Secure Network Coding for Multiple Unicast: On the Case of Single Source

Gaurav Kumar Agarwal, Martina Cardone, and Christina Fragouli

Department of Electrical and Computer Engineering
University of California Los Angeles, Los Angeles, CA 90095, USA *
{gauravagarwal, martina.cardone, christina.fragouli}@ucla.edu

Abstract. This paper considers multiple unicast wireline noiseless networks where a single source wishes to transmit independent messages to a set of legitimate destinations. The primal goal is to characterize the secure capacity region, where the exchanged messages have to be secured from a passive external eavesdropper that has unbounded computational capabilities, but limited network presence. The secure capacity region for the case of two destinations is characterized and it is shown to be a function of only the min-cut capacities and the number of edges the eavesdropper wiretaps. A polynomial-time two-phase scheme is then designed for a general number of destinations and its achievable secure rate region is derived. It is shown that the secure capacity result for the two destinations case is not reversible, that is, by switching the role of the source and destinations and by reversing the directions of the edges, the secure capacity region changes.

1 Introduction

Information theoretical network security is increasingly gaining importance, as we are moving towards a quantum computing era. On the one hand, the computational power at our disposal is continuously increasing and on the other, terabytes of data per seconds are exchanged over communication networks, a large portion of which needs to be secure (e.g., banking, professional, health). However, we still have very limited understanding of information theoretical security bounds and schemes over arbitrary networks.

In this paper, we consider an arbitrary wireline noiseless network with unit capacity edges where a source needs to securely transmit information to one or more receivers. A passive external eavesdropper, Eve, wishes to learn some information about the data exchanged between the legitimate nodes. Eve has unbounded computational capabilities (e.g., a quantum computer), but has limited network presence, namely, she can wiretap at most k edges of her choice. Over such a network, information theoretical network security seeks to design transmission schemes that are unconditionally/perfectly secure.

* The work of the authors was partially funded by NSF under awards 1321120 and 1740047. G. K. Agarwal is also supported by the Guru Krupa Fellowship.

Our first main result is to extend the secure network coding capacity, from the single unicast and multicast cases [1], to the case of two unicast sessions. In particular, in a unicast session, if the min-cut capacity between the source and the receiver is M, then we can securely transmit information at rate $M-k$, where k is the number of edges Eve wiretaps (and the same result extends to the case of multicasting [1]). We prove in this paper that, if the source needs to send two independent messages to two receivers, a surprising direct extension of the single unicast case applies, where again the secure capacity region is uniquely determined by the min-cut capacities $M_{\{1\}}, M_{\{2\}}$ and $M_{\{1,2\}}$ (towards the first, second and the union of the two receivers), reduced by the number of the eavesdropped edges k, and thus the network structure plays no role. This is enabled by the observation that the source can establish secure keys with the two receivers that need not to be independent, i.e, they may share common randomness that can be efficiently multicast using network coding techniques. To the best of our knowledge, this is the first result that provides the secure capacity region characterization for a general network where multiple unicast sessions take place simultaneously.

Our second main result focuses on the case where we have an arbitrary number m of unicast sessions. We first derive an outer bound on the secure capacity region and then design a polynomial-time transmission scheme and derive its achievable secure rate region. In particular, our achievable scheme consists of two phases, where first secure keys are exchanged between the source and the destinations, then messages are encoded with these keys and finally transmitted. Although this scheme is not optimal, it is computationally efficient and it provides a performance guarantee on the secure achievable rate region as a function of any rate m-tuple that is achievable in the absence of the eavesdropper Eve.

Finally, we also show that the secure capacity result is irreversible, i.e., the capacity region of the reverse network (obtained by switching the role of the source and the destinations and by reversing the directions of the edges) is not the same as the one of the original network. This is a surprising result since it implies that – different from the unsecure case where irreversible networks must necessary have non-linear network coding solutions [2, 3] – under security constraints even networks with linear network coding solutions can be irreversible if the traffic is multiple unicast.

Related Work. The benefits of network coding were first shown in the seminal paper by Ahlswede et al. [4], where the authors proved that, in a noiseless network (represented by a directed acyclic graph) with single source and multiple destinations, the source can multicast at a rate equal to the minimum among all the min-cut capacities. Later, Li et al. [5] showed that it suffices to use random linear coding operations to achieve the multicast capacity and, more recently Jaggi et al. [6] designed polynomial-time deterministic algorithms to achieve it. While for the case of single unicast and multicast traffic the capacity is well-known, the same is not true for the case of networks where multiple unicast sessions take place simultaneously and share some of the network resources. For instance, even though the cut-set bound was proved to be tight for some special cases, such as single source with non-overlapping demands and single source with non-overlapping demands and a multicast demand [7], in general it is not tight [8]. It was also recently showed by Kamath et al. [9] that characterizing the capacity of a general network where two unicast sessions take place simultaneously is as hard as characterizing the capacity of a network with general number of unicast sessions. For the case of single source and two destinations with a non-overlapping demand and a multicast demand, Ramamoorthy et. al [10] proposed a nice graph theory based approach to characterize the capacity region.

Cai et al. [1] characterized the secure capacity of a network with multicast traffic, where a passive external eavesdropper wiretaps any k edges of her choice. In particular, the authors showed that a secure multicast communication rate of $M - k$ can always be achieved, where M is the minimum among all the min-cut capacities. Also, for a multicast scenario, Cui et al. [11] designed a secure achievable scheme when Eve can wiretap only some of the edges (i.e., among all possible subsets of k edges, the eavesdropper can wiretap only some of them) and when the edge capacities are non-uniform. Since, even in the absence of the eavesdropper, the capacity of a multiple unicast network is not known in general, very few results are available for security. For instance, recently Agarwal et al. characterized the secure capacity region for some variations of the butterfly network both for noiseless [12] and erasure channels [13]. Although the results in [12] and [13] were the first that provided secure capacity results in multiple unicast scenarios, they are tailored to some specific network topologies. We here extend these results to a general multiple unicast network with single source (for which the capacity in absence of Eve is given by the cut-set bound [7]) and we characterize the secure capacity region for the case of two destinations.

Paper Organization. This paper is organized as follows. In Section 2, we define the setup (i.e., the multiple unicast network with single source and general number of destinations) and we formulate the problem. In Section 3, we focus on the secure capacity region characterization for our setup. In particular, we first derive an outer bound that holds for general number of destinations, we then show that this outer bound is tight for the case of two destinations and we finally design a two-phase secure transmission scheme for general number of destinations and compute its achievable rate region. In Section 4, we analyze and compare our designed schemes in terms of performance and complexity. In Section 4, we also show that the secure capacity result is irreversible and we finally conclude the paper.

2 Setup and Problem Formulation

Notation. Calligraphic letters indicate sets; $\emptyset$ is the empty set and $|\mathcal{A}|$ is the cardinality of $\mathcal{A}$; for two sets $\mathcal{A}_1, \mathcal{A}_2$, $\mathcal{A}_1 \subseteq \mathcal{A}_2$ indicates that $\mathcal{A}_1$ is a subset of $\mathcal{A}_2$, $\mathcal{A}_1 \cup \mathcal{A}_2$ indicates the union of $\mathcal{A}_1$ and $\mathcal{A}_2$, $\mathcal{A}_1 \sqcup \mathcal{A}_2$ indicates the disjoint union of $\mathcal{A}_1$ and $\mathcal{A}_2$, $\mathcal{A}_1 \cap \mathcal{A}_2$ is the intersection of $\mathcal{A}_1$ and $\mathcal{A}_2$ and $\mathcal{A}_1 \backslash \mathcal{A}_2$ is the set of elements that belong to $\mathcal{A}_1$ but not to $\mathcal{A}_2$; $[n_1 : n_2]$ is the set of integers from n_1 to $n_2 \geq n_1$; $[x]^+ := \max\{0, x\}$ for $x \in \mathbb{R}$.

We represent a wireline noiseless network with a directed acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes and $\mathcal{E}$ is the set of directed edges. We further assume that each edge $e \in \mathcal{E}$ is of unit capacity. If an edge $e \in \mathcal{E}$ connects a node i to a node j, we refer to node i as the tail and to node j as the head of e, i.e., $\text{tail}(e) = i$ and $\text{head}(e) = j$. For each node $v \in \mathcal{V}$, we define $\mathcal{I}(v)$ as the set of all incoming edges of node v and $\mathcal{O}(v)$ as the set of all outgoing edges of node v.

In this network (graph), there is one source node S and m destination nodes $D_i, i \in [1:m]$. The source node does not have incoming edges, i.e., $\mathcal{I}(S) = \emptyset$, and each destination node does not have outgoing edges, i.e., $\mathcal{O}(D_i) = \emptyset, \forall i \in [1:m]$. The source S has a message W_i for each destination $D_i, i \in [1:m]$. The m messages are assumed to be independent. Thus, this network consists of multiple unicast traffic where m unicast sessions take place simultaneously and share some of the network resources. A passive eavesdropper Eve is also present in the network and can wiretap any k edges of her choice. We highlight that Eve is an external eavesdropper, i.e., it is not one of the destinations.

Each message $W_i, i \in [1:m]$, is of q-ary entropy rate R_i and each channel is a discrete noiseless channel accepting alphabets over $\mathbb{F}_q$. Over this network, we are interested in finding all possible feasible m-tuples $(R_1, R_2, \ldots, R_m)$ such that each destination $D_i, i \in [1:m]$, reliably decodes the message W_i and Eve receives no information about the messages. In particular, we are interested in information theoretic secure communication, i.e., we consider "perfect secrecy".

The symbol transmitted (respectively, received) over n channel uses on edge $e \in \mathcal{E}$ is denoted as X_e^n (respectively, Y_e^n). Similarly, Z_e^n, $e \in \mathcal{E}$, is the symbol received by Eve on edge $e \in \mathcal{E}$ over n channel uses. Clearly, since the channels are noiseless, $Y_{ei} = Z_{ei} = X_{ei}, \forall i \in [1:n]$; throughout the paper, we use these symbols interchangeably. In addition, for $\mathcal{E}_t \subseteq \mathcal{E}$ we define $X_{\mathcal{E}_t}^n = \{X_e^n : e \in \mathcal{E}_t\}$, $Y_{\mathcal{E}_t}^n = \{Y_e^n : e \in \mathcal{E}_t\}$ and $Z_{\mathcal{E}_t}^n = \{Z_e^n : e \in \mathcal{E}_t\}$. We assume that the source node S has an independent and infinite source of randomness Θ, while the other nodes in the network do not have any randomness.

Definition 1. *A rate m-tuple $(R_1, R_2, \ldots, R_m)$ is said to be securely achievable if there exist a block length n, a set of encoding functions f_e, $\forall e \in \mathcal{E}$, such that*

$$X_e^n = \begin{cases} f_e\left(W_1, W_2, \ldots, W_m, \Theta\right) & \textit{if tail}(e) = \{S\} \\ f_e\left(\{Y_\ell^n : \ell \in \mathcal{I}(\textit{tail}(e))\}\right) & \textit{otherwise} \end{cases},$$

and destination D_i can reliably (with zero error) decode the message W_i i.e., $H\left(W_i | \{Y_e^n : e \in \mathcal{I}(D_i)\}\right) = 0$. Moreover, $\forall\ \mathcal{E}_{\mathcal{Z}} \subseteq \mathcal{E}$, $|\mathcal{E}_{\mathcal{Z}}| \leq k$, $I\left(W_{[1:m]}; Z_{\mathcal{E}_{\mathcal{Z}}}^n\right) = 0$ (strong secrecy requirement). The closure of all such feasible rate m-tuples is the secure capacity region.

3 Secure Capacity

In this section we focus on the secure capacity region characterization for the network described in Section 2, when an eavesdropper Eve wiretaps any k edges of her choice. In particular, we first derive an outer bound for a general number

m of destinations and then design a secure transmission scheme that achieves the outer bound for $m = 2$. This result leads to the secure capacity region characterization for $m = 2$. Finally, we provide the design of a two-phase secure achievable scheme for a general number m of destinations and compute its achievable rate.

3.1 Outer Bound

We here derive an outer bound on the secure capacity region of a multiple unicast network with a single source and m destinations. In particular, the outer bound is provided in the next theorem.

Theorem 1. *An outer bound on the secure capacity region for a multiple unicast network with single source and m destinations is given by*

$$R_{\mathcal{A}} \leq [M_{\mathcal{A}} - k]^+, \quad \forall \mathcal{A} \subseteq [1:m] \ , \tag{1}$$

where $R_{\mathcal{A}} := \sum_{i \in \mathcal{A}} R_i$ and $M_{\mathcal{A}}$ is the min-cut capacity between the source S and the set of destinations $D_{\mathcal{A}} := \{D_i : i \in \mathcal{A}\}$.

Proof. Let $\mathcal{E}_{\mathcal{A}}$ be a min-cut between the source S and $D_{\mathcal{A}}$ and $\mathcal{E}_{\mathcal{Z}} \subseteq \mathcal{E}_{\mathcal{A}}$ be the set of k edges wiretapped by Eve and define $\mathcal{I}(D_{\mathcal{A}}) := \bigcup_{i \in \mathcal{A}} \mathcal{I}(D_i)$. If $|\mathcal{E}_{\mathcal{A}}| < k$, let $\mathcal{E}_{\mathcal{Z}} = \mathcal{E}_{\mathcal{A}}$. We have,

$$\begin{aligned} nR_{\mathcal{A}} = H(W_{\mathcal{A}}) &\overset{\text{(a)}}{=} H(W_{\mathcal{A}}) - H(W_{\mathcal{A}}|X^n_{\mathcal{I}(D_{\mathcal{A}})}) \\ &\overset{\text{(b)}}{=} H(W_{\mathcal{A}}) - H(W_{\mathcal{A}}|X^n_{\mathcal{E}_{\mathcal{A}}}) \\ &\overset{\text{(c)}}{=} I(W_{\mathcal{A}}; X^n_{\mathcal{E}_{\mathcal{Z}}}, X^n_{\mathcal{E}_{\mathcal{A}} \setminus \mathcal{E}_{\mathcal{Z}}}) \\ &= I(W_{\mathcal{A}}; X^n_{\mathcal{E}_{\mathcal{Z}}}) + I(W_{\mathcal{A}}; X^n_{\mathcal{E}_{\mathcal{A}} \setminus \mathcal{E}_{\mathcal{Z}}}|X^n_{\mathcal{E}_{\mathcal{Z}}}) \\ &\overset{\text{(d)}}{=} I(W_{\mathcal{A}}; X^n_{\mathcal{E}_{\mathcal{A}} \setminus \mathcal{E}_{\mathcal{Z}}}|X^n_{\mathcal{E}_{\mathcal{Z}}}) \\ &\overset{\text{(e)}}{\leq} H(X^n_{\mathcal{E}_{\mathcal{A}} \setminus \mathcal{E}_{\mathcal{Z}}}) \\ &\overset{\text{(f)}}{\leq} n[M_{\mathcal{A}} - k]^+ \ , \end{aligned}$$

where $W_{\mathcal{A}} = \{W_i, i \in \mathcal{A}\}$ and where: (i) the equality in (a) follows because of the decodability constraint; (ii) the equality in (b) follows because $X^n_{\mathcal{I}(D_{\mathcal{A}})}$ is a deterministic function of $X^n_{\mathcal{E}_{\mathcal{A}}}$; (iii) the equality in (c) follows from the definition of mutual information and since $\mathcal{E}_{\mathcal{A}} = \mathcal{E}_{\mathcal{Z}} \cup \mathcal{E}_{\mathcal{A} \setminus \mathcal{Z}}$; (iv) the equality in (d) follows because of the perfect secrecy requirement; (v) the inequality in (e) follows since the entropy of a discrete random variable is a non-negative quantity and because of the 'conditioning reduces the entropy' principle; (vi) finally, the inequality in (f) follows since each link has unit capacity and since $|\mathcal{E}_{\mathcal{A}} \setminus \mathcal{E}_{\mathcal{Z}}| = [M_{\mathcal{A}} - k]^+$. By dividing both sides of the above inequality by n we obtain that $R_{\mathcal{A}}$ in (1) is an outer bound on the capacity region of the multiple unicast network with single source and m destinations. This concludes the proof of Theorem 1.

Remark 1. Since the eavesdropper Eve wiretaps any k edges of her choice, intuitively Theorem 1 states that if she wiretaps k edges of a cut with capacity M, we can at most hope to reliably transmit at rate $M-k$. However, this holds only for the case of single source; indeed, as we will see in Section 4.2 through an example, higher rates can be achieved for the case of single destination and multiple sources.

3.2 Secure Capacity Region for $m = 2$

We here prove that the outer bound in Theorem 1 is indeed tight for the case $m = 2$. In particular, our main result is stated in the following theorem.

Theorem 2. *The outer bound in* (1) *is tight for the case* $m = 2$*, i.e., the secure capacity region of a multiple unicast network with single source and* $m = 2$ *destinations is given by*

$$R_1 \leq [M_{\{1\}} - k]^+ \ , \tag{2a}$$

$$R_2 \leq [M_{\{2\}} - k]^+ \ , \tag{2b}$$

$$R_1 + R_2 \leq [M_{\{1,2\}} - k]^+ \ . \tag{2c}$$

Proof. Clearly, from the result in Theorem 1, the rate region in (2) is an outer bound on the capacity region of a multiple unicast network with single source and $m = 2$ destinations. Hence, we now need to prove that the rate region in (2) is also achievable. Towards this end, we start by providing the following definition of *separable* graphs.

Definition 2. *A graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ *with a single source and* m *destinations is said to be* ***separable*** *if its edge set* $\mathcal{E}$ *can be partitioned as* $\mathcal{E} = \sqcup_{\ell=1}^{2^m-1} \mathcal{E}'_\ell$ *such that* $\mathcal{G}'_\ell = (\mathcal{V}, \mathcal{E}'_\ell)$ *and*

$$M_{\mathcal{A}} = \sum_{\substack{\mathcal{J} \subseteq [1:m] \\ \mathcal{J} \cap \mathcal{A} \neq \emptyset}} M^\star_{\mathcal{J}}, \ \forall \mathcal{A} \subseteq [1:m] \ ,$$

where $M_{\mathcal{A}}$ *is the min-cut capacity between the source* S *and the set of destinations* $D_{\mathcal{A}} := \{D_i : i \in \mathcal{A}\}$ *in* $\mathcal{G}$ *and* $M^\star_{\mathcal{J}}$ *is the min-cut capacity between the source* S *and the set of destinations* $D_{\mathcal{B}} := \{D_b : b \in \mathcal{B}\}$, $\forall \mathcal{B} \subseteq \mathcal{J}$ *for the graph* $\mathcal{G}'_\ell$ *with* $\ell \in [1 : 2^m - 1]$ *being the decimal representation of the binary vector of length* m *that has a one in all the positions indexed by* $j \in \mathcal{J}$ *and zero otherwise, with the least significant bit in the first position.*

To better understand the above definition, consider a graph $\mathcal{G}$ with $m = 2$ destinations. Then, the graph $\mathcal{G}$ is separable if it can be partitioned into 3 graphs such that:

- $\mathcal{G}'_1$ has the following min-cut capacities: $M^\star_{\{1\}}$ from S to D_1 and zero from S to D_2,

- $\mathcal{G}_2'$ has the following min-cut capacities: zero from S to D_1 and $M^\star_{\{2\}}$ from S to D_2,
- $\mathcal{G}_3'$ has the following min-cut capacities: $M^\star_{\{1,2\}}$ from S to D_1, $M^\star_{\{1,2\}}$ from S to D_2 and $M^\star_{\{1,2\}}$ from S to $\{D_1, D_2\}$,

where the quantities $M^\star_{\{1\}}$, $M^\star_{\{2\}}$ and $M^\star_{\{1,2\}}$ can be computed using the following set of equations:

$$M_{\{i\}} = M^\star_{\{i\}} + M^\star_{\{1,2\}}, \forall i \in [1:2] \ , \tag{3a}$$

$$M_{\{1,2\}} = M^\star_{\{1\}} + M^\star_{\{2\}} + M^\star_{\{1,2\}} \ . \tag{3b}$$

For example, consider the network $\mathcal{G}_0$ in Fig. 1(a), which has min-cut capacities $M_{\{1\}} = M_{\{2\}} = 3$ and $M_{\{1,2\}} = 4$. It is not difficult to see that $\mathcal{G}_0$ in Fig. 1(a) can be partitioned in three graphs $\mathcal{G}_i', i \in [1:3]$ as shown in Figs. 1(b)-1(d), with min-cut capacities equal to (see (3)) $M^\star_{\{1\}} = M^\star_{\{2\}} = 1$ and $M^\star_{\{1,2\}} = 2$.

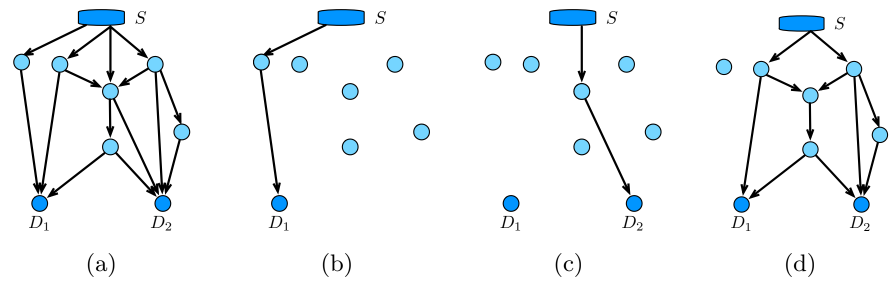

Fig. 1: A 2-destination separable network $\mathcal{G}_0$ in Fig. 1(a) and its partition graphs $\mathcal{G}_i', i \in [1:3]$ in Figs. 1(b)-1(d).

We now state the following lemma, which is a direct consequence of [10, Theorem 1] and we will use to prove the achievability of the rate region in (2).

Lemma 1. *Any graph with a single source and $m = 2$ destinations is separable.*

For completeness we report the proof of Lemma 1 in Appendix A. By leveraging the result in Lemma 1, we are now ready to prove Theorem 2. In particular, we consider two cases depending on the value of k (i.e., the number of edges the eavesdropper wiretaps). Without loss of generality, we assume that $k < \min_{i\in[1:2]} M_i$, as otherwise secure communication to the set of destinations $\{D_i : k \geq M_i\}$ is not possible at any rate, and hence we can just remove this set of destinations from the network.

1. **Case 1:** $k \geq M^\star_{\{1,2\}}$. In this case, by substituting the quantities in (3) into (2), we obtain that the constraint in (2c) is redundant. Thus, we will

now prove that the rate pair $(R_1, R_2) = (M_{\{1\}} - k, M_{\{2\}} - k)$ is securely achievable, which along with the time-sharing argument proves the achievability of the entire region in (2).

We denote with $y_1, y_2, \ldots, y_k$ the k key packets and with $m_i^{(1)}, m_i^{(2)}, \ldots, m_i^{(R_i)}$ (with $i \in [1:2]$) the R_i message packets for D_i. With this, our scheme is as follows:

- We multicast $y_i, \forall i \in [1 : M^\star_{\{1,2\}}]$, to both D_1 and D_2 using $\mathcal{G}'_3$, which has edges denoted by $\mathcal{E}'_3$. This is possible as $\mathcal{G}'_3$ has a min-cut capacity $M^\star_{\{1,2\}}$ to both D_1 and D_2 (see Definition 2).
- We unicast $y_\ell, \forall \ell \in [M^\star_{\{1,2\}} + 1 : k]$, to $D_i, \forall i \in [1 : 2]$, using $k - M^\star_{\{1,2\}}$ paths out of the $M^\star_{\{i\}}$ disjoint paths in $\mathcal{G}'_i$. We denote by $\hat{\mathcal{E}}_i$ the set that contains all the first edges of these paths. Clearly, $|\hat{\mathcal{E}}_i| = k - M^\star_{\{1,2\}}, \forall i \in [1 : 2]$. Notice that $\hat{\mathcal{E}}_i \subseteq \mathcal{E}'_i, \forall i \in [1 : 2]$ (see Definition 2).
- We send the $R_i, \forall i \in [1 : 2]$, encrypted message packets (i.e., encoded with the keys) of D_i on the remaining $M^\star_{\{i\}} - k + M^\star_{\{1,2\}}$ disjoint paths in $\mathcal{G}'_i$. We denote by $\bar{\mathcal{E}}_i$ the set that contains all the first edges of these paths in $\mathcal{G}'_i$. Clearly, $|\bar{\mathcal{E}}_i| = R_i, \forall i \in [1 : 2]$, $\bar{\mathcal{E}}_i \subseteq \mathcal{E}'_i$ and $\bar{\mathcal{E}}_i \cap \hat{\mathcal{E}}_i = \emptyset$ (see Definition 2).

This scheme achieves $R_i = M^\star_{\{i\}} - k + M^\star_{\{1,2\}} = M_{\{i\}} - k, \forall i \in [1 : 2]$, where the second equality follows by using the definitions in (3). Now we prove that this scheme is also secure. We start by noticing that, thanks to Definition 2, the edges $\mathcal{E}'_3$, $\hat{\mathcal{E}}_i$ and $\bar{\mathcal{E}}_i$, with $i \in [1 : 2]$, do not overlap. We write these transmissions in a matrix form (with G and U being the encoding matrices) and we obtain

$$\begin{bmatrix} X_{\mathcal{E}'_3} \\ X_{\hat{\mathcal{E}}_1} \\ X_{\hat{\mathcal{E}}_2} \end{bmatrix} = \underbrace{\begin{bmatrix} g_{11} & g_{12} & \cdots & g_{1k} \\ g_{21} & g_{22} & \cdots & g_{1k} \\ \vdots & \vdots & \ddots & \vdots \\ g_{\ell 1} & g_{\ell 2} & \cdots & g_{\ell k} \end{bmatrix}}_{G} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix}, \ \ell = |\mathcal{E}'_3| + 2\left(k - M^\star_{\{1,2\}}\right) \ ,$$

$$\begin{bmatrix} X_{\bar{\mathcal{E}}_1} \\ X_{\bar{\mathcal{E}}_2} \end{bmatrix} = \underbrace{\begin{bmatrix} u_{11} & u_{12} & \ldots & u_{1k} \\ u_{21} & u_{22} & \ldots & u_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ u_{r1} & u_{r2} & \ldots & u_{rk} \end{bmatrix}}_{U} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} \oplus \begin{bmatrix} m_1^{(1)} \\ \vdots \\ m_1^{(R_1)} \\ m_2^{(1)} \\ \vdots \\ m_2^{(R_2)} \end{bmatrix}, \ r = R_1 + R_2 \ .$$

The eavesdropper Eve wiretaps $k_1 \leq k$ edges from the collection of edges $\{\mathcal{E}'_3, \hat{\mathcal{E}}_1, \hat{\mathcal{E}}_2\}$, over which the linear combinations $X_{\mathcal{E}'_3}$, $X_{\hat{\mathcal{E}}_1}$ and $X_{\hat{\mathcal{E}}_2}$ of keys are transmitted, and $k_2 = k - k_1$ edges from the collection of edges $\{\bar{\mathcal{E}}_1, \bar{\mathcal{E}}_2\}$ over which the messages encoded with the keys $X_{\bar{\mathcal{E}}_1}$ and $X_{\bar{\mathcal{E}}_2}$ are transmitted.

We here note that on the other edges $\mathcal{E}\backslash\{\mathcal{E}_3' \cup \hat{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_1 \cup \hat{\mathcal{E}}_2 \cup \bar{\mathcal{E}}_2\}$, of the network, we either do not transmit any symbol or simply route the symbols from $\{X_{\bar{\mathcal{E}}_1}, X_{\bar{\mathcal{E}}_2}, X_{\hat{\mathcal{E}}_1}, X_{\hat{\mathcal{E}}_2}\}$ (corresponding to the symbols transmitted on disjoint paths). Thus, without loss of generality, we can assume that Eve does not wiretap any of these edges. Since the first $|\mathcal{E}_3'|$ rows of G (i.e., those that correspond to multicasting of they keys) are determined by the network coding scheme for multicasting [4], we assume that we do not have any control over the construction of G.

Thus, we would like to construct the code matrix U such that all the linear combinations of the keys used to encrypt the messages are mutually independent and are independent from the linear combinations of the keys wiretapped on the k_1 edges (notice that this makes the symbols wiretapped by the eavesdropper completely independent from the messages). In particular, since in the worst case Eve wiretaps k_1 edges which are independent linear combinations, we would like that any matrix formed by k_1 independent rows of the matrix G and k_2 rows of the matrix U is full rank. Since there is a finite number of such choices and the determinant of each of these possible matrices can be written in a polynomial form – which is not identically zero – as a function of the entries of U, then we can choose the entries of U such that all these matrices are invertible. Thus, we can always construct the code matrix U such that the edges wiretapped by Eve have an independent key and hence Eve does not get any information about the message packets, i.e., the scheme is secure. This implies that the rate pair $(R_1, R_2) = (M_{\{1\}} - k, M_{\{2\}} - k)$ is securely achievable.

2. **Case 2:** $k < M^\star_{\{1,2\}}$. By substituting the quantities in (3), the rate region in (2) becomes

$$R_i \leq M_{\{i\}} - k = M^\star_{\{i\}} + M^\star_{\{1,2\}} - k, \forall i \in [1:2] \ , \tag{4a}$$

$$R_1 + R_2 \leq M_{\{1,2\}} - k = M^\star_{\{1\}} + M^\star_{\{2\}} + M^\star_{\{1,2\}} - k \ . \tag{4b}$$

We now show that we can achieve the following two corner points i.e., the rate pair

$$\begin{aligned}(R_1, R_2) &= \big(\alpha(M_{\{1,2\}} - M_{\{2\}}) + (1-\alpha)(M_{\{1\}} - k), \\ &\qquad \alpha(M_{\{2\}} - k) + (1-\alpha)(M_{\{1,2\}} - M_{\{1\}})\big) \\ &\overset{\text{(a)}}{=} (M^\star_{\{1\}} + \alpha(M^\star_{\{1,2\}} - k), M^\star_{\{2\}} + (1-\alpha)(M^\star_{\{1,2\}} - k)) \ , \end{aligned} \tag{5}$$

for $\alpha \in \{0, 1\}$, where the equality in (a) follows by using the definitions in (3). This along with the time-sharing argument proves the achievability of the entire region in (4). We recall that we denote with $y_1, y_2, \dots, y_k$ the k key packets and with $m_i^{(1)}, m_i^{(2)}, \dots, m_i^{(R_i)}$ (with $i \in [1:2]$) the R_i message packets for D_i. With this, our scheme is as follows:

- Using the graph $\mathcal{G}_3'$ we multicast to both destinations D_1 and D_2: (i) $y_i, \forall i \in [1:k]$, (ii) $\alpha(M^\star_{\{1,2\}} - k)$ encrypted message packets (i.e., encoded with the keys) for D_1 and (iii) $(1-\alpha)(M^\star_{\{1,2\}} - k)$ encrypted message packets for D_2. Recall that the edges of the graph $\mathcal{G}'_3$ are denoted by $\mathcal{E}'_3$ (see Definition 2). We also highlight that the message packets multicast to the two destinations are encrypted using the key packets, where the encryption is based on the secure network coding result on multicasting [1], which ensures perfect security from an adversary wiretapping any k edges.

- We send $M^\star_{\{i\}}$ encrypted message packets of D_i on the $M^\star_{\{i\}}$ disjoint paths to D_i in the graph $\mathcal{G}'_i$, and denote by $\hat{\mathcal{E}}_i$ the set that contains all the first edges of these paths for $i \in [1:2]$.

This scheme achieves the rate pair in (5). Now we prove that this scheme is also secure. For ease of representation, in what follows we let $R^\star_1 = \alpha(M^\star_{\{1,2\}} - k)$ and $R^\star_2 = (1-\alpha)(M^\star_{\{1,2\}} - k)$. We again notice that, thanks to Definition 2, the edges $\mathcal{E}'_3$, $\hat{\mathcal{E}}_1$ and $\hat{\mathcal{E}}_2$ do not overlap. We write these transmissions in a matrix form (with G, U and S being encoding matrices) and we obtain,

$$X_{\mathcal{E}'_3} = \underbrace{\begin{bmatrix} g_{11} & g_{12} & \cdots & g_{1k} \\ g_{21} & g_{22} & \cdots & g_{1k} \\ \vdots & \vdots & \ddots & \vdots \\ g_{\ell 1} & g_{\ell 2} & \cdots & g_{\ell k} \end{bmatrix}}_{G} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} \oplus \underbrace{\begin{bmatrix} s_{11} & s_{12} & \cdots & s_{1k} \\ s_{21} & s_{22} & \cdots & s_{1k} \\ \vdots & \vdots & \ddots & \vdots \\ s_{\ell 1} & s_{\ell 2} & \cdots & s_{\ell k} \end{bmatrix}}_{S} \begin{bmatrix} m_1^{(1)} \\ \vdots \\ m_1^{(R^\star_1)} \\ m_2^{(1)} \\ \vdots \\ m_2^{(R^\star_2)} \end{bmatrix}, \ \ell = |\mathcal{E}'_3| \ ,$$

$$\begin{bmatrix} X_{\hat{\mathcal{E}}_1} \\ X_{\hat{\mathcal{E}}_2} \end{bmatrix} = \underbrace{\begin{bmatrix} u_{11} & u_{12} & \cdots & u_{1k} \\ u_{21} & u_{22} & \cdots & u_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ u_{r1} & u_{r2} & \cdots & u_{rk} \end{bmatrix}}_{U} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} \oplus \begin{bmatrix} m_1^{(R^\star_1+1)} \\ \vdots \\ m_1^{(R_1)} \\ m_2^{(R^\star_2+1)} \\ \vdots \\ m_2^{(R_2)} \end{bmatrix}, \ r = R_1 + R_2 - (M^\star_{\{1,2\}} - k) \ .$$

The eavesdropper Eve wiretaps $k_1 \leq k$ edges from $\mathcal{E}'_3$, over which the linear combinations $X_{\mathcal{E}'_3}$ of key packets and message packets are sent, and $k_2 = k - k_1$ edges from the collection of edges $\{\hat{\mathcal{E}}_1, \hat{\mathcal{E}}_2\}$ over which the messages encoded with the keys $X_{\hat{\mathcal{E}}_1}$ and $X_{\hat{\mathcal{E}}_2}$ are transmitted. Similar to Case 1, on the other edges $\mathcal{E} \backslash \{\mathcal{E}'_3 \cup \hat{\mathcal{E}}_1 \cup \hat{\mathcal{E}}_2\}$ of the network, we either do not transmit any symbol or simply route the symbols from $\{X_{\hat{\mathcal{E}}_1}, X_{\hat{\mathcal{E}}_2}\}$ (corresponding to the symbols transmitted on disjoint paths). Thus, without loss of generality, we can assume that the eavesdropper does not wiretap any of these edges. Since the matrices G and S are determined by the secure network coding scheme for multicasting [1], we do not have any control over their construction. Thus, we would like to construct the code matrix U in order to ensure security.

Again, similar to the argument used in Case 1, we can create U such that any subset of k_2 rows of U are linearly independent and not in the span of any subset of k_1 rows of G. With this, the keys used to encrypt the messages over any k_2 edges of $\{\hat{\mathcal{E}}_1, \hat{\mathcal{E}}_2\}$ are mutually independent and independent from the keys used over any k_1 edges of $\mathcal{E}_3'$. This, together with the fact that the messages transmitted using $\mathcal{G}_3'$ are already secure, makes our scheme secure. This implies that the rate pair (R_1, R_2) in (5) is securely achievable.

This concludes the proof of Theorem 2.

3.3 A Two-phase Scheme

We now propose the design of a secure transmission scheme that consists of two phases, namely the key generation phase (in which secret keys are generated between the source and the m destinations) and message sending phase (in which the message packets are first encoded using the secret keys and then transmitted to the m destinations). The corresponding achievable rate region is presented in Theorem 3.

Theorem 3. *Let $(\hat{R}_1, \hat{R}_2, \ldots, \hat{R}_m)$ be an achievable rate m-tuple in absence of the eavesdropper Eve. Then, the rate m-tuple $(R_1, R_2, \ldots, R_m)$ with*

$$R_i = \hat{R}_i \left(1 - \frac{k}{M}\right), \forall i \in [1:m] \ , \tag{6}$$

where M is the minimum min-cut between the source and any destination, is securely achievable in the presence of an eavesdropper Eve who wiretaps any k edges of her choice.

Proof. Let M_i be the min-cut capacity between the source and the destination D_i with $i \in [1:m]$. We define M as the minimum among all these individual min-cut capacities, i.e., $M = \min_{i \in [1:m]} M_i$. Let $(\hat{R}_1, \hat{R}_2, \ldots, \hat{R}_m) \in \mathbb{R}^m$ be the unsecure rate m-tuple achieved in the absence of the eavesdropper. We start by approximating this rate m-tuple with rational numbers; notice that this is always possible since the set of rationals $\mathbb{Q}$ is dense in $\mathbb{R}$. Moreover, an information flow through the network (from the source S to an artificial destination D' connected to all the destinations $D_i, i \in [1:m]$ – see also Appendix B) that achieves this rate m-tuple might involve fractional flows over the edges since the rate m-tuple may be fractional. To make the rate m-tuple integral and thereby also the flow over each edge, we multiply the capacity of each edge by a common factor T. This implies that to achieve $(\hat{R}_1, \hat{R}_2, \ldots, \hat{R}_m)$, then $(T\hat{R}_1, T\hat{R}_2, \ldots, T\hat{R}_m)$ is achieved over T instances of the network after which the flow over each edge is an integer. In what follows, we describe our coding scheme and we show that

$$(R_1, R_2, \ldots, R_m) = \left(1 - \frac{k}{M}\right) (\hat{R}_1, \hat{R}_2, \ldots, \hat{R}_m) \tag{7}$$

is achievable. This particular scheme consists of the two following phases.

- *Key generation.* This first phase – in which secure keys are generated between the source and the destinations – consists of k subphases. In each subphase, the source multicasts $M - k$ random packets securely to all destinations. This is possible thanks to the secure network coding result of [1], since the minimum min-cut capacity is M and Eve has access to k edges. Thus, at the end of this phase, a total of $Tk(M - k)$ secure keys are generated, since in each phase we use the network T times.
- *Message sending.* We choose Tk packets out of the $Tk(M-k)$ securely shared (in the key generation phase) random packets. For each choice of Tk packets, we convert the unsecure scheme achieving $(T\hat{R}_1, T\hat{R}_2, \ldots, T\hat{R}_m)$ to a secure scheme achieving the same rate m-tuple. Towards this end, we expand the Tk shared packets into $\sum_{j=1}^{m} T\hat{R}_j$ packets using an MDS code matrix. With this, we have the same number of random packets as the message packets. We then encode the message packets with the random packets and transmit them as it was done in the corresponding unsecure scheme. We repeat this process until we run out of the shared random packets, i.e., we repeat this process $M - k$ times by using T instances of the network each time.

Proof of security. We know that, in absence of security considerations, a time-sharing based scheme is optimal (i.e., capacity achieving) for a multiple unicast network with single source, i.e., network coding is not beneficial [7] (see also Appendix B) Given that we are not using network coding operations and since each edge carries an integer information flow, then the eavesdropper will be able to wiretap at most Tk different messages each encrypted with an independent key. Hence, the eavesdropper will not be able to obtain any information about any of the m messages.

Analysis of the achieved rate m-tuple. The secure scheme described above requires a total of M phases. In particular, in the first k phases we generate the secure keys and in the remaining $M - k$ phases we securely transmit at rates of $(T\hat{R}_1, T\hat{R}_2, \ldots, T\hat{R}_m)$, over T network instances. Thus, the achieved secure message rate $(R_1, R_2, \ldots, R_m)$ is

$$R_j = \frac{M-k}{M}\hat{R}_j = \left(1 - \frac{k}{M}\right)\hat{R}_j, \forall j \in [1:m] \ . \tag{8}$$

This concludes the proof of Theorem 3.

4 Discussion and Conclusions

In this section, we analyze, discuss and compare the results that we have derived in the paper. In particular, we first compare the secure capacity region in (2) with the capacity region of the same network in the absence of the eavesdropper. We then show that the secure capacity result in (2), different from the unsecure counterpart, is irreversible. We also analyze the complexity of the capacity achieving scheme and of the two-phase scheme. Finally, we summarize our main contributions and conclude the paper.

4.1 Secure Capacity versus Unsecure Capacity

For the network with single source and multiple destinations described in Section 2, the unsecure capacity (i.e., in the absence of the eavesdropper) is well known [7, Theorem 9] and given by the following lemma. For completeness we report the proof of the following lemma in Appendix B.

Lemma 2. *The unsecure capacity region for a multiple unicast network with single source node and m destination nodes is given by*

$$R_{\mathcal{A}} \leq M_{\mathcal{A}}, \quad \forall \mathcal{A} \subseteq [1:m] \ , \tag{9}$$

where $R_{\mathcal{A}} := \sum\limits_{i \in \mathcal{A}} R_i$ and $M_{\mathcal{A}}$ is the min-cut capacity between the source S and the set of destinations $D_{\mathcal{A}} := \{D_i : i \in \mathcal{A}\}$.

We now focus on the case of $m = 2$ destinations and compare the secure capacity region in Theorem 2 and the unsecure capacity region in Lemma 2. By comparing (2) with (9) (evaluated for the case $m = 2$), we observe that in the presence of the eavesdropper we lose at most a rate k in each dimension compared to the unsecure case. We notice that the same result holds for the case of $m = 1$ destination and for the case of multicasting the same message to all receivers [1] (i.e., we have a rate loss of k with respect to the min-cut capacity M). However, here it is more surprising since the messages to the $m = 2$ receivers (and potentially the keys) are different.

4.2 Reversibility and Non-reversibility

In order to characterize the unsecure capacity region in (9) network coding is not necessary and routing is sufficient (see also Appendix B). Thus, from the result in [3], it directly follows that the capacity result in (9) is reversible. In particular, let $\mathcal{G}$ be a network with single source and m destinations with a certain capacity region (that can be computed from Lemma 2). Then, the reverse graph $\mathcal{G}'$ is constructed by switching the role of the source and destinations and by reversing the directions of the edges. Thus, $\mathcal{G}'$ will have m sources and one single destination. The result in [3] ensures that $\mathcal{G}$ and $\mathcal{G}'$ will have the same capacity region, i.e., the result in Lemma 2 characterizes also the unsecure capacity region for a multiple unicast network with m sources and single destination.

We now focus on the secure case. In Section 3, we have characterized the secure capacity region for a multiple unicast network with single source and $m = 2$ destinations. In particular, Theorem 2 implies that the secure capacity region does not depend on the specific topology of the network and it can be fully characterized by the min-cut capacities $M_{\{1\}}, M_{\{2\}}$ and $M_{\{1,2\}}$ and by the number k of edges wiretapped by Eve. We now show that this result is irreversible, i.e., the secure capacity region of the reverse network is not the same as the one of the original network. Moreover, we also show that the secure capacity region with 2 sources and single destination cannot anymore be characterized by only the min-cut capacities, i.e., it depends on the specific network topology.

Consider the three networks in Fig. 2 and assume $k = 1$, i.e., Eve wiretaps one edge of her choice. For the network in Fig. 2(a) we have min-cut capacities $\left(M_{\{1\}}, M_{\{2\}}, M_{\{1,2\}}\right) = (1, 2, 2)$ and hence from Theorem 2 it follows that the secure capacity for this network is given by $(R_1, R_2) = (0, 1)$. This point can be achieved by simply using the scheme shown in Fig. 2(a), where y represents the key and W_2 the message for D_2. Now, consider the network in Fig. 2(b) that is obtained from Fig. 2(a) by switching the role of the source and destinations and by reversing the directions of the edges. For this network, which has the same min-cut capacities as the network in Fig. 2(a), the rate pair $(R_1, R_2) = (1, 0)$ is securely achievable using the scheme shown in Fig. 2(b) where W_1 is the message of S_1 and y_1 and y_2 are the keys generated by S_1 and S_2, respectively. The rate pair $(R_1, R_2) = (1, 0)$, which is securely achieved by the network in Fig. 2(b), cannot be securely achieved by the network in Fig. 2(a). This result implies that a secure rate pair that is feasible for one network might not be feasible for the reverse network, i.e., the secure capacity regions can be different and hence cannot be derived from one another. The achievability of the pair $(R_1, R_2) = (1, 0)$ in Fig. 2(b) also shows that the outer bound in (1) does not hold for the case of single destination and multiple sources, in which case it is possible to achieve rates outside this region.

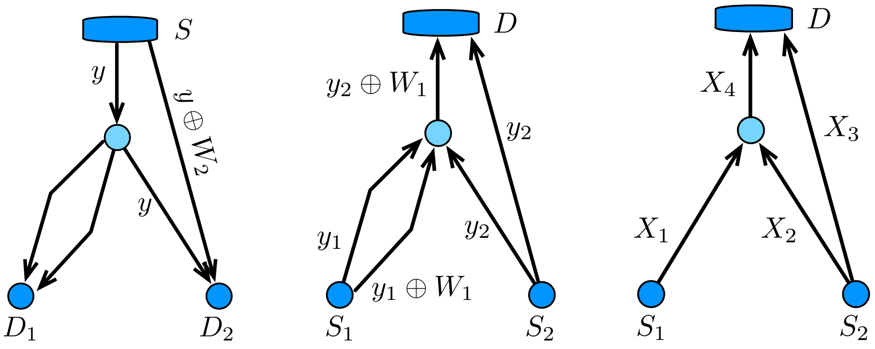

(a) $(R_1, R_2) = (0, 1)$ is capacity.

(b) $(R_1, R_2) = (1, 0)$ is achievable.

(c) $(R_1, R_2) = (1, 0)$ is not achievable.

Fig. 2: Network examples.

Consider now the network in Fig. 2(c), which has the same min-cut capacities $\left(M_{\{1\}}, M_{\{2\}}, M_{\{1,2\}}\right) = (1, 2, 2)$ as the network in Fig. 2(b). We now show that the rate pair $(R_1, R_2) = (1, 0)$, which can be securely achieved in the network in Fig. 2(b), cannot be securely achieved in the network in Fig. 2(c). Let $X_i, i \in$

$[1:4]$, be the transmitted symbols as shown in Fig. 2(c). With this, we have

$$\begin{aligned}
R_1 = H(W_1) \overset{(a)}{=} H(W_1) - H(W_1|X_3, X_4) &\overset{(b)}{\leq} H(W_1) - H(W_1|X_1, X_2, X_3) \\
&= I(W_1; X_1, X_2, X_3) \\
&= I(W_1; X_1) + I(W_1; X_2, X_3|X_1) \\
&\overset{(c)}{=} I(W_1; X_2, X_3|X_1) \\
&= H(X_2, X_3|X_1) - H(X_2, X_3|W_1, X_1) \\
&\overset{(d)}{=} H(X_2, X_3) - H(X_2, X_3) = 0 \ ,
\end{aligned}$$

where: (i) the equality in (a) follows because of the decodability constraint; (ii) the inequality in (b) follows because of the 'conditioning reduces the entropy' principle and since X_4 is a deterministic function of (X_1, X_2); (iii) the equality in (c) follows because of the perfect secrecy requirement; (iv) finally, the equality in (d) follows since (X_2, X_3) is independent of (W_1, X_1). This result shows that the rate pair $(R_1, R_2) = (1, 0)$ is not securely achievable in the network in Fig. 2(c). This implies that, for a network with single destination and multiple sources, we cannot characterize the secure capacity region based only on the min-cut capacities $\left(M_{\{1\}}, M_{\{2\}}, M_{\{1,2\}}\right)$, i.e., the result would depend on the specific network topology.

4.3 Complexity Analysis

The capacity achieving scheme for $m = 2$ destinations that we have proposed (see Section 3.2) first requires that we edge-partition the original graph $\mathcal{G}$ into three graphs (i.e., an edge in $\mathcal{G}$ appears in only one of these three graphs). At this stage, this step requires an exhaustive search over all possible paths in the network, which requires an exponential number of operations in the number of nodes. It therefore follows that the scheme proposed in Section 3.2, even though it allows to characterize the secure capacity region, is of exponential complexity.

Differently, the two-phase scheme proposed in Section 3.3 runs in polynomial time. This is because all the operations that it requires (i.e., find a T such that over T instances all flows are integer, multicast the keys in the key generation phase, encrypt messages at the source (i.e., encode the messages with the keys) and route the encrypted messages) can be performed in polynomial time in the number of edges. However, the two-phase scheme described in Section 3.3 is sub-optimal and does not achieve the outer bound in (1). However, this scheme offers a guarantee on the secure rate region that can always be achieved as a function of any rate m-tuple that is achievable in the absence of the eavesdropper Eve (see (6) in Theorem 3).

One reason behind this is that in the key generation phase some edges in the network are not used. Indeed, when we multicast the M random packets to generate the keys (where M is the minimum of the min-cut capacities and k is the number of edges wiretapped by the eavesdropper) – out of which $M - k$

linear combinations are secure keys – it might have been possible to use the other edges (i.e., those through which the random packets do not flow) to transmit some encrypted message packets. For instance, consider the network example in Fig. 3(a), where the eavesdropper wiretaps $k = 1$ edge of her choice. Our two-phase scheme would multicast $M = \min_{i \in [1:2]} M_{\{i\}} = 2$ random packets y_1 and y_2 (y_1 is transmitted over the solid edges and y_2 over the dashed edges in Fig. 3(a)), out of which $M - k = 1$ is securely received by D_1 and D_2. Hence, the combination $y_1 \oplus y_2$ can be used to securely transmit the message packets. However, we see that in the first phase the dotted edge (i.e., the one that connects S directly to D_2) is not used. This brings to a reduction in the achievable rate region since this edge could have been used to securely transmit a message packet to D_2 by using $W_2 \oplus y_1$ as shown in Fig. 3(a). Given this, we believe that what makes the two-phase scheme suboptimal is the fact that it does not fully leverage all the network resources. In Fig. 3(b), we plotted different rate regions for the network in Fig. 3(a), which has min-cut capacities $M_{\{1\}} = 2$, $M_{\{2\}} = 3$ and $M_{\{1,2\}} = 3$. In particular, the region contained in the solid curve is the unsecure capacity region (given by (9) in Lemma 2), the region inside the dashed curve is the secure capacity region (given by (2) in Theorem 2) and the region contained inside the dotted line is the secure rate region that can be achieved by the two-phase scheme (given by (6) in Theorem 3). From Fig. 3(b), we indeed observe that the rate region achieved by the two-phase scheme is contained inside the secure capacity region.

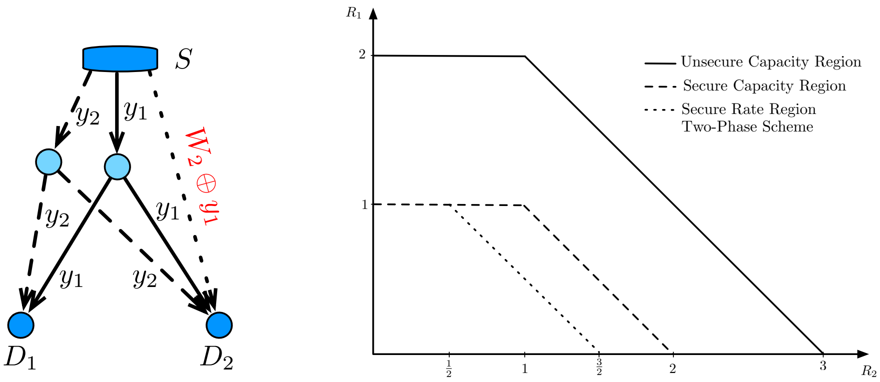

(a) $(R_1, R_2) = (0, 1)$ is capacity.

(b) Rate region for the network in Fig. 3(a).

Fig. 3: Network example for which the two-phase scheme is not optimal.

4.4 Summary

In this paper, we analyzed wireline noiseless networks where a single source would like to convey independent messages to different destinations in the presence of a passive external eavesdropper, who can wiretap any k edges of her choice. We first derived an outer bound on the capacity region that holds for any number of destinations and then showed that this bound is indeed tight for the case of two destinations. To the best of our knowledge, this is the first secure capacity result for a general network where multiple unicast sessions take place simultaneously (i.e., single source and two destinations). We also showed that this secure capacity result, different from the unsecure counterpart, is irreversible. Finally, we have proposed a secure two-phase transmission scheme for general number of destinations and computed its achievable rate region. An appealing feature of this scheme is that, even though it does not achieve the secure capacity region, it can be implemented in polynomial time and it provides a performance guarantee on the secure achievable rate region as a function of any rate tuple that is achievable in the absence of the eavesdropper Eve.

Appendix A

For completeness, we here report the proof of the result in Lemma 1, which is a direct consequence of [10, Theorem 1]. In particular, this result shows that any graph $\mathcal{G}$ with single source and $m = 2$ destinations is separable. The graph $\mathcal{G}$ has min-cut capacity $M_{\{i\}}, i \in [1:2]$, towards destination D_i and min-cut capacity $M_{\{1,2\}}$ towards $\{D_1, D_2\}$, from which $M^\star_{\{i\}}, i \in [1:2]$, and $M^\star_{\{1,2\}}$ can be computed by using the expressions in (3). We represent these min-cut capacities by the triple

$$\left(M_{\{1\}}, M_{\{2\}}, M_{\{1,2\}}\right) = \left(M^\star_{\{1\}} + M^\star_{\{1,2\}}, M^\star_{\{2\}} + M^\star_{\{1,2\}}, M^\star_{\{1\}} + M^\star_{\{2\}} + M^\star_{\{1,2\}}\right) \ ,$$

where the equality follows by using (3). We now prove Lemma 1 in two steps. We first show that the graph $\mathcal{G}$ can be separated into two graphs: $\mathcal{G}_a$ with min-cut capacities $\left(M^\star_{\{1\}}, 0, M^\star_{\{1\}}\right)$ and $\mathcal{G}_b$ with min-cut capacities

$$\left(M^\star_{\{1,2\}}, M^\star_{\{2\}} + M^\star_{\{1,2\}}, M^\star_{\{2\}} + M^\star_{\{1,2\}}\right) \ .$$

Then, by applying the same principle we further separate the graph $\mathcal{G}_b$ into two graphs: $\mathcal{G}_c$ with min-cut capacities $\left(0, M^\star_{\{2\}}, M^\star_{\{2\}}\right)$ and $\mathcal{G}_d$ with min-cut capacities $\left(M^\star_{\{1,2\}}, M^\star_{\{1,2\}}, M^\star_{\{1,2\}}\right)$. This would complete the proof of Lemma 1.

We now prove that we can separate the graph $\mathcal{G}$ into the two graphs $\mathcal{G}_a$ and $\mathcal{G}_b$. Towards this end, from the original graph $\mathcal{G}$, we create a new directed acyclic graph $\mathcal{G}'$ where a new node D' is connected to D_1 through an edge of capacity $M^\star_{\{1\}} + M^\star_{\{1,2\}}$ and to D_2 through an edge of capacity $M^\star_{\{2\}}$. By following similar steps as in the proof of the direct part (achievabiliy) of Lemma 2 (see

Appendix B), it is not difficult to see that in $\mathcal{G}'$ the min-cut capacity between S and D' is $M^\star_{\{1\}} + M^\star_{\{1,2\}} + M^\star_{\{2\}} = M_{\{1,2\}}$, where the equality follows from (3b). From the max-flow min-cut theorem, we can find $M_{\{1,2\}}$ edge-disjoint paths from S to D'; we color the edges in these paths *green*. We can also find $M_{\{2\}}$ edge-disjoint paths from S to D_2; we color the edges in these paths *red*. Notice that, at the end of this process, some of the edges can have both *green* and *red* colors. We also highlight that:

- Out of the $M_{\{1,2\}}$ *green* paths from S to D', $M^\star_{\{1\}}+M^\star_{\{1,2\}}$ paths flow through D_1 and $M^\star_{\{2\}}$ flow through D_2.
- If a path is exclusively *green*, it flows through D_1 since otherwise, in addition to the $M_{\{2\}}$ *red* edge-disjoint paths from S to D_2, we would have also this path and thereby violate the min-cut capacity constraint to D_2.

The second observation above implies that, if there are $M^\star_{\{1\}}$ exclusively *green* paths, then we can separate the graph $\mathcal{G}'$ into two graphs: $\mathcal{G}'_a$ that contains all these $M^\star_{\{1\}}$ exclusively *green* paths and $\mathcal{G}'_b$ that contains all the edges of $\mathcal{G}'$ that are not in $\mathcal{G}'_a$. Given this, by simply removing the node D' and its incoming edges, we get $\mathcal{G}_a$ and $\mathcal{G}_b$. We now show how we can obtain these $M^\star_{\{1\}}$ exclusively *green* paths. Towards this end, we denote with $\mathcal{P}$ the set of all *green* paths from S to D' (notice that these paths might have also some *red* edges). Then, until there exists a path $p \in \mathcal{P}$ such that either it is not exclusively *green* or it does not start with an edge that is both *red* and *green*, we apply the two following steps:

1. Let e be the first edge in p, which is both *green* and *red* and denote with g the *red* path from S to D_2 that contains the edge e. Recall that, since the $M_{\{2\}}$ *red* paths are edge-disjoint, there is only one *red* path g passing through e. We split the path p into two parts as $p_1 - e - p_2$ and similarly we split the path g into $g_1 - e - g_2$.
2. We add the *red* color to p_1 (that before was all *green*) and we remove the *red* color from g_1, i.e., now each edge in g_1 is either *green* or it does not have any color. Note that in this way we replace the *red* path $g_1 - e - g_2$ with $p_1 - e - g_2$ from source S to D_2, which is also disjoint from the rest of $M_{\{2\}} - 1$ *red* paths.

We note that this process will stop only when all the $M_{\{1,2\}}$ paths from S to D' are either exclusively green or start with an edge that is both *red* and *green*. We also note that, since we did not remove any edge, clearly we also did not change any min-cut capacity during this process. Since initially there were $M_{\{2\}}$ *red* edges coming out of S and, in the process of the algorithm, we replaced one *red* by another *red*, then the number of *red* edges outgoing from S still remains the same. Thus, among the $M_{\{1,2\}}$ paths from S to D', only at most $M_{\{2\}}$ paths start with an edge that is both *green* and *red* and therefore, by using (3), at least $M^\star_{\{1\}}$ are exclusively *green* paths. This proves that the original graph $\mathcal{G}$ can be separated into the two graphs $\mathcal{G}_a$ and $\mathcal{G}_b$. By using similar arguments, one can then show that the graph $\mathcal{G}_b$ can be separated into the two graphs $\mathcal{G}_c$ and $\mathcal{G}_d$. This concludes the proof of Lemma 1.

Appendix B

We here give the proof of Lemma 2 (originally proved in [7, Theorem 9]). In particular, we first prove the converse (i.e., the rate region in (9) is an outer bound) and then the direct part (i.e., the rate region in (9) is achievable) of Lemma 2.

Outer Bound: Let $\mathcal{E}_{\mathcal{A}} \subseteq \mathcal{E}$ be a min-cut between the source S and $D_{\mathcal{A}}$ and define $\mathcal{I}(D_{\mathcal{A}}) := \bigcup_{i \in \mathcal{A}} \mathcal{I}(D_i)$. Then, for any $\mathcal{A} \subseteq [1:m]$ we have,

$$\begin{aligned} nR_{\mathcal{A}} = H(W_{\mathcal{A}}) &\overset{\text{(a)}}{=} H(W_{\mathcal{A}}) - H(W_{\mathcal{A}}|X^n_{\mathcal{I}(D_{\mathcal{A}})}) \\ &\overset{\text{(b)}}{=} H(W^n_{\mathcal{A}}) - H(W^n_{\mathcal{A}}|X^n_{\mathcal{E}_{\mathcal{A}}}) \\ &= I(W^n_{\mathcal{A}}; X^n_{\mathcal{E}_{\mathcal{A}}}) \\ &\overset{\text{(c)}}{\leq} H(X^n_{\mathcal{E}_{\mathcal{A}}}) \\ &\overset{\text{(d)}}{\leq} nM_{\mathcal{A}} \ , \end{aligned}$$

where $W_{\mathcal{A}} = \{W_i, i \in \mathcal{A}\}$ and where: (i) the equality in (a) follows because of the decodability constraint; (ii) the equality in (b) follows because $X^n_{\mathcal{I}(D_{\mathcal{A}})}$ is a deterministic function of $X^n_{\mathcal{E}_{\mathcal{A}}}$; (iii) the inequality in (c) follows since the entropy of a discrete random variable is a non-negative quantity; (iv) finally, the inequality in (d) follows since each link has unit capacity and since $|\mathcal{E}_{\mathcal{A}}| = M_{\mathcal{A}}$. By dividing both sides of the above inequality by n we obtain that $R_{\mathcal{A}}$ in (9) is an outer bound on the unsecure capacity region of the multiple unicast network with single source and m destinations.

Achievability: Assume that a rate m-tuple $(R_1, R_2, \ldots, R_m)$ satisfies the constraint in (9). We now prove that this m-tuple is achievable. Towards this end, from the original graph $\mathcal{G}$, we create a new directed acyclic graph $\mathcal{G}'$ where a new node D' is connected to each $D_i, i \in [1:m]$, through an edge $\mathcal{E}'_i$ of capacity R_i. It is not difficult to see that in $\mathcal{G}'$, the min-cut capacity between S and D' is $\sum\limits_{i=1}^{m} R_i$. This can be explained as follows. Suppose that the min-cut from S to D', in addition to a subset of $\mathcal{E}$ (i.e., the set of edges in the original $\mathcal{G}$), also contains some edges $\mathcal{E}'_{\mathcal{J}}$, with $\mathcal{J} \subseteq [1:m]$. This clearly implies that the subset of edges from $\mathcal{E}$ should form a cut between source S and $D_{[1:m]\setminus\mathcal{J}}$, otherwise we would not have a cut between S and D'. Thus, the min-cut has a capacity of at least $\sum\limits_{i \in J} R_i + M_{\left\{D_{[1:m]\setminus\mathcal{J}}\right\}}$ and, since $\sum\limits_{i \in [1:m]\setminus\mathcal{J}} R_i \leq M_{\left\{D_{[1:m]\setminus\mathcal{J}}\right\}}$ (this follows from the outer bound proved above), the min-cut has a capacity of at least $\sum\limits_{i}^{m} R_i$. Then, since the set $\mathcal{E}'_{[1:m]}$ is a cut of capacity $\sum\limits_{i}^{m} R_i$, it follows that the min-cut has a capacity of at most $\sum\limits_{i}^{m} R_i$. This implies that the min-cut capacity between S and D' in $\mathcal{G}'$ is $\sum\limits_{i=1}^{m} R_i$. With this, the achievability of the rate m-tuple $(R_1, R_2, \ldots, R_m)$ that satisfies the constraint in (9) directly follows from the max-flow min-cut theorem. Indeed, since one can communicate a total information of $\sum_{i}^{m} R_i$ from S to D' in $\mathcal{G}'$, then this is possible only if an amount R_i of information flows through $D_i, i \in [1 : m]$, in $\mathcal{G}$. This concludes the proof of Lemma 2. Notice that in order to transmit $\sum_{i=1}^{m} R_i$ message packets from S to D' (single unicast session) network coding is not needed. Thus, there is no need of coding operations to characterize the capacity region of a network with single source and multiple destinations.

References

1. Cai, N., Yeung, R.W.: Secure network coding. In: Proceedings IEEE International Symposium on Information Theory (ISIT),. (July 2002) 323–
2. Koetter, R., Effros, M., Ho, T.: Network codes as codes on graphs. In: Conference on Information Sciences and Systems (CISS). (2004)
3. Riis, S.: Reversible and irreversible information networks. IEEE Transactions on Information Theory **53**(11) (November 2007) 4339–4349
4. Ahlswede, R., Cai, N., Li, S.Y.R., Yeung, R.W.: Network information flow. IEEE Transactions on Information Theory **46**(4) (Jul 2000) 1204–1216
5. Li, S.Y.R., Yeung, R.W., Cai, N.: Linear network coding. IEEE Transactions on Information Theory **49**(2) (February 2003) 371–381
6. Jaggi, S., Sanders, P., Chou, P.A., Effros, M., Egner, S., Jain, K., Tolhuizen, L.M.G.M.: Polynomial time algorithms for multicast network code construction. IEEE Transactions on Information Theory **51**(6) (June 2005) 1973–1982
7. Koetter, R., Medard, M.: An algebraic approach to network coding. IEEE/ACM Transactions on Networking **11**(5) (October 2003) 782–795
8. Kamath, S.U., Tse, D.N.C., Anantharam, V.: Generalized network sharing outer bound and the two-unicast problem. In: International Symposium on Networking Coding (NetCod). (July 2011) 1–6
9. Kamath, S., Tse, D.N.C., Wang, C.C.: Two-unicast is hard. In: IEEE International Symposium on Information Theory (ISIT). (June 2014) 2147–2151
10. Ramamoorthy, A., Wesel, R.D.: The single source two terminal network with network coding. arXiv:0908.2847 (August 2009)
11. Cui, T., Ho, T., Kliewer, J.: On secure network coding with nonuniform or restricted wiretap sets. IEEE Transactions on Information Theory **59**(1) (January 2013) 166–176
12. Agarwal, G.K., Cardone, M., Fragouli, C.: On secure network coding for two unicast sessions: Studying butterflies. In: IEEE Globecom Workshops (GC Wkshps). (December 2016) 1–6
13. Agarwal, G.K., Cardone, M., Fragouli, C.: Coding across unicast sessions can increase the secure message capacity. In: IEEE International Symposium on Information Theory (ISIT). (July 2016) 2134–2138