On Secure Network Coding with Uniform Wiretap Sets

Wentao Huang and Tracey Ho
California Institute of Technology

Michael Langberg[1]
The Open University of Israel

Joerg Kliewer
New Mexico State University

***Abstract*—This paper studies secure unicast communication over a network with uniform wiretap sets and shows that, when network nodes can independently generate randomness, determining the secrecy capacity is at least as difficult as the k-unicast network coding problem. In particular, we show that a general k-unicast problem can be reduced to the problem of finding the secrecy capacity of a corresponding single unicast network with uniform link capacities and any one wiretap link. We propose a low-complexity linear optimization-based achievable strategy involving global random keys that can be generated anywhere in the network, and an efficient greedy algorithm that further improves achieveable rate by exploiting local random keys.**

I. Introduction

The secure network coding problem, introduced by Cai and Yeung [1], concerns information theoretically secure communication over a network where an unknown subset of network links may be wiretapped. A secure code prevents the wiretapper from obtaining information about the message being communicated. The secrecy capacity of a network, with respect to a given collection of possible wiretap sets, is the maximum rate of communication such that for any one of the wiretap sets the secrecy constraints are satisfied. Types of secrecy constraints studied in the literature include perfect secrecy, strong secrecy and weak secrecy. In the uniform setting, i.e. equal capacity links of which any z may be wiretapped, [1] showed that when only the source can generate randomness, the secrecy capacity is given by the cut-set bounds and linear codes suffice to achieve capacity.

This paper considers the problem of finding the secrecy capacity of a network when we allow network nodes in addition to the source to generate independent randomness (i.e. randomness generated at different nodes is statistically independent). We show that a general k-unicast problem can be reduced to a corresponding single unicast secrecy capacity problem with uniform link capacities where any single link can be wiretapped. This implies that determining the secrecy capacity, even in the simple case of a single unicast and uniform wiretap sets of size 1, is at least as difficult as the long-standing open problem of determining the capacity region of multiple-unicast network coding, which is not presently known to be in P, NP or undecidable [2].

The secure network coding problem in the non-uniform setting, i.e. restricted wiretap sets and/or non-uniform link capacities, has been considered by Cui et al. [3], and by Chan and Grant [4], who showed that determining multicast secrecy capacity with restricted wiretap sets is at least as difficult as determining capacity for multiple-unicast network coding. Our reduction is similar to the core ideas appearing in [4] with the following differences which significantly strengthen the result. First, by introducing the idea of key cancellation and replacement at intermediate nodes, our construction does not need to impose restrictions on which links can be wiretapped. Secondly, unlike the reduction in [4] which involves multiple terminals, ours only needs a single destination. Thirdly, while [4] studies perfect secrecy, our results apply to perfect, strong and weak secrecy constraints.

While finding the secure unicast capacity in the uniform setting is difficult, for this case we show a low-complexity linear optimization-based achievable strategy in which any network node may generate global random keys, i.e. random keys that are decoded by both the source and the sink. This approach generalizes the strategy in [1], [5] where only the source generates global keys, and has lower complexity than the linear optimization-based strategies in [3] designed for the non-uniform case. Performance can be further improved by exploiting non-global keys. We propose an efficient algorithm that greedily searches for places where non-global keys can be introduced in place of global keys, thereby increasing the secure communication rate.

II. Model

A network is represented by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of vertices representing network nodes, and $\mathcal{E}$ the set of edges representing network links. Links have unit capacity (unless otherwise specified) and there may be multiple links between a pair of nodes. Each node $i \in \mathcal{V}$ can generate a random variable K_i which is independent of random variables generated at other nodes (the rate of K_i can be upper bounded by the sum of outgoing link capacities of i). Transmissions on the outgoing links of node i can be functions of K_i as well as transmissions received on incoming links of i.

There is a source node $S \in \mathcal{V}$ and a destination node $D \in \mathcal{V}$. S wants to communicate a message M, uniformly drawn from a finite alphabet set $\mathcal{S}_n$, to D using a code with length n. Then the rate of the code is $n^{-1} \log |\mathcal{S}_n|$. We say that a communication rate R is feasible if there exists a sequence of length-n codes such that $|\mathcal{S}_n| = 2^{nR}$ and the probability of decoding error tends to 0 as $n \to \infty$.

[1] This work is done while Michael Langberg was at the California Institute of Technology.

For the secure network coding problem, we specify additionally a collection $\mathcal{A}$ of wiretap link sets, *i.e.*, $\mathcal{A}$ is a collection of subsets of $\mathcal{E}$ such that an eavesdropper can wiretap any one set in $\mathcal{A}$. We consider three kinds of secrecy constraints: the requirement, for all $A \in \mathcal{A}$, that $I(M; X^n(A)) = 0$ corresponds to perfect secrecy; that $I(M; X^n(A)) \to 0$ as $n \to \infty$ corresponds to strong secrecy; and that $\frac{I(M;X^n(A))}{n} \to 0$ as $n \to \infty$ corresponds to weak secrecy, where $X(A) = \{X(a,b) : (a,b) \in A\}$, and $X(a,b)$ is the signal transmitted on the link (a,b). We say a secrecy rate R is feasible if the communication rate R is feasible and the prescribed secrecy condition is satisfied. The secrecy capacity of the network is defined as the supremum of all feasible secrecy rates. In the rest of the paper we study the case that $\mathcal{A}$ is uniform, i.e., $\mathcal{A} = \{A \subset \mathcal{E} : |A| \le z\}$, where z is a specified maximum number of links that can be wiretapped.

III. Multiple Unicast Reduction

A. Reduction of multiple unicast to secure communication

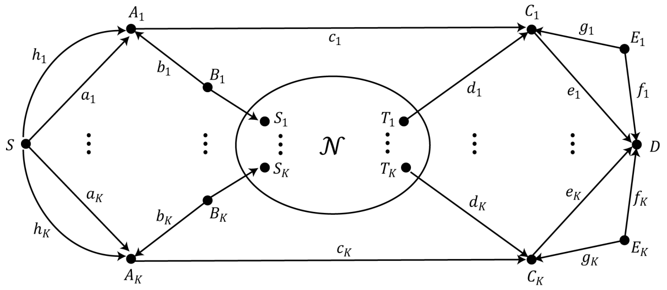

Fig. 1: A secure communication problem with source S and destination D. $\mathcal{N}$ is an arbitrary subnetwork. Links are labeled by the signals transmitted on them. Note there are K branches in total but only the first and the K-th branches are drawn.

The following theorem reduces the multiple unicast network coding problem to the secure network coding problem with only a single unicast and uniform wiretap sets of size 1.

Theorem 1. *Given any unit rate K-unicast problem with source-destination pairs $\{(S_i, T_i),\ i = 1, ..., K\}$ on a network $\mathcal{N}$, the corresponding secure communication problem in Figure 1 with unit capacity links, any one of which can be wiretapped, has secrecy capacity K (under perfect, strong or weak secrecy requirements) if and only if the K-unicast problem is feasible.*

Proof. The secrecy capacity is upper bounded by the capacity K of the min cut from S to D.

"⇒" We show that feasibility of a weak secrecy rate of K implies feasibility of the K-unicast problem. Note that this extends immediately to the cases of strong and perfect secrecy because they imply weak secrecy.

Suppose a secrecy rate of K is achieved by a code with length n. Let M be the source input message, then $H(M) = Kn$. Because there is no shared randomness between different nodes, M is independent with $\{\boldsymbol{d}_1^n,\ \boldsymbol{f}_k^n,\ k = 1, ..., K\}$. Hence

$$H(M|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) = Kn. \tag{1}$$

By the chain rule,

$$\begin{aligned} &H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) + H(\boldsymbol{c}_1^n|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\quad = H(M, \boldsymbol{c}_1^n|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \ge H(M|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n). \end{aligned}$$

So

$$\begin{aligned} H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) &\ge H(M|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\quad - H(\boldsymbol{c}_1^n|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \ge (K-1)n, \end{aligned} \tag{2}$$

where the last inequality holds because of (1) and $H(\boldsymbol{c}_1^n|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \le H(\boldsymbol{c}_1^n) \le n$. Similarly,

$$\begin{aligned} &H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) && (3)\\ &\quad \le H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n, \boldsymbol{e}_2^n, ..., \boldsymbol{e}_K^n) \\ &\qquad + H(\boldsymbol{e}_2^n, ..., \boldsymbol{e}_K^n|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\quad \le n\epsilon_n + H(\boldsymbol{e}_2^n, ..., \boldsymbol{e}_K^n|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) && (4)\\ &\quad \le n\epsilon_n + (K-1)n, && (5) \end{aligned}$$

where $\epsilon_n \to 0$ as $n \to \infty$ and (4) is due to the cut set $\{\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n, \boldsymbol{e}_2^n, ..., \boldsymbol{e}_K^n\}$ from S to D and Fano's inequality. Hence it follows

$$\begin{aligned} H(\boldsymbol{c}_1^n) &\ge H(\boldsymbol{c}_1^n|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\ge H(M|\boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\qquad - H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) && (6)\\ &\ge n - n\epsilon_n, && (7) \end{aligned}$$

where (6) holds because of (2), and (7) follows from (1) and (5). Also notice that

$$\begin{aligned} &H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\quad \ge H(M|\boldsymbol{c}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) - H(\boldsymbol{d}_1^n|\boldsymbol{c}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n), \end{aligned} \tag{8}$$

where

$$H(M|\boldsymbol{c}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) = H(M|\boldsymbol{c}_1^n) \ge Kn - n\delta_n, \tag{9}$$

with $\delta_n \to 0$ as $n \to 0$. Here the first equality holds because $\{M, \boldsymbol{c}_1^n\}$ is independent with $\{\boldsymbol{f}_i^n,\ i = 1, ..., K\}$ and the second inequality holds due to the weak secrecy constraint. Therefore by (5), (8) and (9) we have

$$\begin{aligned} H(\boldsymbol{d}_1^n) &\ge H(\boldsymbol{d}_1^n|\boldsymbol{c}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\ge H(M|\boldsymbol{c}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) - H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) \\ &\ge n - n\epsilon_n - n\delta_n. \end{aligned} \tag{10}$$

Furthermore, by the independency between the sets of $\{M, \boldsymbol{c}_1^n, \boldsymbol{d}_1^n\}$ and $\{\boldsymbol{f}_i^n,\ i = 1, ..., K\}$ we also have

$$H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n, \boldsymbol{f}_2^n, ..., \boldsymbol{f}_K^n) = H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n).$$

According to (2) and (5), it is bounded by

$$(K-1)n \le H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n) \le n\epsilon_n + (K-1)n. \tag{11}$$

Now consider the joint entropy of M, $\boldsymbol{d}_1^n$, $\boldsymbol{c}_1^n$ and expand it in two ways

$$\begin{aligned} H(M, \boldsymbol{d}_1^n, \boldsymbol{c}_1^n) &= H(\boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) + H(M|\boldsymbol{d}_1^n) + H(\boldsymbol{d}_1^n) \\ &= H(M|\boldsymbol{c}_1^n, \boldsymbol{d}_1^n) + H(\boldsymbol{d}_1^n|\boldsymbol{c}_1^n) + H(\boldsymbol{c}_1^n) \\ &\le (K+1)n + n\epsilon_n, \end{aligned}$$

where the last inequality holds because of (11) and $H(\boldsymbol{d}_1^n|\boldsymbol{c}_1^n) \leq n$, $H(\boldsymbol{c}_1^n) \leq n$. Therefore

$$\begin{aligned} H(\boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) &\leq (K+1)n + n\epsilon_n - H(M|\boldsymbol{d}_1^n) - H(\boldsymbol{d}_1^n) \\ &\leq 2n\epsilon_n + n\delta_n, \end{aligned} \tag{12}$$

where (10) and $H(M|\boldsymbol{d}_1^n) = Kn$ (because M and $\boldsymbol{d}_1^n$ are independent by construction) are used to establish the inequality. And so by observing the Markov chain $(M, \boldsymbol{d}_1^n) \rightarrow (M, \boldsymbol{b}_1^n) \rightarrow \boldsymbol{c}_1^n$, it follows

$$\begin{aligned} H(\boldsymbol{c}_1^n|M, \boldsymbol{b}_1^n) &= H(\boldsymbol{c}_1^n|M, \boldsymbol{b}_1^n, \boldsymbol{d}_1^n) \\ &\leq H(\boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) \leq 2n\epsilon_n + n\delta_n. \end{aligned} \tag{13}$$

Then expand the joint entropy of M, $\boldsymbol{b}_1^n$, $\boldsymbol{c}_1^n$ in two ways

$$\begin{aligned} H(M, \boldsymbol{b}_1^n, \boldsymbol{c}_1^n) &= H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n) + H(M|\boldsymbol{c}_1^n) + H(\boldsymbol{c}_1^n) \\ &= H(\boldsymbol{c}_1^n|M, \boldsymbol{b}_1^n) + H(M|\boldsymbol{b}_1^n) + H(\boldsymbol{b}_1^n) \\ &\leq (K+1)n + 2n\epsilon_n + n\delta_n, \end{aligned}$$

where the last inequality holds due to (13) and $H(M|\boldsymbol{b}_1^n) = Kn$, $H(\boldsymbol{b}_1^n) \leq n$. Therefore by (7) and the weak secrecy constraint $H(M|\boldsymbol{c}_1^n) \geq Kn - n\delta_n$, we have

$$\begin{aligned} H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n) &\leq (K+1)n + 2n\epsilon_n + n\delta_n - H(M|\boldsymbol{c}_1^n) - H(\boldsymbol{c}_1^n) \\ &\leq 3n\epsilon_n + 2n\delta_n. \end{aligned} \tag{14}$$

So

$$\begin{aligned} H(\boldsymbol{b}_1^n|M, \boldsymbol{d}_1^n) &\leq H(\boldsymbol{b}_1^n, \boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) \\ &= H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n, \boldsymbol{d}_1^n) + H(\boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) \\ &\leq H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n) + H(\boldsymbol{c}_1^n|M, \boldsymbol{d}_1^n) \\ &\leq 3n\epsilon_n + 2n\delta_n + 2n\epsilon_n + n\delta_n = 5n\epsilon_n + 3n\delta_n, \end{aligned}$$

where the last inequality invokes (14) and (12). Notice that M is independent with $\{\boldsymbol{b}_1^n, \boldsymbol{d}_1^n\}$, so

$$H(\boldsymbol{b}_1^n|\boldsymbol{d}_1^n) = H(\boldsymbol{b}_1^n|M, \boldsymbol{d}_1^n) \leq 5n\epsilon_n + 3n\delta_n. \tag{15}$$

Now we bound the entropy of $\boldsymbol{b}_1^n$. Consider the joint entropy,

$$\begin{aligned} H(M, \boldsymbol{b}_1^n, \boldsymbol{c}_1^n) &= H(\boldsymbol{c}_1^n|M, \boldsymbol{b}_1^n) + H(M|\boldsymbol{b}_1^n) + H(\boldsymbol{b}_1^n) \\ &= H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n) + H(M|\boldsymbol{c}_1^n) + H(\boldsymbol{c}_1^n) \\ &\geq (K+1)n - n\epsilon_n - n\delta_n, \end{aligned}$$

where the last inequality holds because of (7), the secrecy condition $H(M|\boldsymbol{c}_1^n) \geq Kn - n\delta_n$, and $H(\boldsymbol{b}_1^n|M, \boldsymbol{c}_1^n) \geq 0$. So by (13) and because $H(M|\boldsymbol{b}_1^n) = Kn$, we have

$$\begin{aligned} H(\boldsymbol{b}_1^n) &\geq (K+1)n - n\epsilon_n - n\delta_n - H(\boldsymbol{c}_1^n|M, \boldsymbol{b}_1^n) - H(M|\boldsymbol{b}_1^n) \\ &\geq n - 3n\epsilon_n - 2n\delta_n. \end{aligned} \tag{16}$$

Finally, by (15) and (16),

$$I(\boldsymbol{b}_1^n; \boldsymbol{d}_1^n) \geq H(\boldsymbol{b}_1^n) - H(\boldsymbol{b}_1^n|\boldsymbol{d}_1^n) \geq n - 8n\epsilon_n - 5n\delta_n,$$

The above argument extends to all other paths naturally (by renumbering the notations accordingly), so

$$I(\boldsymbol{b}_i^n; \boldsymbol{d}_i^n) \geq n - 8n\epsilon_n - 5n\delta_n, \quad \forall i = 1, ..., K.$$

Therefore $\forall i = 1, ..., K$, by the channel coding theorem, if we employ an outer code of length n by encoding $\boldsymbol{b}_i^n$ as a supersymbol, then there exists an inner code that achieves a rate of $n - 8n\epsilon_n - 8n\delta_n$ from B_i to T_i, and so the overall rate is

$$R_i \geq \frac{n - 8n\epsilon_n - 5n\delta_n}{n} \rightarrow 1 \quad \text{as } n \ \rightarrow \infty.$$

Because B_i can be viewed as a virtual source of S_i, so $\forall i = 1, ..., K$, the unicast from node S_i to T_i of rate 1 is feasible.

"$\Leftarrow$" We show that feasibility of the K-unicast problem implies achievability of secure communication rate K under perfect secrecy, which implies strong and weak secrecy. This rate corresponds to the cut set upper bound on secure communication rate.

Secrecy rate K is achieved by the scheme described in Figure 2, where uniform random local keys are injected by the source and certain intermediate nodes. Denoting by $\epsilon_n^{(i)}$ the probability of error for the unicast from S_i to T_i on network $\mathcal{N}$, then the probability of error from S to D is upper bounded by $K\epsilon_n^* \rightarrow 0$ as $n \rightarrow \infty$, where $\epsilon_n^* = \max_i \epsilon_n^{(i)}$. Note that the scheme achieves perfect secrecy, since links in $\mathcal{N}$ are not downstream of S, and all other links carry uniform random keys of unit rate, or a linear combination involving such a key. □

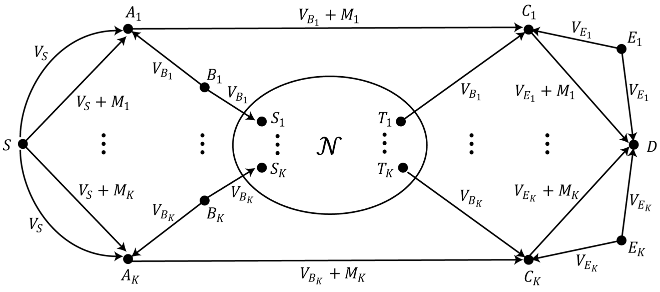

Fig. 2: A scheme to achieve secrecy rate K. V_x is the local key injected by node x, with $H(V_x) = 1$, $\forall x$. M_i, $i = 1, ..., K$ are source input messages, with $H(M_i) = 1$, $i = 1, ..., K$.

The above result can be easily extended to the case of zero error communication with perfect secrecy. In this case, we say a rate R is feasible if there exists a code with finite length n such that $|\mathcal{S}_n| = 2^{nR}$ and the probability of decoding error is strictly zero. Then for the secrecy communication problem in Figure 1, its zero error perfect secrecy capacity is K if and only if the K-unicast for source-destination pairs $\{(S_i, T_i),\ i = 1, ..., K\}$ of unit rate is feasible with zero error. The proof of this claim follows the same outline as the proof of Theorem 1, with the difference that all ϵ_n and δ_n become strictly 0. Then (15) implies that $\boldsymbol{b}_1^n$ is a function of $\boldsymbol{d}_1^n$, and hence that it can be perfectly decoded from $\boldsymbol{d}_1^n$.

B. Reduction of secure communication to multiple unicast

Conversely to the reduction above, for any weakly or strongly secure communication problem, we can construct a communication problem without security constraints (which can in turn be reduced to an equivalent multiple unicast problem [6]) that is feasible if and only if the secure communication problem is feasible. The constructed communication network, called the A-enhanced network and described in [7], [8], adds additional sinks with communication demands. The result for strongly secure communication follows from the equivalence of the capacity region for weak and strong security, shown in [9].

IV. Achievable Strategies

While finding the secure unicast capacity is difficult, below we show a low-complexity linear optimization-based achievable strategy that generalizes the global key strategies in [1], [5] by allowing all network nodes to generate randomness. We further propose an efficient algorithm that greedily searches for places where local keys (keys that not known to both the source and the sink) can be introduced in place of global keys, thereby increasing the secure communication rate.

A. Global Key Schemes

We consider a class of achievable secure coding schemes in which all random keys are global, i.e. decodable with zero error by both the source and the sink, and perfect secrecy is ensured by having global keys of total rate equal to z, the number of wiretapped links. An example is given in Fig 3.

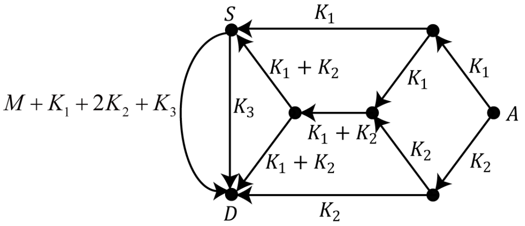

Fig. 3: A network with source S, sink D, unit link capacities and $z = 3$ wiretap links. Global keys K_1 and K_2 are injected by node A. Global key K_3 and message M are injected by S. Illustrated scheme achieves unit rate with perfect secrecy.

Let k_i denote the rate of the random key generated by a node $i \in \mathcal{V}$. A global key scheme corresponds to a multi-source multicast with receiver nodes S and D, and with the following source nodes: node S which has rate $R_S + k_S$, and nodes $i \in \mathcal{V}, i \neq S$ which have rate k_i. Since the capacity region of a multi-source multicast problem is given by cut set bounds [10, Theorem 2.3], the following linear program characterizes the optimal secure message rate achievable using a coding scheme within this class. Here, $\Lambda_t(\mathcal{U})$ denotes the collection of all cuts between $\mathcal{U} \subset \mathcal{V}$ and $t \in \mathcal{V}$ in $\mathcal{G}$, and I_C denotes the capacity of a cut C.

$$\max\ R_S \tag{LP1}$$

$$\text{s.t. } I_C \geq \sum_{i \in \mathcal{U}} r_i, \quad \forall C \in \Lambda_S(\mathcal{U}),\ \forall \mathcal{U} \in 2^{\mathcal{V}\setminus\{S\}} \tag{17}$$

$$I_C \geq \sum_{i \in \mathcal{U}} r_i, \quad \forall C \in \Lambda_D(\mathcal{U}),\ \forall \mathcal{U} \in 2^{\mathcal{V}\setminus\{D\}} \tag{18}$$

$$r_i = k_i, \quad \forall i \in \mathcal{V},\ i \neq S \tag{19}$$

$$r_S = k_S + R_S \tag{20}$$

$$\sum_{i \in \mathcal{V}} k_i \geq z \tag{21}$$

The number of constraints in LP1 is exponential in the size of $\mathcal{V}$. Since [10, Theorem 2.3] implies that a multi-source multicast is feasible if it is feasible for each sink separately, which is equivalent to a multi-commodity flow, we can formulate the following equivalent linear program with number of constraints linear in $|\mathcal{E}|$. In our network model each edge $(u, v) \in \mathcal{E}$ has capacity $c(u, v) = 1$.

$$\max\ R_S \tag{LP2}$$
$$\text{s.t.}$$

$$\sum_{v:(u,v)\in\mathcal{E}} f_D(u,v) - \sum_{v:(v,u)\in\mathcal{E}} f_D(v,u) = \begin{cases} R_S + k_S & u = S \\ -R_S - k_S & u = D \\ 0 & \text{o.w.} \end{cases} \tag{22}$$

$$\sum_{v:(u,v)\in\mathcal{E}} f_i^D(u,v) - \sum_{v:(v,u)\in\mathcal{E}} f_i^D(v,u) = \begin{cases} k_i^D & u = i \\ -k_i^D & u = D \\ 0 & \forall\text{o.w.} \end{cases} \tag{23}$$

$$\sum_{v:(u,v)\in\mathcal{E}} f_S(u,v) - \sum_{v:(v,u)\in\mathcal{E}} f_S(v,u) = \begin{cases} k_D & u = D \\ -k_D & u = S \\ 0 & \forall u \neq S, D \end{cases} \tag{24}$$

$$\sum_{v:(u,v)\in\mathcal{E}} f_i^S(u,v) - \sum_{v:(v,u)\in\mathcal{E}} f_i^S(v,u) = \begin{cases} k_i^S & u = i \\ -k_i^S & u = S \\ 0 & \forall u \neq i, S \end{cases} \tag{25}$$

$$f_S(e),\ f_D(e),\ f_i^D(e),\ f_i^S(e) \geq 0, \forall i, \forall e \in \mathcal{E}$$

$$f_D(e) + \sum_{i \in \mathcal{V}\setminus\{S,D\}} f_i^D(e) \leq c(e), \forall e \in \mathcal{E} \tag{26}$$

$$f_S(e) + \sum_{i \in \mathcal{V}\setminus\{S,D\}} f_i^S(e) \leq c(e), \forall e \in \mathcal{E} \tag{27}$$

$$k_i \leq k_i^S,\quad k_i \leq k_i^D, \quad \forall i \in \mathcal{V}\setminus\{S, D\} \tag{28}$$

$$k_S + k_D + \sum_{i \in \mathcal{V}\setminus\{S,D\}} k_i \geq z \tag{29}$$

Here f_D and $\{f_i^D\}$ represent a multi-commodity flow to D; f_S and $\{f_i^S\}$ represent a multi-commodity flow to S. Equations (22), (23), (24) and (25) are flow conservation constraints; (26) and (27) are link capacity constraints; (28) and (29) ensure secrecy.

Theorem 2. *Let R_S^* be the optimal value of LP2. Then there exists a global key scheme that achieves message rate R_S^* with perfect secrecy.*

Proof. We consider an augmented network $\mathcal{G}'$ obtained from $\mathcal{G}$ as follows. 1) Connect each subset $A \in \mathcal{A}$ to a virtual sink node D^A observing all information on that set of links. Specifically, $\forall (i, j) \in \mathcal{E}$, create node $v_{i,j}$ and replace (i, j) by two edges $(i, v_{i,j})$ and $(v_{i,j}, j)$, then $\forall (i, j) \in A$, create edge $(v_{i,j}, D^A)$, all of unit capacity. 2) Create a virtual message source v_M which is connected to S by edge (v_M, S) of capacity R_S^*. 3) Create a virtual key source v_K which is connected to each node $i \in \mathcal{V}$ by an edge (v_K, i) of capacity k_i corresponding to the solution of LP2. 4) For every $A \in \mathcal{A}$, let $R_{v\to A}$ be the max flow (min cut) capacity from $\{v_M, v_K\}$ to D^A. Note that this equals the max flow capacity from v_K to each virtual sink D^A, because v_M has only one outgoing link (v_M, S), the max flow capacity from v_K to S is R_w (since LP2 ensures that S decodes all keys), and $R_{v\to A} \leq z \leq R_w$. 5) Connect the key source v_K to each virtual sink D^A with edge

(v_K, D^A) of capacity $R_w - R_{v \to A}$, where $R_w = \sum_{i \in \mathcal{V}} k_i$ is the total key rate. Connect the message source v_M to each virtual sink D^A by edge (v_M, D^A) of capacity R_S^*.

Then consider the two-source multicast problem on $\mathcal{G}'$ where S, D and the virtual sinks $\{D^A\}$ each demand rate R_S^* from message source v_M and rate R_w from the key source v_K. The constraints in LP2 guarantee that the required flows exist for S and D, and by construction the additional edges from v_M and v_K guarantee this for each virtual sink D^A. Therefore, the multicast network coding problem is feasible [10]. Furthermore, the total rate observed by each virtual sink D^A is exactly equal to the sum $R_S^* + R_w$ of the message and key rates. This implies that in a capacity-achieving code for this multicast problem, the information received by each virtual sink D^A from set A must be independent of information received on the additional edges from v_M and v_K, which includes the entire source message. Therefore, this corresponds to a code achieving rate R_S^* for the original secrecy problem with perfect secrecy. □

B. Greedy Local Key Enhancement

We describe an algorithm for acyclic graphs that greedily adds local keys to improve a given global key solution.

Let R_S, $\{k_i\}$, f_D, f_S, $\{f_i^D\}$,$\{f_i^S\}$ be the corresponding values from LP2. Consider the residual graph with respect to flows f_D and $\{f_i^D\}$. Any path in this graph from S to a node i with $k_i > 0$ implies that the corresponding keys injected at i can be instead injected at S, by applying the classic augmenting flow algorithm [11] along this path, switching the source of the flow to S. Search for all such paths and switch the key sources to S. If there exists no such path, let Y be the set of all vertices that can be reached from S by a path in the residual graph, Y^c the set of remaining nodes. For any $x \in Y$, $y \in Y^c$, if $(x, y) \in \mathcal{E}$, then this edge must be saturated by f_D and $\{f_i^D\}$; if $(y, x) \in \mathcal{E}$, then it is not used by f_D and $\{f_i^D\}$ (otherwise contradict the fact that $y \in Y^c$). Therefore the total flow from the sources in Y to the sinks in Y^c equals the capacity of the cut (Y, Y^c). But $S \in Y$ is the only source and $D \in Y^c$ is the only sink, therefore the flow from S to D equals the min cut between them. Hence without loss of generality we assume $R_S + k_S$ is the max flow from S to D.

If $k_S = 0$, the global key scheme is already optimal. Otherwise, decompose f_D into $R_S + k_S$ edge-disjoint paths from the S to D, denoted by $P_1, \cdots, P_{R_S+k_S}$. Let $\mathcal{G}_L = (\mathcal{V}, \mathcal{E}_L)$, where $\mathcal{E}_L$ is the set of edges that is not used by any flows in LP2. We apply Algorithm 1 sequentially to each of these paths. Algorithm 1 has a similar structure as the labeling algorithm [11] used to search an augmenting path, but with the difference that it searches for a series of local key gadgets that covers every edge of a path P_k. A *local key gadget* with respect to P_k is either 1) a path from i to j, where i, j are nodes on P_k, in which case local keys can be injected at i and transmitted to j, or 2) a path from a node v to i and a path from v to j, where v is not on P_k while i, j are on P_k. In this case local keys can be injected at v and transmitted to both i, j. Hence each gadget is associated with two nodes in P_k and is able to protect the sub-path of P_k between these two nodes. Furthermore, E_k is introduced to allow the sub-paths to overlap. The virtual local key sources and additional links are introduced for reusing local keys on multiple paths. Intuitively, if Algorithm 1 returns a local key chain L_k on P_k then path P_k is protected by a series of local keys. If Algorithm 1 returns local key chains on R_S+1 paths among $\{P_1, \cdots, P_{R_S+k_S}\}$, then we can transmit one less global key and one unit more message. Without loss of generality we assume the algorithm returns $L_1, \cdots, L_{R_S+1}$ on the first $R_S + 1$ paths $P_1, \cdots, P_{R_S+1}$.

Below we introduce a secure code that achieves rate R_S+1. We first construct a code $\mathcal{C}'$ that transmits $R_S + 1$ units of message and $z - 1$ global keys as described in Section IV-A. Hence $\mathcal{C}'$ achieves perfect secrecy if up to $z - 1$ links are wiretapped. Note that because the graph is acyclic, any edges on P_i are used for the unicast to D and are not upstream of S, hence it suffices for $\mathcal{C}'$ to perform routing on P_i. For example, in the path-based approach in [12], coding is not required at a link that is upstream of only one sink. Next we construct a code $\mathcal{C}$ that is identical to $\mathcal{C}'$ except on $\mathcal{P} = \{P_1, \cdots, P_{R_S+1}\}$. For $P_k \in \mathcal{P}$, consider the gadgets in L_k, each of them generates an independent local key and delivers the key to two nodes $i, j \in P_k$. The upstream node i adds this local key to the incoming signal and sends the sum as the outgoing signal along P_k. This signal travels downstream and the key is canceled when it reaches j. Such cancellation is possible because $\mathcal{C}'$ performs routing on $\mathcal{P}$. Therefore D can successfully decode.

Next we show $\mathcal{C}$ achieves perfect secrecy. Denote by $\mathcal{E}_P$ the set of edges of the paths $\{P_1, \cdots, P_{R_S+k_S}\}$, $\mathcal{E}_P^c$ as the set of edges used by any flows in LP2 except $\mathcal{E}_P$. Recall the augmented graph $\mathcal{G}'$ used to construct $\mathcal{C}'$, and consider the subgraph $\mathcal{G}'_S$ of $\mathcal{G}'$ by deleting all edges in $\mathcal{E}_P$ and all unused edges (edges that carry no flows). Because $R_S + k_S$ is the max flow from S to D, in $\mathcal{G}'_S$ there are no out-going edges from S (otherwise contradicts the fact that $R_S + k_S$ is the max flow). So the min cut from the set of sources $\{v_i : i \in \mathcal{V}\} \cup \{v_M\}$ to any t^A in $\mathcal{G}'_S$ is at most $z - k_S$, and therefore by the construction of $\mathcal{C}'$, any wiretap set A that does not include at least k_S edges in $\mathcal{E}_P$ is not effective, i.e., not leaking source message information. Hence without loss of generality we assume an effective wiretap set A includes at least k_S edges in $\mathcal{E}_P$. By the construction of $\mathcal{C}'$, $\mathcal{E}_P$ consists of R_S+k_S disjoint paths. Clearly A should eavesdrop distinct paths, and therefore by the pigeonhole principle, at least one path in $\mathcal{P}$ is wiretapped. Denote the wiretapped signal on this path as $X_1 + V_1$, where X_1 is the original signal sent in $\mathcal{C}'$, and V_1 is one local key or a combination of them. Denote the remaining wiretapped signals as $X_2 + V_2, \cdots, X_z + V_z$, where V_i, $i \geq 2$ is either zero or a function of local keys. Note V_1 is a random variable uniformly drawn from a finite field and is independent of $\{M, X_2, \cdots, X_z\}$. Therefore $X_1 + V_1$ also follows uniform distribution and is independent of $\{M, X_2, \cdots, X_z\}$. So

$$\begin{aligned} I(M; X_1 + V_1, \cdots, X_z + V_z) &\leq I(M; X_1 + V_1, X_2, \cdots, X_z) \\ &= I(M; X_2, \cdots, X_z) = 0, \end{aligned}$$

Algorithm 1 Local Key Chain Search for P_k

```
▷ Initialization
set $E_k := \{$reversal of the edges of $P_k\}$
unlabel all nodes
set $\text{pred}_{\text{in}}(j) := 0$, $\text{pred}_{\text{out}}(j) := 0$, $\forall j \in \mathcal{V} \backslash P_k$
set $\text{pred}(j) := (0, 0)$, $\forall j \in P_k$
set $\text{LIST}_{\text{in}} := \emptyset$, $\text{LIST}_{\text{out}} := \emptyset$
label $S$ and set $\text{LIST}_P := \{S\}$
▷ Label nodes that can be reached by gadgets
**while** $D$ is unlabeled and $\text{LIST}_P$, $\text{LIST}_{\text{in}}$ and $\text{LIST}_{\text{out}}$ are not all empty **do**
    remove a node $i$ from either $\text{LIST}_P$, $\text{LIST}_{\text{in}}$ or $\text{LIST}_{\text{out}}$
    **if** $i$ is removed from $\text{LIST}_P$ or $\text{LIST}_{\text{out}}$ **then**
        **for all** $(i, j) \in \mathcal{E}_L$, $j \notin P_k$ **do**
            if $j$ is not labeled IN then label $j$ IN, set $\text{pred}_{\text{in}}(j) := i$ and add $j$ to $\text{LIST}_{\text{in}}$
        **end for**
        **for all** $(j, i) \in \mathcal{E}_L$, $j \notin P_k$ **do**
            if $j$ is not labeled OUT then label $j$ OUT, set $\text{pred}_{\text{out}}(j) := i$ add $j$ to $\text{LIST}_{\text{out}}$
        **end for**
        **for all** $(i, j)$ or $(j, i) \in \mathcal{E}_L$, $j \in P_k$ **do**
            **if** $j$ is not labeled P **then** label $j$ P
                **if** $(i, j) \in \mathcal{E}_L$ **then** set $\text{pred}(j) := (i, IN)$
                **else** set $\text{pred}(j) := (i, OUT)$
                **end if**
                add $j$ to $\text{LIST}_P$
            **end if**
        **end for**
        **if** $i \in P_k$ **then**
            **for all** $(i, j) \in E_k$ **do**
                if $j$ is not labeled P then label $j$ P, set $\text{pred}(j) := (i, IN)$ and add $j$ to $\text{LIST}_P$
            **end for**
        **end if**
    **end if**
    **if** $i$ is removed from $\text{LIST}_{\text{in}}$ **then**
        **for all** $(i, j) \in \mathcal{E}_L$, $j \notin P_k$ **do**
            if $j$ is not labeled IN then label $j$ IN, set $\text{pred}_{\text{in}}(j) := i$ and add $j$ to $\text{LIST}_{\text{in}}$
        **end for**
        **for all** $(i, j) \in \mathcal{E}_L$, $j \in P_k$ **do**
            if $j$ is not labeled P then label $j$ P, set $\text{pred}(j) := (i, IN)$ and add $j$ to $\text{LIST}_P$
        **end for**
    **end if**
    **if** $D$ is labeled **then**
        use predecessor pointers to trace back from $D$ to $S$ and obtain a set of edges $L_k$ consisting of local key gadgets
        ▷ Update $\mathcal{G}_L$
        set $L_k := L_k \backslash E_k$
        set $\mathcal{E}_L := \mathcal{E}_L \backslash L_k$
        partition $L_k$ into a collection of local key gadgets $\mathcal{L}_k$, $\forall$ gadget $l \in \mathcal{L}_k$ , create in $\mathcal{G}_L$ virtual local key source $v_l$. $\forall$ node $i$ visited by $l$, create link $(v_l, i)$ of unity capacity.
        **return** local key chain $L_k$
    **end if**
**end while**
```

and perfect secrecy is achieved.

Figure 4 shows an example in which global key schemes achieve at most unit rate. With local key enhancement rate 2 is achieved. For path (S, A, D) the two non-overlapping gadgets are $\{(S, A)\}$ and $\{(B, A), (B, D)\}$; for path (S, C, E, D) the two overlapping gadgets are $\{(E, S)\}$ and $\{(B, C), (B, D)\}$. Local key V_2 and edge (B, D) are reused for both paths.

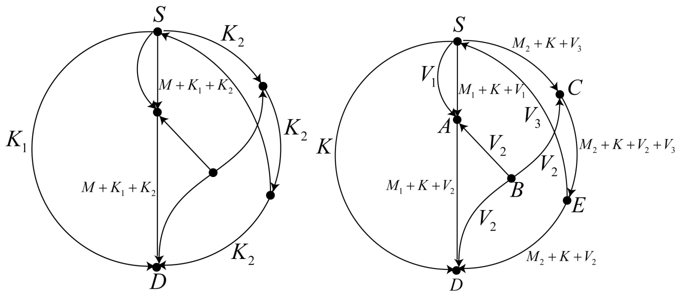

(a) global keys only (b) local keys enhancement

Fig. 4: Network with source S, sink D, unit link capacities and $z = 2$ wiretap links. M_i's are messages, K_i's global keys and V_i's local keys. Both schemes achieve perfect secrecy.

Finally, message rate can be further increased by iteratively applying the above procedure to the resulting solution. At the a-th iteration, $a < k_S$, if Algorithm 1 finds $R_S + a$ local keys chains protecting $R_S + a$ paths, then rate $R_S + a$ is feasible. Due to space limit we defer details to [7].

Acknowledgements

This work has been supported in part by NSF grant CNS 0905615, CCF 0830666 and CCF 1017632, by ISF grant 480/08 and by BSF grant 2010075.

References

[1] N. Cai and R. W. Yeung, "Secure network coding," in *Proc. IEEE Int Information Theory Symp*, 2002.

[2] M. Langberg and M. Medard, "On the multiple unicast network coding conjecture," in *Proc. Allerton 2009*, 2009, pp. 222–227.

[3] T. Cui, T. Ho, and J. Kliewer, "On secure network coding with nonuniform or restricted wiretap sets," *IEEE Trans. Info Theory*, 2012.

[4] T. Chan and A. Grant, "Mission impossible: Computing the network coding capacity region," in *Proc. IEEE ISIT*, July 2008, pp. 320 –324.

[5] R. W. Yeung and N. Cai, "On the optimality of a construction of secure network codes," in *IEEE ISIT*, 2008.

[6] R. Dougherty and K. Zeger, "Nonreversibility and equivalent constructions of multiple unicast networks," *IEEE Transactions on Information Theory*, vol. 52, no. 11, 2005.

[7] W. Huang, T. Ho, M. Langberg, and J. Kliewer, "On secure network coding with uniform wiretap sets," 2013, http://www.its.caltech.edu/~whuang/netcod2013.pdf.

[8] T. Dikaliotis, H. Yao, T. Ho, M. Effros, and J. Kliewer, "Network equivalence in the presence of an eavesdropper," in *Allerton Conference on Communication, Control and Computing*, 2012.

[9] S. Jalali and T. Ho, "On capacity region of wiretap networks," 2012, http://arxiv.org/abs/1212.3859.

[10] T. Ho and D. Lun, *Network Coding: An Introduction*. Cambridge University Press, 2008.

[11] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[12] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. M. G. M. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Info Theory*, 2005.