

Dynamic Algorithms for Multicast with Intra-session Network Coding

Tracey Ho, *Member, IEEE* and Harish Viswanathan, *Senior Member, IEEE*

***Abstract*—The problem of multiple multicast sessions with intra-session network coding in time-varying networks is considered. The network-layer capacity region of input rates that can be stably supported is established. Dynamic algorithms for multicast routing, network coding, power allocation, session scheduling, and rate allocation across correlated sources, which achieve stability for rates within the capacity region, are presented. This work builds on the back-pressure approach introduced by Tassiulas et al., extending it to network coding and correlated sources. In the proposed algorithms, decisions on routing, network coding, and scheduling between different sessions at a node are made locally at each node based on virtual queues for different sinks. For correlated sources, the sinks locally determine and control transmission rates across the sources. The proposed approach yields a completely distributed algorithm for wired networks. In the wireless case, power control among different transmitters is centralized while routing, network coding, and scheduling between different sessions at a given node are distributed.**

***Index Terms*—Back pressure, correlated sources, multi-cast, multi-hop, network coding, scheduling**

This work was done while Tracey Ho was at Lucent Technologies Bell Labs. She is currently with the California Institute of Technology, Pasadena, CA, e-mail: tho@caltech.edu.

Harish Viswanathan is with Alcatel-Lucent, Bell Labs, Murray Hill, NJ, e-mail: harishv@alcatel-lucent.com

I. INTRODUCTION

Network coding has recently been shown to improve performance compared to that of routing for multicasting information over wired and wireless networks [1], [16], [29]. Most of the work in network coding to date assumes a flow model for transmission in which sources generate, at fixed rates, data that is then transmitted over a network with fixed link capacities. However, in real networks, traffic is usually bursty because either the sources generate traffic in bursts or the network nodes employ queuing and scheduling across multiple sessions. In such scenarios, optimal multicasting of information involves not only routing and network coding but also power control and scheduling of different flows on the active links. Furthermore, optimal network coding in this context may depend on the current state of the network - link data rates and buffer occupancy.

Routing, scheduling and power control in networks with bursty traffic has recently received significant attention in the context of wireless networks [2], [15], [21], [22], [23], [24], [28], [30], [31]. Much of the recent work in this area builds on the ideas of [3], [27] that describe algorithms for routing and scheduling flows using queue sizes, or differences in queue size between the queues at the source and the destination of a link, as the metric

to select between different flows. Such an approach is usually said to be *back-pressure* based since heavily loaded nodes downstream push back and slow down the flow coming down from nodes upstream. Such a back-pressure approach is generally optimal in the sense that it allows transmission at the maximum possible arrival rates into the network for which the queues at the various network nodes are still stable. Furthermore, the algorithms are distributed in the wired network case in the sense that decisions are made locally at each node based on feedback from only the immediate destination nodes of the transmission links at the node.

While the back-pressure approach has mostly been applied in the context of unicast transmissions, it has also been extended to the case of multicast transmissions [2], [25]. However, in the multicast case without network coding the algorithms are significantly more complex, even for wired networks.

We extend the above back-pressure based dynamic control algorithms to include network coding and correlated sources. Unlike in the case of a fixed set of flows, for time-varying queuing networks employing the back-pressure approach, network coding also needs to be dynamic and dependent on the state of the network. Random linear network coding [10], introduced for the flow model, extends naturally to a time-varying network with bursty traffic. In this paper we consider dynamic multi-session multicast with network coding in wired and time-varying wireless networks. We consider the case in which each multicast session consists of a set of sources and sinks such that data from all the sources is intended for all the sinks. We establish the capacity region of input rates that can be stabilized with intra-session network coding, and present dynamic algorithms

for routing, network coding, session scheduling and power control that achieve stability for rates within this capacity region. This is a network-layer rather than an information theoretic capacity; it refers to the maximum achievable rates under a given physical layer modulation and coding scheme.

The multicast algorithm described in [25] involves enumeration of all multicast trees used, while that in [2] involves maintaining a virtual queue for every subset of sinks for every session. In our approach, each node has just one virtual queue for each sink of each session (for independent sources) or for each source-sink pair of each session (for correlated sources). Routing, network coding and session scheduling decisions are made locally by comparing, for each link, the difference in length of corresponding virtual queues, summed over each session's queues. For correlated sources, the sinks locally control rate allocation across the sources. This gives a completely distributed algorithm for wired networks; in the wireless case, power control among interfering transmitters is done centrally.

This paper is organized as follows. We present the system model, discuss some network coding considerations and define the notion of stability for multicast in Section II. The characterization of the capacity region and the capacity achieving back-pressure algorithm are presented in Sections IV and V for independent sources on wired and the wireless networks, respectively. In Section VI, we treat the general case where each multicast session consists of multiple correlated sources. We conclude with a summary and discussion of future work in Section VII. Portions of this work have appeared in [14].

II. PRELIMINARIES

A. Network model

Our network model is based on that in [24]. We consider a network comprising a set \mathcal{N} of $N = |\mathcal{N}|$ nodes, with a set \mathcal{L} of communication links between them that are fixed or time-varying according to some specified processes. There is a set of multicast sessions \mathcal{C} sharing the network. Each session $c \in \mathcal{C}$ is associated with a set $\mathcal{S}_c \subset \mathcal{N}$ of source nodes, and an arrival process, at each source node, of exogenous session c packets to be transmitted to each of a set $\mathcal{T}_c \subset \mathcal{N} \setminus \mathcal{S}_c$ of sink nodes. We denote by σ_{max} and τ_{max} the maximum number of sources and sinks respectively of any given multicast session.

Time is assumed to be slotted, with the time unit normalized to the slot length. For simplicity, we assume fixed length packets and link transmission rates that are restricted to integer multiples of the packet-length/time-slot quotient. That is, an integer number of packets can be transmitted in each slot. We assume that the channel conditions are fixed over the duration of a slot, and known at the beginning of the slot. For simplicity of exposition we assume that the exogenous packet arrival and channel processes are independent and identically distributed across slots.

We consider both wired and wireless networks. In our model, for the wired case the network connectivity and the link rates are explicitly specified. For wireless networks, the network connectivity and link transmission rates depend on the transmitted signal power, interference powers, and the channel propagation conditions. This is described in more detail in Section V-A.

In the remainder of this paper, all link, source and flow rates are in packets per unit time unless otherwise stated.

For both wired and wireless cases, we assume upper bounds μ_{max}^{in} and μ_{max}^{out} on the total flow rate into and out of a node respectively.

B. Network coding model

Network coding within a multicast session allows traffic for different sinks of a session to share network capacity [1]. How to determine or achieve the (network-layer) capacity of a network for multiple multicast sessions with coding across sessions is an open question. Thus, we consider the simpler case where coding is done only across packets of the same session.

We use the approach of distributed random linear network coding, described in [6], [10], [13] for independent sources and in [11] for correlated sources, in which each packet transmitted by a node is a random linear combination of the node's previously received input packets, i.e. packets received on incoming links and exogenous source packets, if any. The contents of each packet, as a linear combination of the input packets, are specified by a *coefficient vector* in the packet header, updated by applying to the coefficient vectors the same linear transformations as to the data. The coefficient vector is thus a function of the random code coefficients specifying the linear combinations at intermediate nodes.

For simplicity, in the descriptions and analysis of the policies in this paper, we do not explicitly consider batch restrictions and the overhead of transmitting coding coefficients. An analysis of the capacity loss from having multiple batches of given size would require more detailed source and channel statistics, and is beyond the scope of this paper.

III. INDEPENDENT SOURCES CASE: PROBLEM AND APPROACH

We start with the case where the information at different source nodes is independent. If the sources have correlated information, the policies developed for this case can be used but are suboptimal; higher rates can be achieved using the approach in Section VI which takes the source correlation into account.

Let x_i^c be the average arrival rate of exogenous session c packets at each node i . We describe the capacity region of rates (x_i^c) that are stabilizable in a network with intra-session network coding, and give a dynamic control policy that achieves rates strictly within this region.

A. Network coding and virtual queues

The bits in each packet are grouped into vectors of length m which are viewed as symbols from the finite field \mathbb{F}_q , $q = 2^m$. The k th symbol of a transmitted packet is a scalar linear function, in \mathbb{F}_q , of the k th symbol of each of its constituent packets, and this function is the same for all k . A sink is able to decode when it receives a full set of packets with linearly independent coefficient vectors. We extend the error probability analysis given in [10] for static networks to the case of time-varying packet networks, by viewing source packets and transmitted packets in the time-varying case as analogous to sources and links respectively in the static case. A formal statement of the result is presented in Theorem 3 in Section IV-C.

For network coding within a multicast session, a solution is given by a union of individual flow solutions for each sink [1]. Our approach for developing the capacity region and the capacity-achieving control policy defines *virtual queues* to keep track of the individual sinks' flow

solutions as follows.

Each node i conceptually maintains, for each sink β of each session c , a virtual queue $Q_i^{c\beta}$ whose length $U_i^{c\beta}$ is the number of session c packets queued at node i that are intended for sink β . A single physical session c packet corresponds to a packet in the virtual queue $Q_i^{c\beta}$ of each sink β for which it is intended. An illustrative example is given in Figure 1. Each packet in a virtual queue corresponds to a distinct physical packet; thus there is a one-to-many correspondence between physical packets and packets in virtual queues.

A packet in a virtual queue $Q_i^{c\beta}$ can be transferred over a link (i, j) to the corresponding virtual queue $Q_j^{c\beta}$ at the link's end node j ; this is called a virtual transmission. With network coding, for any subset $\mathcal{T}' \subset \mathcal{T}_c$ of a session's sinks, a single physical transmission of a packet on a link (i, j) can simultaneously accomplish, for each sink $\beta \in \mathcal{T}'$, one virtual transmission from $Q_i^{c\beta}$ to $Q_j^{c\beta}$. The physically transmitted packet is a random linear coded combination of the physical packets corresponding to the virtually transmitted packets. In the case of a wireless broadcast transmission from a node a to a set of nodes Z , although the nodes in Z all receive the transmitted packet, they update their virtual queues selectively, according to the control algorithms described in subsequent sections¹, such that each constituent virtual transmission is point-to-point, i.e. from one queue $Q_a^{c\beta}$ to one queue $Q_b^{c\beta}$ at some end node $b \in Z$, which may differ for different sinks β . An illustration is given in Figure 2. Thus, there is conservation of virtual packets (virtual flows); we can draw an analogy with the no-coding case where physical packets (physical flows) are

¹This can be realized in practice by including in each packet a list of queues that should be incremented by each node in Z .

conserved.

B. Multicast stability condition

Intuitively, a network is stable if all its queues remain bounded. Let $U_i^c(t)$ be the number of physical session c packets queued at node i at time t . We define stability as in [24], in terms of “overflow” functions

$$\gamma_i^c(M) = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t \Pr\{U_i^c(t') > M\} \quad (1)$$

$$\gamma_{sum}(M) = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t \Pr\{\sum_{i,c} U_i^c(t') > M\}. \quad (2)$$

The session c queue at node i is considered stable if $\gamma_i^c(M) \rightarrow 0$ as $M \rightarrow \infty$. A network of queues is considered stable if each individual queue is stable. The following lemma from [24, Lemma 6, page 100] gives a necessary condition for stability of a network:

Lemma 1: (a) As $M \rightarrow \infty$, $\gamma_i^c(M) \rightarrow 0 \forall i, c$ if and only if $\gamma_{sum}(M) \rightarrow 0$.

(b) If the network is stable, then for any $\delta > 0$, there exists a finite value M such that arbitrarily large times \tilde{t} can be found for which $\Pr\{\sum_{i,c} U_i^c(\tilde{t}) \leq M\} > 1 - \delta$. For $\delta = 1/2$, there exists a value M such that the probability that $U_i^c(\tilde{t}) \leq M$ simultaneously for all i, c is greater than $1/2$ infinitely often.

We define a similar overflow function for the sum of the virtual queue lengths $U_i^{c\beta}(t)$:

$$\gamma_{virtualsum}(M) = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t \Pr\{\sum_{i,c,\beta} U_i^{c\beta}(t') > M\} \quad (3)$$

and extend the above lemma to the case with virtual queues.

Lemma 2: (a) As $M \rightarrow \infty$, $\gamma_i^c(M) \rightarrow 0 \forall i, c$ if and only if $\gamma_{virtualsum}(M) \rightarrow 0$.

(b) If the network is stable, then for any $\delta > 0$, there exists a finite value M such that arbitrarily large times

\tilde{t} can be found for which $\Pr[\sum_{i,c,\beta} U_i^{c\beta}(\tilde{t}) \leq M] > 1 - \delta$. For $\delta = 1/2$, there exists a value M such that the probability that $U_i^{c\beta}(\tilde{t}) \leq M$ simultaneously for all i, c, β is greater than $1/2$ infinitely often.

Proof: The proof approach follows that in [20, Lemma 1, page 30].

(a) Since each physical packet included in $U_i^c(t)$ corresponds to a virtual packet in one or more of the corresponding virtual queues whose lengths are $U_i^{c\beta}(t)$, we have, for all i, c, t ,

$$\begin{aligned} U_i^c(t) &\leq \sum_{i,c,\beta} U_i^{c\beta}(t) \leq NK\tau_{max}U_i^c(t) \\ \Rightarrow \Pr\{U_i^c(t) > M\} &\leq \Pr\{\sum_{i,c,\beta} U_i^{c\beta}(t) > M\} \\ &\leq \sum_{i,c} \Pr\{U_i^c(t) > M/NK\tau_{max}\} \end{aligned}$$

where N and K are the number of nodes and sessions respectively and τ_{max} is the maximum number of sinks of a session. The last inequality follows from noting that $\sum_{i,c,\beta} U_i^{c\beta}(t) > M$ implies that $U_i^c(t) > M/NK\tau_{max}$ for some i, c and using the union bound. It follows that for all i, c ,

$$\gamma_i^c(M) \leq \gamma_{virtualsum}(M) \leq \sum_{i,c} \gamma_i^c(M/NK\tau_{max})$$

where the last inequality holds because the lim sup of a sum is less than or equal to the sum of the lim sups. Part (a) follows by taking limits as $M \rightarrow \infty$.

(b) From part (a), if the network is stable, then there exists a value M such that $\gamma_{virtualsum}(M) < \delta$ for arbitrarily small δ . This means that for any arbitrarily large time t_1 , from (3), there exists a value in $\tilde{t} \geq t_1$ for which $\Pr[\sum_{i,c,\beta} U_i^{c\beta}(\tilde{t}) > M]$ is less than δ . ■

C. Redundancy coding at sources

Each source node $i \in \mathcal{S}_c$ forms coded source packets at an average rate $\lambda_i^c = x_i^c + \epsilon$, where x_i^c is its exogenous

packet arrival rate and $\epsilon > 0$. Each coded source packet is formed as an independent random linear combination of previously arrived exogenous packets, and is “added” to each queue $Q_i^{c\beta}, \beta \in \mathcal{T}_c$. Here we provide some intuition for this in advance of the detailed development in subsequent sections. The decoding condition at each sink is that the total number of linearly independent coded packets reaching each sink up to some time t should be equal to the total number of exogenous packets up to that time. The rate of coded source packets is ϵ greater than the rate of exogenous packets to allow for a small fraction of coded packets to be delayed in network queues; this fraction tends to decrease with time since the queue sizes are bounded. Thus, the choice of ϵ trades off between closeness to capacity and the potential decoding delay; for sufficiently large decoding delay t , the decoding condition is satisfied and decoding is successful with high probability.² The number $A_i^c(t)$ of coded session c source packets formed at node i in timeslot t is i.i.d. with expected value

$$\lambda_i^c = E\{A_i^c(t)\}. \quad (4)$$

The second moment of the total number of source packets formed at each node in each timeslot is bounded by a finite maximum value A_{\max}^2 , i.e.

$$E \left\{ \left(\sum_c A_i^c(t) \right)^2 \right\} \leq A_{\max}^2. \quad (5)$$

IV. WIRED NETWORKS

We first describe the capacity region and back pressure policy for independent sources on wired networks, deferring proofs of the results to the next section, which

²If we employ batch coding where each batch contains a fixed number of exogenous packets, feedback from the sinks can be used to signal when the sources should stop forming coded packets of each batch. This determines the effective value of ϵ for each batch.

generalizes these results to the wireless case. We present these results separately for the wired case as they are simpler and provide useful intuition. The main difference between the wired and wireless scenarios is that in a wired network all links are point-to-point links with fixed transmission rates, whereas in a wireless network, links could be point-to-multipoint with mutually dependent transmission rates.

A. Notation

We denote by r_{ab} the capacity of link (a, b) . We use $f_{ab}^{c\beta}$ to denote average virtual flow rate, over link $(a, b) \in \mathcal{L}$, from $Q_a^{c\beta}$ to $Q_b^{c\beta}$. We use g_{ab}^c to denote average physical flow rate for session c over $(a, b) \in \mathcal{L}$. For brevity of notation, we use the convention that any term with subscript ab equals zero unless $(a, b) \in \mathcal{L}$, and any term with superscript $c\beta$ equals zero unless $c \in \mathcal{C}, \beta \in \mathcal{T}_c$.

B. Capacity region with intra-session network coding

Let Λ be the set of all source rate vectors (λ_i^c) such that there exist variables $\{f_{ab}^{c\beta}, g_{ab}^c\}$ satisfying:

$$f_{\beta b}^{c\beta} = 0 \quad \forall b, c, \beta \quad (6)$$

$$f_{ab}^{c\beta} \geq 0 \quad \forall a, b, c, \beta \quad (7)$$

$$\lambda_i^c \leq \sum_b f_{ib}^{c\beta} - \sum_a f_{ai}^{c\beta} \quad \forall i, c, \beta \neq i, \beta \in \mathcal{T}_c \quad (8)$$

$$f_{ab}^{c\beta} \leq g_{ab}^c \quad \forall a, b, c, \beta \quad (9)$$

$$\sum_c g_{ab}^c \leq r_{ab} \quad (10)$$

The variables $\{f_{ab}^{c\beta}\}$ for a (session, sink) pair $(c, \beta \in \mathcal{T}_c)$ define a flow carrying rate at least λ_i^c from each source node i to β (by (7)-(8)), in which virtual flow that is intended for β is not retransmitted away from β (Equation (6)). Network coding allows flows for different

sinks of a common multicast session to share capacity by being coded together [1], so the total usage g_{ab}^c of link (a, b) by session c need only be as large as the maximum virtual usage $f_{ab}^{c\beta}$ by individual sinks $\beta \in \mathcal{T}_c$ of the session (Inequality (9)). The flow constraints given above provide a characterization of the capacity region as shown in Theorem 1 below.

Theorem 1: (a) A necessary condition for stability of multiple multicast sessions with intra-session network coding is $(\lambda_i^c) \in \Lambda$.

(b) A sufficient condition for stability is that (λ_i^c) is strictly interior to Λ .

C. Achievability

The following back pressure policy stabilizes the network for all input rates within the capacity region. It is a special case of the back pressure policy for wireless networks described and analyzed in the next section. The intuition behind the policy is that it chooses, for each link at each time slot, the session with the maximum total weight of virtual transmissions, summed over the session's sinks.

Back-pressure policy for wired networks

For each time slot t and each link (a, b) ,

- Session scheduling: one session

$$c_{ab}^* = \arg \max_c \left\{ \sum_{\beta \in \mathcal{T}_c} \max(U_a^{c\beta} - U_b^{c\beta}, 0) \right\}$$

is chosen.

- Rate allocation: the maximum rate of virtual transmissions from $Q_a^{c\beta}$ to $Q_b^{c\beta}$ is set as

$$\mu_{ab}^{c\beta}(t) = \begin{cases} r_{ab} & \text{if } c = c_{ab}^*, \beta \in \mathcal{T}_c, \\ & \text{and } U_a^{c\beta} - U_b^{c\beta} > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

- Network coding: each session c packet physically transmitted on link (a, b) is a random linear combination, in \mathbb{F}_q , of packets corresponding to a set of virtual transmissions on (a, b) each associated with a different sink in \mathcal{T}_c , as described in Section III-A. The subset of \mathcal{T}_c associated with virtual transmissions consists of the destinations β for which $U_a^{c\beta} - U_b^{c\beta} > 0$.

Let

$$\bar{U} := \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^{t-1} E\{U(t')\}.$$

Theorem 2: If coding rates (λ_i^c) are such that $(\lambda_i^c + \epsilon') \in \Lambda, \epsilon' > 0$, the back-pressure policy stabilizes the system with average total virtual queue length bounded as

$$\sum_{i,c,\beta} \bar{U}_i^{c\beta} = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^{t-1} \sum_{i,c,\beta} E\{U_i^{c\beta}(t')\} \leq \frac{BN}{\epsilon'}$$

where N is the number of network nodes and

$$B = \frac{\tau_{max}}{2} ((A_{max} + \mu_{max}^{in})^2 + (\mu_{max}^{out})^2).$$

Theorems 1 and 2 are special cases of the corresponding theorems for wireless networks that are stated and proved in the next section. While the above result shows that each sink can receive packets at a rate arbitrarily close to the source rate, the packets received contain randomly coded data. In order to retrieve the actual information, each sink decodes the information from the coded packets that it receives. The following theorem shows that the probability that not all sinks are able to decode the information tends to zero exponentially in the coding block length.

Theorem 3: For exogenous arrival rates $x_i^c = \lambda_i^c - \epsilon$, if (λ_i^c) is strictly interior to Λ , then for sufficiently large times t , the probability that not every sink is

able to decode its session's exogenous packets decreases exponentially in the length of the code.

Proof: We can cast a given sequence \mathcal{P} of packet transmissions in the algebraic network coding framework of [16], [10] as follows. We consider a corresponding static network \mathcal{G} with the same node set \mathcal{N} and with links corresponding to transmissions in \mathcal{P} , where for each packet p transmitted from node v to w in \mathcal{P} , \mathcal{G} has one unit-capacity link \tilde{p} from v to w . This has similarity with the notion of a time-expanded network introduced in [1], where a fixed network is replicated in time to analyze a sequence of transmissions in time. The coding coefficient specifying the mapping from link \tilde{p}_1 to \tilde{p}_2 is set to zero if packet p_1 is not among the packets coded together to form p_2 . The following analysis is an extension, based on this correspondence, of the analysis in [16], [10].

Consider any session c . Let $\underline{\xi}$ be the vector consisting of the randomly chosen network coding coefficients associated with the session c packets. Consider any sink $\beta \in \mathcal{T}_c$. It follows from Theorem 2 and Lemma 2 that for any $\delta > 0$, there exists a sufficiently large value V such that

$$\Pr \left(\bigcap_{c,\beta} \left\{ \sum_j U_j^{c\beta}(t) \leq V \text{ i.o.} \right\} \right) \geq 1 - \delta.$$

Since the arrival processes of the exogenous packets and coded packets are ergodic, taking into account the last formula it follows that with arbitrary high probability one can choose arbitrarily large $t \geq V/\epsilon$ so that for all c, β, i it holds that $\sum_j U_j^{c\beta}(t) \leq V$, the number n_i^c of session c exogenous packet arrivals at source i during the time interval t is approximately $(\lambda_i^c - \epsilon)t$ and the total number of session c packets added to the source queue at node i is approximately $\lambda_i^c t$. The total number

m of session c packets arrived by the chosen time t to β from each source i is given by these $\lambda_i^c t$ packets minus the number of them left in the (c, β) virtual queues and hence m is lower bounded by

$$\lambda_i^c t - V \geq (\lambda_i^c - \epsilon)t.$$

The physical packets corresponding to this virtual flow are coded combinations of the n_i^c exogenous packets from i . These exogenous packets are in one-to-one correspondence with the source processes in the corresponding static network \mathcal{G} .

Consider any arbitrary fixed choice of n_i^c of the packets received by β from each session c source node i , corresponding to a choice of n_i^c input links at β in the static network \mathcal{G} . This choice fixes only the sequence of which packets are coded together to form the n_i^c chosen packets; by the Principle of Deferred Decisions (see, e.g. [18]), we can assume that the values of the random coding coefficients are chosen subsequently. We denote by $d^{c\beta}(\underline{\xi})$ the determinant, as a polynomial in $\underline{\xi}$, of the matrix whose rows equal the coefficient vectors of these $\sum_{i \in \mathcal{S}_c} n_i^c$ packets. Consider the physical packet transmissions corresponding to this virtual flow, which are transmissions involving queues $Q_j^{c\beta}$. These physical transmissions would constitute an uncoded physical flow if their originating transmissions from the source nodes were uncoded independent packets and there were no other sinks/virtual flows in the network. We denote by $\tilde{\underline{\xi}}$ the value of $\underline{\xi}$ corresponding to this case, noting that $d^{c\beta}(\tilde{\underline{\xi}}) = 1$ or -1 .³ Thus, $d^{c\beta}(\underline{\xi})$ is not identically zero.

Since the product $\prod_{c,\beta \in \mathcal{T}_c} d^{c\beta}(\underline{\xi})$ as a function of the network code coefficients $\underline{\xi}$ is not identically zero, by the Schwartz-Zippel theorem (see, e.g. [18]), choosing

³For this uncoded flow case, the coefficient vectors of the $(\lambda_i^c - \epsilon)t$ session c packets received by β form the rows of the identity matrix.

the code coefficients uniformly at random from a finite field of size q yields a zero value with probability inversely proportional to q . A nonzero value for the product corresponds to each sink having enough linearly independent combinations to decode all data transmitted up to that point. The result follows since q is exponential in the length of the code. ■

V. WIRELESS NETWORKS

A. Wireless model

The wireless case is considerably more complicated, as we take into account interference among signals transmitted simultaneously by multiple nodes and the fact that a single node's transmission can be received by multiple nodes, referred to as wireless multicast advantage. Furthermore, it is possible for a receive node to combine information received from multiple transmit nodes that are transmitting the same signal.

We assume we are given a set \mathcal{L} consisting of point-to-point links (a, b) and/or point-to-multipoint links (a, Z) where $a, b \in \mathcal{N}, Z \subset \mathcal{N}$. While there are potentially a large number of possible wireless links, in practice we consider a limited set of useful links. For example, the set Z of receivers for a link (a, Z) could be chosen to be some number of closest neighbors of the transmitter a . The selection and size of the set \mathcal{L} of links over which optimization occurs trades off achievable performance against complexity, and is an issue inherent to wireless network optimization in general. It is a pertinent practical problem that is considered in, e.g., [19], [29], but is beyond the scope of this paper.

We assume we are given a rate function $\underline{\mu}(\underline{P}, \underline{S})$ specifying the vector of link rates as a function of the vector of transmit powers \underline{P} and channel states \underline{S} .

Consider first a single link (a, b) . The rate in packets per unit time on the link, $\mu_{ab}(\underline{P}, \underline{S})$, could, for example, be modeled according to the Shannon formula (see, e.g. [7])

$$\mu_{ab}(\underline{P}, \underline{S}) = \frac{1}{J} \log_2 \left(1 + \frac{S_{ab} P_a}{N_0 + \sum_{a' \in \mathcal{N} \setminus \{a\}} S_{a'b} P_{a'}} \right), \quad (12)$$

where J is the packet length in bits, a' denotes nodes transmitting interfering signals, N_0 is the additive white Gaussian noise variance, and S_{ab} is channel gain from transmitter node a to receiver node b and includes the effects of distance loss, large scale and small scale fading effects. This is just one possible approximate but useful formula. Note that the link rates could in general be modeled using any monotonically increasing function of the SINR, or could be restricted to a finite set of possible values corresponding to practical coding and modulation schemes; as in [24], we assume only that $\underline{\mu}(\underline{P}, \underline{S})$ is upper semi-continuous in \underline{P} for all states \underline{S} .

The channel state vector $\underline{S}(t)$ is assumed to be constant over each time slot, i.e., state transitions occur only on slot boundaries, where time t is an integer. We assume that in each time slot the value of $\underline{S}(t)$ is taken i.i.d. from a finite set; we denote by $\pi_{\underline{S}}$ the probability of state \underline{S} . Transmit powers $\underline{P}(t)$ are also held constant over each time slot, and each node a 's transmit power $P_a(t)$ is limited by the maximum transmit power P_{\max} . Hence the set of power allocations Π is a compact set. In the rest of the paper, t is taken to be an integer unless otherwise stated.

For a wireless broadcast link (a, Z) from a node a to a set of nodes Z , the instantaneous link rate $\mu_{aZ}(\underline{P}, \underline{S})$ is also determined by the vector of transmit powers $\underline{P}(t)$ and the channel state vector $\underline{S}(t)$. It is related to the

individual link rates according to

$$\mu_{aZ}(\underline{P}, \underline{S}) = \min_{b \in Z} \mu_{ab}(\underline{P}, \underline{S}) \quad (13)$$

B. Notation

Recall from Section III-A that for a wireless transmission over a link (a, Z) , although the nodes in Z all receive the transmitted packet, they update their virtual queues selectively such that each constituent virtual transmission is point-to-point, i.e. from one queue $Q_a^{c\beta}$ to one queue $Q_b^{c\beta}$ at some end node $b \in Z$. We use $f_{abZ}^{c\beta}$ to denote average virtual flow rate, over link $(a, Z) \in \mathcal{L}$, from $Q_a^{c\beta}$ to $Q_b^{c\beta}$, $b \in Z$. We use g_{aZ}^c to denote the average physical flow rate for session c over $(a, Z) \in \mathcal{L}$. We denote by r_{aZ} the average value of the time-varying rate $\mu_{aZ}(\underline{P}(t), \underline{S}(t))$ of link (a, Z) . Precise definitions of $f_{abZ}^{c\beta}, g_{aZ}^c, r_{aZ}$ are given below with each different context in which the variables are used.

For brevity of notation, we use the convention that any variable term with subscript abZ equals zero unless $(a, Z) \in \mathcal{L}, b \in Z$, and any variable term with superscript $c\beta$ equals zero unless $c \in \mathcal{C}, \beta \in \mathcal{T}_c$.

We use (r_{aZ}) to denote the vector of all link rates $(r_{aZ} : (a, Z) \in \mathcal{L})$.

C. Capacity region with intra-session network coding

Similarly to [24], [20], we define the network graph family Γ consisting of all rate vectors $\underline{r} = (r_{aZ})$ that can be represented as $\underline{r} = \sum_{\underline{S}} \pi_{\underline{S}} \underline{r}_{\underline{S}}$ for some set of rate vectors $\underline{r}_{\underline{S}}$, each of which is in the convex hull of the set of rate vectors $\{\underline{\mu}_{aZ}(\underline{P}, \underline{S}) | \underline{P} \in \Pi\}$. Γ can be viewed as the set of long-term average transmission rates (r_{aZ}) supportable by the network. The assumption that $\underline{\mu}(\underline{P}, \underline{S})$ is upper semi-continuous in \underline{P} for all \underline{S} ensures the compactness of the set of all rate matrices entrywise

less than or equal to some element of Γ , and in turn the compactness of the region Λ defined below [20], which is needed in the proof of Theorem 4.

Let Λ be the set of all input rate vectors (λ_i^c) such that there exist values for $\{f_{abZ}^{c\beta}, g_{aZ}^c\}, (r_{aZ}) \in \Gamma$ satisfying:

$$f_{\beta b Z}^{c\beta} = 0 \quad \forall b, c, \beta, Z \quad (14)$$

$$f_{abZ}^{c\beta} \geq 0 \quad \forall a, b, c, \beta, Z \quad (15)$$

$$\lambda_i^c \leq \sum_{b, Z} f_{ibZ}^{c\beta} - \sum_{a, Z} f_{aiZ}^{c\beta} \quad \forall i, c, \beta \neq i, \beta \in \mathcal{T}_c \quad (16)$$

$$\sum_{b \in Z} f_{abZ}^{c\beta} \leq g_{aZ}^c \quad \forall a, c, \beta, Z \quad (17)$$

$$\sum_c g_{aZ}^c \leq r_{aZ} \quad \forall a, Z. \quad (18)$$

The interpretation of these equations is similar to the wired case: variables $\{f_{abZ}^{c\beta}\}$ define a flow solution from the session c sources to β .

Theorem 4: (a) A necessary condition for stability of multiple multicast sessions with intra-session network coding is $(\lambda_i^c) \in \Lambda$.

(b) A sufficient condition for stability is that (λ_i^c) is strictly interior to Λ .

Proof: (a) Let $X_i^c(t)$ be the total number of session c source packets formed at node i up to time t . For an intra-session network coding and power control policy, let $F_{abZ}^{c\beta}(t)$ be the total virtual flow, up to time t , from queue $Q_a^{c\beta}$ to queue $Q_b^{c\beta}$ over link (a, Z) , and let $G_{aZ}^c(t)$ be the total physical flow of session c packets transmitted over link (a, Z) up to time t . These quantities must satisfy

$$F_{abZ}^{c\beta}(t) \geq 0 \quad \forall a, b, c, \beta, Z \quad (19)$$

$$X_i^c(t) = \sum_{b, Z} F_{ibZ}^{c\beta}(t) - \sum_{a, Z} F_{aiZ}^{c\beta}(t) + U_i^{c\beta}(t) \quad \forall i, c, \beta \neq i, \beta \in \mathcal{T}_c \quad (20)$$

$$\sum_{b \in Z} F_{abZ}^{c\beta}(t) \leq G_{aZ}^c(t) \quad \forall a, b, c, \beta, Z \quad (21)$$

$$\sum_c G_{aZ}^c(t) \leq \sum_{t'=0}^t \mu_{aZ}(\underline{P}(t'), \underline{S}(t')) \quad \forall a, Z. \quad (22)$$

Suppose the system is stabilizable with some intra-session network coding and power control policy. Then there exists a stabilizing policy for which

$$F_{\beta bZ}^{c\beta} = 0 \quad \forall b, Z, c, \beta, \quad (23)$$

i.e. virtual flow that has reached its intended sink is not retransmitted away from the sink.

Let $T_{\underline{S}}(t) \in \{0, 1, \dots, t-1\}$ denote the subset of time slots, up to time $t-1$, during which the channel is in state \underline{S} . As $t \rightarrow \infty$, $\frac{X_i^c(t)}{t}$ and $\frac{|T_{\underline{S}}(t)|}{t}$ converge to their time average values λ_i^c and $\pi_{\underline{S}}$ respectively, and the inequalities

$$\left| \frac{X_i^c(t)}{t} - \lambda_i^c \right| \leq \epsilon' \quad \forall i, c \quad (24)$$

$$\frac{|T_{\underline{S}}(t)|}{t} \leq \pi_{\underline{S}} + \epsilon' \quad \forall \underline{S} \quad (25)$$

are simultaneously satisfied with probability tending to 1 since there are a finite number of arrival processes $X_i^c(t)$ and channel states \underline{S} . By Lemma 2, for the stabilizing policy, there exists some finite value M such that at arbitrarily large times t ,

$$\frac{U_i^{c\beta}(t)}{t} \leq \frac{M}{t} \leq \epsilon' \quad \forall i, c, \beta \quad (26)$$

with probability greater than 1/2. Thus, there exists some value \tilde{t} of t such that (24)-(26) are simultaneously satisfied with nonzero probability.

Choosing one such value \tilde{t} , let $f_{abZ}^{c\beta} = F_{abZ}^{c\beta}(\tilde{t})/\tilde{t}$, $g_{aZ}^c = G_{aZ}^c(\tilde{t})/\tilde{t}$. Considering (22) at $t = \tilde{t}$ and dividing by \tilde{t} , we obtain

$$\begin{aligned} \left(\sum_c g_{aZ}^c \right) &\leq \left(\frac{1}{\tilde{t}} \sum_{t'=0}^{\tilde{t}-1} \mu_{aZ}(\underline{P}(t'), \underline{S}(t')) \right) \\ &= \sum_{\underline{S}} \frac{|T_{\underline{S}}(\tilde{t})|}{\tilde{t}} \left(\mu_{aZ}^{\underline{S}} \right) \end{aligned} \quad (27)$$

where vectors $\left(\mu_{aZ}^{\underline{S}} \right) = \frac{1}{|T_{\underline{S}}(\tilde{t})|} \sum_{t' \in T_{\underline{S}}(\tilde{t})} (\mu_{aZ}(\underline{P}(t'), \underline{S}))$ are elements of the convex hull of rate vectors $\{(\mu(\underline{P}, \underline{S})) | \underline{P} \in \Pi\}$, and the vector inequality is considered entrywise.

Substituting (25) into (27) we obtain

$$\left(\sum_c g_{aZ}^c \right) \leq \sum_{\underline{S}} \pi_{\underline{S}} \left(\mu_{aZ}^{\underline{S}} \right) + \epsilon' (\mu_{aZ}^{max}) \text{card}\{\underline{S}\} \quad (28)$$

where $\text{card}\{\underline{S}\}$ is the number of channel states and (μ_{aZ}^{max}) is the maximum transmission rate of link (a, Z) .

Similarly, considering (19)-(21) at \tilde{t} and dividing by \tilde{t} , and using (24)-(26) we obtain

$$f_{\beta bZ}^{c\beta} = 0 \quad \forall b, c, \beta, Z \quad (29)$$

$$f_{abZ}^{c\beta} \geq 0 \quad \forall a, b, c, \beta, Z \quad (30)$$

$$\begin{aligned} \lambda_i^c &\leq \sum_{b, Z} f_{ibZ}^{c\beta} - \sum_{a, Z} f_{aiZ}^{c\beta} + 2\epsilon' \\ &\quad \forall i, c, \beta \neq i, \beta \in \mathcal{T}_c \end{aligned} \quad (31)$$

$$\sum_{b \in Z} f_{abZ}^{c\beta} \leq g_{aZ}^c \quad \forall a, b, c, \beta, Z. \quad (32)$$

Considering (28)-(32) in the limit as $\epsilon' \rightarrow 0$, we see that (λ_i^c) is a limit point of Λ , which is compact and thus contains its limit points.

(b) is proved constructively in the next section. \blacksquare

D. Control policies

We consider policies that make control decisions at the start of each time slot t and operate as follows.

- Power allocation: A vector of transmit powers $\underline{P}(t) = (P_{aZ}(t))$ is chosen from the set Π of feasible power allocations. This, together with the channel state $\underline{S}(t)$, determines the link rates $\underline{\mu}(t) = (\mu_{aZ}(t))$, assumed constant over the time slot.
- Session scheduling, rate allocation and network coding: For each link (a, Z) , each sink β of each session c is allocated a transmission rate $\mu_{abZ}^{c\beta}(t)$

for each destination node $b \in Z$. These allocated rates must satisfy the overall link rate constraint

$$\mu_{aZ}(t) \geq \sum_{c \in \mathcal{C}} \max_{\beta \in \mathcal{T}_c} \sum_{b \in Z} \mu_{abZ}^{c\beta}(t). \quad (33)$$

$\mu_{abZ}^{c\beta}(t)$ gives the maximum rate of virtual transmissions from $Q_a^{c\beta}$ to $Q_b^{c\beta}$ over (a, Z) . Besides this limit on virtual transmissions for pairs of queues over each link, the total number of virtual transmissions out of $Q_a^{c\beta}$ over all links with start node a is also limited by the queue length $U_a^{c\beta}(t)$ at the start of the time slot. Each session c packet physically transmitted on link (a, Z) is a random linear combination, in \mathbb{F}_q , of packets corresponding to a set of virtual transmissions on (a, Z) , each associated with a different sink in \mathcal{T}_c . Thus, the rate allocated to session c on (a, Z) is the maximum, over sinks $\beta \in \mathcal{T}_c$, of each sink β 's total allocated rate $\sum_{b \in Z} \mu_{abZ}^{c\beta}(t)$, which explains (33).

The following dynamic policy relies on queue length information to make control decisions, without requiring knowledge of the input or channel statistics. The intuition behind the policy is that it seeks to maximize the total weight of virtual transmissions for each time slot, subject to the above constraints.

Back-pressure policy

For each time slot t , transmit powers ($P_{aZ}(t)$) and allocated rates ($\mu_{abZ}^{c\beta}(t)$) are chosen based on the queue lengths ($U_i^{c\beta}(t)$) at the start of the slot, as follows.

- Session scheduling: For each link (a, Z) ,
 - for each session c and sink $\beta \in \mathcal{T}_c$, one end node

$$\begin{aligned} b_{aZ}^{c\beta*} &= \arg \max_{b \in Z} \left(U_a^{c\beta} - U_b^{c\beta} \right) \\ &= \arg \min_{b \in Z} U_b^{c\beta} \end{aligned}$$

is chosen. Let $U_{aZ}^{c\beta*}$ denote $U_{b_{aZ}^{c\beta*}}^{c\beta}$ for brevity.

- one session

$$c_{aZ}^* = \arg \max_c \left\{ \sum_{\beta \in \mathcal{T}_c} \max \left(U_a^{c\beta} - U_{aZ}^{c\beta*}, 0 \right) \right\}$$

is chosen. Let

$$w_{aZ}^* = \sum_{\beta \in \mathcal{T}_{c_{aZ}^*}} \max \left(U_a^{c\beta} - U_{aZ}^{c\beta*}, 0 \right) \quad (34)$$

be the weight of the chosen session.

- Power control: The state $\underline{S}(t)$ is observed, and a power allocation

$$\underline{P}(t) = \arg \max_{\underline{P} \in \Pi} \sum_{a,Z} \mu_{aZ}(\underline{P}, \underline{S}(t)) w_{aZ}^* \quad (35)$$

is chosen.

- Link rate allocation: For each link (a, Z) ,

$$\mu_{abZ}^{c\beta}(t) = \begin{cases} \mu_{aZ}(t) & \text{if } c = c_{aZ}^*, \beta \in \mathcal{T}_c, \\ & b = b_{aZ}^{c\beta*} \text{ and } U_a^{c\beta} - U_b^{c\beta} > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (36)$$

The complexity of the optimization (35) depends on the physical network model. If there are enough channels so that different nodes' transmissions do not interfere, the optimization can be done locally at each node. If transmissions by different nodes interfere, solving the optimization requires coordination or a centralized solution. It is shown in subsequent work [9] that for some network models, e.g. an interference set model which specifies sets of links that interfere with each other, there are distributed approximation algorithms that achieve a guaranteed fraction γ of the optimal solution of (35), and stabilize the system for any rate vector (λ_i^c) such that $(\frac{1}{\gamma} \lambda_i^c + \epsilon') \in \Lambda$ (an extension of [17, Proposition 3] for the non-multicast routing case). Alternatively, the optimization (35) can be done heuristically, e.g. by a greedy approach along the lines of that in [19], [29] but

with the added guidance of weights w_{aZ}^* for prioritization among candidate links (a, Z) . However, we do not have performance guarantees for these heuristics.

The stability of the back pressure policy for $(\lambda_i^c + \epsilon') \in \Lambda$ is shown by comparison with a randomized policy that assumes knowledge of the long-term input and channel statistics. We will show that the randomized policy is stable, and that stability of the randomized policy implies stability of the back pressure policy.

Randomized policy

Assume given values of rate vector $(r_{aZ}) \in \Gamma$ and flow variables $\{f_{abZ}^{c\beta}, g_{aZ}^{c\beta}\}$ satisfying:

$$f_{\beta bZ}^{c\beta} = 0 \quad \forall b, c, \beta, Z \quad (37)$$

$$f_{abZ}^{c\beta} \geq 0 \quad \forall a, b, c, \beta \in \mathcal{T}_c, Z \quad (38)$$

$$\lambda_i^c + \epsilon' = \sum_{b,Z} f_{ibZ}^{c\beta} - \sum_{a,Z} f_{aiZ}^{c\beta} \quad \forall c \in \mathcal{C}, \beta, i \neq \beta \quad (39)$$

$$\sum_{b \in Z} f_{abZ}^{c\beta} \leq g_{aZ}^c \quad \forall a, Z, c, \beta, Z \quad (40)$$

$$\sum_c g_{aZ}^c \leq r_{aZ} \quad \forall a, Z. \quad (41)$$

The following lemma, adapted from [24, Lemma 1], shows that for any rate vector $(r_{aZ}) \in \Gamma$, power can be allocated according to the time-varying channel state $\underline{S}(t)$ such that the time average link rates converge to (r_{aZ}) .

Lemma 3: Consider a rate vector $(r_{aZ}) \in \Gamma$. There exists a stationary randomized power allocation policy which gives link rates $\mu_{aZ}(t)$ satisfying

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t \mu_{aZ}(t') = r_{aZ}$$

with probability 1 for all $(a, Z) \in \mathcal{L}$, where, for each time slot t in which channel state $\underline{S}(t)$ takes value \underline{S} , the power allocation is chosen randomly from a finite set $\{\underline{P}_{\underline{S},1}, \dots, \underline{P}_{\underline{S},m}\}$ according to stationary probabilities

$\{q_{\underline{S},1}, \dots, q_{\underline{S},m}\}$.

Proof: The proof follows that of [20, Lemma 8, page 125]. From the definition of Γ in Section V-C and by Carathéodory's Theorem [5], $(r_{aZ}) = \sum_{\underline{S}} \pi_{\underline{S}} \underline{r}_{\underline{S}}$ for some set of rate vectors $\underline{r}_{\underline{S}}$, each of which is a convex combination of vectors in $\{\underline{\mu}_{aZ}(\underline{P}, \underline{S}) | \underline{P} \in \Pi\}$. For each state \underline{S} , the probabilities of the stationary randomized power allocation policy are chosen according to the weights of the corresponding convex combination, which gives $E\{\mu(\underline{P}(t'), \underline{S}(t')) | \underline{S}(t') = \underline{S}\} = \underline{r}_{\underline{S}}$. By the law of large numbers, $\frac{1}{|T_{\underline{S}}(t)|} \sum_{t' \in T_{\underline{S}}(t)} \mu(\underline{P}(t'), \underline{S}(t')) \rightarrow \underline{r}_{\underline{S}}$ as $t \rightarrow \infty$, where $T_{\underline{S}}(t) \subset \{0, 1, \dots, t-1\}$ denotes the subset of time slots, up to time $t-1$, during which the channel is in state \underline{S} . The result follows since $|T_{\underline{S}}(t)|/t \rightarrow \pi_{\underline{S}}$ as $t \rightarrow \infty$. ■

The randomized policy is designed such that

$$E\{\mu_{abZ}^{c\beta}(t)\} = f_{abZ}^{c\beta}. \quad (42)$$

For each time slot t , transmit powers $(P_{aZ}(t))$ and allocated rates $(\mu_{abZ}^{c\beta}(t))$ are chosen based on the given values of $(r_{aZ}), \{f_{abZ}^{c\beta}, g_{aZ}^{c\beta}\}$ as well as the channel state $\underline{S}(t)$, as follows.

- Power allocation: The channel state $\underline{S}(t)$ is observed, and power is allocated according to the algorithm of Lemma 3, giving instantaneous link rates $\mu_{aZ}(t)$ and long-term average rates r_{aZ} .
- Session scheduling and rate allocation: For each link (a, Z) , one session $c = c_{aZ}$ is chosen randomly with probability $\frac{g_{aZ}^c}{\sum_c g_{aZ}^c}$. Each of its sinks $\beta \in \mathcal{T}_c$ is chosen independently with probability $\frac{\sum_b f_{abZ}^{c\beta}}{g_{aZ}^c}$. Let $\mathcal{T}_{aZ} \subset \mathcal{T}_c$ denote the set of chosen sinks. For each $\beta \in \mathcal{T}_{aZ}$, one destination node $b = b_{aZ}^{c\beta}$ in Z is chosen with probability $\frac{f_{abZ}^{c\beta}}{\sum_b f_{abZ}^{c\beta}}$. The corresponding

allocated rates are

$$\mu_{abZ}^{c\beta}(t) = \begin{cases} \frac{\sum_c g_{aZ}^c}{r_{aZ}} \mu_{aZ}(t) & \text{if } c = c_{aZ}, \beta \in \mathcal{T}_{aZ} \\ & \text{and } b = b_{aZ}^{c\beta} \\ 0 & \text{otherwise.} \end{cases} \quad (43)$$

Theorem 5: If input rates (λ_i^c) are such that $(\lambda_i^c + \epsilon') \in \Lambda$, both the randomized policy and the back pressure policy stabilize the system with average total virtual queue length bounded as

$$\sum_{i,c,\beta} \overline{U_i^{c\beta}} = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^{t-1} \sum_{i,c,\beta} E\{U_i^{c\beta}(t')\} \leq \frac{BN}{\epsilon'} \quad (44)$$

where N is the number of network nodes and

$$B = \frac{\tau_{max}}{2} ((A_{max} + \mu_{max}^{in})^2 + (\mu_{max}^{out})^2).$$

The proof of this theorem uses the following result which is adapted from [24, Theorem 3]:

Theorem 6: Let $\underline{U}(t) = (U_1(t), \dots, U_n(t))$ be a vector of queue lengths, for some positive integer n . Define the Lyapunov function $L(\underline{U}(t)) = \sum_{j=1}^n [U_j(t)]^2$. If for all t

$$E\{L(\underline{U}(t+1)) - L(\underline{U}(t)) | \underline{U}(t)\} \leq C_1 - C_2 \sum_{j=1}^n U_j(t) \quad (45)$$

for some positive constants C_1, C_2 , and if $E\{L(\underline{U}(0))\} < \infty$, then

$$\sum_{j=1}^n \overline{U_j} = \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^{t-1} \sum_{j=1}^n E\{U_j(t')\} \leq \frac{C_1}{C_2} \quad (46)$$

and each queue is stable.

Proof: Summing over $t = 0, 1, \dots, T-1$ the expectation of (45) over the distribution of $\underline{U}(t)$, we have

$$E\{Ln(\underline{U}(T)) - L(\underline{U}(0))\} \leq TC_1 - C_2 \sum_{t=0}^{T-1} \sum_{j=1}^n E\{U_j(t)\}.$$

Since $L(\underline{U}(T)) > 0$,

$$\frac{1}{T} \sum_{t=0}^{T-1} \sum_{j=1}^n E\{U_j(t)\} \leq \frac{C_1}{C_2} + \frac{1}{TC_2} E\{L(\underline{U}(0))\}.$$

Taking the lim sup as $T \rightarrow \infty$ gives (46). Each queue is stable since

$$\begin{aligned} \gamma_j(M) &= \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t \Pr\{U_j(t') > M\} \\ &\leq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{t'=0}^t E\{U_j(t')\} / M \\ &\leq \frac{C_1}{C_2 M} \\ &\rightarrow 0 \text{ as } M \rightarrow \infty. \end{aligned}$$

■

Proof of Theorem 5: The queue lengths evolve according to:

$$\begin{aligned} U_i^{c\beta}(t+1) &\leq \max \left\{ U_i^{c\beta}(t) - \sum_{b,Z} \mu_{ibZ}^{c\beta}(t), 0 \right\} \\ &\quad + \sum_{a,Z} \mu_{aiZ}^{c\beta}(t) + A_i^c(t) \end{aligned} \quad (47)$$

which reflects the policy that the total number of virtual transmissions out of $Q_i^{c\beta}$ is limited by the queue length $U_i^{c\beta}(t)$.

Define the Lyapunov function $L(\underline{U}) = \sum_{i,c,\beta} (U_i^{c\beta})^2$. Squaring (47) and dropping some negative terms from the right hand side, we obtain

$$\begin{aligned} [U_i^{c\beta}(t+1)]^2 &\leq [U_i^{c\beta}(t)]^2 + \left(A_i^c + \sum_{a,Z} \mu_{aiZ}^{c\beta} \right)^2 + \\ &\quad \left(\sum_{b,Z} \mu_{ibZ}^{c\beta} \right)^2 - 2U_i^{c\beta}(t) \left[\sum_{b,Z} \mu_{ibZ}^{c\beta} - \sum_{a,Z} \mu_{aiZ}^{c\beta} - A_i^c \right] \end{aligned} \quad (48)$$

where the time dependencies of $\mu_{abZ}^{c\beta}$ and A_i^c are not shown for brevity, since these remain constant over the considered time slot.

Taking expectations of the sum of (48) over all i, c, β , noting that

$$\sum_{i,c,\beta} \left(\sum_{b,Z} \mu_{ibZ}^{c\beta} \right)^2$$

$$\begin{aligned}
&\leq \sum_{i,c} \tau_{max} \left(\max_{\beta \in \mathcal{T}_c} \sum_{b,Z} \mu_{ibZ}^{c\beta} \right)^2 \\
&\leq \sum_i \tau_{max} \left(\sum_c \left[\max_{\beta \in \mathcal{T}_c} \sum_{b,Z} \mu_{ibZ}^{c\beta} \right] \right)^2 \quad (49) \\
&\leq N \tau_{max} (\mu_{max}^{out})^2, \\
&\sum_{i,c,\beta} \left(A_i^c + \sum_{a,Z} \mu_{aiZ}^{c\beta} \right)^2 \\
&\leq \sum_{i,c} \tau_{max} \left(A_i^c + \max_{\beta \in \mathcal{T}_c} \sum_{a,Z} \mu_{aiZ}^{c\beta} \right)^2 \\
&\leq \sum_i \tau_{max} \left(\sum_c \left[A_i^c + \max_{\beta \in \mathcal{T}_c} \sum_{a,Z} \mu_{aiZ}^{c\beta} \right] \right)^2 \quad (50) \\
&\leq N \tau_{max} (A_{max} + \mu_{max}^{in})^2,
\end{aligned}$$

where the Cauchy-Schwarz inequality is used in steps (49) and (50), and using (4), (5), we obtain the drift expression

$$\begin{aligned}
E\{L(\underline{U}(t+1)) - L(\underline{U}(t)) | \underline{U}(t)\} &\leq 2BN - \\
2 \sum_{i,c,\beta} U_i^{c\beta}(t) &\left[E \left\{ \sum_{b,Z} \mu_{ibZ}^{c\beta} - \sum_{a,Z} \mu_{aiZ}^{c\beta} \middle| \underline{U}(t) \right\} - \lambda_i^c \right] \quad (51)
\end{aligned}$$

For randomized policy, substituting (39) and (42) into (51) gives

$$E\{L(\underline{U}(t+1)) - L(\underline{U}(t)) | \underline{U}(t)\} \leq 2BN - 2\epsilon' \sum_{i,c,\beta} U_i^{c\beta}(t). \quad (52)$$

Applying Theorem 6 gives

$$\sum_{i,c,\beta} \overline{U_i^{c\beta}} \leq \frac{BN}{\epsilon'}. \quad (53)$$

Thus the randomized policy satisfies the queue occupancy bound (44).

For the back pressure policy, $E\{\mu_{abZ}^{c\beta}(t) | \underline{U}(t)\}$ is dependent on $\underline{U}(t)$. The drift expression (51) can be

expressed as

$$\begin{aligned}
&E\{L(\underline{U}(t+1)) - L(\underline{U}(t)) | \underline{U}(t)\} \\
&\leq 2BN - 2 \left[D - \sum_{i,c,\beta} U_i^{c\beta}(t) \lambda_i^c \right]
\end{aligned}$$

where

$$D = \sum_{i,c,\beta} U_i^{c\beta}(t) \left[E \left\{ \sum_{b,Z} \mu_{ibZ}^{c\beta} - \sum_{a,Z} \mu_{aiZ}^{c\beta} \middle| \underline{U}(t) \right\} \right], \quad (54)$$

which is the portion of the drift expression that depends on the policy, can be rewritten as

$$D = \sum_{a,b,Z} \sum_{c,\beta} E \left\{ \mu_{abZ}^{c\beta} | \underline{U}(t) \right\} (U_a^{c\beta}(t) - U_b^{c\beta}(t)). \quad (55)$$

We compare the values of (55) for the two policies, giving

$$\begin{aligned}
D_{rand} &= \sum_{a,b,Z} \sum_{c,\beta} f_{abZ}^{c\beta} (U_a^{c\beta} - U_b^{c\beta}) \\
&\leq \sum_{a,Z} \sum_c g_{aZ}^c \sum_{\beta} \max_{b \in Z} (U_a^{c\beta} - U_b^{c\beta}) \\
&\leq \sum_{a,Z} \sum_c g_{aZ}^c w_{aZ}^* \\
&\leq \sum_{a,Z} r_{aZ} w_{aZ}^* \\
&= \sum_{a,Z} \left(\sum_{\underline{S}} \pi_{\underline{S}} r_{aZ}^{\underline{S}} \right) w_{aZ}^* \\
&\leq \sum_{\underline{S}} \pi_{\underline{S}} \max_{\underline{P} \in \Pi} \sum_{a,Z} \mu_{aZ}(\underline{P}, \underline{S}) w_{aZ}^* \\
&= D_{backpressure}.
\end{aligned}$$

where the last step follows from (35)-(36). Since the Lyapunov drift for the back-pressure policy is more negative than the drift for the randomized policy, the bound (53) also applies for the back-pressure policy. This completes the proof. ■

This also proves Theorem 4(b).

Proof of Theorems 1 and 2: The results follow from specializing Theorems 4 and 5 to the case where each

link (a, Z) has a destination set Z of size 1 and a capacity that does not depend on \underline{P} or \underline{S} . ■

Theorem 3 on decoding applies to both the wired and wireless cases.

VI. CORRELATED SOURCES

A. Model, approach and notation

We consider the case of source nodes with correlated information. For simplicity, we assume that the exogenous source processes consist of bits, or have been converted to bits by appropriate source coding.

Let x_α^c be the rate of exogenous bits which are to be transmitted from each source $\alpha \in \mathcal{S}_c$. The exogenous bits at $\alpha \in \mathcal{S}_c$ are partitioned into groups of x_α^c bits which we refer to as symbols in this paragraph. As in the classical distributed source coding problem formulation of [26], we consider coding over a block of n symbols present at each source⁴, and assume that for each session c , successive tuples of source symbols are drawn i.i.d. from some joint distribution Q_c .

Following [11], we consider intra-session random vector linear coding in \mathbb{F}_2 across the exogenous bits in each block. Each source packet contains n/T coded source bits from each of a set of z blocks, where T is a constant whose value will be determined later. Each of the coded source bits is formed as an independent random linear combination, in \mathbb{F}_2 , of the nx_α^c exogenous bits in the corresponding block. The code description overhead is amortized over the set of z blocks by using the same linear combinations to form the coded source bits for

⁴If the exogenous bits arrive sequentially at the sources, an initial startup delay is incurred in waiting to accumulate n symbols before coding begins. This differs from the model for independent sources in the previous sections, where information is encoded incrementally without needing a full block of information to be present.

each block in each packet. A set of source packets formed in this way from a set of z blocks constitutes a batch across which network coding occurs; as discussed in Section II-B there is some capacity loss from having multiple batches, the analysis of which is beyond the scope of this paper. The rate at which coded source packets are formed at each source is determined by the control policy.

Each packet formed subsequently at a node by coding together a set of input packets contains n/T coded bits for each block, each an independent random linear combination of the bits from that block in the input packets. The same linear combination is used for each of the blocks in the packet. An illustration is given in Figure 3.

Each sink decodes by mapping its block of received bits to a block of decoded bits that has minimum entropy or maximum Q -probability among all possible source values consistent with the received bit values.

The virtual queues in the correlated sources case are defined similarly as in Section III-A, except that they are additionally indexed by sources: each node i maintains, for each source α and sink β of each session c a virtual queue $Q_i^{c\alpha\beta}$ whose length $U_i^{c\alpha\beta}(t)$ is the number of packets from source α queued at node i that are intended for sink β . At each source $\alpha \in \mathcal{S}_c$, each coded source packet formed is added to some subset of the virtual queues $Q_\alpha^{c\alpha\beta}$, $\beta \in \mathcal{T}_c$ according to the control policy. A virtual transmission transfers a packet from a queue $Q_\alpha^{c\alpha\beta}$ to a queue $Q_b^{c\alpha\beta}$ over a link (a, Z) where $b \in Z$. At a sink node $\beta \in \mathcal{T}_c$, packets are removed from queue $Q_\beta^{c\alpha\beta}$ at a rate controlled by the policy.

For simplicity, in the rest of the paper all rates are given in packets per unit time, for a fixed packet length

of zn/T bits. We use $\lambda^{c\alpha\beta}$ to denote the average rate at which packets are added to $Q_\alpha^{c\alpha\beta}$, and $f_{abZ}^{c\alpha\beta}$ to denote average virtual flow rate, over link $(a, Z) \in \mathcal{L}$, from $Q_a^{c\alpha\beta}$ to $Q_b^{c\alpha\beta}$, $b \in Z$. We use g_{aZ}^c to denote the average physical flow rate for session c over $(a, Z) \in \mathcal{L}$.

For a set \mathcal{S} of sources, we denote by $H(\mathcal{S})$ their joint entropy in bits per unit time divided by the packet length zn/T . Conditional entropies are similarly denoted.

Let $\lambda_\alpha^c = H(\alpha) + \epsilon$, $\epsilon > 0$, for each source $\alpha \in \mathcal{S}_c$. We assume that λ_α^c is less than or equal to the maximum outflow rate μ_{max}^{out} of a node.

For brevity of notation, we use the convention that any variable term with superscript $c\alpha\beta$ equals zero unless $c \in \mathcal{C}$, $\alpha \in \mathcal{S}_c$, $\beta \in \mathcal{T}_c$, and any variable term with subscript aZ equals zero unless $(a, Z) \in \mathcal{L}$.

B. Capacity region with intra-session network coding

A transmission problem with correlated sources is considered *achievable* with intra-session network coding if there exists a sequence of codes such that the probability of decoding any session c source symbol in error at any sink in \mathcal{T}_c tends to zero.

Let Λ be the set of all entropy vectors $(H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) | \mathcal{S}' \subset \mathcal{S}_c, c \in \mathcal{C})$ such that there exist variables $(r_{aZ}) \in \Gamma$ and $\{f_{abZ}^{c\alpha\beta}, g_{aZ}^c, \lambda^{c\alpha\beta}\}$ satisfying:

$$f_{\beta b Z}^{c\alpha\beta} = 0 \quad \forall b, Z, c, \alpha, \beta \quad (56)$$

$$f_{abZ}^{c\alpha\beta} \geq 0 \quad \forall a, b, Z, c, \alpha, \beta \quad (57)$$

$$\lambda^{c\alpha\beta} \leq \sum_{b, Z} f_{\alpha b Z}^{c\alpha\beta} - \sum_{a, Z} f_{a \alpha Z}^{c\alpha\beta} \quad \forall c, \alpha, \beta \quad (58)$$

$$0 = \sum_{b, Z} f_{i b Z}^{c\alpha\beta} - \sum_{a, Z} f_{a i Z}^{c\alpha\beta} \quad \forall c, \alpha, \beta, i \notin \{\alpha, \beta\} \quad (59)$$

$$\sum_{\alpha, \beta} f_{abZ}^{c\alpha\beta} \leq g_{aZ}^c \quad \forall a, Z, c, \beta \quad (60)$$

$$\sum_c g_{aZ}^c \leq r_{aZ} \quad \forall a, Z \quad (61)$$

$$\sum_{\alpha \in \mathcal{S}'} \lambda^{c\alpha\beta} > H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) \quad \forall c, \mathcal{S}' \subseteq \mathcal{S}_c, \beta \in \mathcal{T}_c \quad (62)$$

where Γ is defined as in Section V-C.

Analogously to the independent sources case, variables $\{f_{abZ}^{c\alpha\beta}\}$ define a session c flow of size at least $\lambda^{c\alpha\beta}$ from source α to sink β . For each pair (c, β) , (62) corresponds to the Slepian-Wolf region [26].

Theorem 7: A necessary condition for achievability of a transmission problem is that the source entropy vector satisfies $(H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) | \mathcal{S}' \subset \mathcal{S}_c, c \in \mathcal{C}) \in \Lambda$.

Proof: For an intra-session network coding and power control policy, let $Y_{in}^{c\alpha\beta}(t)$ be the total number of packets added to queue $Q_\alpha^{c\alpha\beta}$ up to time t , and $Y_{out}^{c\alpha\beta}(t)$ the total number of packets removed from $Q_\beta^{c\alpha\beta}$ in that time. Let $F_{abZ}^{c\alpha\beta}(t)$ be the total virtual flow (in packets) from queue $Q_a^{c\alpha\beta}$ to $Q_b^{c\alpha\beta}$ over link (a, Z) up to time t , and $G_{aZ}^c(t)$ the total number of session c packets physically transmitted over link (a, Z) up to time t . By definition these quantities satisfy

$$F_{abZ}^{c\alpha\beta}(t) \geq 0 \quad \forall a, b, c, \alpha, \beta, Z \quad (63)$$

$$Y_{in}^{c\alpha\beta}(t) = \sum_{b, Z} F_{\alpha b Z}^{c\alpha\beta}(t) - \sum_{a, Z} F_{a \alpha Z}^{c\alpha\beta}(t) + U_\alpha^{c\alpha\beta}(t) \quad \forall c, \alpha, \beta \quad (64)$$

$$0 = \sum_{b, Z} F_{i b Z}^{c\alpha\beta}(t) - \sum_{a, Z} F_{a i Z}^{c\alpha\beta}(t) + U_i^{c\alpha\beta}(t) \quad \forall c, \alpha, \beta, i \notin \{\alpha, \beta\} \quad (65)$$

$$\sum_{\alpha, \beta} F_{abZ}^{c\alpha\beta}(t) \leq G_{aZ}^c(t) \quad \forall a, c, \beta, Z \quad (66)$$

$$\sum_c G_{aZ}^c(t) \leq \sum_{t'=0}^t \mu_{aZ}(\underline{P}(t'), \underline{S}(t')) \quad (67)$$

$$Y_{in}^{c\alpha\beta}(t) - Y_{out}^{c\alpha\beta}(t) = \sum_i U_i^{c\alpha\beta}(t) \quad \forall c, \alpha, \beta \quad (68)$$

Suppose the transmission problem is achievable with some intra-session network coding and power control policy. Then there exists some achievable policy satisfy-

ing

$$F_{\beta bZ}^{c\alpha\beta} = 0 \quad \forall b, Z, c, \alpha, \beta \quad (69)$$

i.e. virtual flow that has reached its intended sink is not retransmitted away from the sink. Also, for each session c and sufficiently large t , the flow rates from the sources in \mathcal{S}_c to each sink $\beta \in \mathcal{T}_c$ must be in the Slepian-Wolf region:

$$\sum_{\alpha \in \mathcal{S}'} Y_{out}^{c\alpha\beta}(t) > H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) t \quad \forall c, \mathcal{S}' \subseteq \mathcal{S}_c, \beta \in \mathcal{T}_c \quad (70)$$

Let $T_{\underline{S}}(t) \subset \{0, 1, \dots, t-1\}$ denote the subset of time slots, up to time $t-1$, during which the channel is in state \underline{S} . By similar reasoning as in the independent sources case (Equations (25)-(26)), there exists some finite M and some time \tilde{t} for which

$$0 \leq \frac{U_i^{c\alpha\beta}(\tilde{t})}{\tilde{t}} \leq \frac{M}{\tilde{t}} \leq \epsilon' \quad \forall i, c, \alpha, \beta \quad (71)$$

$$\frac{|T_{\underline{S}}(\tilde{t})|}{\tilde{t}} \leq \pi_{\underline{S}} + \epsilon' \quad \forall \underline{S} \quad (72)$$

are simultaneously satisfied with nonzero probability.

Let

$$f_{abZ}^{c\alpha\beta} = F_{abZ}^{c\alpha\beta}(\tilde{t})/\tilde{t} \quad (73)$$

$$g_{aZ}^c = G_{aZ}^c(\tilde{t})/\tilde{t} \quad (74)$$

$$\lambda^{c\alpha\beta} = Y_{in}^{c\alpha\beta}(\tilde{t})/\tilde{t} \quad (75)$$

As in the the independent sources case (Equations (27)-(28)), from (67), (72) and (74) we obtain

$$\left(\sum_c g_{aZ}^c(t) \right) \leq \sum_{\underline{S}} \pi_{\underline{S}} \left(\mu_{aZ}^{\underline{S}} \right) + \epsilon' (\mu_{aZ}^{max}) \text{card}\{\underline{S}\} \quad (76)$$

where vectors $\left(\mu_{aZ}^{\underline{S}} \right)$ are elements of the convex hull of rate vectors $\{(\mu(\underline{P}, \underline{S})) | \underline{P} \in \Pi\}$, $\text{card}\{\underline{S}\}$ is the number of channel states and (μ_{aZ}^{max}) is the maximum transmission rate of link (a, Z) .

Considering (63)-(66) and (68)-(70) at \tilde{t} , dividing by \tilde{t} , and using (71) and (73)-(75), we obtain

$$f_{abZ}^{c\alpha\beta} \geq 0 \quad \forall a, b, c, \beta \in \mathcal{T}_c, Z \quad (77)$$

$$\lambda^{c\alpha\beta} \leq \sum_{b,Z} f_{abZ}^{c\alpha\beta} - \sum_{a,Z} f_{a\alpha Z}^{c\alpha\beta} + \epsilon' \quad \forall c, \alpha, \beta \quad (78)$$

$$\sum_{b,Z} f_{ibZ}^{c\alpha\beta} - \sum_{a,Z} f_{aiZ}^{c\alpha\beta} \leq 0 \leq \sum_{b,Z} f_{ibZ}^{c\alpha\beta} - \sum_{a,Z} f_{aiZ}^{c\alpha\beta} + \epsilon' \quad \forall c, \alpha, \beta, i \notin \{\alpha, \beta\} \quad (79)$$

$$\sum_b f_{abZ}^{c\alpha\beta} \leq g_{aZ}^c \quad \forall a, b, c, \beta \in \mathcal{T}_c, Z \quad (80)$$

$$0 \leq \lambda^{c\alpha\beta} - \frac{Y_{out}^{c\alpha\beta}(\tilde{t})}{\tilde{t}} \leq N\epsilon' \quad (81)$$

$$f_{\beta bZ}^{c\alpha\beta} = 0 \quad \forall b, Z, c, \alpha, \beta \quad (82)$$

$$\sum_{\alpha \in \mathcal{S}'} \frac{Y_{out}^{c\alpha\beta}(\tilde{t})}{\tilde{t}} > H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) \quad \forall c, \mathcal{S}' \subseteq \mathcal{S}_c, \beta \in \mathcal{T}_c \quad (83)$$

Considering (76)-(83) in the limit as $\epsilon' \rightarrow 0$, we see that $(H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) | \mathcal{S}' \subset \mathcal{S}_c, c \in \mathcal{C})$ is a limit point of Λ , which is compact and thus contains its limit points. ■

C. Policies and achievability

We next give sufficient conditions and a back-pressure policy for achievability. Similarly to the case of independent sources, the conditions involve, informally speaking, the network capacity and source data rates being slightly higher than the joint source entropy rates.

The back-pressure policy for correlated sources differs from that for independent sources primarily in the operation at the sinks and the sources. The rates at which packets are injected into the network by the different sources of a session may have to be traded off against each other as the total information rate from all the sources may be larger than the joint entropy rate.

We propose a mechanism in which the different sinks monitor the amount of information received from each of

the sources and provide feedback implicitly through back-pressure to throttle the source rates. This is accomplished by maintaining virtual queues on a per source basis at each of the sinks and emptying these queues at appropriate rates. The information in these virtual queues creates the necessary gradient in queue sizes that then propagates back to the sources. The sources compress the information stream and transmit packets into the network at rates limited by the gradients and thus each source in the set of correlated sources transmits at the appropriate rate.

We can draw the following intuitive analogy between this policy and the back-pressure policy for independent sources described earlier. In the independent sources case, the system is driven by packet arrivals at the sources, which create “positive pressure” at the sources. In the correlated sources case, the system is driven by removal of packets at the sinks, which creates “negative pressure” or “vacuum” at the sinks. The policies differ in the details of the operation of sources and sinks, but the operation of interior network nodes according to the pressure gradient is the same in both cases. Thus, we call the policy for correlated sources a reverse back-pressure policy.

Unlike the back-pressure policy for independent sources, the policy for correlated sources ensures that the length of each virtual queue never exceeds a maximum length M whose value we will derive below. The virtual queues are, conceptually, initialized as containing M dummy packets each. Dummy packets are treated like other virtual packets by the control policy, and can be viewed as all zero packets for the purposes of network coding. They serve only an accounting purpose, allowing queue length gradients to be set up by the sinks when

initially there are no real data packets in the network. Initially dummy packets are removed at the sinks, but the number of dummy packets as a proportion of the total number of packets removed decreases with time since the total number of dummy packets is finite.

For ease of explanation, we describe the policies below in terms of the difference

$$V_i^{c\alpha\beta}(t) = M - U_i^{c\alpha\beta}(t) \quad (84)$$

between the maximum queue length and the virtual queue lengths. A negative gradient in $U_i^{c\alpha\beta}(t)$ is equivalent to a positive gradient in $V_i^{c\alpha\beta}(t)$. Thus, the reverse back-pressure policy for correlated sources can be viewed as similar to the back-pressure policy for independent sources with the roles of sources and sinks exchanged, and $V_i^{c\alpha\beta}(t)$ in place of $U_i^{c\alpha\beta}(t)$.

Each sink attempts to remove packets from the various virtual queues at rates within the Slepian-Wolf region. However, at times there may be insufficient packets at a sink node owing to channel variations. Each sink $\beta \in \mathcal{T}_c$ maintains, for each $\alpha \in \mathcal{S}_c$, a variable $W^{\alpha\beta}(t) \geq 0$ which keeps track of the net number of packets “owed” when there are insufficient packets to remove from $Q_\beta^{c\alpha\beta}$. This is analogous to the use of overflow buffers for positive flow in [4].

Specifically, suppose there exists a rate vector $(r_{aZ}) \in \Gamma$ and flow variables $\{f_{abZ}^{c\alpha\beta}, g_{aZ}^c, \lambda^{c\alpha\beta}\}$ satisfying, for some $\epsilon > 0$:

$$f_{\beta bZ}^{c\alpha\beta} = 0 \quad \forall b, Z, c, \alpha, \beta \quad (85)$$

$$f_{abZ}^{c\alpha\beta} \geq 0 \quad \forall a, b, Z, c, \alpha, \beta \quad (86)$$

$$\lambda^{c\alpha\beta} \leq \sum_{b,Z} f_{abZ}^{c\alpha\beta} - \sum_{a,Z} f_{a\alpha Z}^{c\alpha\beta} \quad \forall c, \alpha, \beta \quad (87)$$

$$0 = \sum_{b,Z} f_{ibZ}^{c\alpha\beta} - \sum_{a,Z} f_{aiZ}^{c\alpha\beta} \quad \forall c, \alpha, \beta, i \notin \{\alpha, \beta\} \quad (88)$$

$$\sum_{\alpha,b} f_{abZ}^{c\alpha\beta} \leq g_{aZ}^c \quad \forall a, Z, c, \beta \quad (89)$$

$$\sum_c g_{aZ}^c \leq r_{aZ} \quad \forall a, Z \quad (90)$$

$$\sum_{\alpha \in \mathcal{S}'} \lambda^{\alpha\beta} \geq H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) + \epsilon \\ \forall c, \mathcal{S}' \subseteq \mathcal{S}_c, \beta \in \mathcal{T}_c \quad (91)$$

Then, by the structure of the Slepian-Wolf region (91), there exists $(r_{aZ}) \in \Gamma$ and $\{f_{abZ}^{c\alpha\beta}, g_{aZ}^c, \lambda^{\alpha\beta}\}$ satisfying

$$\lambda^{\alpha\beta} \leq \lambda_\alpha^c = H(\alpha) + \epsilon \quad \forall c, \alpha, \beta \quad (92)$$

in addition to (85)-(91).

At time $t = 0$, $V_i^{c\alpha\beta}(0) = W^{c\alpha\beta}(t) = 0 \quad \forall i, c, \alpha, \beta$.

We consider policies that carry out the following sequence of steps in each time slot t .

- **Rebalancing:** Let t^- and t^+ denote the time instant just before and just after rebalancing respectively. For each $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$, let

$$\hat{W}^{c\alpha\beta}(t) = \min \left\{ V - V_\beta^{c\alpha\beta}(t^-), W^{c\alpha\beta}(t^-) \right\}$$

where $V \leq M$ is a constant whose value will be derived later. V represents an upper bound on the sink's "vacuum" $V_\beta^{c\alpha\beta}(\tau), \tau \in \mathbb{R}^+$, enforced by the policy; this follows inductively from $V_\beta^{c\alpha\beta}(0) = 0$ and the steps of the policy, as we will show. $\hat{W}^{c\alpha\beta}(t)$ packets are removed from $Q_\beta^{c\alpha\beta}$ and $\hat{W}^{c\alpha\beta}(t)$ is subtracted from $W^{c\alpha\beta}(t)$, i.e.

$$V_\beta^{c\alpha\beta}(t^+) = V_\beta^{c\alpha\beta}(t^-) + \hat{W}^{c\alpha\beta}(t) \\ W^{c\alpha\beta}(t^+) = W^{c\alpha\beta}(t^-) - \hat{W}^{c\alpha\beta}(t)$$

Note that

$$0 \leq V - V_\beta^{c\alpha\beta}(t^-) \leq M - V_\beta^{c\alpha\beta}(t^-) = U_\beta^{c\alpha\beta}(t^-) \quad (93)$$

where the first step follows from the induction hypothesis and the last step follows from (84). Thus,

$$0 \leq \hat{W}^{c\alpha\beta}(t) \leq U_\beta^{c\alpha\beta}(t^-) \quad (94)$$

and there are sufficient packets in $Q_\beta^{c\alpha\beta}$ to remove.

Also note that $V_\beta^{c\alpha\beta}(t^+) \leq V$ and $W^{c\alpha\beta}(t^+) \geq 0$.

- **Outflow rate allocation:** Let ϵ' be any positive constant less than ϵ . For each sink $\beta \in \mathcal{T}_c, c \in \mathcal{C}$, the policy chooses outflow variables $\{A_{out}^{c\alpha\beta}(t) | \alpha \in \mathcal{S}_c\}$ subject to

$$\sum_{\alpha \in \mathcal{S}'} A_{out}^{c\alpha\beta}(t) \geq H(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) + \epsilon - \epsilon' \quad \forall c, \mathcal{S}' \subseteq \mathcal{S}_c. \quad (95)$$

It then removes

$$\hat{A}_{out}^{c\alpha\beta}(t) = \min \left\{ V - V_\beta^{c\alpha\beta}(t^+), A_{out}^{c\alpha\beta}(t) \right\} \quad (96)$$

packets from queue $Q_\beta^{c\alpha\beta}$ and $A_{out}^{c\alpha\beta}(t) - \hat{A}_{out}^{c\alpha\beta}(t)$ is added to $W^{c\alpha\beta}(t^+)$. Note that by similar reasoning as in the rebalancing step,

$$0 \leq V - V_\beta^{c\alpha\beta}(t^+) \leq U_\beta^{c\alpha\beta}(t^+).$$

- **Inflow rate control:** For each source $\alpha \in \mathcal{S}_c, c \in \mathcal{C}$, the policy chooses inflow variables $A_{in}^{c\alpha\beta}(t)$ subject to

$$A_{in}^{c\alpha\beta}(t) \leq \lambda_\alpha^c. \quad (97)$$

Let

$$\hat{A}_{in}^{c\alpha\beta}(t) = \min \left\{ A_{in}^{c\alpha\beta}(t), V_\alpha^{c\alpha\beta}(t) \right\}.$$

$\max_\beta \hat{A}_{in}^{c\alpha\beta}(t)$ coded source packets are formed at α and $\hat{A}_{in}^{c\alpha\beta}(t)$ of them are added to queue $Q_\alpha^{c\alpha\beta}$ for each sink $\beta \in \mathcal{T}_c$.

- **Power allocation:** The policy chooses a vector of transmit powers $\underline{P}(t) = (P_{aZ}(t))$ from the set Π of feasible power allocations. This, together with the channel state $\underline{S}(t)$, determines the link rates $\underline{\mu}(t) = (\mu_{aZ}(t))$, assumed constant over the time slot.
- **Session scheduling, rate allocation and network coding:** For each link (a, Z) , each (source, sink) pair (α, β) of each session c , and each destination

node $b \in Z$, the policy allocates a transmission rate $\mu_{abZ}^{c\alpha\beta}(t)$, subject to

$$\mu_{\beta bZ}^{c\alpha\beta}(t) = 0 \quad (98)$$

as well as the overall link rate constraint

$$\mu_{aZ}(t) \geq \sum_{c \in \mathcal{C}} \max_{\beta \in \mathcal{T}_c} \sum_{b, \alpha} \mu_{abZ}^{c\alpha\beta}(t).$$

$\mu_{abZ}^{c\alpha\beta}(t)$ gives the maximum rate of virtual transmissions from $Q_a^{c\alpha\beta}$ to $Q_b^{c\alpha\beta}$ over (a, Z) . Besides this limit on virtual transmissions for pairs of queues over each link, the total number of virtual transmissions into $Q_a^{c\alpha\beta}$ over all links with destination node a is also limited by the value

$$\begin{cases} V_a^{c\alpha\beta}(t^+) & a \neq \alpha \\ \max\{V_a^{c\alpha\beta}(t) - A_{in}^{c\alpha\beta}(t), 0\} & a = \alpha. \end{cases}$$

Each session c packet physically transmitted on link (a, Z) is a random linear combination, in \mathbb{F}_q , of packets corresponding to a set of virtual transmissions on (a, Z) , each associated with a different sink in \mathcal{T}_c . Thus, the rate allocated to session c on (a, Z) is the maximum, over sinks $\beta \in \mathcal{T}_c$, of each sink β 's total allocated rate $\sum_{b, \alpha} \mu_{abZ}^{c\alpha\beta}(t)$. Note that by (96) and (98),

$$V_{\beta}^{c\alpha\beta}(\tau) \leq V \quad \forall t^+ \leq \tau \leq (t+1)^-.$$

The following reverse back-pressure algorithm uses the state of the virtual queues described above to allow receiver nodes to draw appropriate rates from among various correlated sources.

Reverse back-pressure policy

- **Outflow rate allocation:** Each sink $\beta \in \mathcal{T}_c, c \in \mathcal{C}$ chooses outflow variables $\{A_{out}^{c\alpha\beta}(t) | \alpha \in \mathcal{S}_c\}$ to minimize

$$\sum_{\alpha} V_{\beta}^{c\alpha\beta}(t^+) A_{out}^{c\alpha\beta}(t)$$

subject to (95). This optimization problem has a simple greedy solution [8, Theorem 1]: let $\{\alpha_1, \dots, \alpha_{|\mathcal{S}_c|}\}$ be the set of sources \mathcal{S}_c in ascending order of the value of $V_{\beta}^{c\alpha\beta}(t^+)$, i.e.

$$V_{\beta}^{c\alpha_1\beta}(t^+) \leq \dots \leq V_{\beta}^{c\alpha_{|\mathcal{S}_c|}\beta}(t^+).$$

The solution of the optimization problem is

$$A_{out}^{c\alpha_1\beta}(t) = H(\alpha_1)$$

$$A_{out}^{c\alpha_2\beta}(t) = H(\alpha_2 | \alpha_1)$$

...

$$A_{out}^{c\alpha_{|\mathcal{S}_c|}\beta}(t) = H(\alpha_{|\mathcal{S}_c|} | \alpha_{|\mathcal{S}_c|-1}, \dots, \alpha_1) \quad (99)$$

- **Inflow rate control:** For each source $\alpha \in \mathcal{S}_c, c \in \mathcal{C}$, $A_{in}^{c\alpha\beta}(t) = \lambda_{\alpha}^c$.
- **Session scheduling:** For each link (a, Z) , one session

$$c_{aZ}^* = \arg \max_c \left\{ \sum_{\beta \in \mathcal{T}_c} \max \left(\max_{\alpha \in \mathcal{S}_c} \left(\max_{b \in Z} \left(V_b^{c\alpha\beta} - V_a^{c\alpha\beta} \right) \right), 0 \right) \right\}$$

is chosen. For each $\beta \in \mathcal{T}_{c_{aZ}^*}$, let

$$\alpha_{aZ}^{\beta*} = \arg \max_{\alpha \in \mathcal{S}_{c_{aZ}^*}} \left(\max_{b \in Z} \left(V_b^{c_{aZ}^* \alpha \beta} - V_a^{c_{aZ}^* \alpha \beta} \right) \right)$$

$$b_{aZ}^{\beta*} = \arg \max_{b \in Z} \left(V_b^{c_{aZ}^* \alpha_{aZ}^{\beta*} \beta} - V_a^{c_{aZ}^* \alpha_{aZ}^{\beta*} \beta} \right)$$

$$w_{aZ}^* = \max_c \left\{ \sum_{\beta \in \mathcal{T}_c} \max \left(\max_{\alpha \in \mathcal{S}_c} \left(\max_{b \in Z} \left(V_b^{c\alpha\beta} - V_a^{c\alpha\beta} \right) \right), 0 \right) \right\}.$$

- **Power control:** The state $\underline{S}(t)$ is observed, and a power allocation

$$\underline{P}(t) = \arg \max_{\underline{P} \in \Pi} \sum_{a, Z} \mu_{aZ}(\underline{P}, \underline{S}(t)) w_{aZ}^* \quad (100)$$

is chosen.

- Link rate allocation: For each link (a, Z) ,

$$\mu_{abZ}^{c\alpha\beta}(t) = \begin{cases} \mu_{aZ}(t) & \text{if } c = c_{aZ}^*, \beta \in \mathcal{T}_c, \\ & \alpha = \alpha_{aZ}^{\beta*}, b = b_{aZ}^{\beta*} \text{ and} \\ & U_a^{c\alpha\beta} - U_b^{c\alpha\beta} > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (101)$$

In analyzing the reverse back-pressure algorithm, we compare it with the following randomized policy, which bases its control decisions on the assumed solution to (85)-(92).

Randomized policy

- Outflow rate allocation: For each $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$, $A_{out}^{c\alpha\beta}(t) = \lambda^{c\alpha\beta} - \epsilon'$.
- Inflow rate control: For each source $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$, $A_{in}^{c\alpha\beta}(t) = \lambda^{c\alpha\beta}$.
- Power allocation: The channel state \underline{S} is observed, and power is allocated according to the algorithm of Lemma 3, giving instantaneous link rates $\mu_{aZ}(t)$ and long-term average rates r_{aZ} .
- Session scheduling and network coding: For each link (a, Z) , one session $c = c_{aZ}$ is chosen randomly with probability $\frac{g_{aZ}^c}{\sum_c g_{aZ}^c}$. Each of its sinks β is independently chosen with probability $\frac{\sum_{\alpha, b} f_{abZ}^{c\alpha\beta}}{g_{aZ}^c}$. Let \mathcal{T}_{aZ} denote the set of chosen sinks. For each chosen sink, one (source, destination node) pair $(\alpha \in \mathcal{S}_c, b \in Z) = (\alpha_{aZ}^{c\beta}, b_{aZ}^{c\beta})$ is chosen with probability $\frac{f_{abZ}^{c\alpha\beta}}{\sum_{\alpha, b} f_{abZ}^{c\alpha\beta}}$. The corresponding allocated rates are

$$\mu_{abZ}^{c\alpha\beta}(t) = \begin{cases} \frac{\sum_c g_{aZ}^c}{r_{aZ}} \mu_{aZ}(t) & \text{if } c = c_{aZ}, \beta \in \mathcal{T}_{aZ}, \\ & \alpha = \alpha_{aZ}^{c\beta} \text{ and } b = b_{aZ}^{c\beta} \\ 0 & \text{otherwise.} \end{cases} \quad (102)$$

Theorem 8: Suppose the source statistics satisfy $(\{H(S') + \epsilon | S' \subset \mathcal{S}_c, c \in \mathcal{C}\}, \{\lambda_\alpha^c | \alpha \in \mathcal{S}_c, c \in \mathcal{C}\}) \in$

- Λ. Consider the reverse back-pressure algorithm with $V = \frac{BN}{\epsilon'}$ and $M = V + N\mu_{max}^{out}$, where

$$B = \frac{\tau_{max}}{2} \left(\frac{1}{N} \sum_{i,c} E \left\{ \left(A_{in}^{c\alpha\beta} \right)^2 \right\} + \frac{2}{N} \sigma_{max} \mu_{max}^{out} \mu_{max}^{in} + (\mu_{max}^{out})^2 + (\mu_{max}^{in})^2 \right),$$

and σ_{max} and τ_{max} denote the maximum number of sources and sinks respectively of a multicast session.

- (a) The algorithm is stable with

$$U_i^{c\alpha\beta}(\tau) \leq M \quad \forall i, c, \alpha, \beta, \tau \quad (103)$$

$$\sum_{c, \alpha, \beta} E \{ W^{c\alpha\beta}(\tau) \} \leq \sum_c |\mathcal{S}_c| |\mathcal{T}_c| \left(\frac{NM^2}{2V} + \mu_{max}^{out} \right) \quad \forall \tau. \quad (104)$$

- (b) For sufficiently large T , the probability of error decreases exponentially in n .

Proof: The proof is presented in the appendix. ■

VII. SUMMARY AND FUTURE WORK

We presented dynamic algorithms with network coding for multicast in wired and time-varying wireless networks. We showed that random network coding can be applied in such a dynamic setting. In our algorithms, feedback to nodes upstream is achieved through back-pressure. In particular, for the correlated sources case, the source rate allocation is also achieved through back-pressure by control of per-source virtual queues at the sinks modulating the relative gradients between different sources. This is in contrast to the Internet where source control is achieved through explicit feedback such as in the transmission control protocol (TCP). Combining network coding with the methods that are currently in widespread use for flow control and scheduling would be an important area for future research.

In wireless networks, network coding results in nodes transmitting different information that may interfere with

each other resulting in lower transmission rates compared to the case without network coding where the same information is broadcast by different nodes. Thus there is an inherent trade-off between network coding and reduced interference in combination with larger combining gain in the wireless case. Our approach provides a way to combine both techniques by optimizing over the different transmit scenarios. Understanding the balance between network coding and interference reduction are interesting topics for investigation.

Another line of work is to investigate the performance of heuristics for transmitter scheduling that take into account the product of link rate and queue size difference, rather than link rate alone.

APPENDIX

Proof of Theorem 8a: Consider a time slot t . For both policies,

$$V_i^{c\alpha\beta}(t^-) = V_i^{c\alpha\beta}(t^+) = V_i^{c\alpha\beta}(t) \quad \forall t \in \mathbb{Z}, c, \alpha, \beta, i \neq \beta.$$

For each $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$, define the Lyapunov function

$$L^{c\alpha\beta}(\underline{V}^{c\alpha\beta}, W^{c\alpha\beta}) = \sum_i (V_i^{c\alpha\beta})^2 + 2VW^{c\alpha\beta}$$

and let

$$L^{c\alpha\beta}(t) = L^{c\alpha\beta}(\underline{V}^{c\alpha\beta}(t), W^{c\alpha\beta}(t)). \quad (105)$$

In the rebalancing step, the change in $L^{c\alpha\beta}(t)$ is

$$\begin{aligned} & L^{c\alpha\beta}(t^+) - L^{c\alpha\beta}(t^-) \\ &= \left(V_\beta^{c\alpha\beta}(t^-) + \hat{W}^{c\alpha\beta}(t) \right)^2 - \left(V_\beta^{c\alpha\beta}(t^-) \right)^2 \\ &\quad - 2V\hat{W}^{c\alpha\beta}(t) \\ &= \hat{W}^{c\alpha\beta}(t) \left(\hat{W}^{c\alpha\beta}(t) + 2V_\beta^{c\alpha\beta}(t^-) - 2V \right) \\ &\leq 0. \end{aligned} \quad (106)$$

where the last step follows from (93) and (94).

In the rest of the time slot $(t, t+1)$, $V_i^{c\alpha\beta}(t)$ and $W^{c\alpha\beta}(t)$ evolve according to:

$$\begin{aligned} V_\alpha^{c\alpha\beta}(t+1) &\leq \max \left\{ V_\alpha^{c\alpha\beta}(t) - A_{in}^{c\alpha\beta}(t) \right. \\ &\quad \left. - \sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta}(t), 0 \right\} + \sum_{b,Z} \mu_{abZ}^{c\alpha\beta}(t) \end{aligned} \quad (107)$$

$$\begin{aligned} V_\beta^{c\alpha\beta}((t+1)^-) &\leq \max \left\{ V_\beta^{c\alpha\beta}(t^+) - \sum_{a,Z} \mu_{a\beta Z}^{c\alpha\beta}(t), \right. \\ &\quad \left. 0 \right\} + \hat{A}_{out}^{c\alpha\beta}(t) \end{aligned} \quad (108)$$

$$\begin{aligned} V_i^{c\alpha\beta}(t+1) &\leq \max \left\{ V_i^{c\alpha\beta}(t) - \sum_{a,Z} \mu_{aiZ}^{c\alpha\beta}(t), 0 \right\} \\ &\quad + \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta}(t) \quad \forall i \notin \{\alpha, \beta\} \end{aligned} \quad (109)$$

$$\begin{aligned} W^{c\alpha\beta}((t+1)^-) &= W^{c\alpha\beta}(t^+) + A_{out}^{c\alpha\beta}(t) \\ &\quad - \hat{A}_{out}^{c\alpha\beta}(t). \end{aligned} \quad (110)$$

Squaring (107)-(109) and dropping some negative terms from the right hand sides, we obtain

$$\begin{aligned} [V_\alpha^{c\alpha\beta}(t+1)]^2 &\leq [V_\alpha^{c\alpha\beta}(t)]^2 + \left(A_{in}^{c\alpha\beta} \right)^2 + \left(\sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right)^2 \\ &\quad + \left(\sum_{b,Z} \mu_{abZ}^{c\alpha\beta} \right)^2 + 2A_{in}^{c\alpha\beta} \left(\sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) \\ &\quad - 2V_\alpha^{c\alpha\beta}(t) \left[\sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} - \sum_{b,Z} \mu_{abZ}^{c\alpha\beta} + A_{in}^{c\alpha\beta} \right] \end{aligned} \quad (111)$$

$$\begin{aligned} [V_\beta^{c\alpha\beta}((t+1)^-)]^2 &\leq [V_\beta^{c\alpha\beta}(t^+)]^2 + \left(\hat{A}_{out}^{c\alpha\beta} \right)^2 \\ &\quad + \left(\sum_{a,Z} \mu_{a\beta Z}^{c\alpha\beta} \right)^2 + \left(\sum_{b,Z} \mu_{b\beta Z}^{c\alpha\beta} \right)^2 + 2\hat{A}_{out}^{c\alpha\beta} \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right) \\ &\quad - 2V_\beta^{c\alpha\beta}(t^+) \left[\sum_{a,Z} \mu_{a\beta Z}^{c\alpha\beta} - \sum_{b,Z} \mu_{b\beta Z}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right] \end{aligned} \quad (112)$$

$$[V_i^{c\alpha\beta}(t+1)]^2 \leq [V_i^{c\alpha\beta}(t)]^2 + \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right)^2$$

$$+ \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right)^2 - 2V_i^{c\alpha\beta}(t) \left[\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} - \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right] \quad \forall i \notin \{\alpha, \beta\} \quad (113)$$

where the time index t of $\mu_{abZ}^{c\alpha\beta}$, $\hat{A}_{out}^{c\alpha\beta}$ and $A_{in}^{c\alpha\beta}$ are not shown for simplicity of notation.

From (110)-(113), we have

$$\begin{aligned} L^{c\alpha\beta}((t+1)^-) - L^{c\alpha\beta}(t^+) &\leq \left(A_{in}^{c\alpha\beta} \right)^2 + \left(\hat{A}_{out}^{c\alpha\beta} \right)^2 \\ &+ \sum_i \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right)^2 + \sum_i \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right)^2 \\ &+ 2A_{in}^{c\alpha\beta} \left(\sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) + 2\hat{A}_{out}^{c\alpha\beta} \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right) \\ &- 2 \left[\sum_i V_i^{c\alpha\beta}(t) \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} - \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right) \right. \\ &+ V_\alpha^{c\alpha\beta}(t^+) \left(A_{in}^{c\alpha\beta} \right) - V_\beta^{c\alpha\beta}(t^+) \left(\hat{A}_{out}^{c\alpha\beta} \right) \left. \right] \\ &+ 2V \left(A_{out}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right) \\ &\leq \left(A_{in}^{c\alpha\beta} \right)^2 + \left(\hat{A}_{out}^{c\alpha\beta} \right)^2 + \sum_i \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right)^2 \\ &+ \sum_i \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right)^2 + 2A_{in}^{c\alpha\beta} \left(\sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) \\ &+ 2\hat{A}_{out}^{c\alpha\beta} \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right) - 2 \left[\sum_i V_i^{c\alpha\beta}(t) \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right. \right. \\ &- \left. \left. \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right) + V_\alpha^{c\alpha\beta}(t^+) \left(A_{in}^{c\alpha\beta} \right) \right. \\ &\left. \left. - V_\beta^{c\alpha\beta}(t^+) \left(\hat{A}_{out}^{c\alpha\beta} \right) \right] + \frac{\left(A_{out}^{c\alpha\beta} \right)^2}{2}, \end{aligned}$$

since, from (96), either $\hat{A}_{out}^{c\alpha\beta} = A_{out}^{c\alpha\beta}$, in which case

$$V_\beta^{c\alpha\beta}(t^+) \hat{A}_{out}^{c\alpha\beta} + V \left(A_{out}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right) = V_\beta^{c\alpha\beta}(t^+) A_{out}^{c\alpha\beta},$$

or $\hat{A}_{out}^{c\alpha\beta} = V - V_\beta^{c\alpha\beta}(t^+)$, in which case

$$\begin{aligned} V_\beta^{c\alpha\beta}(t^+) \hat{A}_{out}^{c\alpha\beta} + V \left(A_{out}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right) \\ = V_\beta^{c\alpha\beta}(t^+) A_{out}^{c\alpha\beta} + \left(A_{out}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right) \left(V - V_\beta^{c\alpha\beta}(t^+) \right) \end{aligned}$$

$$\begin{aligned} &= V_\beta^{c\alpha\beta}(t^+) A_{out}^{c\alpha\beta} + \left(A_{out}^{c\alpha\beta} - \hat{A}_{out}^{c\alpha\beta} \right) \hat{A}_{out}^{c\alpha\beta} \\ &\leq V_\beta^{c\alpha\beta}(t^+) A_{out}^{c\alpha\beta} + \left(\frac{A_{out}^{c\alpha\beta}}{2} \right)^2. \end{aligned}$$

Summing over all $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$ and taking expectations, we obtain

$$\begin{aligned} &\sum_{c,\alpha,\beta} E \{ L^{c\alpha\beta}((t+1)^-) - L^{c\alpha\beta}(t^+) | \underline{V}(t^+), \underline{W}(t^+) \} \\ &\leq 2BN - 2 \sum_{c,\alpha,\beta} \left[\sum_i V_i^{c\alpha\beta}(t) E \left\{ \sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right. \right. \\ &\quad \left. \left. - \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \right. \\ &\quad \left. + V_\alpha^{c\alpha\beta}(t) E \left\{ A_{in}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \right. \\ &\quad \left. - V_\beta^{c\alpha\beta}(t^+) E \left\{ A_{out}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \right] \quad (114) \end{aligned}$$

where

$$\begin{aligned} B &= \frac{\tau_{max}}{2} \left(\frac{5}{2N} \sum_{c,\alpha} (\lambda_\alpha^c)^2 + \frac{4\sigma_{max} \mu_{max}^{in} \mu_{max}^{out}}{N} \right. \\ &\quad \left. + (\mu_{max}^{out})^2 + (\mu_{max}^{in})^2 \right), \end{aligned}$$

by noting that

$$\begin{aligned} &\sum_{i,c,\alpha,\beta} \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right)^2 \\ &\leq \sum_{i,c} \tau_{max} \max_{\beta \in \mathcal{T}_c} \sum_{\alpha \in \mathcal{S}_c} \left(\sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} \right)^2 \\ &\leq \sum_i \tau_{max} \left(\sum_c \left[\max_{\beta \in \mathcal{T}_c} \sum_{\alpha \in \mathcal{S}_c, a, Z} \mu_{aiZ}^{c\alpha\beta} \right] \right)^2 \quad (115) \\ &\leq N \tau_{max} (\mu_{max}^{out})^2, \\ &\sum_{i,c,\alpha,\beta} \left(\sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \right)^2 \leq N \tau_{max} (\mu_{max}^{in})^2, \\ &\sum_{c,\alpha,\beta} \left(A_{in}^{c\alpha\beta} \sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) \\ &\leq \sum_c \sigma_{max} \tau_{max} \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} \left(A_{in}^{c\alpha\beta} \sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) \end{aligned}$$

$$\begin{aligned}
&\leq \sigma_{max} \tau_{max} \left(\sum_c \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} A_{in}^{c\alpha\beta} \right) \\
&\quad \left(\sum_c \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} \sum_{a,Z} \mu_{a\alpha Z}^{c\alpha\beta} \right) \quad (116) \\
&\leq \sigma_{max} \tau_{max} \mu_{max}^{in} \mu_{max}^{out} \\
&\sum_{c,\alpha,\beta} \left(A_{out}^{c\alpha\beta} \sum_{b,Z} \mu_{\beta b Z}^{c\alpha\beta} \right) \\
&\leq \sum_c \sigma_{max} \tau_{max} \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} \left(A_{out}^{c\alpha\beta} \sum_{b,Z} \mu_{\beta b Z}^{c\alpha\beta} \right) \\
&\leq \sigma_{max} \tau_{max} \left(\sum_c \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} A_{out}^{c\alpha\beta} \right) \\
&\quad \left(\sum_c \max_{\alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c} \sum_{b,Z} \mu_{\beta b Z}^{c\alpha\beta} \right) \quad (117) \\
&\leq \sigma_{max} \tau_{max} \mu_{max}^{in} \mu_{max}^{out},
\end{aligned}$$

where the Cauchy-Schwarz inequality is used in the numbered steps (115)-(117). The drift expression (114) can be expressed as

$$\begin{aligned}
&\sum_{c,\alpha,\beta} E\{L^{c\alpha\beta}((t+1)^-) - L^{c\alpha\beta}(t^+) | \underline{V}(t^+), \underline{W}(t^+)\} \\
&\leq 2BN - 2D(\underline{V}(t^+), \underline{W}(t^+)) \quad (118)
\end{aligned}$$

where

$$\begin{aligned}
&D(\underline{V}(t^+), \underline{W}(t^+)) \\
&= \sum_{i,c,\alpha,\beta} V_i^{c\alpha\beta}(t) E \left\{ \sum_{a,Z} \mu_{aiZ}^{c\alpha\beta} - \sum_{b,Z} \mu_{ibZ}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \\
&\quad + \sum_{c,\alpha,\beta} \left[V_\alpha^{c\alpha\beta}(t) E \left\{ A_{in}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \right. \\
&\quad \left. - V_\beta^{c\alpha\beta}(t) E \left\{ A_{out}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+) \right\} \right] \quad (119)
\end{aligned}$$

is the portion of (114) that depends on the policy.

For the randomized policy,

$$\begin{aligned}
E\{\mu_{abZ}^{c\alpha\beta}(t)\} &= f_{abZ}^{c\alpha\beta} \quad \forall a, b, Z, c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c \\
A_{in}^{c\alpha\beta} &= \lambda^{c\alpha\beta} \quad \forall c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c \\
A_{out}^{c\alpha\beta} &= \lambda^{c\alpha\beta} - \epsilon' \quad \forall c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c
\end{aligned}$$

independently of $\underline{V}(t), \underline{W}(t)$. Substituting this into (119) and using (85)-(88) gives

$$D_{randomized}(\underline{V}(t^+), \underline{W}(t^+)) \geq \epsilon' \sum_{c,\alpha,\beta} V_\beta^{c\alpha\beta}(t^+).$$

Next, we consider the back-pressure policy. Variables $E\{\mu_{abZ}^{c\alpha\beta}(t) | \underline{V}(t^+), \underline{W}(t^+)\}$, $E\left\{A_{in}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+)\right\}$ and $E\left\{A_{out}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+)\right\}$ are dependent on $\underline{V}(t^+)$. (119) can be rewritten as

$$\begin{aligned}
&D(\underline{V}(t^+), \underline{W}(t^+)) \\
&= \sum_{a,b,Z} \sum_{c,\beta} E\left\{\mu_{abZ}^{c\alpha\beta} | \underline{V}(t^+), \underline{W}(t^+)\right\} (V_a^{c\alpha\beta}(t) \\
&\quad - V_b^{c\alpha\beta}(t)) + \sum_{c,\alpha,\beta} V_\alpha^{c\alpha\beta}(t) E\left\{A_{in}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+)\right\} \\
&\quad - \sum_{c,\alpha,\beta} V_\beta^{c\alpha\beta}(t) E\left\{A_{out}^{c\alpha\beta} \middle| \underline{V}(t^+), \underline{W}(t^+)\right\}. \quad (120)
\end{aligned}$$

The three terms of the expression above involve disjoint sets of policy-dependent variables, so the three terms can be considered separately. The reverse back-pressure policy maximizes each of them subject to constraints which are also satisfied by the randomized policy: for the first term this is shown in detail in the proof of Theorem 5. For the second term, this follows from the inflow rate control step where the constraint is given by (97). For the third term, this follows from the outflow rate allocation step where the constraint is given by (95). Thus,

$$\begin{aligned}
&D_{backpressure}(\underline{V}(t^+), \underline{W}(t^+)) \\
&\geq D_{randomized}(\underline{V}(t^+), \underline{W}(t^+)) \\
&\geq \epsilon' \sum_{c,\alpha,\beta} V_\beta^{c\alpha\beta}(t^+). \quad (121)
\end{aligned}$$

Since flow is transmitted only from a longer to a shorter queue, and since $V_\beta^{c\alpha\beta}(\tau) \leq V \quad \forall \tau$,

$$V_i^{c\alpha\beta}(\tau) \leq V + N\mu_{max}^{out} \quad \forall i, c, \alpha, \beta, \tau.$$

Thus, the maximum virtual queue length M can be set to $V + N\mu_{max}^{out}$.

From (96) and (99), the increase in $W^{c\alpha\beta}(\tau)$ over the time slot (which occurs in the outflow rate allocation step) is at most $A_{out}^{c\alpha\beta}(t) \leq H(\alpha) + \epsilon \leq \mu_{max}^{out}$. If for all (c, α, β) , $W^{c\alpha\beta}(t^+) = 0$ then

$$\sum_{c,\alpha,\beta} L^{c\alpha\beta}((t+1)^-) \leq \sum_c |\mathcal{S}_c| |\mathcal{T}_c| (NM^2 + 2V\mu_{max}^{out}). \quad (122)$$

If on the other hand $W^{c\alpha\beta}(t^+) > 0$ for some (c, α, β) , then $V_\beta^{c\alpha\beta}(t^+) = V$,⁵ and

$$\sum_{c,\alpha,\beta} E \left\{ L^{c\alpha\beta}((t+1)^-) - L^{c\alpha\beta}(t^+) \middle| V_\beta^{c\alpha\beta}(t^+) = V \right\} \leq 2BN - 2\epsilon'V$$

from (118) and (121); setting $V = \frac{BN}{\epsilon'}$ and using (106) gives $E \{ L^{c\alpha\beta}((t+1)^-) - L^{c\alpha\beta}(t^-) \} \leq 0$. By induction on the number of time slots, using (122), we have

$$\sum_{c,\alpha,\beta} E \{ L^{c\alpha\beta}(t^-) \} \leq \sum_c |\mathcal{S}_c| |\mathcal{T}_c| (NM^2 + 2V\mu_{max}^{out}) \quad \forall t \in \mathbb{Z}.$$

This in turn gives, using (105),

$$\sum_{c,\alpha,\beta} E \{ W^{c\alpha\beta}(t^-) \} \leq \sum_c |\mathcal{S}_c| |\mathcal{T}_c| \left(\frac{NM^2}{2V} + \mu_{max}^{out} \right) \quad \forall t \in \mathbb{Z}. \quad (123)$$

Since $E \{ W^{c\alpha\beta}(\tau) \} \leq E \{ W^{c\alpha\beta}(t^-) \} \quad \forall \tau \in (t-1, t]$, (104) follows. ■

The proof of Theorem 8b uses the following theorem which is a straightforward generalization of [12, Theorem 6] from a static network model to the packet network model described in Section VI-A.

⁵In the rebalancing step, either $\hat{W}^{c\alpha\beta}(t) = V - V_\beta^{c\alpha\beta}(t^-)$, in which case $V_\beta^{c\alpha\beta}(t^+) = V$, or else $\hat{W}^{c\alpha\beta}(t) = W^{c\alpha\beta}(t^-)$, in which case $W^{c\alpha\beta}(t^+) = 0$.

Theorem 9: Consider a block of n exogenous symbols from each source and a sink node $\beta \in \mathcal{T}_c$. If the corresponding virtual flow from each subset $\mathcal{S}' \subset \mathcal{S}_c$ of source nodes to β is greater than $TH(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}'))$ packets (where each packet contains n/T bits from the block), then the decoding error probability at β decreases exponentially in n .

Proof: Analogously to the proof of Theorem 3, for a given sequence \mathcal{P} of packet transmissions corresponding to the block, we can define a corresponding static network \mathcal{G} which has the same node set, unit capacity links corresponding to transmissions in \mathcal{P} , and each subset $\mathcal{S}' \subset \mathcal{S}_c$ of session c sources has conditional entropy $TH(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}'))$ conditioned on the other sources. With this correspondence, the virtual flow from \mathcal{S}' to β in the packet model corresponds to the minimum cut between \mathcal{S}' and β in the static graph \mathcal{G} .

We can then apply [12, Theorem 6] to obtain that the probability that a sink decodes sources in \mathcal{S}' wrongly and sources in $\mathcal{S}_c \setminus \mathcal{S}'$ correctly decreases exponentially in n . [12, Theorem 6] is stated for the case of two sources, but for sessions with more than two sources, for each subset $\mathcal{S}' \subset \mathcal{S}_c$ of sources, the same proof applies with \mathcal{S}' taking the place of one source and $\mathcal{S}_c \setminus \mathcal{S}'$ the other. ■

Proof of Theorem 8b: Over the time period $0 \leq t \leq T$, there is a virtual flow of $\sum_{t=1}^T A_{out}^{c\alpha\beta}(t) - W^{c\alpha\beta}(T)$ packets from α to β , of which at most NM are dummy packets. By Theorem 8a, for sufficiently large T , with high probability there is, for each $(c, \alpha \in \mathcal{S}_c, \beta \in \mathcal{T}_c)$, a virtual flow of more than $\sum_{t=1}^T (A_{out}^{c\alpha\beta}(t) - \frac{\epsilon - \epsilon'}{|\mathcal{S}_c|})$ non-dummy packets from α to β . From (95), we have

$$\sum_{t=1}^T \left(\sum_{\alpha \in \mathcal{S}'} A_{out}^{c\alpha\beta}(t) - (\epsilon - \epsilon') \right) \geq TH(\mathcal{S}' | (\mathcal{S}_c \setminus \mathcal{S}')) \quad \forall c, \mathcal{S}' \subseteq \mathcal{S}_c.$$

Applying Theorem 9 gives the result. ■

ACKNOWLEDGMENT

The authors would like to thank the associate editor and anonymous reviewers for their very thorough reading and many valuable comments and suggestions, especially one of the reviewers who provided Figure 3 and improvements to the proof of Theorem 3.

REFERENCES

- [1] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, July 2000.
- [2] B. Awerbuch, A. Brinkmann, and C. Scheideler. Anycasting and multicasting in adversarial systems: routing and admission control, preliminary technical report, 2002.
- [3] B. Awerbuch and T. Leighton. A simple local control approximation algorithm for multicommodity flow. In *Proc. 34th IEEE Symposium on Foundations of Computer Science*, 1993.
- [4] B. Awerbuch and T. Leighton. Improved approximation algorithms for the multicommodity flow problem and local competitive routing in dynamic networks. In *Proc. 26th Ann. ACM Symp. Theory of Computing*, 1994.
- [5] D. Bertsekas, A. Nedic, and A. E. Ozdaglar. *Convex analysis and optimization*. Athena Scientific, 2003.
- [6] P. A. Chou, Y. Wu, and K. Jain. Practical network coding. In *Proc. 41st Annual Allerton Conference on Communication, Control, and Computing*, 2003.
- [7] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [8] R. Cristescu, B. Beferull-Lozano, and M. Vetterli. Networked slepian-wolf: theory, algorithms, and scaling laws. *IEEE Transactions on Information Theory*, 51(12):4057–4073, Dec. 2005.
- [9] T. Cui, L. Chen, and T. Ho. Distributed optimization in wireless networks using broadcast advantage. In *Proc. IEEE Conference on Decision and Control*, 2007, to appear.
- [10] T. Ho, R. Koetter, M. Médard, D. R. Karger, and M. Effros. The benefits of coding over routing in a randomized setting. In *Proc. IEEE International Symposium on Information Theory*, 2003.
- [11] T. Ho, M. Médard, M. Effros, R. Koetter, and D. R. Karger. Network coding for correlated sources. In *Proc. Conference on Information Sciences and Systems*, 2004.
- [12] T. Ho, M. Médard, R. Koetter, D. Karger, M. Effros, J. Shi, and B. Leong. A random linear network coding approach to multicast. *IEEE Transactions on Information Theory*, 52(10):4413–4430, October 2006.
- [13] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger. On randomized network coding. In *Proc. 41st Annual Allerton Conference on Communication, Control, and Computing*, 2003.
- [14] T. Ho and H. Viswanathan. Dynamic algorithms for multicast with intra-session network coding. In *Proc. 43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005.
- [15] T. Klein and H. Viswanathan. Centralized power control in multi-hop wireless networks. In *Proc. IEEE International Symposium on Information Theory*, 2003.
- [16] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 11(5):782–795, October 2003.
- [17] X. Lin and N. B. Shroff. The impact of imperfect scheduling on cross-layer congestion control in wireless networks. *IEEE/ACM Transactions on Networking*, 14(2):302–315, April 2006.
- [18] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, 1995.
- [19] S. Mukherjee and H. Viswanathan. Throughput range tradeoff of wireless mesh backhaul networks. *IEEE Journal on Selected Areas in Communications*, 24(3):593 – 602, March 2006.
- [20] M. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time-Varying Channels*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2003.
- [21] M. Neely. Energy optimal control for time varying networks. In *Proc. IEEE INFOCOM*, 2005.
- [22] M. Neely, E. Modiano, and C.-P. Li. Fairness and optimal stochastic control for heterogeneous networks. In *Proc. IEEE INFOCOM*, 2005.
- [23] M. Neely, E. Modiano, and C. Rohrs. Packet routing over parallel time-varying queues with application to satellite and wireless networks. In *Proc. Conference on Information Sciences and Systems*, 2002.
- [24] M. Neely, E. Modiano, and C. Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications*, 23(1):89–103, Jan. 2005.
- [25] S. Sarkar and L. Tassiulas. A framework for routing and congestion control for multicast information flows. *IEEE Transactions on Information Theory*, 48(10):2690– 2708, October 2002.
- [26] D. Slepian and J. K. Wolf. Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory*, 25:471 – 480, July 1973.

- [27] L. Tassiulas and A. F. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, Dec. 1992.
- [28] H. Viswanathan and K. Kumaran. Rate scheduling for multiple antenna downlink wireless systems. In *Proc. Thirty-Ninth Annual Allerton Conference on Communications and Control*, 2001.
- [29] Y. Wu, P. A. Chou, Q. Zhang, K. Jain, W. Zhu, and S.-Y. Kung. Network planning in wireless ad hoc networks: A cross-layer approach. *IEEE Journal on Selected Areas in Communications, Special Issue on Wireless Ad Hoc Networks*, 23(1):136 – 150, January 2005.
- [30] E. Yeh and A. Cohen. Throughput and delay optimal resource allocation in multiaccess fading channels. In *Proc. IEEE International Symposium on Information Theory*, 2003.
- [31] E. Yeh and A. Cohen. Throughput optimal power and rate control in queued multiaccess and fading channels. In *Proc. IEEE International Symposium on Information Theory*, 2004.




Fig. 1. An example illustrating virtual queues. Each physical session c packet at source node s is intended to be multicast to the two sink nodes y, z , and so corresponds to one packet in each virtual queue Q_s^{cy} and Q_s^{cz} .

Tracey Ho (M) is an Assistant Professor in Electrical Engineering and Computer Science at the California Institute of Technology. She received a Ph.D. (2004) and B.S. and M.Eng degrees (1999) in Electrical Engineering and Computer Science (EECS) from the Massachusetts Institute of Technology (MIT). Her primary research interests are in information theory, network coding and communication networks.

Harish Viswanathan (SM) received the B. Tech. degree from the Department of Electrical Engineering, Indian Institute of Technology, Chennai, India in 1992 and the M.S. and Ph.D. degrees from the School of Electrical Engineering, Cornell University, Ithaca, NY in 1995 and 1997, respectively. He was recipient of the Cornell Sage Fellowship. He joined Lucent Technologies Bell Labs, NJ in 1997 and has worked on multiple antenna technology and other enhancements for third generation wireless, distributed radio network architecture, and next generation wireless technologies. He is currently a Director at Alcatel-Lucent focusing on Machine-to-machine communications and sensor networks strategy. His research interests include communication theory, wireless networks, information theory, and signal processing.




Fig. 2. An example illustrating a physical broadcast transmission with two virtual transmissions, for a multicast session with two sinks. Each oval corresponds to a node. The left node broadcasts a physical packet received by the two right nodes, one of which adds the packet to the virtual queue for sink 1, and the other, to the virtual queue for sink 2.

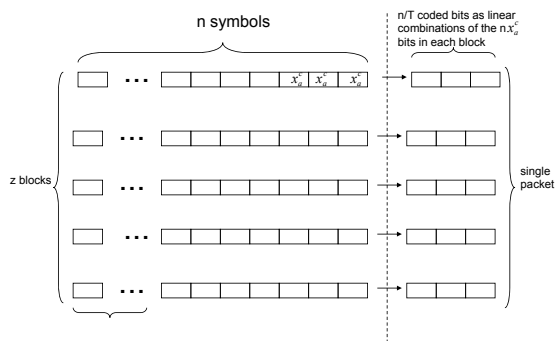


Fig. 3. Illustration of the packet structure for the case of correlated sources, courtesy of an anonymous reviewer.