On Secure Network Coding With Nonuniform or Restricted Wiretap Sets

Tao Cui, *Student Member, IEEE*, Tracey Ho, *Senior Member, IEEE*, and Jörg Kliewer, *Senior Member, IEEE*

***Abstract*—The secrecy capacity of a network, for a given collection of permissible wiretap sets, is the maximum rate of communication such that observing links in any permissible wiretap set reveal no information about the message. This paper considers secure network coding with nonuniform or restricted wiretap sets, for example, networks with unequal link capacities where a wiretapper can wiretap any subset of k links, or networks where only a subset of links can be wiretapped. Existing results show that for the case of uniform wiretap sets (networks with equal capacity links/packets where any k can be wiretapped), the secrecy capacity is given by a cut-set bound if random keys are injected at the source (and decoded at the sink), whether or not the communicating users have information about the choice of wiretap set. In contrast, we show that for the nonuniform case, this secrecy rate is achievable for the case of known but not unknown wiretap set. We give achievable linear optimization-based strategies where random keys are canceled at intermediate nonsink nodes or injected at intermediate nonsource nodes. Finally, we show that determining the secrecy capacity is an NP-hard problem.**

***Index Terms*—Information-theoretic security, network coding, network interdiction, NP-hard, secrecy capacity.**

I. Introduction

INFORMATION-THEORETICALLY secure communication uses coding to ensure that an adversary that wiretaps a subset of network links obtains no information about the secure message. The secrecy capacity of a network, for a given collection of permissible wiretap sets, is defined as the maximum rate of communication such that any one of the permissible wiretap sets reveals no information about the message. In general, the choice of wiretap set is unknown to the communicating users, though we also discuss the case of known wiretap set where the encoding and decoding functions are allowed to depend on the choice of wiretap set, in which case the secrecy capacity is the maximum rate achievable under the worst case wiretap set.

A theoretical basis for information-theoretic security was given in the seminal paper by Wyner [1] using Shannon's notion of perfect secrecy [2], where a coset coding scheme based on a linear maximum distance separable code was used to achieve security for a wiretap channel. More recently, information-theoretic security has been studied in networks with general topologies. The secure network coding problem, where a wiretapper observes an unknown set of links, was introduced by Cai and Yeung [3], [4]. They proposed a coding strategy, which we refer to as the global key strategy, in which the source injects random key symbols that are decoded at the sink along with the message. They showed achievability of this strategy in the nonuniform case where a wiretapper can observe one of an arbitrary given collection of wiretap link sets, and optimality of this strategy for multicast in the uniform case, where each link has equal capacity and a wiretapper can observe up to k links, and randomness is generated only at the source. For this case, various constructions of secure linear network codes have been proposed in, e.g., [5]–[7]. In [8], Cai and Yeung showed by an example that allowing nonsource nodes to generate randomness also can be advantageous. Other related work on secure network communication includes weakly secure codes [9] and wireless erasure networks [10].

In this paper, we consider secure communication over wireline networks in the nonuniform case. In the case of throughput optimization without security requirements, the assumption that all links have unit capacity is made without loss of generality, since links of larger capacity can be modeled as multiple unit capacity links in parallel. However, in the secure communication problem, such an assumption cannot be made without loss of generality. For the case of uniform wiretap sets, the global key strategy achieves a multicast secrecy rate given by the minimum, over source–sink cuts and wiretap sets, of the capacity of the cut minus the wiretapped links, whether or not the wiretap set is known [3], [4]. In contrast, in the nonuniform case, this secrecy rate is achievable for known but not unknown wiretap sets, even for a single source and sink. We give new linear optimization-based achievable strategies where random keys are canceled at intermediate nonsink nodes or injected at intermediate nonsource nodes, and show ways in which they can outperform the global key strategy. Finally,

Manuscript received October 29, 2009; revised September 06, 2011; accepted April 28, 2012. Date of publication August 31, 2012; date of current version December 19, 2012. This work was supported in part by subcontract #069144 issued by BAE Systems National Security Solutions, Inc.; in part by the Defense Advanced Research Projects Agency and the Space and Naval Warfare System Center, San Diego, CA, under Contracts N66001-08-C-2013 and W911NF-07-1-0029; in part by the National Science Foundation under Grants CNS 0905615, CCF 0830666, and CCF 1017632; and in part by Caltech's Lee Center for Advanced Networking. The material in this paper was presented in part at the 2010 Information Theory and Applications Workshop, La Jolla, CA, and in part at the 2010 IEEE Information Theory Workshop, Dublin, Ireland.

T. Cui and T. Ho are with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: taocui@caltech.edu; tho@caltech.edu).

J. Kliewer is with the Klipsch School of Electrical and Computer Engineering, New Mexico State University, Las Cruces, NM 88003 USA (e-mail: jkliewer@nmsu.edu).

Communicated by S. Ulukus, Associate Editor for Communication Networks.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2012.2216936

we show that determining the secrecy capacity is an NP-hard problem.[1]

II. Network Model and Problem Formulation

In this paper, we focus on acyclic graphs for simplicity; we expect that our results can be generalized to cyclic networks using the approach in [13] and [14] of working over fields of rational functions in an indeterminate delay variable.

We model a wireline network by a directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the vertex set and $\mathcal{E}$ is the directed link set. Each link $(i, j) \in \mathcal{E}$ is a noise-free bit-pipe with a given capacity $c_{i,j}$. We denote the set of incoming links (w, v) of a node v by $\mathcal{I}(v)$ and the set of outgoing links (v, w) of v by $\mathcal{O}(v)$.

A source node $s \in \mathcal{V}$ wishes to communicate a message M securely to a sink node $d \in \mathcal{V}$. An eavesdropper can wiretap a set $\mathcal{A}$ of links chosen from a known collection $\mathcal{W}$ of possible wiretap sets. Without loss of generality, we can restrict our attention to maximal wiretap sets, i.e., no set in $\mathcal{W}$ is a subset of another. The choice of wiretap set $\mathcal{A}$ is unknown to the communicating nodes, except where otherwise specified in this paper. In the case of a known wiretap set, the wiretapper can choose an arbitrary wiretap set $\mathcal{A}$ in $\mathcal{W}$ which is then revealed to the communicating nodes.

Let R_s denote a source–sink secure communication rate (where the precise meaning of "security" is defined below) and $R_{w,v}$ a rate of local randomness generated at node $v \in \mathcal{V}$, which is independent across different nodes v. A block code of blocklength n is defined by a mapping

$$f_{e,s}^{(n)} : \{1, \ldots, 2^{nR_s}\} \times \{1, \ldots, 2^{nR_{w,s}}\} \to \{1, \ldots, 2^{nc_e}\}, \quad e \in \mathcal{O}(s)$$

from the message $M \in \{1, \ldots, 2^{nR_s}\}$ and randomness generated at the source s to the vector transmitted on each outgoing link e of the source, a mapping

$$f_{e,v}^{(n)} : \prod_{d \in \mathcal{I}(v)} \{1, \ldots, 2^{nc_d}\} \times \{1, \ldots, 2^{nR_{w,v}}\} \to \{1, \ldots, 2^{nc_e}\}, \quad e \in \mathcal{O}(v)$$

from the vectors received by a nonsource node v and local randomness generated at v to the vectors transmitted on each outgoing link e of v, and a mapping

$$g_d^{(n)} : \prod_{d \in \mathcal{I}(d)} \{1, \ldots, 2^{nc_d}\} \to \mathcal{X}_s^n$$

from the vectors received by the sink d to the decoded output. Node mappings are applied in topological order; each node receives input vectors from all its incoming links before applying the mappings corresponding to its outgoing links.

[1]While this paper shows NP-hardness of determining secrecy capacity for unequal link capacities or restricted wiretap sets, after submitting the paper, we found that the problem may be hard even when all links have equal capacity and any single link can be wiretapped. Specifically, we show a reduction from the multiple unicast network coding problem, which is not currently known to be in P, NP or undecidable [11]. This result will be presented formally in an upcoming paper. In disseminating this new result, we also learned of related work [12] showing a reduction from the multiple unicast network coding problem to the secure network coding problem with unequal link capacities and restricted wiretap sets.

The secrecy capacity is defined as the highest possible source–sink communication rate R_s for which there exists a sequence of block codes such that the sink decodes message M reliably and, for any choice of $\mathcal{A} \in \mathcal{W}$, the message communicated is kept information theoretically secret from the eavesdropper who observes $\mathbf{Z}^n = [Z_1, \ldots, Z_n]$. As we will see, our results in Sections III–V are not affected by whether we impose a zero or asymptotically negligible decoding error requirement at the sink, or by whether we impose a weak secrecy requirement where for any $\epsilon_n > 0$ there exists n such that $\frac{1}{n} I(M; \mathbf{Z}^n) < \epsilon_n$, a strong secrecy requirement where for any $\epsilon_n > 0$ there exists n such that $I(M; \mathbf{Z}^n) < \epsilon_n$, or a perfect secrecy requirement $I(M; \mathbf{Z}^n) = 0$ for all n.

We provide achievable strategies in Section III for this general problem; in Sections IV and V, we show unachievability and NP-hardness results in that hold even for the following special cases:

1) Scenario 1 is a wireline network with *equal* link capacities, where the wiretapper can wiretap an unknown subset of k links from a known collection of vulnerable network links.
2) Scenario 2 is a wireline network with *unequal* link capacities, where the wiretapper can wiretap an unknown subset of k links from the entire network.

It is convenient to show these results for Scenario 1 first, and then show the corresponding results for Scenario 2, by converting the Scenario 1 networks considered into corresponding Scenario 2 networks for which the same result holds.

Although, for the sake of simplicity, we focus on single-source single-sink networks, Strategy 2 in Section III can be easily extended to multicast networks, whereas the unachievability and NP-hardness results discussed in Sections IV and V directly apply to both multicast and nonmulticast cases since the single-source single-sink case represents a special case for both.

III. Cut-Set Bound and Achievable Strategies

In this section, we consider the general wireline problem with unequal link capacities where the eavesdropper can wiretap an unknown set $\mathcal{A}$ of links chosen from a known collection $\mathcal{W}$ of possible wiretap sets. We consider a simple cut-set bound and give two linear optimization-based achievable strategies and show cases in which they outperform the existing global key strategy. The intuition and results in this section are subsequently used in showing the results in Sections IV and V.

A. Simple Cut-Set Bound

Let $\mathcal{S}^c$ denote the set complement of a set $\mathcal{S}$. A cut for $x, y \in \mathcal{V}$ is a partition of $\mathcal{V}$ into two sets $\mathcal{V}_x$ and $\mathcal{V}_x^c$ such that $x \in \mathcal{V}_x$ and $y \in \mathcal{V}_x^c$. For the $x - y$ cut given by $\mathcal{V}_x$, the cut-set $[\mathcal{V}_x, \mathcal{V}_x^c]$ is the set of links going from $\mathcal{V}_x$ to $\mathcal{V}_x^c$, i.e.,

$$[\mathcal{V}_x, \mathcal{V}_x^c] = \{(u, v) | (u, v) \in \mathcal{E},\ u \in \mathcal{V}_x,\ v \in \mathcal{V}_x^c\}. \quad (1)$$

The set $\{(v, u) | (v, u) \in \mathcal{E},\ u \in \mathcal{V}_x,\ v \in \mathcal{V}_x^c\}$ comprises the reverse links of the cut.

Theorem 1: Consider a network of point-to-point links, where link (i,j) has capacity $c_{i,j}$. The secrecy capacity is upper bounded by

$$\min_{\{\mathcal{V}_s:\,\mathcal{V}_s \text{ is an } s-d \text{ cut}\}} \min_{\mathcal{A}\in\mathcal{W}} \sum_{(i,j)\in[\mathcal{V}_s,\mathcal{V}_s^c]\cap\mathcal{A}^c} c_{i,j} \tag{2}$$

if randomness is generated only at the source, or if there are no reverse links on the minimizing cut. This holds whether or not the communicating nodes have knowledge of the chosen wiretap set $\mathcal{A}$.

Proof: Consider any source–sink cut $\mathcal{V}_s$ and any wiretap set $\mathcal{A}\in\mathcal{W}$. Denote by $\mathbf{X}$ the transmitted signals from nodes in $\mathcal{V}_s$ over links in $[\mathcal{V}_s,\mathcal{V}_s^c]$ and denote by $\mathbf{Y}$ and $\mathbf{Z}$ the observed signals from links in $[\mathcal{V}_s,\mathcal{V}_s^c]$ and in $[\mathcal{V}_s,\mathcal{V}_s^c]\cap\mathcal{A}$, respectively. We consider block coding with blocklength n and secret message rate R_s. By the weak secrecy requirement $H(M|\mathbf{Z}^n)/n \geq H(M)/n-\epsilon_n$, we have

$$\begin{aligned} nR_s &\leq H(M) \leq H(M|\mathbf{Z}^n)+n\epsilon_n \\ &\overset{(a)}{\leq} H(M|\mathbf{Z}^n)-H(M|\mathbf{Y}^n)+2n\epsilon_n \\ &= H(M|\mathbf{Z}^n)-H(M|\mathbf{Y}^n,\mathbf{Z}^n)+2n\epsilon_n \\ &= I(M;\mathbf{Y}^n|\mathbf{Z}^n)+2n\epsilon_n \\ &\overset{(b)}{\leq} I(\mathbf{X}^n;\mathbf{Y}^n|\mathbf{Z}^n)+2n\epsilon_n \\ &\leq n\sum_{(i,j)\in[\mathcal{V}_s,\mathcal{V}_s^c]\cap\mathcal{A}^c} c_{i,j}+2n\epsilon_n \end{aligned} \tag{3}$$

where $\epsilon_n\to 0$ as $n\to+\infty$ and

(a) is due to Fano's inequality;

(b) is due to the data processing inequality and the fact that $M\to\mathbf{X}^n\to\mathbf{Y}^n\to\mathbf{Z}^n$ forms a Markov chain.

■

Note that the proof addresses the case of asymptotically negligible decoding error but clearly also applies to the zero error case (where the capacity can be no greater) since it is a converse result. For the same reason, it also includes the strong secrecy case $H(M|\mathbf{Z}^n)\geq H(M)-\epsilon_n$ and the perfect secrecy case $H(M|\mathbf{Z}^n)=H(M)$. Further, the bound in (2) is always achievable in the uniform link-capacity case (by the global key strategy [3], [4]), or when the wiretap set is known (by not transmitting on the wiretapped links, where achievability follows from the Ford–Fulkerson theorem [15]).

B. Achievable Strategies for Unknown Wiretap Set

In the global key strategy, the source transmits R_s secret information symbols and R_w random key symbols, where R_s+R_w is equal to the min-cut of the network. This scheme does not achieve capacity in general. One reason is that a node with connections to two other nodes can generate common randomness between them as in [8]. We show in the following that in the nonuniform case, there is another advantage of local random keys in improving utilization of cuts that have larger capacity links and larger overall cut capacity compared to other cuts in the network.

In particular, we show that capacity can be improved not only by local key generation but also by key cancellation at intermediate nodes. We consider combinations of local and global random keys, where a local key is injected at a nonsource node and/or canceled at a nonsink node. However, it is complicated to optimize over all possible combinations of nodes at which keys are injected and canceled. Thus, we propose the following more tractable constructions, which we will use to develop further results in subsequent sections.

Strategy 1: Random Keys Injected by Source and Possibly Canceled at Intermediate Nodes: Our first construction achieves secrecy with random keys injected only at the source. The source carries out precoding so that random keys are canceled at intermediate nodes and the sink receives the intended message without interference from the random keys. As such, it applies in the single-source, single-sink case and is useful in networks where the incoming capacity of the sink is too small to accommodate the message plus all the keys needed in the network. An example is given in Fig. 1, where each link has unit capacity, the number of wiretapped links is $k=2$, and only the first layer of the three layer network is allowed to be wiretapped. The secret message rate $R_s=3$ is achievable by using the strategy in Fig. 1, where the operation is on a finite field $GF(5)$. In Fig. 1, a,b,c are secret messages and f,g are keys. The message on the ith link in the first layer is denoted as x_i, $i=1,2,3,4,5$. The key f is canceled at the second layer and the key cancelation scheme is labeled on the last layer links. It is easy to see that $H(x_i,x_j|a,b,c)=2,\forall i\neq j$ which means perfect secrecy is achieved. At the same time, the sink d can decode a,b,c and the key g. When key cancelation is not applied, let R_s and R_w be the secrecy rate and the random key rate at the source, respectively. Let z be the total rate of transmission on the first layer. To achieve secrecy, we must have $R_w\geq\frac{2}{5}z$, where the cut-set condition on the first layer requires $R_s+R_w\leq z$. Since the sink needs to decode both message and random key symbols from the source, the cut-set condition on the last layer requires $R_s+R_w\leq 4$. Combining these we obtain $R_s\leq\max_z\min(4-\frac{2}{5}z,\ \frac{3}{5}z)=\frac{12}{5}$, which is strictly less than 3.

To formally develop the Strategy 1 construction, we will use the following result:

Claim 1 [16, Corollary 19.21]: Given an acyclic network, there exists, for a sufficiently large finite field, a linear network code in which the dimension of the received subspace at each nonsource node t is $\min(\omega,\text{maxflow}(t))$, where ω is the dimension of the message subspace.

Let R_s denote the secret message rate and $z_{i,j}$ the transmission rate on each network link $(i,j)\in\mathcal{E}$, whose values we will discuss how to choose below. Consider the graph $\mathcal{G}$ with the capacity of each link $(i,j)\in\mathcal{E}$ set as $z_{i,j}\leq c_{i,j}$. As illustrated in Fig. 2, augment the graph as follows.

1) Connect each subset of links $\mathcal{A}\in\mathcal{W}$ to a virtual node $t^{\mathcal{A}}$: more precisely, for each directed link $(i,j)\in\mathcal{E}$ in the network, create a node $v_{i,j}$ and replace (i,j) by two links $(i,v_{i,j})$ and $(v_{i,j},j)$ of capacity $z_{i,j}$, and for each $(i,j)\in\mathcal{A}$ create a link $(v_{i,j},t^{\mathcal{A}})$ of capacity $v_{i,j}$. Let $R_{s\to\mathcal{A}}$ be the max flow/min cut capacity between s and $t^{\mathcal{A}}$.

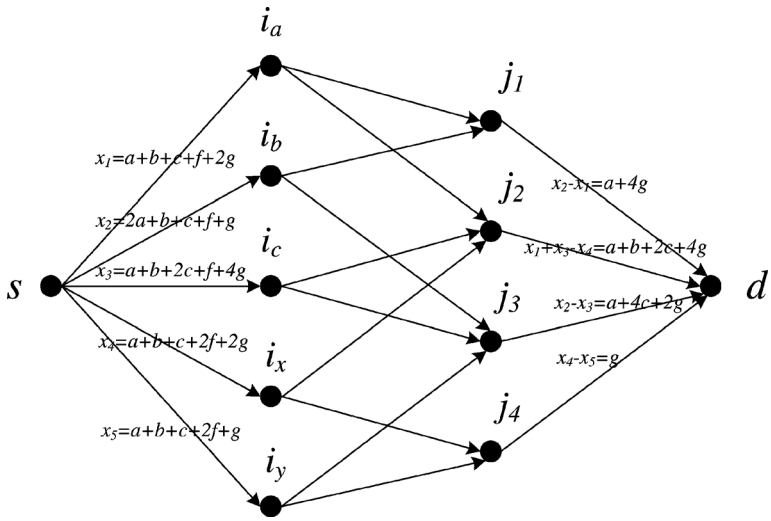

Fig. 1. Example of Strategy 1, where any two of the five links in the first layer can be wiretapped. Secrecy rate 3 is achieved by canceling at the second layer one of the two random keys injected by the source. The operation is on a finite field $GF(5)$.

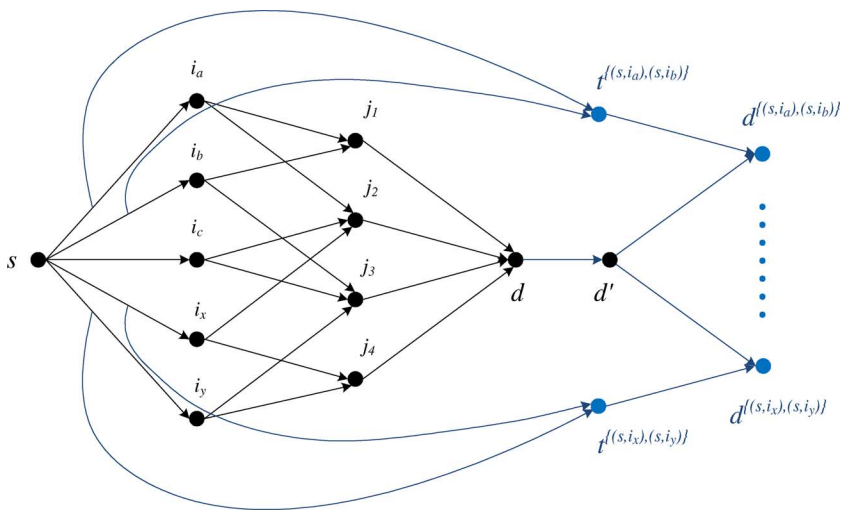

Fig. 2. Illustration of Strategy 1, an achievable construction where random keys are injected by the source and possibly canceled at intermediate nodes. In this figure, $k = 2$ and only the five links in the first layer can be wiretapped.

2) Add a virtual sink node d' and join the actual sink d to d' by a link (d, d') of capacity of R_s.
3) Connect both $t^{\mathcal{A}}$ and the virtual sink d' to a virtual sink $d^{\mathcal{A}}$ by adding a link $(t^{\mathcal{A}}, d^{\mathcal{A}})$ of capacity $R_{s\to\mathcal{A}}$ and a link $(d', d^{\mathcal{A}})$ of capacity R_s, respectively.

The source sends a secret message $\mathbf{v} = [v_1, \ldots, v_{R_s}]^T$ along with R_w random key symbols $\mathbf{w} = [w_1, \ldots, w_{R_w}]^T$.[2] The values of R_s, R_w, and $z_{i,j}$ are chosen such that each virtual sink $d^{\mathcal{A}}$ can decode $R_s + R_{s\to\mathcal{A}}$ linear combinations of message and random key symbols, and the sink d can decode the R_s message symbols. Specifically, if for each $\mathcal{A}$, the rate $R_s + R_{s\to\mathcal{A}}$ equals the min-cut capacity between the source and the virtual sink $d^{\mathcal{A}}$ and $R_{s\to\mathcal{A}} \leq R_w$, by using Claim 1, there exists a network code such that each $d^{\mathcal{A}}$ receives $R_s + R_{s\to\mathcal{A}}$ linearly independent combinations of $\mathbf{v}$ and $\mathbf{w}$ when the finite field size is sufficiently large ($q > \binom{|\mathcal{E}|}{k}$). Let the signals received at a particular virtual sink $d^{\mathcal{B}}$ be denoted as $\mathbf{M}_{\mathcal{B}}[\mathbf{v}^T, \mathbf{w}^T]^T$, where $\mathbf{M}_{\mathcal{B}}$ is an $(R_s + R_{s\to\mathcal{B}}) \times (R_s + R_w)$ received coding matrix with full row rank. We can add $R_w - R_{s\to\mathcal{B}}$ rows to $\mathbf{M}_{\mathcal{B}}$ to get a full rank $(R_s + R_w) \times (R_s + R_w)$ square matrix $\tilde{\mathbf{M}}_{\mathcal{B}}$. We then precode the secret message and keys using $\tilde{\mathbf{M}}_{\mathcal{B}}^{-1}$, i.e., the source transmits $\tilde{\mathbf{M}}_{\mathcal{B}}^{-1}[\mathbf{v}^T, \mathbf{w}^T]^T$, so that link $(d', d^{\mathcal{B}})$ carries $\mathbf{v}$.

Claim 2: Strategy 1 allows the sink to decode the message $\mathbf{v}$ and achieves perfect secrecy.

Proof: Since (d, d') is the only incoming link of $(d', d^{\mathcal{B}})$, and both links have capacity R_s which is equal to the rate of the message $\mathbf{v}$, link (d, d') carries exactly $\mathbf{v}$. This implies that sink d receives $\mathbf{v}$. Furthermore, for any virtual sink $d^{\mathcal{A}}$, the received coding matrix with precoding is $\mathbf{M}_{\mathcal{A}}\tilde{\mathbf{M}}_{\mathcal{B}}^{-1}$, which is a full row rank matrix. As $\mathbf{M}_{\mathcal{A}}\tilde{\mathbf{M}}_{\mathcal{B}}^{-1}$ is a full row rank matrix, the coding vectors of the received signals from the set $\mathcal{A}$ of wiretapping links span a rank $R_{s\to\mathcal{A}}$ subspace that is linearly independent of the set of coding vectors corresponding to message $\mathbf{v}$ received on $(d', d^{\mathcal{A}})$. Therefore, the signals received

[2]We assume that R_s and $R_{s\to\mathcal{A}}$ are integers, which can be approximated arbitrarily closely by scaling the capacity of all links by the same factor.

on $\mathcal{A}$ are independent of the message $\mathbf{v}$, and perfect secrecy is achieved. ■

Note that applying $\tilde{\mathbf{M}}_{\mathcal{B}}^{-1}$ causes the random keys injected by the source to be either canceled at intermediate nodes or decoded by the sink.

It remains to optimize over values of R_s, R_w, and $z_{i,j}$ such that for each $\mathcal{A}$, the rate $R_s + R_{s\to\mathcal{A}}$ equals the min-cut capacity between s and $d^{\mathcal{A}}$ and $R_{s\to\mathcal{A}} \le R_w$. Since computing $R_{s\to\mathcal{A}}$ (the min-cut capacity between s and $t^{\mathcal{A}}$) for arbitrary $z_{i,j}$ involves a separate max flow computation, to simplify the optimization, we can constrain $R_{s\to\mathcal{A}}$ to be equal to some upper bound $U_{\mathcal{A}}$ on $R_{s\to\mathcal{A}}$ and thereby obtain an achievable secrecy rate using Strategy 1. For instance, we can take $U_{\mathcal{A}}$ to be $\sum_{(i,j)\in\mathcal{A}} z_{i,j}$, or alternatively take $U_{\mathcal{A}}$ to be the min-cut capacity between s and $t^{\mathcal{A}}$ on the graph with the original link capacities $c_{i,j}$. We can write a linear program (LP) for this key cancelation strategy as follows:

$$\begin{aligned}
&\max R_s \\
&\text{subject to } \sum_{(i,j)\in\mathcal{E}} f_{i,j}^{\mathcal{A}} - \sum_{(i,j)\in\mathcal{E}} f_{j,i}^{\mathcal{A}} \\
&\qquad = \begin{cases} R_s + U_{\mathcal{A}}, & \text{if } i = s \\ -R_s - U_{\mathcal{A}}, & \text{if } i = d^{\mathcal{A}} \\ 0, & \text{otherwise} \end{cases} \\
&\qquad\qquad\qquad\qquad\qquad\qquad \forall \mathcal{A} \in \mathcal{W} \\
&f_{i,j}^{\mathcal{A}} \le z_{i,j} \le c_{i,j} \quad \forall (i,j) \in \mathcal{E} \qquad (4)
\end{aligned}$$

where $f_{i,j}^{\mathcal{A}}$ represents the rate of flow on link (i,j) that is intended for the virtual sink $d^{\mathcal{A}}$. By defining and imposing flow conservation conditions on the virtual flow $f^{\mathcal{A}} = (f_{i,j}^{\mathcal{A}} : (i,j) \in links)$, we ensure that the min cut between the source and $d^{\mathcal{A}}$ is at least $R_s + U_{\mathcal{A}}$. Since the only incoming links of $d^{\mathcal{A}}$ are $(t^{\mathcal{A}}, d^{\mathcal{A}})$ of capacity $R_{s\to\mathcal{A}}$ and $(d', d^{\mathcal{A}})$ of capacity R_s, this implies that $R_{s\to\mathcal{A}}$ equals the upper bound $U_{\mathcal{A}}$. Thus, the optimal value of (4) gives an achievable secrecy rate.

Strategy 2: Random Keys Injected by Source And/Or Intermediate Nodes and Decoded at Sink: In Strategy 2, any node in the network can inject random keys. The sink is required to decode both the secret message and the random keys from all nodes, i.e., keys are not canceled within the network, while the random key rates must be sufficient to "fill" each wiretap set (in a sense that is made precise below). Although for simplicity of notation the algorithm description below is for the single-source, single-sink case, this strategy applies directly to multiple-source multicast case. If random keys are injected only at the source, the strategy reduces to the global key strategy in [3]. Note that under the assumption that only the source knows the message and different nodes do not have common randomness a priori, here we cannot apply the key cancelation and precoding idea from Strategy 1, since after applying the precoding matrix, each node may potentially be required to transmit a mixture of the source message and other nodes' random keys.

Let $R_{w,v}$ be the random key injection rate at node v. As before, R_s denotes the secret message rate at the source and $z_{i,j}$ the transmission rate on link (i,j). We will address the choice

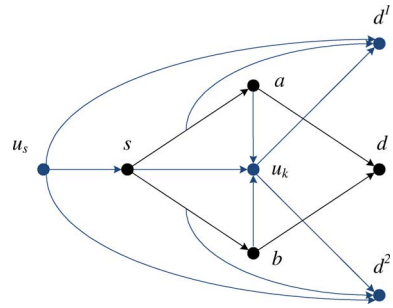

Fig. 3. Example of the augmented network construction for the proof of correctness of Strategy 2, where s, a, b, d are nodes of the original graph, and only one of the two links (s,a) and (s,b) can be wiretapped.

of these rates below. Consider the graph $\mathcal{G}$ with the capacity of each link $(i,j) \in \mathcal{E}$ set as $z_{i,j}$. Connect each subset of links $\mathcal{A} \in \mathcal{W}$ to a virtual node $d^{\mathcal{A}}$: more precisely, for each directed link $(i,j) \in \mathcal{E}$ in the network, create a node $v_{i,j}$ and replace (i,j) by two links $(i, v_{i,j})$ and $(v_{i,j}, j)$ of capacity $z_{i,j}$, and for each $(i,j) \in \mathcal{A}$ create a link $(v_{i,j}, d^{\mathcal{A}})$ of capacity $v_{i,j}$. Intuitively, we want the max flow/min cut capacity from the message and random key sources to $d^{\mathcal{A}}$ to be equal to that in the absence of the message. Similarly to Strategy 1, we can simplify the optimization by constraining this max flow/min cut capacity to be equal to an upper bound, $\sum_{(i,j)\in\mathcal{A}} z_{i,j}$. Specifically, we have the following LP:

$$\begin{aligned}
&\max R_s \\
&\text{subject to} \\
&\qquad \sum_j f_{i,j}^{\mathcal{A}} - \sum_j f_{j,i}^{\mathcal{A}} \begin{cases} = -\sum_{(i',j')\in\mathcal{A}} z_{i',j'}, & \text{if } i = d^{\mathcal{A}} \\ \le R_{w,i}, & \text{otherwise} \end{cases} \\
&\qquad \forall \mathcal{A} \in \mathcal{W}, \\
&\qquad \sum_j f_{i,j}^{d} - \sum_j f_{j,i}^{d} \\
&\qquad = \begin{cases} -\left(R_s + \sum_{v\in\mathcal{V}, v\neq d} R_{w,v}\right), & \text{if } i = d \\ R_s + R_{w,s}, & \text{if } i = s \\ R_{w,i}, & \text{otherwise} \end{cases} \\
&\qquad f_{i,j}^{\mathcal{A}} \le z_{i,j}, \quad f_{i,j}^{d} \le z_{i,j}, \quad z_{i,j} \le c_{i,j} \quad \forall (i,j) \in \mathcal{E} \qquad (5)
\end{aligned}$$

where the first set of equations corresponds to the requirement that the network accommodates a flow $f^{\mathcal{A}} = (f_{i,j}^{\mathcal{A}} : (i,j) \in \mathcal{E})$ of size $\sum_{(i,j)\in\mathcal{A}} z_{i,j}$ from the random key sources to $d^{\mathcal{A}}$, the second set of equations corresponds to the requirement that the network accommodates a flow $f^{d} = (f_{i,j}^{d} : (i,j) \in \mathcal{E})$, of size equal to the sum of the message and random key rates, from the message and random key sources to the sink d, and the third set of inequalities corresponds to the link capacity constraints.

Claim 3: Strategy 2 allows the sink to decode the message $\mathbf{v}$, and achieves perfect secrecy.

Proof: As illustrated in the example of Fig. 3, consider an augmented network with

1) a virtual source node u_s connected to the source node s by a directed link (u_s, s) of capacity R_s, and connected to

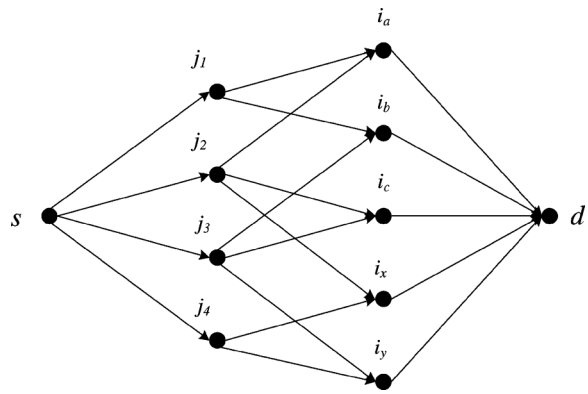

Fig. 4. Example of the usefulness of Strategy 2.

each virtual sink $d^{\mathcal{A}}$ by a directed link $(u_s, d^{\mathcal{A}})$ of capacity R_s, and

2) a virtual node u_k connected to each node v by a directed link (v, u_k) of capacity $R_{w,v}$, and connected to each virtual sink $d^{\mathcal{A}}$ by a directed link $(u_k, d^{\mathcal{A}})$ of capacity $\sum_v R_{w,v} - \sum_{(i,j)\in\mathcal{A}} z_{i,j}$.

The source information enters the network at the virtual source node u_s and is transmitted to each virtual sink $d^{\mathcal{A}}$. Consider a multisource multicast problem on this network, where the actual sink node and the virtual sinks $d^{\mathcal{A}}$ each demand the source message and all the random keys. By the first constraint of the LP, the max flow from the random key sources to $d^{\mathcal{A}}$ in the original network equals $\sum_{(i,j)\in\mathcal{A}} z_{i,j}$; together with the additional capacity in the augmented network ($\sum_v R_{w,v} - \sum_{(i,j)\in\mathcal{A}} z_{i,j}$ from the random key sources and R_s from u_s), the max flow from the message and random key sources to each virtual sink $d^{\mathcal{A}}$ is sufficient to ensure that the multicast problem is feasible [17]. A capacity-achieving code for this multicast problem in the transformed graph corresponds to a code for the original secrecy problem, since the information received by each virtual sink $d^{\mathcal{A}}$ from the set $\mathcal{A}$ of original network links must be independent of information received from the additional links, which includes the entire source message. ■

An example where this strategy is useful is given in Fig. 4, which is obtained by interchanging the source and the sink as well as reversing all the links in Fig. 2. At most three links in the last layer can be wiretapped. By injecting one local key at node j_2 and two global keys at the source, Strategy 2 can achieve secrecy rate 2. On the other hand, if random keys are only injected at the source, the secrecy rate is at most $\frac{8}{5}$. Let R_s and R_w be the secrecy rate and the random key rate at the source, respectively. Let z be the total rate of transmission on the last layer. To achieve secrecy, we must have $R_w \geq \frac{3}{5}z$, where the min-cut condition on the last layer requires $R_s + R_w \leq z$. Since the source injects all the random keys, the min-cut condition on the first layer requires $R_s + R_w \leq 4$. Combining these we obtain $R_s \leq \frac{8}{5}$, which is strictly less than 2.

From the examples, we see that Strategies 1 and 2 are useful in complementary cases. In general, these two strategies can be combined to obtain a higher secrecy rate. We use these strategies in the following sections to develop theoretical results. However, for numerical computation of achievable rates in scenarios 1 and 2, we note that the number of possible wiretapping sets, and thus the size of the LPs, are exponential in the size k of each wiretap set, so they are most useful for small k.

IV. Unachievability in the Nonuniform Case

In the case of unrestricted wiretapping sets, unit link capacities, and all random keys injected at the source, the secrecy capacity is equal to the simple cut-set bound in (2) [3]. The rate in (2) is always achievable in the uniform case, or in the nonuniform case when the wiretap set is known. In contrast, when the wiretap set is unknown, we show that if the set of wiretappable links is restricted (Scenario 1), the rate (2) is not achievable in general even with local random keys, by considering the example in Fig. 5. We give an analytical proof that the rate in (2) is not achievable and further use the program Information-Theoretic Inequalities Prover (Xitip) [18] to show that the secrecy capacity is bounded away from the rate in (2). We then convert the example into one with unequal link capacities (Scenario 2) and show unachievability for this case also.

A. Restricted Wiretap Set (Scenario 1)

Consider the example in Fig. 5, where all links have unit capacity and any three of the five middle layer links can be wiretapped. Let the middle layer links be 1–5 (from top to bottom) and the last layer links be 6–8 (from top to bottom). Let the signal carried by link i be called signal i, or S_i, and let the source information be denoted X. We can easily observe that the rate in (2) equals 2.

To provide intuition for the case when the wiretap set is unknown, we first show that secrecy rate 2 cannot be achieved by using scalar linear coding. Then, the argument is converted to an information-theoretic proof that secrecy rate 2 cannot be achieved by any coding scheme.

Suppose secrecy rate 2 is achievable with a scalar linear network code. First note that the source cannot inject more than unit amount of random key; otherwise, the first layer cannot carry two units of source data. Let the random key injected by the source be denoted K. For the case when the source injects a unit amount of random key, we first have the following observations. Signal 6 must be a function of signal 1; otherwise, if the adversary sees the signals 2–4, then he knows signals 6–7. Also, signal 8 must be a function of signal 5; otherwise, if the adversary sees signals 1, 2, and 4, then he knows signals 7–8. Similarly, we can show that signal 8 must be a function of signal 1, and signal 7 must be a function of signal 2. We consider the following two cases.

Case 1: signal 5 is a linear combination of signals present at the source node. To achieve the full key rank condition on links 1, 2, and 5, node *a* must put two independent local keys k_1 and k_2 on links 1 and 2, respectively. Link 7, whose other input is independent of k_2, is then a function of k_2. Similarly, Link 8 is a function of k_1. This means that the last layer has two independent local keys on it.

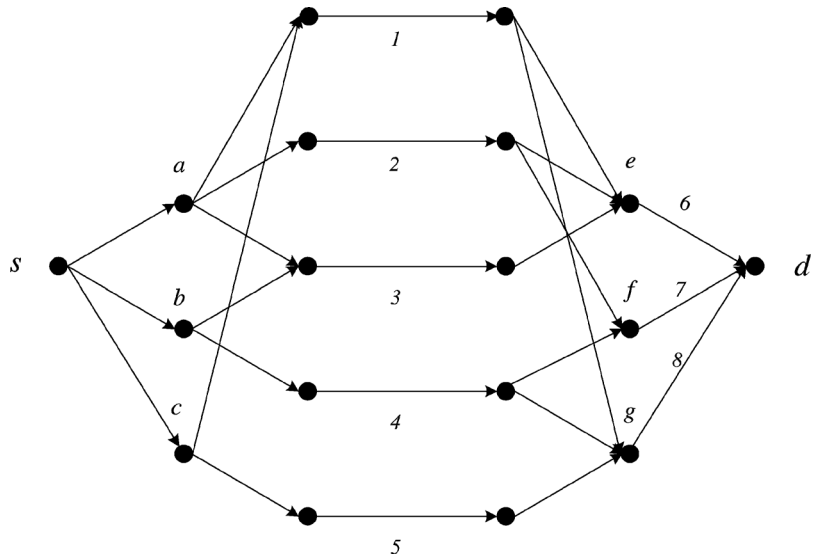

Fig. 5. Example to show that the secrecy rate without knowledge of wiretapping set is smaller than that with such knowledge. The wiretapper can wiretap any three of the five links in the middle layer.

Case 2: signal 5 is a linear combination of signals present at the source node as well as a local key k injected by node c.
Case 2a: k is also present in signal 1. Then, k is present in signal 6 and is independent of the key present in signal 7.
Case 2b: k is not present in signal 1. Then, k is present in signal 8 and is independent of the key present in signal 7.

In all three cases 1, 2a, and 2b, there is a pair of last layer links which are functions of two independent random keys, leaving capacity for only one unit of secret message. Thus, we conclude that the secrecy rate without knowledge of the wiretapping set by using only linear network coding is less than 2.

We can extend the above argument to analytically show this result for any coding scheme in the case of zero decoding error and perfect secrecy.

Theorem 2: For the wireline network in Fig. 5, a secrecy rate of 2 is not achievable if any three out of the five links (1–5) in the middle layer are wiretapped and the location of those links is unknown.

Proof: See the Appendix. ■

We can also show that the secrecy rate is bounded away from 2, by using the framework in [19], where linear optimization is used to check whether an information inequality is implied by set of linear information inequalities together with Shannon-type information inequalities. The derived bound holds for either zero or asymptotically negligible decoding error, and for perfect, strong, or weak secrecy, since it is derived by linear optimization over a set of linear constraints whose constant terms can be scaled by the code blocklength with a corresponding scaling in the objective function[3] and by the continuity of the optimal value of a linear program with respect to its constraints (see, e.g., [20]). Let X be the message sent from the source and $Z_i, i = 1, 2, 3$, be the signals on the links adjacent to the source. We find an upper bound on $H(X)$ that is implied by

$$
\begin{aligned}
&(1)\quad H(Z_i) \le 1,\ H(S_j) \le 1,\ i = 1,2,3,\ j = 1,\ldots,8,\\
&(2)\quad H(X|S_6, S_7, S_8) = 0,\\
&(3)\quad I(X, Z_1, Z_2, Z_3, S_4, S_5, S_7, S_8; S_6|S_1, S_2, S_3) = 0,\\
&(4)\quad I(X, Z_1, Z_2, Z_3, S_1, S_3, S_5, S_6, S_8; S_7|S_2, S_4) = 0,\\
&(5)\quad I(X, Z_1, Z_2, Z_3, S_2, S_3, S_6, S_7; S_8|S_1, S_4, S_5) = 0,\\
&(6)\quad I(X; S_1, S_2, S_3) = 0,\ I(X; S_1, S_2, S_4) = 0,\\
&(7)\quad I(X; S_1, S_2, S_5) = 0,\ I(X; S_1, S_3, S_4) = 0,\\
&(8)\quad I(X; S_1, S_3, S_5) = 0,\ I(X; S_1, S_4, S_5) = 0,\\
&(9)\quad I(X; S_2, S_3, S_4) = 0,\ I(X; S_2, S_3, S_5) = 0,\\
&(10)\quad I(X; S_2, S_4, S_5) = 0,\ I(X; S_3, S_4, S_5) = 0,\\
&(11)\quad I(S_1; Z_2|Z_1, Z_3) = 0,\ I(S_2; Z_2, Z_3|Z_1) = 0,\\
&(12)\quad I(S_3; Z_3|Z_1, Z_2) = 0,\ I(S_4; Z_1, Z_3|Z_2) = 0,\\
&(13)\quad I(S_5; Z_1, Z_2|Z_3) = 0,\ I(S_1; S_4|Z_1, Z_2, Z_3) = 0,\\
&(14)\quad I(S_2; S_4, S_5|Z_1, Z_2, Z_3) = 0,\ I(S_3; S_5|Z_1, Z_2, Z_3) = 0,\\
&(15)\quad I(S_4; S_1, S_2, S_5|Z_1, Z_2, Z_3) = 0,\\
&\qquad\ I(S_5; S_2, S_3, S_4|Z_1, Z_2, Z_3) = 0,\\
&(16)\quad I(S_1, S_2, S_3, S_4, S_5; X|Z_1, Z_2, Z_3) = 0,
\end{aligned}
\tag{6}
$$

where the first inequality is the capacity constraint, the second constraint shows that the sink can decode X, constraints 3–5 mean that the signals in the last layer are independent of other signals given the incoming signals from the middle layer, constraints 6–10 represent the secrecy constraints when any three links in the middle layer are wiretapped, and constraints 11–16 represent the conditional independence between the signals in the first layer and those in the middle layer. In particular, constraint 16 shows that $X \to (Z_1, Z_2, Z_3) \to (S_1, \ldots, S_5)$ forms

[3]The linear optimization maximizes $H(X)$ over the constraints in (6), which correspond to the zero decoding error and perfect secrecy case, and the closure of the set of Shannon-type inequalities, which forms a convex cone.

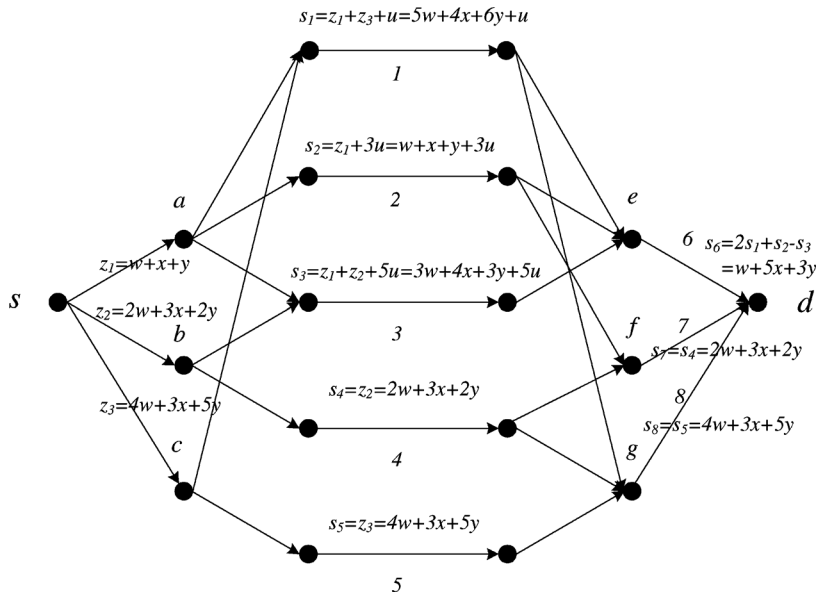

Fig. 6. Coding scheme achieving secrecy rate 1 without knowledge of the wiretap set for the network in Fig. 5, where any three of the five middle layer links can be wiretapped. w is the secret message, x and y are keys injected at the source, and u is a key injected at node a and canceled at node e. The operations are over a finite field $GF(7)$.

a Markov chain. Note that constraints 3–5 and 11–16 implicitly allow some randomness to be injected at the corresponding nodes. We use the Xitip program [18], which relies on the framework in [19], to show that $H(X) \leq 5/3$ is implied by the set of equalities (6). Therefore, 5/3 is an upper bound on the secrecy rate when the location of wiretapper is unknown, which is less than the secrecy rate 2 achievable when such information is known. Therefore, there is a strict gap between the secrecy capacity and the achievable rate in (2).

On the other hand, the secrecy rate for the wireline network in Fig. 5 is at least 1 which is shown by the example in Fig. 6, where a finite field $GF(7)$ is used. In this example, a combination of Strategies 1 and 2 is used, where keys are injected inside the network and are also canceled at intermediate nodes.

B. Unequal Link Capacities (Scenario 2)

We have restricted the wiretapped links to be in the middle layer in Fig. 5. We next show that the rate in (2) also holds for the secure network coding problem with unequal link capacities (Scenario 2). We convert the example of Fig. 5 by partitioning each nonmiddle layer link into $\frac{1}{\epsilon}$ parallel small links each of which has capacity ϵ. Any three links can be wiretapped in the transformed graph. We prove the unachievability of the rate in (2) in the transformed network.

First, we show a lower bound on the min-cut between the source and the sink in the transformed network when three links are deleted. Note that deleting any k' ($k' \leq 3$) nonmiddle layer links reduces the min-cut by at most $k'\epsilon$. When $k' = 0$, the min-cut is 2. When $k' = 1$ or at most two middle layer links are deleted, the min-cut is at least 2 after deleting these middle layer links, and the min-cut is at least $2 - k'\epsilon \geq 2 - \epsilon$ after further deleting the $k' = 1$ nonmiddle layer link. When $k' = 2$ or at most one middle layer link is deleted, the min-cut between the source and the sink is 3 after deleting this middle layer link, and the min-cut is at least $3 - k'\epsilon \geq 3 - 3\epsilon$ after further deleting the k' nonmiddle layer links. Therefore, the cut-set bound is at least $\min(2 - \epsilon, 3 - 3\epsilon)$.

For the case where the location of the wiretap links is unknown, we prove the unachievability of the rate in (2) in the transformed network. First, consider the transformed network with the restriction that the wiretapper can only wiretap any three links in the middle layer. The optimal solution is exactly the same as for the original network of the previous section and achieves secrecy rate at most $5/3$. Now, consider the transformed network without the restriction on wiretapping set, i.e., the wiretapper can wiretap any three links in the entire network. As wiretapping only the middle layer links is a subset of all possible strategies that the wiretapper can have, the secrecy rate in the transformed network is less than or equal to that in the former case, which is strictly smaller than the rate in (2) for ϵ strictly smaller than $1/4$. Therefore, this rate is still unachievable when the wiretap links are unrestricted in the transformed graph.

V. NP-Hardness

We show in the following that determining the secrecy capacity is NP-hard by reduction from the clique problem, which determines whether a graph contains a clique[4] of a given size r.

The network interdiction problem, which is to minimize the maximum flow of the network when a given number of links in the network are removed, is shown in [21] to be NP-hard, by showing that for any clique problem on a given graph $\mathcal{H}$, there exists a corresponding network $\hat{\mathcal{G}}^{\mathcal{H}}$ whose interdiction capacity is r if and only if $\mathcal{H}$ contains a clique of size r. We use a similar construction and show that for this family of networks $\hat{\mathcal{G}}^{\mathcal{H}}$, the secrecy capacity is equal to the interdiction capacity.

[4]A clique of size r in a graph is a set of r pairwise adjacent vertices, or in other words, an induced subgraph which is a complete graph.

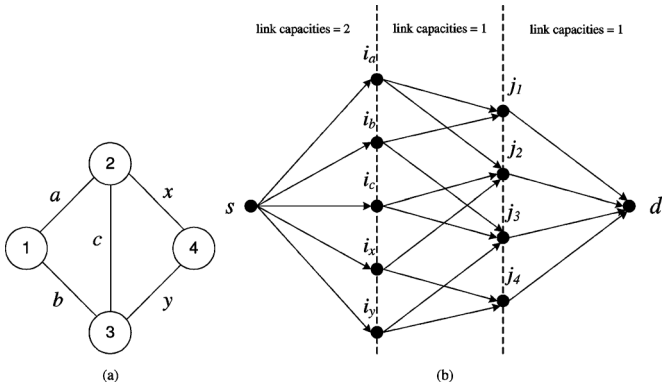

Fig. 7. Example of NP-hardness proof for the case with knowledge of the wiretapping set. (a) Original Graph $\mathcal{H}$. (b) Transformed Graph $\mathcal{G}^{\mathcal{H}}$.

Since the upper and lower bounds used to establish the secrecy capacity are the same regardless of whether the wiretap set is known or unknown, and whether capacity is defined with either zero or asymptotically negligible decoding error and either perfect, strong, or weak secrecy, it follows that determining secrecy capacity under these different definitions is NP-hard.

We briefly describe the approach in [21] in the following. Given an undirected graph $\mathcal{H} = (\mathcal{V}_h, \mathcal{E}_h)$, we will define a capacitated directed network $\hat{\mathcal{G}}^{\mathcal{H}}$ such that there exists a set of links $\hat{\mathcal{A}}'$ in $\hat{\mathcal{G}}^{\mathcal{H}}$ containing less than or equal to $|\mathcal{E}_h| - \binom{r}{2}$ links such that $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'$ has a maximum flow of r if and only if $\mathcal{H}$ contains a clique of size r. For a given undirected graph $\mathcal{H} = (\mathcal{V}_h, \mathcal{E}_h)$ without parallel links and self loops, we create a capacitated, directed graph $\mathcal{G}^{\mathcal{H}} = (\mathcal{N}, \mathcal{A})$ as follows: For each link $e \in \mathcal{E}_h$, create a node i_e in a node set $\mathcal{N}_1$, and for each vertex $v \in \mathcal{V}_h$, create a node j_v in a node set $\mathcal{N}_2$. In addition, create source node s and destination node d. For each link $e \in \mathcal{E}_h$, direct a link in $\mathcal{G}^{\mathcal{H}}$ from s to i_e with capacity 2 and call this set of links $\mathcal{A}_1$. For each link $e = (u, v) \in \mathcal{E}_h$, direct two links in $\mathcal{G}^{\mathcal{H}}$ from i_e to j_v and j_u with capacity 1, respectively, and call this set of links $\mathcal{A}_2$. For each vertex $v \in \mathcal{V}_h$, direct a link with capacity 1 from j_v to d. Let this be the set of links $\mathcal{A}_3$. This completes the construction of $\mathcal{G}^{\mathcal{H}} = (\mathcal{N}, \mathcal{A}) = (\{s\} \cup \{d\} \cup \mathcal{N}_1 \cup \mathcal{N}_2, \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3)$. In Fig. 7, we give an example of the graph transformation, where $\mathcal{H} = (\{1, 2, 3, 4\}, \{a, b, c, x, y\})$. We use the following result from [21].

Lemma 1 [21, Lemma 2]: Let $\mathcal{G}^{\mathcal{H}}$ be constructed from $\mathcal{H}$ as above. Then, there exists a set of links $\mathcal{A}'_1 \subseteq \mathcal{A}_1$ with $|\mathcal{A}'_1| = |\mathcal{E}_h| - \binom{r}{2}$ such that the maximum flow from s to d in $\mathcal{G}^{\mathcal{H}} - \mathcal{A}'_1$ is r if and only if $\mathcal{H}$ contains a clique of size r. ■

After obtaining $\mathcal{G}^{\mathcal{H}}$, we generate $\hat{\mathcal{G}}^{\mathcal{H}}$ by replacing each link (i_e, j_v) with $|\mathcal{E}_h|$ parallel links each with capacity $1/|\mathcal{E}_h|$ and call this link set $\hat{\mathcal{A}}_2$. We carry out the same procedure for links (j_v, d) and call this link set $\hat{\mathcal{A}}_3$. Then, $\hat{\mathcal{G}}^{\mathcal{H}} = (\mathcal{N}, \mathcal{A}) = (\{s\} \cup \{d\} \cup \mathcal{N}_1 \cup \mathcal{N}_2, \mathcal{A}_1 \cup \hat{\mathcal{A}}_2 \cup \hat{\mathcal{A}}_3)$. It is shown in [21] that the worst case set $\hat{\mathcal{A}}'$ is a subset of $\mathcal{A}_1$; therefore, by Lemma 1 and the NP-hardness of the clique problem, finding the interdiction capacity is NP-hard.

Now, we consider the secrecy capacity when $k = |\mathcal{E}_h| - \binom{r}{2}$. From Lemma 1, the condition that $\mathcal{H}$ contains a clique of size r is equivalent to the condition that the max-flow to the sink in $\mathcal{G}^{\mathcal{H}}$ after removing a set of k links from $\mathcal{A}_1$ is r. We now show that the latter condition is equivalent to the condition that the secrecy capacity of $\mathcal{G}^{\mathcal{H}}$ when the wiretapper accesses any unknown subset of k links from $\mathcal{A}_1$ (Scenario 1) is r. By Theorem 1, the secrecy capacity is upper bounded by r whether or not the wiretap set is known. To show achievability, for each subset $\mathcal{A}'$ of k links from $\mathcal{A}_1$, we create nodes $t^{\mathcal{A}'_1}$ and $d^{\mathcal{A}'_1}$ with their corresponding incident links as described in Strategy 1. As the wiretapped links each have capacity 2 and are connected to the source directly, the min-cut between the source and each virtual sink $d^{\mathcal{A}'_1}$ is at least $2k + r$. Then, by using Strategy 1, the secrecy rate r is achievable even if the wiretap set is unknown.

Finally, we show that the same condition is also equivalent to the condition that the secrecy capacity of $\hat{\mathcal{G}}^{\mathcal{H}}$ when any k links are wiretapped (Scenario 2) is r. By the same arguments as in the previous scenario, the secrecy capacity is upper bounded by r, and when only first layer links are wiretapped, the min-cut between the source and each virtual sink $d^{\mathcal{A}'_1}$ is at least $2k + r$. Since each second layer link has a single first layer link as its only input, wiretapping a second layer link yields no more information to the wiretapper than wiretapping a first layer link. When some links in the third layer are wiretapped, let the wiretap set be $\hat{\mathcal{A}}' = \hat{\mathcal{A}}'_1 \cup \hat{\mathcal{A}}'_3$ where $|\hat{\mathcal{A}}'_3| \geq 1$ and $|\hat{\mathcal{A}}'_1| \leq k - 1$. Thus, $\mathcal{A}_1 - \hat{\mathcal{A}}'_1$ contains at least $\binom{r}{2} + 1$ links. We create nodes $t^{\hat{\mathcal{A}}'}$ and $d^{\hat{\mathcal{A}}'}$ with their corresponding incident links as described in Strategy 1. Since removing links in $\mathcal{A}_1$ corresponds to removing links in $\mathcal{H}$, after removing links in $\mathcal{H}$ corresponding to $\hat{\mathcal{A}}'_1$, $\mathcal{H}$ contains a subgraph $\mathcal{H}_1$ containing $\binom{r}{2}$ links plus at least one link $e = (u, v)$.

Case 1: $\mathcal{H}_1$ is a clique of size r. In this case, the number of vertices with degree greater than 0 in $\mathcal{H}_1 \cup e$ is $r + 2$.

Case 2: $\mathcal{H}_1$ is not a clique. $\mathcal{H}_1$ contains at least $r+1$ vertices with degree greater than 0.

According to [21, Lemma 1], the max-flow in $\mathcal{G}^{\mathcal{H}}$ is equal to the number of vertices in $\mathcal{H}$ with degree greater than 0. In both cases, the max-flow of $\mathcal{G}^{\mathcal{H}}$ after removing links in $\hat{\mathcal{A}}'_1$ is at least $r + 1$. Let $\tilde{R}_{s \to \hat{\mathcal{A}}'_3}$ be the max-flow capacity from the source to $\hat{\mathcal{A}}'_3$ in $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'_1$.

We can use a variant of the Ford–Fulkerson (augmenting paths) algorithm, e.g., [22], as follows to construct a max-flow subgraph $\mathcal{D}$ from s to $\hat{\mathcal{A}}'_3$ in $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'_1$ satisfying the property that after removing $\mathcal{D}$ from $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'_1$, the min-cut between s and d is at least

$$\begin{aligned} r + 1 - \tilde{R}_{s \to \hat{\mathcal{A}}'_3} &\geq r + 1 - |\hat{\mathcal{A}}'_3|/|\mathcal{E}_h| \\ &\geq r + 1 - (|\mathcal{E}_h| - 1)/|\mathcal{E}_h| \\ &> r \end{aligned} \tag{7}$$

where we have used $|\hat{\mathcal{A}}'_3| \leq |\mathcal{E}_h| - 1$. Considering the network $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'_1$ with all link directions reversed, we construct augmenting paths via depth first search from d to s, starting first by constructing augmenting paths via links in $\hat{\mathcal{A}}'_3$, until we obtain a set of paths corresponding to a max flow of capacity $\tilde{R}_{s \to \hat{\mathcal{A}}'_3}$ between s and $\hat{\mathcal{A}}'_3$. We add further augmenting paths until we obtain a max flow (of capacity at least $r + 1$) between s and d, which may cause some of the paths traversing links in $\hat{\mathcal{A}}'_3$ to be redefined but without changing their total capacity.

The subgraph $\mathcal{D}$ consists of the final set of paths traversing links in $\hat{\mathcal{A}}'_3$. Thus, the paths remaining after removing $\mathcal{D}$ have a total capacity lower bounded by (7). Therefore, the min-cut between the source and $d^{\hat{\mathcal{A}}'}$ in $\hat{\mathcal{G}}^{\mathcal{H}} - \hat{\mathcal{A}}'_1 - \mathcal{D}$ is at least r, and the min-cut between the source and $d^{\hat{\mathcal{A}}'}$ in $\hat{\mathcal{G}}^{\mathcal{H}}$ is at least $r + R_{s\to\hat{\mathcal{A}}'_1} + \tilde{R}_{s\to\hat{\mathcal{A}}'_3} = r + R_{s\to\hat{\mathcal{A}}'}$. By using Strategy 1, a secure rate of r is achievable. We have thus proved the following theorem.

Theorem 3: For a single-source single-sink network consisting of point-to-point links, computing the secrecy capacity is NP-hard.

VI. Conclusion

In this paper, we have addressed the secrecy capacity of wireline networks where different links have different capacities and restricted wiretapping sets. For the case of networks with equal capacity links where any k links can be wiretapped, the secrecy capacity is given by a cut-set bound if random keys are injected at the source, whether or not the communicating users have information about the choice of the wiretap set. In contrast, we have shown that this rate is unachievable in general in the case of unknown wiretap sets or unequal capacity links, even if nonsource nodes can generate randomness. Further, we have proposed achievable linear optimization-based strategies where random keys are canceled at intermediate nonsink nodes or injected at intermediate nonsource nodes. Finally, we have shown that determining the secrecy capacity is an NP-hard problem.

Appendix
Proof of Theorem 2

We prove Theorem 2 by contradiction. Suppose that a secrecy rate of 2 is achievable for the network in Fig. 5. As before, let X and K denote, respectively, the secret message and random key injected by the source node, and S_i the signal on link i. Then, each triple of links in the middle layer has zero mutual information with the source data, and each pair of links in the middle layer has joint conditional entropy 2 given the other three links.

Since the message X is decodable from information on the last layer, we have $I(S_6, S_7, S_8; X) = 2$. Since $I(S_1, S_2, S_3; X) = 0$, by the data processing inequality $I(S_6; X) = 0$, therefore, $I(S_7, S_8; X|S_6) = 2$ and $H(S_7|S_6) = I(S_7; X|S_6) = 1$. Then, $H(S_7|X, S_6) = H(S_7|S_6) - I(S_7; X|S_6) = 0$. This implies that S_7 does not depend on random keys injected by node f or $\mathrm{head}(4)$ which would be independent of X, S_6. Similarly, $I(S_8; X) = 0$, implying $H(S_7|S_8) = I(S_7; X|S_8) = 1$ and $H(S_7|X, S_8) = 0$. Thus, S_7 does not depend on random keys injected by node $\mathrm{head}\ (2)$ which would be independent of X and S_6. In a similar manner, we can show that S_6 and S_8 also do not depend on any random keys injected after the middle layer. Also, since $H(S_7, S_8|S_6) \geq I(S_7, S_8; X|S_6) = 2$ and $H(S_6) \geq H(S_6|S_8) = 1$, therefore $H(S_6, S_7, S_8) = 3$. Let S_A denote the adversary's observations. By the secrecy requirement, $H(S_6, S_7, S_8|S_A) = 2$, which implies $I(S_6, S_7, S_8; S_A) = H(S_6, S_7, S_8) - H(S_6, S_7, S_8|S_A) = 1$.

Then, the mutual information $I(S_6; S_2, S_3) = 0$; otherwise, if the adversary sees signals 2–4, his mutual information with signals 6–7 is greater than 1. The mutual information $I(S_8; S_1, S_4) = 0$; otherwise, if the adversary sees signals 1, 2, and 4, his mutual information with signals 7–8 is greater than 1. The mutual information $I(S_8; S_4, S_5) = 0$; otherwise, if the adversary sees signals 2, 4, and 5, his mutual information with signals 7–8 is greater than 1. The mutual information $I(S_7; S_4, S_5) = 0$; otherwise, if the adversary sees signals 1, 4, and 5, his mutual information with signals 7–8 is greater than 1.

Case 1: signal 5 is a function of only signals present at the source node, i.e., $H(S_5|X, K) = 0$. By the zero mutual information condition for links 1, 2, and 5, $H(S_1, S_2, S_5|X) = 3$, so

$$H(S_1, S_2, S_5|X, K) = H(S_1, S_2|X, K, S_5) = 2. \quad (8)$$

Since S_4 is conditionally independent of S_1, S_2 given X and K, we have $H(S_1, S_2|X, K, S_4, S_5) = 2$, $I(S_1, S_2; X, K, S_4, S_5) = 0$ and $I(S_1, S_2; X, K|S_4, S_5) = 0$. Now

$$\begin{aligned} &I(S_1, S_2, S_7, S_8; X, K|S_4, S_5) \\ &= I(S_7, S_8; X, K|S_4, S_5) + I(S_1, S_2; X, K|S_7, S_8, S_4, S_5) \\ &= I(S_1, S_2; X, K|S_4, S_5) + I(S_7, S_8; X, K|S_1, S_2, S_4, S_5). \end{aligned} \quad (9)$$

Since S_7, S_8 is conditionally independent of X, K given S_1, S_2, S_4, S_5, we have

$$I(S_7, S_8; X, K|S_1, S_2, S_4, S_5) = 0. \quad (10)$$

Then, by the nonnegativity of conditional mutual information

$$I(S_7, S_8; X, K|S_4, S_5) \leq I(S_1, S_2; X, K|S_4, S_5) = 0. \quad (11)$$

Next, note that S_1 and S_2 are conditionally independent given S_4 and S_5, since $H(S_1|S_4, S_5) = H(S_2|S_1, S_4, S_5) = 1$. Therefore, S_7 and S_8 are conditionally independent given S_4 and S_5, i.e., $I(S_7; S_8|S_4, S_5) = 0$. Since $H(S_7|S_4, S_5) = H(S_7) - I(S_7; S_4, S_5) = 1$, it follows that $H(S_7|S_8, S_4, S_5) = 1$. Then, we have

$$\begin{aligned} &I(S_7, S_8; S_4, S_5) \\ &= I(S_8; S_4, S_5) + I(S_7; S_4, S_5|S_8) \\ &= I(S_8; S_4, S_5) + H(S_7|S_8) - H(S_7|S_4, S_5, S_8) \\ &= 0 + 1 - 1 = 0. \end{aligned} \quad (12)$$

So, $I(S_7, S_8; X, K, S_4, S_5) = I(S_7, S_8; X, K|S_4, S_5) + I(S_7, S_8; S_4, S_5) = 0$, and therefore, $H(S_7, S_8|X) \geq H(S_7, S_8|X, K, S_4, S_5) = 2$, which contradicts the requirement that there is at most 1 unit of random key on the last layer.

Case 2: signal 5 is not a function only of signals present at the source

Case 2a: signal 1 has nonzero mutual information with some random key injected at node c. Then,

$H(S_1|X,K,S_2,S_3,S_4) > 0$. For brevity, let $A = (S_2,S_3)$ and $Y = (X,K,S_4)$. Since $I(S_6;A) = 0$ and $H(S_6|S_1,A) = 0$, we have $H(A)+H(S_6) = H(A,S_6) \leq H(A,S_1) = H(S_1) + H(A|S_1)$. Since $H(S_6) = H(S_1)$, we have $H(A) = H(A|S_1)$ and so $H(S_1|A) = H(S_1)$. Then, from $H(S_1,S_6|A) = H(S_1|A,S_6) + H(S_6|A) = H(S_6|A,S_1) + H(S_1|A)$, we have $H(S_1|A,S_6) = 0$. Since $H(S_1|A,Y,S_6) \leq H(S_1|A,S_6) = 0$ and $H(S_6|A,Y,S_1) \leq H(S_6|A,S_1) = 0$, from

$$\begin{aligned} I(S_1;S_6|Y,A) &= H(S_1|A,Y) - H(S_1|A,Y,S_6) \\ &= H(S_6|A,Y) - H(S_6|A,Y,S_1) > 0 \end{aligned} \tag{13}$$

we have $H(S_6|A,Y) = H(S_1|A,Y) > 0$. Then, since $H(S_7|S_2,S_4) = 0$, we have $H(S_6|S_7,X) > 0$. Also, since $H(S_7|X) = 1$, we have $H(S_6,S_7|X) > 1$.

Case 2b: signal 1 has zero mutual information with any random key injected at node c. Then, $H(S_5|X,K,S_1,S_2,S_4) > 0$. Similar reasoning as for case 2a applies with $A = (S_1,S_4)$, $Y = (X,K,S_2)$, S_5 in place of S_1, and S_8 in place of S_6.

From Cases 1, 2a, and 2b, we conclude that the secrecy rate without knowledge of the wiretapping set by using any nonlinear or linear coding strategy is smaller than 2, obtained for the case where such knowledge is present at the source.

REFERENCES

[1] L. H. Ozarow and A. D. Wyner, "Wire-tap channel II," in *Proc. Eurocrypt*, Apr. 1984, pp. 33–50.

[2] C. E. Shannon, "Communication theory of secrecy systems," *Bell Syst. Tech. J.*, vol. 28, pp. 656–715, 1948.

[3] N. Cai and R. Yeung, "Secure network coding," in *Proc. IEEE Int. Symp. Inf. Theory*, Jun. 2002, p. 323.

[4] R. W. Yeung and N. Cai, "On the optimality of a construction of secure network codes," in *Proc. IEEE Int. Symp. Inf. Theory*, Toronto, ON, Canada, Jul. 2008, pp. 166–170.

[5] J. Feldman, T. Malkin, R. Servedio, and C. Stein, "On the capacity of secure network coding," presented at the Allerton Conf. Commun., Control, Comput., Sep. 2004.

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

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

[8] N. Cai and R. W. Yeung, "A security condition for multi-source linear network coding," in *Proc. IEEE Int. Symp. Inf. Theory*, Nice, France, Jun. 2007, pp. 561–565.

[9] K. Bhattad and K. R. Nayayanan, "Weakly secure network coding," presented at the presented at the WINMEE, RAWNET, NETCOD Workshops, Riva del Garda, Italy, 2005.

[10] A. Mills, B. Smith, T. Clancy, E. Soljanin, and S. Vishwanath, "On secure communication over wireless erasure networks," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2008, pp. 161–165.

[11] M. Langberg and M. Médard, "On the multiple unicast network coding conjecture," in *Proc. 47th Annu. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, Sep. 2009, pp. 222–227.

[12] T. Chan and A. Grant, "Mission impossible: Computing the network coding capacity region," in *Proc. IEEE Int. Symp. Inf. Theory*, Toronto, ON, Canada, Jul. 2008, pp. 320–324.

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

[14] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger, "On randomized network coding," presented at the Allerton Conf. Commun., Control, Comput., Sep. 2003.

[15] L. R. Ford and D. R. Fulkerson, *Flows in Networks*. Princeton, NJ: Princeton Univ. Press, 1962.

[16] R. W. Yeung, *Information Theory and Network Coding*. New York: Springer, Aug. 2008.

[17] 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.

[18] Xitip – Information-Theoretic Inequalities Prover [Online]. Available: http://xitip.epfl.ch/

[19] R. W. Yeung, "A framework for linear information inequalities," *IEEE Trans. Inf. Theory*, vol. 43, no. 6, pp. 1924–1934, Nov. 1997.

[20] V. Bohm, "On the continuity of the optimal policy set for linear programs," *SIAM J. Appl. Math.*, vol. 28, no. 2, pp. 303–306, 1975.

[21] R. K. Wood, "Deterministic network interdiction," *Math. Comput. Model.*, vol. 17, no. 2, pp. 1–18, 1993.

[22] B. C. Dean, M. X. Goemans, and N. Immorlica, "Finite termination of "augmenting path" algorithms in the presence of irrational problem data," *Lecture Notes Comput. Sci.*, vol. 4168, pp. 268–279, 2006.

Tao Cui (S'04) received the M.Sc. degree in the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada, in 2005, and the M.S. degree and Ph.D. degree from the Department of Electrical Engineering, California Institute of Technology, Pasadena, USA, in 2006 and 2009, respectively. His research interests are in the interactions between networking theory, communication theory, and information theory.

Tracey Ho (M'06–SM'11) is an assistant professor of electrical engineering and computer science at the California Institute of Technology. She received the B.S., M.Eng., and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in 1999, 1999, and 2004 respectively. She was a co-recipient of the IEEE Communications Society and Information Theory Society Joint Paper Award in 2009 for "A Random Linear Network Coding Approach to Multicast." Her primary research interests are in information theory, network coding, and communication networks.

Jörg Kliewer S'97–M'99–SM'04) received the Dipl.-Ing. (M.Sc.) degree in Electrical Engineering from Hamburg University of Technology, Germany, in 1993 and the Dr.-Ing. degree (Ph.D.) in Electrical Engineering from the University of Kiel, Germany, in 1999, respectively.

From 1993 to 1998, he was a Research Assistant at the University of Kiel, and from 1999 to 2004, he was a Senior Researcher and Lecturer with the same institution. In 2004, he visited the University of Southampton, Southampton, U.K., for one year, and from 2005 until 2007, he was with the University of Notre Dame, Notre Dame, IN, as a Visiting Assistant Professor. In August 2007, he joined New Mexico State University, Las Cruces, NM, as an Assistant Professor. His research interests include information theory, error correcting codes, network coding, and communication networks.

Dr. Kliewer was the recipient of a Leverhulme Trust Award and a German Research Foundation Fellowship Award in 2003 and 2004, respectively. He was a Member of the Editorial Board of the EURASIP *Journal on Advances in Signal Processing* from 2005–2009 and is Associate Editor of the IEEE TRANSACTIONS ON COMMUNICATIONS since 2008.