problem of communicating information secretly and reliably over a network containing a malicious eavesdropping and jamming adversary. Under the assumptions that vanishingly small probabilities of error and block coding are allowed, we substantially improve on the best achievable rates in prior work [26] and also prove the optimality of our achievable rates. A key component of our code design is a scheme that allows a small amount of information to be transmitted secretly and reliably over the network, as long as the total number of packets that the adversary can either eavesdrop on or jam is less than the communication capacity of the network. In proving this scheme, we correct an error in the proof of prior work [21] by a subset of the authors of this work.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[2] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inf. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.

[3] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. Netw.*, vol. 11, no. 5, pp. 782–795, Oct. 2003.

[4] S. Jaggi, P. Sanders, P. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1973–1982, Jun. 2005.

[5] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.

[6] M. Médard and A. Sprintson, *Network Coding: Fundamentals and Applications*. Amsterdam, The Netherlands: Elsevier, 2011.

[7] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Médard, and M. Effros, "Resilient network coding in the presence of Byzantine adversaries," *IEEE Trans. Inf. Theory*, vol. 54, no. 6, pp. 2596–2603, Jun. 2008.

[8] H. Yao, D. Silva, S. Jaggi, and M. Langberg, "Network codes resilient to jamming and eavesdropping," in *Proc. IEEE Int. Symp. Netw. Coding*, Toronto, ON, Canada, Jun. 2010, pp. 1–6.

[9] N. Cai and R. W. Yeung, "Secure network coding," in *Proc. IEEE Int. Symp. Inf. Theory*, Lausanne, Switzerland, Jun. 5, 2002, pp. 323–323.

[10] J. Feldman, T. Malkin, C. Stein, and R. A. Servedio, "On the capacity of secure network coding," in *Proc. 42nd Annu. Allerton Conf. Commun., Control, Comput.*, Sep. 2004.

[11] S. Y. E. Rouayheb and E. Soljanin, "On wiretap networks II," in *Proc. IEEE Int. Symp. Inf. Theory*, Nice, France, Jun. 24–29, 2007, pp. 551–555.

[12] D. Silva and F. R. Kschischang, "Security for wiretap networks via rank-metric codes," in *Proc. IEEE Int. Symp. Inf. Theory*, Toronto, ON, Canada, Jul. 6–11, 2008, pp. 176–180.

[13] D. Silva and F. R. Kschischang, "Universal secure network coding via rank-metric codes," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 1124–1135, Feb. 2011.

[14] P. F. Oliveira and J. Barros, "A network coding approach to secret key distribution," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 414–423, Sep. 2008.

[15] T. Ho, B. Leong, R. Koetter, M. Medard, M. Effros, and D. R. Karger, "Byzantine modification detection in multicast networks using randomized network coding," *IEEE Trans. Inf. Theory*, vol. 54, no. 6, pp. 2798–2803, Jun. 2008.

[16] R. W. Yeung and N. Cai, "Network error correction, part I: basic concepts and upper bounds," *Commun. Inf. Syst*, vol. 6, no. 1, pp. 19–36, 2006.

[17] N. Cai and R. W. Yeung, "Network error correction, part II: Lower bounds," *Commun. Inf. Syst*, vol. 6, no. 1, pp. 37–54, 2006.

[18] R. Matsumoto, "Construction algorithm for network error-correcting codes attaining the singleton bound," ArXiv:cs.IT/0610121, Oct. 2006.

[19] D. Silva, F. R. Kschischang, and R. Kötter, "A rank-metric approach to error control in random network coding," *IEEE Trans. Inf. Theory*, vol. 54, no. 9, pp. 3951–3967, Sep. 2008.

[20] R. Kötter and F. R. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3579–3591, Aug. 2008.

[21] S. Jaggi and M. Langberg, "Resilient network codes in the presence of eavesdropping Byzantine adversaries," in *Proc. IEEE Int. Symp. Inf. Theory*, Jun. 24–29, 2007, pp. 541–545.

[22] D. Charles, K. Jain, and K. Lauter, "Signatures for network coding," *Int. J. Inf. Coding Theory*, vol. 1, no. 1, pp. 3–14, 2009.

[23] F. Zhao, T. Kalker, M. Medard, and J. K. Han, "Signatures for content distribution with network coding," in *Proc. ISIT*, 2007, pp. 556–560.

[24] D. Silva, "Error control for network coding," Ph.D. dissertation, University of Toronto, Toronto, ON, Canada, 2009.

[25] C.-K. Ngai and S. Yang, "Deterministic secure error-correcting (SEC) network codes," in *Proc. IEEE Inf. Theory Workshop*, Tahoe City, CA, USA, Sep. 2–6, 2007, pp. 96–101.

[26] C.-K. Ngai and R. W. Yeung, "Secure error-correcting (SEC) network codes," in *Proc. Workshop Netw. Coding Theory Appl.*, Lausanne, Switzerland, Jun. 15–16, 2009, pp. 98–103.

[27] D. Silva and F. R. Kschischang, "Universal secure error-correcting schemes for network coding," in *Proc. IEEE Int. Symp. Inf. Theory*, 2010, pp. 2428–2432.

[28] E. M. Gabidulin, "Theory of codes with maximum rank distance," *Probl. Inf. Transm.*, vol. 21, no. 1, pp. 1–12, 1985.

Hongyi Yao graduated from Tsinghua University, Beijing, China.

He is a member of Technical Staff with Oracle, Inc., Redwood City, CA, USA. He was a Postdoctoral Scholar with the California Institute of Technology, Pasadena, CA, USA. His thesis advisor at Tsinghua University was Prof. Andrew Yao. His research interests include secure network transmission, reliable network control, and secure wireless communications.

Danilo Silva received the B.Sc. degree from the Federal University of Pernambuco (UFPE), Recife, Brazil, in 2002, the M.Sc. degree from the Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rio de Janeiro, Brazil, in 2005, and the Ph.D. degree from the University of Toronto, Toronto, ON, Canada, in 2009, all in electrical engineering.

From 2009 to 2010, he was a Postdoctoral Fellow with the University of Toronto; the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland; and the State University of Campinas (UNICAMP), Campinas, Brazil. In 2010, he joined the Department of Electrical Engineering, Federal University of Santa Catarina (UFSC), Florianópolis, Brazil, where he is currently an Assistant Professor. His research interests include channel coding, information theory, and network coding.

Dr. Silva was a recipient of a CAPES Ph.D. Scholarship in 2005, the Shahid U. H. Qureshi Memorial Scholarship in 2009, and a FAPESP Postdoctoral Scholarship in 2010.

Sidharth Jaggi received the B.Tech. degree from the Indian Institute of Technology, Bombay, Mumbai, India, in 2000, and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, CA, USA, in 2005, all in electrical engineering.

He was a Postdoctoral Associate with the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, until 2006. Since 2007, he has been an Assistant Professor with the Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong. His research interests include network coding and network error-correcting algorithms, coding theory, steganography, group testing, and compressive sensing.

Michael Langberg received the B.Sc. degree in mathematics and computer science from Tel-Aviv University, Tel-Aviv, Israel, in 1996, and the M.Sc. and Ph.D. degrees in computer science from the Weizmann Institute of Science, Rehovot, Israel, in 1998 and 2003, respectively.

Between 2003 and 2006, he was a Postdoctoral Scholar with the Electrical Engineering and Computer Science departments, California Institute of Technology, Pasadena, CA, USA. He is currently an Associate Professor with the Electrical Engineering Department, State University of New York at Buffalo, Buffalo, NY, USA, and with the Mathematics and Computer Science departments, The Open University of Israel, Ra'anana, Israel. His research addresses the algorithmic and combinatorial aspects of information in communication, management, and storage—focusing on the study of information theory, coding theory, network communication and network coding, big data in the form of succinct data representation, and probabilistic methods in combinatorics

Dr. Langberg is an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY.