

Computing Bounds on Network Capacity Regions as a Polytope Reconstruction Problem

Anthony Kim

Oracle Corporation

500 Oracle Parkway, Redwood Shores, CA 94065

Email: tonyekim@yahoo.com

Muriel Médard

Research Laboratory of Electronics

MIT, Cambridge, MA 02139

Email: medard@mit.edu

Abstract—We define a notion of network capacity region of networks that generalizes the notion of network capacity defined by Cannons et al. and prove its notable properties such as closedness, boundedness and convexity when the finite field is fixed. We show that the network routing capacity region is a computable rational polytope and provide exact algorithms and approximation heuristics for computing the region. We define the semi-network linear coding capacity region, with respect to a fixed finite field, that inner bounds the corresponding network linear coding capacity region, show that it is a computable rational polytope, and provide exact algorithms and approximation heuristics. We show connections between computing these regions and a polytope reconstruction problem and some combinatorial optimization problems, such as the minimum cost directed Steiner tree problem. We provide an example to illustrate our results. The algorithms are not necessarily polynomial-time.

I. INTRODUCTION

In a seminal work in 2000, Ahlswede et al. [1] introduced the network coding model to the problem of communicating information in networks. They showed that the extended capabilities of intermediate nodes to code on incoming packets give greater information throughput than in the traditional routing model and that the capacity of any multiple multicast network with a single source node is equal to the minimum of min-cuts between the source and receiver nodes.

A few explicit outer bounds of capacity regions of networks exist; the max-flow/min-cut bounds is one such set of bounds, which were sufficient in the case of single-source multiple multicast networks. Harvey et al.[2] combined information theoretic and graph theoretic techniques to provide a computable outer bound on the network coding capacity regions. Yan et al.[3] gave an explicit outer bound that improves upon the max-flow/min-cut bounds and showed its connection to a minimum cost network coding problem. They used their results to compute the capacity region of a special class of 3-layer networks. Thakor et al.[4] gave a new computable outer bound, based on characterizations of all functional dependencies in networks. Yan et al.[5] provided an exact characterization of the capacity region for general multi-source multi-sink networks by bounding the constrained region in the entropy space. However, they noted that explicitly evaluating the obtained capacity regions remains difficult in general. In a related work, Chan and Grant[6] showed that even the explicit characterization of capacity region for single-source networks

can be difficult since the computation of a capacity region involves the non-polyhedral set of entropy functions and that linear programming bounds do not suffice.

The routing capacity region of networks is better understood via linear programming approaches. Cannons et al.[7] defined a notion of network routing capacity that is computable with a linear program and showed that every rational number in $(0, 1]$ is the routing capacity of some solvable network. Yazdi et al.[8] and [9] extended a special case of Farkas Lemma called the “Japanese Theorem” to reduce an infinite set of linear constraints to a finite set in terms of minimal Steiner trees and applied the results to obtain the routing capacity region of undirected ring networks. In a subsequent work, Kakhbod and Yazdi[10] provided complexity results on the description size of the inequalities obtained and applied them to the undirected ring networks.

In this paper, we study the network capacity region of networks along the lines of work by Cannons et al. [7]. We define the network capacity region of networks analogously to the rate regions in information theory and show its notable properties when the finite field is fixed: closedness, boundedness and convexity. In the case of routing, we prove that the network routing capacity region is a computable rational polytope and provide exact algorithms and approximation heuristics for computing the region. In the case of linear network coding, we define an auxiliary region, called the semi-network linear coding capacity region, which is a computable rational polytope that inner bounds the network linear coding capacity region, and provide exact algorithms and approximation heuristics, with respect to a fixed finite field. The main idea is to reduce computation of the capacity regions to a polytope reconstruction problem and use linear programming techniques on associated combinatorial optimization problems. Our results generalize to directed networks with cycles and undirected networks. This present work partially addresses two problems proposed by Cannons et al. [7]: whether there exists efficient algorithms for computing their notions of network routing capacity and network linear coding capacity. It follows from our work that there exist combinatorial approximation algorithms, not polynomial-time, for computing the network routing capacity and for computing a lower bound of the network linear coding capacity. In addition, we provide an example to illustrate our results.

We note that algorithms and heuristics we present are not polynomial-time schemes. We do not distinguish the vertex and hyperplane descriptions of polytopes, but note that converting one description into another can be computationally expensive.

The paper is organized as follows. In Section II, we give a fractional network coding model. In Section III, we define the network coding capacity region and prove its notable properties. In Section IV and V, we prove notable properties of network coding capacity regions, or inner bounds thereof, in the case of routing and linear coding and discuss algorithmic aspects of the regions. In Section VI, we give an example that illustrates our results. Finally, in Section VII, we discuss further extensions and conclude. Due to the page limit, we omit some details and refer to [11] for more detailed exposition.

II. FRACTIONAL NETWORK CODING MODEL

We define a fractional network coding model. Most definitions are adapted from Cannons et al. [7]. We write vectors or points in a multi-dimensional space with a hat as in \hat{k} and let \hat{k}_i (also $\hat{k}(i)$) denote the i -th coordinate of the vector. We shall omit the hat when convenient.

A *capacitated network* \mathcal{N} is a finite, directed, acyclic multigraph given by a 7-tuple $(\nu, \epsilon, \mu, c, \mathcal{A}, S, R)$ where ν is a node set, ϵ is an edge set, μ is a message set, $c : \epsilon \rightarrow \mathbb{Z}^+$ is an edge capacity function, \mathcal{A} is an alphabet, $S : \nu \rightarrow 2^\mu$ is a source mapping, and $R : \nu \rightarrow 2^\mu$ is a receiver mapping.

We define fractional edge function, fractional decoding function, and fractional message assignment with respect to a finite field F , where $|F| \geq |\mathcal{A}|$, a *source dimension vector* \hat{k} , and an *edge dimension* n . Let $\mathcal{N} = (\nu, \epsilon, \mu, c, \mathcal{A}, S, R)$ be a capacitated network and $m_1, \dots, m_{|\mu|}$ be the messages. Let $\hat{k} = (k_1, \dots, k_{|\mu|})$ be a vector of positive integers and n be a positive integer. For each edge $e = (x, y)$, a *fractional edge function* is a map $f_e : (F^{k_{i_1}}) \times \dots \times (F^{k_{i_\alpha}}) \times (F^{nc(e_{\alpha+1})}) \times \dots \times (F^{nc(e_{\alpha+\beta})}) \rightarrow F^{nc(e)}$. For each node $x \in \nu$ and message $m_j \in R(x)$, a *fractional decoding function* is a map $f_{x, m_j} : (F^{k_{i_1}}) \times \dots \times (F^{k_{i_\alpha}}) \times (F^{nc(e_{\alpha+1})}) \times \dots \times (F^{nc(e_{\alpha+\beta})}) \rightarrow F^{k_j}$, where $m_{i_1}, \dots, m_{i_\alpha}$ are α messages generated by x and $e_{\alpha+1}, \dots, e_{\alpha+\beta}$ are β in-edges of x . We denote the collections of fractional edge and fractional decoding functions by $\mathcal{F}_e = \{f_e : e \in \epsilon\}$ and $\mathcal{F}_{x, m} = \{f_{x, m} : x \in \nu, m \in R(x)\}$. Note k_i is the source dimension for message m_i .

A *fractional message assignment* is a collection of maps $\bar{a} = (a_1, \dots, a_{|\mu|})$ where a_i is a message assignment for m_i , $a_i : m_i \rightarrow F^{k_i}$. A *fractional network code* on \mathcal{N} is a 5-tuple $(F, \hat{k}, n, \mathcal{F}_e, \mathcal{F}_d)$ where F is a finite field, with $|F| \geq |\mathcal{A}|$, \hat{k} is a source dimension vector, n is an edge dimension, \mathcal{F}_e is a collection of fractional edge functions, and \mathcal{F}_d is a collection of fractional decoding functions. We have different classes of fractional network codes corresponding to \mathcal{F}_e and \mathcal{F}_d being routing, linear, and nonlinear functions over the field F : *fractional routing/linear/nonlinear network codes*. A fractional network code is a *fractional network code solution* if for every fractional message assignment \bar{a} , we have $f_{x, m_j}(a_{i_1}(m_{i_1}), \dots, a_{i_\alpha}(m_{i_\alpha}), s(e_{\alpha+1}), \dots, s(e_{\alpha+\beta})) =$

$a_j(m_j)$, for all $x \in \nu$ and $m_j \in R(x)$, where $m_{i_1}, \dots, m_{i_\alpha}$ are α messages generated by x and $e_{\alpha+1}, \dots, e_{\alpha+\beta}$ are β in-edges of x . If the above equation holds for a particular $x \in \nu$ and message $m \in R(x)$, then we say *node x 's demand m is satisfied*. We have classes of fractional network code solutions: *fractional routing/linear/nonlinear network code solutions*.

If $(F, \hat{k}, n, \mathcal{F}_e, \mathcal{F}_d)$ is a fractional network code solution for \mathcal{N} , source node $x \in \nu$ sends a vector of k_i symbols from F , representing alphabets in \mathcal{A} , for each message $m_i \in S(x)$; each receiver node $x \in \nu$ demands the original vector of k_i symbols corresponding to m_i for each $m_i \in R(x)$; and each edge e carries a vector of $c(e)n$ symbols. We refer to coordinates of the symbol vector corresponding to m_i as *message m_i 's coordinates* and coordinates of the symbol vector on edge e as *edge e 's coordinates*. Note that a coordinate of edge e can be *active* or *inactive*, depending on whether it actively carries a symbol in the fractional network code solution or not. A fractional network code solution is *minimal* if the set A of all active coordinates of edges is minimal, i.e., there exists no fractional network code solution with the same F , \hat{k} and n and the set of active coordinates that is a strict subset of A .

In this paper, we consider non-degenerate networks; for each demand of a message at a node, there is a path from a source node generating the message to the receiver node and no message is both generated and demanded by the same node. We shall abridge network code descriptions and solution concepts when convenient. In this work, a multicast network has exactly one message produced by a source node and demanded by more than one receiver nodes and multi-source multi-sink networks are the most general network instances.

III. NETWORK CAPACITY REGIONS

A vector of nonnegative numbers $(\frac{k_1}{n}, \dots, \frac{k_{|\mu|}}{n}) \in \mathbb{Q}_+^{|\mu|}$ is an *achievable coding rate vector* for \mathcal{N} if there exists a fractional network code solution $(F, \hat{k}, n, \mathcal{F}_e, \mathcal{F}_d)$ where $\hat{k} = (k_1, \dots, k_{|\mu|})$. The *network capacity region* of \mathcal{N} is the closure of all achievable coding rate vectors in $\mathbb{R}^{|\mu|}$. There are different classes of network capacity regions: the *network routing capacity region*, \mathcal{C}_r , which is the closure of all achievable routing rate vectors; the *network linear coding capacity region*, \mathcal{C}_l , which is the closure of all achievable linear coding rate vectors; and the *network coding capacity region*, \mathcal{C} , which is the closure of all achievable coding rate vectors.

Theorem 1: Let \mathcal{N} be a capacitated network and F a fixed finite field. The corresponding network capacity region \mathcal{C} is a closed, bounded and convex set in $\mathbb{R}_+^{|\mu|}$.

Proof: (Closedness) By definition, \mathcal{C} is a closed set. (Boundedness) By symmetry, it suffices to show that $\frac{k_1}{n}$ is bounded in any achievable coding rate vector. Let n be the edge dimension, ν_1 be the set of source nodes in ν that generate message m_1 and γ be the sum of capacities of out-edges of nodes in ν_1 . Then, $k_1 \leq \gamma n$ as we cannot send more than γn independent coordinates of message m_1 and expect receivers to recover all the information. Hence, $\frac{k_1}{n} \leq \gamma$. (Convexity) Let $x_0, x_1 \in \mathcal{C}$ and $\lambda \in [0, 1]$. It is straightforward to show that $x = (1 - \lambda)x_0 + \lambda x_1 \in \mathcal{C}$ by rate sharing; we

use sequences of achievable coding rate vectors converging to x_0 and x_1 to produce a sequence of achievable coding rate vectors converging to x . ■

Similarly, the routing capacity region \mathcal{C}_r and linear coding capacity region \mathcal{C}_l are closed, bounded and convex regions when the finite field F is fixed. However, the rate-sharing argument seems to break down in the case of linear coding if the finite field is not fixed.

IV. NETWORK ROUTING CAPACITY REGIONS

We show the network routing capacity region \mathcal{C}_r is the image of a higher-dimensional rational polytope under an affine map and consider the computation of \mathcal{C}_r as the polytope reconstruction problem with a ray oracle. Since multi-source multi-sink networks can be reduced to multiple multicast networks, it suffices to show the results with respect to the multiple multicast networks.

A. Properties

Theorem 2: The network routing capacity region \mathcal{C}_r is a bounded rational polytope in $\mathbb{R}_+^{|\mu|}$ and is computable.

Proof: (Polytope) It suffices to consider minimal fractional routing solutions, which consist of routing messages along Steiner trees, since any fractional routing solution can be reduced to a minimal one. Let \mathcal{T}_i be the finite set of all Steiner trees rooted at the source node of message m_i and spanning all receiver nodes that demand m_i , and \mathcal{T} be the union, $\mathcal{T} = \mathcal{T}_1 \cup \dots \cup \mathcal{T}_{|\mu|}$. Then, any minimal fractional routing solution satisfies the following constraints:

$$\begin{aligned} \sum_{T \in \mathcal{T}} T(e) \cdot x(T) &\leq c(e)n, \quad \forall e \in \epsilon \\ \sum_{T \in \mathcal{T}_i} x(T) &= k_i, \quad \forall 1 \leq i \leq |\mu| \\ x &\geq 0, \end{aligned}$$

where $x(T)$ is the number of times Steiner tree T is used in the solution and $T(e)$ is an indicator that is 1 if T uses edge e , or 0 otherwise. After scaling $x(T)$ by n , any minimal fractional routing solution satisfies

$$\begin{aligned} \sum_{T \in \mathcal{T}} T(e) \cdot x(T) &\leq c(e), \quad \forall e \in \epsilon \\ x &\geq 0. \end{aligned}$$

As the coefficients are in \mathbb{Q} , the above set of inequalities defines a bounded rational polytope \mathcal{P}_r , with rational extreme points. The polytope is bounded, because edge capacities are finite and no Steiner tree can be used to route for infinitely many times. Each minimal fractional routing solution reduces to a rational point inside the polytope \mathcal{P}_r , and each rational point x inside \mathcal{P}_r has a minimal fractional routing solution $(F, \hat{k}, n, \mathcal{F}_e, \mathcal{F}_d)$ that reduces to it, such that $\frac{k_i}{n} = \sum_{T \in \mathcal{T}_i} x(T)$ for all i . It follows that the network routing capacity region \mathcal{C}_r is the image of \mathcal{P}_r under the affine map $\psi_r : (x(T))_{T \in \mathcal{T}} \mapsto \left(\sum_{T \in \mathcal{T}_1} x(T), \dots, \sum_{T \in \mathcal{T}_{|\mu|}} x(T) \right)$. It follows that \mathcal{C}_r is a bounded rational polytope.

(Computability) We first compute the vertices v_1, \dots, v_h of polytope \mathcal{P}_r by any vertex enumeration algorithm. Then we compute the images of the vertices of \mathcal{P}_r under the affine map ψ_r . The network routing capacity region is given by the

vertices of the convex hull of points $\psi_r(v_1), \dots, \psi_r(v_h)$ in $\mathbb{R}_+^{|\mu|}$. ■

B. Algorithms

We now provide exact algorithms and approximation heuristics for computing the network routing capacity region \mathcal{C}_r . We already provided an exact algorithm in the proof of Theorem 2. Since the polytope \mathcal{P}_r is defined in a high-dimensional space, the exact algorithm may not be efficient in practice. Using results in [12] and [13] on polytope reconstruction problems, we derive different kinds of algorithms that might be more efficient.

In our polytope reconstruction problem, we compute the facet description of a given polytope by making calls to the ray oracle \mathcal{O}_{Ray} which, given a ray, returns the intersection point on the polytope surface and the ray. We reduce the computation of \mathcal{C}_r to the polytope reconstruction problem by 1) reflecting \mathcal{C}_r around the origin to get a symmetric polytope \mathcal{Q} that contains the origin in its interior and 2) solving the linear programs similar to the one in Cannons et al. [7] to implement the ray oracle \mathcal{O}_{Ray} . To reflect \mathcal{C}_r , we map all calls to the ray oracle to equivalent calls with rays defined in $\mathbb{R}_+^{|\mu|}$. We use the algorithm outlined in Section 5 of Gritzmann et al. [13] to compute all the facets of the resulting polytope \mathcal{Q} and, therefore, \mathcal{C}_r . The main idea of the algorithm is to first find a polytope \mathcal{Q}' that contains \mathcal{Q} and whose facet-defining hyperplanes are a subset of those for \mathcal{Q} (Theorem 5.3 in [13]), and then successively add more facet-defining hyperplanes of \mathcal{Q} to \mathcal{Q}' by using \mathcal{O}_{Ray} . By Theorem 5.5 in Gritzmann et al. [13] and the symmetries around the origin, we need at most $f_0(\mathcal{C}_r) + (|\mu| - 1)f_{|\mu|-1}^2(\mathcal{C}_r) + (5|\mu| - 4)f_{|\mu|-1}(\mathcal{C}_r)$ calls to the ray oracle where $f_i(\mathcal{C}_r)$ denotes the number of i -dimensional faces of \mathcal{C}_r that do not contain the origin (the 0-th dimensional faces being the points).

If we use an exact algorithm for the ray oracle \mathcal{O}_{Ray} , we get an exact hyperplane description of the network routing capacity region. If instead we use an approximation algorithm that computes some point r such that the actual intersection point lies between r and Ar , then we obtain an approximation heuristic that computes a set of points r such that the boundary of \mathcal{C}_r lies between points r and Ar . We note that an approximation algorithm for \mathcal{O}_{Ray} does not necessary yield an approximation algorithm for \mathcal{C}_r , where an A -approximation of \mathcal{C}_r would be a polytope \mathcal{P} such that $\mathcal{P} \subset \mathcal{C}_r \subset A\mathcal{P}$. While an approximation algorithm for the oracle \mathcal{O}_{Ray} does not necessarily lead to a polytope description of \mathcal{C}_r , it might be faster and more efficient than exact algorithms and, hence, more applicable to compute a quick “sketch” of \mathcal{C}_r . For instance, we can find approximate intersection points on a sufficiently large number of rays evenly spread apart.

C. Implementations of Oracle \mathcal{O}_{Ray}

Given the hyperplane description of the polytope \mathcal{P}_r and a ray with a rational slope, $\hat{x} = \hat{q}t, t \geq 0$, we want to compute the rational intersection point of the ray and the boundary of

\mathcal{C}_r . Note that the intersection point is exactly $\lambda_{max}\hat{q}$ where λ_{max} is the optimal value to the linear program

$$\begin{aligned} \max \quad & \lambda \\ \text{s. t.} \quad & \sum_{T \in \mathcal{T}} T(e) \cdot x(T) \leq c(e), \quad \forall e \in \epsilon \\ & \sum_{T \in \mathcal{T}_i} x(T) \geq \lambda q_i, \quad \forall i \\ & x, \lambda \geq 0. \end{aligned} \quad (\text{IV.1})$$

We can use any linear programming algorithm, such as the ellipsoid algorithm, to solve the linear program exactly and, thus, obtain an exact oracle \mathcal{O}_{Ray} . For the approximate ray oracles, we can employ algorithms that solve the linear program within a provable approximation guarantee. Note that as the network routing capacity region \mathcal{C}_r is a rational polytope, it suffices to consider rays with a rational slope.

Using the results for the multicommodity flow and related problems by Garg and Könemann [14], we can derive a combinatorial approximation algorithm that solves the linear program (IV.1). It computes a point \hat{r} such that $\lambda_{max}\hat{q}$ is on the line segment between \hat{r} and $(1 + \omega)A\hat{r}$ for some numbers $\omega > 0$ and $A \geq 1$. The main idea is to view solving (IV.1) as concurrently packing Steiner trees according to the ratio defined by \hat{q} and use the results for the minimum cost directed Steiner tree problem: given an acyclic directed multigraph $\mathcal{G} = (\nu, \epsilon)$, a length function $l : \epsilon \rightarrow \mathbb{R}^+$, a source node s and receiver nodes n_1, \dots, n_k , find a minimum cost (defined by $\sum_{e \in \epsilon'} l(e)$) subset of edges ϵ' such that there is a directed path from s to each n_i in ϵ' . We assume we have an oracle $\mathcal{O}_{DSteiner}$ that solves the minimum cost directed Steiner tree problem, which is well-known to be NP-hard, within an approximation guarantee A . Then we have the following theorem:

Theorem 3: There exists an $(1 + \omega)A$ -approximation algorithm for the linear program (IV.1) in time $O(\omega^{-2} (|\mu| \log A |\mu| + |\epsilon|) A \log |\epsilon| \cdot T_{DSteiner})$, where $T_{DSteiner}$ is the time required to solve the minimum cost directed Steiner tree problem with oracle $\mathcal{O}_{DSteiner}$ within an approximation guarantee A .

Note that computing the network routing capacity region is not the same as the multicommodity flow, as in our problem, a message can be demanded by multiple nodes and can be duplicated at intermediate nodes. There are algorithms for $\mathcal{O}_{DSteiner}$; the approximation using shortest paths yields $A = O(|\nu|)$ in time $O((|\nu|^2 + |\nu||\epsilon|) \log |\nu|)$ and Charikar et al. [15] gives a set of algorithms with $A = i(i - 1)|\nu|^{1/i}$ and time $O(|\nu|^{3i})$ for any integer $i > 1$. We can also use any brute-force algorithm that yields $A = 1$. Note that further improvements in the minimum cost directed Steiner tree problem translates to improvements to our approximation algorithm by above theorem.

V. NETWORK LINEAR CODING CAPACITY REGIONS

We define a computable polytope \mathcal{C}'_l , which we call the *semi-network linear coding capacity region*, that is contained in the network linear coding capacity region \mathcal{C}_l . It is unknown how good of an approximation the polytope \mathcal{C}'_l is to \mathcal{C}_l . Unlike in the last section, the finite field is important in the computation of \mathcal{C}'_l because of linearity of functions and

termination of algorithms. In this section, we assume that a network \mathcal{N} is a multi-source multi-sink network and that the finite field F is fixed.

A. Definitions

The *weight vectors associated with \mathcal{N}* are vectors w in $\{0, 1\}^{|\mu|}$ such that there exists a scalar-linear network code solution when only messages m_i with $w_i = 1$ are considered and all the edges are assumed to be of unit capacity. We refer to the scalar-linear network code solutions corresponding to these weight vectors as *partial scalar-linear network code solutions*. A fractional network code is a *simple fractional linear network code solution* if the fractional network code is linear over the finite field F and can be written as an aggregate of partial scalar-linear solutions (when considered with unit edge capacities). We define the *semi-network linear coding capacity region* \mathcal{C}'_l of \mathcal{N} as the closure of all coding rate vectors achievable by simple fractional linear network code solutions. Clearly, the network linear coding capacity region \mathcal{C}_l contains the semi-network linear coding capacity region \mathcal{C}'_l . Note that scalar-linear network code solution and simple fractional linear network code solution are two different solution concepts.

B. Results

By the same line of reasoning as in the case of routing, we get the following theorems.

Theorem 4: Assume a finite field F is given. The semi-network linear coding capacity region \mathcal{C}'_l , with respect to F , is a bounded rational polytope in $\mathbb{R}_+^{|\mu|}$ and is computable.

Instead of the minimum cost directed Steiner tree problem, we have two associated problems: minimum cost scalar-linear network code problem where given a network with unit edge capacities, a finite field F , and a length function $l : \epsilon \rightarrow \mathbb{R}_+$, we want to compute the minimum cost (defined to be the sum of lengths of the edges used) scalar-linear network code solution; and fractional covering with box constraints problem which is to solve

$$\begin{aligned} \min \quad & \sum_{j=1}^m c(j)x(j) \\ \text{s. t.} \quad & \sum_j A(i, j)x(j) \geq b(i), \quad \forall 1 \leq i \leq n \\ & x(j) \leq u(j), \quad \forall 1 \leq j \leq m \\ & x \geq 0, \end{aligned}$$

where A is an $n \times m$ nonnegative integer matrix, b a non-negative vector, c a positive vector, and u a nonnegative integer vector. Assuming we have two oracles, $\mathcal{O}_{SLinear}$ and T_{FCover} , for these problems, we get:

Theorem 5: There exists an $(1 + \omega)B$ -approximation algorithm for oracle \mathcal{O}_{Ray} , for \mathcal{C}'_l , with time $O(\omega^{-2} (\log A |\mu| + |\epsilon|) B \log |\epsilon| \cdot (T_{FCover} + k' T_{SLinear}))$, where T_{FCover} is the time required to solve the fractional covering problem by \mathcal{O}_{FCover} within an approximation guarantee B , $T_{SLinear}$ is the time required to solve the minimum cost scalar-linear network code problem exactly by $\mathcal{O}_{SLinear}$, and k' is the total number of weight vectors associated with \mathcal{N} .




Fig. 1. Network \mathcal{N} and its network routing capacity region and semi-network linear capacity region, each computed with an exact \mathcal{O}_{Ray} and an approximate \mathcal{O}_{Ray} .

For \mathcal{O}_{FCover} , we can use any linear programming algorithm or the combinatorial approximation algorithm by Fleischer [16]. For $\mathcal{O}_{SLinear}$, there is a dynamic programming algorithm in [11]. Without the minimum cost condition, the scalar-linear network code problem reduces to the decidability problem of determining whether or not a network has a scalar-linear solution, which is NP-hard [17]. Without the fixed finite field F , the decidability of the problem is unknown; this justifies in part our assumption of the fixed finite field F .

VI. AN EXAMPLE

For network \mathcal{N} , we computed the network routing capacity region (two inner curves) and the semi-network linear coding capacity region (two outer curves) with both exact and approximate ray oracles in Figure 1. Network \mathcal{N} has two source nodes at the top and two receiver nodes at the bottom. For exact ray oracles, we hard-coded the corresponding linear programs and used a linear program solver, linprog, in MATLAB. To obtain approximate intersections points, we simply used an approximate oracle \mathcal{O}_{Ray} in place of the linear program solver in an implementation of a 2D-polytope reconstruction algorithm. We note that the approximate oracle \mathcal{O}_{Ray} worked well in this example and led to its successful termination, but this may not hold for arbitrary networks in general.

VII. FURTHER DISCUSSION

Our results have a few extensions: We can design membership testing algorithms that given a rate vector, determines whether or not there exists a fractional network code solution that achieves the rates from algorithms we have provided, for the network routing capacity region and semi-network linear coding capacity region. We can also generalize the results to directed networks with cycles and undirected networks straightforwardly.

Note that the network routing capacity defined by Cannons et al. [7] corresponds to a point on the boundary of polytope \mathcal{C}_r ; it is exactly the intersection point between the (outer) boundary of \mathcal{C}_r and the ray $\hat{x} = (1, \dots, 1)t, t \geq 0$. Hence, our work partially addresses two problems proposed by Cannons et al.: whether there exists efficient algorithms for computing their

notions of network routing capacity and network linear coding capacity. It follows from our work that there exist combinatorial approximation algorithms, albeit not polynomial-time, for computing the network routing capacity and for computing a lower bound of the network linear coding capacity.

For details omitted in this paper, we refer to [11].

ACKNOWLEDGMENT

The authors would like to thank Michel Goemans and the anonymous reviewers for helpful comments. This material is based upon work supported by the Air Force Office of Scientific Research (AFOSR) under award No. 016974-002.

REFERENCES

- [1] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung, "Network information flow," *Information Theory, IEEE Transactions on*, vol. 46, no. 4, pp. 1204–1216, Jul 2000.
- [2] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, pp. 2345–2364, 2006.
- [3] X. Yan, J. Yang, and Z. Zhang, "An outer bound for multisource multisink network coding with minimum cost consideration," *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, pp. 2373–2385, 2006.
- [4] S. Thakor, A. Grant, and T. Chan, "Network coding capacity: A functional dependence bound," in *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, June 2009, pp. 263–267.
- [5] X. Yan, R. W. Yeung, and Z. Zhang, "The capacity region for multi-source multi-sink network coding," in *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, 24–29 2007, pp. 116–120.
- [6] T. Chan and A. Grant, "Mission impossible: Computing the network coding capacity region," in *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, 6–11 2008, pp. 320–324.
- [7] J. Cannons, R. Dougherty, C. Freiling, and K. Zeger, "Network routing capacity," in *Information Theory, 2005. ISIT 2005. Proceedings. International Symposium on*, 4–9 2005, pp. 11–13.
- [8] S. M. S. T. Yazdi, S. A. Savari, F. Farnoud, and G. Kramer, "A multi-message capacity region for undirected ring networks," in *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, 24–29 2007, pp. 1091–1095.
- [9] S. Yazdi, S. Savari, K. Carlson, and G. Kramer, "The capacity region of a collection of multicast sessions in an undirected ring network," in *Parallel Processing Workshops, 2007. ICPPW 2007. International Conference on*, 10–14 2007, pp. 40–40.
- [10] A. Kakhbod and S. M. S. T. Yazdi, "On describing the routing capacity regions of networks," *Mathematical Methods of Operations Research*, pp. 1432–2994, May 2010.
- [11] A. Kim, "On network coding capacity - matroidal networks and network capacity regions," Master's thesis, MIT, Cambridge, MA 02139, Sept. 2010, available at <http://www.mit.edu/~tonyekim/>. Also at arXiv.
- [12] R. Cole and C. K. Yap, "Shape from probing," *J. Algorithms*, vol. 8, no. 1, pp. 19–38, 1987.
- [13] P. Grötszmann, V. Klee, and J. Westwater, "Polytope Containment and Determination by Linear Probes," *Proc. London Math. Soc.*, vol. s3-70, no. 3, pp. 691–720, 1995.
- [14] N. Garg and J. Könemann, "Faster and simpler algorithms for multi-commodity flow and other fractional packing problems," *SIAM Journal on Computing*, vol. 37, no. 2, pp. 630–652, 2007.
- [15] M. Charikar, C. Chekuri, T.-y. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li, "Approximation algorithms for directed steiner problems," in *SODA '98: Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1998, pp. 192–200.
- [16] L. Fleischer, "A fast approximation scheme for fractional covering problems with variable upper bounds," in *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2004, pp. 1001–1010.
- [17] A. R. Lehman and E. Lehman, "Complexity classification of network information flow problems," in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, Jan. 2004, pp. 142–150.