

Switched Local Schedules for Diamond Networks

Siddhartha Brahma*, Ayan Sengupta* and Christina Fragouli†

*EPFL, Switzerland, †UCLA, USA

Abstract—We consider a Gaussian diamond network where a source communicates with the destination through n non-interfering half-duplex relays. We focus on half-duplex schedules that utilize only local channel state information, i.e., each relay has access to its incoming and outgoing channel realizations. We demonstrate that random independent switching, resulting in multiple listen-transmit sub cycles at each relay, while still respecting the overall locally optimal listen-transmit fractions, enables to approximately achieve at least $3/4$ of the capacity of the 2-relay diamond network. With a single listen-transmit cycle, this fraction drops from $3/4$ to $1/2$. We also provide simulation results that point to the same fractions of capacity being retained over networks with more than 2 relays.¹

I. INTRODUCTION

To optimally operate half duplex relay networks we need to solve a hard problem: decide what fraction of time each relay should listen or transmit (i.e., decide the half-duplex strategy) so as to optimize the network capacity. Solving this problem requires both global channel knowledge as well as use of exponential complexity operations; in this paper, we take first steps in understanding how much of the global capacity we can achieve, if we instead assume local channel knowledge and use of low-complexity randomized scheduling strategies.

Developing and understanding the performance of low complexity half-duplex schedules is an important component in making relaying practical. Physical layer cooperation is fast gaining traction as a means to achieve high spectral-efficiency communication over networks; recent approximation approaches to characterizing the capacity of relay networks [1] as well as coding schemes and first implementations provide a promise of seeing such schemes in practice. Yet, in half-duplex networks, standard assumptions for scheduling are (i) knowledge of *global* channel state information of *all* links in the network, and (ii) use of an optimization routine that assigns probabilities to a possibly *exponential* (in the number of relays) set of relaying states. Both of these assumptions are practically prohibitive: the first, since it would require dedicated inter-relay communication and exchanging global CSI over wireless channels, and the second, simply from computational complexity constraints. Moreover, such centralized relay scheduling strategies do not scale for fast-changing networks, where if a relay joins the network, the entire schedule needs to be re-derived.

In this paper, we start exploring the performance we can achieve if we restrict ourselves to only *local* relay operations. By local, we mean that two conditions are satisfied: (i)

each relay in the network only has access to its incoming and outgoing channel realizations, and (ii) there can be no communication between the relays to share CSI, and hence, no node in the network can solve a centralized optimization problem. We focus our attention on the half-duplex n -relay diamond network and use the approximate Quantize-Map-and-Forward rate expressions in [2].

We propose the following approach for the network operation: every relay in the n -relay network uses its incoming and outgoing links to derive the half-duplex listen and transmit fractions that would be optimal in the absence of all other relays in the network. Having derived this fraction locally, we allow the relays to *switch* multiple times (say κ) between listen and transmit modes *independently at random* in the duration of operation, while still respecting the overall listen-transmit fractions at each relay. Having every relay do this independently of each other induces a *global* schedule over the resulting *linear* (in n) number of states.²

We demonstrate that with this approach, for the 2 relay diamond network, we can provably achieve at least $3/4$ of the optimal (encompassing global CSI sharing and centralized optimization over schedules) objective function value as the number of switches $\kappa \rightarrow \infty$. We also prove that the performance of the *deterministic* local strategy with only *one* listen-transmit cycle in the duration of operation drops the above fraction down to $1/2$. Numerical evaluation of our strategies over networks with larger number of relays suggest that the same fractions of capacity are retained over them as well. Interestingly, we see that incorporating only a few switches already enables to leverage most of the benefits that we prove in the limiting $\kappa \rightarrow \infty$ case: for instance, over 2-relay networks, we extract 71% of the gain (in expectation) proved for $\kappa \rightarrow \infty$ by incorporating two random listen-transmit sub cycles. We also observe that in expectation, the switching gain over the baseline deterministic local strategy increases as the network size increases—up from 11% for 2-relay networks to 24% for 5-relay networks using only two sub cycles.

Related Work: In previous work, the authors in [5] show for a specific case that schedules using only local information and one switch at each relay can achieve rates close to capacity. The authors in [2], [3] approach low complexity relaying in half-duplex diamond networks by reducing the number of relaying states or using a subset of relays to approximately achieve a constant fraction of the capacity. We would like to note that the use of randomization in our work

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

²We note that in practice, the patterns for mode switching at the relays can be *pseudorandom* and the seed is shared with the destination for decoding.

is fundamentally different from the use of random switching in [6] and [4]. In their work, the randomness in the switching sequence is used to convey additional information from the source to the destination while we use randomness to generate schedules that achieve higher rates.

II. PROBLEM FORMULATION




Fig. 1. The Gaussian n -relay half-duplex diamond network with a possible relaying state and cut.

A. Network Model

We consider the Gaussian n -relay diamond network, where a source \mathcal{S} transmits information to a destination \mathcal{D} with the help of half-duplex relays. At any given time t , each relay \mathcal{R}_i can either listen (L) or transmit (T), but not both; we denote its state by $M_i[t] \in \{L, T\}$.

Let $X_s[t]$ denote the signal transmitted by \mathcal{S} at time t and $X_i[t]$ denote that transmitted by relay \mathcal{R}_i . Let $Y_d[t]$ and $Y_i[t]$ denote the signals received by \mathcal{D} and \mathcal{R}_i respectively. The signal flow over the network is then written as:

$$\begin{aligned} Y_i[t] &= h_{is}X_s[t] + Z_i[t] \text{ when } M_i[t] = L \\ &= 0 \text{ when } M_i[t] = T \\ X_i[t] &= 0 \text{ when } M_i[t] = L \\ Y_d[t] &= \sum_{i=1}^n h_{id}X_i[t] + Z[t] \end{aligned}$$

where h_{is}, h_{id} are the complex channel coefficients from \mathcal{S} to \mathcal{R}_i and \mathcal{R}_i to \mathcal{D} respectively. $Z_i[t]$ and $Z[t]$ are i.i.d white Gaussian random processes with $\mathcal{CN}(0, 1)$ distributions. The power constraints for the source and all the relays are normalized to unity.

We can then calculate the individual link capacities from \mathcal{S} to \mathcal{R}_i (denoted by l_i) and from \mathcal{R}_i to \mathcal{D} (denoted by r_i) as

$$l_i = \log(1 + |h_{is}|^2), \quad r_i = \log(1 + |h_{id}|^2) \quad (1)$$

For any positive integer α , let $[1 : \alpha]$ denote the set $\{1, 2, \dots, \alpha\}$. For $s \in [1 : 2^n]$, let $m_s \in \{L, T\}^n$ be a distinct relaying state of the network. The fraction of time the relays spend in state m_s will be denoted by f_s , where $\sum_{s \in [1:2^n]} f_s = 1$. We will use $\mathcal{L}(m_s)$ and $\mathcal{T}(m_s)$ to denote the set of indices of the relays in listening and transmitting state in m_s , respectively. Also, for $j \in [1 : 2^n]$, $\Lambda_j \subseteq [1 : n]$ denotes the cut separating $\mathcal{S} \cup (\cup_{i \in \Lambda_j} \mathcal{R}_i)$ from $\mathcal{D} \cup (\cup_{i \in \bar{\Lambda}_j} \mathcal{R}_i)$. Unless otherwise stated, the term “constant” will mean a quantity that is a function of n , independent of the channel SNRs.

B. An Approximation to the Capacity

Let C_{hd}^n denote the capacity of the n -relay half-duplex diamond network; to achieve it, we need to optimize over $\{f_s\}_{s \in [1:2^n]}$, the fractions of time that the relays spend in each state m_s . From the work of [1], [7], C_{hd}^n can be approximated up to an additive constant by a quantity C_{opt}^n that is only a function of the link capacities $\{l_i, r_i\}$ defined in (1).

Theorem 2.1: For an n relay half-duplex diamond network, there exist constants $G(n)$ and $G'(n)$ such that

$$C_{opt}^n - G'(n) \leq C_{hd}^n \leq C_{opt}^n + G(n) \quad (2)$$

where

$$C_{opt}^n = \max_{\{f_s\}} \min_{j \in [1:2^n]} \sum_{s=1}^{2^n} f_s \left(\max_{i \in \bar{\Lambda}_j \cap \mathcal{L}(m_s)} l_i + \max_{i \in \Lambda_j \cap \mathcal{T}(m_s)} r_i \right)$$

The minimization is over all the cuts $\{\Lambda_j\}_{j \in [1:2^n]}$ and the maximization is over all schedules $\{f_s\}_{s \in [1:2^n]}$ such that $\sum_{s \in [1:2^n]} f_s = 1$. $\bar{\Lambda}_j$ is $[1 : n] \setminus \Lambda_j$.

It is to be noted that [1] guarantees C_{hd}^n can be achieved with QMF relaying within a constant; hence, from the above theorem, we can conclude that C_{opt}^n is *also* achievable by QMF relaying within an additive constant, which positively motivates working with the C_{opt}^n expression in the remainder. C_{opt}^n is also the optimum solution of the following LP:

$$\text{LP : Maximize } C \quad (3)$$

$$\begin{aligned} \sum_{s=1}^{2^n} f_s \left(\max_{i \in \bar{\Lambda}_j \cap \mathcal{L}(m_s)} l_i + \max_{i \in \Lambda_j \cap \mathcal{T}(m_s)} r_i \right) &\geq C; j \in [1 : 2^n] \\ \sum_{s=1}^{2^n} f_s &= 1; f_s \geq 0 \forall s \in [1 : 2^n]; C \geq 0 \end{aligned}$$

where we have $2^n + 1$ variables (C and $\{f_s\}_{s \in [1:2^n]}$).

C. Local Scheduling Strategy

Achieving C_{opt}^n requires global knowledge of the link strengths in order to solve the optimization problem in (3) and determine the fraction of time spent in each scheduling state m_s . In practice, this may be expensive and will require inter-relay communication as well as a central node (eg. the source) that performs the optimization. The optimization itself consists of $2^n + 1$ variables and $2^n + 1$ constraints; solving it explicitly becomes prohibitive even for moderately large values of n . Instead, as a more practical approach, we look at what can be achieved by using only local information at the relays; we assume that each relay only knows the channel strengths of its incoming and outgoing links, i.e., \mathcal{R}_i knows l_i and r_i .

1) Multiple Independent Switches: At a local level, each relay can control when it switches from an L state to a T state and vice-versa, and it can do so multiple times within the duration of operation. We normalize the duration to unity and hence, all the switches are made at points in the interval $[0, 1]$. For purposes of counting, a switch will always denote a *transition from L to T* . We will also assume that each relay always starts in a L state and ends in a T state. Thus, if a relay makes κ switches, there will be κ transitions from L to T and $\kappa - 1$ transitions from T to L . Choosing a local




Fig. 2. Example illustration of randomized switching and induced states in a 3 relay network with each relay using 2 switches.

switching strategy then amounts to choosing $2\kappa - 1$ points on the unit interval for each relay independently (see Fig. 2). For \mathcal{R}_i , let the points (in ascending order) at which the transitions from L to T happen be denoted by $p_{i,1}^{L \rightarrow T}, \dots, p_{i,\kappa}^{L \rightarrow T}$ and let the points (in asc. order) at which the transitions from T to L happen be denoted by $p_{i,1}^{T \rightarrow L}, \dots, p_{i,\kappa-1}^{T \rightarrow L}$. Together, they define the *switching sequence* for \mathcal{R}_i , which is denoted by $\mathcal{P}_i(\kappa)$.

$$\mathcal{P}_i(\kappa) = \{p_{i,1}^{L \rightarrow T}, p_{i,1}^{T \rightarrow L}, \dots, p_{i,\kappa-1}^{T \rightarrow L}, p_{i,\kappa}^{L \rightarrow T}\}$$

For ease of notation, we set $p_{i,\kappa}^{T \rightarrow L} = 1$ and $p_{i,0}^{L \rightarrow T} = 0$ for all $i \in [1 : n]$. We use $\mathcal{P}(\kappa)$ to denote the union of all the switching sequences, i.e., $\mathcal{P}(\kappa) = \cup_{i \in [1:n]} \{\mathcal{P}_i(\kappa)\}$. The total time spent by \mathcal{R}_i in state L (denoted by $F_{i,L}(\mathcal{P}(\kappa))$) and in state T (denoted by $F_{i,T}(\mathcal{P}(\kappa))$), can then be computed as

$$F_{i,L}(\mathcal{P}(\kappa)) = \sum_{k=1}^{\kappa} (p_{i,k}^{L \rightarrow T} - p_{i,k-1}^{T \rightarrow L})$$

$$F_{i,T}(\mathcal{P}(\kappa)) = \sum_{k=1}^{\kappa} (p_{i,k}^{T \rightarrow L} - p_{i,k}^{L \rightarrow T})$$

For a given $\mathcal{P}(\kappa)$, the rate achieved by the network is denoted by $C^n(\mathcal{P}(\kappa))$. In this paper, we focus on *randomized switching*, i.e. each relay \mathcal{R}_i chooses the positions in $\mathcal{P}_i(\kappa)$ randomly³. Note that the above strategy is linear (in n) in terms of the *state complexity*, i.e., the number of relaying states with non-zero probabilities. More precisely, the number of active states for κ switches is at most $\min\{2^n, n(2\kappa - 1) + 1\}$.

2) *Local Optimality of Overall Fractions*: We have each \mathcal{R}_i listen and transmit for an (overall) fraction of time that is *optimal* for an isolated single half-duplex relay \mathcal{R}_i (essentially a one-hop line network). In the absence of any other information about the strengths of links connecting the other relays, this is a reasonable strategy to follow. It is easy to see that for \mathcal{R}_i , the optimal listening and transmitting fractions that our strategy should choose are as follows

$$F_{i,L}(\mathcal{P}(\kappa)) = \frac{r_i}{l_i + r_i} \text{ and } F_{i,T}(\mathcal{P}(\kappa)) = \frac{l_i}{l_i + r_i} \quad (4)$$

The quantity we will be interested in is the *expected rate* achieved by the network for κ switches $C_{rnd}^n(\kappa) = \mathbb{E}[C^n(\mathcal{P}(\kappa))]$, where the expectation is taken over the choices of $\mathcal{P}_i(\kappa)$'s that satisfy the criterion in (4).

³This can practically be thought of as a pseudorandom pattern that is shared with the destination.

3) *Varying the Number of Switches*: We also wish to analyze the performance of our strategy as we (progressively) increase the number of switches that each relay employs. In that regard, we have the following extremes cases:

Deterministic Switching: When $\kappa = 1$, the set $\mathcal{P}(\kappa)$ is uniquely determined. Each \mathcal{R}_i makes only one switch from L to T at the point $\frac{r_i}{l_i + r_i}$ (see Fig. 3 for illustration). Notice that this corresponds to a *deterministic* switching strategy and hence we denote the achieved rate by $C_{det}^n \equiv C_{rnd}^n(1)$.




Fig. 3. Deterministic switching and induced states in a 3 relay network. The local L and T fractions are shown in red and blue color, respectively.

Continuous Switching: In the limit of κ becoming very large, the relays switch in a manner such that at each instant, \mathcal{R}_i is in state L with probability $\frac{r_i}{l_i + r_i}$ and in state T with probability $\frac{l_i}{l_i + r_i}$. We denote the limiting rate by $C_{lim}^n \equiv \lim_{\kappa \rightarrow \infty} C_{rnd}^n(\kappa)$.

In order to further understand the performance limits of our local randomized switching strategy, we also look at an *upper bound* to the rates achieved by any strategy that follows the local optimality criterion (4). This can be computed by adding constraints to the LP (3) as shown in what follows. The optimal rate so obtained is denoted by C_{lopt}^n . Table 1 summarizes the quantities defined in this section.

κ	Number of switches from L to T in each relay.
C_{opt}^n	Capacity of the network.
$C_{rnd}^n(\kappa)$	Expected rate achieved for κ switches.
C_{det}^n	Rate achieved for $\kappa = 1$.
C_{lim}^n	Rate achieved in the limit of large κ .
C_{lopt}^n	Upper bound to rates for strategies satisfying (4).

Table 1. Summary of quantities considered.

D. Computation of Rates

A particular switching sequence $\mathcal{P}(\kappa)$ uniquely induces a *global schedule* $\{f_s(\mathcal{P}(\kappa))\}_{s \in [1:2^n]}$. The fraction of time the network is in state $m_s \in \{L, T\}^n$ is the union of (possibly disjoint) intervals where each \mathcal{R}_i is in the state implied by m_s . More formally,

$$f_s(\mathcal{P}(\kappa)) = \left| \bigcup_{k_{[1:n]} \in [1:\kappa]^n} \left\{ \bigcap_{i \in \mathcal{L}(m_s)} [p_{i,k_i}^{T \rightarrow L}, p_{i,k_i}^{L \rightarrow T}] \cap \bigcap_{i \in \mathcal{T}(m_s)} [p_{i,k_i}^{L \rightarrow T}, p_{i,k_i}^{T \rightarrow L}] \right\} \right|$$

where $k_{[1:n]}$ is the n -tuple $\{k_1, \dots, k_n\}$ and the union is over all possible κ^n n -tuples in $[1 : \kappa]^n$. In practice, each term $f_s(\mathcal{P}(\kappa))$ can be computed efficiently by first sorting all the points (irrespective of the relay index) in $\mathcal{P}(\kappa)$ and then traversing the sorted sequence, keeping track of the relay

states. Once this is accomplished, $C^n(\mathcal{P}(\kappa))$ can be computed by performing a min-cut computation as follows:

$$C^n(\mathcal{P}(\kappa)) = \min_{j \in [1:2^n]} \sum_{s=1}^{2^n} f_s(\mathcal{P}(\kappa)) \left(\begin{array}{c} \max_{i \in \Lambda_j \cap \mathcal{L}(m_s)} l_i \\ + \max_{i \in \Lambda_j \cap \mathcal{T}(m_s)} r_i \end{array} \right) \quad (5)$$

The expected rate $C_{rnd}^n(\kappa)$ can then be computed numerically by taking a large enough sample of random $\mathcal{P}(\kappa)$'s.

The above discussion trivially holds for C_{det}^n . For C_{lim}^n , as the random switches at each relay occur independently of each other, in the limit of large κ as discussed in Section II-C3, the fraction of time spent in a state m_s is given by:

$$f_s^{lim} = \prod_{i \in \mathcal{L}(m_s)} \frac{r_i}{l_i + r_i} \prod_{i \in \mathcal{T}(m_s)} \frac{l_i}{l_i + r_i}$$

C_{lim}^n can then be computed by setting $f_s(\mathcal{P}(\kappa)) = f_s^{lim}$ in (5). Finally, C_{lopt}^n can be computed by adding the following constraints to LP (3):

$$\sum_{s: i \in \mathcal{L}(m_s)} f_s = \frac{r_i}{l_i + r_i} \forall i \in [1:n] \quad (6)$$

The quantity on the left of (6) represents the total fraction of time \mathcal{R}_i is in state L for a schedule $\{f_s\}_{s \in [1:2^n]}$. The remainder of the LP remains the same and the optimum represents the maximum rate achievable by strategies (not necessarily without coordination) that follow (4).

In the following sections, we present results that illustrate the performance of our local random switching strategy.

III. PERFORMANCE OVER THE 2-RELAY NETWORK

For brevity, in this section we use the following substitution for the 2-relay network: $a \leftarrow l_1, b \leftarrow l_2, c \leftarrow r_1, d \leftarrow r_2$.

A. Comparison with C_{opt}^n

1) *Theorems:* For $n = 2$ relays, the linear program for C_{opt}^n can be solved to obtain a closed form expression [2]. Four cases arise depending on whether $a \geq b, c \geq d$ and the value of $\delta = ab - cd$. For space constraints, we will only show the proofs for the case $a \geq b, c \leq d$ and $\delta \leq 0$; the other cases are similar. For this case, we have (from [2]):

$$C_{opt}^2 = \frac{ac(b+d) + bd(a-b)}{(b+d)(a+c-b)}$$

Theorem 3.1: For a 2-relay half-duplex diamond network,

$$\frac{C_{lim}^2}{C_{opt}^2} \geq \frac{3}{4}$$

Proof: For deriving the expression for C_{lim}^2 , we order the relaying states as $\{LL, LT, TL, TT\}$ and let the corresponding time fractions be f_1, f_2, f_3, f_4 . From the previous section, for the limiting (continuous) random local schedule, we have

$$\begin{aligned} f_1 &= \frac{cd}{(a+c)(b+d)}, f_2 = \frac{cb}{(a+c)(b+d)} \\ f_3 &= \frac{ad}{(a+c)(b+d)}, f_4 = \frac{ab}{(a+c)(b+d)} \end{aligned} \quad (7)$$

By definition

$$C_{lim}^2 = \min \left\{ \begin{array}{l} af_1 + af_2 + bf_3, af_1 + (a+d)f_2 + df_4 \\ bf_1 + (b+c)f_3 + cf_4, df_2 + cf_3 + df_4 \end{array} \right\}$$

which, for our case, simplifies to

$$\begin{aligned} C_{lim}^2 &= \frac{a(bc + bd + cd)}{(a+c)(b+d)} & \text{if } a \leq d \\ &= \frac{d(ab + ac + bc)}{(a+c)(b+d)} & \text{if } a > d \end{aligned}$$

In the case $a \leq d$, showing $C_{lim}^2 \geq \frac{3}{4}C_{opt}^2$ is equivalent to

$$4a(a-b+c)(cd+b(c+d)) - 3(a+c)(-b^2d + a(cd + b(c+d))) \geq 0$$

Denoting the l.h.s by ϕ , the following inequalities hold.

$$\phi/c \stackrel{(i)}{\geq} a^2b - 4ab^2 + abc + a^2d - 3abd + 3b^2d + acd$$

$$\stackrel{(ii)}{\geq} -3ab^2 + abc + a^2d - 3abd + 3b^2d + acd$$

$$\stackrel{(iii)}{\geq} -3ab^2 + a^2b^2/d + a^2d - 3abd + 3b^2d + a^2b$$

$$= \frac{1}{d}(a^2(b^2 + d^2 + bd) + a(-3b^2d - 3bd^2) + 3b^2d^2)$$

(i) and (ii) hold as $a \geq b$; (iii) holds because $c \geq ab/d$. The numerator in the last quantity is a quadratic expression in the variable a and its discriminant is

$$\Delta = 9(bd^2 + b^2d)^2 - 12b^2d^2(b^2 + d^2 + bd) = -3b^2(b-d)^2d^2 \leq 0$$

Therefore, $\phi \geq 0$, which establishes our claim. Clearly, equality is attained when a, b, c, d are equal. The proof for $a > d$ follows similarly. ■

For deterministic switching, the worst case ratio drops to $1/2$, as shown below.

Theorem 3.2: For a 2-relay half-duplex diamond network,

$$\frac{C_{det}^2}{C_{opt}^2} \geq \frac{1}{2}$$

Proof: To derive the expression for C_{det}^2 , notice that relay 1 listens for $\frac{c}{a+c}$ of time and relay 2 listens for $\frac{d}{b+d}$ fraction. Depending on which one is larger, we will have the states $\{LL, TL, TT\}$ or $\{LL, LT, TT\}$. For the first case, we have

$$f_1 = \frac{c}{a+c}, f_2 = 0, f_3 = \frac{ad-bc}{(a+c)(b+d)}, f_4 = \frac{b}{b+d}$$

Using these, we can derive

$$\begin{aligned} C_{det}^2 &= \frac{-b^2c + a(cd + b(c+d))}{(a+c)(b+d)} & \text{if } a+c \leq b+d \\ &= \frac{acd + b(-c^2 + ad + cd)}{(a+c)(b+d)} & \text{if } a+c \geq b+d \end{aligned}$$

For the first case, proving our claim is equivalent to showing

$$\left\{ \begin{array}{l} 2(-b^2c + a(cd + b(c+d)))(a+c-b) \\ -(a+c)(ac(b+d) + bd(a-b)) \end{array} \right\} \geq 0$$

Denoting the l.h.s by ϕ , the following inequalities hold.

$$\begin{aligned} \phi/c &= -3ab^2 + 2b^3 - b^2c + a^2d - abd + b^2d + acd \\ &\quad + (a-b)(ab + bc) \end{aligned}$$

$$\stackrel{(i)}{\geq} -3ab^2 + 2b^3 + a^2d - abd + b^2(d-c) + acd$$

$$\stackrel{(ii)}{\geq} a^2(b+d) - a(3b^2 + bd) + 2b^3 = (a-b)(ab + ad - 2b^2)$$

$$\stackrel{(iii)}{\geq} (a-b)(b(a-b) + b^{\frac{1}{2}}(a^{\frac{3}{2}} - b^{\frac{3}{2}}))$$

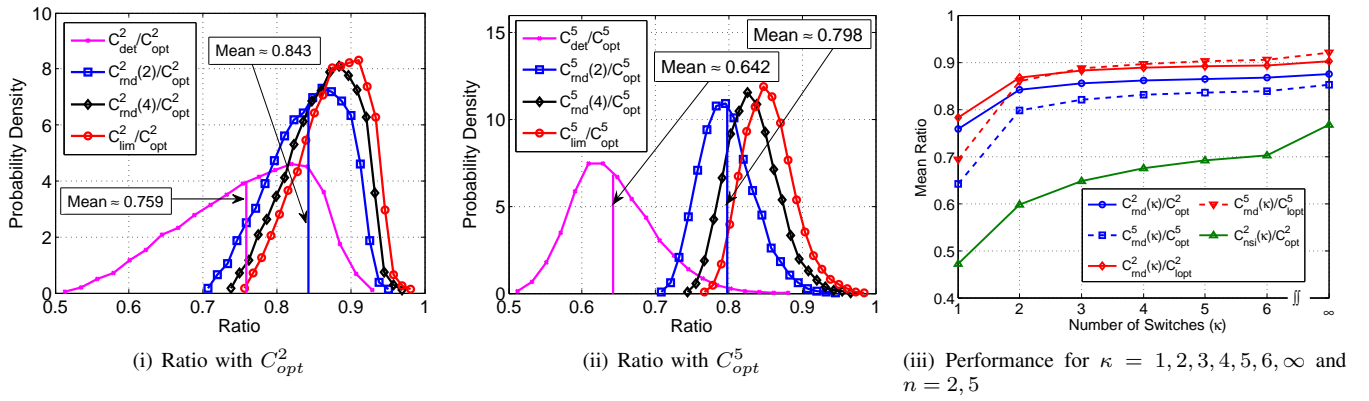


Fig. 4. Numerical evaluations for the 2 and 5 relay network with channel strengths sampled uniformly and independently from $[0, 30]$ dB

Here, (i) is true because $a \geq b$, (ii) is true because $c \geq d$ and $cd \geq ab$ and finally (iii) is true because $d^2 \geq cd \geq ab \implies d \geq \sqrt{ab}$. Again, equality holds when a, b, c, d are equal. The proof for the other three cases is similar. ■

2) *Numerical Evaluation*: The above two theorems show that the worst case performance of randomized switching (in the limit of large κ) is much better than that of deterministic switching. For finite values of $\kappa > 1$, it is difficult to derive general closed form lower bounds on the ratio of $C_{rnd}^2(\kappa)/C_{opt}^2$, though it is likely to lie between the two extremes of 0.5 and 0.75. To investigate its performance, we perform the following simulation. We select the channel strengths for the four links in the network, independently and uniformly at random, in the range $[0, 30]$ dB. For each configuration, we compute the quantities C_{det}^2/C_{opt}^2 , C_{lim}^2/C_{opt}^2 and $C_{rnd}^2(\kappa)/C_{opt}^2$ for $\kappa = 2, 4$ and then plot the p.d.f of these quantities as shown in Fig. 4(i).

The plot shows that there is a **significant jump (about 11%)** in the average performance of $C_{rnd}^2(2)/C_{opt}^2$ (0.843) over C_{det}^2/C_{opt}^2 (0.759). Thereafter, as we increase κ , the average performance slowly saturates to 0.877, which is the mean of C_{lim}^2/C_{opt}^2 . Thus, randomization and increasing the number of switches **increases the average performance by about 15.5%**, a large chunk of which (**approximately 71% of the difference**) is leveraged by using two random switches⁴.

B. Comparison with C_{lopt}^n

C_{lopt}^2 represents the upper bound of the rates achievable by strategies that adhere to the local optimality criterion (4). It is interesting to see how our switching strategies perform with respect to this bound. We reiterate that C_{lopt}^n also encompasses strategies that allow inter-relay communication, and hence is strictly an upper bound for the types of *distributed* scheduling strategies we propose, albeit a tighter one than C_{opt}^n .

For $n = 2$ relays, the linear program for C_{lopt}^n can be solved to obtain a closed form expression. It is as follows

⁴It can be shown that $C_{rnd}^2(2)/C_{opt}^2$ is 0.7 when a, b, c, d are equal. Numerical evidence suggests that this is also the minimum.

$$\begin{aligned} C_{lopt}^2 &= \frac{b(d+a)}{b+d} \text{ if } ab \leq cd, a \leq b \\ &= \frac{a(b+c)}{a+c} \text{ if } ab \leq cd, a \geq b \\ &= \frac{c(a+d)}{a+c} \text{ if } ab \geq cd, c \geq d \\ &= \frac{d(b+c)}{b+d} \text{ if } ab \geq cd, c \leq d \end{aligned}$$

Using this, we prove the following results.

Theorem 3.3: For a 2-relay half-duplex diamond network,

$$\frac{C_{lim}^2}{C_{lopt}^2} \geq \frac{3}{4} \text{ and } \frac{C_{det}^2}{C_{lopt}^2} \geq \frac{1}{2}$$

Proof: Using the expression for C_{lim}^2 for $a \leq d$ (and for $a \geq b, d \geq c$ and $ab \leq cd$) the claim is equivalent to showing

$$\phi = 4a(bc + bd + cd) - 3a(b+c)(b+d) \geq 0$$

The following sequence of inequalities hold

$$\begin{aligned} \phi &= cd + bc + bd - 3b^2 \stackrel{(i)}{\geq} b(c+d-2b) + b(a-b) \stackrel{(ii)}{\geq} 0 \\ (i) \text{ is true because } ab &\geq cd \text{ and } (ii) \text{ is true because } a \geq b \text{ and } c+d \geq 2\sqrt{cd} \geq 2\sqrt{ab} \geq 2b. \end{aligned}$$

Next, using the expression for C_{det}^2 for $a+c \leq b+d$ (and for $a \geq b, d \geq c$ and $ab \leq cd$) the claim is equivalent to showing

$$\phi = 2(-b^2c + a(cd + b(c+d))) - a(b+c)(b+d) \geq 0$$

The following sequence of inequalities hold

$$\begin{aligned} \phi &= -ab^2 + abc - 2b^2c + abd + acd \\ &= bc(a-b) + a(cd - b^2) + b(ad - bc) \stackrel{(i)}{\geq} 0 \\ (i) \text{ is true because } a &\geq b, d \geq c \implies ad \geq bc \text{ and } cd \geq ab \geq b^2. \end{aligned}$$

As in the previous case, it is difficult to obtain closed form bounds for $C_{rnd}^2(\kappa)/C_{lopt}^2$ for finite $\kappa > 1$. When the simulations are repeated for the ratios C_{det}^2/C_{lopt}^2 , C_{lim}^2/C_{lopt}^2 and $C_{rnd}^2(\kappa)/C_{lopt}^2$ for $\kappa = 2, 4$, a trend similar to Fig. 4(i) is observed. Although the worst case ratios are the same as those with respect to C_{opt}^2 , the average performance is better.

IV. PERFORMANCE OVER LARGER NETWORKS

For diamond networks of larger size, the performance trends observed in the previous section are essentially similar with a few interesting caveats. Fig. 4(ii) plots the p.d.f of C_{det}^5/C_{opt}^5 , C_{lim}^5/C_{opt}^5 and $C_{rnd}^5(\kappa)/C_{opt}^5$ for $\kappa = 2, 4$ for random instances of a 5-relay network. In this case, the gain of mean performance going from $\kappa = 1$ to $\kappa = 2$ is significantly more: $0.642 \rightarrow 0.798$ —an **increase of about 24.3%**. This also shows that deterministic switching performs worse for larger number of relays, but even $\kappa = 2$ greatly boosts performance.

In Fig. 4(iii), we plot the mean performance ratios of different schemes as a function of the number of switches for $n = 2, 5$ relays. This plot essentially conveys three messages: (i) C_{lopt}^n is a more useful outer bound than C_{opt}^n when comparing local scheduling strategies; (ii) increasing the number of switches has highly diminishing returns for larger κ and very quick saturation towards the asymptotic value is observed, which is practically important as too many switches can have significant network overhead, and (iii) Local CSI helps over not using any CSI: the performance of $C_{nst}^n(\kappa)$, which incorporates κ random switches *without* respecting (4) performs significantly worse, especially for small κ .

REFERENCES

- [1] A. S. Avestimehr, S. N. Diggavi, and D. N. C. Tse, “Wireless network information flow: a deterministic approach”, *IEEE Trans. on Inform. Theory*, vol. 57(4), pp. 1872–1905, 2011.
- [2] S. Brahma, A. Ozgur and C. Fragouli, “Simple schedules for half-duplex networks”, *IEEE ISIT*, 2012.
- [3] S. Brahma and C. Fragouli, “A simple relaying strategy for diamond networks”, *IEEE ISIT*, 2014.
- [4] M. Cardone, D. Tuninetti, R. Knopp and U. Salim, “On the Gaussian Half-Duplex Relay Channel”, *IEEE Trans. on Inf. Theory*, vol. 60(5), pp. 2542–2562, 2014.
- [5] M. Jorgovanovic, M. Weiner, D. Tse, B. Nikolic, I.-H. Wang and V. Nagpal, “Relay scheduling and interference cancellation for Quantize-Map-and-Forward cooperative relaying”, *IEEE ISIT*, 2013.
- [6] G. Kramer, “Models and Theory for Relay Channels with Receive Constraints”, *42nd Allerton Conf.*, 2004.
- [7] C. Nazeroglu, A. Ozgur and C. Fragouli, “Wireless network simplification: the Gaussian N-relay diamond network”, *IEEE ISIT*, 2011.