Cross-Layer Optimizations for Intersession Network Coding on Practical 2-Hop Relay Networks

Chih-Chun Wang, Abdallah Khreishah
Center of Wireless Systems and Applications (CWSA)
School of ECE, Purdue University

Ness B. Shroff
Departments of ECE and CSE
The Ohio State University

***Abstract*—Full characterization of Intersession Network Coding (INC), i.e., coding across multiple unicast sessions, is notoriously challenging. Nonetheless, the problem can be made tractable when considering practical constraints that restrict the types of INC schemes of interest. This paper characterizes the INC capacity of *2-session wireless 2-hop relay networks* with a packet erasure channel model and a round-based feedback schedule motivated by the usage of "reception reports" in practical protocols such as COPE. The capacity regions are formulated as linear programming problems, which admit simple concatenation with other competing techniques such as opportunistic routing (OpR), and cross-layer (CL) optimization. Extensive numerical evaluation is conducted on 1000 random topologies, which compares and quantifies the throughput benefits of INC, OpR, and CL, and their arbitrary combinations. The results show that by combining all three techniques of INC, OpR, and CL, the throughput of a wireless 2-hop relay network can be improved by 100–500% over the benchmark single-path routing solution depending on the number of sessions to be coded together.**

I. INTRODUCTION

Network coding (NC) has promised strict throughput improvement over non-coding solutions [1], [2], [3]. Although the benefits of NC within a single multicast session is elegantly characterized by the min-cut/max-flow theorem, to realize the throughput benefits for multiple unicast sessions, NC must be performed across different sessions, termed intersession network coding (INC). Full characterization of INC is notoriously challenging and existing investigation shows that the existence of a closed-form solution is highly unlikely for general network settings [4]. Nonetheless, the problem can be made tractable when considering practical constraints that restrict the types of INC schemes of interest. In this work, we are interested in a special wireless 2-hop relay network that is motivated by empirical INC protocols such as COPE [3].

Wireless mesh networks (WMN) have been considered as a cost-effective solution to the *last mile problem*. Among many techniques that improve the throughput of a WMN, the following three have demonstrated significant throughput enhancement: 1-hop INC (the COPE protocol [3]); Intrasession-coding-based opportunistic routing (OpR) (the MORE protocol [1], [5]) that takes advantage of the spatial diversity of wireless channels by opportunistically deciding the packet route; Cross-layer (CL) optimization [6] that allows distributed resource allocation that jointly takes into account the traffic demands and the interference constraints of wireless channels.

This work was supported by the NSF Grant CNS-0905331.

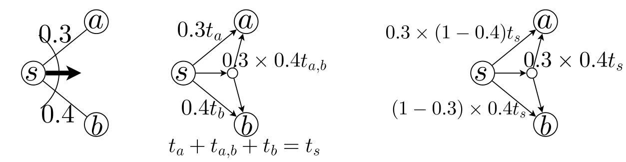

(a) Wireless broadcast (b) PrR-based averaging (c) Opportunistic-scheduling-based averaging

Fig. 1. Two different wireless-to-wireline conversion methods.

Most existing results on joint INC and CL are based on converting the wireless networks to the corresponding wireline counterpart (see [7] and the reference therein). Two conversion methods are commonly used: the predetermined-routing-based (PrR-based) conversion and the opportunistic-scheduling-based (OpS-based) conversion. Take Fig. 1(a) for example, source s broadcasts a packet wirelessly and nodes a and b can hear the packet with probabilities 0.3 and 0.4, respectively and independently. The PrR-based conversion predetermines the packet route and uses error control coding or Automatically Repeat reQuest (ARQ) to average the random overhearing events into deterministic average rates as in Fig. 1(b). The OpS-based conversion assumes that s knows before transmission which of the receivers a and b will successfully overhear the to-be-transmitted packet. The deterministic rates to a, to b, and to (a, b) simultaneously, are thus expressed as a percentage of t_s, the amount of time for which s is broadcasting, as described in Fig. 1(c). Neither PrR nor OpS model the practical schemes closely, as random overhearing is not considered in the PrR conversion while the assumption of perfectly knowing the channel condition is too optimistic for WMNs.[1] In this work we model the wireless channel by a more realistic *packet erasure channel* (PEC) and characterize the capacity of a 2-session 2-hop relay network under round-based feedback.

II. THE SETTING

We use X (and sometimes Y) to denote a symbol in $\mathsf{GF}(q)$. A 1-to-K packet erasure channel (PEC) takes an input $X \in \mathsf{GF}(q)$ and outputs a K-dimensional vector in $(\{X\} \cup \{*\})^K$,

[1] For a single multicast session, random linear network coding achieves the capacity of OpS even without the knowledge about the channel condition [8]. Nonetheless, the capacity of INC when not knowing the channel condition is strictly smaller than the capacity of OpS.

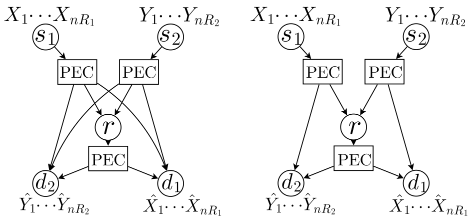

(a) With OpR (b) Without OpR

Fig. 2. Illustration of 2-session relay networks.

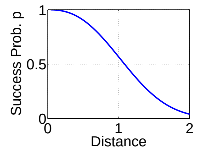

Fig. 3. The overhearing probability versus the distance.

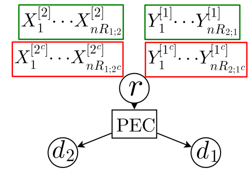

Fig. 4. A broadcast PEC problem with side information.

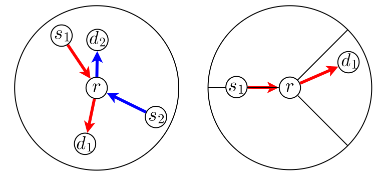

(a) (b)

Fig. 5. (a) An example of random node placement for $M = 2$ sessions. (b) The constraint on the topological relationship between s_i and d_i.

where the k-th coordinate being "$*$" denotes that the input symbol X is erased for the k-th receiver. We consider only i.i.d. PECs of which the erasure pattern is independent for each channel usage. For example, a 1-to-2 PEC can be described by four parameters $p_{s;12}$, $p_{s;12^c}$, $p_{s;1^c2}$, and $p_{s;1^c2^c}$, which denote the probabilities that X is received by both receivers (rx) 1 and 2, by only rx 1, by only rx 2, and by neither rx 1 nor 2. This notation can be easily extended to the marginal success probability $p_{s;1} \triangleq p_{s;12} + p_{s;12^c}$ and to the union $p_{s;1\cup 2} \triangleq p_{s;12} + p_{s;12^c} + p_{s;1^c2}$.

A 2-session relay network is described in Fig. 2(a), in which source s_1 would like to send nR_1 packets $X_1, \cdots, X_{nR_1}$ to destination d_1, and s_2 would like to send nR_2 packets $Y_1, \cdots, Y_{nR_2}$ to d_2. r is a relay node. Each of s_1, s_2, and r can use the corresponding PEC n times, respectively, and we are interested in the largest achievable rate pair (R_1, R_2) that guarantees decodability at d_1 and d_2 with close-to-1 probability for sufficiently large n and the finite field size q.

To model the "reception report" suggested in COPE, we enforce the following sequential, round-based feedback schedule: Each of s_1 and s_2 transmits n symbols, respectively. After the transmission of $2n$ symbols, two reception reports are sent from d_1 and d_2, respectively, back to the relay r so that r knows which packets have successfully arrived which destinations. After the reception reports, no further feedback is allowed and the relay r has to make its own decision how to use the available n PEC usages to guarantee decodability at d_1 and d_2. In our setting, we also assume that the success probability parameters of all PECs and all the coding operations are known to all nodes. The only unknown part is the values of the X and Y symbols. For the purpose of illustration, a simplified network setting is also depicted in Fig. 2(b), in which the packets sent by s_i will not be overheard by the 2-hop-away destination d_i. In this simplified setting, the question thus becomes given the following parameters: $p_{s_1;rd_2}, p_{s_1;rd_2^c}, p_{s_1;r^cd_2}$, and $p_{s_1;r^cd_2^c}$; $p_{s_2;rd_1}, p_{s_2;rd_1^c}, p_{s_2;r^cd_1}$, and $p_{s_2;r^cd_1^c}$; $p_{r_1;d_1d_2}$, $p_{r;d_1d_2^c}$, $p_{r;d_1^cd_2}$, and $p_{r;d_1^cd_2^c}$, what is the maximum achievable rate under the round-based feedback model. For future reference, we say Fig. 2(a) admits OpR as the packets can be overheard by the two-hop destinations while Fig. 2(b) does not admit OpR. Since the settings are symmetric, we sometimes assume that $p_{r;d_1} \geq p_{r;d_2}$ which can be achieved by relabelling the sessions.

III. The Capacity Results

Consider Fig. 2(b) and the scenario in which s_1 and s_2 have finished transmission and the reception reports have been sent to r. The question now becomes a broadcast PEC problem with side information (SI) as depicted in Fig. 4. That is, the packets $X_1^{[2]}, \cdots, X_{nR_{1;2}}^{[2]}$ have been overheard by d_2 and the packets $X_1^{[2^c]}, \cdots, X_{nR_{1;2^c}}^{[2^c]}$ have not. Similarly, $Y_1^{[1]}, \cdots, Y_{nR_{2;1}}^{[1]}$ have been overheard by d_1 and $Y_1^{[1^c]}, \cdots, Y_{nR_{2;1^c}}^{[1^c]}$ have not. $X^{[2]}$ and $Y^{[1]}$ packets can later serve as SI when decoding at d_2 and d_1, respectively. The total rates for each session are $R_1 = R_{1;2} + R_{1;2^c}$ and $R_2 = R_{2;1} + R_{2;1^c}$. We then have the following theorem:

Theorem 1: For any rate vectors $(R_{1;2}, R_{1;2^c}, R_{2;1}, R_{2;1^c})$ for the broadcast PEC with SI in Fig. 4 and assuming that $p_{r;d_1} \geq p_{r;d_2}$, one can communicate the values of all X and Y packets to d_1 and d_2, respectively, within n channel usage if and only if $(R_{1;2}, R_{1;2^c}, R_{2;1}, R_{2;1^c})$ satisfies

$$R_{1;2} + R_{1;2^c} + R_{2;1^c} \leq p_{r;d_1} \quad (1)$$

$$R_{2;1} + R_{2;1^c} + \frac{p_{r;d_2}}{p_{r;d_1}} R_{1;2^c} \leq p_{r;d_2}. \quad (2)$$

Remark: The broadcast capacity of *Gaussian channel* with SI has been considered in many papers (see [9] as a representative work). In addition to the difference of the settings of PECs and Gaussian channels, this paper will also consider the best coding strategy at sources s_1 and s_2 and characterize the capacity as a linear programming problem. It is worth noting that for 3-session Gaussian broadcast channels with general SI, the capacity remains an open problem. On the other hand, in [10] we show that the capacity of 3-session broadcast PEC can be computed in a similar way as outlined in this paper. Recently, the capacity of 2-session broadcast PECs with instant packet-by-packet feedback (but without SI) is studied in [11].

Sketch of the proof of Theorem 1:

Achievability: For any vector $(R_{1;2}, R_{1;2^c}, R_{2;1}, R_{2;1^c})$ satisfying (1) and (2), perform the following 2-staged coding scheme sequentially.

Stage 1: Whenever we can use the broadcast PEC, we mix the packets of the three groups: $X_1^{[2]}$ to $X_{nR_{1;2}}^{[2]}$, $Y_1^{[1]}$ to $Y_{nR_{2;1}}^{[1]}$, and $Y_1^{[1^c]}$ to $Y_{nR_{2;1^c}}^{[1^c]}$ by random linear network coding (RLNC) [8] and generate one outgoing symbol. Repeat this random packet generation until we have sent out the following amount of randomly generated packets:

$$\max\left(\frac{nR_{1;2}+nR_{2;1^c}}{p_{r;d_1}},\frac{nR_{2;1}+nR_{2;1^c}}{p_{r;d_2}}\right). \quad (3)$$

Stage 2: Whenever we can use the broadcast PEC, we mix the packets of $X_1^{[2^c]}$ to $X_{nR_{1;2^c}}^{[2^c]}$ by RLNC and generate one outgoing symbol. Repeat this random packet generation until we have sent out the following amount of packets:

$$\frac{nR_{1;2^c}}{p_{r;d_1}}. \quad (4)$$

The decodability proof is straightforward as from rx 1's perspective, d_1 has received a sufficiently large amount of coded packets to decode the information packets in the three groups: $X^{[2]}$, $X^{[2^c]}$, and $Y^{[1^c]}$ (since d_1 has already heard the $Y^{[1]}$ packets). Similarly, d_2 simply ignores the packets transmitted in Stage 2 and use the packets in Stage 1 to decode the desired two groups of packets $Y^{[1^c]}$ and $Y^{[1]}$. It remains to show that one can finish Stages 1 and 2 within the allowed n channel usages, which is equivalent to proving that (3) plus (4) is no larger than n. By noticing that (3)+(4)$\leq n$ is satisfied for any $(R_{1;2}, R_{1;2^c}, R_{2;1}, R_{2;1^c})$ vector satisfying (1) and (2), the proof of the achievability is thus complete.

The converse part can be proven by deriving the lower bound on the amount of interference that session 1 will impose on destination d_2. (Stage 2 of the achievability scheme carefully limits the amount of interference to the minimal possible amount, $\frac{p_{r;d_2}}{p_{r;d_1}}R_{1;2^c}$, one can hope to achieve.) ∎

Theorem 1 can be used to derive the 2-session relay network capacity by noting that if we let each of s_1 and s_2 send n packets by RLNC, we can maximize the overheard SI at d_2 and d_1, respectively. Based on this observation, the capacity of a 2-session relay network becomes:

Theorem 2: Assuming that $p_{r;d_1} \geq p_{r;d_2}$, for a 2-session relay network without OpR (Fig. 2(b)), the rate pair (R_1, R_2) is achievable if and only if they satisfy

$$R_1 \leq \min\left(p_{s_1;r}, p_{r;d_1} - (R_2 - p_{s_2;d_1})^+\right)$$
$$R_2 \leq \min\left(p_{s_2;r}, p_{r;d_2} - \frac{p_{r;d_2}}{p_{r;d_1}}(R_1 - p_{s_1;d_2})^+\right),$$

where $(\cdot)^+$ is the projection to a non-negative number.

In the following, we generalize Theorem 2 for OpR (Fig. 2(a)) and for the case in which s_1, s_2, and r have different amount of available time slots: nt_{s_1}, nt_{s_2}, and nt_r, respectively.

Theorem 3: Assuming that $p_{r;d_1} \geq p_{r;d_2}$ and variable transmission time nt_{s_1}, nt_{s_2}, and nt_r, for a 2-session relay network with OpR (Fig. 2(b)), the rate pair (R_1, R_2) is achievable if and only if they satisfy

$$R_1 \leq t_{s_1}p_{s_1;d_1} + \min(t_{s_1}p_{s_1;rd_1^c}, t_r p_{r;d_1} - (R_2 - t_{s_2}p_{s_2;d_1\cup d_2})^+) \quad (5)$$

$$R_2 \leq t_{s_2}p_{s_2;d_2} + \min(t_{s_2}p_{s_2;rd_2^c}, t_r p_{r;d_2} - \frac{p_{r;d_2}}{p_{r;d_1}}(R_1 - t_{s_1}p_{s_1;d_1\cup d_2})^+). \quad (6)$$

IV. Comparison to Other Schemes

A. The Variants of Optimal INC

The optimal capacity of INC in Theorem 3 is cast as a linear-programming problem and contains three orthogonal components: (i) INC via the new coding-based argument in Theorem 1, (ii) OpR via modelling the overhearing opportunities as the 1-to-3 PECs of s_1 and s_2 in Fig. 2(a), and (iii) CL design via variable transmission time t_{s_1}, t_{s_2}, and t_r. One can thus quantify the improvement of each component. For example, with CL we are allowed to optimally choose t_{s_1} to t_r with $t_{s_1}+t_{s_2}+t_r = 1$ assuming that the node exclusive model is used and s_1, s_2, and r cannot transmit simultaneously. If without CL, we simply set $t_{s_1}=t_{s_2}=t_r=1/3$ for fair time-sharing between the sources and the relay. If OpR is not allowed, we simply use the extended version of Theorem 2 with added time variables t's. We denote the four variants of the INC schemes, depending on whether we use OpR or CL or not, by (INC, OpR, CL), (INC, OpR, ×CL), (INC, ×OpR, CL), (INC, ×OpR, ×CL).

B. The Baseline Non-Coding Scheme and Its Variants

If we use a single-hop routing scheme without OpR, then the capacity region can be described as

$$R_i \leq \min(t_{s_i}p_{s_i;r}, t_r^{[i]}p_{r;d_i}), \forall i \in \{1,2\}, \quad (7)$$

where $t_r^{[i]}$ is the amount of time that r sends the session-i packets. If we allow OpR, for which the packets can directly reach its 2-hop-away destinations, the capacity region becomes

$$R_i \leq t_{s_i}p_{s_i;d_i} + \min(t_{s_i}p_{s_i;rd_i^c}, t_r^{[i]}p_{r;d_i}), \forall i \in \{1,2\}. \quad (8)$$

Whether to use CL or not will impose different time sharing constraints on t_{s_i} and $t_r^{[i]}$ in a similar way as discussed in Section IV-A. (7) and (8) can thus be used to generate four variants of non-coding schemes, depending on whether we use OpR or CL or not. We denote the four variants by (×INC, OpR, CL), (×INC, OpR, ×CL), (×INC, ×OpR, CL), (×INC, ×OpR, ×CL). The simplest scheme (×INC, ×OpR, ×CL) will be used as the baseline of all our numerical comparison.

C. Schemes for More Than 2 Sessions

The ×INC schemes described in (7) to (8) can be easily generalized to the case when there are $M > 2$ sessions, for which (7) and (8) hold for $i = 1$ to M; and the time sharing constraints without CL become

$$\forall i \in \{1,\cdots,M\} \quad t_{s_i} = \frac{1}{M+1}, \text{ and } \sum_{i=1}^{M} t_r^{[i]} = \frac{1}{M+1},$$

$$\text{or when CL is used:} \quad \sum_{i=1}^{M} t_{s_i} + \sum_{i=1}^{M} t_r^{[i]} = 1.$$

Theorem 3 describes the INC-based schemes for $M = 2$. In [10], Theorem 3 has been generalized for $M = 3$ and new upper and lower bounds are provided for general $M > 3$, which are empirically tight for most PEC channel parameters when $M \leq 5$. In our simulation results, we also include the curves for $M = 3$ to 5 for better comparison.

D. Practical, Low-Complexity Schemes for $M > 2$

Optimally performing INC over multiple sessions, say $M = 5$, is an exponentially complicated process as it requires carefully taking advantages of the various overheard portions of the data stream. In this paper, we thus consider some suboptimal but more practical schemes as follows.

1) Multipath Routing: Unlike the single-path routing, each source s_i can choose whether to directly send the packets to the 2-hop-away destination at the cost of higher drop rate (since $p_{s_i;d_i}$ is generally less than $p_{s_i;r}$), or to send the packets through the relay r. The corresponding capacity is thus characterized by the following linear program:

$$R_i \leq t_{s_i}^{[\text{direct}]} p_{s_i;d_i} + \min\left(t_{s_i}^{[\text{relay}]} p_{s_i;r}, t_r^{[i]} p_{r;d_i}\right).$$

With CL, we enforce that $\sum_{i=1}^{M} t_{s_i}^{[\text{direct}]} + t_{s_i}^{[\text{relay}]} + t_r^{[i]} \leq 1$.

2) 2-INC and 3-INC: To reduce complexity of INC over multiple sessions, we can limit the number of sessions to be coded together and use linear programming to optimally identify which sessions to encode. For example, 2-INC denotes a scheme that allows INC only over two sessions while optimally allocating the corresponding time-sharing percentage. The following linear constraints describe the capacity region of 2-INC without OpR.

$$R_i \leq \sum_{j:j\neq i} R_i^{\{i,j\}} \quad \forall i \in \{1, \cdots, M\}$$

$$\forall i, j, i \neq j, \quad R_i^{\{i,j\}} \leq \min\Big(t_{s_i}^{\{i,j\}} p_{s_i;r},$$

$$t_r^{\{i,j\}} p_{r;d_i} - \min\left(1, \frac{p_{r;d_i}}{p_{r;d_j}}\right)\left(R_j^{\{i,j\}} - t_{s_j}^{\{i,j\}} p_{s_j;d_i}\right)^+\Big),$$

where $\{i, j\}$ is an unordered pair of distinct indices with $1 \leq i \neq j \leq M$. For example, both notations $t_{s_1}^{\{1,2\}}$ and $t_{s_1}^{\{2,1\}}$ refer to the same variable. There are $\binom{M}{2}$ different unordered $\{i, j\}$ pairs. Therefore, there are totally $2 \cdot \binom{M}{2}$ variables of $R_i^{\{i,j\}}$, $2 \cdot \binom{M}{2}$ variables of $t_{s_i}^{\{i,j\}}$, $\binom{M}{2}$ variables of $t_r^{\{i,j\}}$, and M variables of R_i. With CL, we require the sum of all $t_{s_i}^{\{i,j\}}$ and $t_r^{\{i,j\}}$ to be ≤ 1. Similar formulation can be made by combining the 3-session INC capacity in [10] to derive a 3-INC scheme for general $M \geq 3$.

3) A multicast-based scheme in [2]: In Section 4 of [2] an INC scheme, denoted as the CCH scheme, is proposed based on using several multicast sessions (totally $2^M - 1$ multicast sessions) to serve the need of M unicast sessions. Each multicast session is indexed by $C \in 2^{\{1,\cdots,M\}}$, which contains all sessions participate in this multicast session. The corresponding capacity region is described by the following linear constraints:

$$R_i \leq \sum_{C:i\in C} R_i^C, \quad \forall i \in \{1, \cdots, M\}$$

$$R_i^C \leq t_{s_i}^C p_{s_i;r}, \quad \forall i \in \{1, \cdots, M\}, \forall C \in 2^{\{1,\cdots,M\}}$$

$$R_i^C + \sum_{j:j\in A\backslash i} R_j^C \leq t_r^C p_{r;d_i} + \sum_{j:j\in A\backslash i} t_{s_j}^C p_{s_j;d_i},$$

$$\forall i \in \{1, \cdots, M\}, \forall C \in 2^{\{1,\cdots,M\}}, \forall A \subseteq C,$$

where R_i^C, $t_{s_i}^C$, and t_r^C are the rate of session i, and the time allocations of s_i and r for the multicast session C, respectively. With CL, we require the sum of all $t_{s_i}^C$ and t_r^C to be ≤ 1.

E. Remark on the practical COPE protocol

COPE has many features that are designed for practical implementation and quantifying its performance remains an open problem. The main differences between our INC formulation and COPE are: (i) COPE requires instant decoding after receiving each packet while INC allows delayed decoding after receiving the entire batch. (ii) COPE suggests using *reception reports* broadcast at a random frequency while our formulation focuses on strict feedback scheduling. The former difference reduces the achievable rate of COPE when compared to INC, while the latter difference may enhance or reduce the rate of COPE due to the random nature of the feedback schedule.

V. Numerical Experiments

We compare the throughput of various schemes on randomly constructed 2-hop relay networks. Our construction starts from a unit circle with the relay r placed at the center. We then uniformly randomly place M source nodes s_i and M destination nodes d_i in the circle (see Fig. 5(a)). The only condition we impose is that for each (s_i, d_i) pair, d_i must be in the 90-degree pie area opposite to s_i (see Fig. 5(b)). For each randomly constructed network, we use the Euclidean distance between each node to determine the overhearing probability. More explicitly, for any two nodes separated by distance D we use the Rayleigh fading model to decide the overhearing probability: $p = \int_{T^*}^{\infty} \frac{2x}{\sigma^2} e^{-\frac{x^2}{\sigma^2}} dx$, where we choose $\sigma^2 \triangleq \frac{1}{(4\pi)^2 D^\alpha}$, the path loss order $\alpha = 2.5$, and the decodable SNR threshold $T^* = 0.06$. Fig. 3 plots the overhearing probability p versus distance D. We assume that the overhearing event is independent among different receivers.

For each randomly generated network, we compute the overhearing probabilities and use the corresponding linear constraints on the time-sharing variables t's and the rate variables R's to compute the achievable rate of each scheme. A common linear objective function is used for all different schemes, which is described in the following:

$$\max_{t\text{'s}, R\text{'s}, \beta} \quad \sum_{i=1}^{M} R_i \tag{9}$$

$$\text{subject to} \quad R_i = \beta \min\left(p_{s_i;r\cup d_i}, p_{s_i;d_i} + p_{r;d_i}\right), \quad \forall i, \tag{10}$$

linear ineq. for the scheme of interest.

Namely, we are interested in maximizing the sum rate while (10) enforces that the rates are proportional to the unicast capacity with fair scheduling (assuming all other sessions $j \neq i$ remain silent). Constraint (10) thus ensures fairness proportional to the inherent interference-free capacity.

Given a randomly generated network, the achievable sum rates are computed for all the schemes. We then repeat this computation for 1000 randomly generated networks. Let $\zeta^*_{\text{scheme},k}$ denote the achievable sum rate of the given scheme for the k-th randomly chosen topology. We are

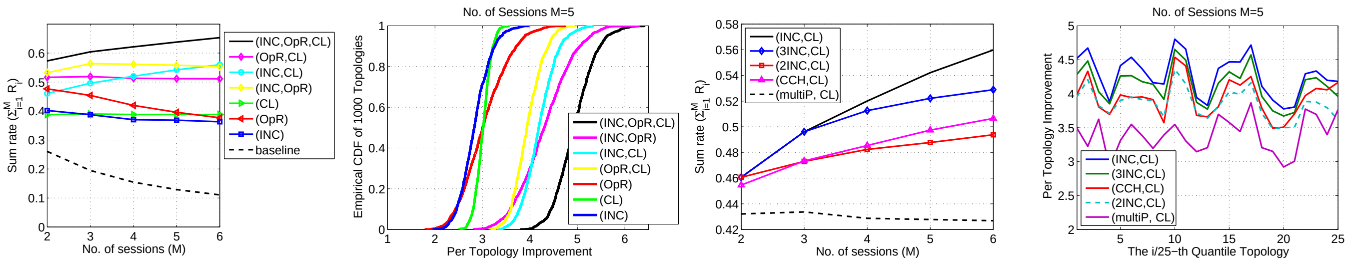

(a) (b) (c) (d)

Fig. 6. (a) The average sum rates for different schemes combining INC, OpR, and CL. (b) The cumulative distribution function (CDF) of per topology improvement (PTI). The schemes in the legend corresponds to the curves from the rightmost to the leftmost when focusing on the intersection of CDF=0.6. (c) The average sum rates for different practical schemes without OpR, (d) The PTI for 25 representative topologies. These 25 topologies have their sum rates being the 25-th quantile points when sorting the sum rates in the ascending order. The list of schemes in the legend corresponds to the curves from the topmost to the bottommost when focusing on the performance of of the 15/25-th quantile point.

interested in the following two performance metrics. The average sum rate over 1000 topologies $\frac{1}{1000}\sum_{k=1}^{1000}\zeta^*_{\text{scheme},k}$ and per topology improvement $\triangleq \frac{\zeta^*_{\text{scheme},k}}{\zeta^*_{\text{baseline},k}}$.

A common baseline scheme, the vanilla ($\times$INC, $\times$OpR, $\times$CL) scheme as described in Section IV-B, is used when computing the Per Topology Improvement (PTI). The above procedure is then repeated for different M values.

Fig. 6(a) plots the sum rates of different schemes with respect to (w.r.t.) the number of sessions. To satisfy the proportional fairness constraint in (10), the sum rate of the baseline scheme decreases w.r.t. M due to the increased congestion at relay. On the other hand, schemes using CL and INCs successfully mitigate the congestion as their rum rates remain flat (even after satisfying (10)) for large M. When we jointly incorporate all three techniques: INC, OpR, and CL, not only the congestion is resolved but sum rate also increases w.r.t. M by taking full advantage of the spatial diversity, which mounts to x6.5 improvement over the baseline single-hop routing scheme when $M = 6$. We also plot the empirical cumulative distribution function (CDF) of PTI for the 1000 random topologies with $M = 5$. The (INC,OpR,CL) scheme achieves x3.8 to x6.5 throughput improvement when compared to the baseline scheme.

We also compare the practical schemes discussed in Section IV-D. Fig. 6(c) plots the average sum rates versus M. Among all practical schemes, 3-INC realizes the largest percentage of the gains of the optimal INC scheme. 2-INC schemes has similar performance to the CCH scheme for $M = 2$–5. The 2-INC scheme has lower complexity as it considers only $\binom{M}{2}$ pairs of sessions, while the CCH scheme jointly considers all $2^M - 1$ different multicast sessions. All INC schemes do not allow direct s_i to d_i communication (i.e., without OpR) but can still outperform the multipath-routing schemes with direct (s_i, d_i) links by an additional 10–20% for $M = 5$. Fig. 6(d) plots the PTI for 25 representative topologies whose sum rates are the $i/25$-th quantile points when sorting the 1000 topologies according to the ascending order of the sum rates of the baseline scheme. As shown in Fig. 6(d), the practical INC schemes and the optimal INC scheme demonstrate uniform PTI (3.5–4.7) for all topologies. In general, the more congested the original topology is, the higher PTI we have when using INC schemes (other than the multipath-routing scheme) to resolve congestion, which is illustrated by the negative trend of the PTI when moving from the 1st to the 25-th quantile point topology.

VI. Conclusion

In this paper, we have constructed and compared various inter-session coding schemes for practical 2-hop relay networks both theoretically and numerically, which have demonstrated significant throughput benefits (x3.8–x6.5) when compared to the baseline single-path routing solution.

References

[1] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, "Trading structure for randomness in wireless opportunistic routing," in *Proc. ACM Special Interest Group on Data Commun. (SIGCOMM)*. Kyoto, Japan, August 2007.

[2] T. Cui, L. Chen, and T. Ho, "Energy efficient opportunistic network coding for wireless networks," in *Proc. 27th IEEE Conference on Computer Communications (INFOCOM)*. Phoenix, USA, April 2008.

[3] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, "XORs in the air: Practical wireless network," in *Proc. ACM Special Interest Group on Data Commun. (SIGCOMM)*, 2006.

[4] R. Dougherty, C. Freiling, and K. Zeger, "Linear network codes and systems of polynomial equations," *IEEE Trans. Inform. Theory*, vol. 54, no. 5, pp. 2303–2316, May 2008.

[5] D. Koutsonikolas, C.-C. Wang, and Y. Hu, "CCACK: Efficient network coding based opportunistic routing through cumulative coded acknowledgments," in *Proc. 29th IEEE Conference on Computer Communications (INFOCOM)*. San Diego, USA, March 2010.

[6] X. Lin, N. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Selected Areas of Commun.*, vol. 24, no. 8, pp. 1452–1463, August 2006.

[7] A. Khreishah, C.-C. Wang, and N. Shroff, "Cross-layer optimization for wireless multihop networks with pairwise intersession network coding," *IEEE J. Selected Areas of Commun.*, vol. 27, no. 5, pp. 606–621, June 2009.

[8] T. Ho, M. Médard, R. Koetter, D. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Trans. Inform. Theory*, vol. 52, no. 10, pp. 4413–4430, October 2006.

[9] Y. Wu, "Broadcasting when receivers know some messages a priori," in *Proc. IEEE Int'l Symp. Inform. Theory*. Nice, France, June 2007, pp. 1141–1145.

[10] C.-C. Wang and N. Shroff, "On the capacity region of 1-hop intersession network coding — a broadcast packet erasure channel analysis with general message side information," to be submitted to ISIT 2010.

[11] L. Georgiadis and L. Tassiulas, "Broadcast erasure channel with feedback — capacity and algorithms," in *Proc. 5th Workshop on Network Coding, Theory, & Applications (NetCod)*. Lausanne, Switzerland, June 2009, pp. 54–61.