

Full length article

The role of coding in the choice between routing and coding for wireless unicast



Ramakrishna Gummadi^{a,*}, Laurent Massoulié^b, Ramavarapu Sreenivas^c

^a Department of Management Science and Engineering, Stanford University, United States

^b Technicolor Paris Research Labs, Paris, France

^c Department of Industrial and Systems Engineering, University of Illinois at Urbana Champaign, United States

ARTICLE INFO

Article history:

Received 16 February 2011

Received in revised form 7 March 2012

Accepted 3 May 2012

Available online 11 May 2012

Keywords:

Erasures

Feedback

Network Coding

Routing

Capacity

ABSTRACT

We consider the benefits of coding in wireless networks, specifically its role in exploiting the local broadcast property of the wireless medium. We first argue that for unicast, the throughput achieved with network coding is the same as that achieved without any coding. This argument highlights the role of a general max-flow min-cut duality and is more explicit than previous proofs. The maximum throughput can be achieved in multiple ways without any coding, for example, using backpressure routing, or using some centralized flow scheduler that is aware of the network topology. However, all such schemes, in order to take advantage of the local broadcast property, require dynamic routing decisions for choosing the next hop for each packet from among the nodes where it is successfully received. This choice seems to depend critically on feedback signaling information like queue lengths, or ARQ. In contrast, note that the use of network coding can achieve the same without such feedback, in exchange for decoding overhead.

A key issue to be resolved in making a comparison between routing and coding would be how critical feedback signaling is, for the throughput of routing policies. With this motivation, we first explore how feedback at a given node affects its throughput, with arbitrary rates of its one-hop neighbors to the destination. *Static* routing policies which are essentially *feedback independent*, are considered. An explicit characterization of the optimal policies under such a feedback constraint is obtained, which turns out to be a natural generalization of both flooding and traditional routing (which does not exploit local broadcast, because the next hop is fixed prior to the transmission). When losses at the receivers are independent (still allowing for dependencies on transmissions by two different nodes, to model interference), the reduction in capacity due to constraining the feedback is limited to a constant fraction ($e^{-1} = 37\%$) of the coding capacity, and gets arbitrarily close to optimal as the unconstrained capacity goes to zero. We also extend this analysis to a layered multihop network and also compare the throughput of flooding to backpressure via simulations for a layered network assuming independent losses. Finally, if there are dependencies in the losses seen by receivers from a single broadcast, the reduction could be arbitrarily bad, even with just two hops.

Published by Elsevier B.V.

1. Introduction

We consider wireless erasure networks, where local broadcast influences the role of coding beyond that of

merely dealing with lossy transmissions, namely erasures. We show that feedback signaling is a critical factor that defines the role of coding in this situation, in the sense that it is one way to avoid the extensive feedback signaling that is necessary for routing policies. To characterize this more precisely, we consider a formal notion of restricted feedback signaling and derive the throughput of routing

* Corresponding author.

E-mail address: gsk84@gmail.com (R. Gummadi).

policies with restricted feedback on a two-hop network. This allows us to obtain a lower bound on the throughput when the losses are independent, and also to show that it is possible to have arbitrary degradation of throughput with dependent losses.

Local broadcast refers to the property of wireless systems where a single transmission can potentially address a group of clients. Although this is implicit for any wireless transmission, abstractions like the protocol model assume a graph based model in which local broadcast is essentially treated as interference (e.g. [1–4] and many others). In reality, local broadcast represents a diversity gain associated with the channel that could be exploited by either a well-adapted routing scheme or by network coding. In general, coding in the network has two distinct roles in the absence of this local broadcast feature. The first arises from the fact that even without any losses, information flow is fundamentally different from commodity flow, an aspect that points to the field of network coding [5], starting with the pioneering work of [6]. In wireline settings, network coding addresses strictly more general settings than that of a unicast transmission. On the other hand, forward error correction (FEC) addresses another aspect that stems from lossy transmissions. The role of coding in the context of wireless broadcast is something that needs to be distinguished from the previous two concerns, which are relatively better understood. Notably, in this context, the actual performance of optimal coding schemes is better articulated than are the limits of routing policies, both with and without restrictions on feedback signaling. This creates an avenue to discuss the role of coding by way of a better understanding about the limits of optimal routing schemes and then contrasting it with optimal coding schemes.

In this paper, we consider the role of coding, specifically in exploiting the local broadcast property of the wireless medium. We first argue that for unicast, the throughput achieved with network coding is the same as that achieved without any coding. This argument highlights the role of a general max-flow min-cut duality and is more explicit than previously published proofs of this fact. The maximum throughput can be achieved in multiple ways without any coding, for example, using backpressure routing, or using some centralized flow scheduler that is aware of the network topology. However, all such schemes, in order to utilize the local broadcast property, require dynamic routing decisions for choosing the next hop for each packet from among the nodes where it is successfully received. This choice seems to depend critically on feedback signaling information like queue lengths, or ARQ. In contrast, the use of network coding can achieve the same without such feedback, in exchange for decoding overhead.

A key task in comparing routing and coding is to assess the importance of feedback signaling to the throughput of routing policies. With this motivation, we explore how feedback at a given node affects its throughput, for an arbitrary vector of given rates of its one-hop neighbors to the destination. *Static* routing policies which are essentially *feedback independent*, are considered. An explicit characterization of the optimal policies under such

a feedback constraint is obtained, which can be interpreted as a natural generalization of both flooding and traditional routing (which does not exploit local broadcast, because the next hop is fixed prior to the transmission). When losses at the receivers are independent (still allowing for dependencies on transmissions by two different nodes, to model interference), the reduction in capacity due to constraining the feedback is limited to a constant fraction ($e^{-1} = 37\%$) of the coding capacity, and gets arbitrarily close to optimal as the capacity itself is low. This result is also extended to a more general version on feed-forward networks without any assigned rates of the one-hop neighbors to the destination. However, if there are dependencies in the losses seen by receivers from a single broadcast, the reduction could be arbitrarily bad, even with just two hops.

Network coding was originally introduced in [6] as a general framework to achieve the optimal multicast rate from a data source to a set of receivers in wired networks. In contrast to “store-and-forward”, also called a “routing” operation, network coding performs recombinations of data packets at network nodes, while the former operation never alters original packets. Since then, many other applications of network coding have been identified. In particular, at a practical level, it has been considered in the context of wireless networks [7] for unicast communications, where such a wireless setting was identified as an especially good candidate for network coding because of the local broadcast property. The wireless scenario of [7] features lossy transmissions as well as local broadcast, although it is not specific to the erasure channel. The experimental evaluations of [7] show benefits of network coding over routing policies in terms of the transmission rates achieved. However, Smith and Hassibi [8] prove that in any erasure network, there exists a policy that can achieve capacity with no coding. This raises the question: Are there any benefits of network coding over routing in wireless unicast communications? If so, where do they stem from, and how large can they be?

In a nutshell, we show that the benefits of coding over routing depend on the extent of available feedback signaling and characterize the relation. Our first contribution is a simpler argument, using linear programming (LP) duality, that the maximum routing throughput in the case of wireless unicast is the same as the rate achieved with coding. Using this, it is shown that backpressure routing achieves not only the optimal throughput among routing policies, but also more generally across policies that involve coding. However, backpressure policy performs a very dynamic routing of each packet through the network, and requires exchanging queue lengths from all neighbors at every step. This motivates an investigation into the fundamental limits of *static* routing policies which do not need extensive feedback signaling. We study this tradeoff and quantify the limits imposed by an appropriately defined notion of restricted feedback.

Consider the following extreme case to motivate the analysis (see Fig. 1): There are m distributed agents with each agent holding a copy of a set of m distinct packets (which they all received from the source). With full coordination, they can schedule all the m packets in one

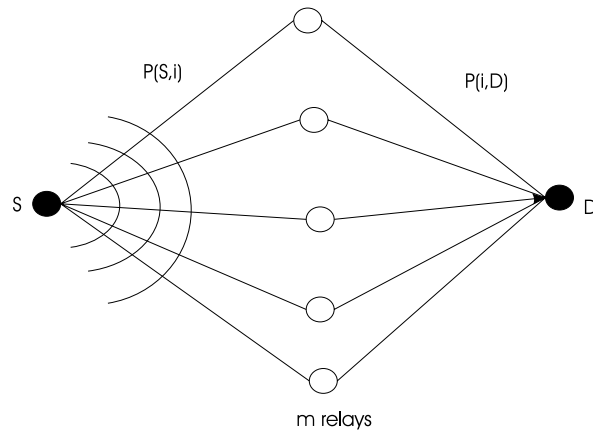


Fig. 1. The relay network example.

transmission each by avoiding duplication. However, if we restrict their choice to be made in a distributed manner without coordination, the total number of distinct packets covered by a random choice at each relay approaches $1 - (1 - 1/m)^m \rightarrow 1 - e^{-1}$, thus leading to only 63% throughput. 100% throughput could be achieved by either (i) making coordinated choices among relay nodes through conferencing among all the relays, or (ii) letting each node send an independent random linear combination of all m packets. In other words, network coding seems to be essentially solving a distributed synchronization problem without the need for any feedback signaling.

The above synchronization constraint is modeled as the *feedback independent routing* (FIR) restriction (see Definition 6): informally, this says that the decision of whether a node chooses to forward a packet should not depend on erasures at other nodes. In all examples of policies that achieve the capacity without any coding, this information is implicitly utilized based on feedback signaling. Under the FIR restriction, we show that the optimal policies are characterized as *tagging policies* where each packet being broadcast is assigned multiple next hops. *Tagging policies* can be considered a generalization of both flooding (where every broadcast packet is routed by everyone that receives it), and traditional routing (where each packet is routed only by a specific receiver chosen prior to the broadcast). When transmissions are subject to mutually independent losses¹ at the receivers (but still possibly allowing for dependences of losses for packets transmitted by two different nodes), even when restricted to feedback independent routing, it is possible to achieve at least 63% of the capacity. In fact, as the capacity, $C^* \rightarrow 0$, the throughput achieved becomes 100% of C^* in the limit. For a general feed-forward network with $h + 1$ hops in which all nodes are restricted to operate under a similar routing constraint, we show under a similar independence condition for link losses (but allowing dependence across different broadcasts) that the reduction in throughput is lower bounded as $f^h(C^*)$, where $f(x) =$

$1 - e^{-x}$. Thus, for a limited number of hops, and when the capacity is low to begin with, one can achieve close to optimal throughput without actually making dynamic routing choices. One might imagine that dependencies in link losses supply implicit information about other link losses, and therefore allow better throughput compared to independent losses. But this is not true, as we show a counterexample with dependent link losses for which static feedback independent routing capacity is arbitrarily bounded away from the network coding capacity even for a 2-hop network. One of the implications in this conclusion is that, when feedback is constrained, coding in the network could be unavoidable to achieve non-vanishing throughput.

2. Related work

The merits of routing versus coding have been extensively studied in the context of wireless, both theoretically and by experiments. The following two lines of work deal with the ways in which local broadcast can be exploited: (i) By using coding: This has been investigated in [9,10,7]. The advantage of using network coding to exploit local broadcast is the lack of a need for sophisticated coordination and/or routing choices among the nodes. The price to pay for this is the decoding complexity. (ii) By making dynamic routing/flooding choices along with rich feedback signaling: This was the path taken in [11,8,12]. The maximum throughput in an information theoretic sense for the wireless erasure network (WEN) model was studied in [10] and was shown to be equal to the appropriate hypergraph min-cut, denoted C^* . In [8], for the unicast wireless setting, a flooding based policy with network wide instantaneous broadcast of the identity of each packet received at the destination was shown to achieve C^* . Using this, [8] make an important observation that coding is in fact not necessary for achieving a throughput equal to the capacity in wireless unicast settings. The model used by Smith and Hassibi differs slightly from ours, in the sense that they consider a continuous time data arrival model, but it is not surprising that the same result holds for the discrete arrival model. As one of our contributions, we provide a more fundamental reason based on max-flow min-cut equivalence. In [12],

¹ i.e., the random variables representing the losses events are mutually independent

backpressure routing (requiring rich feedback signaling to perform routing) was shown to be an optimal routing plus scheduling policy in a much more general context taking into account interference effects. In general (i.e. for multicast), [13] propose the use of feedback together with network coding for online decoding. In our work, since we consider a single unicast flow throughout, our results are relevant only to intra-session coding. In contrast, when using network coding across multiple parallel flows as in [14], the situation corresponds to inter-session coding.

3. The wireless erasure network model

Let $G = (V, E)$ be a directed graph. Let $\mathcal{N}(i) = \{j \in V : (i, j) \in E\}$. We consider a wireless network that operates on this graph over time $t \in \{0, 1, 2, \dots\}$. For each t , a node can “broadcast” to its neighbors. The network is subject to probabilistic constraints on the successes of these broadcasts. Specifically, at any given time, t , for each $i \in V$, $Z \subseteq \mathcal{N}(i)$, let $\chi(i, Z, t)$ denote a $\{0, 1\}$ random variable that represents the following:

$$\chi(i, Z, t) = \begin{cases} 1, & \text{if broadcast from node } i \text{ at time} \\ & t \text{ is successful to } Z, \\ & \text{and fails to } \mathcal{N}(i) \setminus Z \\ 0, & \text{otherwise.} \end{cases}$$

The random variables χ can be arbitrarily correlated across the argument i , which allows for modeling arbitrary interference constraints, but we will assume that they are independent across t . For example, this could model the situation in random access scheduling, where at each time slot, there is a probability that a given node actually transmits, with the transmission being successful if and only if no other node in its interference radius is simultaneously active. Note that, for Z , we have: $\sum_{Z \subseteq \mathcal{N}(i)} \chi(i, Z, t) = 1 \forall i, t$. We define: $c(i, Z) = E[\chi(i, Z, t)] \forall t$. Thus, we are given a network topology along with $c(i, Z)$ as the capacities. Our analysis will be impervious to the correlations across i , so we will not explicitly specify them. An interesting special case of the above model is to have link (i, j) successful with probability $p(i, j)$, with losses from i to each of the neighbors in $\mathcal{N}(i)$ being independent. We then have:

$$c(i, Z) = \prod_{j \in Z} p(i, j) \prod_{j \in \mathcal{N}(i) \setminus Z} (1 - p(i, j)).$$

We consider a single unicast flow. However, the insights obtained are also general enough for multiple competing flows, with each flow assigned a fixed fraction of the link capacities and a comparison with intra-session network coding in such a scenario. Thus, without loss of generality we assume a source, S , and destination, D . The source has an infinite set of packets indexed over the integers, intended for replication at D . For any $v \in V$, let $\alpha_v(t)$ denote the number of distinct packets that were replicated at node v till time slot t .

Definition 1 (Routing Policy, \mathcal{P}). A routing policy \mathcal{P} decides for each node $i \in V$, and time $t \in \{0, 1, \dots\}$, a packet to be transmitted from among the $\alpha_i(t-1)$ choices in its possession.

Definition 2 (Capacity). The throughput of a policy \mathcal{P} is defined as:

$$C(\mathcal{P}) = \liminf_{t \geq 0} \frac{E[\alpha_D(t)]}{t}.$$

The capacity is the highest possible throughput $C \triangleq \sup_{\mathcal{P}} C(\mathcal{P})$.

4. A max-flow characterization

A key observation is that any policy which uniquely routes each packet without keeping multiple copies can be represented by a valid flow on the throughput constrained graph. In the linear program defined below, each feasible solution represents a policy that routes a fraction proportional to $r(i, j, Z)$ of successful broadcasts from i to Z uniquely to $j \in Z$. The term $\sum_{\{Z \subseteq \mathcal{N}(i); j \in Z\}} r(i, j, Z)$ represents the net flow from i to j . The optimum value of the linear program (LP) then represents the throughput achieved.

Definition 3 (F , The Maxim Flow Value). Let P denote the set of all $S - D$ paths in G . We define the F for a given broadcast capacitated graph as the optimum value of the following LP:

$$F = \max \sum_{p \in P} x_p \quad (1)$$

Subject to: $x_p \geq 0 \quad \forall p \in P$

$$r(i, j, Z) \geq 0 \quad \forall \{(i, j, Z) : (i, j) \in E, j \in Z \subseteq \mathcal{N}(i)\}$$

$$\sum_{\{p \in P: (i, j) \in p\}} x_p - \sum_{\{Z \subseteq \mathcal{N}(i); j \in Z\}} r(i, j, Z) \leq 0 \quad \forall (i, j) \in E$$

$$\sum_{j \in Z} r(i, j, Z) \leq c(i, Z) \quad \forall \{(i, Z) : Z \subseteq \mathcal{N}(i)\}$$

Let x_p^* , $r^*(i, j, Z)$ denote the optimum solution to the above LP. It is straightforward to show that this capacity can be achieved by the following policy.

Definition 4 (\mathcal{P}_{fs} , The Flow Splitting Policy). Any packet transmitted by node i , and received by the set $Z \subseteq \mathcal{N}(i)$ of its neighbors is “routed” uniquely to $j \in \mathcal{N}(i)$ with probability $\frac{r^*(i, j, Z)}{\sum_{k \in Z} r^*(i, k, Z)}$ (thus ensuring that at most one copy of each distinct packet is being transmitted at any point).

We now consider the min-cut appropriate for the model by considering the probability that at least one of the nodes across the cut receives a transmission.

Definition 5 (Minimum Cut, C^*). A cut is a disjoint partition of V into A and \bar{A} with $S \in A$ and $D \in \bar{A}$. The capacity of the cut is then:

$$C(A) = \sum_{i \in A, Z \subseteq \mathcal{N}(i), Z \cap \bar{A} \neq \emptyset} c(i, Z).$$

The minimum cut is: $C^* = \min_A C(A)$.

It is easy to argue that C^* is an upper bound on the throughput for any scheme even with coding, and was

in fact shown to be equal to the information theoretic capacity of the WEN in [10]. Thus, $C \leq C^*$, where we are reminded that C stands for the routing capacity and C^* stands for the min-cut, which is equal to the capacity with coding, namely the information theoretic capacity. Based on a duality argument analogous to the classical max-flow min-cut theorem, the following can be shown.

Theorem 1. $F = C^* = C$.

Proof. Consider the dual program for the LP at Eq. (1) with dual variables $b(i, Z)$ for each $i \in V$, $Z \subseteq \mathcal{N}(i)$ and $y(i, j)$ for each $(i, j) \in E$. We have:

$$DUAL^* = \min \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z)b(i, Z) \quad (2)$$

Such that:

$$\sum_{(i,j) \in p} y(i, j) \geq 1 \quad \forall p \in P \quad (3)$$

$$-y(i, j) + b(i, Z) \geq 0$$

$$\forall \{(i, j, Z) : (i, j) \in E, j \in Z \subseteq \mathcal{N}(i)\}$$

$$y(i, j) \geq 0 \quad \forall (i, j) \in E$$

$$b(i, Z) \geq 0 \quad \forall (i, Z) : i \in V, Z \subseteq \mathcal{N}(i). \quad (4)$$

Consider the min-cut as written in Eq. (5), and let A^* , \bar{A}^* denote this cut. Let $y^*(i, j) = 1$ if $i \in A^*$, $j \in \bar{A}^*$ and 0 otherwise. Similarly, let $b^*(i, Z) = 1$ if $i \in A^*$, $Z \cap \bar{A}^* \neq \phi$ and 0 otherwise. Then, it can be verified that this defines a feasible solution to the dual LP above, and that $\sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z)b^*(i, Z) = C^*$. Thus,

$$C^* \geq DUAL^*. \quad (5)$$

Now consider an integral constrained version of the above dual:

$$DUAL_* = \min \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z)b(i, Z) \quad (6)$$

Such that:

$$\sum_{(i,j) \in p} y(i, j) \geq 1 \quad \forall p \in P \quad (7)$$

$$-y(i, j) + b(i, Z) \geq 0$$

$$\forall \{(i, j, Z) : (i, j) \in E, j \in Z \subseteq \mathcal{N}(i)\} \quad (8)$$

$$y(i, j) \in \{0, 1\} \quad \forall (i, j) \in E$$

$$b(i, Z) \in \{0, 1\} \quad \forall (i, Z) : i \in V, Z \subseteq \mathcal{N}(i). \quad (9)$$

Let y_* , b_* define the optimal solution to the above integral constrained LP. Then, define $A_* = \{i \in V : \exists \text{ a path } p, \text{ from } S \text{ to } i \text{ such that } \sum_{(i,j) \in p} y_*(i, j) = 0\}$. Then, $D \notin A_*$ due to Eq. (7), and thus A_* defines an (S, D) cut. Furthermore, from Eqs. (6) and (8), we get $C(A_*) = DUAL_*$. This implies that

$$C^* \leq DUAL_*. \quad (10)$$

To summarize, we so far have:

$$F = DUAL^* \leq C \leq C^* \leq DUAL_*. \quad (11)$$

If we are able to argue that the constrained LP indeed achieves the optimum (the details of this are given in Lemma 2), we would then have $DUAL^* = DUAL_*$, implying that all quantities in Eq. (11) are the same. \square

Lemma 2. $DUAL^* = DUAL_*$.

Proof. We use an analogous argument employed in showing the corresponding statement for the classical max-flow min-cut theorem. The argument considers a probability distribution on the set of all possible cuts and argues that the expected value of C^* thus obtained is no more than $DUAL^*$, which in turn implies that $DUAL_* \leq DUAL^*$, thus completing the proof. Consider the dual LP in Eq. (2), and let $y^*(i, j)$, $b^*(i, Z)$ denote the optimal solution which achieves $DUAL^*$. Consider a graph with edge lengths given by $y^*(i, j)$ and let $d(i) \triangleq$ length of the shortest path from S , with edge lengths given by $y^*(i, j)$. Let $\lambda \in (0, 1)$ be chosen uniformly and define:

$$A^* = \{i \in V : d(i) \leq \lambda\}.$$

This defines a cut with probability 1, since $d(S) \geq 1$ from Eq. (3). Then:

$$\begin{aligned} E[C(A^*)] &= E \left[\sum_{i \in A^*, Z \subseteq \mathcal{N}(i), Z \cap \bar{A}^* \neq \phi} c(i, Z) \right] \\ &= E \left[\sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) \mathbb{1}\{i \in A^*, Z \subseteq \mathcal{N}(i), Z \cap \bar{A}^* \neq \phi\} \right] \\ &= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) P(i \in A^*, j \in \bar{A}^* \text{ for some } j \in Z) \\ &= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) P(d(i) \leq \lambda, d(j) > \lambda \text{ for some } j \in Z) \\ &= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) \left(\max_{j \in Z} d(j) - d(i) \right)_+ (\because \lambda \text{ is uniform}) \\ &\leq \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) \left(\max_{j \in Z} y^*(i, j) \right) (\text{triangle inequality}) \\ &\leq \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i, Z) b^*(i, Z) (\text{from Eq. (4)}) = DUAL^* \quad \square \end{aligned}$$

Thus, coding in the network is not necessary to achieve the optimal throughput in unicast, assuming unconstrained feedback signaling. This fact was also argued by Smith and Hassibi [8] in a closely related continuous time data arrival model (and conjectured and verified by simulation for the discrete time model that we consider). This was accomplished by considering a policy which involves flooding the network with each packet until at least one copy reaches the destination and subsequently using a network wide feedback signal to delete these copies every time the destination receives a new packet. It was shown that such a policy stabilizes the network for all rates below C^* using the Lyapunov stability argument. Lemma 3 shows that a distributed backpressure scheme which routes each packet to the least loaded neighbor (thus, avoiding multiple copies) also achieves the information theoretic min-cut capacity, C^* . This is also suggested by Theorem 1 in conjunction with Neely's result of optimality of backpressure routing schemes in a context that involves power control and with a different interference model [12].

Lemma 3. Consider a Markov chain defined on the network as follows: Each node has a queue of packets. New packets arrive to the queue at the source according to a Bernoulli process of rate λ . For any t, i, Z such that $\chi(i, Z, t) = 1$, the backpressure policy routes a packet from node i to a node j such that queue size difference is maximized (subject to being positive). If $\lambda < C^*$, the Markov chain thus obtained is stable (positive recurrent).

Proof. The proof follows along the lines of [15] using a quadratic Lyapunov function and Foster's condition to show that the Markov chain is positive recurrent. Packets are injected at the source node according to a Bernoulli i.i.d. process with mean $\lambda < C^*$. We shall use the potential, $V(t) = \sum_{i=1}^n q_i^2(t)$, where i represents an index over the n nodes in the network and $q_i(t)$ represents the number of packets that are held at node i . Let $\{\mathcal{F}_t\}_{t \geq 0}$ be a filtration adapted to the queue length process. By Foster's theorem (see [15] and the references therein), a sufficient condition for stability is to show that the Lyapunov drift $E[V(t+1) - V(t)/\mathcal{F}_t] < 0$ when $V(t)$ is sufficiently large. Let $\mathbf{q}(t)$ denote the n dimensional row vector with $q_i(t)$ being the i^{th} element. Let \mathbf{R} be the adjacency matrix of dimension $n \times L$ (where $L = |E|$, the number of links) for the given graph (i.e., the $(i, l)^{\text{th}}$ element, $r(i, l)$ is 1(−1) iff the link l starts(ends) at i , and 0 otherwise). Let $\mathbf{E}(t)$ denote the L dimensional indicator vector denoting the subset of links on which packets were routed by the backpressure policy, and let $\mathbf{A}(t)$ denote the arrival process indicator vector, i.e. the element of $\mathbf{A}(t)$ corresponding to the source is the Bernoulli random variable with mean λ and all other elements are 0. Then, the queue lengths evolve according to $\mathbf{q}(t+1) = \mathbf{q}(t) + \mathbf{R}\mathbf{E}(t) + \mathbf{A}(t)$. Thus (with \cdot denoting the usual dot product of vectors):

$$\begin{aligned} V(t+1) - V(t) &= \mathbf{q}(t+1) \cdot \mathbf{q}(t+1) - \mathbf{q}(t) \cdot \mathbf{q}(t) \\ &= (\mathbf{R}\mathbf{E}(t) + \mathbf{A}(t)) \cdot (2\mathbf{q}(t) + \mathbf{R}\mathbf{E}(t) + \mathbf{A}(t)) \\ &\leq (n+1)^2 + 2(\mathbf{q}(t) \cdot \mathbf{R}\mathbf{E}(t) + \mathbf{q}(t) \cdot \mathbf{A}(t)). \end{aligned}$$

Since the first term above is constant over time, showing that the second term has a large negative drift for large queue lengths is sufficient (since large potential implies large queue length under connectivity assumptions on the graph). Let $\mathbf{W}(t)$ denote the L dimensional vector where the l^{th} element denotes the queue length difference for the l^{th} link. Then, $\mathbf{q}(t) \cdot \mathbf{R}\mathbf{E}(t) = -\mathbf{W}(t) \cdot \mathbf{E}(t)$. Hence, we only need to argue that $E[-\mathbf{W}(t) \cdot \mathbf{E}(t) + \mathbf{q}(t) \cdot \mathbf{A}(t)/\mathcal{F}_t]$ has a negative drift. $E[\mathbf{q}(t) \cdot \mathbf{A}(t)/\mathcal{F}_t] = \mathbf{q}(t) \cdot \mathbf{A}$ where \mathbf{A} is a vector where the element corresponding to the source is λ and all other elements are 0. Since $\lambda < C^*$, it follows from the max-flow interpretation of Theorem 1 that there exist variables $r^*(i, j, Z)$ satisfying the constraints of the linear program given by Eq. (1), and an $\epsilon > 0$ such that $\mathbf{q}(t) \cdot \mathbf{A} \leq (1 - \epsilon)\mathbf{W}(t) \cdot \mathbf{f}^*$, where the component of \mathbf{f}^* corresponding to link $l = (i, j)$ is $f_l^* = \sum_{\{Z \in \mathcal{N}(i): j \in Z\}} r^*(i, j, Z)$. Let $\mathbf{e}^*(t)$ denote the indicator vector of dimension L , for the routing selected by the flow splitting policy, \mathcal{P}_{f_s} given in

Definition 4. Then, $\mathbf{W}(t) \cdot \mathbf{f}^* = E[\mathbf{W}(t) \cdot \mathbf{e}^*(t)/\mathcal{F}_t]$. Thus,

$$\begin{aligned} E[-\mathbf{W}(t) \cdot \mathbf{E}(t) + \mathbf{q}(t) \cdot \mathbf{A}(t)/\mathcal{F}_t] &= E[-\mathbf{W}(t) \cdot \mathbf{E}(t) + (1 - \epsilon)\mathbf{W}(t) \cdot \mathbf{e}^*(t)/\mathcal{F}_t] \\ &= E[-\mathbf{W}(t) \cdot \mathbf{E}(t) + \mathbf{W}(t) \cdot \mathbf{e}^*(t)/\mathcal{F}_{t+} / \mathcal{F}_t] \\ &\quad - \epsilon E[\mathbf{W}(t) \cdot \mathbf{e}^*(t)/\mathcal{F}_t]. \end{aligned}$$

Here, \mathcal{F}_{t+} denotes the sigma algebra that contains information about the successes on links at time t in addition to the queue length process until time t , and thus, $\mathcal{F}_t \subseteq \mathcal{F}_{t+} \subseteq \mathcal{F}_{t+1}$. The first term is non-positive because the backpressure policy minimizes $\mathbf{W}(t) \cdot \mathbf{E}(t)$ among all possible options conditioned on the information available on the link successes and the queue lengths (which is what conditioning on \mathcal{F}_{t+} denotes). The second term takes arbitrarily large negative values for large $V(t)$ for any fixed $\epsilon > 0$, and thus, we have the required negative drift. \square

5. Formalizing a notion of restricted feedback

While we have so far discussed schemes that can achieve the unicast capacity in a wireless network without the need for coding, employing coding in the network can achieve the capacity without using feedback. This calls for an understanding of the inherent limits to the achievable throughput when we restrict the exploitation of feedback in the choice of routing policies. The decision of whether to forward a packet further at a given node should ideally be made without considering the erasure events on other links (this is not the case with any of the routing schemes we have discussed so far which achieve the maximum throughput). In this context, we first fix the vector of rates that the one-hop neighbors can support to the destination simultaneously and study how feedback constrained routing affects the overall throughput from the source. The most obvious visualization of this is a network in which the source and destination are assisted by relays which do not communicate within themselves (Fig. 1).

Let $c(Z) = E[\chi(S, Z, t)]$ for $Z \subseteq [m]$. Let $p(S, i) = \sum_{Z \subseteq [m]: i \in Z} c(Z)$ denote the probability that a packet transmitted from S is received at i successfully. We shall let $p(i, D)$ denote the long-term throughput that relay i can support to the sink D . Given a generalized routing/flooding policy, we shall use the $\{0, 1\}$ random variables, $r_i(p)$, $r_i^*(p)$ which denote the following events:

1. $\{r_i(p) = 1\} \iff$ Packet transmitted in time slot p from the source (which shall henceforth be referred to as packet p) was received successfully by relay i . We assume that the source attempts broadcast of distinct packets at each time slot without any loss of generality (e.g., by employing appropriate source coding, or because it has an unlimited stream of useful packets).
2. $\{r_i^*(p) = 1\} \iff$ Packet p is routed to D via relay i . Note that: (1) $\{r_i^*(p) = 1\} \Rightarrow \{r_i(p) = 1\}$. (2) It is possible that $r_i^*(p) = 1$ for multiple i .

We will now explicitly describe what we mean by feedback independent routing.

Definition 6 (*Feedback Independent Routing (FIR)*). $\forall p \in \{0, 1, \dots\}$, $\forall A, B \subseteq [m]$ such that $A \cap B = \phi$, the given routing policy satisfies the FIR restriction if, conditioned on $\{r_i(p) = 1 \forall i \in A\}$, the following two collections of random variables: $\{r_i^*(p)\}_{i \in A}$ and $\{r_i(p)\}_{i \in B}$, are mutually independent.

This condition is essentially equivalent to assuming a lack of feedback to the broadcasting node. Technically, lack of feedback is sufficient but not entirely necessary to satisfy this. While this distinction is subtle, it might be noted that one does not violate FIR by using rudimentary feedback for purposes other than routing. Nevertheless, source coding is a convenient way to completely eliminate feedback. As for the first hop nodes, we merely consider them as black-boxes that support some arbitrary vector of simultaneous rates to the destination. For a 2-hop relay network, these simultaneous rates could be achieved by employing any capacity achieving scheme for a single erasure channel, i.e. either by (1) feedback to the relays, or (2) forward erasure coding (FEC) at each relay.²

We have so far discussed three policies which achieve capacity without any coding: (1) the flow splitting policy \mathcal{P}_{fs} in Definition 4, (2) the backpressure policy in Lemma 3 and (3) the policy of [8]; but they all involve a heavy interaction between feedback and routing, and hence fail to satisfy the given constraint.

6. Capacity under feedback independent routing constraint

Definition 7. The capacity with feedback independent routing, C_R is defined as the maximum throughput as in Definition 2, restricted to policies \mathcal{P} satisfying Definition 6.

For FIR, one extreme is to tag each packet with a single relay (this is how routing is done in practice in the 802.11 protocol). This could be suboptimal because it does not exploit the local broadcast advantage. The other extreme is to flood every packet to all relays. This could be suboptimal when the relays do not have enough capacity to forward all packets they receive to the destination. They will then have to make distributed decisions on which packets to forward from among the received packets, leading to redundant transmissions and a hence a decreased throughput. As we will show, when we restrict to FIR, the maximum throughput is achieved within a subclass of policies that we shall call the *tagging policies*. A tagging policy assigns to each packet a subset of relays, $Z \subseteq [m]$, which is independent of any feedback. We represent the fraction of packets that are tagged with Z as $t(Z)$. A relay i that receives a packet successfully routes the packet without dropping it if and only if $i \in Z$. These packets are retransmitted until the destination receives

them. The capacity of such policies can be expressed via the linear program below, where each feasible solution corresponds to a specific tagging policy. The last constraint in the LP states that the arrival rate of packets to any relay's queue has to be less than its forwarding capacity to the destination.

Definition 8 (*Tag Capacity, C_T*). C_T is defined as the optimum value of the following LP with variables $t(Z) \geq 0$ where $Z \subseteq [m]$.

$$C_T = \max \sum_{Z, Z' \subseteq [m]: Z \cap Z' \neq \phi} t(Z)c(Z') \quad (12)$$

$$\text{Subject to: } \sum_{Z \subseteq [m]} t(Z) \leq 1; \quad (13)$$

$$p(S, i) \left(\sum_{Z \subseteq [m]: i \in Z} t(Z) \right) < p(i, D) \quad \forall i \in [m]. \quad (14)$$

All packets that reach a relay successfully *and* have the relay in the tag will eventually reach D because of the constraint in Eq. (14) (which implies that the queue of packets at each relay is stable). Since the tags are chosen independent of the losses, the probability that a packet is successfully transmitted to some relay which is also included in its tag is $\sum_{Z \cap Z' \neq \phi} t(Z)c(Z')$, which can be readily shown to be equal to the throughput of the policy as per Definition 2.

Lemma 4. $C_R \geq C_T$

Lemma 4 holds due to the fact that any feasible solution to the LP gives us a tagging scheme whose throughput is equal to the value of the LP.

7. Optimality of tagging policies under FIR

Remarkably, these policies will now also be shown to be optimal in general under the FIR constraint, thus giving us an explicit characterization of C_R in the form of Theorem 5.

Theorem 5. $C_R \leq C_T$, implying $C_T = C_R$

Proof of Theorem 5. We look at any policy that satisfies FIR and show that it is possible to define appropriate variables that define a feasible solution to the LP in Definition 8, and for which the throughput is upperbounded by the objective function.

Consider any arbitrary policy. We will show that the expected number of distinct packets that reach the destination in k time slots cannot be more than kC_T as long as FIR in Definition 6 is satisfied, thus implying that the C_R is at most C_T . Define:

$$t(Z) = \frac{1}{k} \sum_{p=1}^k P \left(\frac{r_i^*(p) = 1 \forall i \in Z \text{ and } 0 \forall i \in [m] \setminus Z}{r_i(p) = 1 \forall p \in [m]} \right) \quad (15)$$

where we use the standard convention of writing $\frac{P(A \cap B)}{P(B)}$ as $P\left(\frac{A}{B}\right)$ for an event B with positive probability. It is easily

² The key aspect that distinguishes such FEC from network coding is that in the case of FEC, the destination has to be able to decode the data being encoded by each relay independently from the transmissions received from that specific relay alone, whereas with network codes, the relay only needs to collect the packets from all relays and jointly decode them.

verified that $t(Z) \geq 0$ and $\sum_{Z \subseteq [m]} t(Z) = 1$. We will now verify that the third constraint (Eq. (14)) holds for any $i \in [m]$.

$$\begin{aligned}
 kp(i, D) &\geq E \left[\sum_{p=1}^k \mathbb{1}\{r_i^*(p) = 1\} \right] \\
 &= \sum_{p=1}^k P(r_i(p) = 1) P \left(\frac{r_i^*(p) = 1}{r_i(p) = 1} \right) \\
 &= p(S, i) \sum_{p=1}^k P \left(\frac{r_i^*(p) = 1}{r_i(p) = 1} \right) \\
 &= p(S, i) \sum_{p=1}^k P \left(\frac{r_i^*(p) = 1}{r_j(p) = 1 \forall j \in [m]} \right) \quad (\text{using Definition 6}) \\
 &= p(S, i) \sum_{p=1}^k \sum_{Z \ni i} P \\
 &\quad \times \left(\frac{r_j^*(p) = 1 \forall j \in Z \text{ and } 0 \forall j \in [m] \setminus Z}{r_j(p) = 1 \forall j \in [m]} \right) \\
 &= kp(S, i) \sum_{Z \ni i} t(Z) \quad (\text{by definition of } t(Z) \text{ in Eq. (15)}).
 \end{aligned}$$

We now calculate the number of distinct packets replicated at D in k time slots.

$$\begin{aligned}
 E[\alpha_D(k)] &= \sum_{p=1}^k P \left(\left\{ \bigcup_{i \in [m]} \{r_i^*(p) = 1\} \right\} \right) \\
 &= \sum_{p=1}^k \sum_{Z \subseteq [m]} P(r_i(p) = 1 \forall i \in Z \text{ and } 0 \forall i \in [m] \setminus Z) \\
 &\quad \times P \left(\frac{r_i^*(p) = 1 \text{ for some } i \in [m]}{r_i(p) = 1 \forall i \in Z \text{ and } 0 \forall i \in [m] \setminus Z} \right) \\
 &= \sum_{p=1}^k \sum_{Z \subseteq [m]} c(Z) P \left(\frac{r_i^*(p) = 1 \text{ for some } i \in Z}{r_i(p) = 1 \forall i \in Z \text{ and } 0 \forall i \in [m] \setminus Z} \right) \\
 &\quad \times (\because \forall i \in [m] \setminus Z, \{r_i(p) = 0\} \Rightarrow \{r_i^*(p) = 0\}) \\
 &= \sum_{p=1}^k \sum_{Z \subseteq [m]} c(Z) P \left(\frac{r_i^*(p) = 1 \text{ for some } i \in Z}{r_i(p) = 1 \forall i \in [m]} \right) \\
 &\quad \times (\text{using FIR Definition 6}) \\
 &= \sum_{p=1}^k \sum_Z c(Z) \sum_{Z' \cap Z \neq \emptyset} P \left(\frac{r_i^*(p) = 1 \forall i \in Z' \text{ and } 0 \text{ else}}{r_i(p) = 1 \forall i \in [m]} \right) \\
 &= k \sum_{Z \subseteq [m]} c(Z) \sum_{Z' \subseteq [m]: Z' \cap Z \neq \emptyset} t(Z) \\
 &= k \sum_{Z, Z' \subseteq [m]: Z \cap Z' \neq \emptyset} t(Z) c(Z'). \quad \square
 \end{aligned}$$

In general, it is not obvious if C_R is strictly less than the general min-cut, C^* . Indeed, in many cases, these two quantities match, implying that in such cases, not only do we not need any coding, but the capacity can be

achieved by optimized tagging schemes that satisfy FIR. For example, consider a network with two relays with success probabilities $p(S, 1), p(S, 2), p(1, D), p(2, D)$ where link losses for the same transmission are independent.

Example 1. If $p(S, 1), p(S, 2), p(1, D), p(2, D)$ are all $1/2$, we have: $C(\{1\}) = C(\{2\}) = C(\{1, 2\}) = 1/4$. From Fig. 2(a), $C^* = 3/4$ and from the LP (Eq. (12)), we can calculate that $C_R = 3/4$. In fact, a flooding policy achieves this rate of $3/4$. So there is no reduction in the capacity under FIR.

Example 2. Consider an example where the cuts are more finely balanced: $p(S, 1), p(S, 2) = 1/2$; and $p(1, D), p(2, D) = 3/8$. Again, $C^* = 3/4$, as explained in Fig. 2(b). Set $t(\{1\}) = x, t(\{2\}) = y, t(\{1, 2\}) = z$. Then: $C_R = \text{Max}(\frac{1}{2}x + \frac{1}{2}y + \frac{3}{4}z)$ Subject to: $x, y, z \geq 0, x + y + z \leq 1, x + z \leq 3/4, y + z \leq 3/4$. The optimal value can be verified to be $5/8 < 3/4 = C^*$.

This optimum throughput of $5/8$ under FIR is achieved by routing 25% of the packets exclusively to relay 1, 25% exclusively to relay 2 and the remaining 50% of the packets to both the relays.

8. Quantifying the loss of throughput under FIR

In this section, we will use the theory from previous sections to obtain results on the throughput attainable under FIR. We will first consider the case where the link losses from a given transmitter seen at its various receivers are independent. Note that this involves no assumptions on interference between two different transmissions. As a critical tool, we will use a *flooding policy*, which is defined below.

Definition 9 (Flooding Policy, \mathcal{P}_F). Each relay blindly chooses each received packet from the source to be forwarded to the destination with probability $\min(1, p(i, D)/p(S, i))$. In other words, the relay effectively makes a uniformly random selection of a $p(i, D)/p(S, i)$ fraction of its received packets to be forwarded to the destination whenever $p(i, D) \leq p(S, i)$.

We named the above policy as flooding in a general sense, because every packet that is successfully received is considered for being forwarded at any relay. It is possible that the total number of received packets at each relay could be more than the rate it can support to the destination, in which case \mathcal{P}_F essentially makes a random selection of packets corresponding the maximum throughput it can support. The throughput of \mathcal{P}_F can be calculated by evaluating the probability that a packet transmitted by the source in any given time slot reaches the destination along at least one of the m relays. Since these events are independent under our current hypothesis, we have Lemma 6, which is obtained from calculating the probability that any given packet reaches the destination via *at least* one of the relays.

Lemma 6.

$$C(\mathcal{P}_F) = 1 - \prod_{i \in [m]} (1 - \min(p(S, i), p(i, D))).$$




Fig. 2. Examples.

9. A lower bound for independent erasures

Using this flooding policy, we now prove a lower bound on C_R for a network with arbitrary fixed rates from the first hop forward. This argument explicitly bounds the loss of throughput because of the redundant transmissions arising out of making distributed decisions at each of the relay nodes with arbitrary min cuts when we have independence among link successes. The bound we provide also implies that when the capacity is low, the flooding is almost optimal for independent losses to the first hops. We will provide the proof on a 2-hop network for easier visualization, but the same claim holds for any fixed rates from the first hop to the destination.

Theorem 7. Consider the relay network of Fig. 3 with independent losses from the source to the relays.

$$C_R \geq 1 - e^{-C^*}.$$

It follows from this bound that, as $C^* \rightarrow 0$, $C_R/C^* \rightarrow 1$.

Proof. The argument is based on comparing the throughput of the flooding policy $C(\mathcal{P}_F)$ with C^* , since \mathcal{P}_F is a policy that satisfies FIR. We will show that $C(\mathcal{P}_F) \geq 1 - e^{-C^*}$, which in turn implies the theorem since \mathcal{P}_F clearly satisfies FIR (Definition 6).

Recall that:

$$C(\mathcal{P}_F) = 1 - \prod_{i \in [m]} (1 - \min(p(S, i), p(i, D))).$$

Applying the definition of the min-cut (see Fig. 3 for illustration), C^* , to the network under consideration, we see that

$$C^* = \min_{A \subseteq [m]} \left(1 - \prod_{i \in [m]/A} (1 - p(S, i)) + \sum_{i \in A} p(i, D) \right).$$

Let A^* be the arg min over $A \subseteq [m]$ for the above equation, so that:

$$C^* = 1 - \prod_{i \in [m]/A^*} (1 - p(S, i)) + \sum_{i \in A^*} p(i, D). \quad (16)$$

Consider any $i \in A^*$. By definition of A^* , we have:

$$\begin{aligned} 1 - \prod_{j \in [m]/A^*} (1 - p(S, j)) + \sum_{j \in A^*} p(j, D) \\ \leq 1 - (1 - p(S, i)) \prod_{j \in [m]/A^*} (1 - p(S, j)) \\ + \sum_{j \in A^*} p(j, D) - p(i, D) \end{aligned}$$

which implies:

$$p(i, D) \leq p(S, i) \prod_{j \in [m]/A^*} (1 - p(S, j)) \leq p(S, i). \quad (17)$$

Thus:

$$\begin{aligned} C(\mathcal{P}_F) &= 1 - \prod_{i \in [m]} (1 - \min(p(S, i), p(i, D))) \\ &\geq 1 - \prod_{i \in A^*} (1 - \min(p(S, i), p(i, D))) \\ &= 1 - \prod_{i \in A^*} (1 - p(i, D)) \\ &\quad \times (\because p(i, D) \leq p(S, i) \forall i \in A^* \text{ from Eq. (17)}) \\ &\geq 1 - \prod_{i \in A^*} e^{-p(i, D)} (\because 1 - x \leq e^{-x} \forall x \geq 0) \\ &= 1 - e^{-\sum_{i \in A^*} p(i, D)} \\ &\geq 1 - e^{-C^*} (\because \sum_{i \in A^*} p(i, D) \leq C^* \text{ from Eq. (16)}). \quad \square \end{aligned}$$

10. Extension to layered multi-hop networks with independent losses

The prior analysis characterizes the degradation through the effects of feedback at a single node, given the rates achievable from the subsequent hops in any general network. In this section, we build upon this argument to obtain a bound for arbitrary feed-forward networks where every node is restricted to static routing and without any assumption of achievable rates from the one-hop




Fig. 3. Illustration of a min-cut and the set defining A^* for this min-cut.




Fig. 4. An illustration of the layered feed-forward networks being considered.

neighbors. Further, we also simulate both flooding and backpressure policies for a layered network with ten hops and plot the results obtained.

As illustrated in Fig. 4, the feed-forward networks have nodes that are divided into layers where broadcasts from any node in layer i are heard only by nodes in layer $i + 1$.

Theorem 8. For a general feed-forward network with $h + 1$ hops, subject to feedback independent routing, the capacity is at least $f^h(C^*)$ where $f(x) = 1 - e^{-x}$ and C^* is the min cut capacity.

Proof. The argument builds on the analysis in Theorem 7. First, we note that feed-forward networks can be reduced to general layered networks where only nodes at subsequent levels are connected, by introducing dummy nodes with unit capacity links. For such a layered network, we will adopt the convention that the sink is at level 0, and the source is at level $h + 1$. We will define (i, j) to be the node indexed j at level i . In this notation, the source is assumed to be $(h + 1, 0)$, and the sink is $(0, 0)$. Given any policy, \mathcal{P} , we define $\mathcal{P}(i, j)$ as the rate at which distinct packets are streaming to the sink through (i, j) . For example, $\mathcal{P}(h + 1, 0)$ is the throughput of the policy.

Consider ψ , the flooding policy satisfying FIR, and compare it with the splitting policy, \mathcal{P}_{fs} in Definition 4: A node b queues the packets that it receives from node a in proportion to the rate $r^*(a, b) \doteq \sum_{Z \subseteq \mathcal{N}(a)} r^*(a, b, Z)$ with r^* as in Definition 4. Note that this is a static calculation and involves no feedback unlike \mathcal{P}_{fs} or the backpressure policy, which requires queue length information to decide every routing step. For such a policy, we make the following claim, which can be shown using induction on i .

Claim 1. $\psi(i, j) \geq f^{i-1}(\mathcal{P}_{fs}(i, j))$

The claim follows easily for $i = 1$. Assume that it holds for some i . We will look at a node $(i + 1, j)$ which is connected to $(i, j_1), \dots, (i, j_k)$. Consider two different 2-hop networks both with $(i + 1, j)$ as the source and $(i, j_1), \dots, (i, j_k)$ as relays with the same link characteristics as the feed-forward network for the first hop. To describe the second hop capacities, we will use the assumption that both \mathcal{P}_{fs} and ψ choose the same fraction of the incoming capacities from various nodes for the forward throughput. Let α_t be the fraction of incoming rate that is chosen by node (i, j_t) for the throughput from $(i + 1, j)$. For the second hop, the first and second networks have capacities $(\alpha_1 \psi(i, j_1), \dots, \alpha_k \psi(i, j_k))$ and $(\alpha_1 \mathcal{P}_{fs}(i, j_1), \dots, \alpha_k \mathcal{P}_{fs}(i, j_k))$ respectively. Let C_1^*, C_2^* be the respective C^* 's. Now, we can use Theorem 7 to argue that $\psi(i + 1, j) \geq f(C_1^*)$. Using the induction hypothesis, we know that $\alpha_t \psi(i, j_t) \geq \alpha_t f^{i-1}(\mathcal{P}_{fs}(i, j_k))$. It can be shown under this condition (using Lemma 9) that $C_1^* \geq f^{i-1}(C_2^*)$, which implies that $\psi(i + 1, j) \geq f(f^{i-1}(C_2^*)) = f^i(C_2^*) \geq f^i(\mathcal{P}_{fs}(i, j))$, thus completing the induction hypothesis. \square

Lemma 9. Following the notation used in Section 5, consider a two-hop relay network with parameters $c(Z)$ for the first hop and $r_i : i \in [m]$ as the rates for the second hop. Let the min-cut of this network be C_1^* . Replace the second hop rates by $\alpha_i g(r_i/\alpha_i)$ where $g(x) = f^n(x)$ for some n with $f(x) = 1 - e^{-x}$. The new min-cut $C_2^* \geq g(C_1^*)$.

Proof. First note that for each $\alpha > 0$, for the given g , we have $\alpha_i g(r_i/\alpha) \geq g(r_i)$, which can again be shown using induction on the exponent of f corresponding to the given g . Therefore, it is sufficient to show that the network with second hop capacities, $g(r_i)$, has a min-cut, $C \geq g(C_1^*)$. Let A^* be the set defining the min-cut for the network defining C_1^* in the same sense as it was used in proof of Theorem 7. It can further be shown that $g(x_1 + \dots + x_n) \leq g(x_1) + \dots + g(x_n)$ by using induction on the exponent of f in the representation of g . Using this fact (and denoting $\Theta = 1 - \prod_{i \in [m]/A^*} (1 - p(S, i))$ for convenience below), we then have: $C \geq \Theta + \sum_{i \in A^*} g(r_i) \geq g(\Theta) + \sum_{i \in A^*} g(r_i) \geq g(\Theta + \sum_{i \in A^*} r_i) = g(C_1^*)$. \square

10.1. Simulation results

In this section, we ran a simulation of the flooding policy and the backpressure policy on a layered network with ten layers in total (see Fig. 5). At each layer, there are ten individual nodes. When a node in layer i broadcasts, each of the nodes in layer $i + 1$ receives the packet with an acceptance probability p , and these events are assumed to be independent across the nodes. In the case of a flooding policy, every successfully received packet is accepted into the queue of any node, without regard to whether a duplicate of the same packet exists elsewhere in the network. In contrast, the backpressure policy ensures that at most one copy of any given packet is present somewhere in the network, by exchanging queue length information to elect the node with the smallest queue

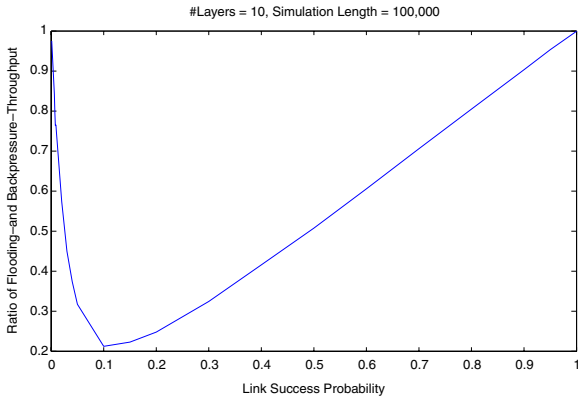


Fig. 5. The ratio of flooding to backpressure throughputs on a layered network example with 10 layers, each with 10 nodes as a function of the link success probability, which is set to be identical across all node pairs.

as the unique forwarding node for any given packet after its broadcast event. For convenience, we simulate a network with an identical acceptance probability at all links and plot the ratio of the resultant throughput of the flooding scheme with the throughput of the backpressure policy, as a function of the acceptance probability p . From the plot, it is seen that this ratio converges to one as the $p \rightarrow 0$. While this simulation is specific to the network in question, the theory implies that a similar qualitative relation is expected to hold irrespective of the number of nodes in each layer and variable acceptance probabilities, subject to the constraint that all loss events are independent. In the next section, we study the effects of relaxing the independence constraint. Remarkably as we will see in Section 11, even in a two-hop network, the ratio of throughput under feedback constraints to that of the unconstrained capacity could be arbitrarily bad, unlike the independent losses case.

A notable observation from the above plot is that the performance of a flooding based policy approaches that of the backpressure throughput (which is equal to the min-cut capacity) as the link success probability goes to zero.

11. Dependent losses

From the previous section, one wonders if we are always assured of a constant factor reduction capacity in general, without either feedback or network coding on at least bounded diameter networks. Naively speaking, this seems plausible because correlations in the losses can only supply the relays with implicit information that might help them synchronize their routing more efficiently while satisfying FIR. This intuition is flawed, as the following counterexample shows that $C_R \rightarrow 0$ even though the capacity $C^* = 1$. This counterexample can be practically motivated as an extreme case of a situation where most of the time (with probability p to be explicated later) the broadcast channel is bad enough to be actually useful to a very small set of relays but occasionally experiences high strength, in which case most of the relays receive the broadcast. We make use of the characterization of C_R from Theorem 5 in showing this fact.

Definition 10 (Network, \mathcal{N} of Size m). Let $\epsilon > 0 = \frac{1}{\sqrt{m}}$. When the source broadcasts, with probability ϵ , all relays receive the packet and with probability $1 - \epsilon$, exactly one of the relays receives the packet.

$$c(Z) = \begin{cases} \frac{1 - \epsilon}{m} & \text{if } |Z| = 1 \\ \epsilon & \text{if } Z = [m] \\ 0 & \text{otherwise} \end{cases}$$

and for the relay to destination, we have:

$$p(i, D) = 1/m \quad \forall i \in [m].$$

Theorem 10. For the network \mathcal{N} , $C^* = 1$, but $\lim_{m \rightarrow \infty} C_R = 0$.

Proof. By computing the min-cut, we have $C^* = 1$. We shall now derive an upper bound on the C_R which tends to 0 as $m \rightarrow \infty$ using Theorem 5. This implies the claim because of Theorem 5. Since $c(Z)$ depends only on $|Z|$, and the LP (12) is symmetric over i , one can set without loss of generality, $t(Z) = \phi(|Z|)$. This implies:

$$\begin{aligned} & \sum_{Z, Z' \subseteq [m]: Z \cap Z' \neq \emptyset} t(Z)c(Z') \\ &= \sum_{k=1}^m \binom{m}{k} \phi(k) \left(\sum_{Z': Z' \cap [k] \neq \emptyset} c(Z') \right) \\ &= \sum_{k=1}^m \phi(k) \binom{m}{k} \left(\epsilon + (1 - \epsilon) \frac{k}{m} \right). \end{aligned}$$

Therefore, the constraint in Eq. (13) becomes:

$$\sum_{k=1}^m \binom{m}{k} \phi(k) \leq 1. \quad (18)$$

As for the constraint in Eq. (14), we have:

$$\begin{aligned} p(S, i) &= \sum_{Z \subseteq [m]: i \in Z} c(Z) = \frac{1 - \epsilon}{m} + \epsilon \sum_{Z \subseteq [m]: i \in Z} t(Z) \\ &= \sum_{k=1}^m \binom{m-1}{k-1} \phi(k) = \frac{1}{m} \sum_{k=1}^m \binom{m}{k} k \phi(k). \end{aligned}$$

Thus the second constraint becomes:

$$\sum_{k=1}^m \binom{m}{k} k \phi(k) \leq \frac{m}{1 - \epsilon + m\epsilon}. \quad (19)$$

Thus,

$$\begin{aligned} C_T &= \sum_{k=1}^m \phi(k) \binom{m}{k} \left(\epsilon + (1 - \epsilon) \frac{k}{m} \right) \\ &= \epsilon \left(\sum_{k=1}^m \binom{m}{k} \phi(k) \right) + \frac{1 - \epsilon}{m} \left(\sum_{k=1}^m \binom{m}{k} k \phi(k) \right) \\ &\leq \epsilon + \frac{1 - \epsilon}{1 - \epsilon + m\epsilon} \quad (\text{from Eqs. (18) and (19)}) \\ &= o(1/\sqrt{m}) \quad \left(\text{since } \epsilon = \frac{1}{\sqrt{m}} \right). \quad \square \end{aligned}$$

12. Conclusion

While network coding is necessary to achieve the maximum throughput for multicast connections, this is not the case with wireless unicast. Rather, network coding is a convenient way to solve the distributed routing problem without having to depend on feedback signaling to make complicated routing choices for achieving the maximum throughput. In this context, we analyzed a relay network and quantitatively characterized the limitations of ‘static’ routing policies that operate in a feedback independent manner. Our characterization allows for explicitly identifying situations when there is no loss of throughput by restricting to such simple routing policies. Further, we show that the reduction in the throughput is controlled when the link losses for a given transmission are independent, and could even be minimal when the capacity is low, on a general feed-forward network. At worst, this is 63% and gets progressively close to 100%, as the capacity itself goes to 0. Thus, in such a situation, network coding delivers no benefit over simpler blind routing policies in the limit of unreliable communication. Nevertheless, highly correlated losses could lead to an unbounded loss even on a 2-hop network. In such a situation, network coding might be unavoidable if we need to be conservative with the feedback.

References

- [1] E. Arikian, Some complexity results about packet radio networks, *IEEE Transactions on Information Theory* 30 (1984) 681–685.
- [2] P. Gupta, P.R. Kumar, The capacity of wireless networks, *IEEE Transactions on Information Theory* 46 (2000) 388–404.
- [3] R. Gummadi, K. Jung, D. Shah, R. Sreenivas, Feasible rate allocation in wireless networks, in: *Proc. IEEE International Conference on Computer Communications*, 2008, pp. 995–1003.
- [4] R. Gummadi, K. Jung, D. Shah, R. Sreenivas, Computing the capacity region of a wireless network, in: *Proc. IEEE International Conference on Computer Communications*, 2009, pp. 1341–1349.
- [5] R. Yeung, *Information Theory and Network Coding*, Springer, 2008.
- [6] R. Ahlswede, N. Cai, S. Li, R. Yeung, Network information flow, *IEEE Transactions on Information Theory* 46 (2000) 1204–1216.
- [7] S. Chachulski, M. Jennings, S. Katti, D. Katabi, Trading structure for randomness in wireless opportunistic routing, in: *Proc. ACM Special Interest Group Conference on Communications*, SIGCOMM, 2007.
- [8] B. Smith, B. Hassibi, Wireless erasure networks with feedback, in: *Proc. IEEE International Symposium on Information Theory*, 2008.
- [9] D. Lun, M. Medard, R. Koetter, M. Effros, On coding for reliable communication over packet networks, *Physical Communication* 1 (2008) 3–20.
- [10] A. Dana, R. Gowaikar, R. Palanki, M. Effros, Capacity of wireless erasure networks, *IEEE Transactions on Information Theory* 52 (2006) 789–804.
- [11] S. Biswas, R. Morris, Opportunistic routing in multi-hop wireless networks, in: *Proc. ACM Special Interest Group Conference on Communications*, SIGCOMM, 2005.
- [12] M. Neely, R. Urgaonkar, Optimal Backpressure Routing for Wireless Networks with Multi-Receiver Diversity AD HOC Networks, vol. 7, Elsevier, 2009, pp. 862–881.
- [13] J. Sundararajan, D. Shah, M. Medard, Feedback-based online network coding April 2009. URL <http://arxiv.org/abs/0904.1730>.
- [14] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, J. Crowcroft, XORs in the air: practical wireless network coding, *IEEE/ACM Transactions on Networking* 16 (3) (2008) 497–510.
- [15] L. Tassiulas, A. Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks, *IEEE Transactions on Automatic Control* 37 (1992) 1936–1948.



Ramakrishna Gummadi is a Post-Doctoral researcher in the Department of Management Science at Stanford University. He received a Ph.D. in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in December 2011 with a dissertation on coding and scheduling in networks. He was the recipient of a Vodafone fellowship in 2007 and the Mavis Fellowship award in 2009 from the College of Engineering at the University of Illinois for academic excellence. During his graduate study, he also completed internships at the Technicolor research lab in Paris, Microsoft Research Cambridge and at Bell Labs Murray Hill, NJ.



Laurent Massoulié is a Distinguished Scientist and Fellow at Technicolor. He is a member of the Technicolor Research and Innovation Paris Lab. He graduated from Ecole Polytechnique, France in 1991, obtained his Ph.D. from Université Paris Sud in 1995 and his “HDR” from Université Paris 7 in 2010. He has held positions at France Telecom R&D, Issy-les-Moulineaux France (1995–1999) and at Microsoft Research, Cambridge United Kingdom (1999–2006). He has co-authored the Best Paper Award-winning papers of IEEE INFOCOM’99, ACM SIGMETRICS 05 and ACM Conext’07. He has served as associate editor of *Queueing Systems: Theory and Applications* (2000–2006), *IEEE/ACM ToN* (2008) and *Stochastic Systems Journal* (2011-).



Sreenivas received his B.Tech. Degree in electrical engineering from the Indian Institute of Technology, Madras, India, in 1985, and his M.S. and Ph.D. Degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1987 and 1990, respectively. He was a Post-Doctoral Fellow in Decision and Control at the Division of Applied Sciences, Harvard University, Cambridge, MA, before he joined the University of Illinois at Urbana-Champaign in September 1992. He is currently an Associate Professor of Industrial and Enterprise Systems Engineering, Research Associate Professor with the Coordinated Science Laboratory, Information and Trust Institute, and the Illinois Center for Transportation.