Understanding the Benefits of Opportunism and Cooperation in Multihop Wireless Networks

Anand Seetharam†, Chi-Kin Chau§, Jim Kurose†, Don Towsley†
†Department of Computer Science, University of Massachusetts, Amherst, MA 01003, USA
§Computer Laboratory, University of Cambridge, UK
Email: {anand, kurose, towsley}@cs.umass.edu, Chi-Kin.Chau@cl.cam.ac.uk

***Abstract*—Opportunism and cooperation have emerged as powerful techniques for increasing throughput in a wireless network by exploiting the broadcast nature of the medium. In this paper we use simple Markov Chain models, augmented by simulation, to investigate the benefits of using opportunism and cooperation rather than traditional hop-by-hop, fixed-path routing in a multi-hop wireless mesh network. We consider several different topologies, including a linear network and a diamond network with a group of relays between source and destination. In the latter case, we consider both a single source-destination pair, as well as two competing source-destination pairs. We quantify the throughput gains and observe that both opportunism and cooperation can provide a significant improvement in throughput over the hop-by-hop case for these simple topologies.**

I. Introduction

Until recently, multihop wireless networks have traditionally used fixed-path, hop-by-hop forwarding techniques conceptually similar to those used in their wired counterparts. However, opportunism [1], [9], [10], [6], [12], [4], [11], [2] and cooperation [8], [7], [3] have recently emerged as powerful techniques for increasing throughput in a wireless network by exploiting the broadcast nature of the medium. Informally, in an opportunistic approach, several nodes may receive a copy of a packet transmission, with that node that is best able to forward the packet further towards its destination taking responsibility for the next transmission of that packet. In cooperative approaches, multiple nodes may receive a copy of the packet and will then concurrently transmit the packet, effectively "combining" their signals to enhance downstream reception. Intuitively, one expects that cooperation and opportunism will result in throughput gains when compared to traditional routing protocols.

In this paper we use simple Markov chain models and analysis, augmented by simulation, to investigate the performance gains obtained by opportunism and cooperation. We consider several different topologies, including a linear network and a diamond network with a group of relays between source and destination. In the latter case, we consider both a single source-destination pair, as well as two competing source-destination pairs that must "share" the relays, and compare opportunistic and cooperative forwarding with conventional fixed-route, hop-by-hop transmission.

[1]This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defense and was accomplished under Agreement Number W911NF-06-3-0002.l. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

Our work differs from previous efforts in that our goal is to compare idealized and representative hop-by-hop, opportunistic, and cooperative routing approaches using simple models and under common assumptions, rather than to investigate a specific opportunistic or cooperative transmission protocol in detail. We also consider the case that a packet must be transmitted several times among intervening nodes before reaching the destination, and the case in which multiple flows must share the intervening relays. Our evaluation demonstrates and discusses the extent to which throughput gains can be achieved by using opportunism and cooperation over traditional routing, for the simple scenarios studied in this paper.

Specifically, we highlight the contributions in the paper:

1) We derive analytical results to quantify the throughput of various opportunistic and cooperative forwarding approaches on linear and diamond topologies.
2) Using these results, we compare the throughput gains of these opportunistic and cooperative forwarding approaches, observing that the higher-overhead cooperative forwarding yields only marginally improved performance, with the lower-overhead optimally planned opportunistic forwarding can give nearly comparable performance.
3) Considering multiple competing flows, we show through simulation that unbridled cooperation within individual flows, without considering other competing flows, can reduce the throughput of the cooperative scheme to below that of an opportunistic approach.

The remainder of the paper is organized as follows. §II provides an overview of related work. We outline various network scenarios (network topologies and underlying assumptions) in §III. We define the fixed-route, opportunistic and cooperative policies, and our models in §IV and §V respectively. The evaluation results are presented in §VI and §VII and we conclude the paper in §VIII.

II. Related Work

The desire to opportunistically or cooperatively exploit the broadcast nature of wireless transmissions has recently spurred a large body of research on routing and forwarding in multi-hop wireless networks. The seminal idea describing opportunistic routing was proposed by Biswas and Morris [1]. Since then, a number of papers have been published that practically explored the benefits of opportunism [9], [10], [6], [12]. Research efforts have also theoretically analyzed the benefits of opportunism, including [4], where the authors performed a Markov Chain analysis to determine the expected number of network-wide link-layer transmissions (ETX) needed to transfer a packet from source to destination in a wireless mesh network. That work assumes that link success probabilities are provided *a priori* and does not consider channel fading, an important component of our models. [5] provides a recursive relation and Linear Programming formulation for estimating the minimum number of opportunistic transmissions required. Similarly [11] proposes an analytical model to study the performance (ETX) of opportunistic routing protocols. Baccelli *et al.* [2] quantify the end-to-end delay obtained by using opportunistic schemes based on geographic routing and demonstrate that it is about two times less than that obtained by typical shortest path routing. They also provide a theoretical study demonstrating how to tune the MAC parameters to minimize the number of time slots required to transmit packet from source to destination.

There is also a considerable amount of research describing the advantages of cooperation in wireless networks. [7] and [8] summarize much of this prior work in cooperative diversity and demonstrate how cooperation improves network performance. However, most past research on cooperation has been in the context of the physical layer, with only a few efforts exploring how cooperation interacts with the higher layers [7]. In [3] the authors discuss how to effectively schedule cooperative transmissions for multiple access scenarios by helping sources with poor channels to the destination using relays that have better channel quality. Our work differs from prior work in that we address primarily a network-layer concern (multi-hop routing), with the goal of comparing idealized and representative hop-by-hop, opportunistic, and cooperative routing approaches - using simple models, under common assumptions, and in a multi-hop setting - rather than investigating a specific opportunistic or cooperative transmission protocol in detail.

III. Network Model

The two network topologies we consider are a linear topology and a diamond network having multiple equidistant relays between a source and destination.

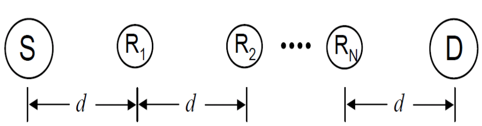

Fig. 1. Linear Network

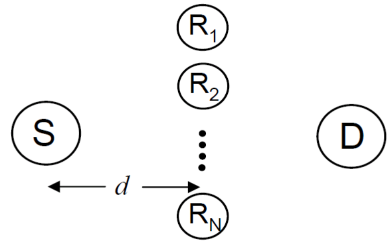

Fig. 2. Diamond Network

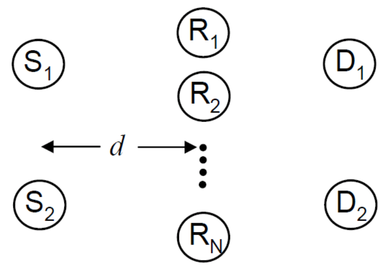

Fig. 3. Double-Diamond Network

In Figure 1 we illustrate the scenario of a single source-destination pair and N intervening relays. The distance between any two successive nodes in the linear array is identical and denoted by d. Figure 2 depicts a setting with a group of N relays between the source and destination. The distance between the source and the destination is $2d$. The distance between any two other nodes in this topology is d; while this latter assumption is an idealization, it nonetheless allows us to model the fact that channel fading occurs between relay nodes without having to condition on a specific set of distances among nodes. In Figure 3, there are 2 source-destination pairs, again with N intervening relays. The distance between any source and any destination is $2d$, while the distance between any two other nodes in this topology is again d.

We consider time to be partitioned into slots. For all topologies (Figures 1, 2 and 3) we assume that for each destination there is always only one packet in the network, although multiple copies of that packet may be present at the various network nodes at a given point in time. The goal is to transfer the packet from source to destination, using the intervening nodes in an opportunistic or cooperative manner, as discussed below. When a packet reaches the destination, remaining copies are removed from the network and the source can then transmit a new packet in the next time slot. Let N_{pack} denote the number of packets received at the destination per time slot. Our evaluation metric is throughput (T), which is defined as

$$T = E[N_{pack}] \tag{1}$$

We will model the channel as a Rayleigh fading channel, and assume that the fading coefficients are iid and exponentially distributed with mean 1. For a single source-destination pair, a transmission is considered successful if the received Signal to Noise Ratio (SNR) is above some threshold β. Let P_0 ,

N_0 and α be the transmitted power, received noise, and path loss coefficient respectively. Thus, for an individual packet transmission, the received SNR at a receiver a distance d away is exponentially distributed and is given by:

$$P[SNR > \beta] = e^{\frac{-\beta N_0 d^\alpha}{P_0}} \quad (2)$$

Let p_i denote the probability of successfully receiving a transmitted packet at a distance id. Let q_i be the corresponding loss probability. Hence $q_i = 1 - p_i$. For the topology in Figure 1 if S transmits a packet, p_1 denotes the probability that R_1 will receive it while p_2 and p_3 denote the probability that R_2 and D will be able to successfully receive it. Using (1) and considering α = 2 we have: $p_2 = p_1^4$ and $p_3 = p_1^9$. In general the probability of successful reception k-hops away is $p_1^{k^2}$.

In Figure 2, the probability of the source reaching the destination directly is p_2. For the remaining nodes the probability of reception of any transmitted packet is p_1. We will make extensive use of p_1, p_2, p_3 and q_1, q_2, q_3 throughout this paper.

For the multiple source-destination pair depicted by the double-diamond in Figure 3, we must also take interference into account, since there can be two different packets in the network at the same time. When different nodes transmit different packets, these different transmissions can interfere at the receiver. We assume that a packet will be successfully received by the receiver if the Signal to Interference and Noise Ratio ($SINR$) is above the threshold β. We investigate this scenario primarily through simulation.

IV. Policies Investigated

In this section we provide a brief overview of four different routing and forwarding policies to be compared in this paper namely, the hop-by-hop, the greedy opportunistic, the fully opportunistic and the cooperative case.

Hop-by-Hop Case: This case represents the routing strategy in which a packet is transferred from source to destination over a predetermined sequence of nodes. In this scenario, we do not take advantage of the fact that other nodes might also overhear the transmission due to broadcast nature of the medium. In Figure 1, this predetermined path is $< S, R_1, R_2 .... R_N, D >$. Source S first sends the packet to R_1, which in turn sends it to R_2. R_2 then sends the packet to R_3. This process is continued until the packet reaches destination D. In Figure 2, the source sends the packet to a specific relay which in turn sends the packet to the destination. In Figure 3, a source again sends a packet to its corresponding destination through a specific relay, with each source-destination pair choosing a different relay.

Greedy Opportunistic Case: In this case, a myopic view is adopted. Among the nodes that have a copy of the packet, that node that has the highest probability of reaching the destination is entrusted with transmitting the packet in the next time slot. Thus, for the linear network, if multiple nodes have the packet, this greedy approach results in the node closest to the destination always being selected to transmit the packet. In the diamond topology, in the first time slot, the source initially transmits the packet, as it alone has the packet. If the destination successfully receives the packet, then the next time slot begins with the source transmitting a new packet. If the source fails to reach the destination or any of the relays, the source retransmits the packet in the next time slot. However if some of the relays receive the packet and the destination does not, then one of the relays is randomly selected to transmit the packet until it is received at the destination. As all relays are a distance d from the destination, the greedy opportunistic case randomly chooses one of these relays.

The greedy policy for the double diamond scenario is similar to that of the diamond topology. Each source acts independently. If none of the relays possess the packet the source itself transmits the packet. However, once one of the relays has the packet, one of these relays is randomly selected to transmit the packet to its destination. We note that in our idealized model, we do not consider the overhead associated with selecting a transmitting relay.

Fully Opportunistic Case: Here we assume the presence of an oracle that is aware of the prevalent channel conditions between all pairs of nodes at the beginning of any time slot. The oracle chooses that transmission to occur which can either get the packet through to the destination or can move the system to a state where a larger number of relay nodes have a copy of the packet.

Consider, for example, a linear network having 2 relays R_1 and R_2 between source S and destination D. Suppose that S and R_2 have a copy of the packet. Let us assume that at the beginning of the next time slot S has a good [2]SNR to D, while R_2 has a bad SNR to D. In the next time slot S will thus transmit the packet. Similarly, consider the case where R_2 has a bad SNR to D and R_1, S has a bad SNR to D, but S has a good SNR to R_1. In the fully opportunistic case, S will transmit the packet as the system will transition to a state where S, R_1 and R_2 have a copy of the packet. Note the difference between the greedy and the fully opportunistic scenarios: in the greedy scenario, R_2 being closer to the destination will be always given the opportunity to transmit the packet while in the fully opportunistic case,the oracle selects the best node after determining channel conditions.

For the diamond topology, we also illustrate the fully opportunistic case with an example. At the beginning of a time slot let us assume that the source s and i relays have the packet. Among these nodes (i.e., source s and the i relays) if only one them can reach the destination, that transmission takes place. If multiple nodes can reach the destination any one of these nodes is allowed to transmit. However, if none of the nodes having the packet can reach the destination, the node that can reach the largest number of relays that do not currently have the packet is allowed to transmit. The number of relays not having the packet is $(N - i)$. Let us assume that the j^{th} relay

[2]by a good SNR we mean that the SNR is above the threshold β

can reach k_j of the remaining $(N-i)$ nodes while the source can reach k_s of these nodes. The oracle selects the node to transmit which has $max(k_s, k_1......k_i)$. Such a transmission occurs because it enables the system to transition to a state where a larger number of relays have a copy of the packet.

For the topology considered in Figure 3 two different packets can exist in the network. As mentioned earlier, when nodes possessing different packets transmit simultaneously, interference can prevent the successful reception of the packet at a receiver. For successful reception, the $SINR$ must be above a required threshold. Therefore in the fully opportunistic scenario in Figure 3, the transmissions that occur are determined as follows (again, assuming an idealized oracle-based capability). If there are two nodes (one possessing packet 1 and the other possessing packet 2) such that their simultaneous transmissions will enable both the packets to reach their respective destinations such a set of transmissions will occur. However, if no such combination is available the oracle determines which of the nodes having either packet can reach its corresponding destination in the absence of other interfering transmissions. Let us assume that a node R has a packet for destination D_1 and is capable of reaching it in the absence of interference. Along with this transmission, a simultaneous transmission from another node having a packet for D_2 is allowed if the interference caused by this second transmission is low enough for D_1 to successfully receive the packet from R. If a situation arises where multiple nodes can reach their respective destinations in the absence of interference, all possible combinations of simultaneous transmissions are considered. This search selects that combination of two nodes to transmit such that one transmission enables one of the destinations to receive its packet while the other one ensures that maximum number of relays not having any one of the two packets can receive it. Following a similar argument, if none of the nodes having either packet is able to reach the appropriate destination, even in the absence of interfering transmissions, one node among those having packet 1 and one node from those having packet 2 transmit such that the largest number of nodes not having any packet currently, can receive any one of the two packets.

Cooperative Case: Cooperation is based on the premise that if multiple nodes having a copy of the same packet synchronize their transmissions, a receiver will be able to receive the packet if the additive sum of their SNR is above a required threshold. Therefore by exploiting cooperation it is possible for a node to successfully decode a packet, even if individually the nodes have a bad SNR to that receiver. For Figures 1 and 2 having a single source-destination pair, our cooperative policy allows all nodes having the packet to transmit and their signal strengths add up and aid in getting the packet through to the destination. For the multiple flow situation depicted in Figure 3 if we blindly allow all nodes having a packet to transmit it is likely that the interference experienced at a receiver will be large, resulting in it being unable to decode either packet. To study if cooperation will be beneficial in presence of multiple sources we adopt the following strategy: among the nodes possessing packet for D_1 those two nodes having the maximum SNR to D_1 will cooperate in sending the packet. A similar rule is adopted by nodes having packet destined for D_2. We consider this simple strategy because it is the minimum amount of cooperation that is possible when multiple flows exist.

V. Analytical Framework

We now present an analytical model to determine the throughput for the two topologies considering the various policies described in the previous section. We propose a discrete Markov chain model that produces a closed form expression for the throughput. We consider that each node can be in any one of the two states - 0 or 1, where 0 and 1 denote the absence and presence of the packet at a node respectively.

A. Linear Network

For the line network we perform the analysis considering that there are 2 relays between source and destination. Later, we derive an expression for throughput for the general line topology for the greedy opportunistic case. The throughput of the two-relay linear network is simple and given by $p_1/3$. We perform a Markov chain analysis to determine the throughput of the system in the opportunistic and cooperative cases. In both cases, the state of the Markov chain is (S, R_1, R_2, D).

Greedy Opportunistic Case: The initial state of the linear network under greedy opportunistic transmission is $A-(1,0,0,0)$ with the source S attempting to transmit a new packet. Apart from this initial state, the system can be in two other possible states $B-(0,1,0,0)$ and $C-(0,0,1,0)$. As soon as the destination receives the packet the system immediately reverts back to state A with zero delay. The corresponding transition probability matrix $\{P_M\}$ is shown in (2).

$$P_M = \begin{pmatrix} p_3 + q_1q_2q_3 & p_1q_2q_3 & p_2q_3 \\ p_2 & q_1q_2 & p_1q_2) \\ p_1 & 0 & q_1 \end{pmatrix} \quad (3)$$

Note that no transition from state C to B is possible. The steady state probabilities are denoted by π_A, π_B and π_C for states A, B and C then the throughput of the system under the greedy opportunistic policy is given by,

$$T = p_3\pi_A + p_2\pi_B + p_1\pi_C \quad (4)$$

Fully Opportunistic Case: The initial state of the linear network under fully opportunistic transmission is $A-(1,0,0,0)$ and from there the system can transition to any one of the following three states, namely $B-(1,1,0,0)$, $C-(1,0,1,0)$ and $D-(1,1,1,0)$. (5) provides the transition matrix under the fully opportunistic policy.

$$P_M = \begin{pmatrix} p_3 + q_1q_2q_3 & p_1q_2q_3 & q_1p_2q_3 & p_1p_2q_3 \\ 1-q_2q_3 & q_1q_2^2q_3 & 0 & (1-q_1q_2)q_2q_3 \\ 1-q_1q_3 & 0 & q_1^3q_3 & (1-q_1^2)q_1q_3 \\ 1-q_1q_2q_3 & 0 & 0 & q_1q_2q_3 \end{pmatrix} \quad (5)$$

The state transition probabilities are different from the greedy opportunistic case because we take into account the fact that all nodes that receive the packet may be potential candidates for transmitting the packet in the next time slot. Similar to the greedy scenario, the steady state probabilities are denoted by π_A, π_B, π_C and π_D and thus the throughput can be expressed as,

$$T = p_3\pi_A + (1-q_2q_3)\pi_B + (1-q_1q_3)\pi_C + (1-q_1q_2q_3)\pi_D \tag{6}$$

Cooperative Case: The states of the Markov chain for the linear network under cooperative transmission are identical to those of the fully opportunistic case. The combined SNR at a receiver is the sum of exponential random variables. However the distribution of the SNR is not an Erlang Distribution because of the different distances between the various transmitters and receivers. For example, considering that the system is in state B both S and R_1 have a copy of the packet. The distance between S and R_2 and that between R_1 and R_2 are $2d$ and d respectively and hence the combined SNR at R_2 is the sum of two exponential random variables with different rates. The transition matrix for the cooperative case is shown in (7).

$$P_M = \begin{pmatrix} p_3 + q_1q_2q_3 & p_1q_2q_3 & q_1p_2q_3 & p_1p_2q_3 \\ d_1 & g_1 & 0 & a_1 \\ e_1 & 0 & h_1 & b_1 \\ f_1 & 0 & 0 & c_1 \end{pmatrix} \tag{7}$$

where,
$a_1 = [1 - \frac{4}{3}(q_1 - \frac{q_2}{4})][\frac{36}{5}(\frac{q_2}{4} - \frac{q_3}{9})]$
$b_1 = [p_1 - ln(p_1)p_1][\frac{9}{8}(q_1 - \frac{q_3}{9})]$
$c_1 = \frac{36}{5}[\frac{1}{3}(q_1 - \frac{q_2}{4}) - \frac{1}{8}(q_1 - \frac{q_3}{9})]$
$d_1 = [1 - \frac{36}{5}(\frac{q_2}{4} - \frac{q_3}{9})]$
$e_1 = [1 - \frac{9}{8}(q_1 - \frac{q_3}{9})]$
$f_1 = 1 - c_1$
$g_1 = [\frac{4}{3}(q_1 - \frac{q_2}{4})][\frac{36}{5}(\frac{q_2}{4} - \frac{q_3}{9})]$
$h_1 = [1 - (p_1 - ln(p_1)p_1)][\frac{9}{8}(q_1 - \frac{q_3}{9})]$
Following similar reasoning as before we can determine the throughput as,

$$T = p_3\pi_A + d_1\pi_B + e_1\pi_C + f_1\pi_D \tag{8}$$

Analyzing the Greedy Opportunistic Case for a General Line Topology: Thus far we have used Markov chains to analyze the linear topology. However extending this approach for more than 4 nodes is challenging. So here we present a simple recursive relation for studying the greedy opportunistic case for a general linear network in Figure 1, having a number of relays between the source and destination, each relay being d distance away from the preceding one. As noted earlier the probability of successful reception k-hops away is $p_k = p_1^{k^2}$.

Let $L[k]$ be the expected number of hops a packet can go in one transmission given that the destination is k hops away. It is straightforward to obtain the following recurrence equation for $L[k]$:

$$L[k] = kp_1^{k^2} + (1-p_1^{k^2})L[k-1], \qquad L[1] = p_1 \tag{9}$$

That is, with probability $p_1^{k^2}$, it can reach the destination in one transmission (namely, going k hops), otherwise it will reach some node before the destination as if the destination is $(k-1)$ hops away (namely, going $L[k-1]$ hops).

Theorem 1: The solution to $L[k]$ (Eqn. (9)) is:

$$L[k] = \Big(\prod_{i=1}^{k} 1 - p_1^{i^2}\Big) \sum_{j=1}^{k} \frac{jp_1^{j^2}}{\prod_{l=1}^{j} 1 - p_1^{l^2}} \tag{10}$$

Proof: First, it is easy to see $L[1] = p_1$ from Eqn. (10). Next, we substitute Eqn. (10) into Eqn. (9).

$$\begin{aligned} & kp_1^{k^2} + (1-p_1^{k^2})L[k-1] \\ &= kp_1^{k^2} + (1-p_1^{k^2})\Big(\prod_{i=1}^{k-1} 1 - p_1^{i^2}\Big) \sum_{j=1}^{k-1} \frac{jp_1^{j^2}}{\prod_{l=1}^{j} 1-p_1^{l^2}} \\ &= kp_1^{k^2} + \Big(\prod_{i=1}^{k} 1 - p_1^{i^2}\Big) \sum_{j=1}^{k-1} \frac{jp_1^{j^2}}{\prod_{l=1}^{j} 1-p_1^{l^2}} \\ &= \Big(\prod_{i=1}^{k} 1 - p_1^{i^2}\Big)\Big(\frac{kp_1^{k^2}}{\prod_{l=1}^{k} 1-p_1^{l^2}} + \sum_{j=1}^{k-1} \frac{jp_1^{j^2}}{\prod_{l=1}^{j} 1-p_1^{l^2}}\Big) \\ &= L[k] \end{aligned}$$

Hence, Eqn. (10) is a solution to Eqn. (9). ∎

In fact, we can approximate as $L[k] \approx p_1 + 2p_1^4 - p_1^5$, which is independent of k when p_1 is not large.

Theorem 2: When $k \geq 2$ and p_1 is small,

$$L[k] = p_1 + 2p_1^4 - p_1^5 + O(p_1^9) \tag{11}$$

Proof: By Theorem 1, we obtain:

$$\begin{aligned} L[k] &= \Big(\prod_{i=1}^{k} 1 - p_1^{i^2}\Big) \sum_{j=1}^{k} \frac{jp_1^{j^2}}{\prod_{l=1}^{j} 1-p_1^{l^2}} \\ &= \sum_{j=1}^{k} jp^{j^2} \Big(\prod_{i=j+1}^{k} 1 - p_1^{i^2}\Big) \\ &= p_1\Big(1 - p_1^4 + O(p_1^9)\Big) + 2p_1^4\Big(1 + O(p_1^6)\Big) + O(p_1^9) \\ &= p_1 + 2p_1^4 - p_1^5 + O(p_1^9) \end{aligned}$$

Using the above recursive relation one can deduce the throughput to be

$$T \approx \frac{p_1 + 2p_1^4 - p_1^5}{k} \tag{12}$$

∎

B. Diamond Network

In this subsection, we evaluate the diamond network in which there are in general N relays between the source and the destination. For this diamond topology obtaining the throughput for the conventional hop-by-hop case is simple, and is given by $p_1/2$. Determining the throughput for the opportunistic and cooperative policies is more challenging and we once again use Markov Chains to analyze this scenario. As before, each node can be in either of the two states 0 and 1. As relays are equidistant from one another, we are not concerned as to which relay actually has a copy of the packet; instead we are only interested in the number of relays which have a copy. Let us assume that there are N relays in the system. If i relays have the packet, then the system is said to be in state i. There are $(N+1)$ states - starting from 0 and extending up to N. The source always retains a copy of the packet and thus the system being in state i indicates that there are $(i+1)$ nodes in the system that have a copy of the packet (i.e. the i relays and the source). From the discussion above we observe that the number of entries in the transition matrix increases as $O(N^2)$. Hence we devise a general strategy to determine the transition matrix. P_{ij} denotes the probability of transitioning from state i to j.

Greedy Opportunistic Case: We illustrate the procedure for determining the transition matrix entries for the greedy opportunistic case below.

$$P_{ij} = \begin{cases} \binom{N}{j} p_1^j q_1^{N-j} q_2, & \text{if } i = 0, j \neq 0, j \geq i \\ \binom{N-i}{j-i} p_1^{j-i} q_1^{N-j+1}, & \text{if } i \neq 0, j \neq 0, j \geq i \\ 0, & \text{if } j \neq 0, j < i \\ 1 - \sum_{k=1}^{N} P_{ik}, & \text{if } j = 0 \end{cases}$$

Let us consider one of these cases $(i \neq 0, j \neq 0, j \geq i)$ in detail. Such a transition will only occur if the destination fails to receive the packet and any $j-i$ relays among the remaining $N-i$ relays receive the packet. The probability of the destination not receiving the packet in the greedy opportunistic case (one of the relays is randomly chosen here) is q_1 and the probability that $j-i$ relays among the remaining $N-i$ relays receive the packet is given by a binomial distribution with parameters $(N-i, p_1)$. Combining these two probabilities we obtain the value $\binom{N-i}{j-i} p_1^{j-i} q_1^{N-j+1}$.

Fully Opportunistic Case: For the fully opportunistic case if none of nodes that has a copy of the packet can reach the destination, then the node that can reach the largest subset of nodes that currently do not have a copy of the packet is allowed to transmit. Let us assume that the system is currently in state i. Let $X_1, X_2 .... X_i$ be random variables denoting the number of relays (which currently don't have a copy of the packet) that each of the i relays can reach. Let X_s be the corresponding random variable for the source. Each of these random variables will follow a binomial distribution with parameters $(N-i, p_1)$. Let Y be the random variable denoting $max(X_1, X_2 .... X_i, X_s)$. The strategy for determining the transition matrix is provided below.

$$P_{ij} = \begin{cases} \binom{N-i}{j-i} P[Y = j-i] q_1^i q_2, & \text{if } j \neq 0, j \geq i \\ 0, & \text{if } j \neq 0, j < i \\ 1 - \sum_{k=1}^{N} P_{ik}, & \text{if } j = 0 \end{cases}$$

We consider the case $(j \neq 0, j \geq i)$ as an example to provide an intuition. As in the greedy opportunistic case, P_{ij} can be expressed as the probability that the packet fails to reach the destination $(q_1^i q_2)$ and the probability that the node that can reach the largest number of nodes that currently don't have a copy of the packet is allowed to transmit $(\binom{N-i}{j-i} P[Y = j-i])$.

Cooperative Case: For the cooperative scenario, one can approximately determine the transition probabilities. The strategy for filling the matrix is described below.

$$P_{ij} = \begin{cases} \binom{N}{j} p\prime_i^j q\prime_i^{N-j} q_2, & \text{if } i = 0, j \neq 0, j \geq i \\ \binom{N-i}{j-i} p\prime_i^{j-i} q\prime_i^{N-j} q_2 q\prime_{i-1}, & \text{if } i \neq 0, j \neq 0, j \geq i \\ 0, & \text{if } j \neq 0, j < i \\ 1 - \sum_{k=1}^{N} P_{ik}, & \text{if } j = 0 \end{cases}$$

where $p\prime_i$ is an Erlang distribution and with the following form: $p\prime_i = \sum_{n=0}^{i} \frac{p_1(-ln(p_1))^n}{n!}$. $q\prime_i$ is given as $q\prime_i = 1 - p\prime_i$. As before for the case $(j \neq 0, j \geq i)$ the probability of not being able to reach the destination can be approximated by $q_2 q\prime_{i-1}$ while the probability of reaching the $j-i$ relays can be expressed as $\binom{N-i}{j-i} p\prime_i^{j-i} q\prime_i^{N-j}$.

C. Upper Bound Calculations

Given a particular value of p_1 the maximum throughput attainable for the diamond topology (as $N \to \infty$) for the various policies can be determined. For the Hop-by-Hop scenario the upper bound is $p_1/2$. For the remaining policies we determine it in the following way. Let us assume that the source transmits M packets in all. For the fully opportunistic and cooperative schemes the $p_2 M$ packets will get through in 1 time slot. The remaining $q_2 M$ packets will get through in 2 time slots because of the large value of N. Therefore the total number of time slots required to get the M packets through to the destination is $(p_2 + 2q_2)M$. Hence the maximum throughput attainable (T_u) for these two scenarios can be expressed as $\frac{1}{p_2 + 2q_2}$. For the greedy opportunistic scenario the total number of time slots required to get the M packets through to the destination is $(p_2 + (p_1 + 1)\frac{q_2}{p_1})M$ and thus $T_u = \frac{1}{p_2 + (p_1+1)\frac{q_2}{p_1}}$. For example, using the above-mentioned expressions and considering $p_1 = 0.5$ the values of T_u for the greedy opportunistic, fully opportunistic and cooperative cases can be determined as 0.3478, 0.5161 and 0.5161 respectively.

VI. Analytic Results

In this section we plot the analytic results for the linear and diamond topologies studied in this paper for the various forwarding strategies (hop-by-hop, greedy opportunistic, fully opportunistic, cooperative). Figure 4 shows the throughput of the linear network for the various policies as function of p_1. As expected, we observe that the opportunistic and the cooperative schemes perform considerably better than the traditional hop-by-hop strategy. We note that at low reception probabilities, cooperation and opportunism can provide only a small gain in throughput by taking advantage of relays that are far away.

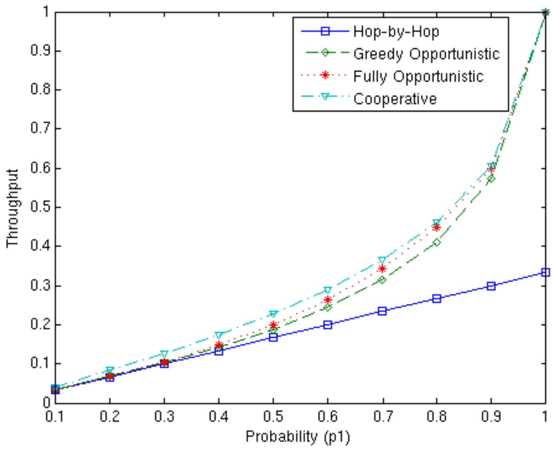

Fig. 4. Throughput of Line Network

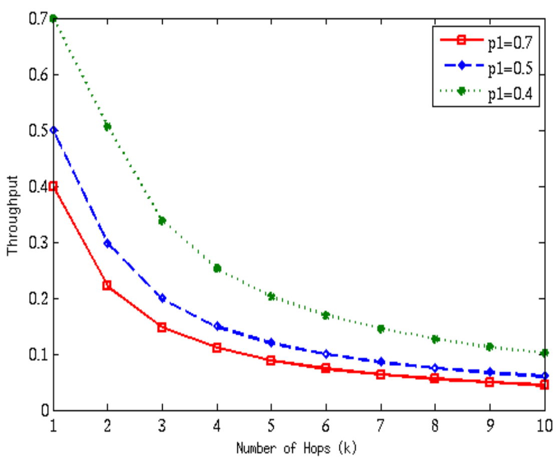

Fig. 5. Throughput for a Greedy Policy in a general Line Network

However as the reception probability increases, the chance of getting a good SNR to nodes far away increases and thus under these circumstances opportunism and cooperation outperform the conventional routing scheme. We observe that when $p_1 \rightarrow 1$, a single transmission will be able to successfully get the packet through to the destination and thus throughput of opportunistic and cooperative schemes converge to 1.

Figure 5 illustrates how throughput scales in the linear network using the recurrence in (9) for the greedy opportunistic case. As expected the graph shows a downward trend as the number of nodes increases. If we consider the case of 3 hops we find that the results correspond with those shown in Figure 4.

Results analyzing the diamond network are depicted in Figure 6 and 7. From Figure 6 the cooperative and opportunistic schemes have a distinct advantage over the traditional hop-by-hop scenario. Another interesting observation is that significant improvements in throughput can be obtained by using the

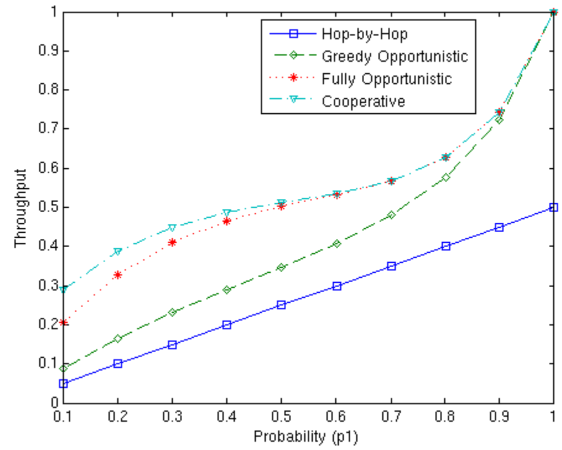

Fig. 6. Throughput of Diamond Network for varying values of p_1:-N=10

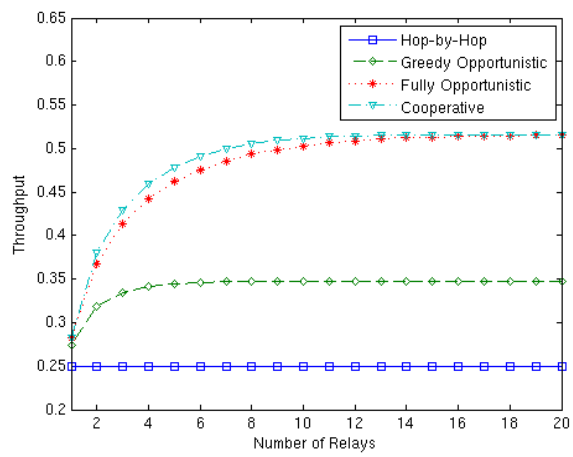

Fig. 7. Throughput of Diamond Network for varying number of relays: p_1=0.5

fully opportunistic and cooperative schemes over the greedy opportunistic scheme, especially at low values p_1. We also observe that when the reception probability is low, the throughput obtained by cooperation is higher than the fully opportunistic scenario. This suggests that if there are multiple relays close to the transmitting node, a combined effort taking advantage of these nodes can significantly increase the throughput at lower reception probabilities. Figure 7, shows the benefits of cooperation and opportunism as a function of the number of relays for $p_1 = 0.5$. The graph also shows that when there are a fewer number of relays, cooperation has a marginal throughput improvement over opportunism. However, with an increasing number of relays this gain diminishes and finally their performance becomes equal. Further, we also observe from Figure 7 that the value of T_u matches with the result upper bound results obtained in the previous section.

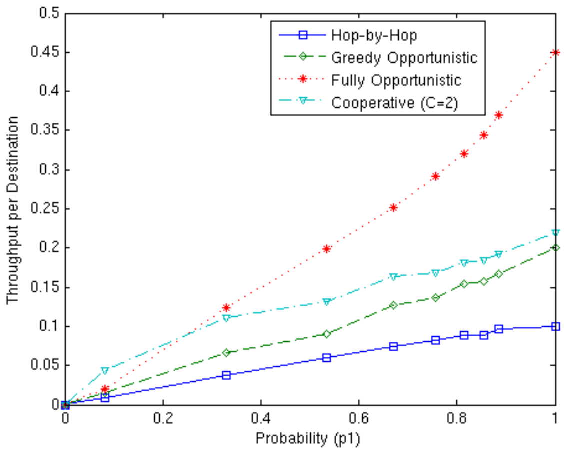

Fig. 8. Throughput per node for 2 source-destination pair

VII. Simulation Results

The analytic results in the previous section suggest that opportunism and cooperation can provide a significant increase in performance. However, how does this conclusion change when multiple flows are considered? In this section we present simulation results for the scenario described in Figure 3 where we have two simultaneous flows. We consider that there are ($N = 10$) relays between the sources and destinations. Figure 8 presents the results. Cooperative (C=2) denotes that for each destination among the nodes having the packet, two nodes having the best SNR cooperate and transmit. We observe that a greedy opportunistic or a naive cooperative method gives poorer throughput and that the gains over the hop-by hop scenario are small. Choosing the transmitter after considering all the prevalent channel conditions can provide significant throughput gains as evidenced by observing the curve for the fully opportunistic case. With multiple transmitters the throughput gains of cooperation and opportunism are likely to diminish further. Our work suggests that adopting a strategy that takes the channel conditions between the various nodes into account is useful while making decisions locally will not be beneficial. This result further highlights the point that to reap the benefits of cooperation, intelligent scheduling strategies are needed to combat the interference from other ongoing transmissions.

VIII. Conclusion and Futre Work

In this paper we have used simple modeling and analysis, augmented by simulation, to investigate the performance benefits of using opportunism and cooperation rather than traditional hop-by-hop, fixed-path routing in a multi-hop wireless mesh network. We consider several different topologies, including a linear network and a diamond network with a group of relays between source and destination, with one or two source-destination pairs. Our analysis has quantified the throughput gains that are possible and provided insight into when opportunism does (and does not) provide gains over traditional hop-by-hop, fixed-path routing. Our future research will focus on the case of multiple network flows, and more general topologies.

References

[1] Sanjit Biswas and Robert Morris. EXOR: Opportunistic Routing for Wireless Networks. In*Proc. of Sigcomm'05.*

[2] Francois Baccelli , Baromiej Baszczyszyn and Paul Muhlethaler On the Performance of Time-Space Opportunistic Routing in Multihop Mobile Ad Hoc Networks In*Proc. of WiOPT'08.*

[3] Zhiguo Ding, Kin. K. Leung, Dennis. L. Goeckel and Don Towsley A Relay Assisted Cooperative Transmission Protocol for Wireless Multiple Access Systems In*Proc. of IEEE Transactions on Communications, 2008.*

[4] Mei-Hsuan Lu, Peter Steenkiste and Tsuhan Chen. Video Transmission over Wireless Multihop Networks Using Opportunistic Routing. In *Proc. of Packet Video Workshop(PV' 07), 2007.*

[5] Szymon Chachulski and Dina Katabi. Trading Structure for Randomness in Wireless Opportunistic Routing. http://people.csail.mit.edu/szym/sm.pdf

[6] Szymon Chachulski, Michael Jennings, Sachin Katti and Dina Katabi. Trading Structure for Randomness in Wireless Opportunistic Routing. In *Proc. of Sigcomm'07*

[7] Anna Scaglione, Dennis L. Geockel and J. Nicholas Laneman. Cooperative Communications in Mobile Ad-Hoc Networks: Rethinking the Link Abstraction. *IEEE Signal Processing Mag., vol. 23, no. 5, pp. 18-29, Sept. 2006.*

[8] J. Nicholas Laneman. Cooperative Diversity: Models, Algorithms and Architectures. In *Cooperation in Wireless Networks: Principles and Applications. Springer, 2006, ch6.*

[9] M. Zorzi and R. R. Rao. Geographic Random Forwarding(GeRaF) for ad hoc and sensor networks: multihop performance. *IEEE Trans. on Mobile Computing, Volume 2, Issue 4, Oct.-Dec. 2003 Page(s): 349 - 365.*

[10] Mei-Hsuan Lu, Peter Steenkiste and Tsuhan Chen. A theoretical model for opportunistic routing in ad hoc networks. In *Proc. of Mobicom'09.*

[11] Angela Sara Cacciapuoti, Marcello Caleffi and Luigi Paura. Design, Implementation and Evaluation of an Efficient Opportunistic Retransmission Protocol. In *Proc. of IEEE International Conference on Ultra Modern Telecommunications (ICUMT'09).*

[12] Eric Rozner, Jayesh Sheshadri, Yogita Ashok Mehta and Lili Qiu. SOAR: Simple Opportunistic Adaptive Routing Protocol for Wireless Mesh Networks. *IEEE Transactions on Mobile Computing 2009.*