2009 IEEE 34th Conference on Local Computer Networks (LCN 2009)
Zürich, Switzerland; 20-23 October 2009

Distributed Strategies for Minimum-Latency Cooperative Retransmission in Wireless Networks

Lixiang Xiong* Lavy Libman†,‡ Guoqiang Mao*,‡
*School of Electrical and Information Engineering, University of Sydney, NSW 2006, Australia
†School of Information Technologies, University of Sydney, NSW 2006, Australia
‡Networked Systems research group, NICTA, 13 Garden St., Eveleigh, NSW 1430, Australia
xlx@ee.usyd.edu.au, llibman@it.usyd.edu.au, guoqiang@ee.usyd.edu.au

***Abstract*—We consider cooperative retransmission strategies in wireless networks, where the retransmission of a failed frame is handled not by the original source but rather by common neighbors overhearing the transmission. The majority of existing literature in this space focuses on opportunistic mechanisms for choosing a single "best" neighbor, with a goal of minimizing the number of required (re-)transmissions. However, the coordination overhead of such mechanisms renders them unsuitable in general for scenarios involving delay-critical control or sensing applications, where the delivery latency, rather than number of retransmissions, is the dominant performance criterion. Accordingly, we study a distributed uncoordinated setting, where each neighbor that successfully overhears a frame decides independently whether to retransmit it in subsequent time slots, considering that multiple simultaneous such retransmissions will cause a collision. We employ a Bayesian approach to analyze the evolution of the system state view from the perspective of each cooperative neighbor, and derive a strategy of finding a sequence of retransmission probabilities for every neighbor in each time slot to minimize the expected delivery latency. We demonstrate for a wide variety of scenarios that this strategy achieves a significantly lower expected latency than either traditional retransmission or two-hop routing to the destination.**

I. INTRODUCTION

There has been an increasing interest in *opportunistic routing (OR)* methods in wireless networks in recent years [1]–[4]. Unlike the traditional layered approach, where the route is pre-selected by a network-layer protocol and failures are coped with by link-layer retransmission, OR chooses the next hop on the path to the destination on-the-fly. As a result, this approach is more flexible in responding to temporary variations in wireless link quality. Indeed, in the event of a temporary degratation of a particular link (e.g. due to slow fading), the problem is overcome locally, without involving a costly recalculation of paths by a routing protocol.

The choice of next hop in OR depends on the instantaneous state of the respective links, and is generally performed by obtaining feedback from all the candidates, though the exact details of different methods vary. In [1], each packet is broadcast with a list of next-hop candidates; each neighbor on that list that hears the packet acknowledges it in its turn (i.e. the ACKs are staggered), and the source then chooses the best next hop using information contained in the ACKs. A similar procedure is proposed in [2], except that the choice is based on Clear-to-Send (CTS) responses to a Request-to-Send (RTS) frame, rather than ACKs of an actual data packet. In the approach of [3], potential candidates are identified by location rather than an explicit list, and the CTS replies are not staggered, leading to collisions that are resolved via a binary exponential backoff until a single CTS response is obtained. A common feature of all these schemes is that they always select a *single* next hop, inevitably incurring a coordination overhead in the form of additional latency until the next hop's identity is agreed upon. In another variation that is gaining increasing popularity [4], the next hop is not chosen explicitly; rather, packets are transmitted in large batches without acknowledgments, and upon completion of a batch, each neighbor forwards the packets it has overheard, after waiting a random delay and skipping any packets overheard from other neighbors in the meantime. In this case, the (implicit) overhead latency is amortized among all packets in the batch.

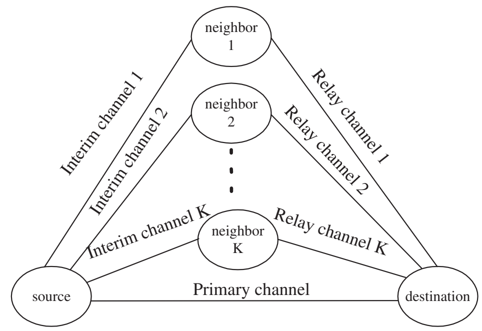

Fig. 1. A network model of cooperative neighbors.

Existing works on performance evaluation of OR methods have predominantly considered the throughput [4], number of transmissions [5], or energy [6], rather than latency, as the performance criterion, thereby ignoring the impact of the coordination overhead. Conversely, we argue that, due to this overhead, OR methods may be unsuitable for delay-critical application scenarios, characterized by a low rate of messages but with stringent delay requirements. Examples range from common event detection or monitoring applications in sensor networks, to more specialized scenarios such as adaptive urban traffic control [7]. This observation motivates us to

‡NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

consider the prospects of a different class of routing methods, involving cooperative forwarding by multiple *uncoordinated* neighbors, that do not attempt to agree on just one forwarder for a packet. For example, consider the network depicted in Figure 1, with $K = 10$ neighbors and a probability of 20% for a transmission to be successful over any individual channel. Then, any traditional or opportunistic routing scheme using a single neighbor will have an expected latency of at least $\frac{1}{0.2} = 5$ transmission slots, whereas an uncoordinated retransmission by all neighbors overhearing the source can be successful with a very high probability after just one slot.

Only a handful of past studies have attempted to consider such uncoordinated cooperation, over a single wireless hop. In [8], a fixed time-division multiple-access (TDMA) scheme is assumed, so that any node overhearing a neighbor's unsuccessful frame may retransmit it in its own allocated slot.† In [9], neighbors persistently retransmit overheard frames until acknowledgment (in a stop-and-wait regime), and a frame is successfully received when at least one neighbor has a 'good' relay channel. Both of these limited models sidestep the possibility of *collision* among the cooperative retransmissions: in [8], such collisions cannot occur by virtue of the fixed TDMA allocation, while [9] implicitly assumes physical-layer multiplexing that allows the frame to be decoded in the event of multiple simultaneous retransmissions. Clearly, without these facilitating assumptions, multiple simultaneous retransmissions may result in collisions; consequently, unlike [8], [9], in general it may not be best for all neighbors to always retransmit their overheard frames.

Accordingly, in this paper we assume a slotted time model, and define a *cooperation strategy* to be a sequence of retransmission probabilities in subsequent time slots, applied by every node that has overheard a copy of the frame. This model naturally gives rise to the problem of finding a strategy profile that minimizes the expected latency, i.e. number of slots until the frame is received collision-free. While an exact solution to this problem is extremely difficult to obtain in general, we show that a significant improvement of the expected frame latency can be achieved using a straightforward strategy, based on a greedy maximization of the success probability in each slot in turn, followed by a Bayesian re-evaluation of the system state from the perspective of each neighbor. As in [8], [9], the focus of the analysis in this paper is on the single-hop setting, with the multi-hop extension left for subsequent work.

Our current study builds upon and significantly extends our earlier work in [10], where a similar optimization problem was defined in a very simple context of memoryless channels, such that the probability of a channel to be in a "good" state was identical and independent in every slot. This memoryless model, which essentially trivialized the respective analysis, may be reasonable for noisy or fast-fading channels, but is not applicable to slow-fading channels (such as, e.g., in outdoor environments), which, crucially, is precisely when cooperative and opportunistic routing methods offer the most significant improvement over traditional (layered) ones. Therefore, in this paper, we specifically consider the 2-state Markov model (also known as an *order-1* Markov or *Gilbert* model), which has been shown by numerous studies to provide an adequate description of the bursty frame loss process in practical wireless fading channels [9], [11], [12]. In particular, a comprehensive experimental study of typical IEEE 802.11 channels [12] has confirmed that, while more complex models are required for an accurate representation of the *bit-level* error process, a 2-state model is quite sufficient at time scales of packets.

The rest of this paper is organized as follows. Section II presents our system model and formally defines the optimal cooperation strategy problem. Section III conducts the detailed analysis of our proposed method, and Section IV presents a numerical study to evaluate the resulting strategy and compare it to traditional alternatives. Section V concludes the paper.

II. System Model and Problem Formulation

We consider a wireless network consisting of a source node, a destination node, and a fixed number K of cooperative neighbor nodes in their vicinity (Figure 1). A frame transmitted by the source may or may not arrive successfully over the direct channel, and may also be overheard by some of the neighbor nodes via the interim channels. Only the intended destination returns an acknowledgment (ACK) upon successful reception; neither the source nor any neighbor can tell which other neighbors, if any, have obtained a copy of the frame. We define a time *slot* to have a sufficient duration for a frame transmission plus the wait for an ACK; it is assumed that slots are of fixed duration and synchronized among the nodes. In the subsequent slots after a frame's first transmission, any node possessing a copy of the frame may decide to make a cooperative retransmission. We assume a Stop-and-Wait regime similar to [9], where only one frame is handled at a time from its first transmission by the source until its eventual acknowledgment. The notable difference that sets our model apart from that of [9] is that, for successful reception, the destination must receive exactly *one* collision-free transmission in a slot.

We assume a two-state Markov channel fading model, in which any channel can be in one of two states: either "on", in which the transmitted signal arrives with sufficient power to be decoded without error (barring a collision), or "off", in which a transmitted signal does not arrive at all. A collision occurs if and only if a node receives two or more transmissions simultaneously with the respective channels being "on" (in other words, a transmission over an "off" channel does not cause any interference).‡

We denote the transition probability between the "off" (*bad*) state to the "on" (*good*) state in every slot (and vice versa) by P_{bg_sd} (respectively P_{gb_sd}) for the direct channel

†We consider the terms *frame* and *packet* to have equivalent meaning for our purpose, though we use "packet" to denote messages in multi-hop networks, and "frame" in the context of a single hop.

‡In a real fading channel, the transition between the two states does not occur instantly, and an intermediate state exists wherein the received signal power is insufficient for correct decoding, but enough to cause interference to other transmissions. For simplicity, we ignore this intermediate state and assume that the time any channel spends in it is negligible; however, we point out that our analysis can be extended from a two-state to a three-state model in a straightforward manner.

(between *source* and *destination*); P_{bg_sn}, P_{gb_sn} for any of the interim channels (between *source* and *neighbor*); and P_{bg_nd}, P_{gb_nd} for any of the relay channels (between *neighbor* and *destination*). Thus, initially (i.e. before the first transmission), the states of all channels are sampled according to their respective steady-state probabilities:

$$P_{ss\_xx} \triangleq \frac{P_{bg\_xx}}{P_{bg\_xx} + P_{gb\_xx}}, \tag{1}$$

where 'xx' $\in$ {'sd','sn','nd'} is substituted for the corresponding channel type. The channel state transitions are assumed to be mutually independent among different node pairs, which is realistic in most practical scenarios where nodes are spaced sufficiently far apart. We emphasize that, for reasons of tractability, our analysis assumes identical channel parameters for all neighbors (thereby resulting in a symmetrical strategy); we discuss the impact of this assumption and possible ways to alleviate it in more detail later in the paper.

We now proceed to formally define the cooperation strategy and formulate the optimization problem that is the focus of the rest of the paper. A *cooperation strategy* of a node is a sequence of probabilities for retransmitting the frame copy, denoted by $\tau_{s[1]}, \tau_{s[2]}, \ldots$ for the source node and $\tau_{n[1]}, \tau_{n[2]}, \ldots$ for any of the neighbor nodes. Each element in the sequence dictates the probability that the node will transmit the frame in the corresponding slot, provided it already has a copy of the frame. The sequence starts from slot 1, corresponding by definition to the first transmission by the source; hence, we require $\tau_{s[1]} = 1$, $\tau_{n[1]} = 0$. The sequence length is infinite in principle, though in practice one may terminate the sequence after a limited number of slots, beyond which the frame is dropped. We emphasize that a strategy is a *fixed* sequence of transmission probabilities, which may not depend on the outcome of any earlier events at the node. This limitation captures the uncoordinated nature of cooperation, which requires every node's strategy to be known to all other nodes in advance, and therefore, be independent of any information that is only known locally. For example, the source node may not base its action in the third slot on whether it overheard any cooperative retransmissions in the second slot, as that information is not known to its neighbors, and they will not be able to compute their own optimal strategies in that slot.

For a strategy profile (i.e. a collection of the individual node strategies), the performance metric used to evaluate it is the expected number of slots until successful delivery of the frame to the destination, i.e. $\sum_{i=1}^{\infty} i \cdot \Pr\{\textit{delivery occurs in slot } i\}$; a strategy profile is *optimal* if it minimizes the above metric. Unfortunately, even though the only factors determining the probability of delivery in slot i are the channel quality parameters and the strategy sequence up to that slot, it is not possible to express the relation with a simple explicit formula. Indeed, this probability depends on the *system state* in slot i, which includes the number of neighbors possessing a copy of the frame, as well as the individual states of all channels; in turn, however, the probability distribution of the system state is dependent on the outcomes in all previous slots. Therefore, we defer the explicit formulation of the performance metric until a later stage in our analysis, in section III.

III. Analysis and Solution Method

A. Definitions and preliminary analysis

We begin by defining more rigorously the notion of *system state* mentioned in the previous section. Clearly, the system state should include all the quantities that impact its future dynamics, and ultimately the strategy performance. In our case, these are: for each of the K neighbors, a binary value indicating whether it has already got a copy of the frame; and for each of the $2K + 1$ channels, a binary value indicating whether it is "on" or "off". This implies that, in principle, the system state consists of a vector of $3K + 1$ binary elements.

Fortunately, the assumption of symmetry among the neighbors allows the relevant state information to be considerably reduced. Accordingly, instead of tracking the state of every individual neighbor and channel, our main idea is to focus on the following probability distributions, henceforth referred to as the *state distributions*:

- $P_{[i]}\{k\}$ is the distribution of the number of neighbors, k, that have overheard a copy of the frame before slot i;
- $\underset{rel[i]}{P}\{r, d|k\}$, where $d \in \{0, 1\}$, is the conditional probability, given that k neighbors have the frame, that exactly r out of them have relay channels in the "on" state *and* the direct channel is "off" ($d = 0$) or "on" ($d = 1$), respectively, in slot i;
- $\underset{int[i]}{P}\{c|k\}$ is the conditional probability, given that k neighbors have the frame, that c out of the *remaining* $K - k$ ones have an interim channel in the "on" state in slot i.

Indeed, we are not interested in the states of interim channels of neighbors that have already received the frame, as they do not impact the future system dynamics in any way. Similarly, the state of the relay channel of any neighbor not yet in possession of the frame does not impact anything else in the system; hence, when a neighbor overhears the frame for the first time, its relay channel state is still governed by the steady-state probability. Therefore, the state distributions defined above capture all the relevant information for the future system dynamics. Thus, the probability of the system in slot i to have k neighbor nodes with the frame, r of them to have relay channels that are "on", c of the *remaining* $K - k$ neighbors to have interim channels that are "on", and the direct channel to be in state d, is $P_{[i]}\{k\}\underset{rel[i]}{P}\{r, d|k\}\underset{int[i]}{P}\{c|k\}$. We henceforth refer to the tuple (k, r, c, d) as the *system state vector*. As will become apparent below, the decoupling of the system state into the above separate distribution functions (but keeping the relay and direct channel states coupled in a single distribution) is necessary to facilitate the calculations involved.

We now proceed to derive an expression for the strategy performance metric. To that end, we first compute the conditional success probability in a generic slot, provided that the system is in a particular state (k, r, c, d), and the strategy values (i.e. retransmission probabilities of the source and neighbors) are τ_s, τ_n in that slot. To be successful, the slot must have exactly one transmission by a node (source or neighbor) whose channel to the destination is "on"; thus,

$$P^{suc}(\tau_s,\tau_n|k,r,d) \triangleq \begin{cases} r\tau_n(1-\tau_n)^{r-1} & d=0 \\ (1-\tau_s)r\tau_n(1-\tau_n)^{r-1}+ & \\ \quad \tau_s(1-\tau_n)^r & d=1 \end{cases} \quad (2)$$

Note that the conditional probability in expression (2) is denoted $P^{suc}(\tau_s,\tau_n|k,r,d)$ rather than $P^{suc}(\tau_s,\tau_n|k,r,c,d)$, since it does not directly depend on the interim channel states. Consequently, the total probability of success in slot i is

$$P_{[i]}^{suc} \triangleq \sum_{k=0}^{K} \sum_{\substack{0\le r\le k \\ d\in\{0,1\}}} P^{suc}(\tau_{s[i]},\tau_{n[i]}|k,r,d)\cdot P_{[i]}\{k\} \underset{rel[i]}{P}\{r,d|k\}, \quad (3)$$

and, finally, the expected frame latency is

$$\sum_{i=1}^{\infty} i\cdot P_{[i]}^{suc}\cdot \prod_{j=1}^{i-1}\left(1-P_{[j]}^{suc}\right). \quad (4)$$

The deceptively simple form of expression (4) may lead to the wrong conclusion that, in order to minimize the expected latency, one must simply find the values of $\tau_{s[i]}$, $\tau_{n[i]}$ that maximize (3) for every slot i separately. Generally, this may not yield the optimal strategy, since it ignores the impact of the strategy choice on the system state distributions in future slots. For instance, the strategy of the source in any slot affects the number of neighbors that overhear the frame for the first time in that slot, and, consequently, the future distribution of k. Due to this dependency, a straightforward exact minimization of (4) is unfeasible. In the following, we analyze an approximate heuristic, based on an iterative greedy maximization of (3) for each slot in turn, ignoring the impact of the strategy choice on the future dynamics.

B. The proposed method

Our proposed heuristic approach to approximate the optimal cooperation strategy operates on each slot iteratively. First, the state distributions are initialized before the first slot as follows:

$$P_{[1]}\{k=0\}=1, \quad P_{[1]}\{k>0\}=0;$$

$$\underset{int[1]}{P}\{c|k=0\} = \binom{K}{c}(P_{ss\_sn})^c(1-P_{ss\_sn})^{K-c};$$

$$\underset{rel[1]}{P}\{r=0,d=1|k=0\} = P_{ss\_sd},$$

$$\underset{rel[1]}{P}\{r=0,d=0|k=0\} = 1-P_{ss\_sd}.$$

Indeed, before the first slot, the number of neighbors with the frame is obviously zero, and the number of interim channels that are "on" during the first transmission is distributed binomially, with a parameter that is the interim channels' steady-state probability. Since $k=0$ with probability 1 during the first slot, it is not necessary to initialize the interim and relay state distributions for other possible values of k.

After the initialization, our heuristic proceeds for each slot i (starting from $i=1$) iteratively. Thus, we assume that the system state distributions for slot i are given; the calculation for that slot then yields the strategy elements $\tau_{s[i]}$, $\tau_{n[i]}$, as well as the state distributions for slot $i+1$. More specifically, the following calculation steps are performed in slot i:

Step 1: the optimal $\tau_{s[i]}^*$ and $\tau_{n[i]}^*$ are solved numerically to maximize expression (3) (except for slot $i=1$, where $\tau_{s[1]}^*=1$ and $\tau_{n[0]}^*=0$ are required by definition);

Step 2: assuming that the slot nevertheless results in a failure, the system state distributions for slot i are revised *a posteriori* (we denote this by attaching a superscript of 'ap' to the respective distributions);

Step 3: finally, the state distributions for slot $i+1$ are computed, accounting for the new neighbors that overhear the frame in slot i and the transitions in the channel states.

The need for an *a posteriori* revision in calculation step 2, which may not be readily apparent, is explained by the following example.

Example: Consider a network with only $K=1$ cooperating neighbor, $P_{bg\_sd}=P_{gb\_sd}=0.5$, $P_{bg\_sn}=0.99$, $P_{gb\_sn}=0.01$, $P_{bg\_nd}=1$, $P_{gb\_nd}=0$, i.e. a direct channel that is "on" half the time, interim channel "on" 99% of the time, and relay channel that is always perfect. Assuming the initial transmission by the source fails, the neighbor has got the frame at the start of slot 2 with a probability of $P_{[2]}\{k=1\}=0.99$. Hence, clearly, the optimal strategy in slot 2 is to allow the neighbor transmit the frame uninterrupted ($\tau_{s[2]}^*=0$, $\tau_{n[2]}^*=1$); indeed, a simultaneous retransmission by the source would far more likely cause a collision than result in a delivery. However, if this strategy is applied and the slot still ends up in a failure, then the distribution must be revised *a posteriori* to $P_{[2]}^{ap}\{k=1\}=0$, as a failure can only occur if the neighbor did not hear the frame after all. This implies $P_{[3]}\{k=1\}=0$ as well, and, therefore, the optimal strategy in slot 3 is $\tau_{s[3]}^*=1$ (the value of $\tau_{n[3]}$ is immaterial). Clearly, ignoring the *a posteriori* revision step and attempting to calculate the strategy in slot 3 independently of the outcome of slot 2 would result in a wasted slot, calling for retransmission by a neighbor that cannot possibly have the frame if the slot is reached at all.

Similarly, consider the case of $K=1$, $P_{bg\_sd}=P_{gb\_sd}=0.5$, $P_{bg\_sn}=1$, $P_{gb\_sn}=0$, $P_{bg\_nd}=0.09$, $P_{gb\_nd}=0.01$. Here, the interim channel is perfect, while the relay channel is "on" 90% of the time in steady-state. Now, the probability of the neighbor having the frame in slot 2 is $P_{[2]}\{k=1\}=1$; also, $\underset{rel[2]}{P}\{r=1,d=0|k=1\} = \underset{rel[2]}{P}\{r=1,d=1|k=1\}=0.45$ while $\underset{rel[2]}{P}\{r=0,d=0|k=1\}=\underset{rel[2]}{P}\{r=0,d=1|k=1\}=0.05$. Thus, again, the best strategy in this slot is a retransmission by the neighbor only. However, a failure in slot 2 will not change the distribution of $P_{[2]}\{k\}$ *a posteriori*, since the neighbor is known to have the frame with certainty. Rather, a failure implies that the relay channel must have been "off" after all, i.e. the relay state distribution is revised to $\underset{rel[2]}{P^{ap}}\{r=1,d=0|k=1\}=\underset{rel[2]}{P^{ap}}\{r=1,d=1|k=1\}=0$, $\underset{rel[2]}{P^{ap}}\{r=0,d=0|k=1\}=\underset{rel[2]}{P^{ap}}\{r=0,d=1|k=1\}=0.5$. After accounting for a single Markov transition step, the relay channel will be "on" with a probability of only 0.09 by slot 3, and it can be verified that the optimal strategy then is $\tau_{s[3]}^*=1$, $\tau_{n[3]}^*=0$.

Finally, we mention that if the direct channel parameters are set to $P_{bg\_sd}=0$, $P_{gb\_sd}=1$, then the optimal strategy trivially becomes $\tau_s^*=1, \tau_n^*=1$ in every slot, for any setting of the interim and relay channel parameters. Indeed, if there is no direct channel between the source and destination (i.e. it is always "off"), then a simultaneous transmission by the source can never interfere with the neighbor, yet it may save time if the neighbor has not heard the frame yet. This shows that, in general, the optimal strategy may have both $\tau_s > 0$ and $\tau_n > 0$ in the same slot. □

C. Detailed analysis

We proceed to elaborate the details of the calculation steps in slot i, outlined above. We do not consider the implementation of step 1 any further; it is beyond the scope of the paper to suggest any specific numerical solution method for the respective maximization. We merely point out that, since expression (3) is in general non-concave, care must be taken to avoid choosing τ_s^*, τ_n^* that only attain a local maximum.†

In step 2, the revision of the state distribution *a posteriori* is achieved using Bayes' formula, i.e. by scaling the probability of every possible state by the likelihood that the strategy (τ_s^*, τ_n^*) would have failed in that state. For convenience, we define, in a similar fashion to (2),

$$P^{suc}(\tau_s, \tau_n|k) = \sum_{\substack{0\le r\le k\\ d\in\{0,1\}}} P^{suc}(\tau_s, \tau_n|k, r, d) \quad (5)$$

and obtain

$$P^{ap}_{[i]}\{k\} = \frac{P_{[i]}\{k\}(1 - P^{suc}(\tau_s^*, \tau_n^*|k))}{\sum_{k'=0}^{K} P_{[i]}\{k'\}(1 - P^{suc}(\tau_s^*, \tau_n^*|k'))} \quad (6)$$

$$\underset{rel_{[i]}}{P^{ap}}\{r, d|k\} = \frac{\underset{rel_{[i]}}{P}\{r, d|k\}(1 - P^{suc}(\tau_s^*, \tau_n^*|k, r, d))}{\sum_{\substack{0\le r\le k\\ d\in\{0,1\}}} (1 - P^{suc}(\tau_s^*, \tau_n^*|k, r, d))} \quad (7)$$

Clearly, there is no *a posteriori* revision of the interim channel state distribution, as it has no impact on the success probability in the slot.

Finally, we proceed to consider calculation step 3, namely finding the system state distributions that are in effect at the beginning of slot $i+1$. This is the least straightforward step, due to the various interactions between the number of new neighbors overhearing the frame in slot i and the channel state transitions, which require a detailed and careful consideration. Throughout this subsection, we use $\hat{k}$, $\hat{r}$, $\hat{c}$, and $\hat{d}$ to denote the system state in slot $i+1$, reserving k, r, c, and d to denote the state variables in slot i.

1) The distribution $P_{[i+1]}\{\hat{k}\}$: A new neighbor will overhear the frame in slot i if and only if the source has transmitted in that slot, and the corresponding interim channel is "on". Therefore, we can define the probability of the system to have $\hat{k}$ frame copies in slot $i+1$ if it had k of them in slot i:

$$\Pi_{[i]}\{k, \hat{k}\} = \begin{cases} \tau_s^* \cdot \underset{int_{[i]}}{P}\{\hat{k} - k|k\} & k < \hat{k} \\ \tau_s^* \cdot \underset{int_{[i]}}{P}\{0|k\} + (1 - \tau_s^*) & k = \hat{k} < K \end{cases} \quad (8)$$

and, consequently,

$$P_{[i+1]}\{\hat{k}\} = \sum_{k=0}^{\hat{k}} P^{ap}_{[i]}\{k\} \cdot \Pi_{[i]}\{k, \hat{k}\}. \quad (9)$$

†Since there is no possibile ambiguity, henceforth we denote the strategy chosen in step 1 of slot i by τ_s^*, τ_n^*, without mentioning the index i explicitly.

2) The distribution $\underset{rel_{[i+1]}}{P}\{\hat{r}, \hat{d}|\hat{k}\}$: To calculate the relay channel state distribution in slot $i+1$, we distinguish between the k "old" channels corresponding to nodes that already had the frame in slot i (whose state distribution is given by $\underset{rel_{[i]}}{P^{ap}}\{r, d|k\}$), and the $\hat{k} - k$ "new" channels of nodes that obtained the frame copy in slot i for the first time. As the state of these "new" channels is independent of the system's history so far, they are still governed by their steady-state distribution; thus, the number thereof that are "on" in slot $i+1$ will be distributed binomially with a parameter of P_{ss_nd}.

To capture the state transitions in the "old" channels, we define an auxiliary function $\underset{old}{\Pi}\{r, r'|k\}$, which is the probability to have r' relay channels (out of the "old" k) in the "on" state given that r of them were "on" in the previous slot. This requires some j out of the r channels to remain "on", plus $r' - j$ additional channels to have a transition from "off" to "on". Thus,

$$\underset{old}{\Pi}\{r, r'|k\} = \sum_{j=\max[0, r'-(k-r)]}^{\min(r,r')} \binom{r}{j}(1 - P_{gb\_nd})^j (P_{gb\_nd})^{r-j} \cdot \binom{k-r}{r'-j}(P_{bg\_nd})^{r'-j}(1 - P_{bg\_nd})^{k-r-(r'-j)} \quad (10)$$

To combine this with the "new" channels, we define another auxiliary function $\underset{rel}{\Pi}\{r, \hat{r}|k, \hat{k}\}$, which is the probability to have a total of $\hat{r}$ relay channels (out of $\hat{k}$) in the "on" state, provided that there were k "old" channels, of which r were "on", in the previous slot. We obtain this by summing over all possible values of r' (thereby requiring $\hat{r} - r'$ of the "new" channels to be "on" in the new slot):

$$\underset{rel}{\Pi}\{r, \hat{r}|k, \hat{k}\} = \sum_{r'=\max[0, \hat{r}-(\hat{k}-k)]}^{\min(\hat{r},k)} \underset{old}{\Pi}\{r, r'|k\} \binom{\hat{k}-k}{\hat{r}-r'} \cdot (P_{ss\_nd})^{\hat{r}-r'}(1 - P_{ss\_nd})^{\hat{k}-k-(\hat{r}-r')} \quad (11)$$

Expression (11) is a conditional probability, given that the number of frame copies in slot i was k. Summing the total probability and taking into account the state transition of the direct channel, we finally obtain

$$\underset{rel_{[i+1]}}{P}\{\hat{r}, 0|\hat{k}\} = \sum_{k=0}^{\hat{k}} \left(\frac{\Pi_{[i]}\{k, \hat{k}\}}{\sum_{k'=0}^{\hat{k}} \Pi_{[i]}\{k', \hat{k}\}} \right) \cdot \sum_{r=0}^{k} \underset{rel}{\Pi}\{r, \hat{r}|k, \hat{k}\} \cdot \Big[(1 - P_{bg\_sd}) \underset{rel_{[i]}}{P^{ap}}\{r, 0|k\} + P_{gb\_sd} \underset{rel_{[i]}}{P^{ap}}\{r, 1|k\} \Big] \quad (12)$$

and

$$\underset{rel_{[i+1]}}{P}\{\hat{r}, 1|\hat{k}\} = \sum_{k=0}^{\hat{k}} \left(\frac{\Pi_{[i]}\{k, \hat{k}\}}{\sum_{k'=0}^{\hat{k}} \Pi_{[i]}\{k', \hat{k}\}} \right) \cdot \sum_{r=0}^{k} \underset{rel}{\Pi}\{r, \hat{r}|k, \hat{k}\} \cdot \Big[P_{bg\_sd} \underset{rel_{[i]}}{P^{ap}}\{r, 0|k\} + (1 - P_{gb\_sd}) \underset{rel_{[i]}}{P^{ap}}\{r, 1|k\} \Big]. \quad (13)$$

3) The distribution $\underset{int[i+1]}{P}\{\hat{c}|\hat{k}\}$: For the sake of simplicity and due to space limits, the analysis derived here will assume isolated neighbors that are unable to overhear their peers' transmissions, and can only overhear the frame from the original source; we refer the reader to an extended version of this work [13] where this assumption is alleviated. Accordingly, it follows that if the source transmitted in slot i, all interim channels of neighbors that do not have the frame by the end of that slot must be "off"; consequently, in slot $i+1$ each such channel has a probability of P_{bg_sn} to be "on", and their total number is distributed binomially. On the other hand, if the source was silent, the number of neighbors without a frame copy does not change, and the distribution of their interim channel states makes a single Markov transition. Accordingly, we define an auxiliary function $\underset{int}{\Pi}\{c,\hat{c}|\hat{k}\}$, which is the probability of $\hat{c}$ interim channels (out of $K-\hat{k}$) to be "on" after that number was c in the previous slot (in a similar manner to (10)):

$$\underset{int}{\Pi}\{c,\hat{c}|\hat{k}\} = \sum_{j=\max[0,\hat{c}-(K-\hat{k}-c)]}^{\min(c,\hat{c})} \binom{c}{j}(1-P_{gb\_sn})^j(P_{gb\_sn})^{c-j}\cdot \binom{K-\hat{k}-c}{\hat{c}-j}(P_{bg\_sn})^{\hat{c}-j}(1-P_{bg\_sn})^{K-\hat{k}-c-(\hat{c}-j)} \quad (14)$$

With the help of this auxiliary function, we obtain

$$\underset{int[i+1]}{P}\{\hat{c}|\hat{k}\} = \Bigg[\tau_s^* \sum_{k=0}^{\hat{k}} P_{[i]}^{ap}\{k\} \underset{int[i]}{P}\{\hat{k}-k|k\} \underset{int}{\Pi}\{0,\hat{c}|\hat{k}\}+ (1-\tau_s^*)P_{[i]}^{ap}\{\hat{k}\}\sum_{c=0}^{\hat{k}} \underset{int[i]}{P}\{c|\hat{k}\}\underset{int}{\Pi}\{c,\hat{c}|\hat{k}\}\Bigg] / P_{[i+1]}\{\hat{k}\}, \quad (15)$$

where $P_{[i+1]}\{\hat{k}\}$ has been calculated in (9). Expression (15) obtains the required conditional probability of $\hat{c}|\hat{k}$ by dividing the total probability of moving into state $\hat{c},\hat{k}$ in slot $i+1$ (in brackets) by the probability of having $\hat{k}$ copies in that slot. The total probability in the brackets is a sum of two terms. The first term corresponds to the cases where the source transmitted in slot i; thus, the interim channels of those neighbors who still do not possess a copy must have been "off" in slot i. The second term corresponds to the case where the source was silent and no additional neighbors received the frame (i.e. $\hat{k}=k$), regardless of the state of their interim channels. In both terms, the auxiliary function $\underset{int}{\Pi}(\cdot)$ is then used to capture a single Markov transition of the interim channel states.

IV. Evaluation of the Cooperation Strategies

As mentioned in the Introduction, it is generally accepted that the popular 2-state Markov model used in this paper is adequate in most practical scenarios [9], [11], and in particular with typical 802.11 channels [12], [14]. To achieve a realistic performance evaluation of the cooperation strategy, we set the channel parameter values based on the results from [14], which explored the correspondence between a channel's average signal-to-noise ratio (SNR) and its Markov model parameters (i.e. average duration in "good" and "bad" states, which is readily converted to P_{bg}, P_{gb}), for various combinations of frame size and transmission rate. Specifically, we use the results of [14] for 802.11b/g channels with 1500-byte frames transmitted at 11Mbps, and with the average SNR of a channel set to $SNR[dB] = -20\log L$, where L is the distance between its endpoint nodes (this corresponds to free-space propagation). Due to space limits, we only present a small representative set of scenarios here, though we observed very similar qualitative effects in a much wider range of cases.

TABLE I
Channel parameter settings for numerical evaluation.

Scenario	L_{sn}	L_{nd}	Interim channel	Relay channel
1	$\frac{1}{2}L_{sd}$	$\frac{1}{2}L_{sd}$	SNR=27.5dB, P_{bg_sn}=0.23, P_{gb_sn}=0.02	SNR=27.5dB, P_{bg_sn}=0.23, P_{gb_sn}=0.02
2	$\frac{\sqrt{3}}{3}L_{sd}$	$\frac{\sqrt{3}}{3}L_{sd}$	SNR=26.3dB, P_{bg_sn}=0.20, P_{gb_sn}=0.04	SNR=26.3dB, P_{bg_sn}=0.20, P_{gb_sn}=0.04
3	$\frac{\sqrt{2}}{2}L_{sd}$	$\frac{\sqrt{2}}{2}L_{sd}$	SNR=24.5dB, P_{bg_sn}=0.16, P_{gb_sn}=0.13	SNR=24.5dB, P_{bg_sn}=0.16, P_{gb_sn}=0.13
4	$\frac{\sqrt{3}}{3}L_{sd}$	$\frac{\sqrt{3}}{3}L_{sd}$	SNR=21.5dB, P_{bg_sn}=0.11, P_{gb_sn}=0.99	SNR=21.5dB, P_{bg_sn}=0.11, P_{gb_sn}=0.99
5	$\frac{1}{2}L_{sd}$	$\frac{\sqrt{3}}{2}L_{sd}$	SNR=27.5dB, P_{bg_sn}=0.23, P_{gb_sn}=0.02	SNR=22.7dB, P_{bg_sn}=0.13, P_{gb_sn}=0.44
6	$\frac{\sqrt{3}}{2}L_{sd}$	$\frac{1}{2}L_{sd}$	SNR=22.7dB, P_{bg_sn}=0.13, P_{gb_sn}=0.44	SNR=27.5dB, P_{bg_sn}=0.23, P_{gb_sn}=0.02

Since a typical cooperative retransmission scenario arguably consists of a poor direct channel with better interim and relay channels, we set L_{sd} such that the direct channel has $\text{SNR}_{sd} = 21.5$dB, which corresponds to $P_{bg\_sd} = 0.11$ and $P_{gb\_sd} = 0.99$; thus, in this case, the direct channel is "on" one-tenth of the time on average. For the neighbors, we consider channel parameters corresponding to several possible locations, as listed in Table I.

It is important to comment that, since our analysis has assumed all neighbors to have identical channel parameters for tractability purposes (resulting in a common sequence of retransmission probabilities), the results reported below will only hold precisely in an idealized setting where all neighbors are located at the same distance to the source and destination, yet have mutually independent channels (e.g. in a circle perpendicular to the source-destination axis). Clearly, in a real network, neighbors will be scattered around in space and have a range of different distances and channel qualities to the source and destination; thus, the average latency attained in practice will generally lie in between the values obtained here for these idealized scenarios. The extension of our method to find strategies that explicitly take advantage of unequal neighbor distances and channel qualities remains an important direction for future work. Nevertheless, the results presented here suffice to obtain an insight to the potential benefits of the uncoordinated cooperative retransmission approach.

The cooperation strategy sequences resulting in the above scenarios are listed in Table II (due to space limits, only the first few elements of each strategy sequence are shown, for $K = 1, \ldots, 5$), and their respective performance is plotted

TABLE II
EXAMPLE STRATEGY SEQUENCES.

K	Strategy for $P_{nn} = 0$	Strategy for $P_{nn} = 1$
	Scenario 1	
1	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.5907,0.5808,0.5817,0.5894)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.5907,0.5560,0.5487,0.5465)$
3	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.3938,0.3936,0.3944,0.3959)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.3938,0.3779,0.3710,0.3677)$
4	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2954,0.2955,0.2960,0.2966)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2954,0.2845,0.2791,0.2762)$
5	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2363,0.2364,0.2367,0.2370)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2363,0.2280,0.2236,0.2212)$
	Scenario 2	
1	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.6531,0.6241,0.6311,0.6496)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.6531,0.5826,0.5752,0.5729)$
3	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.4354,0.4342,0.4361,0.4403)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.4354,0.4056,0.3945,0.3897)$
4	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.3265,0.3268,0.3280,0.3297)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.3265,0.3068,0.2978,0.2933)$
5	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2612,0.2615,0.2622,0.2631)$	$\tau_s = (1,0,0,0,0)$, $\tau_n = (0,0.2612,0.2463,0.2389,0.2349)$
	Scenario 3	
1	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,0.9226,0.8527,0.8564)$
3	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,0.7824,0.8420,0.7881)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,0.6237,0.6589,0.6220)$
4	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.8213,0.6605,0.6256,0.5998)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.8213,0.5143,0.5004,0.4772)$
5	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.6569,0.5704,0.4980,0.4744)$	$\tau_s = (1,0,1,1,0)$, $\tau_n = (0,0.6569,0.4339,0.4029,0.4229)$
	Scenario 4	
1	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
3	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
4	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
5	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
	Scenario 5	
1	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
3	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
4	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,1,1,0.9833)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
5	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.9532,0.8031,0.8170,0.8070)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.9532,0.8028,0.8456,0.8482)$
	Scenario 6	
1	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,1,1)$
2	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,1,0.9427,0.8433)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,0.8163,0.7305,0.6957)$
3	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,0.7976,0.7054,0.6496)$	$\tau_s = (1,1,1,1,1)$, $\tau_n = (0,1,0.6024,0.5112,0.4616)$
4	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,0.5110,0.6707,0.5541)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,1,0.4239,0.5551,0.3617)$
5	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.7943,0.5004,0.4380,0.3676)$	$\tau_s = (1,0,1,1,1)$, $\tau_n = (0,0.9532,0.3637,0.3688,0.2807)$

in Figure 2. Each graph shows the expected latency as a function of the number of cooperative neighbors K. For each scenario, we compare the performance of our strategy to that of simple retransmissions over a direct connection and over a 2-hop connection via one of the neighbors.† In addition, we evaluate the performance under two extreme settings of the inter-neighbor channel quality: one where it is "always-off", i.e. neighbors cannot overhear each other's retransmissions (as calculated in section III of this paper); and the other where it is "always-on", i.e. neighbors always overhear their peers' retransmissions perfectly (the corresponding analysis can be found in the extended version [13]). These cases are labelled in Figure 2 as $P_{nn} = 0$ and $P_{nn} = 1$, respectively. Our motivation for this comparison comes from [10], where it was shown (albeit for the simpler case of memoryless channels) that the inter-neighbor channel quality has only a minor impact on the overall performance of the cooperation strategy.

These results lead us to observe several important insights. First, we note that the simulation results match the analytical results well; furthermore, the difference in performance between the two inter-neighbor channel quality extremes is indeed negligible (intuitively, the ability to overhear the frame from a peer, and not just from the source, is offset by the additional collisions when the source and another neighbor transmit together), leading us to conjecture that the performance will be similar for any non-extreme setting as well.

More importantly, in all scenarios, we observe that the cooperation strategy obtained via our heuristic method (with a proper choice of K) achieves a substantially better expected latency than both the direct connection and the 2-hop routing alternatives. In scenarios 1 and 2, where the interim and relay channels are of very good quality, two-hop routing already comes close to the best possible latency, and one cooperative neighbor is best (more uncoordinated neighbors merely increase the rate of collisions). As the interim and relay channels get worse, the optimal neighbor retransmission probabilities increase, and, furthermore, it becomes better to involve a larger number of cooperative neighbors; this is similar to the effect observed in [10]. Thus, the potential benefits of uncoordinated, simultaneous cooperation by multiple neighbors, which is the subject of this paper, are clearly demonstrated, especially for wireless environments with low-quality channels.

V. CONCLUSION

We considered distributed strategies for cooperative forwarding of packets by multiple uncoordinated nodes. Similar to opportunistic routing (OR), such strategies are designed to cope locally with temporary degradation of wireless links without escalating to a full routing path recalculation. However, unlike most OR methods in existing literature, they avoid the overhead of coordinating the choice of a single forwarder

†For simple retransmissions, the expected latency of successful delivery of a frame over a link with a 2-state Markov channel is

$$\frac{P_{bg}}{P_{bg}+P_{gb}} \cdot 1 + \frac{P_{gb}}{P_{bg}+P_{gb}} \cdot \left(\frac{1}{P_{bg}} + 1\right), \qquad (16)$$

since $\frac{1}{P_{bg}}$ is the expected number of slots for the channel to turn "on" if it was "off" during the initial transmission.

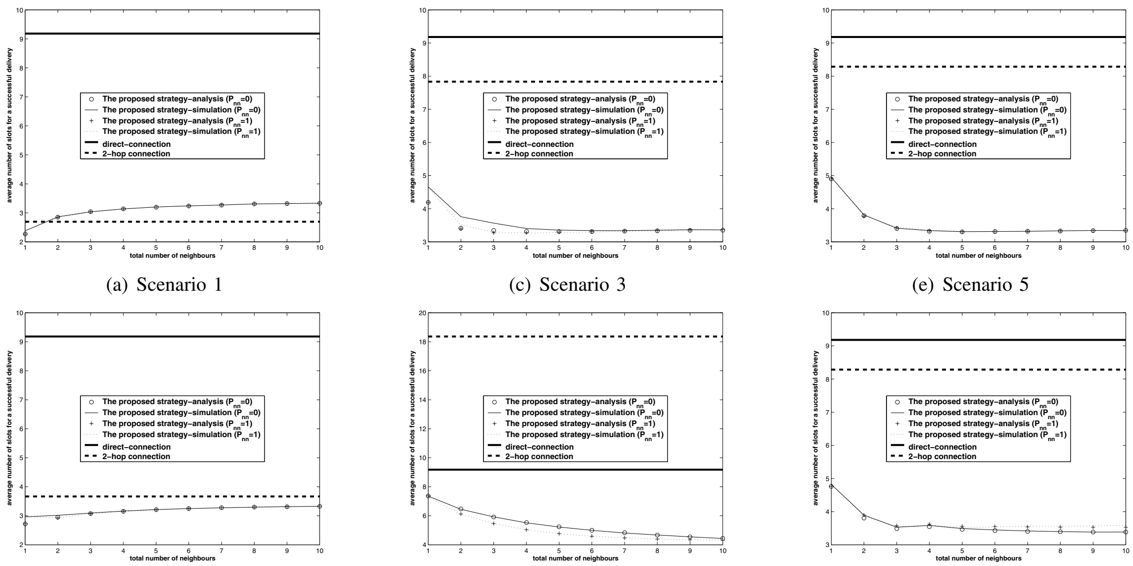

(a) Scenario 1 (c) Scenario 3 (e) Scenario 5

(b) Scenario 2 (d) Scenario 4 (f) Scenario 6

Fig. 2. Performance of uncoordinated cooperative retransmission strategies.

at every hop, making them more suitable in delay-critical applications. In this paper, we defined the fundamental problem of finding a strategy that minimizes the expected delivery latency with multiple cooperative neighbors. We conducted a detailed analysis of a heuristic solution, and demonstrated that the resulting strategies achieve a far superior performance to traditional methods. We emphasize that, even though the calculations in our method may appear to involve complex expressions, their implementation is straightforward and they need only be performed infrequently, i.e. upon discovering a change in the neighbor set of a node. The extension of our results to a generic multi-hop setting is left for further work.

While the Markov model used in our analysis is adequate for describing the frame loss process in fading channels, the utility of our results is still limited by the assumption of identical channel parameters for all neighbors. There are two ways to extend our analysis in this regard. The obvious one is to assume an asymmetric model, assigning individual Markov parameters to each interim and relay channel. However, apart from making the analysis intractable, it would only be relevant in scenarios where the individual channel parameters are known in advance, e.g. in fixed mesh networks, where the nodes are able to measure their long-term channel statistics and recalculate the cooperation strategies upon any changes. An alternative approach that would be practical for common scenarios of mobile ad-hoc networks with small nodes, unable to afford to measure the channel statistics upon every change of their neighbor set, is to assume that neighbors are equivalent *a priori* but scattered in space, and, therefore, their channels' Markov model parameters are not constants but are themselves randomly sampled from a range of possible values. This promising approach is the subject of ongoing work.

REFERENCES

[1] P. Larsson. Selection diversity forwarding in a multihop packet radio network with fading channel and capture. *ACM Mobile Computing and Communications Review*, 5(4):47–54, October 2001.

[2] S. Jain and S.R. Das. Exploiting path diversity in the link layer in wireless ad hoc networks. In *Proc. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Taormina, Italy, June 2005.

[3] M. Zorzi and R.R. Rao. Geographic random forwarding (GeRaF) for ad hoc and sensor networks: Energy and latency performance. *IEEE Transactions on Mobile Computing*, 2(4):349–365, October 2003.

[4] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In *Proc. ACM SIGCOMM*, pages 133–144, Philadelphia, PA, August 2005.

[5] Z. Zhong and S. Nelakuditi. On the efficacy of opportunistic routing. In *Proc. IEEE SECON*, pages 441–450, San Diego, CA, June 2007.

[6] P. Casari, A. Marcucci, M. Nati, C. Petrioli, and M. Zorzi. A detailed simulation study of geographic random forwarding (GeRaF) in wireless sensor networks. In *Proc. IEEE MILCOM*, pages 59–68, Atlantic City, NJ, October 2005.

[7] K-C. Lan, Z. Wang, R. Berriman, T. Moors, M. Hassan, L. Libman, M. Ott, B. Landfeldt, Z. Zaidi, A. Seneviratne, and D. Quail. Implementation of a wireless mesh network testbed for traffic control. In *Proc. Workshop on Wireless Mesh and Ad Hoc Networks (WiMAN)*, Honolulu, HI, August 2007.

[8] I. Cerutti, A. Fumagalli, and G. Ho. Saturation throughput gain in fixed multiplexing radio networks with cooperative retransmission protocols. In *Proc. IEEE International Conference on Communications (ICC)*, volume 10, pages 4489–4494, Istanbul, Turkey, June 2006.

[9] M. Dianati, X. Ling, K. Naik, and X. Shen. A node-cooperative ARQ scheme for wireless ad-hoc networks. *IEEE Transactions on Vehicular Technology*, 55(3):1032–1044, May 2006.

[10] L. Xiong, L. Libman, and G. Mao. Optimal strategies for cooperative MAC-layer retransmission in wireless networks. In *Proc. IEEE Wireless Communications and Networking (WCNC)*, Las Vegas, NV, April 2008.

[11] M. Zorzi, R.R. Rao, and L.B. Milstein. ARQ error control for fading mobile radio channels. *IEEE Transactions on Vehicular Technology*, 46(2):445–455, May 1997.

[12] S.A. Khayam and H. Radha. Markov-based modeling of wireless local area networks. In *Proc. ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pages 100–107, San Diego, CA, September 2003.

[13] L. Xiong, L. Libman, and G. Mao. Distributed strategies for minimum-latency cooperative retransmission in wireless networks. NICTA Technical Report ATP-93, March 2008. Available from http://www.it.usyd.edu.au/~llibman/papers/TR2008-ATP-93.pdf.

[14] J. Aráuz and P. Krishnamurthy. Markov modeling of 802.11 channels. In *Proc. IEEE Vehicular Technology Conference (VTC)*, pages 771–775, Orlando, FL, October 2003.