

Throughput Optimal Network Coding and Scheduling in Wireless Networks

Prasanna Chaporkar and Alexandre Proutiere

Abstract

Recently, network coding (NC) emerged as a promising technology for significantly improving throughput and energy efficiency of wireless networks, even for unicast communication. Often, NC schemes are designed as an autonomous layer, independent of the underlying Phy and MAC capabilities and algorithms. Consequently, these schemes are greedy, in the sense that all opportunities of broadcasting combinations of packets are exploited. *We demonstrate that this greedy design principle may in fact reduce the network throughput.* This begets the need for adaptive NC schemes. We further show that designing appropriate MAC scheduling algorithms is critical for achieving the throughput gains expected from NC. In this paper, we propose a general framework to develop optimal and adaptive joint NC and scheduling schemes. Optimality is shown for various Phy and MAC constraints. We apply this framework to two different NC architectures: COPE, a scheme recently proposed in [11], and XOR-Sym, a new scheme we present here. XOR-Sym is designed to achieve a lower implementation complexity than that of COPE, and yet to provide similar throughput gains.

Index Terms

Multi-hop wireless networks, network coding, throughput optimality

I. INTRODUCTION

Recently, network coding (NC) emerged as a promising technology for designing power efficient and scalable schemes that provide optimized usage of the available bandwidth. Our

P. Chaporkar is with Indian Institute of Technology Bombay, Mumbai, India. A. Proutiere is with Microsoft Research, Cambridge, UK. (e-mails: chaporkar@ee.iitb.ac.in, alexandre.proutiere@microsoft.com).

Parts of this paper are presented in MOBICOM '07.

aim is to investigate possible performance gains through NC and the optimal way of using inter-session NC for unicast communication in multi-hop wireless networks.

Though NC was first applied mainly in the context of multicast in wired networks [1], [13], and subsequently in wireless networks [16], [20], it is found to be particularly amicable for enhancing the throughput (the number of packets delivered to the destination per unit time) of wireless networks even for unicast applications [10], [11], [14], [22], [25], [28]. This is mainly due to the broadcast property of wireless channel, meaning that a transmission from a node can potentially be intercepted by all its neighbors.

The throughput gain via NC in case of unicast sessions is typically illustrated using the network shown in Figure 1. Without NC, four transmissions are required to deliver one packet from each of the sessions. Thus, a throughput λ is achievable if and only if $\lambda \leq 1/4$, i.e., if $\lambda \leq (>, \text{resp.}) 1/4$, then there exists (does not exist, resp.) a scheduling scheme that arbitrates transmissions in various slots such that the throughput λ is provided to each of the sessions. Now, with NC, m XORs two packets, one from each session, and then broadcasts the XOR-ed packet. Now, nodes a and b recover the desired packet by XOR-ing the received packet from m with their own packet. Thus, only 3 transmissions are required to deliver one packet from each of the sessions. Clearly, λ is achievable iff $\lambda \leq 1/3$. The throughput gain of NC is therefore $4/3$ in this example.

The promise of potential throughput gain has instigated significant research in designing efficient NC schemes for unicast communication in wireless networks. Following are the two key features of the schemes proposed in the literature: (i) They advocate the use of NC each time an opportunity to combine and broadcast packets is available. Indeed, the schemes are designed to increase the number of NC opportunities through better routing [25] and through opportunistic listening [11]. Distributed algorithms are also designed to identify certain structures in the network topology so as to determine and exploit most of the NC opportunities [6], [19], [29], [7], [21]. (ii) Network coding and scheduling schemes are designed separately. We advocate caution in using these features. The main motivation of this paper stems from the following observation regarding the schemes with at least one of these features: *Systems with NC may have smaller throughputs than those without it.*

This observation may seem counter-intuitive as previous work shows that one can only gain by using NC, and the gain can only increase if more opportunities to combine packets are used. We show that if NC is used each time an opportunity arises or if the scheduling scheme does

not account for NC, then the system throughput may be smaller than that achieved without NC. This calls for a joint design of scheduling and NC strategies. This paper aims at developing a framework that enables this joint design. Specifically, our contributions are as follows:

- We first explain why NC can deteriorate throughput when it is not jointly designed with scheduling, or when NC opportunities are all exploited (see Section II).
- We then propose a general framework that allows us to characterize the throughput region (the set of achievable throughputs of the various sessions) of networks with NC, and to design optimal and adaptive joint NC and scheduling schemes. The schemes are optimal as they provide the required throughputs, whenever possible. The schemes are adaptive as they take the scheduling and NC decisions based on the current system state only, and do not require the knowledge of channel and arrival statistics a priori (see Section IV).
- We show how our framework can be applied to COPE, a NC scheme recently proposed for unicast sessions in wireless networks [11] (see Section V).
- We also propose a novel NC scheme, XOR-Sym, which exhibits a lower computational complexity than that in COPE. Under XOR-Sym, packets have to be decoded at their destinations only, not at intermediate nodes (see Section VI). In spite of this additional constraint, we show that XOR-Sym and COPE may provide similar throughput gains (see Section VII).

Because of the space constraints, proofs for all the results will be presented in [4]. However, in this submission, we include the proofs in the appendix to facilitate the review process.

II. CAN NC DETERIORATE THE NETWORK THROUGHPUT?

Here, we illustrate, using three representative examples, the fact that if NC is used each time an opportunity arises or if the scheduling scheme does not account for NC, then the throughput can be smaller than that achieved without NC. In the first example, we fix the scheduling scheme (it provides maximum throughput when NC is not implemented) and demonstrate how applying NC reduces the system throughput. This indicates that NC and scheduling should be jointly considered. In the last two examples, we compare: (a) the throughput under an optimal scheduling without NC; (b) the throughput of the same system under an optimal scheduling adapted to NC. The scheduling in (b) is optimal subject to using NC at each opportunity. We show again that the throughput decreases when NC is used. This conclusion is more striking than that of the first example as here the scheduling scheme is aware of the NC capabilities.

Example 1 : Consider the network of Figure 1. Let the links experience random fading. Consequently, their rates oscillate randomly and independently between 1 and N : $R_1(t)$ and $R_2(t)$ are independent and identically distributed (i.i.d.), and equal to 1 with probability (w.p.) $1/2$, and to N w.p. $1/2$. With NC, for correct reception at both a and b , m has to broadcast at rate $\min\{R_1(t), R_2(t)\}$. First, consider the system without NC and with the following optimal opportunistic scheduling: If $R_1(t) = N = R_2(t)$, schedule link (a, m) w.p. $1/2$ and (b, m) w.p. $1/2$; if $R_1(t) = 1 = R_2(t)$, schedule each link w.p. $1/4$; if $(R_1(t), R_2(t)) = (N, 1)$, schedule (a, m) w.p. $1/4$ and (m, a) w.p. $3/4$; if $(R_1(t), R_2(t)) = (1, N)$, schedule (b, m) w.p. $1/4$ and (m, b) w.p. $3/4$. With this scheme, a throughput λ is achievable iff $\lambda \leq (1 + 3N)/16$. When NC is implemented, node m broadcasts XOR-ed packets whenever either (m, a) or (m, b) is scheduled in the above scheme. For the above scheduling scheme with NC, λ is achievable iff $\lambda \leq 1/2$ as m always transmits at rate 1 when scheduled. Note that applying NC strictly reduces the throughput if $N > 7/3$.

We make the following two observations on Example 1. (1) Assume that the rates $R_1(t) = r_1$ and $R_2(t) = r_2$ are not time varying, and without loss of generality, let $r_1 \leq r_2$. Then, irrespective of the scheduling used, NC provides a higher throughput than that without it. This is because with NC, packets from m to b (faster link) are transmitted along with packets from m to a (slower link). Since the transmissions from m to a have to happen in any case, NC saves transmissions from m to b . (2) With NC, there exists a scheduling scheme that can provide a throughput λ iff $\lambda \leq (1 + 3N)/12$ (which is higher than the achievable throughput without NC). The optimal scheme is as follows: If $(R_1(t), R_2(t)) = (N, 1)$, then schedule (a, m) ; if $(R_1(t), R_2(t)) = (1, N)$, then schedule (b, m) ; if $(R_1(t), R_2(t)) = (N, N)$, then broadcast XOR-ed packets from m ; if $(R_1(t), R_2(t)) = (1, 1)$, then schedule transmissions from nodes uniformly at random.

From the first observation, it may seem that if the link rates are constant, then NC improves the throughput performance for any topology. And from the second observation, it may seem that if an optimal scheduling with NC is used, then again the throughput increases. But, in the following example, we show that both statements do not hold.

Example 2 : We now provide an example illustrating why taking all opportunities to combine packets may result in throughput reduction, even when an optimal scheduling is used. Consider, in Figure 2, a simple extension of the network shown in Figure 1. Let $R_1(t) = 2$ and $R_2(t) = 1$ for all t , i.e., the rates are fixed but different. Now, for correct receptions at both a and b , m

has to broadcast combined packets at rate 1. Let the throughput requirements be $\lambda = \lambda_1 = 2/3$ and $\lambda_2 = 1/3$. We claim that the desired throughputs can be provided if NC is not used, while they can not be guaranteed if NC is used. Without NC, to provide the desired throughputs, we can use a scheduling scheme that activates the links (a_1, a_2) and (m, b) simultaneously and (a_1, a_2) and (b, m) simultaneously in $1/3$ fraction of slots each, and activates (b_1, b_2) and (m, a) simultaneously and (b_1, b_2) and (a, m) simultaneously in $1/6$ fraction of slots each. Now we prove that these throughputs can not be achieved using NC. Indeed, since $\lambda_1 = 2/3$, (a_1, a_2) has to be active in at least $2/3$ fraction of slots. As a consequence, (a, m) and (m, a) can be active in at most $1/3$ fraction of slots. Thus, to provide a throughput of $1/3$ to each of the two sessions that use (a, m) or (m, a) , these links must transmit at a rate no less than 2 when active. This is impossible if NC is used as then m broadcasts XOR-ed packets at rate 1 only.

From the above example it may seem that the main reason for reduction in throughput with NC is that the links have different capacities, and hence a broadcasted packet has to be transmitted at a lower rate (in this example, m broadcasts at the rate of 1 packet/slot, while the capacity of the link (m, b) is 2 packets/slot). In the next example, we demonstrate that the throughput with NC can be less than that without NC even when all the links have the same capacity.

Example 3 : Consider a wireless network shown in Figure 3. Let the throughput requirements be $\lambda = 1/4$ and $\lambda_1 = \lambda_2 = 1/2$. We claim that the desired throughputs can be provided if NC is not used, while they can not be guaranteed if NC is used. Without NC, to provide the desired throughputs, we can use a scheduling scheme that activates the links (a_1, a_2) and (m, b) simultaneously and (a_1, a_2) and (b, m) simultaneously in $1/4$ fraction of slots each, and activates (b_1, b_2) and (m, a) simultaneously and (b_1, b_2) and (a, m) simultaneously in $1/4$ fraction of slots each. Now we prove that these throughputs can not be achieved using NC. Since, $\lambda_1 = \lambda_2 = 1/2$ and (a_1, a_2) and (b_1, b_2) can not be active simultaneously, at least one them has to be active in each slot to guarantee the required throughput (necessary condition). But if NC is used then both a_1 and b_1 have to be silent. Thus, the required throughput can not be guaranteed.

In all previous examples, we have considered networks where NC does not use *opportunistic listening* (OL). OL refers to the ability of nodes to overhear packets transmitted in their neighborhood even when these packets are not meant for them. To see how OL helps, refer to Figure 4. Here, when a (resp. b) transmits a packet P_a (resp. P_b) to m , b' (resp. a') can overhear it. Thus as in Figure 1, node m can broadcast the XOR-ed packet $P_a \oplus P_b$ to both a' and b' , who recover

their respective packets by XOR-ing the broadcasted packet with the overheard packet. Using similar arguments as those used in Example 3, it can be shown that the network of Figure 4 can support the throughputs $\lambda = 1/4$ and $\lambda_1 = \lambda_2 = 1/2$ without NC, while it cannot when NC and OL are used. Thus, NC does not guarantee throughput improvement even when OL is used.

A key feature used in the above examples is that when an XOR-ed packet is transmitted to multiple receivers, all the other nodes in the neighborhood of the receivers have to remain silent: the use of NC reduces the spatial reuse in the network. Hence, for deciding whether to use NC, one has to evaluate the trade-off between the reduction in capacity due to the reduction in the spatial reuse and the capacity improvement due to the broadcast of XOR-ed packets.

Summarizing the insights from the above examples, Example 1 shows that NC and scheduling should be jointly designed, since using NC with arbitrary scheduling may result in performance losses. Examples 2 and 3 show that the decision to use NC has to be a function of many parameters including the network topology, the link rates and the throughput requirements of the various sessions. This calls for the design of joint NC and scheduling schemes that adapt to the network topology and link rates and provide the required throughput to each session, if doing so is at all possible.

III. A FRAMEWORK FOR DESIGNING JOINT NC AND SCHEDULING

A. Network Topology and Sessions

Consider a multi-hop wireless network, represented as a directed graph $G = (V, E)$, where V and E denote the set of nodes and links, respectively. The network is used by sessions to transport data packets. A session A is characterized by a doublet $(s(A), d(A)) \in V \times V$, where $s(A)$ and $d(A)$ denote the source and the destination, respectively, of A . Let \mathcal{S} denote the set of all sessions. Time is slotted.

We assume that the exogenous packets corresponding to the session A arrive at $s(A)$ as per a stochastic process $\{\lambda_A(t)\}_{t \geq 1}$, where $\lambda_A(t)$ denote the number of packets arriving in slot t . Packets have the same length. Exogenous arrivals across the slots are assumed to be i.i.d. Moreover, assume that $\lambda_A(1) \leq c < \infty$ for every A and define $\lambda_A = \mathbb{E}[\lambda_A(1)]$. Packets are stored in infinite buffers until served.

Packets of session $A \in \mathcal{S}$ are routed from $s(A)$ to $d(A)$ in, possibly, multiple hops. We consider fixed routing, and denote by \mathcal{R}_A the route for session A . This route is an ordered

subset of V , $\mathcal{R}_A = \{a_0, a_1, \dots, a_{N_A}\}$, such that $a_0 = s(A)$ and $a_{N_A} = d(A)$. Let us denote by $e_k^A = (a_k, a_{k+1})$ for every $k \in \{0, \dots, N_A - 1\}$. Furthermore, for every $i \in \mathcal{R}_A$ and $i \neq s(A)$, let $s_i(A)$ denote the node preceding node i on the route of session A , i.e., packets of session A use link $(s_i(A), i)$. Similarly, for every $i \neq d(A)$, $d_i(A)$ denotes the node after node i on route of session A . For each session $A \in \mathcal{S}$, each node $i \in V$ maintains a queue $q_{i,A}$ to store packets corresponding to this session. All the queues are served in First In First Out (FIFO) order. At the beginning of slot t , the queue length of $q_{i,A}$ is denoted by $Q_{i,A}(t)$, and its Head of Line (HoL) packet by $P_{i,A}(t)$. Finally, we say that sessions A and A' are symmetric sessions if $s(A) = d(A')$ and $d(A) = s(A')$. Note that the sets of links traversed by the packets of symmetric sessions may not be the same. The notations are illustrated in Figure 5.

B. MAC Layer and Scheduling Policies

In networks without NC, a scheduling policy at MAC layer decides, in each slot, which links should be activated and which sessions should be served on these links. In networks with NC, a scheduling policy has to additionally decide whether and how NC should be used. In other words, the policy imposes which nodes should use NC, and which packets should be encoded at these nodes. In this paper, we restrict our attention to NC schemes that allow bit-wise XOR of packets only. Thus, the NC scheme defines the set of possible XORs at each node; but, it is the scheduling scheme that decides whether and when to perform these XORs. For illustration, consider Figure 2. Here, NC allows XOR-ing the packets from nodes a and b at node m , but the scheduling policy will arbitrate whether and when to use this facility. For example, if $\lambda = \lambda_1 = 2/3$ and $\lambda_2 = 1/3$, then a scheduling policy that provides the required throughputs to all the sessions will not XOR packets at m (refer to Example 2); but if $\lambda = 4/3$ and $\lambda_1 = \lambda_2 = 1/3$, then a scheduling policy that provides the required throughputs to all the sessions will XOR packets at m .

Let \mathcal{L} denote the set of L feasible scheduling decisions, or schedules. Each element of \mathcal{L} defines (1) the links that are activated, (2) the sessions that are served on these links, and (3) the sessions whose packets are XOR-ed together.

Assumption 1: If $\ell \in \mathcal{L}$, then every ℓ_1 such that the set of active links under ℓ_1 is a subset of that under ℓ also belongs to \mathcal{L} .

The exact nature of \mathcal{L} depends on the MAC and Phy layer constraints, and also on the NC

scheme used. We provide the description of \mathcal{L} after presenting the NC schemes considered in Sections V and VI. But, for illustration, let us assume that the NC scheme XORs packets corresponding to symmetric sessions only. Then, each schedule $\ell \in \mathcal{L}$ is a subset of $E \times \overline{\mathcal{S}}$, where $\overline{\mathcal{S}} = \mathcal{S} \cup \{A \oplus A', A \in \mathcal{S}\}$. Notation $(e, A) \in \ell$ means that the link $e = (i, j)$ is active and serves queue $q_{i,A}$; $(e, A \oplus A') \in \ell$ means that $e = (i, j)$ is active and serves XOR-ed packets from queues $q_{i,A}$ and $q_{i,A'}$. The MAC and Phy layer constraints further restrict the choice of valid schedules. For example, if the RTS/CTS mechanism is used in IEEE 802.11-based networks and link $e = (i, j)$ is scheduled, then no node in the *neighborhood* of i and j can be scheduled. Thus, \mathcal{L} can not contain a schedule that allows nodes in the neighborhood of i and j to transmit, while simultaneously activating link (i, j) . Finally, the set of feasible schedules in a given slot has to reflect the fact that the transmissions from empty queues can not be scheduled.

Definition 1 (Scheduling Policy): A scheduling policy Δ is an algorithm that chooses a feasible schedule $\ell \in \mathcal{L}$ in each slot t .

To describe the system states under policy Δ , we use the superscript Δ : for example, $\ell^\Delta(t)$ will denote the schedule chosen by Δ in slot t ; $Q_{i,A}^\Delta(t)$ will denote the length of $q_{i,A}$ in slot t under Δ . Let $\mathcal{C}_\mathcal{L}$ denote the class of scheduling policies Δ such that $\ell^\Delta(t) \in \mathcal{L}$ for all t . The class $\mathcal{C}_\mathcal{L}$ also includes the *off-line* policies that arbitrate scheduling by taking into account past, present and even future network states.

C. The Phy Layer

We categorize the wireless systems into two classes, namely, systems with fixed link rates and systems with adaptive link rates.

1) *Fixed Rate Systems:* In such systems, the transmitter and receiver of each link negotiate the link rate during network set-up, and then always use this rate to communicate. Examples of such systems are networks based on the IEEE802.11 standards, where the rate control is performed rarely (at much longer time scale than that of packet transmissions). Let R_e denote the rate negotiated on link e . The variations in channel quality induced by fading and interference can be captured through packet error probabilities (PEP). Specifically, the PEP is the probability that the SINR is above certain level. We denote by $p_{e_k^A}(\ell)$ the PEP on link e_k^A for session A under schedule ℓ . The PEP also depends on t if the model accounts for fading. We assume that

$p_{e_k^A}(\ell) = 1$, if session A is not scheduled on e_k^A under ℓ . Now, we give an example to show how the PEP is related to the interference model.

Example 4: The Protocol Model: This model is a generalization of that considered in [9]. A transmission on link $e = (i, j)$ at the negotiated rate is successful if none of the nodes in the set K_e is transmitting. Typically $k \in K_e$, if the distance from k to j is sufficiently small. As a consequence, $p_e(\ell) = 0$, if all nodes in K_e are inactive under ℓ ; and $p_e(\ell) = 1$ otherwise.

2) *Adaptive Rate Systems:* In systems with a more elaborated Phy layer, link rates are adapted to the channel conditions and interference (e.g., by using advanced coding capabilities such as Hybrid ARQ). We denote by $R_{e_k^A}(\ell)$ the rate of link e_k^A for session A under schedule ℓ . The link rate also depends on t if the model has to account for fading. We assume that $R_{e_k^A}(\ell) = 0$ if session A is not scheduled on e_k^A under ℓ . Here is an example to show how the link rates relate to the interference model.

Example 5: The SINR-rate Model: Usually the link rate is related to the SINR at the receiver, and it is often well approximated by Shannon formula (up to a multiplicative constant). For example, consider the network of Figure 5, and assume that all nodes transmit at full power, say 1, when scheduled. If links $e_1^A = (1, 2)$ and $e_1^{A'} = (3, 4)$ are active under ℓ in slot t , then, the rate on link e_1^A is: $R_{e_1^A}(\ell, t) = W \log \left(1 + \frac{G_{12}(t)}{N_0 + G_{41}(t)} \right)$, where $G_{ij}(t)$ is the channel gain from i to j in slot t , N_0 is the noise power, and W is the bandwidth. Now, if node 2 broadcasts XOR-ed packet to nodes 1 and 3 under ℓ in slot t , then the rates on these links are: $R_{e_2^A}(\ell, t) = R_{e_3^{A'}}(\ell, t) = W \log \left(1 + \min \left\{ \frac{G_{21}(t)}{N_0}, \frac{G_{23}(t)}{N_0} \right\} \right)$.

Assumption 2: Let ℓ_1 be such that the set of active links in ℓ_1 is a subset of that in ℓ . Then, the rate (PEP, resp.) on every active link in ℓ_1 is greater (smaller, resp.) than or equal to that on the same link in ℓ . (Assumption 2 is typically valid in wireless networks as activating fewer links reduces interference.)

D. Design Objectives

Our aim is to propose *optimal* joint adaptive NC and scheduling schemes. Next, we introduce various definitions and then state this optimization problem.

Recall that the set of valid schedules \mathcal{L} accounts for the possible NC opportunities, i.e., for the NC scheme. Most of the proposed NC schemes, e.g. COPE, are designed under the constraint that XOR-ed packets must to be decoded at the next hop. Here, we relax this constraint. Thus,

encoded packets can be further XOR-ed with other, possibly encoded, packets. Hence, we have to carefully study the decodability of packets. For scalability, we impose that packets are decoded *on the fly*: if the NC scheme decides that an XOR-ed packet P has to be decoded at node i , then i should be able to decode P immediately after it receives P . Thus, the NC schemes considered have to be correct in the following sense.

Definition 2 (Correctness): Let P be a packet of session A . It is created in $q_{s(A),A}$ at time t . Assume that the packets of A are to be decoded at node $i \in \mathcal{R}_A$. Also, let $\mathcal{G} = \{\ell(u)\}_{u \geq t}$ denote a sequence of valid schedules after time t such that the first packet containing P (say P') arrives at i in slot $t_{\mathcal{G}}$. Then, we say that the NC scheme is correct, if i can decode P' to recover P immediately upon arrival of P' for every valid scheduling sequence \mathcal{G} .

Intuitively, the notion of correctness decouples the NC scheme and the scheduling strategy. Note that NC scheme only affects the set of valid schedules \mathcal{L} . But, once \mathcal{L} is defined, NC oblivious scheduling policy can be designed (see Definition 1). Now, if the NC scheme is correct, then each packet of every session can be recovered at its respective destination irrespective of the scheduling decisions as the packets of each session A must be decoded at $d(A)$. Next, we define the performance measures of interest.

Definition 3 (Stability): The system is stable under Δ , if $\sup_{t \geq 1} \{\mathbb{E}[Q_{i,A}^\Delta(t)]\} < \infty$ for every $i \in V$ and $A \in \mathcal{S}$. An arrival rate vector $\lambda = [\lambda_A : A \in \mathcal{S}]$ is said to be stabilizable by Δ , if the system is stable under Δ for λ .

Stability ensures finite expected delay for every packet. Moreover, in practice, the buffer capacity is finite, though large. Here, stability guarantees limited losses due to buffer overflow.

Definition 4 (Throughput Region): The throughput region of Δ is the set Λ^Δ of all the stabilizable rate vectors by Δ . The throughput region of the class of scheduling policies $\mathcal{C}_{\mathcal{L}}$ is $\Lambda_{\mathcal{L}} = \cup_{\Delta \in \mathcal{C}_{\mathcal{L}}} \Lambda^\Delta$.

Definition 5 (Throughput Optimality): A policy Δ is said to be throughput optimal in class $\mathcal{C}_{\mathcal{L}}$, if $\Lambda^\Delta = \Lambda_{\mathcal{L}}$.

IV. OPTIMAL SCHEDULING THEOREM

Now, we propose a throughput optimal policy within the class $\mathcal{C}_{\mathcal{L}}$, for any given set of schedules \mathcal{L} . In fact, we obtain a more general result: we provide a throughput optimal policy that minimizes certain cost. The cost may, for example, reflect the power consumption in the system, or as

explained in Section VII, may also be used to control the packet header size. We use the results derived here to obtain the throughput optimality of the NC and scheduling schemes considered in Sections V and VI.

Let $f(\ell)$ denote the cost if schedule ℓ is chosen. We assume that this cost function satisfies:

Assumption 3: The function $f(\cdot)$ is bounded, and for every ℓ_1 such that the set of activated links under ℓ_1 is a subset of that under ℓ , $f(\ell_1) \leq f(\ell)$.

Clearly, Assumption 3 holds if $f(\ell)$ is the total power required when schedule ℓ is chosen. Now let the arrival rate vector be λ . Then, the cost incurred under scheduling policy Δ is: $F^\Delta(\lambda) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T f(\ell^\Delta(t))$. Let $\mathcal{C}_\mathcal{L}(\lambda)$ denote the set of all policies that stabilizes λ using schedules in \mathcal{L} . Then, define $F_{\min}^{\mathcal{C}_\mathcal{L}(\lambda)} = \inf_{\Delta \in \mathcal{C}_\mathcal{L}(\lambda)} \{F^\Delta(\lambda)\}$.

Definition 6 (ϵ -Optimality): A policy Δ is said to be ϵ -optimal for a given λ , if $\Delta \in \mathcal{C}_\mathcal{L}(\lambda)$, and $F^\Delta(\lambda) \leq F_{\min}^{\mathcal{C}_\mathcal{L}(\lambda)} + \epsilon$.

We propose a policy that is both throughput optimal and ϵ -optimal. Due to space limitations, we obtain the results only for adaptive rate systems. Similar results can be obtained for fixed rate systems by replacing $R_{e_k^A}(\ell)$ with $R_{e_k^A}(1 - p_{e_k^A}(\ell))$ in the following. We analyze systems without random fading. The analysis can be generalized to account for fading, see [4].

A. Throughput Region

We first characterize the throughput region of $\mathcal{C}_\mathcal{L}$. Let $\mathcal{X}_\mathcal{L}$ denote the set of all arrival rate vectors λ for which there exists a vector $\alpha = [\alpha_1 \cdots \alpha_L]$ such that for all ℓ , $\alpha_\ell \geq 0$, $\sum_{\ell \in \mathcal{L}} \alpha_\ell = 1$, and, $\sum_{\ell \in \mathcal{L}} \alpha_\ell R_{e_k^A}(\ell) \geq \lambda_A$, $\forall k < N_A$ and $\forall A \in \mathcal{S}$. Let $\mathcal{X}_\mathcal{L}^\circ$ be the set of ν such that there exists $\lambda \in \mathcal{X}_\mathcal{L}$ with $\nu < \lambda$ coordinate-wise. Next, we characterize the throughput region of $\mathcal{C}_\mathcal{L}$.

Theorem 1: The throughput region $\Lambda_\mathcal{L}$ satisfies $\mathcal{X}_\mathcal{L}^\circ \subseteq \Lambda_\mathcal{L} \subseteq \mathcal{X}_\mathcal{L}$. In words, if $\lambda \in \mathcal{X}_\mathcal{L}^\circ$, then there exists $\Delta \in \mathcal{C}_\mathcal{L}$ such that $\lambda \in \Lambda^\Delta$, but if $\lambda \notin \mathcal{X}_\mathcal{L}$, then $\lambda \notin \Lambda^\Delta$ for every $\Delta \in \mathcal{C}_\mathcal{L}$.

B. Optimal Policy

Now, we define a parameterized back-pressure based policy denoted by $\Delta^*(\kappa)$, and prove its throughput optimality and ϵ -optimality. Let $\partial Q_{k,A}(t)$ denote the back-pressure along e_k^A , i.e., $\partial Q_{k,A}(t) = Q_{a_k,A}(t) - Q_{a_{k+1},A}(t)$. At time t , $\Delta^*(\kappa)$ chooses the schedule defined by:

$$\ell^{\Delta^*(\kappa)}(t) = \arg \max_{\ell \in \mathcal{L}} \left\{ \sum_{A,k} R_{e_k^A}(\ell) \partial Q_{k,A}(t) - \kappa f(\ell) \right\}. \quad (1)$$

Theorem 2: For all $\kappa < \infty$, $\Delta^*(\kappa)$ is throughput optimal in $\mathcal{C}_{\mathcal{L}}$. Moreover, for all $\epsilon > 0$, there exists $\widehat{\kappa} > 0$ such that for all $\kappa > \widehat{\kappa}$, $\Delta^*(\kappa)$ is ϵ -optimal.

The problem of minimizing cost subject to stability has been studied previously in [5], [18], [26]. However, our result is not a consequence of the results derived there. In [5], [18], the authors analyze one hop sessions only. So, the queueing process is driven primarily by the exogenous arrivals that are independent of the scheduling decisions. Here, however, the queueing process is affected by the chosen schedule as the arrivals in $q_{a_k,A}$ are the departures from $q_{a_{k-1},A}$. In [26], the author has studied multi-hop networks, but under the following assumption: if the set of active links under ℓ_1 is a subset of that under ℓ , then for all links e_k^A activated under both ℓ_1 and ℓ , $R_{e_k^A}(\ell_1) = R_{e_k^A}(\ell)$. This assumption does not hold in typical wireless networks as the link rates depend on the interference caused by the transmissions on other active links. Thus, typically, $R_{e_k^A}(\ell_1) > R_{e_k^A}(\ell)$. In view of these differences, though the nature of our optimal policy $\Delta^*(\kappa)$ is similar to those proposed earlier, the proofs from [5], [18], [26] do not hold here.

Like many other back-pressure based policies proposed in literature [5], [18], [27], [26], $\Delta^*(\kappa)$ is centralized and has high computational complexity. Fortunately, back-pressure based policies are extensively studied, and many schemes for reducing their complexity [5] and for distributed implementations [3], [8], [17], [24] have been proposed. Similar approaches can be developed for the joint NC and scheduling scheme proposed here.

V. OPTIMAL SCHEDULING FOR COPE

Here, we apply the general framework developed in Section IV to provide an optimal scheduling strategy adapted to COPE, a NC scheme recently introduced in [11].

A. Overview of COPE

COPE is a practical NC scheme designed for improving the throughput of unicast sessions in networks with arbitrary topology. In COPE, nodes send XOR-ed combinations of packets that can be decoded at the next hop: a node i sends an XOR-ed packet $P_1 \oplus \dots \oplus P_m$ only to nodes that already have $m - 1$ of m packets P_1, \dots, P_m . When a node j receives an encoded packet, it immediately decodes it. A node j possesses the $m - 1$ required packets in two possible scenarios: (i) these packets have been transmitted by j or (ii) j has intercepted these packets by listening to the transmissions (not meant for j) from its neighboring nodes; this is referred

to as *opportunistic listening* (OL). Scenario (ii) is possible because of the broadcast nature of the wireless channel. Here, we do not consider OL. A detailed discussion on advantages and limitations of OL is presented in [4].

Locally Symmetric Sessions: Since, we do not allow OL, a node can have the packets required to decode an encoded packet only if (i) is satisfied. Let packets of sessions A and B be routed through i . These sessions are *locally symmetric* at i if $d_i(A) = s_i(B)$ and $s_i(A) = d_i(B)$. Here, i can XOR packets from sessions A and B , and broadcast the XOR-ed packet to $d_i(A)$ and $d_i(B)$. The latter nodes will be able to decode the XOR-ed packet as (i) holds.

As illustrated in Example 1, COPE, associated with an arbitrary scheduling policy, may not provide any throughput gain. This calls for the design of a joint NC and scheduling policy that will guarantee that the gains expected from COPE can actually be met. To this aim, we apply the framework of Section IV and derive a throughput optimal policy adapted to COPE.

B. An Optimal Scheduling for COPE

Let us first characterize the set of valid schedules $\mathcal{L}_{\text{COPE}}$ compatible with COPE. Note that COPE is correct only if at most two packets corresponding to locally symmetric sessions are XOR-ed (Theorem 4.1 of [11]). Hence, the set of schedules compatible with COPE is defined as follows: A schedule $\ell \in \mathcal{L}_{\text{COPE}}$ is defined as a subset of $\mathcal{E} = \cup_{e \in E} (e \times \overline{\mathcal{S}(e)})$, where $\overline{\mathcal{S}(e)} = \mathcal{S} \cup \{A \oplus B : A, B \text{ locally symmetric at } i, e = (i, d_i(A))\}$. Notation $(e, A) \in \ell$ means that the link $e = (i, j)$ is active and serves queue $q_{i,A}$; $(e, A \oplus B) \in \ell$ means that link $e = (i, j)$ is active and serves XOR-ed packets from locally symmetric sessions A and B . Now, schedule ℓ belongs to $\mathcal{L}_{\text{COPE}}$ if it satisfies the following constraints: $\forall A \in \mathcal{S}$ and $\forall e = (i, d_i(A))$,

- if $(e, A) \in \ell$, then for all B , $(e, A \oplus B) \notin \ell$;
- if $(e, A \oplus B) \in \ell$, then $(e', A \oplus B) \in \ell$ where $e' = (i, d_i(B))$, and $(e, A) \notin \ell$, $(e', B) \notin \ell$.

In addition to the above constraints, any schedule ℓ in $\mathcal{L}_{\text{COPE}}$ has to satisfy the Phy and MAC constraints as illustrated in Section III-B.

Consider the scheduling policy Δ_{COPE}^* that depending on the queue lengths and link rates, selects, in slot t , schedule ℓ defined as follows:

$$\ell^{\Delta_{\text{COPE}}^*}(t) = \arg \max_{\ell \in \mathcal{L}_{\text{COPE}}} \left\{ \sum_{A,k} R_{e_k^A}(\ell) \partial Q_{k,A}(t) \right\}.$$

We prove that Δ_{COPE}^* has the largest throughput region within the class \mathcal{C}_1 of the joint NC and scheduling policies with correct NC and that do not use OL.

Theorem 3: The policy Δ_{COPE}^* is throughput optimal in \mathcal{C}_1 .

Proof: Since any correct NC scheme without OL can XOR two packets from locally symmetric sessions only, any $\Delta \in \mathcal{C}_1$ selects schedules from $\mathcal{L}_{\text{COPE}}$. Thus, from Theorem 2, Δ_{COPE}^* is throughput optimal in \mathcal{C}_1 . ■

C. Throughput gains of COPE

In general, quantifying the throughput gain achieved with NC is difficult as it depends on many parameters that include the network topology, the underlying Phy and MAC layers and the relative throughput requirements of the sessions. We define the throughput gain by comparing the throughput region of the set of scheduling policies with NC, and the throughput region Λ_0 of policies without NC. The gain achieved by COPE for the network of Figure 5 with a Phy layer satisfying the Protocol model is illustrated in Figure 6. There, $G(u)$ is the gain in direction u , where u is the unit vector representing the relative throughput requirements of the sessions A and A' . The throughput gain is then defined as $\max_u G(u)$.

In [22], [23], the authors characterize the maximum throughput region of 1D networks with NC. In [15], [12], [2], upper bounds on the throughput gains for large random networks are derived. Characterizing the throughput gain with NC for more general topologies is quite challenging. However, even for an arbitrary network, one can use Theorem 1 to characterize the throughput region with or without NC, and numerically compute the throughput gain. A similar approach is used in [25].

VI. XOR-SYM: A SIMPLIFIED NC SCHEME

In this section, we design a NC scheme that requires a minimal change in the present network architecture and yet provides similar performance benefits as COPE. To this aim, we enforce the following constraint on the type of NC used in the network.

C1: Decoding at Destination Only: A packet corresponding to session A is decoded at $d(A)$ only, and not at any other node.

Many of the NC strategies proposed in the literature (e.g., COPE) require that packets are decoded at each node. Thus, each node has to maintain the packets received and transmitted

successfully in the past in order to decode the packets that will arrive in the future. Moreover, whenever an encoded packet arrives, in order to decode it, a node has to perform look-up in its buffer for all but one packets that compose the incoming packet. The look-up may be computationally expensive. We eliminate this potential bottleneck for scalability of NC schemes by imposing the constraint C1. Intermediate nodes can then remain simple: they only need to perform bit-wise XOR of HoL packets; the required additional functionality can be incorporated without adversely affecting scalability. In the following, we propose XOR-Sym, a correct NC scheme satisfying the constraint C1 and yet providing throughput benefits.

A. The XOR-Sym coding scheme

Figures 7 and 8 provide the pseudo codes for XOR-Sym in the cases of fixed and adaptive rate systems. The key feature of XOR-Sym is that it XORs packets corresponding to symmetric sessions only. Contrast this with COPE which XORs packets corresponding to *locally* symmetric sessions at each node. Due to space limitations, we only describe XOR-Sym for fixed rate systems. Consider the network of Figure 5, whose Phy layer follows the Protocol model and with negotiated link rates all equal to 1 packet/slot (refer to Figure 7 for systems with heterogeneous rates). If the scheduling scheme decides to serve session A only on link $(2, 3)$ in slot t , then node 2 transmits $P_{2,A}(t)$. If $P_{2,A}(t)$ is successfully received at node 3, then node 2 discards this packet and replace it with a new packet at the HoL position in $q_{2,A}$. $P_{2,A}(t)$ is queued at the end of $q_{3,A}$. If $P_{2,A}(t)$ is not successfully received at node 3, then it is retained at the HoL position in $q_{2,A}$. Now, suppose that the scheduling scheme decides to broadcast an XOR-ed packet from node 2 on links $(2, 1)$ and $(2, 3)$. Then, 2 broadcasts $P = P_{2,A}(t) \oplus P_{2,A'}(t)$. Three cases arise. (i) Both 1 and 3 receive P successfully. Then, 2 discards these packets, and new packets come to the HoL positions in $q_{2,A}$ and $q_{2,A'}$. P is decoded at node 1, while it is queued at the end of $q_{3,A}$. (ii) Only one of the intended recipients, say node 3, receives P correctly. Then, $P_{2,A}(t)$ is discarded from $q_{2,A}$ and is replaced by a new packet at the HoL position of $q_{2,A}$, while $P_{2,A'}(t)$ is retained at the HoL position in $q_{2,A'}$. P is queued at the end of $q_{3,A}$. The case when only 1 receives P correctly is similar. (iii) Both 1 and 3 do not receive P correctly. Then, both $P_{2,A}(t)$ and $P_{2,A'}(t)$ are retained at HoL positions in $q_{2,A}$ and $q_{2,A'}$.

Since intermediate nodes do not decode packets, encoded packets can be XOR-ed again. For example, in (ii) above, the XOR-ed packet P is queued in $q_{3,A}$ as it is. When P comes to the

HoL position in $q_{3,A}$, it can be XOR-ed again with a packet from $q_{3,A'}$. Thus, it is not clear whether the destinations can decode the received packets. In the following lemma, we show that XOR-Sym is correct given that: for each session A , $s(A)$ keeps all the packets of A that it has already transmitted, and $d(A)$ keeps all the packets of A that it could correctly decode.

Lemma 1: The NC scheme XOR-Sym is correct.

B. An Optimal Scheduling for XOR-Sym

Since XOR-Sym combines packets only from symmetric sessions, the set of all possible schedules $\mathcal{L}_{\text{XOR-Sym}}$ is as follows. A schedule $\ell \in \mathcal{L}_{\text{XOR-Sym}}$ is a subset of $E \times \overline{\mathcal{S}}$, where $\overline{\mathcal{S}} = \mathcal{S} \cup \{A \oplus A', A \in \mathcal{S}\}$. In addition, if $\ell \in \mathcal{L}_{\text{XOR-Sym}}$, it satisfies the following constraints:

- if $(e, A) \in \ell$, then $(e, A \oplus A') \notin \ell$;
- if $(e, A \oplus A') \in \ell$, then $(e', A \oplus A') \in \ell$ where $e' = (i, d_i(A'))$, and $(e, A) \notin \ell$, $(e', A') \notin \ell$.

Now, consider the scheduling policy $\Delta_{\text{XOR-Sym}}^*$ that, depending on the queue lengths and link rates, selects, in slot t , schedule ℓ defined as follows:

$$\ell_{\text{XOR-Sym}}^*(t) = \arg \max_{\ell \in \mathcal{L}_{\text{XOR-Sym}}} \left\{ \sum_{A,k} R_{e_k^A}(\ell) \partial Q_{k,A}(t) \right\}.$$

Now, we prove that $\Delta_{\text{XOR-Sym}}^*$ has the largest throughput region within the class \mathcal{C}_2 of the joint NC and scheduling schemes with correct NC and that satisfies the constraint C1. Note that \mathcal{C}_2 also contains off-line policies.

Theorem 4: The policy $\Delta_{\text{XOR-Sym}}^*$ is throughput optimal in \mathcal{C}_2 .

In view of Theorem 2, the above result follows from the fact that any scheme in \mathcal{C}_2 chooses schedules from $\mathcal{L}_{\text{XOR-Sym}}$ in each slot, which is a consequence of the following lemma.

Lemma 2: Consider a NC scheme satisfying the constraint C1, and assume that it XORs packets from sessions A and B , where $B \neq A'$. Then the NC scheme is not correct.

C. Throughput gains of XOR-Sym

Note that for any network, $\mathcal{L}_0 \subseteq \mathcal{L}_{\text{XOR-Sym}} \subseteq \mathcal{L}_{\text{COPE}}$, where \mathcal{L}_0 is the set of all feasible schedules without NC. Thus, $\Lambda_0 \subseteq \Lambda_{\text{XOR-Sym}} \subseteq \Lambda_{\text{COPE}}$: the throughput gain achieved with XOR-Sym over policies that do not use NC is greater than 1, but it may be less than that achieved with COPE. The scalability of XOR-Sym compared to that of COPE is obtained at the expense of a smaller throughput region. Note however, that the maximum gain achieved by

XOR-Sym and COPE are identical, and are achieved in the 1D network as described at the end of Section V. Moreover, under XOR-Sym, the computational complexity at intermediate nodes and the throughput gain can be traded by splitting sessions into several logical sessions. For example, consider a network where packets of sessions A and B follow the routes $\mathcal{R}_A = \{1, 2, 3, 4, 5\}$ and $\mathcal{R}_B = \{6, 4, 3, 2, 1\}$. A and B are not symmetric as $d(A) \neq s(B)$, but both these sessions traverse through nodes 1, 2, 3 and 4. Now, let us split each of these sessions into two logical sessions as follows: $A_1 = (1, 4)$, $A_2 = (4, 5)$ and $B_1 = (4, 1)$, $B_2 = (6, 4)$. Note that now A_1 and B_1 are symmetric and their packets can be XOR-ed under XOR-Sym. Thus, splitting sessions will provide a larger throughput region. But, now the intermediate node 4 has to decode packets, increasing its complexity. Note that XOR-Sym and COPE are identical if the sessions are split into several logical sessions, each traversing exactly one link. A technical difficulty with this approach is that the arrivals at the sources of the logical sessions are not i.i.d.; however, the analysis in Section IV can be extended to this case. Finally, we believe that creating 1-hop logical sessions everywhere (as in COPE) is not necessary to ensure optimal throughput, because most often only few links are bottlenecks in the network. It may be sufficient to define logical sessions so as to maximize the NC opportunities around these links. The logical sessions may also be created adaptively based on the queue length information.

D. Limitation of XOR-Sym

In NC schemes, to ensure decodability, the header of each packet contains the identities of all the packets XOR-ed in this packet. For a packet $P = P_1 \oplus \dots \oplus P_m$, we say that its packet header size is m . Now, if two packets of header sizes m and n are XOR-ed, then the header length of the resulting packet is at most $m + n$. With XOR-Sym, since packets are decoded at destinations only, the header sizes can be quite large. Theoretically, it is possible to construct an example where the header size can become arbitrarily large even for networks with simple topologies as in Figure 5; however, as shown in Section VII, we have verified using simulations that in fact, the header size remains modest unless the network becomes heavily loaded. In Section VII, we also propose some solutions to limit the header sizes.

VII. NUMERICAL EXPERIMENTS

In this section, we present some numerical experiments verifying the analytical results of the previous sections. We give the performance of XOR-Sym and of the associated optimal scheduling policy $\Delta_{\text{XOR-Sym}}^*$. Due to space limitations, we present results in the case of simple 1D networks. Refer to [4] for results on networks with more general topologies.

Consider a 1D network as depicted in Figure 5 but with N nodes. Interference follows the protocol model, and we assume that the reception at a node is interfered by the transmission of the 1-hop neighbors, i.e., for instance, using the notation of Section III-C.1, $K_{(i,i+1)} = \{i+2\}$. The negotiated link rates are all equal to 1. It is then easy to prove (see [4]) that the throughput regions with and without XOR-Sym are independent of N and represented in Figure 6. In this example, the NC gain is maximized when the arrival rates of the two symmetric sessions are equal, $\lambda_A = \lambda_{A'}$, and COPE and XOR-Sym provide similar throughput gains.

Figure 9 (top-left) provides the mean end-to-end packet delay as a function of the session rate for $\Delta_{\text{XOR-Sym}}^*$. The results are compared with those obtained without NC, but with a throughput optimal policy. Note that as expected, these schemes achieve maximum throughput, i.e., the mean packet delay is finite for all $\lambda_A < 1/3$ with XOR-Sym, and for all $\lambda_A < 1/4$ without NC. In Figure 9 (top-right) we present the mean packet header size using XOR-Sym. When the network size is small, e.g. $N = 4$, the mean header size remains small unless the system load approaches the stability limit. The header size increases with N .

To reduce the number of packets XOR-ed into a single packet, we associate a cost to the XOR-ing procedure: for any schedule ℓ chosen at time slot t , we denote by $f(\ell, t)$ the total number of packets involved in XORs under ℓ (e.g., if under ℓ , only packets $P_1 \oplus P_2$ and P_3 are XOR-ed, the cost is 3). Note that this cost function does not strictly correspond to the framework of Section IV; but the latter can be readily modified to account for this kind of costs. Figure 9 (Bottom) presents the mean packet header size using the optimal policy $\Delta_{\text{XOR-Sym}}^*(\kappa)$, for different values of κ in a network of $N = 8$ nodes. The choice of κ allows us to tune the trade-off between packet header size and delay.

VIII. CONCLUSION

We have investigated the use of network coding (NC) in wireless multi-hop networks for unicast sessions. Surprisingly, we could build simple and realistic examples of networks where

NC reduces the throughput performance. This happens when the NC schemes are greedy in the sense that all opportunities to combine and broadcast packets are exploited. We have also observed that if NC and scheduling are designed separately, then the throughput gain expected from NC may not be achieved.

These observations have emphasized the need for adaptive schemes that use NC opportunities only when they can provide performance benefits. It seems also critical that the scheduling choices and the NC decisions should be coupled. Hence, we have developed a generic framework to design joint optimal NC and scheduling schemes. We have applied this framework to propose an optimal scheduling scheme adapted to COPE, a recently introduced NC scheme. We have also designed XOR-Sym, a new NC scheme, and its associated optimal scheduling scheme. XOR-Sym exhibits a lower complexity than that of COPE but yet offers similar performance gains. The proposed framework can be extended to account for random fading, and also to design rate control mechanisms to maximize certain network utility. Due to the space constraints, we present the extensions in [4].

REFERENCES

- [1] R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network information flow. *IEEE Trans. on Information Theory*, 46(4):1204–1216, 2000.
- [2] S. Aly, V. Kapoor, J. Meng, and A. Klappenecker. Bounds on the network coding capacity for wireless random networks. In *NetCod '07: Third Workshop on Network Coding, Theory and Applications*, Jan 2007.
- [3] A. Brzezinski, G. Zussman, and E. Modiano. Enabling distributed throughput maximization in wireless mesh networks - a partitioning approach. In *Proc. of ACM Mobicom*, 2006.
- [4] P. Chaporkar and A. Proutiere. Adaptive network coding and scheduling for maximizing throughput in wireless networks. Technical report, Indian Inst. of Tech., Mumbai, India, 2008. <http://www.ee.iitb.ac.in/wiki/faculty/chaporkar>.
- [5] P. Chaporkar and S. Sarkar. Stable scheduling policies for maximizing throughput in generalized constrained queueing. In *Proc. of IEEE Infocom*, 2006.
- [6] Q. Dong, J. Wu, W. Hu, and J. Crowcroft. Practical network coding in wireless networks. In *MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, pages 306–309, Montréal, Québec, Canada, Sep 2007.
- [7] A. Eryilmaz and D. Lun. Control for intersession network coding. In *NetCod '07: Third Workshop on Network Coding, Theory and Applications*, Jan 2007.
- [8] A. Gupta, X. Lin, and R. Srikant. Low-complexity distributed scheduling algorithms for wireless networks. In *INFOCOM '07: 26th IEEE International Conference on Computer Communications*, pages 1631–1639, May 2007.
- [9] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, 46(2):388–404, 2000.
- [10] T. Ho and R. Koetter. Online incremental network coding for multiple unicasts. In *Proc. of DIMACS Working Group on Network Coding*, 2005.

- [11] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. Xors in the air: Practical wireless network coding. In *Proc. of ACM Sigcomm*, 2006.
- [12] J. Le, J. Lui, and D. Chiu. How many packets can we encode? - an analysis of practical wireless network coding. In *INFOCOM '08: IEEE the 27th Conference on Computer Communications*, pages 371–375, Phoenix, AZ, Apr 2008.
- [13] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Trans. on Information Theory*, 49(2):371–381, 2003.
- [14] Z. Li and B. Li. Network coding: The case of multiple unicast sessions. In *Proc. of Allerton Conference*, 2004.
- [15] J. Liu, D. Goeckel, and D. Towsley. Bounds of the gain of network coding and broadcasting in wireless networks. In *Proc. of IEEE Infocom*, 2007.
- [16] D. S. Lun, N. Ratnakar, R. Koetter, M. Medard, E. Ahmed, and H. Lee. Achieving minimum-cost multicast: A decentralized approach based on network coding. In *Proc. of IEEE Infocom*, 2005.
- [17] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossiping. In *Proc. of ACM Sigmetrics / IFIP Performance*, 2006.
- [18] M. Neely. Energy optimal control for time varying wireless networks. In *IEEE Infocom*, 2005.
- [19] S. Omiwade, R. Zheng, and C. Hua. Butterflies in the mesh: Lightweight localized wireless network coding. In *NetCod '08: Fourth Workshop on Network Coding, Theory and Applications*, pages 1–6, Jan 2008.
- [20] A. Ramamoorthy, J. Shi, and R. Wesel. On the capacity of network coding for random networks. *IEEE Trans. on Information Theory*, 51(8):2878–2885, 2005.
- [21] S. Rayanchu, S. Sen, J. Wu, S. Sengupta, and S. Banerjee. Loss-aware network coding for unicast wireless sessions: Design, implementation, and performance evaluation. In *SIGMETRICS '08*, 2008.
- [22] Y. E. Sagduyu and A. Ephremides. Network coding in wireless queueing networks: Tandem network case. In *Proc. of IEEE ISIT*, 2006.
- [23] Y. E. Sagduyu and A. Ephremides. Some optimization trade-offs in wireless network coding. In *Proc. of CISS*, 2006.
- [24] S. Sanghavi, L. Bui, and R. Srikant. Distributed link scheduling with constant overhead. In *ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 313–324, San Diego, CA, 2007.
- [25] S. Sengupta, S. Rayanchu, and S. Banerjee. An analysis of wireless network coding for unicast sessions: The case for coding -aware routing. In *Proc. of IEEE Infocom*, 2007.
- [26] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, 50(4):401–457, 2005.
- [27] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multi-hop radio networks. *IEEE Trans. on Automatic Control*, 37(12):1936–1949, 1992.
- [28] Y. Wu, P. A. Chou, and S.-Y. Kung. Information exchange in wireless networks with network coding and physical-layer broadcast. Technical Report MSR-TR-2004-78, 2004.
- [29] H. Yomo and P. Popovski. Opportunistic scheduling for wireless network coding. In *ICC '07: IEEE International Conference on Communications*, pages 5610–5615, Jun 2007.

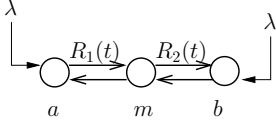


Fig. 1. A 3-node network topology handling two sessions, one from a to b and another from b to a . Packets for both the sessions are routed through relay m . The network is symmetric, i.e., the required throughputs for sessions from a to b and from b to a are the same (λ packets/slot), and the rates on the links (a, m) and (m, a) ((b, m) and (m, b) , resp.) are equal to $R_1(t)$ ($R_2(t)$, resp.) packets/slot in slot t . $R_1(t) = R_2(t) = 1$ at all time t . Because of interference, only one of nodes a , b and m can transmit in a slot.

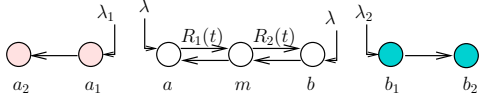


Fig. 2. An extension of the network shown in Figure 1. Here, two sessions from a_1 to a_2 and b_1 to b_2 are added, and these require throughputs of λ_1 and λ_2 , respectively. The maximum transmission rate on (a_1, a_2) and (b_1, b_2) is 1 packet/slot in each slot. We assume that a_1 (b_1 , resp.) can not transmit when a (b , resp.) is either transmitting or receiving. This interference model arises if IEEE 802.11 MAC with RTS and CTS is used.

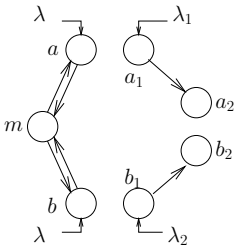


Fig. 3. The network shown in Figure 1 with different topology. Here, the maximum transmission rates on all the links is 1 packet/slot in each slot. We assume that a_1 (b_1 , resp.) can not transmit when a (b , resp.) is either transmitting or receiving. Moreover, a_1 and b_1 can not simultaneously transmit.

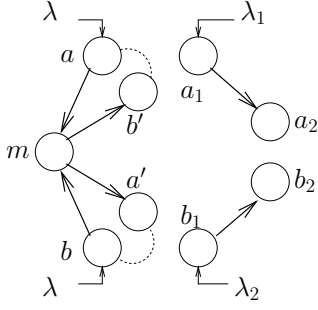


Fig. 4. A network topology with four sessions, viz. from a to a' , from b to b' , from a_1 to a_2 and from b_1 to b_2 . Paths traversed by each of the sessions is shown using the directed arrows. The dashed arc between the nodes a and b' (b and a' , resp.) indicates that the transmissions from a (b , resp.) can be intercepted by b' (a' , resp.). The maximum transmission rates on all the links is 1 packet/slot in each slot. We assume that a_1 (b_1 , resp.) can not transmit when either link (a,m) or link (m,b') ((b,m) or (m,a') , resp.) is active. Moreover, a_1 and b_1 can not simultaneously transmit.

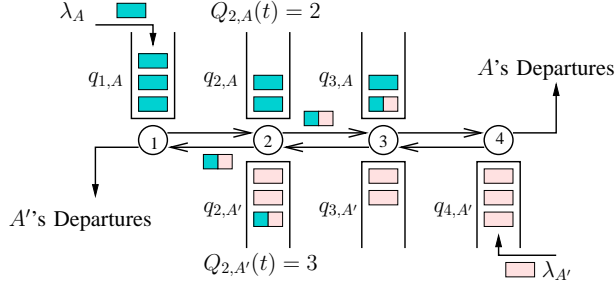


Fig. 5. A 4-node linear network with symmetric sessions A and A' . $s(A) = d(A') = 1$ and $s(A') = d(A) = 4$. At $s(A)$ and $s(A')$ new packets arrive at rate λ_A and $\lambda_{A'}$, respectively. Here, $\mathcal{R}_A = \{1, 2, 3, 4\}$, while $\mathcal{R}_{A'} = \{4, 3, 2, 1\}$. Regarding node 2 for example: $s_2(A) = 1$ and $d_2(A) = 3$; 2 FIFO queues, $q_{2,A}$ and $q_{2,A'}$, corresponding to both sessions are maintained; packets from these queues can be XOR-ed and broadcasted to nodes 1 and 3.

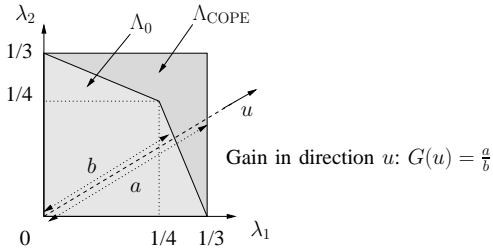


Fig. 6. The throughput regions with or without COPE for the network of Figure 5 - The Phy layer follows the Protocol model, and interfering nodes are the 1-hop neighbors only.

XOR-Sym for Fixed Rate Systems

begin

```

1: Assume  $(e, A) \in \ell^\Delta(t)$ , where  $e = (i, d_i(A))$ ;
2: for  $j = 1, \dots, R_e$  do
3:   Transmit  $P_{i,A}(t)$ ;
4:   if  $d_i(A)$  successfully receives the packet then
5:     Discard  $P_{i,A}(t)$ ;
6:      $P_{i,A}(t)$  is replaced by the next in line packet in  $q_{(i,A)}$ ;
7:   end if
8: end for
9: Assume  $(e, A \oplus A') \in \ell^\Delta(t)$ , where  $e = (i, d_i(A))$ ;
10: for  $j = 1, \dots, \min\{R_e, R_{e'}\}$  do
11:    $P \leftarrow P_{i,A}(t) \oplus P_{i,A'}(t)$ ;
12:   Broadcast  $P$ ;
13:   if Both  $d_i(A)$  and  $d_i(A')$  successfully receive  $P$  then
14:     Discard  $P$ ;
15:      $P_{i,A}(t)$  and  $P_{i,A'}(t)$  are replaced by the next in line packets in  $q_{(i,A)}$  and  $q_{(i,A')}$  respectively;
16:   else if Only  $d_i(A)$  successfully receives  $P$  then
17:     Retain  $P_{i,A'}(t)$  at HoL position in  $q_{(i,A')}$ ;
18:      $P_{i,A}(t)$  is replaced by the next in line packet in  $q_{(i,A)}$ ;
19:   else if Only  $d_i(A')$  successfully receives  $P$  then
20:     Retain  $P_{i,A}(t)$  at HoL position in  $q_{(i,A)}$ ;
21:      $P_{i,A'}(t)$  is replaced by the next in line packet in  $q_{(i,A')}$ ;
22:   else
23:     Retain both the packets at HoL positions in their respective queues;
24:   end if
25: end for

```

end

Fig. 7. Pseudo code of XOR-Sym NC scheme with scheduling policy Δ for fixed rate systems - These tasks are performed in each slot.

XOR-Sym for Adaptive Rate Systems

begin

- 1: Assume $(e, A) \in \ell^\Delta(t)$, where $e = (i, d_i(A))$;
- 2: Transmit $R_e(\ell(t))$ packets from $q_{(i,A)}$;
- 3: Discard all the transmitted packets;
- 4: Assume $(e, A \oplus A') \in \ell^\Delta(t)$, where $e = (i, d_i(A))$;
- 5: $R \leftarrow \min\{R_e(\ell(t)), R_{e'}(\ell(t))\}$, where $e' = (i, d_i(A'))$;
- 6: **for** $j = 1, \dots, R$ **do**
- 7: Broadcast $P_{i,A}(t) \oplus P_{i,A'}(t)$ (at rate R);
- 8: Discard $P_{i,A}(t)$ and $P_{i,A'}(t)$;
- 9: $P_{i,A}(t)$ and $P_{i,A'}(t)$ are replaced by the next in line packets in $q_{(i,A)}$ and $q_{(i,A')}$ respectively;
- 10: **end for**

end

Fig. 8. Pseudo code of XOR-Sym NC scheme with policy Δ for adaptive rate systems - These tasks are performed in each slot.

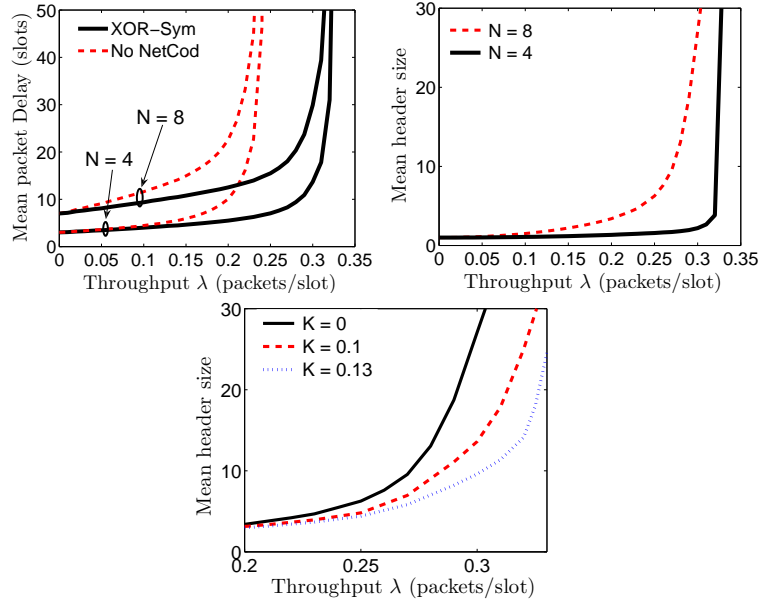


Fig. 9. [Top] Mean packet delay as a function of session throughputs $\lambda = \lambda_A = \lambda_{A'}$ with or without XOR-Sym and policy $\Delta_{\text{XOR-Sym}}^*$ - [Middle] Mean packet header size using $\Delta_{\text{XOR-Sym}}^*$ - [Bottom] Mean packet header size using $\Delta_{\text{XOR-Sym}}^*(\kappa)$ for $\kappa = 0, 0.1$, and 0.13 , $N = 8$.

APPENDIX I

PROOFS OF THEOREMS 1 AND 2

First we state the supporting lemmas that we use to prove Theorems 1 and 2. Due to space constraints, proofs for some of the lemmas are omitted.

Lemma 3: The throughput region satisfies $\Lambda_{\mathcal{L}} \subseteq \mathcal{X}_{\mathcal{L}}$.

Let $\overline{\mathcal{C}_{\mathcal{L}}}$ denote the space of all the policies that are allowed to schedule transmission from $q_{i,A}$ in slot t even when $Q_{i,A}(t) = 0$. Scheduling a transmission from an empty queue corresponds to transmitting a pseudo packet. The pseudo packets are immediately discarded by the receiving node. We note that we use such policies only to obtain a compact proof of the optimality of Δ^* , and we do not allow policies to schedule transmissions from empty queues. Note that $\mathcal{C}_{\mathcal{L}} \subseteq \overline{\mathcal{C}_{\mathcal{L}}}$.

Lemma 4: If $\lambda \in \mathcal{X}_{\mathcal{L}}^{\circ}$, then there exists $\delta_1 > 0$ such that for every $0 < \delta < \delta_1$, there exists $\alpha = [\alpha_1 \cdots \alpha_L]$ that satisfies

$$\begin{aligned} \sum_{\ell=1}^L \alpha_{\ell} R_{e_k^A, A}(\ell) &= \lambda_A + (k+1)\delta, \quad \forall k < N_A \text{ and } A \in \mathcal{S} \\ \sum_{\ell=1}^L \alpha_{\ell} &= 1, \text{ and } \alpha_{\ell} \geq 0 \text{ for every } \ell. \end{aligned}$$

Proof: The result follows from Assumption 1, and the fact that for every $\lambda \in \mathcal{X}_{\mathcal{L}}^{\circ}$, there exists $\delta_1 > 0$ and such that for every $0 < \delta < \delta_1$, $\lambda + \delta \in \mathcal{X}_{\mathcal{L}}^{\circ}$. ■

For rest of the section, fix any $\lambda \in \mathcal{X}_{\mathcal{L}}^{\circ}$. Also, fix $\delta_1 > 0$ that satisfies the conditions in Lemma 4. Now, we define a randomized policy $\Delta_1(\delta) \in \overline{\mathcal{C}_{\mathcal{L}}}$ as follows. Under $\Delta_1(\delta)$, $\ell^{\Delta_1(\delta)}(t) = \ell$ w.p. α_{ℓ} in every slot t independent of the queue lengths and the decisions in the previous slots. The vector α is a solution of the following Linear Program (LP) for $0 < \delta < \delta_1$.

LP(δ) :- Minimize: $U(\delta) = \sum_{\ell=1}^L \alpha_{\ell} f(\ell)$

Subject to:

- 1) $\sum_{\ell=1}^L \alpha_{\ell} R_{e_k^A, A}(\ell) = \lambda_A + (k+1)\delta$ for every $k < N_A$ and $A \in \mathcal{S}$
- 2) $\sum_{\ell=1}^L \alpha_{\ell} = 1$, and $\alpha_{\ell} \geq 0$ for every ℓ .

Lemma 5: For every $\epsilon > 0$ there exists $\delta_{\epsilon} > 0$ such that for every $\delta < \delta_{\epsilon}$, $F_{\min}^{\mathcal{C}_{\mathcal{L}}(\lambda)} \geq U(\delta) - \epsilon$ w.p. 1.

Proof: First, we show that $U(0) \leq F_{\min}^{\mathcal{C}_{\mathcal{L}}(\lambda)}$ w.p. 1. Let $\Delta \in \overline{\mathcal{C}_{\mathcal{L}}}(\lambda)$. Such Δ exists because of Lemma 4. Fix any non-trivial sample path. Let γ_{ℓ} denote the fraction of time ℓ is scheduled

by Δ . Clearly, by stationarity of Δ ,

$$\sum_{\ell=1}^L \gamma_\ell = 1, \text{ and } \gamma_\ell \geq 0 \text{ for every } \ell. \quad (2)$$

Then, by stability of Δ , we also know that

$$\lambda_A = \sum_{\ell=1}^L \gamma_\ell R_{e_k^A, A}(\ell), \forall k < N_A \text{ and } A \in \mathcal{S}, \quad (3)$$

$$F^\Delta = \sum_{\ell=1}^L \gamma_\ell f(\ell). \quad (4)$$

Equations (2) and (3) show that γ is a feasible solution for $\mathbf{LP}(\mathbf{0})$. Moreover, (4) is the objective function for $\mathbf{LP}(\mathbf{0})$. Since Δ is an arbitrary policy in $\overline{\mathcal{C}_\mathcal{L}}(\boldsymbol{\lambda})$, we conclude that

$$U(0) \leq F_{\min}^{\overline{\mathcal{C}_\mathcal{L}}(\boldsymbol{\lambda})} \text{ w.p. } 1. \quad (5)$$

Since the feasible set of $\mathbf{LP}(\delta)$ is convex and compact, and $f(\ell)$ is a bounded function, by continuity, we conclude that $U(\delta) \rightarrow U(0)$ as $\delta \rightarrow 0$. Thus, for every $\epsilon > 0$ there exists $\delta_\epsilon > 0$ such that for every $\delta < \delta_\epsilon$, $U(\delta) \leq U(0) + \epsilon$. From (5), we conclude that

$$U(\delta) \leq F_{\min}^{\overline{\mathcal{C}_\mathcal{L}}(\boldsymbol{\lambda})} + \epsilon \text{ w.p. } 1. \quad (6)$$

Now, since $\mathcal{C}_\mathcal{L}(\boldsymbol{\lambda}) \subseteq \overline{\mathcal{C}_\mathcal{L}}(\boldsymbol{\lambda})$, the result follows. \blacksquare

Note that $F^{\Delta_1(\delta)} = U(\delta)$. Thus, Lemma 5 forges the first link between the costs under policies in $\mathcal{C}_\mathcal{L}(\boldsymbol{\lambda})$ and that under $\overline{\mathcal{C}_\mathcal{L}}(\boldsymbol{\lambda})$. Let us denote the drift in the backlog of $q_{k,A}$ in slot t under Δ as $\partial R_{k,A}^\Delta(t)$, i.e.,

$$\partial R_{k,A}^\Delta(t) = \begin{cases} R_{e_{k-1}^A, A}(\ell^\Delta(t)) - R_{e_k^A, A}(\ell^\Delta(t)) & : k \neq 0, N_A \\ \Lambda_A(t) - R_{e_k^A, A}(\ell^\Delta(t)) & : k = 0 \\ 0 & : k = N_A. \end{cases}$$

Now, consider any $\Delta \in \mathcal{C}_\mathcal{L}$ and observe that

$$Q_{a_k, A}^\Delta(t+1) = \max \{Q_{a_k, A}^\Delta(t) + \partial R_{k,A}^\Delta(t), 0\}. \quad (7)$$

$$\text{Let } \xi^\Delta(t) \stackrel{\text{def}}{=} \sum_{k,A} \left[(Q_{a_k, A}^\Delta(t+1))^2 - (Q_{a_k, A}^\Delta(t))^2 \right].$$

With (7) and some elementary algebra, it follows that

$$\begin{aligned} \mathbb{E}[\xi^\Delta | \mathbf{Q}] &\leq Z + 2 \sum_A Q_{a_0, A} \lambda_A - 2\kappa \mathbb{E}[f(\ell^\Delta) | \mathbf{Q}] \\ &\quad - 2\mathbb{E} \left[\sum_{k, A} R_{e_k^A, A}(\ell^\Delta) \partial Q_{k, A} - \kappa f(\ell^\Delta) | \mathbf{Q} \right], \end{aligned} \quad (8)$$

where $Z = |\mathcal{S}|(c^2 + R_{\max}^2)$. We have omitted t for notational simplicity. Then,

Lemma 6: Given the queue lengths, $\Delta^*(\kappa)$ maximizes the last term in (8) among all the policies in $\overline{\mathcal{C}}_{\mathcal{L}}$.

Now, from Lemma 6 and (8), we conclude that

$$\begin{aligned} \mathbb{E}[\xi^{\Delta^*(\kappa)} | \mathbf{Q}] &\leq Z + 2 \sum_A Q_{a_0, A} \lambda_A - 2\kappa \mathbb{E}[f(\ell^{\Delta^*(\kappa)}) | \mathbf{Q}] \\ &\quad - 2\mathbb{E} \left[\sum_{k, A} R_{e_k^A, A}(\ell^{\Delta_1(\delta)}) \partial Q_{k, A} - \kappa f(\ell^{\Delta_1(\delta)}) | \mathbf{Q} \right]. \end{aligned} \quad (9)$$

Since the choice of schedule is independent of the queue lengths under $\Delta_1(\delta)$, it follows that

$$\begin{aligned} \mathbb{E}[f(\ell^{\Delta_1(\delta)}) | \mathbf{Q}] &= \mathbb{E}[f(\ell^{\Delta_1(\delta)})] = \sum_{\ell} \alpha_{\ell} f(\ell) = U(\delta). \\ \mathbb{E}[R_{e_k^A, A}(\ell^{\Delta_1(\delta)}) | \mathbf{Q}] &= \sum_{\ell} \alpha_{\ell} R_{e_k^A, A}(\ell) = \lambda_A + (k+1)\delta. \end{aligned}$$

Substituting the above quantities in (9), we obtain

$$\begin{aligned} \mathbb{E}[\xi^{\Delta^*(\kappa)} | \mathbf{Q}] &\leq Z - 2\kappa \mathbb{E}[f(\ell^{\Delta^*(\kappa)}) | \mathbf{Q}] - 2\delta \sum_{k, A} Q_{a_k, A} + 2\kappa U(\delta). \end{aligned} \quad (10)$$

Note that the process $\{\mathbf{Q}^{\Delta^*(\kappa)}(t)\}_{t \geq 1}$ is a Markov chain. Thus, to show stability under $\Delta^*(\kappa)$, it suffices to show that the queue length process is positive recurrent.

Lemma 7: For every $\boldsymbol{\lambda} \in \mathcal{X}_{\mathcal{L}}^{\circ}$, $\{\mathbf{Q}^{\Delta^*(\kappa)}(t)\}_{t \geq 1}$ is positive recurrent for every $\kappa < \infty$.

Proof: Note that $\mathbb{E}[\xi^{\Delta^*(\kappa)} | \mathbf{Q}]$ denote the expected Lyapunov drift. From (10), $\kappa < \infty$ and finite support of $f(\cdot)$, it follows that $\mathbb{E}[\xi^{\Delta^*(\kappa)} | \mathbf{Q}] < \infty$ for every \mathbf{Q} . Moreover, $\mathbb{E}[\xi^{\Delta^*(\kappa)} | \mathbf{Q}] < -1$ whenever $\sum_{k, A} Q_{a_k, A} > (Z + \kappa U(\delta) + 1)/2\delta$. Thus, the positive recurrence follows from Foster's Theorem. \blacksquare

Lemma 8: For all $\boldsymbol{\lambda} \in \mathcal{X}_{\mathcal{L}}^{\circ}$, $F^{\Delta^*(\kappa)}(\boldsymbol{\lambda}) \leq \frac{Z}{2\kappa} + U(\delta)$ w.p. 1.

Proof: From Lemma 7, for every $\lambda \in \mathcal{X}_{\mathcal{L}}^{\circ}$, the queue length is stationary under $\Delta^*(\kappa)$. Thus, the result follows by taking the expectation in (10) with respect to stationary distribution of the queue length process, and observing from the Renewal Reward Theorem that $\mathbb{E}[f(\ell^{\Delta^*(\kappa)})] = F^{\Delta^*(\kappa)}(\lambda)$ w.p. 1. ■

Now, we prove Theorems 1 and 2.

A. Proof of Theorem 1

Proof: (Theorem 1) From Lemma 7, $\mathcal{X}_{\mathcal{L}}^{\circ} \subseteq \Lambda^{\Delta^*(\kappa)} \subseteq \Lambda_{\mathcal{L}}$ for every $\kappa \geq 0$. Thus, the result follows from Lemma 3. ■

B. Proof of Theorem 2

Proof: (Theorem 2) From Lemma 7, $\Delta^*(\kappa)$ is throughput optimal for every $\kappa < \infty$. Now, we show ϵ -optimality. Fix $\epsilon > 0$. From Lemma 5, choose $\delta > 0$ such that $U(\delta) - F_{\min}^{\mathcal{C}_{\mathcal{L}}}(\lambda) \leq \epsilon/2$. Also, choose $\hat{\kappa}$ such that $Z/\hat{\kappa} = \epsilon$. Now, ϵ -optimality follows from Lemma 8 for every $\kappa > \hat{\kappa}$. ■

APPENDIX II

PROOF OF LEMMAS 1 AND 2

A. Proof of Lemma 1

Proof: We sketch the proof. Let $\{P_k^A\}_{k \geq 1}$ denote the ordered sequence of session A 's packets, i.e., P_k^A is transmitted by $s(A)$ before P_{k+1}^A and after P_{k-1}^A . Let $P(k) = P_1 \oplus \dots \oplus P_{m(k)}$ denote the first packet containing P_k^A arriving at $d(A)$ in slot $t(k)$. Because of FIFO service, clearly, $t(k) \leq t(k+1)$ for every $k \geq 1$. Now, the result follows from the following claim: For every $k \geq 1$, there does not exist $u \in \{1, \dots, m(k)\}$ and $v > k$ such that $P_v^A = P_u$. Thus, $P(k)$ can only contain packets from session A' or the packets transmitted before P_k^A . The correctness follows by induction on k . ■

B. Proof of Lemma 2

Proof: Consider a NC scheme that allows XOR-ing of packets from sessions A and B at node i , where $B \neq A'$ and $i \in \mathcal{R}_A \cap \mathcal{R}_B$. Also, without loss of generality, let $s(B) \neq d(A)$, and $\mathcal{R}_A = \{a_0, \dots, a_{N_A}\}$ and $\mathcal{R}_B = \{b_0, \dots, b_{N_B}\}$ with $a_k = b_j = i$ for some k and j . Furthermore,

let P_A be the first packet arriving in $q_{s(A),A}$ at time t . Now, our aim is to construct a sequence of valid schedules such that the first packet containing P_A arriving at $d(A)$ can not be decoded correctly. Then, the result will follow from Definition 2. The construction is as follows: First, let $q_{i,B}$ is empty at t . Then, find the largest $u < j$ such that $q_{b_u,B}$ is non-empty. If no such u exists, then do not schedule any link until a packet (say P_B) arrives in $q_{s(B),B}$. Note that a new packet will arrive at $s(B)$ in finite time w.p. 1 as $\lambda_B > 0$. When a packet arrives in $q_{s(B),B}$, we can choose $u = 0$. Once the value of u is determined, schedule e_u^B, e_{u+1}^B and so on, one at a time until a packet arrives in $q_{i,B}$. Next, schedule a sequence of links e_0^A, \dots, e_{k-1}^A one at a time so that P_A arrives in $q_{i,A}$. Now, multicast $P_A \oplus P_B$ on e_k^A and e_j^B , which is possible as NC allows XOR-ing of these packet at i . Finally, schedule $e_{k+1}^A, \dots, e_{N_A-1}^A$ one at a time so that the XOR-ed packet $P_A \oplus P_B$ arrives at $d(A)$. Now, note that because of the FIFO service, P_B is not available at $d(A)$. Thus, P_A can not be recovered at $d(A)$ upon its arrival. This proves the required. ■