The Role of Feedback in the choice between Routing and Coding for Wireless Unicast

Ramakrishna Gummadi
CSL and ECE,
University of Illinois, Urbana, USA
gummadi2@illinois.edu

Laurent Massoulie
Thomson Corporate Research,
Paris, France
laurent.massoulie@thomson.net

Ramavarapu Sreenivas
CSL and IESE,
University of Illinois, Urbana, USA
rsree@illinois.edu

***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 provide a natural argument 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 in such a result, and is more explicit than previous proofs. This maximum throughput can be achieved either by a flow scheduler with knowledge of the network topology, or by a blind back-pressure algorithm, however all such policies require dynamic routing decisions which depend critically on rich feedback signalling information. We then investigate how feedback at a single node affects its throughput to a destination, with fixed rates of its one hop neighbors to the destination. We first analyze *static* routing policies which are in a sense, *feedback independent*. Under such a constraint, we obtain an explicit characterization of the optimal policy, which turns out to be a natural generalization of both flooding and traditional routing. For independent losses at the receivers (still possibly allowing for dependencies for transmitted by two different nodes to model interference), the reduction is limited to a constant fraction of the capacity, and could even be arbitrarily close to optimal depending on the network. However, if there are dependencies in the losses seen by receivers from a single broadcast, the reduction could be arbitrarily bad, even on a 2-hop network. Our analysis provides several new insights on the scenarios in which coding in the network can (and can not) be avoided and how it relates to the feedback.**

I. Introduction

Network coding was originally introduced in [1] as a solution 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 "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 it has been considered in the context of wireless networks [3] for unicast communications, where such a wireless setting was identified as an especially good candidate for network coding because of the local broadcast. The wireless scenario of [3] features lossy transmissions (a drawback of wireless) as well as local broadcast (an advantage of wireless). The experimental evaluations of [3] show benefits of network coding over routing policies in terms of the transmission rates achieved. This leaves something to be explained because of its contrast with recent work, [6] and [7], which imply that routing policies can achieve the capacity for unicast communications, in a large class of wireless network models. This begs the following 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, our analysis shows that the benefits of coding over routing are dependent on the richness of the feedback signalling and the network in question, and quantitatively argues about their critical relation. Our first contribution is a natural argument that highlights the role of LP duality in understanding why the maximum routing throughput in the case of wireless unicast is the same as the capacity achieved with coding. As we will see, this implies that backpressure routing achieves not only the optimal throughput among routing policies, but also more generally across policies that involve coding. While being conceptually elegant and distributed, backpressure policy performs a highly 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 limitations faced by *static* routing policies which do not need rich feedback signalling. Our analysis brings forth several new insights on this tradeoff.

To begin, let us discuss an extreme case (see Fig I): consider m distributed agents with one packet each from a set of m distinct packets (which all of them have received from the source) to forward to the destination. With full coordination, of course, they can together cover all the m packets in one transmission each. However, if we restrict their choice to be made in a distributed manner without coordination (note the analogy with the coupon collectors problem), 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. Note that 100% throughput could be achieved by either (i) making coordinated choices among relay nodes through conferencing among all the relays, or (ii) by 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 signalling.

The above synchronization constraint is modeled as the *Feedback Independent Routing* (FIR) restriction (see Defini-

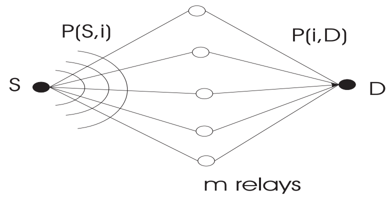

Fig. 1. The relay network example

tion 4.1): informally, this says that the decision of whether a node chooses to forward a packet it received does not depend on loss events to other nodes. In all examples of policies that achieve the capacity without any coding, this information is implicitly exploited via extensive feedback signalling among nodes. 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 before the broadcast). When transmissions are subject to mutually independent losses at the receivers (but still possibly allowing for dependencies 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. Lest we get too optimistic about the previous achievability argument (and imagine that dependencies in link losses supply implicit information about other link losses, and thus should allow to improve over than independent losses), we show a rather surprising counter example 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.

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 [3]–[5]. 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 signalling: This was the path taken in [2], [6], [7]. The maximum throughput in an information theoretic sense for the Wireless Erasure Network (WEN) model was studied in [4] and was shown to be equal to a generalized min cut (GMC). In [7], 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 GMC. Using this, [7] makes an important observation that coding is in fact not necessary for achieving a throughput equal to the capacity in wireless unicast settings. In [6], backpressure routing (requiring rich feedback signalling to perform routing) was shown to be an optimal routing cum scheduling policy in a much more general context taking into account interference effects. In general (i.e. for multicast), [8] proposes the use of feedback together with network coding for online decoding.

II. The Wireless Erasure Network (WEN) Model

Let $G = (V, E)$ be a directed graph. Let $\mathcal{N}(i) = \{j \in V : (i,j) \in E\}$. We consider an information network that operates on this graph over time $t \in \{0, 1, 2, \ldots\}$. 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 $i \in V, Z \subseteq N(i)$, let $\chi(i, Z, t)$ denote a $\{0, 1\}$ random variable that represents the following:

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

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 stationary across t (not even necessarily independent). Although, in real deployments of wireless networks, the characteristics of the channels might not quite satisfy the stationarity assumption across t, the model we adopt provides an excellent compromise between its tractability towards obtaining any non trivial insights, and its faithfulness to reality. Nevertheless, there exists at least one real situation for which our model is exact: the random access scheduling model, 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: $\sum_{Z \subseteq \mathcal{N}(i)} \chi(i, Z, t) = 1 \quad \forall i, t$. We define: $c(i, Z) = E[\chi(i, Z, t)] \quad \forall t$. 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. Then, we would have:

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

We will 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 2.1 (A General routing policy, $\mathcal{P}$): A general routing Policy $\mathcal{P}$ decides for each node $i \in V$, and time $t \in \{0, 1, \ldots\}$, a packet to be transmitted from among the $\alpha_i(t-1)$ choices in its possession.

Definition 2.2 (Capacity $C(\mathcal{P})$): 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 then, $C^* = \sup_{\mathcal{P}} C(\mathcal{P})$

III. 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 $\sum_{\{Z \in \mathcal{N}(i): j \in Z\}} r(i, j, Z)$ represents the net flow from i to j. The value of the LP then represents the throughput that is achieved.

Definition 3.1 (MFC, the 'Max Flow Capacity'): Let P denote the set of all $S - D$ paths in G. We define the MFC for a given broadcast capacitated graph as the optimum value of the following LP:

$$MFC = \max \sum_{p \in P} x_p \qquad (1)$$

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

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

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

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

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 3.2 ($\mathcal{P}_{fs}$, 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 define a notion of min cut that is generalized to the wireless setting by considering the probability that *at least* one of the nodes across the cut receives a transmission.

Definition 3.3 (General Min Cut, GMC): 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:

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

The General Min Cut, is then: $GMC = \min_A C(A)$.

It is easy to argue that GMC 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 [4]. Thus, $C^* \leq GMC$. Based on a duality argument analogous to the classical max flow min cut theorem, we make the following claim, the proof of which is provided in the appendix:

Theorem 3.4: $GMC = MFC = C^*$

Thus, coding in the network is not necessary to achieve the optimal throughput in unicast, assuming unconstrained feedback signalling. This fact was also argued by [7] in a closely related continuous time 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 the GMC using Lyapunov stability argument. The below corollary 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, GMC. This is, in a sense, implied by theorem 3.4 in conjunction with Neely's result of optimality of back-pressure routing schemes in a context that also involves power control, with a different interference model ([6]). We give an outline of our proof in the appendix.

Corollary 3.5: 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 $\lambda < GMC$, the Markov Chain thus obtained is stable (positive recurrent).

IV. Routing and 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 (Figure I).

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 the 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 may assume that the source attempts distinct broadcast of distinct packets at each instance without any loss of generality (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 4.1 (Feedback Independent Routing (FIR)): $\forall p \in \{0,1,\ldots\}, \forall A, B \subseteq [m]$ such that $A \cap B = \phi$, the given routing policy satisfies the FIR restriction if, conditioned on $\{r_i(p) = 1 \quad \forall \quad i \in A\}$, 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 achieveing scheme for a single erasure channel, i.e. either by (1) feedback to the relays, or (2) Forward Erasure Coding (FEC) at each relay. [1]

We have thus far discussed 3 policies which achieve capacity without any coding: (1) the flow splitting policy, $\mathcal{P}_{fs}$ defined in 3.2, (2) the backpressure policy in lemma 3.5 and (3) the policy of [7] but none of them satisfies the above constraint and involve a heavy interaction between feedback and routing, unlike network coding.

[1] 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.

A. Characterization of the Optimal Policies

Definition 4.2: The Capacity with Feedback Independent Routing, C_{FIR} is defined as the maximum throughput as in def 2.1, restricted to policies $\mathcal{P}$ satisfying def 4.1.

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 received 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 below LP, 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 4.3 (Tag Capacity, TC): TC is defined as the optimum value of the following LP with variables $t(Z) \geq 0$ where $Z \subseteq [m]$.

$$TC = \max \sum_{Z, Z' \subseteq [m]: Z \cap Z' \neq \phi} t(Z) c(Z') \tag{2}$$

$$\text{Subject to: } \sum_{Z \subseteq [m]} t(Z) \leq 1; \tag{3}$$

$$p(S,i) \left(\sum_{Z \subseteq [m]: i \in Z} t(Z) \right) < p(i,D) \qquad \forall i \in [m] \tag{4}$$

All packets that reach a relay successfully *and* have the relay in the tag will eventually reach D because of the constraint in Eq (4) (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 2.1.

Lemma 4.4: $C_{FIR} \geq TC$

Proof: Omitted. The argument involves observing that any feasible solution to the LP gives us a tagging scheme whose throughput is equal to the value of the LP. ∎

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_{FIR} in the form of the below theorem, whose proof is provided after a brief discussion.

Theorem 4.5: $C_{FIR} \leq TC$, implying $TC = C_{FIR}$

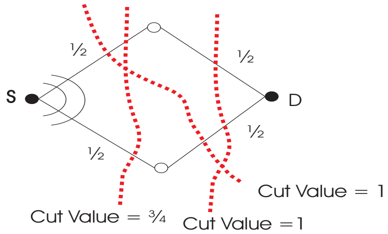

(a) Example 1

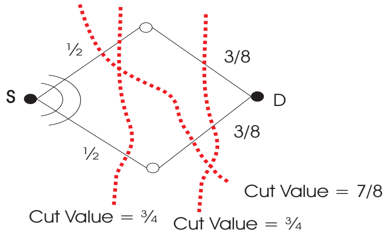

(b) Example 2

Fig. 2. Examples

In general, it is not obvious if C_{FIR} is strictly less than the general min cut capacity, GMC. Indeed, in many cases, these two quantities match, implying that in such cases, not only do we not need any coding, but the capaciy 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 the Fig. 2(a), $GMC = 3/4$ and from the LP (Eq (2)), we can calculte that $C_{FIR} = 3/4$. 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, $GMC = 3/4$, as explained in Fig 2(b). Set $t(\{1\}) = x, t(\{2\}) = y, t(\{1,2\}) = z$. Then: $C_{FIR} = Max \quad (\frac{1}{2}x + \frac{1}{2}y + \frac{3}{4}z)$ Subject to: $x, y, z \geq 0, \quad x + y + z \leq 1, \quad x + z \leq 3/4, \quad y + z \leq 3/4$. The optimal value can be verified to be $5/8 < 3/4 = GMC$. 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.

Proof of Theorem 4.5: Consider any arbitrary policy. We will show that the expected number of distinct packets that reach the destination in k time slots can not be more than $TC.k$ as long as FIR in definition 4.1 is satisfied, thus implying that the C_{FIR} is at most TC. Define:

$$t(Z) = \frac{1}{k}\sum_{p=1}^{k} P\left(\frac{r_i^*(p) = 1 \ \forall \ i \in Z \text{ and } 0 \quad \forall i \in [m]/Z}{r_i(p) = 1 \quad \forall \ p \in [m]}\right) \tag{5}$$

where we use the standard convention of writing $\frac{P(A \cap B)}{P(B)}$ as $P(\frac{A}{B})$ for an event B with positive probability. It is easily verified that $t(Z) \geq 0$ and $\sum_{Z \subseteq [m]} t(Z) = 1$. We will now verify that the third constraint (Eq (4)) holds for any $i \in [m]$.

$$kp(i,D) \geq E[\sum_{p=1}^{k} \mathbb{1}\{r_i^*(p) = 1\}]$$

$$= \sum_{p=1}^{k} P(r_i(p) = 1) P\left(\frac{r_i^*(p) = 1}{r_i(p) = 1}\right)$$

$$= p(S,i) \sum_{p=1}^{k} P\left(\frac{r_i^*(p) = 1}{r_i(p) = 1}\right)$$

$$= p(S,i) \sum_{p=1}^{k} P\left(\frac{r_i^*(p) = 1}{r_j(p) = 1 \ \forall j \in [m]}\right) \text{ (using def.4.1)}$$

$$= p(S,i) \sum_{p=1}^{k} \sum_{Z \ni i} P\left(\frac{r_j^*(p) = 1 \ \forall \ j \in Z \text{ and } 0 \ \forall j \in [m]/Z}{r_j(p) = 1 \ \forall j \in [m]}\right)$$

$$= kp(S,i) \sum_{Z \ni i} t(Z) \text{ (by definition of } t(Z) \text{ in Eq (5))}$$

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

$$E[\alpha_D(k)] = \sum_{p=1}^{k} P(\{\bigcup_{i \in [m]} \{r_i^*(p) = 1\}\})$$

$$= \sum_{p=1}^{k} \sum_{Z \subseteq [m]} P\Big(r_i(p) = 1 \ \forall \ i \in Z \text{ and } 0 \quad \forall i \in [m]/Z\Big) \times$$

$$P\Big(\frac{r_i^*(p) = 1 \text{ for some } i \in [m]}{r_i(p) = 1 \ \forall \ i \in Z \text{ and } 0 \quad \forall i \in [m]/Z}\Big)$$

$$= \sum_{p=1}^{k} \sum_{Z \subseteq [m]} c(Z) P\left(\frac{r_i^*(p) = 1 \text{ for some } i \in Z}{r_i(p) = 1 \ \forall \ i \in Z \text{ and } 0 \quad \forall i \in [m]/Z}\right)$$

$$(\because \forall i \in [m]/Z, \{r_i(p) = 0\} \Rightarrow \{r_i^*(p) = 0\})$$

$$= \sum_{p=1}^{k} \sum_{Z \subseteq [m]} c(Z) P\left(\frac{r_i^*(p) = 1 \text{ for some } i \in Z}{r_i(p) = 1 \ \forall \ i \in [m]}\right)$$

(using FIR definition 4.1)

$$= \sum_{p=1}^{k} \sum_{Z} c(Z) \sum_{Z' \cap Z \neq \phi} P\left(\frac{r_i^*(p) = 1 \ \forall i \in Z' \text{ and } 0 \text{ else}}{r_i(p) = 1 \ \forall \ i \in [m]}\right)$$

$$= k \sum_{Z \subseteq [m]} c(Z) \sum_{Z' \subseteq [m] : Z' \cap Z \neq \phi} t(Z)$$

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

■

B. Analysis of the Loss of throughput

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 4.6 (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 claim:

Lemma 4.7:

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

Proof: Omitted. The argument relies on calculating the probability that any given packet reaches the destination via *at least* one of the relays. ■

Using this flooding policy, we now prove a lower bound on C_{FIR} 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 4.8: Consider the relay network of figure I with independent losses from the source to the relays.

$$C_{FIR} \geq 1 - e^{-C^*}$$

Thus, $C_{FIR}/C^* \geq 1 - e^{-1}$. It follows from this bound that, as $C^* \to 0$, $C_{FIR}/C^* \to 1$

Proof: We will specifically show that $C(\mathcal{P}_F) \geq 1 - e^{-GMC}$, which in turn implies the theorem since $\mathcal{P}_F$ clearly satisfies FIR (definition 4.1), and $GMC = C^*$.

Recall that:

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

Applying the definition of GMC to the network under consideration, we see that

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

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

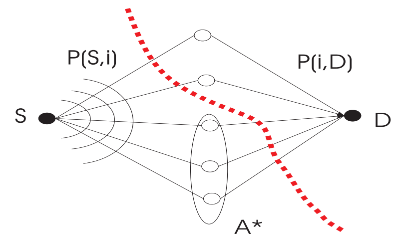

Fig. 3. An Illustration of a general min cut and the set defining A^* in this min cut.

$$GMC = 1 - \prod_{i\in[m]/A^*} (1-p(S,i)) + \sum_{i\in A^*} p(i,D) \quad (6)$$

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

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

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

which implies:

$$p(i,D) \leq p(S,i) \prod_{j\in[m]/A^*} (1-p(S,j)) \leq p(S,i) \quad (7)$$

Thus:

$$C(\mathcal{P}_F) = 1 - \prod_{i\in[m]} (1 - \min(p(S,i), p(i,D)))$$
$$\geq 1 - \prod_{i\in A^*} (1 - \min(p(S,i), p(i,D))$$
$$= 1 - \prod_{i\in A^*} (1 - p(i,D))$$
$$(\because\ p(i,D) \leq p(S,i)\ \ \forall\ i \in A^* \text{ from Equation 7})$$
$$\geq 1 - \prod_{i\in A^*} e^{-p(i,D)} \quad (\because\ 1 - x \leq e^{-x}\ \ \forall\ x \geq 0)$$
$$= 1 - e^{-\sum_{i\in A^*} p(i,D)}$$
$$\geq 1 - e^{-GMC} \quad (\because \sum_{i\in A^*} p(i,D) \leq GMC \text{ from Equation 6})$$

■

The above 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. We next build upon this argument to obtain a bound for arbitrary feedforward networks where *every* node is restricted to static routing and without any assumption of achievable rates from the one hop neighbors.

Theorem 4.9: For a general feedforward network with $h+1$ hops, it is possible to achieve a throughput of at least $f^h(C^*)$ where $f(x) = 1 - e^{-x}$ and C^* is the min cut capacity.

Proof Outline: First, we note that feedforward networks can be reduced to general layered networks where only nodes at subsequent levels are connected, but 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.

We define, ψ, a generalized flooding policy satisfying FIR and compare it with the splitting policy, $\mathcal{P}_{fs}$ in Definition 3.2: 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 3.2. 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 $(i, j_1), \ldots, (i, j_k)$. Consider two different 2-hop networks both with (i+1,j) as the source and $(i, j_1), \ldots, (i, j_k)$ as relays with the same link characteristics as the feedforward 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 $(\alpha_1 \psi(i, j_1), \ldots, \alpha_k \psi(i, j_k))$ and $(\alpha_1 \mathcal{P}_{fs}(i, j_1), \ldots, \alpha_k \mathcal{P}_{fs}(i, j_k))$ respectively. Let C_1^*, C_2^* be the respective GMC's. Now, we can use theorem 4.8 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 6.2 of the appendix) 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. ■

From the above result, 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 below counter example shows that $C_{FIR} \to 0$ even though the capacity $C^* = 1$. This counter example could also have practical relevance since it can be 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_{FIR} from theorem 4.5 in showing this fact.

Definition 4.10 ("Confusion Network" of size m): When the source broadcasts, with probability p, exactly one of the relays receive the packet, and with probability $1-p$, all relays receive the packet, where $\Omega(\frac{1}{m}) \leq 1-p \leq o(1)$. More explicitly, we can choose $p = 1 - \frac{1}{\sqrt{m}}$ and then have:

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

and for the relay to destination, we have:

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

Theorem 4.11: For the "Confusion network", $C^* = 1$, but $\lim_{m \to \infty} C_{FIR} = 0$

Proof: It is easy to verify that $GMC = 1$. We shall now derive an upper bound on the optimum value of the LP in definition 4.3 which tends to 0 as $m \to \infty$. This implies the claim because of theorem 4.5. Since $c(Z)$ depends only on $|Z|$, and the LP 2 is symmetric over i, one can set without loss of generality, $t(Z) = \phi(|Z|)$, in which case: $\sum_{Z, Z' \subseteq [m]: Z \cap Z' \neq \phi} t(Z) c(Z') = \sum_{k=1}^m \binom{m}{k} \phi(k) (\sum_{Z': Z' \cap [k] \neq \phi} c(Z')) = \sum_{k=1}^m \phi(k) \binom{m}{k} (1 - p + p\frac{k}{m})$

The constraint in Eq (3) becomes:

$$\sum_{k=1}^m \binom{m}{k} \phi(k) \leq 1 \tag{8}$$

As for the constraint in Eq (4), we have:

$$p(S, i) = \sum_{Z \subseteq [m]: i \in Z} c(Z) = \frac{p}{m} + 1 - p$$

and

$$\sum_{Z \subseteq [m]: i \in Z} t(Z) = \sum_{k=1}^m \binom{m-1}{k-1} \phi(k) = \frac{1}{m} \sum_{k=1}^m \binom{m}{k} k\phi(k)$$

Thus the second constraint becomes:

$$\sum_{k=1}^m \binom{m}{k} k\phi(k) \leq \frac{m}{p + m(1-p)} \tag{9}$$

Thus,

$$\begin{aligned}
TC &= \sum_{k=1}^{m} \phi(k) \binom{m}{k} (1-p+p\frac{k}{m}) \\
&= (1-p)\left(\sum_{k=1}^{m}\binom{m}{k}\phi(k)\right) + \frac{p}{m}\left(\sum_{k=1}^{m}\binom{m}{k}k\phi(k)\right) \\
&\leq 1-p+\frac{p}{p+m(1-p)} \quad (\text{ from Eqs } (8),(9)) \\
&= O(1/\sqrt{m}) \quad (\text{ since } p = 1-\frac{1}{\sqrt{m}})
\end{aligned}$$

■

V. 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 signalling 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 feedforward 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] R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network information flow. *IEEE Trans. on Information Theory*, 46: 4:1204–1216, July 2000.

[2] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. *ACM SIGCOMM*, 2005.

[3] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. *ACM SIGCOMM*, 2007.

[4] A. Dana, R. Gowaikar, R. Palanki, and M. Effros. Capacity of wireless erasure networks. *IEEE Trans. on Information Theory*, 52: 3:789–804, March 2006.

[5] D. Lun, M. Medard, R. Koetter, and M. Effros. On coding for reliable communication over packet networks. *Physical Communication*, 1:3–20, March 2008.

[6] M. Neely and R. Urgaonkar. Optimal backpressure routing for wireless networks with multi-receiver diversity. *AD HOC NETWORKS (ELSEVIER)*, 7: 5:862–881, July 2009.

[7] B. Smith and B. Hassibi. Wireless erasure networks with feedback. *On Arxiv: http://arxiv.org/pdf/0804.4298v1*, 2008.

[8] J. Sundararajan, D. Shah, and M. Medard. Feedback-based online network coding. *On Arxiv: http://arxiv.org/abs/0904.1730*, April 2009.

[9] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. on Automatic Control*, 37:1936–1948, December 1992.

VI. Appendix

Proof of Theorem 3.4: 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:

$$DUAL^* = \min \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z)b(i,Z) \tag{10}$$

Such that:

$$\sum_{\{(i,j)\in p\}} y(i,j) \geq 1 \quad \forall p \in P \tag{11}$$

$$-y(i,j) + b(i,Z) \geq 0 \ \ \forall \{(i,j,Z) : (i,j) \in E, j \in Z \subseteq \mathcal{N}(i)\}$$

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

$$b(i,Z) \geq 0 \ \ \forall (i,Z) : i \in V, Z \subseteq \mathcal{N}(i) \tag{12}$$

Consider the mincut as described in (3.3), 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 $\sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z)b^*(i,Z) = GMC$. Thus,

$$GMC \geq DUAL^* \tag{13}$$

Now consider an integral constrained version of the above dual:

$$DUAL_* = \min \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z)b(i,Z) \tag{14}$$

Such that:

$$\sum_{\{(i,j)\in p\}} y(i,j) \geq 1 \ \ \forall p \in P \tag{15}$$

$$-y(i,j) + b(i,Z) \geq 0 \ \ \forall \{(i,j,Z) : (i,j) \in E, j \in Z \subseteq \mathcal{N}(i)\} \tag{16}$$

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

$$b(i,Z) \in \{0,1\} \ \ \forall (i,Z) : i \in V, Z \subseteq \mathcal{N}(i) \tag{17}$$

Let y_*, b_* define the optimal solution to the above integral constrained LP. Then, define $A_* = \{i \in V : \exists$ a path, p, from S to i s.t. $\sum_{\{(i,j)\in p\}} y_*(i,j) = 0\}$. Then, clearly, $D \notin A_*$ because of constraint (15), and thus A_* defines an (S,D) cut. Furthermore, by interpreting (16) and (14), it can be verified that $C(A_*) = DUAL_*$. This implies that

$$GMC \leq DUAL_* \tag{18}$$

To summarize, we can so far claim:

$$MFC = DUAL^* \leq C^* \leq GMC \leq DUAL_* \tag{19}$$

If we are able to argue that the constrained LP indeed achieves the optimum (the details of this are given in Lemma 6.1), we would then have $DUAL^* = DUAL_*$, implying that all quantities in Eq (19) are the same. ■

Lemma 6.1: $DUAL^* = DUAL_*$

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 GMC thus obtained is no more than $DUAL^*$, which in turn implies that $DUAL_* \leq DUAL^*$ completing the proof. Consider the dual LP at Eq (10), 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 Eqn 11. Then:

$$E[C(A^*)] = E[\sum_{i \in A^*, Z \subseteq \mathcal{N}(i), Z \cap \bar{A}^* \neq \phi} c(i,Z)]$$

$$= E[\sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) \mathbb{1}\{i \in A^*, Z \subseteq \mathcal{N}(i), Z \cap \bar{A}^* \neq \phi\}]$$

$$= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) P\left(i \in A^*, j \in \bar{A}^* \text{ for some } j \in Z\right)$$

$$= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) P\left(d(i) \leq \lambda, d(j) > \lambda \text{ for some } j \in Z\right)$$

$$= \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) (\max_{j \in Z} d(j) - d(i))_+ \quad (\because \lambda \text{ is uniform})$$

$$\leq \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) (\max_{j \in Z} y^*(i,j)) \quad (\text{ traingle inequality })$$

$$\leq \sum_{i \in V, Z \subseteq \mathcal{N}(i)} c(i,Z) b^*(i,Z) \quad (\text{ from Eq. (12)}) = DUAL^*$$

■

Proof Outline for Corollary 3.5: The proof follows along the lines of [9] 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 $\lambda < GMC$. 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 [9] and the references therein), a sufficient condition for stability is to show that 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 $(i,l)^{th}$ 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 bdenoting 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 $\mathbf{q}(t+1) = \mathbf{q}(t) + \mathbf{R}\ \mathbf{E} + \mathbf{A}(t)$. Thus (with . denoting the usual dot product of vectors):

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

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, $\mathbf{q}(t).\mathbf{RE}(t) = -\mathbf{W}(t).\mathbf{E}(t)$. Hence, we only need to argue that $E[-\mathbf{W}(t).\mathbf{E}(t) + \mathbf{q}(t).\mathbf{A}(t)/\mathcal{F}_t]$ has a negative drift. $E[\mathbf{q}(t).\mathbf{A}(t)/\mathcal{F}_t] = \mathbf{q}(t).\mathbf{A}$ where A is a vector where the element corresponding to the source is λ and all other elements 0. Since $\lambda < GMC$, it follows from the Max flow interpretation of theorem 3.4 that there exist variables $r^*(i,j,Z)$ satisfying the constraints of LP at Eq (1), and an $\epsilon > 0$ such that $\mathbf{q}(t).\mathbf{A} \leq (1-\epsilon)\mathbf{W}(t).\mathbf{f}^*$, 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, $\mathcal{P}_{fs}$ given in Definition 3.2. Then, we claim that $\mathbf{W}(t).\mathbf{f}^* = E[\mathbf{W}(t).\mathbf{e}^*(t)/\mathcal{F}_t]$. Thus, $E[-\mathbf{W}(t).\mathbf{E}(t) + \mathbf{q}(t).\mathbf{A}(t)/\mathcal{F}_t] = E[-\mathbf{W}(t).\mathbf{E}(t) + (1-\epsilon)\mathbf{W}(t).\mathbf{e}^*(t)/\mathcal{F}_t] = E[-\mathbf{W}(t).\mathbf{E}(t) + \mathbf{W}(t).\mathbf{e}^*(t)/\mathcal{F}_{t^+}/\mathcal{F}_t] - \epsilon E[\mathbf{W}(t).\mathbf{e}^*(t)/\mathcal{F}_t]$. 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 $\mathbf{W}(t).\mathbf{E}(t)$ 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. ■

Lemma 6.2: Following the notation used in Section IV, consider a 2 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 we 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 mincut, $C \geq g(C_1^*)$. 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 4.8. It can further be shown that $g(x_1 + \ldots + x_n) \leq g(x_1) + \ldots g(x_n)$ 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: $C \geq \Theta + \sum_{i \in A^*} g(r_i) \geq g(\Theta) + \sum_{i \in A^*} g(r_i) \geq g(\Theta + \sum_{i \in A^*} r_i) = g(C_1^*)$ ■