

Efficient Subnetwork Selection in Relay Networks

Siddhartha Brahma^{*1}, Ayan Sengupta^{*1}, Christina Fragouli²

¹EPFL, Switzerland, ²UCLA, USA

^{*}Co-First Authors

Abstract—We consider a source that would like to communicate with a destination over a layered Gaussian relay network. We present a *computationally efficient* method that enables to select a near-optimal (in terms of throughput) subnetwork of a *given size* connecting the source with the destination. Our method starts by formulating an integer optimization problem that maximizes the rates that the Quantize-Map-and-Forward relaying protocol can achieve over a selected subnetwork; we then relax the integer constraints to obtain a non-linear optimization over reals. For diamond networks, we prove that this optimization over reals is concave, while for general layered networks we give empirical demonstrations of near-concavity, paving the way for efficient algorithms to solve the relaxed problem. We then round the relaxed solution to select a specific subnetwork. Simulations using off-the-shelf non-linear optimization algorithms demonstrate excellent performance with respect to the true integer optimum for both diamond networks as well as multi-layered networks. Even with these non-customized algorithms, significant time savings are observed *vis-à-vis* exhaustive integer optimization¹.

I. INTRODUCTION

Consider a source that would like to communicate with a destination over a (potentially large) wireless Gaussian relay network. Using all the network relays for this purpose can be wasteful in terms of resources; instead, we would like to only use a certain number of relays, i.e., a subnetwork of a fixed size. To this end, this paper addresses the following question: How do we select the *best* (in terms of throughput) subnetwork of a *given size* in a *computationally efficient* manner?

In a sense, we are interested in generalizing routing over physical layer cooperation networks: In routing, we select the best one (or k) path(s) over which to forward the information from a source to a destination; over physical layer cooperation networks, the corresponding operation would be selecting the best subnetwork (of a given size). This is a conceptually simple, but computationally a hard optimization problem.

Our objective function should facilitate the selection of a subnetwork that enables the highest throughput. This would be the subnetwork that has the largest information-theoretic capacity; however, determining the capacity explicitly remains an unsolved problem till date. We resort instead to *approximate* schemes, such as Quantize-Map-and-Forward (QMF) [2] and Noisy Network Coding (NNC) [7], that are proved to achieve all rates within a constant additive gap (independent of the channels in the network) from the cutset upper bound on capacity for *arbitrary* single-source single-destination wireless relay networks. We are encouraged in

this choice by the fact that initial concerns about the additive gap affecting QMF/NNC performance at moderate SNRs have also been addressed in [11], [4], with appropriate choices of quantizer distortions demonstrating excellent moderate SNR performance. We thus focus our attention on QMF/NNC-based relay networks and hence, use the *achievable rate* expressions for these schemes [7], or indeed their close siblings—the cutset upper bound expressions on network capacity (which differ from the QMF/NNC rates by a mere constant), as the network objective functions for our problem.

At its core, subnetwork selection is an integer optimization problem: To every relay \mathcal{R} in a network \mathcal{W} having \mathcal{N} relays, we assign an *integer* selection variable $\theta_{\mathcal{R}}$ that takes value 1 if \mathcal{R} is selected and 0 if \mathcal{R} is not selected. In the corresponding objective function, the incoming and outgoing channels to and from the relay \mathcal{R} are weighted by $\theta_{\mathcal{R}}$. An optimal subnetwork thus corresponds to the integer tuple $\{\theta_{\mathcal{R}}\}_{\mathcal{R} \in \mathcal{W}}$ of size \mathcal{N} that maximizes the weighted objective function subject to a subnetwork size constraint of the form $\sum_{\mathcal{R} \in \mathcal{W}} \theta_{\mathcal{R}} = K$. A trivial solution to the above integer optimization problem requires an exhaustive search over the possibly *exponentially many* tuples that satisfy the constraint.

Our main contribution is a method that enables to find an approximate solution to this problem efficiently, and consists of the following steps: (i) We *relax* the integer constraints on the selection parameters $\theta_{\mathcal{R}}$, allowing them to be reals in the interval $[0, 1]$, while all other constraints remain unchanged. (ii) We prove that the relaxed optimization problem thus obtained is always *concave* for diamond networks and *almost concave* for general layered networks, facilitating efficient algorithms for solving the relaxed problem. (iii) We round the relaxed solution by choosing the top K values in the relaxed optimum to obtain a selection tuple $\{\theta_{\mathcal{R}}\}_{\mathcal{R} \in \mathcal{W}}$.

Numerical evaluations with random channel sets demonstrate: (a) $\geq 98\%$ of the integer optimum value is attained by our algorithm (using off-the-shelf non-linear optimization solvers) with a probability of 0.97 for networks of size 20 nodes, and (b) The time taken by our algorithm is less than that of an exhaustive integer optimization by factors of more than 450 for a network of $\mathcal{N} = 30$ nodes. The empirically observed running time is polynomial in \mathcal{N} for diamond networks and sub-exponential for multilayer networks.

It is to be noted that this relaxation approach outlined above is not merely restricted to the QMF/NNC-rate type objective functions, but can be used as a *general tool for selection* with any given metric; for example, we also apply our algorithm on the capacity approximation expression for diamond networks

¹This work was supported by the European Research Council grant NOWIRE ERC-2009-StG-240317.

proposed in [8] and demonstrate the results obtained.

Related Work: Previous work on relay selection in wireless networks can be divided into three categories: (i) The work of [12] and the references therein propose algorithms to select the *single best* relay (in terms of cooperative diversity) in one-layer networks. (ii) The work of [3] and the references therein analyse the performance of heuristics for selecting a subset of relays (again, from a single layer of relays) for Amplify-and-Forward (AF) based protocols. (iii) Recent work by [1] proves upper bounds on multiplicative and additive gaps for AF-based relay selection, primarily for diamond networks. The work of [8] proves general multiplicative lower bounds on the the rate achievable by a subset of relays in a diamond network.

This paper is organized as follows: Section II details the communication model and explains the problem formulation. Section III illustrates our solution strategy for diamond networks, while Section IV deals with the more involved problem of multilayer networks. Finally, Section V provides numerical evaluations that illustrate the benefits of our approach.

II. MODEL AND FORMULATION

A. Communication Model

We consider a full-duplex layered wireless network \mathcal{W} containing L layers of nodes. The source is the singleton node in layer 0, while the destination is the singleton node in layer $L - 1$. For ease of exposition, all the intermediate layers are assumed to have exactly n nodes, although our techniques can handle any configuration. As shown in Fig. 1, each signal path from the source to the destination in a layered network gets relayed by exactly the same number of hops. We assume all nodes to have a single antenna. The signal flow over this network can then be written as:

$$Y_i^{l+1} = \sum_{j=1}^n h_{ij}^l X_j^l + Z_i^{l+1} \quad (1)$$

where Y_i^{l+1} denotes the received signal at node $i \in [1 : n]$ in layer $l + 1$ ($l \in [0 : L - 2]$), h_{ij}^l denotes the complex channel coefficient from node j in layer l to node i in layer $l + 1$, X_j^l denotes the transmitted signal from node $j \in [1 : n]$ in layer l and Z_i^{l+1} denotes the i.i.d zero mean complex Gaussian noise at the receiver i in layer $l + 1$. For our network, we have a *per node* power constraint, given by $\mathbb{E}[|X_j^l|^2] \leq 1$. We also normalize the noise powers to unity, i.e., $Z_i \sim \text{i.i.d } \mathcal{CN}(0, 1)$. Notice that these *per node* signal flow equations can be coalesced into *per layer* equations involving vectors and matrices as follows:

$$\mathbf{Y}^{l+1} = \mathbf{H}^l \mathbf{X}^l + \mathbf{Z}^{l+1} \quad (2)$$

where $\mathbf{Y}^{l+1} = [Y_1^{l+1}, Y_2^{l+1}, \dots, Y_n^{l+1}]^T$, $\mathbf{X}^l = [X_1^l, X_2^l, \dots, X_n^l]^T$, \mathbf{H}^l is the MIMO channel matrix from \mathbf{X}^l to \mathbf{Y}^{l+1} with $\mathbf{H}^l(i, j) = h_{ij}^l$ and $\mathbf{Z}^{l+1} = [Z_1^{l+1}, Z_2^{l+1}, \dots, Z_n^{l+1}]^T$.

B. Capacity Outer bounds and Rate Expressions

Since the capacity of such networks cannot be characterized, the gap from the traditional cutset upper bound on capacity




Fig. 1. The Gaussian full-duplex layered network with L layers having n relays each, except the first and last one.

[2] is often used as a metric to evaluate the performance of a given relaying protocol over such networks.

A standard practice is to use inputs X_j at every network node that are picked from an i.i.d complex Gaussian distribution, i.e., $X_j \sim \text{i.i.d } \mathcal{CN}(0, 1)$. This is precisely the strategy used to prove the constant gap performance of the QMF and NNC schemes in [2], [7], and in view of these results, the following modified version of the cutset upper bound, termed \bar{C}_{iid} in [2], with the above-mentioned inputs is of interest:

$$\bar{C}_{\text{iid}} = \min_{\Omega} \sum_{l=0}^{L-1} \log \det \left(I + \mathbf{H}_{\Omega}^l \mathbf{H}_{\Omega}^{l\dagger} \right) \quad (3)$$

Here, \mathbf{H}_{Ω}^l denotes the MIMO channel matrix from \mathbf{X}_{Ω}^l to $\mathbf{Y}_{\Omega^c}^{l+1}$ and Ω denotes a *cut* in the network as shown in Fig. 1. For uniformity in dimensions, we choose to represent \mathbf{X}_{Ω}^l , $\mathbf{Y}_{\Omega^c}^{l+1}$ and \mathbf{H}_{Ω}^l as $n \times 1$, $n \times 1$ and $n \times n$ matrices respectively, by inserting zeroes in the appropriate rows and columns as dictated by the index of elements in Ω or Ω^c .

It is important to note that the achievable rate for the NNC scheme in [11] using i.i.d Gaussian vector quantizers of distortion Δ at each node is fundamentally related to the \bar{C}_{iid} expression, while additionally factoring in the quantization loss. The NNC rate for our network as per [11] is given by

$$R_{\text{NNC}} = \min_{\Omega} \left\{ \sum_{l=0}^{L-1} \log \det \left(I + \frac{\mathbf{H}_{\Omega}^l \mathbf{H}_{\Omega}^{l\dagger}}{1 + \Delta} \right) - |\Omega^c| \log \left(\frac{1 + \Delta}{\Delta} \right) \right\} \quad (4)$$

We will later argue that the differences in (3) and (4) do not substantially affect our optimization algorithm, and for the sake of clarity, we will proceed to define the subnetwork selection problem using the expressions for \bar{C}_{iid} in the remainder of this section.

C. Subnetwork Selection

We want to select an optimal subnetwork $\mathcal{W}_S^{\text{opt}}$ of \mathcal{W} that maximizes the subnetwork's \bar{C}_{iid} (or R_{NNC}) expression over all subnetworks \mathcal{W}_S with the following size constraint: \mathcal{W}_S contains K_l relays ($K_l \geq 1$) in layer $l \in [1 : L - 2]$. For this purpose, we first define a set of $n \times n$ diagonal selection matrices, \mathbf{S}_l , for each layer $l \in [1 : L - 2]$ of the form $\mathbf{S}_l = \text{diag}(\sqrt{\theta_{l1}}, \sqrt{\theta_{l2}}, \dots, \sqrt{\theta_{ln}})$ such that $\theta_{li} \in \{0, 1\}$ and

$\sum_{i=1}^n \theta_{li} = K_l$. For a given subnetwork, the matrices \mathbf{S}_l will have $\theta_{li} = 1$ iff relay $i \in [1 : n]$ in layer l is in the subnetwork and $\theta_{li} = 0$ otherwise. The reason for using a square root in the expressions will become clear in the next section.

For a given \mathcal{W}_S , the subnetwork's $\bar{C}_{S,\text{iid}}$ is given as:

$$\bar{C}_{S,\text{iid}} = \min_{\Omega} I_{S,\Omega} \quad (5)$$

where

$$I_{S,\Omega} = \sum_{l=0}^{L-1} \log \det \left(I + \mathbf{H}_{S,\Omega}^l \mathbf{H}_{S,\Omega}^{l\dagger} \right) \quad (6)$$

Here, $\mathbf{H}_{S,\Omega}^l$ denotes the (still $n \times n$) MIMO channel matrix from $\mathbf{X}_{S,\Omega}^l$ to $\mathbf{Y}_{S,\Omega^c}^{l+1}$ where $\mathbf{X}_{S,\Omega}^l = \mathbf{S}_l \mathbf{X}_{\Omega}^l$ and $\mathbf{Y}_{S,\Omega^c}^{l+1} = \mathbf{S}_{l+1} \mathbf{Y}_{\Omega^c}^{l+1}$. In this setting, $\mathbf{H}_{S,\Omega}^l$ can be related to \mathbf{H}_{Ω}^l as:

$$\mathbf{H}_{S,\Omega}^l = \begin{cases} \mathbf{S}_l \mathbf{H}_{\Omega}^l \mathbf{S}_{l+1}, & l \in [1 : L-2] \\ \mathbf{H}_{\Omega}^0 \mathbf{S}_1, & l = 0 \\ \mathbf{S}_{L-1} \mathbf{H}_{\Omega}^{L-1}, & l = L-1 \end{cases} \quad (7)$$

Essentially, $\mathbf{H}_{S,\Omega}^l$ is obtained from \mathbf{H}_{Ω}^l by replacing with 0, the rows (resp. columns) indexed by the relays in layers l (resp. $l+1$) that are not selected. This way, we retain an $n \times n$ channel matrix at each layer l that is equivalent in terms of singular values to the $K_l \times K_{l+1}$ channel matrix at that layer.

Now, the problem of finding $\mathcal{W}_S^{\text{opt}}$ essentially reduces to optimally selecting the set of diagonal integer selection matrices $\{\mathbf{S}_l\}_{l \in [1:L-2]}$. This integer optimization problem can be stated as:

$$\{\mathbf{S}_l^{\text{opt}}\}_{l \in [1:L-2]} = \arg \max_{\{\mathbf{S}_l\}_{l \in [1:L-2]}} \bar{C}_{S,\text{iid}}(\{\mathbf{S}_l\}) \quad (8)$$

$\text{tr}(\mathbf{S}_l^2) = K_l$

where the trace condition is equivalent to $\sum_{i=1}^n \theta_{li} = K_l$.

III. RELAXED APPROXIMATION-DIAMOND NETWORKS

We first illustrate our relaxation approach for (approximately) solving the above optimization problem in (8) by taking the simplest example of an n -relay diamond network. In this network, there is a set of n relays in layer 1, while layer 0 and layer 2 contain the source and destination nodes respectively. For this network, we essentially have only 1 selection matrix, $\mathbf{S}_1 = \text{diag}(\sqrt{\theta_{11}}, \sqrt{\theta_{12}}, \dots, \sqrt{\theta_{1n}})$ corresponding to the relays in layer 1 that we have to optimize.

For purposes of simplification, we make the following abuse of notation in the remainder of this section: (i) $\theta_{1i} \leftarrow \theta_i$ since there is only 1 layer of relays; (ii) $h_{i1}^0 \leftarrow h_i^0$ to denote the channels from the source (indexed as node 1 in layer 0) to relay i ; (iii) $h_{1i}^1 \leftarrow h_i^1$ to denote the channels from relay i (in layer 1) to the destination (indexed as node 1 in layer 2).

Specializing (8) for the diamond network, where we wish to select the optimal subnetwork having K_1 relays, we have:

$$\{\theta_i^{\text{opt}}\}_{i \in [1:n]} = \arg \max_{\{\theta_i\}_{i \in [1:n]} \in \{0,1\}: \sum_{i=1}^n \theta_i = K_1} \bar{C}_{S,\text{iid}}^{\text{dia}}(\{\theta_i\}) \quad (9)$$

where

$$\bar{C}_{S,\text{iid}}^{\text{dia}}(\{\theta_i\}) = \min_{\Omega} \left\{ \log \left(1 + \sum_{i \in \Omega} \theta_i |h_i^0|^2 \right) + \log \left(1 + \sum_{i \in \Omega^c} \theta_i |h_i^1|^2 \right) \right\} \quad (10)$$

A. Relaxing the Integer Program

For an approximate solution to the integer program in (9), we first *relax* the constraints in the problem as follows: Instead of using the integer θ_i 's lying in the discrete set $\{0,1\}$, we replace them with *real* variables $\tilde{\theta}_i$'s that lie in the interval $[0,1]$. With this relaxation, the following theorem holds:

Theorem 3.1: The optimization problem, defined as:

$$\{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]} = \arg \max_{\{\tilde{\theta}_i\}_{i \in [1:n]} \in [0,1]: \sum_{i=1}^n \tilde{\theta}_i = K_1} \bar{C}_{S,\text{iid}}^{\text{dia}}(\{\tilde{\theta}_i\}) \quad (11)$$

where $\bar{C}_{S,\text{iid}}^{\text{dia}}(\{\cdot\})$ is defined in (10), is *concave maximization* problem in $\{\tilde{\theta}_i\}_{i \in [1:n]}$

Proof: Observe that the constraints on $\tilde{\theta}_i$ are *linear*. Hence, it remains to show that $\bar{C}_{S,\text{iid}}^{\text{dia}}$ is concave in $\{\tilde{\theta}_i\}_{i \in [1:n]}$.

To this end, observe that for a given cut Ω , $(1 + \sum_{i \in \Omega} \tilde{\theta}_i |h_i^0|^2)$ and $(1 + \sum_{i \in \Omega^c} \tilde{\theta}_i |h_i^1|^2)$ are *affine* functions of $\{\tilde{\theta}_i\}_{i \in [1:n]}$. Hence, $\log(1 + \sum_{i \in \Omega} \tilde{\theta}_i |h_i^0|^2)$ and $\log(1 + \sum_{i \in \Omega^c} \tilde{\theta}_i |h_i^1|^2)$ are concave in $\{\tilde{\theta}_i\}_{i \in [1:n]}$, and so is their sum. Moreover, since the point wise minimum of concave functions is also concave, we can conclude that $\bar{C}_{S,\text{iid}}^{\text{dia}}$ is concave in $\{\tilde{\theta}_i\}_{i \in [1:n]}$, which proves the theorem. Notice here the significance of using a square root in the diagonal entries of the selection matrices, which lead to the affine functions inside the log terms. ■

Theorem 3.1 ensures the *existence* of a *polynomial time* algorithm (in the number of relays n) that solves the relaxed optimization problem in (11), for example using the interior point method for concave maximization, provided there exists a polynomial time algorithm to find $\bar{C}_{S,\text{iid}}^{\text{dia}}$.

Finding $\bar{C}_{S,\text{iid}}^{\text{dia}}$ *a priori* consists of evaluating 2^n terms corresponding to the cuts Ω and then taking a minimum, which takes exponential time. However, it was shown in [10], that the mutual information terms inside the minimization of $\bar{C}_{S,\text{iid}}^{\text{dia}}$ are *submodular* in the sets Ω . Since submodular minimization can be accomplished using only a polynomial (in n) number of evaluations of the mutual information terms [9], (11) can be solved in polynomial time.

B. Rounding the Relaxed $\tilde{\theta}_i$'s

Clearly, since the feasible set for (11) is a superset of (9), the relaxed optimal value will be greater. However, the $\{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]}$ can have fractional values that do not correspond to an actual subnetwork of size K_1 . The next step then is to *round* the fractional solution of (11) to a discrete solution that represents a subnetwork selection. Mathematically, a rounding is a map $f_R : \{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]} \mapsto \{\theta_i^{\text{sel}}\}_{i \in [1:n]}$ such that $\{\theta_i^{\text{sel}}\}_{i \in [1:n]} \in \{0,1\}$ and $\sum_{i=1}^n \theta_i^{\text{sel}} = K_1$. An intuitive way to perform the rounding in this case would be to set $\theta_i^{\text{sel}} = 1$ iff $\tilde{\theta}_i^{\text{opt}}$ is among the maximum K_1 values in the set $\{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]}$ and set $\theta_i^{\text{sel}} = 0$ otherwise.

C. Using R_{NNC} instead of \bar{C}_{iid} for optimization

The relaxation approach illustrated above extends naturally to the QMF/NNC rate expressions R_{NNC} . For this, specializ-

ing (4) to diamond networks, and using relaxed $\{\tilde{\theta}_i\}$'s, the optimization problem can be stated as:

$$\{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]} = \arg \max_{\substack{\{\tilde{\theta}_i\}_{i \in [1:n]} \in [0,1]: \\ \sum_{i=1}^n \tilde{\theta}_i = K_1}} \bar{R}_{S,\text{NNC}}^{\text{dia}}(\{\tilde{\theta}_i\}) \quad (12)$$

where

$$\bar{R}_{S,\text{NNC}}^{\text{dia}}(\{\tilde{\theta}_i\}) = \min_{\Omega} \left\{ \begin{array}{l} \log(1 + \sum_{i \in \Omega} \tilde{\theta}_i \frac{|h_i^0|^2}{K_1}) + \log(1 + \sum_{i \in \Omega^c} \tilde{\theta}_i |h_i^1|^2) \\ - \sum_{i \in \Omega^c} \tilde{\theta}_i \log(\frac{K_1}{K_1-1}) \end{array} \right\} \quad (13)$$

using $\Delta = K_1 - 1$ as in [11]. From Theorem 3.1, we can straightaway conclude the concavity of (12), since the only additional terms, $\sum_{i \in \Omega^c} \tilde{\theta}_i \log(\frac{K_1}{K_1-1})$, are linear in $\{\tilde{\theta}_i\}_{i \in [1:n]}$.

D. Applications in other capacity approximations

For n -relay diamond networks, a simpler and more approximate expression based on point-to-point link capacities has been proposed for capacity approximation in [8], given by:

$$\bar{C}_{\text{P2P}}^{\text{dia}} = \min_{\Omega} \left\{ \max_{i \in \Omega} \log(1 + |h_i^0|^2) + \max_{i \in \Omega^c} \log(1 + |h_i^1|^2) \right\} \quad (14)$$

The inherent advantage of working with (14) is that $\bar{C}_{\text{P2P}}^{\text{dia}}$ can be evaluated in $O(n \log(n))$ time [8], which is faster than the polynomial time submodular minimization algorithms needed to evaluate $\bar{C}_{\text{iid}}^{\text{dia}}$ or $R_{\text{NNC}}^{\text{dia}}$. On the flip side, this approximation is not good for low SNRs and it does not generalize to multi-layered networks beyond the diamond topology. However, for diamond networks, we can still apply our relaxation framework on the $\bar{C}_{\text{P2P}}^{\text{dia}}$ expression to get a set of relays that (approximately) maximizes $\bar{C}_{\text{P2P}}^{\text{dia}}$ and see how that selected set of relays perform in terms of $\bar{C}_{\text{iid}}^{\text{dia}}$ or $R_{\text{NNC}}^{\text{dia}}$. In this case, the relaxed optimization problem can take the following form:

$$\{\tilde{\theta}_i^{\text{opt}}\}_{i \in [1:n]} = \arg \max_{\substack{\{\tilde{\theta}_i\}_{i \in [1:n]} \in [0,1]: \\ \sum_{i=1}^n \tilde{\theta}_i = K_1}} \bar{C}_{S,\text{P2P}}^{\text{dia}}(\{\tilde{\theta}_i\}) \quad (15)$$

where

$$\bar{C}_{S,\text{P2P}}^{\text{dia}}(\{\tilde{\theta}_i\}) = \min_{\Omega} \left\{ \begin{array}{l} \max_{i \in \Omega} \tilde{\theta}_i \log(1 + |h_i^0|^2) \\ + \max_{i \in \Omega^c} \tilde{\theta}_i \log(1 + |h_i^1|^2) \end{array} \right\} \quad (16)$$

Note that unlike (11), (15) is *not* a concave optimization problem and in general, non-linear optimization algorithms can potentially get stuck in local maximas. Nevertheless, owing to the faster speed in computing $\bar{C}_{S,\text{P2P}}^{\text{dia}}$, it is worth giving this expression a try, and surprisingly, we show in Section V that off-the-shelf non-linear optimizers do give good results with the $\bar{C}_{S,\text{P2P}}^{\text{dia}}$ expression used for selection.

IV. RELAXED APPROXIMATION–MULTILAYER NETWORKS

For multilayer networks, the procedure is similar to the one outlined for diamond networks. The integer optimization problem over the $n(L-2)$ variables $\{\theta_{li}\}_{i \in [1:n], l \in [1:L-2]}$, as given in (8), can be relaxed to the corresponding continuous problem in $\{\tilde{\theta}_i\}$'s. Once the relaxed optimization problem is solved, the optimal fractional solution is rounded to an integer solution representing a subnetwork of appropriate size.

The objective function in the relaxed version of (8) is the minimum of $2^{n(L-2)}$ terms, each of which is a sum of terms of the form:

$$I_{S,\Omega}^l = \log \det(I + \tilde{\mathbf{S}}_l \mathbf{H}_{\Omega}^l \tilde{\mathbf{S}}_{l+1}^l \mathbf{H}_{\Omega}^{l+1} \tilde{\mathbf{S}}_l)$$

where $\tilde{\mathbf{S}}_l = \text{diag}(\sqrt{\tilde{\theta}_{l1}}, \dots, \sqrt{\tilde{\theta}_{ln}})$, $\text{tr}(\tilde{\mathbf{S}}_l^2) = K_l$ and $\text{tr}(\tilde{\mathbf{S}}_{l+1}^2) = K_{l+1}$. In general, the above term is not a concave function of $\{\tilde{\theta}_{li}\}_{i \in [1:n]}$ and $\{\tilde{\theta}_{l+1,i}\}_{i \in [1:n]}$. Empirical evidence however suggests that it is *almost* concave.

Firstly, we denote by \mathbf{x} , the vector of variables

$$\mathbf{x} = \{\tilde{\theta}_{l1}, \dots, \tilde{\theta}_{ln}, \tilde{\theta}_{l+1,1}, \dots, \tilde{\theta}_{l+1,n}\}$$

Thus $I_{S,\Omega}^l(\mathbf{x})$ defines a hyper-surface. For a fixed set of channel coefficients, if this surface is exactly concave, then for any pair of points $\mathbf{x}^1, \mathbf{x}^2 \in [0,1]^{2n}$ such that $\sum_{i=1}^n x_i^j = K_l$ and $\sum_{i=n+1}^{2n} x_i^j = K_{l+1}$ for $j = 1, 2$ and for every $\lambda \in [0,1]$, the following quantity must be always non-negative:

$$D(\mathbf{x}^1, \mathbf{x}^2, \lambda) = I_{S,\Omega}^l(\lambda \mathbf{x}^1 + (1-\lambda)\mathbf{x}^2) - \lambda I_{S,\Omega}^l(\mathbf{x}^1) - (1-\lambda)I_{S,\Omega}^l(\mathbf{x}^2)$$

In our experiments, we compute the probability $P_{ccv}(\mathbf{H}_{\Omega}^l)$ that a random pair of $\mathbf{x}^1, \mathbf{x}^2$ satisfies $D(\mathbf{x}^1, \mathbf{x}^2, \lambda) \geq 0$ for all values of $\lambda \in [0,1]$ picked with sufficient granularity (10^{-3} to be precise). The individual channel coefficients h_{ij}^l in each instance were picked i.i.d as follows: $10 \log_{10}(|h_{ij}^l|^2) \sim \mathcal{U}[0, 35]$ and $\angle h_{ij}^l \sim \mathcal{U}[0, 2\pi]$. Ω was fixed such that all nodes in layer l were in Ω and all nodes in layer $l+1$ were in Ω^c (but the results do not depend on the choice of Ω). The empirically observed value of this probability, averaged over several random channel-set instantiations, for different values of n and $K_l = K_{l+1} = n/2$ are as follows:

n	2	4	≥ 6
$\mathbb{E}_{\mathbf{H}_{\Omega}^l} [P_{ccv}(\mathbf{H}_{\Omega}^l)]$	0.9876	0.9997	≈ 1

This demonstrates that $I_{S,\Omega}^l(x)$ is *almost concave*, implying that the relaxed version of (8) also has similar properties.

V. NUMERICAL EVALUATIONS

Algorithms: We evaluate three algorithms for subnetwork selection. (i) RLX-FULL: the main algorithm corresponding to the relaxed problem (11) for diamond networks and the relaxed version of (8) for multilayer networks. (ii) RLX-SMPL: specific to diamond networks, corresponding to the relaxed problem (15). (iii) RND: baseline algorithm where a subnetwork of appropriate size is selected at random. For both RLX-FULL and RLX-SMPL, the fractional optimum is rounded by picking the top K_i values in each layer and accordingly selecting the subnetwork (as described in Section III-B).

Implementation: The implementations (in C++) require two main modules: (i) A submodular minimization routine that evaluates $\bar{C}_{S,\text{iid}}(\{\tilde{\mathbf{S}}_l\})$ for a specific set $\{\tilde{\mathbf{S}}_l\}$. For this, we used the C implementation of an algorithm based on the minimum-norm base [5], shown to be the most efficient general purpose routine in this regard. (ii) A routine that solves (11) and the relaxed version of (8). In view of Theorem 3.1, suitable interior-point based methods can solve (11) in polynomial time. However, off-the-shelf open source libraries to this end gave less than satisfactory results in practice. Instead, a Nelder-Mead simplex-based general purpose non-linear optimization

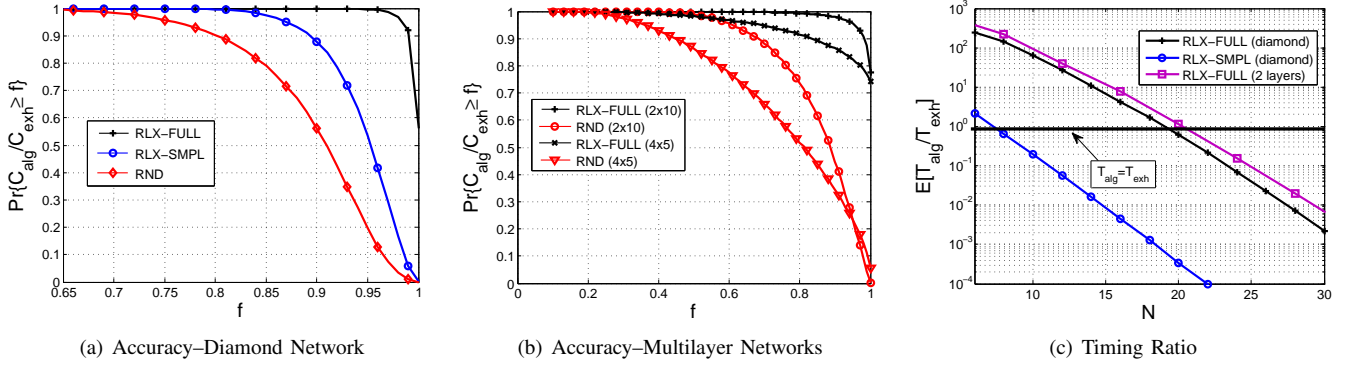


Fig. 2. Accuracy and Timing Performance of Algorithms

routine from the NLOpt library is used [6]. Specifically, the NLOPT_LN_NELDERMEAD function, combined with the augmented Lagrangian method in NLOPT_AUGLAG_EQ (to encode the size constraints) was used.

In all our experiments, for a given (n, K_1, \dots, K_{L-2}) size tuple, we ran each algorithm for several (greater than 10^9) random channel-set instantiations of the network. The individual channel coefficients h_{ij}^l in each instance were picked i.i.d as follows: $10 \log_{10}(|h_{ij}^l|^2) \sim \mathcal{U}[0, 35]$ and $\angle h_{ij}^l \sim \mathcal{U}[0, 2\pi]$.

A. Accuracy Results

For each random channel-set instantiation, we compute the ratio C_{alg}/C_{exh} , where C_{alg} (resp. C_{exh}) denotes $\bar{C}_{S, \text{iid}}(\{\tilde{S}_l\})$ of the optimal subnetwork selected by our algorithms (resp. by exhaustive search). Naturally, with this metric, the higher the C_{alg}/C_{exh} ratio, the better is the accuracy of the algorithm.

Fig. 2(a) plots the complementary c.d.f. of C_{alg}/C_{exh} for the three algorithms over diamond networks. The number of relays is $n = 20$ and the subnetwork size is $K_1 = 10$. Clearly, RLX-FULL produces subnetworks that have throughput equal or very close to the exhaustive optimal most of the time. Some sample values read off from the curve are as follows

	RLX-FULL	RLX-SMPL	RND
$Pr\{C_{alg}/C_{exh} \geq 0.98\}$	0.9697	0.1679	0.0345
$Pr\{C_{alg}/C_{exh} \geq 0.94\}$	0.9983	0.6349	0.2724
$Pr\{C_{alg}/C_{exh} \geq 0.90\}$	0.9995	0.8782	0.5617

It is clear from the above that both RLX-FULL and RLX-SMPL perform significantly better than RND.

For multilayer networks (performance shown in Fig. 2(b)), experiments were performed for two different configurations: In the first one, marked as 2×10 in Fig. 2(b), there are two intermediate layers of $n = 10$ nodes each and $K_1 = K_2 = 5$. In the second configuration, marked as 4×5 in Fig. 2(b), there are 4 intermediate layers of 5 nodes each with a staggered size constraint of $K_1 = 2, K_2 = 4, K_3 = 2, K_4 = 3$. Again, in both configurations, we see that the complementary c.d.f. of RLX-FULL consistently outperforms that of RND and the benefits increase significantly at higher accuracies (i.e., higher C_{alg}/C_{exh} ratios).

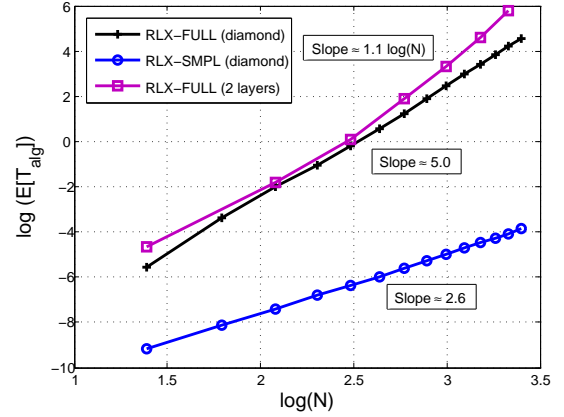


Fig. 3. $\log(\mathbb{E}[T_{alg}])$ vs $\log(N)$. A slope of $\delta \Rightarrow$ running time of $O(n^\delta)$

B. Time Complexity

To measure the time efficiency of our algorithms, we construct the following configurations: (i) A diamond network with N relays, from which we select $N/2$ relays, and (ii) A layered network with having 2 intermediate layers of $N/2$ relays each, and we select $N/4$ relays from each layer.

For each N , we perform a large number of experiments with random channel-set instantiations and plot the average value of T_{alg}/T_{exh} in Fig. 2(c), where T_{alg} (resp. T_{exh}) denotes the running time of our algorithms (resp. exhaustive search).

For the diamond network with $N = 30$, RLX-FULL is more efficient than an exhaustive search by a factor of **460**, while for RLX-SMPL, this factor is more than 2.1×10^6 . This is primarily due to the much faster $O(N \log N)$ computation time of $\bar{C}_{P2P}^{\text{dia}}$ (in RLX-SMPL) w.r.t that of the submodular minimization routine for $\bar{C}_{S, \text{iid}}^{\text{dia}}$. For the multilayer configuration, RLX-FULL gives time saving factors of more than **50** and **440** for $N = 28$ and 32 respectively.

While it is difficult to theoretically analyze the time complexity of a Nelder-Mead simplex-based algorithm for our problem, in Fig. 3 we give an empirical demonstration of time complexity for our implementations, where we plot $\log(T_{alg})$ (averaged over random channel-set instances) vs $\log(N)$ for the two configurations above.

For the diamond network, a fairly linear behavior is obtained, with slopes of approximately $\delta = 5.0$ and $\delta = 2.6$

for RLX-FULL and RLX-SMPL, implying that their running time is approximately $O(\mathcal{N}^{5.0})$ and $O(\mathcal{N}^{2.6})$ respectively. For the two-layered configuration, the slope is not constant, but a slowly growing function of \mathcal{N} (about $1.1 \log \mathcal{N}$ to a first approximation). Nevertheless, this is still the first systematic sub-exponential complexity ($\approx O(\mathcal{N}^{1.1 \log \mathcal{N}})$) algorithm for (approximately) solving the original integer optimization problem for multilayer networks, providing significant time savings w.r.t an exhaustive search (Fig. 2(c)). Also, with customized solvers (as opposed to the general purpose routines used here), further complexity gains are expected.

REFERENCES

- [1] S. Agnihotri, S. Jaggi, and M. Chen, "Analog network coding in general SNR regime: performance of network simplification", *Proc. IEEE Information Theory Workshop (ITW)*, pp. 632–636, Sept. 2012.
- [2] A. S. Avestimehr, S. N. Diggavi, and D. N. C. Tse, "Wireless network information flow: a deterministic approach", *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 1872–1905, Apr. 2011.
- [3] J. Cai, S. Shen, J. W. Mark, and A. S. Alfa, "Semi-distributed user relaying algorithm for amplify-and-forward wireless relay networks", *IEEE Transactions on Wireless Communications*, vol. 7, no. 4, pp. 1348–1357, Apr. 2008.
- [4] B. Chern, and A. Ozgur, "Achieving the capacity of the N -relay gaussian diamond network within $\log N$ bits", *Proc. IEEE Information Theory Workshop (ITW)*, pp. 377–380, Sept. 2012.
- [5] S. Fujishige, and S. Isotani, "A submodular function minimization algorithm based on the minimum-norm base", *Pacific Journal of Optimization*, vol. 7, no. 1, pp. 3–17, 2011.
- [6] "NLOpt nonlinear optimization library", Available at <http://ab-initio.mit.edu/wiki/index.php/NLOpt>.
- [7] S. Lim, Y-H. Kim, A. El Gamal, and S-Y. Chung, "Noisy network coding", *IEEE Transactions on Information Theory*, vol. 57, no. 5, pp. 3132–3152, May 2011.
- [8] C. Nazeroglu, A. Ozgur, and C. Fragouli, "Wireless network simplification: the Gaussian N -relay diamond network", Available at <http://arxiv.org/abs/1103.2046>.
- [9] J. Orlin, "A faster strongly polynomial time algorithm for submodular function minimization", *Mathematical Programming*, vol. 118, no. 2, pp. 237–251, 2009.
- [10] F. Parvaresh, and R. H. Etkin, "On computing the capacity of relay networks in polynomial time", *Proc. IEEE Int. Symp. Inf. Theory*, pp. 1342–1346, 2011.
- [11] A. Sengupta, I-H. Wang, and C. Fragouli, "Optimizing Quantize-Map-and-Forward Relaying for Gaussian Diamond Networks", *Proc. IEEE Information Theory Workshop (ITW)*, pp. 381–385, Sept. 2012.
- [12] R. Tannious, and A. Nosratinia, "Spectrally-efcient relay selection with limited feedback", *IEEE Journal on Selected Areas in Communication*, vol. 26, no. 8, pp. 1419–1428, 2008.