This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the WCNC 2008 proceedings.

Optimal Strategies for Cooperative MAC-Layer Retransmission in Wireless Networks

Lixiang Xiong*,† Lavy Libman† Guoqiang Mao*,†

*School of Electrical and Information Engineering
University of Sydney
Sydney, NSW 2006, Australia
{xlx,guoqiang}@ee.usyd.edu.au

†Networked Systems research group, NICTA
Bay 15, Australian Technology Park
Eveleigh, NSW 1430, Australia
Lavy.Libman@nicta.com.au

***Abstract*—The concept of *cooperative retransmission* in wireless networks has attracted considerable research attention. The basic idea is that when a receiver cannot decode a frame, the retransmission is handled not by its original source but rather by a neighbour that overheard the transmission successfully, and may have a better channel to the destination. However, the majority of existing literature tackles the issue from the physical layer perspective, with either a single cooperating neighbour, or a multiple-neighbour setting where the receiver is capable of combining and decoding the signal from several simultaneous retransmissions. In this paper, we consider the case of multiple cooperating neighbours from a MAC-layer perspective. Thus, we assume a receiver that can only decode one transmission at a time, while multiple simultaneous retransmissions (by several neighbours that had overheard the frame successfully) will cause a collision. As a result, each neighbour with a successfully overheard copy of the frame faces a tradeoff between helping with a cooperative retransmission and possibly causing a collision. Accordingly, we pose the optimization problem of finding a distributed randomized strategy for the cooperating neighbours, which assigns a certain retransmission probability to every neighbour in each time slot, so as to minimize the expected latency until successful reception. We analyse the performance achieved by two approaches: one where the original source is silent while the neighbours conduct their cooperative retransmissions, and another where both the source and the neighbours may have a nonzero retransmission probability simultaneously. We show that the latter approach offers a significant performance improvement over the former one, as well as either traditional retransmission or two-hop routing to the destination.**

I. Introduction

In the traditional layered design approach to wireless networks, the route between a source and a destination is selected by a network-layer protocol, and each node along the route is responsible for transmitting (and retransmitting, if necessary) the respective packets to the next hop. This approach is inflexible in responding to variations in wireless link quality. Thus, in the event of a temporary degradation or outage of a particular link (e.g. due to fading or local interference), there is no way for packets to be delivered via an alternative route, short of re-running the route selection protocol. The concept of *cooperative retransmission*, which has attracted increasing research attention in recent years, allows the above limitation to be overcome on a local basis. The idea is that, when a packet cannot be decoded by its intended next-hop receiver but is successfully overheard by a common neighbour, that neighbour may retransmit the packet on behalf of the sender. The packet is effectively relayed via a neighbour with a better channel to the receiver, without escalating the problem to the routing layer.

While the idea of cooperation between wireless neighbours to improve the effective channel quality to a receiver is not new, the majority of related research has taken place in the *physical* layer; see, e.g., [1], [2] for a detailed survey of the solutions used with a single neighbour (relay), and the growing literature on *cooperative diversity* (e.g., [3]–[5]) for the more general case of multiple relays. Essentially, this can be seen as an extension of the *spatial diversity* concept of multiple-input multiple-output (MIMO), where the multiple antennas are located at the cooperative nodes. Such cooperative-diversity methods require the receiver to be able to combine the cooperative signals and decode them jointly; therefore, they are unsuitable for simple receivers, with a traditional single-antenna/single-user decoder.

At higher layers, much of the related research has focused on the concept of *opportunistic forwarding*, where the next hop of each packet is determined on-the-fly (rather than pre-selected by a routing protocol). One of the earliest proposals of the idea is [6], where a packet is broadcast with a list of candidate next-hop forwarders; neighbours that hear the packet and are on the list reply with acknowledgments that (optionally) contain their link quality information, which is then used by the source to choose the best next hop. Some later proposals, such as [7] (link-layer anycast), [8] (geographic random forwarding), and [9] (hybrid-ARQ retransmissions), can be seen as variations on the same theme. Another interesting variation is ExOR [10], where the source does not choose a single next hop at all; rather, packets are transmitted unacknowledged in large batches, where upon completion of a batch, each neighbour forwards the packets it has successfully heard, after waiting a random delay (that depends on its position in a predetermined priority list) and skipping any packets it could overhear being forwarded by another neighbour in the meantime.

While opportunistic forwarding methods work well in multi-hop (and especially dense) network settings, their excessive overheads due to the selection process of the next-hop neighbour (or the additional random wait in the case of [10]) make them unsuitable for delay-critical applications in a single-hop

*Lixiang Xiong is supported by Australian Research Council discovery grant, grant number: DP0559248.

†NICTA is funded by the Australian Government's *Backing Australia's Ability* initiative, in part through the Australian Research Council.

1525-3511/08/$25.00 ©2008 IEEE

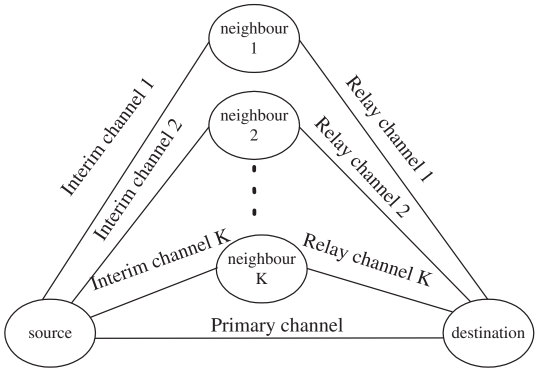

Fig. 1. The wireless network considered in this paper.

setting. Indeed, consider a simple setting of a sender, receiver, and $K > 1$ neighbours that can serve as potential relays (Figure 1). Assume a symmetrical system where all neighbours have identical and independent *a priori* interim and relay channel characteristics. If the direct transmission from the source to destination fails, any of the aforementioned opportunistic forwarding methods will iterate over the neighbours before choosing one to forward the packet, resulting in a high packet latency of $O(K)$ slots. On the other hand, a strategy where any neighbour overhearing the packet simply retransmits it immediately in the next time slot — bypassing any attempt to agree first on a single node to do the retransmission, and instead handling any resulting collisions with some collision resolution scheme — may be able to achieve a lower expected latency, particularly in a setting where the individual overhearing probability is low for each neighbour, yet the aggregate probability of overhearing by at least one neighbour is reasonably high.

We point out that a latency analysis of the network of Figure 1 with cooperative retransmissions has been undertaken in the past. In [11], a fixed time-division multiple-access (TDMA) scheme is assumed, so that any node overhearing a neighbour's unsuccessful packet may retransmit that packet in its own allocated slot provided the queue of its own packets is empty, and a latency analysis is conducted based on independent Poisson arrivals. In [12], the system is assumed to operate in a stop-and-wait regime with neighbours continuously retransmitting overheard packets until acknowledgment, and the channels are modelled as independent, 2-state (on-off) Markov chains, such that a single 'good' channel from any retransmitting neighbour suffices to decode the packet at the destination. Both of these works sidestep the possibility of *collision* among the cooperative neighbours' retransmissions: in [11], such collisions cannot occur by virtue of the fixed TDMA allocation, while in [12], it is implicitly assumed that the packet can be decoded in the presence of multiple simultaneous retransmissions, thereby requiring a cooperative diversity-enabled receiver.

Motivated by the above observations, in this paper we consider cooperative retransmission by uncoordinated neighbours (that do not attempt to agree on one forwarder for a packet, unlike the opportunistic forwarding methods), and focus specifically on simple single-antenna receivers, for which multiple simultaneous retransmissions will result in a collision; as a result, clearly it may not be best for all neighbours to always retransmit their overheard packets. Accordingly, we define a cooperation *strategy* as a sequence of retransmission probabilities to be followed by each neighbour in subsequent slots after the source's original transmission. We pose the problem of finding a strategy that minimizes the expected latency of a packet (i.e. the expected number of slots until the packet is successfully received at the destination, without collision), and explore some possible approaches for its solution.

It is important to emphasize that, while it is relatively easy to calculate the optimal retransmission probabilities for a *single* slot (i.e. to maximize the likelihood of successful packet reception after a single retransmission attempt), the problem of finding optimal strategies for *subsequent* slots is far more challenging. Our contribution towards its solution is twofold. First, we explore the subset of strategies where the cooperative neighbours keep retransmitting in subsequent slots exclusively, while the original sender is silent. For such strategies, we derive the expected packet latency analytically and are able to characterize the optimum precisely. We then remove the constraint for the original sender to remain silent during the cooperative retransmissions, allowing the retransmission probabilities of both the original sender and the neighbours to be nonzero simultaneously. For such strategies, computing an analytical optimal solution is difficult. Instead, we propose a particular heuristic, based on an iterative Bayesian estimation of the distribution of neighbours possessing a copy of the packet. We study its performance in detail and demonstrate that it attains a significantly lower expected packet latency than either traditional routing or opportunistic forwarding approaches, for a broad range of channel quality settings.

The rest of this paper is organized as follows. Section II introduces the system model and discusses the assumptions made in our work. Section III derives the optimal strategy for a single retransmission attempt, and Section IV explains our approaches for extending the retransmission strategy to subsequent time slots. Section V presents a numerical study to evaluate and compare the performance of the proposed strategies. Finally, the paper is concluded in Section VI.

II. System model and Assumptions

We consider a wireless network consisting of a source node, a destination node, and a fixed number K of cooperative neighbour nodes in their vicinity (Figure 1). When the source transmits a packet, it may arrive successfully over the primary (direct) channel with some probability, and may also be overheard by some of the neighbour nodes via the interim channels. Only the intended destination acknowledges the packet upon successful reception; neither the source nor any neighbour can tell which other neighbours, if any, have also obtained a copy of the packet. We define the time from a packet transmission to the completion of its acknowledgment as a *slot*; it is assumed that slots are of fixed duration and synchronized at all the nodes. In the subsequent slots after the original transmission, any node possessing a copy of the packet may decide to make a cooperative retransmission over its relay channel. For successful reception, the destination must receive exactly one error-free transmission in a slot, without collision with other simultaneous transmissions. We assume that one packet is active at a time, i.e. no other packets are transmitted between the first transmission by the source until the eventual delivery to the receiver (or, possibly, until reaching a maximum number of retransmission attempts); we further assume, for simplicity, that the feedback (acknowledgment) channels to the source and all neighbours are error-free. Thus, our model is similar to the node-cooperative stop-and-wait (NCSW) setting of [12], with the notable difference in our case being the possibility of collision if multiple retransmissions occur at the same time.

We define a cooperation *strategy* of a node to be a sequence of probabilities for retransmitting the packet copy. The elements of the sequence correspond to time slots, and its length should be equal to the maximum number of slots after which the packet is dropped (infinite if a 100% reliability is required). Thus, at every time slot, each node that had overheard a copy of the packet transmits it with the probability dictated by the strategy for that slot. We define an *optimal* strategy profile as a collection of individual node strategies that minimizes the expected number of slots until successful reception of the packet by the destination. We emphasize that a strategy consists of a fixed sequence of numbers, and the transmission probabilities may not depend on any earlier events at the node. The reason for this limitation is to allow a distributed implementation of the strategies; this requires every node's transmission probabilities to be known to all other nodes, and therefore, be independent of any information that is only known locally. For example, the source node may not base its transmission decision in the third slot on whether it overheard any cooperative retransmissions in the second slot, as that information is not known to its neighbours, and they will not be able to compute their own optimal strategies in that slot.

While the above definitions can be interpreted generically, in this paper we use a concrete propagation model to study cooperation strategies. We assume a simple fading model, similar to [12], in which any channel can be in one of two states: either "on" (no fade), in which the transmitted signal arrives with sufficient power to be decoded without error (barring a collision), or "off" (deep fade), 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). Finally, we assume that channel states between different pairs of nodes are mutually independent, which is realistic in most practical scenarios with nodes spaced sufficiently far apart.

In this paper, in the interest of simplicity, we analyse the optimal strategy problem and present our framework for its solution only for memoryless channels, where the probability of being in the "on" state in a transmission slot is fixed and independent of the state in previous slots. We denote this probability by P_{sd} for the primary channel (between source and destination); P_{sn} for any of the interim channels (between source and neighbour); and P_{nd} for any of the relay channels (between neighbour and destination). We make two observations regarding this model. First, even though the work that led to this paper has also considered the more sophisticated Gilbert model, where the transitions between states are dictated by a Markov process (as in [12]), we are unable to present the full results here due to space constraints. Nevertheless, the optimal strategy problem is sufficiently non-trivial even for memoryless channels that its solution approaches, and the results and insights obtained from their evaluation (explained in section V), remain qualitatively valid even when extended to the Gilbert model. Second, our model assumes a symmetric system, where all neighbours are equivalent *a priori*. This does not mean that the underlying physical characteristics of all neighbours (e.g. their distances from the source and destination) must be identical; it only means that they are indistinguishable for the purpose of computing the optimal strategy. Alternatively, an asymmetric model, assigning individual (and generally different) probabilities to each interim and relay channel, might seem more realistic; however, it would only be relevant if the entity computing the optimal cooperation strategies were informed of these different probabilities. This may be feasible in scenarios such as fixed mesh networks with powerful nodes, which know their locations and/or are able to measure their long-term channel statistics and recalculate the optimal strategies upon any changes; however, in most typical scenarios involving ad-hoc networks of small mobile nodes, it is unfeasible for them to be aware of their location and distance from the destination, or afford to measure the channel statistics upon every change in their neighbour set. We leave the asymmetric extension of our model to future work.

III. The First Retransmission

We begin our analysis by considering the optimal cooperative retransmission strategy for a single time slot. Thus, the strategy simply boils down to a single number, namely the probability of retransmission for any node that had successfully overheard the packet; we denote this probability by τ. The optimal strategy, then, is the value of τ that maximizes the probability of successful delivery. In order to find it, we first consider the probability distribution of the number of neighbours k that have successfully overheard the packet from the source's original transmission. Since the interim channels are symmetric, this distribution is binomial:

$$P\{k\} = \binom{K}{k} P_{sn}^k \left(1 - P_{sn}\right)^{K-k}. \tag{1}$$

For a successful delivery, out of these k nodes, there must be exactly one that both makes a retransmission *and* has a good channel. Consequently,

$$P^{suc} = \sum_{k=1}^{K} P\{k\} \cdot k\tau P_{nd}(1 - \tau P_{nd})^{k-1}, \tag{2}$$

which, after a straightforward simplification, becomes

$$P^{suc} = KP_{sn}\tau P_{nd}(1 - P_{sn}\tau P_{nd})^{K-1}. \tag{3}$$

The optimal τ^* is now obtained by equating the first derivative of (3) to zero, which yields

$$\tau^* = \frac{1}{KP_{sn}P_{nd}}. \tag{4}$$

The above expression for τ^*, of course, is only valid if $\frac{1}{KP_{sn}P_{nd}} \leq 1$. Otherwise, i.e. if $P_{sn}P_{nd} < \frac{1}{K}$, the probability of successful delivery (3) is monotonically increasing in τ, and its optimum is then achieved with $\tau^* = 1$.

We now assign the optimal strategy τ^* back into (3), to evaluate the maximum success probability that can be obtained after one cooperative retransmission slot. If $K > \frac{1}{P_{sn}P_{nd}}$, then τ^* is given by (4), and

$$P^{suc^*} = \left(1 - \frac{1}{K}\right)^{K-1}. \quad (5)$$

Curiously, we observe that this expression is *decreasing* in K (it tends to $\frac{1}{e}$ for $K \to \infty$); in other words, having too many neighbours in the cooperation group may, in fact, degrade the performance of cooperative retranmission. It is easily verified that, in the range $1 \leq K < \frac{1}{P_{sn}P_{nd}}$ (such that $\tau^* = 1$), the probability of successful delivery is increasing in K, as intuitively expected. Hence, we conclude that the best size of the cooperation group is around $\frac{1}{P_{sn}P_{nd}}$; if the number of neighbour nodes is larger than that, it is better to voluntarily choose a smaller cooperation group (and thereby keep τ^* close to 1), rather than use all the available neighbours with a smaller retransmission probability.†

IV. Subsequent Retransmissions

A. Strategy 1: neighbours repeat attempts with silent source

We now consider the cooperation strategies beyond the first retransmission slot. We first focus on the strategy where, if the first attempt fails, the neighbours continue to make additional retransmission attempts while the source remains silent. Since the source does not transmit again, the number of neighbours with a copy of the packet k, and its distribution, remains unchanged from the first slot; consequently, the optimal τ^* that maximizes the probability of exactly one neighbour transmitting with a good relay channel is the same as for the first slot, derived in section III.‡ Since we assume the channel states to be independent between slots, the retransmission attempts will form a Bernoulli process with a success probability of

$$P^{suc}_{k,\tau^*} \triangleq k\tau^* P_{nd}(1 - \tau^* P_{nd})^{k-1} \quad (6)$$

in each slot.

However, due to the possibility of the case $k = 0$ (i.e. all interim channels were "off" during the original transmission and no neighbour overheard the packet, in which case $P^{suc}_{k,\tau^*} = 0$), the above strategy is not guaranteed to succeed after a finite number of attempts. Therefore, we define a maximum number of cooperative attempts before the retransmission process restarts again with the original source node, and denote it by $m - 1$. Thus, we are considering a periodic strategy with a period of m slots, where each period starts with a transmission by the source, followed by $m - 1$ cooperative retransmissions by the neighbours. The choice of m reflects a tradeoff between time wasted on cooperative attempts in the case of $k = 0$ and that wasted on a source retransmission otherwise. Generally, the better the interim channels (P_{sn}) and the worse the relay channels (P_{nd}), the higher the optimal value of m.

To find the optimal m analytically, we write the following recursive expression for the expected number of slots until success, E:

$$E = P_{sd} \cdot 1 + (1 - P_{sd}) \sum_{k=0}^{K} P\{k\} \cdot \left[\sum_{i=1}^{m-1} P^{suc}_{k,\tau^*} \left(1 - P^{suc}_{k,\tau^*}\right)^{i-1} \cdot (i+1) + \left(1 - P^{suc}_{k,\tau^*}\right)^{m-1} \cdot (E + m) \right]. \quad (7)$$

This expression accounts for the probability of success in the direct transmission over the primary channel, or after $i \in \{1, \ldots, m-1\}$ attempts of cooperative retransmission in the first period, or, if all $m-1$ such attempts prove unsuccessful, the entire first period of m slots is wasted and the expected number of additional slots until success is the same as originally. Using the geometric-sum formula $\sum_{i=1}^{N} ir^{i-1} = \frac{1-r^{N+1}}{(1-r)^2} - \frac{(N+1)r^N}{1-r}$ and grouping together the coefficients of E (we omit the straightforward details to save space), we arrive at the formula

$$E = \frac{P_{sd} + (1 - P_{sd}) \sum_{k=0}^{K} P\{k\} \left[\frac{1 - \left(1 - P^{suc}_{k,\tau^*}\right)^{m-1}}{P^{suc}_{k,\tau^*}} + 1 \right]}{1 - (1 - P_{sd}) \sum_{k=0}^{K} P\{k\} \left(1 - P^{suc}_{k,\tau^*}\right)^{m-1}}, \quad (8)$$

where $P\{k\}$ is given by (1), and the expression in brackets in the numerator for $k = 0$ (i.e. $P^{suc}_{k,\tau^*} = 0$) should be taken as equal to m. In the subsequent performance evaluation in Section V, we use (8) to manually find the optimal period m for this strategy, for any instance of P_{sd}, P_{sn}, P_{nd}, and K.

B. Strategy 2: simultaneous source+neighbour transmissions

In the strategy described in the previous subsection, the parameter m reflected the tradeoff between extending the chance to retransmissions by the cooperative neighbours (which normally have better channels to the receiver), and wasting the time in case no neighbours had overheard the packet (which, as a direct consequence of the fact that the sender is silent, is not rectified during the entire $m - 1$ slots). In order to overcome this disadvantage, we now describe a heuristic strategy in which the transmission probabilities of both the sender (τ_s) and the neighbours (τ_n) are allowed to be greater than zero simultaneously. As a result, the number of neighbours with a copy of the packet continues to grow over time (up to K).

The heuristic is based on a greedy approach that attempts to maximize the probability of successful reception in each slot in turn. To assist in the calculation, we maintain and update the distribution $P\{k\}$, i.e. the number of neighbours with a copy of the packet so far. Thus, in every slot i, the following calculation steps are made:

1) the optimal τ_s and τ_n are solved numerically to maximize the probability of success in this slot, given by the

†Since $\frac{1}{P_{sn}P_{nd}}$ is, in general, not a whole number, the optimal cooperation group size may be the integer to either side of it.

‡In the analysis of strategy 1, we ignore the possibility of overhearing the packet from another neighbour's transmission; we consider this possibility later in the discussion of strategy 2.

expression

$$P^{suc} = \sum_{k=0}^{K} P_i\{k\} \cdot P^{suc}_{k,\tau_s,\tau_n}, \quad (9)$$

where $P_i\{k\}$ denotes the distribution of k *before* the start of slot i, and

$$P^{suc}_{k,\tau_s,\tau_n} \triangleq (1-\tau_s P_{sd})k\tau_n P_{nd}(1-\tau_n P_{nd})^{k-1} + \tau_s P_{sd}(1-\tau_n P_{nd})^k \quad (10)$$

2) assuming that the slot nevertheless results in a failure, the distribution $P_i\{k\}$ is revised *a posteriori* using Bayes' formula, as follows:

$$P^{rev}_i\{k\} = \frac{P_i\{k\}(1-P^{suc}_{k,\tau^*_s,\tau^*_n})}{\sum_{k'=0}^{K} P_i\{k'\}(1-P^{suc}_{k',\tau^*_s,\tau^*_n})}, \quad (11)$$

where τ^*_s, τ^*_n are the optimal strategy values obtained in step 1;

3) finally, $P_i\{k\}$ is updated to account for the new neighbours that overhear the packet from the source in this slot, yielding

$$P_{i+1}\{k\} = (1-\tau^*_s)P^{rev}_i\{k\} + \tau^*_s \sum_{k'=0}^{k} P^{rev}_i\{k'\} \cdot \binom{K-k'}{k-k'} P^{k-k'}_{sn}(1-P_{sn})^{K-k+k'}. \quad (12)$$

Note that expression (12) considers only the possibility of overhearing a transmission from the source, not from another neighbour. One can also consider the case where a channel between two neighbours can be "on" with a probability $P_{nn} > 0$; then, a new neighbour may overhear the packet from either the source or another neighbour, provided there is no collision. The extension of (12) to this case is straightforward and omitted here.

Example: Consider a network with only $K = 1$ cooperating neighbour, $P_{sd} = 0.5$, $P_{sn} = 0.99$, $P_{nd} = 1$. If the first transmission by the source fails, the neighbour has a probability of $P_2\{k=1\} = 0.99$ to have the packet at the start of the second slot. Therefore, clearly, the optimal strategy in this slot is to allow it to transmit the packet uninterrupted ($\tau^*_s = 0, \tau^*_n = 1$). Indeed, a simultaneous transmission by the source would interfere with the neighbour's one with a probability of $0.99 \cdot 0.5$, and would only be helpful with a probability of $0.01 \cdot 0.5$. However, if this strategy is applied and still fails, then we have $P^{rev}_2\{k=1\} = P_3\{k=1\} = 0$, as failure can only occur if the neighbour did not have the packet after all. Accordingly, the optimal strategy in the third slot is $\tau^*_s = 1$ (the strategy of the neighbour is immaterial).

However, in the same system but with $P_{sd} = 0$, the optimal strategy trivially becomes $\tau^*_s = 1, \tau^*_n = 1$ in every slot. The source does not have a channel to the receiver and therefore cannot interfere with the neighbour; meanwhile, the simultaneous transmission by the source saves time if the neighbour still has not heard the packet. □

V. Evaluation of the Cooperation Strategies

In this section, we investigate the performance of the proposed strategies numerically, under various combinations of channel quality for the direct, interim and relay channels. In each scenario, we examine the impact of the number of cooperative neighbours on the expected latency, and compare it with traditional one-hop and two-hop routing as well.

We begin with an arguably typical cooperative retransmission scenario: a poor primary channel, with better interim and relay channels. Accordingly, we demonstrate the performance of our strategies with $P_{sd} = 0.1$, $P_{sn} = P_{nd} = 0.5$. In this scenario, retransmissions over the direct hop require on average $\frac{1}{P_{sd}} = 10$ slots until success, while two-hop routing over any of the neighbours (with retransmission in each hop) achieves an average latency of $\frac{1}{P_{sn}} + \frac{1}{P_{nd}} = 4$ slots.

The results are shown in Figure 2(a). The performance of Strategy 1 is obtained with the optimal period m (manually found using expression (8) for each K). In addition, we test the performance of Strategy 2 under three values of P_{nn}, i.e. the neighbour-to-neighbour channel quality (see the comment following expression (12)). These range from $P_{nn} = 0$ (neighbours cannot overhear each other at all), to $P_{nn} = 1$ (neighbours always overhear transmissions from their peers). As Figure 2(a) clearly shows, even though the performance of Strategy 1 improves as the number of cooperative nodes increases and it attains a lower expected latency than two-hop routing for $K \geq 3$, Strategy 2 outperforms it consistently. We also observe that the impact of the neighbour-to-neighbour channel quality on the overall performance is negligible; intuitively, the ability of a neighbour to overhear the packet from other neighbours (and not just from the source) is counterbalanced by the additional collisions that occur when the source and another neighbour transmit together.

Figure 2(b) shows the results for the case where the quality of the direct channel is improved to $P_{sd} = 0.3$. Now, the direct channel is better than the two-hop route, requiring only $\frac{1}{0.3} \approx 3.33 < 4$ slots. Nevertheless, our cooperative retransmission strategies are still able to significantly improve the expected latency, mainly because they take advantage of the better relay channel when the original transmission over the direct channel fails. Of course, as the direct channel becomes better, this effect diminishes; once $P_{sd} = 0.5$ on par with the interim and relay channels, the optimal strategy trivially reduces to retranmission over the direct channel, ignoring the neighbours.

In Figures 2(c) and 2(d), we examine the effect of reducing the interim channel quality to $P_{sn} = 0.3$ and $P_{sn} = 0.1$, respectively; Figures 2(e) and 2(f) do the same for the relay channel. As expected, the performance of all strategies deteriorates monotonically with the channel quality; nevertheless, there is merit in using cooperative retransmission as long as either the interim or relay channel is better than the direct one. Interestingly, these figures also demonstrate that the same effect we observed in the analysis of the first slot — namely, that the optimal number of cooperating neighbours increases as the quality of the channels deteriorates — holds for the overall performance of the greedy heuristic strategy as well.

We point out that we have conducted the performance evaluation on a wide variety of scenarios, and only presented a small representative sample here due to space limits. However, our observations indicate that the main effects seen in Figure 2 — namely, that Strategy 1 improves with increasing neighbour population, that Strategy 2 consistently outperforms it (and is never worse than one-hop or two-hop routing), and that

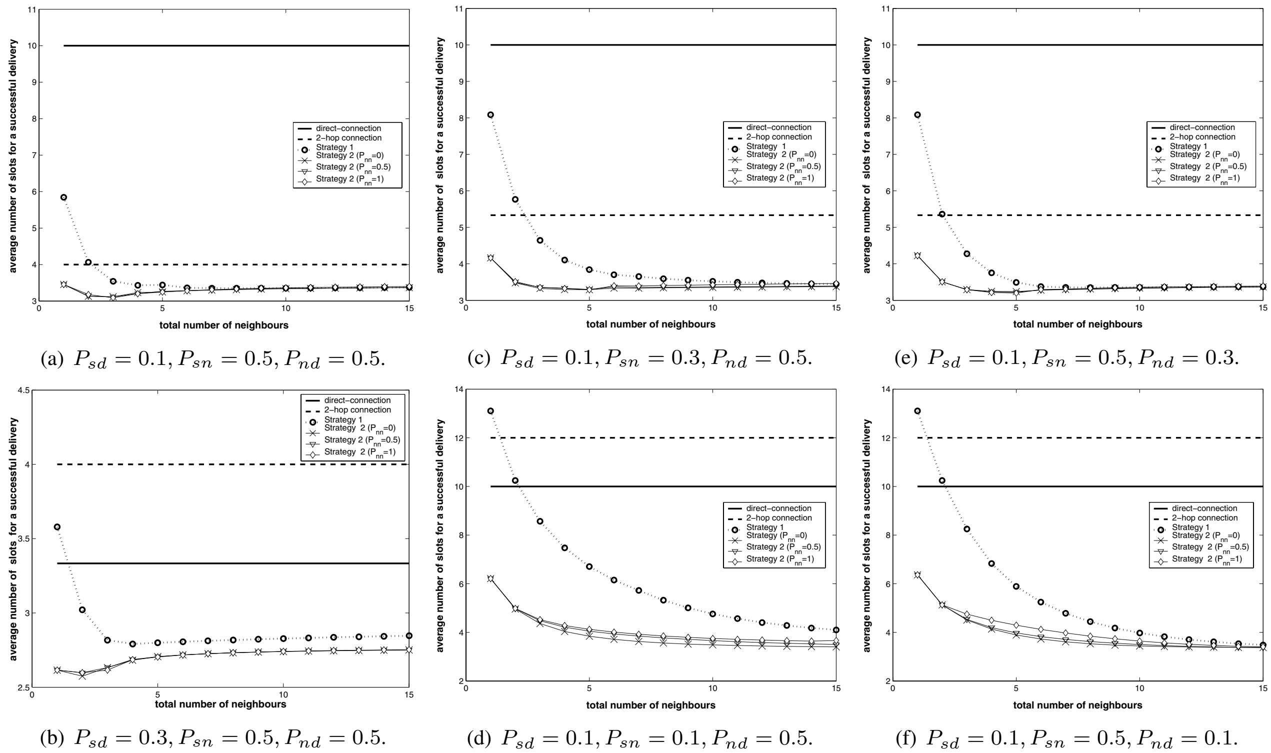

(a) $P_{sd} = 0.1, P_{sn} = 0.5, P_{nd} = 0.5$. (c) $P_{sd} = 0.1, P_{sn} = 0.3, P_{nd} = 0.5$. (e) $P_{sd} = 0.1, P_{sn} = 0.5, P_{nd} = 0.3$.

(b) $P_{sd} = 0.3, P_{sn} = 0.5, P_{nd} = 0.5$. (d) $P_{sd} = 0.1, P_{sn} = 0.1, P_{nd} = 0.5$. (f) $P_{sd} = 0.1, P_{sn} = 0.5, P_{nd} = 0.1$.

Fig. 2. Numerical results.

the neighbour-to-neighbour channel quality has a very minor impact on the performance — occur consistently throughout our evaluations.

VI. Conclusion

We have studied the problem of optimal MAC-layer cooperative retransmission strategies in single-hop wireless networks, which employ probabilistic retransmission by neighbours overhearing the original packet to minimize the expected packet latency. We considered two kinds of strategies; one where the cooperative neighbours retransmit exclusively while the original sender is silent, and another where both may have a non-zero transmission probability simultaneously. We derived analytically the best strategy possible in the first kind. For the second kind, we considered a heuristic that combines a greedy maximization of success probability in each slot with a Bayesian re-estimation of the distribution of the number of neighbours with a copy of the packet following every slot. We demonstrated that, in general, the latter strategy achieves a superior performance and considerably reduces the expected time until the packet is successfully received. Our results were obtained for memoryless channels, but can be extended in a straightforward manner to a Markov (e.g. Gilbert) channel model as well. The extension of our solution approaches to the case of asymmetric neighbours is left for future work.

References

[1] A. Nosratinia, T.E. Hunter, and A. Hedayat. Cooperative communication in wireless networks. *IEEE Communications Magazine*, 42(10):74–80, October 2004.

[2] L. Lai, K. Liu, and H. El Gamal. The three-node wireless network: Achievable rates and cooperation strategies. *IEEE Transactions on Information Theory*, 52(3):805–828, March 2006.

[3] A. Sendonaris, E. Erkip, and B. Aazhang. User cooperation diversity (parts I,II). *IEEE Transactions on Communications*, 51(11):1927–1948, November 2003.

[4] J.N. Laneman, D.N.C. Tse, and G.W. Wornell. Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *IEEE Transactions on Information Theory*, 50(12):3062–3080, December 2004.

[5] B. Aazhang, R.S. Blum, J.N. Laneman, K.J.R. Liu, W. Su, and A. Wittneben. Guest editorial: Cooperative communications and networking. *IEEE Journal on Selected Areas in Communications*, 25(2):241–244, February 2007.

[6] 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.

[7] 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)*, pages 22–30, Taormina, Italy, June 2005.

[8] 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.

[9] B. Zhao and M.C. Valenti. Practical relay networks: A generalization of hybrid-ARQ. *IEEE Journal on Selected Areas in Communications*, 23(1):7–18, January 2005.

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

[11] 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.

[12] 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.