2014 International Conference on Computing, Networking and Communications, Network Algorithm & Performance Evaluation Symposium

Optimal Routing for Bidirectional Flows with Network Coding in Asymmetric Wireless Networks

Hooman Reisi Dehkordi Lavy Libman
School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia
{*hreisi,llibman*}@*cse.unsw.edu.au*

***Abstract*—*Wireless network coding* has attracted considerable research interest due to its promise of delivering increased capacity with high energy efficiency via the transmission of packet combinations from multiple flows. Much of this interest is based on the approach of the COPE protocol [1], which uses broadcasts of XOR combinations of packets to reduce the number of transmissions required, as long as they can be immediately decoded at the next hop of the respective flows. In this paper, we focus on the question of optimal routing of a bidirectional flow between a single pair of endpoints in the presence of wireless network coding, where the links are imperfect and have positive, and possibly asymmetric, packet erasure probabilities. Unlike previous studies, we explicitly allow packets to be forwarded in the network in a coded form, to be ultimately decoded by the destinations. We pose the optimization problem of finding the optimal path in each direction, so as to minimize the expected overall number of transmissions required to successfully deliver one packet in each direction of the flow, and present a polynomial-time algorithm for its solution. In particular, we demonstrate that, counterintuitively, it is possible for the optimal pair of paths for the bidirectional flow to have an overlapping section where packets are transmitted in the *same* direction, and illustrate the importance of this possibility by evaluating the performance of optimal network coding paths in random and asymmetric mesh topologies of between 9 and 36 nodes.**

I. Introduction

Wireless network coding has been proposed in recent years as a promising technique that takes advantage of the inherent broadcast nature of the wireless medium and achieves a high capacity with a reduced number of transmissions. The idea was first proposed by the COPE protocol [1], which showed that, in a three-node network consisting of two endpoints and an intermediate relay, the number of transmissions required to communicate one packet in each direction can be reduced from 4 to 3, with the relay transmitting a XOR combination of both packets that can be then decoded at the respective destinations. The same idea can be used more generally when multiple flows intersect at a common intermediate node; then, a single transmission of a XOR combination of packets from all flows can be used, provided that the next hop of each flow has been able to overhear the packets from the other flows' respective sources. Since originally proposed, the COPE protocol has become a subject of immense attention, with a large volume of literature devoted to its performance analysis and various extensions and improvements [2].

The original COPE proposal focused on the network coding action in each node and did not consider its interaction with routing, which was assumed to be given by a separate protocol. The need to account for the potential coding opportunities in the routing protocol was first pointed out in [3], which introduced a wireless link metric based on expected number of coded transmissions (ECX). A more thorough discussion of the desirable properties of coding-aware routing metrics was provided in [4], together with a distributed mechanism for discovering potential network coding mechanisms on the available paths. Subsequently, the work in [5] developed a complete on-demand coding-aware routing protocol (OCAR), based on a metric that balances between the desire to maximize coding opportunities among flows (thereby attracting routing paths towards the same nodes) and to minimize interference (which pushes routes away from each other). From a theoretical perspective, the most detailed treatment of optimal coding-aware routing to date is given in [6], where the trade-off between coding opportunities and interference is formulated as a linear programming problem. The main limitation of all of the above studies, due to their focus on routing of multiple simultaneous flows with different source-destination pairs, is that they only consider coding opportunities where the combined packet can immediately be decoded into the original native packets by each of the next hops on every flow. In other words, simple forwarding of a coded packet over more than one hop is not permitted.

In this paper, we focus on the case of COPE applied to a single bidirectional flow between a pair of endpoints; consequently, we allow network coding to occur at any point in the network, and the resulting packet to be subsequently routed (in coded form) to the destinations and ultimately decoded there, rather than decoded immediately at the next hop. We note that, for this case, it has been shown that a coding gain (i.e. ratio between the number of transmissions required with and without network coding) of 2 can be achieved in a chain topology, when the number of hops tends to infinity [1]. This limit, however, is achievable only under the assumption that the individual network links are perfect, i.e. transmitted packets are never lost. In this work, we are interested in optimal coding-aware routing of a bidirectional flow in the presence of lossy links, and especially when the loss rates in the links are asymmetric. In the latter case, the optimal routing may not necessarily correspond to a bidirectional chain, i.e. the optimal paths may not consist of the same set of links traversed in opposite directions. In fact, we will later show that the optimal pair of paths may share some overlapping links in the *same* direction, where the packets from both sides are routed together in combined form (see Figure 1d). This fact highlights that any optimal coding-aware routing algorithm, which aims to take advantage of overlapping between the forward and reverse paths, must account for the possibility that coded packets may be used in path sections that overlap either in coinciding or in opposing directions.

Our contribution is as follows. For a general wireless network with a pair of endpoints and any number of intermediate nodes, with given erasure probabilities of packets on each link, we define the problem of finding the pair of paths that minimizes the expected number of transmissions required to transmit one packet in each direction. We then derive an explicit formula for the optimization target, and present a polynomial-time solution algorithm to the above optimization problem, which follows an approach similar to the Floyd-Warshall shortest path algorithm while enumerating over the possible nodes between which the shared overlapping section of the pair of paths may run. We evaluate the algorithm, and the reduction in the required number of transmissions enabled by overlapping paths and forwarding of packets in coded form, using simulation of random mesh topologies consisting of up to 36 nodes.

The rest of this paper is structured as follows. Section II defines the network model and assumptions and formally states the optimization problem. The optimal routing solution algorithm is presented in section III, and evaluated in simulation using random mesh network topologies in Section IV. Finally, Section V concludes the paper.

II. Network Model and Problem Definition

We consider a wireless network with two endpoints, henceforth denoted by A and B, attempting to send a data packet to each other. We assume that A and B are out of each other's direct communication range, and their packets must be delivered in a multi-hop fashion via a number of intermediate relays, denoted by $R_1, R_2, \ldots$. The wireless links in the networks are lossy, such that a packet transmission from a node u can be successfully received by another node v with a probability $P_{u,v}$. In other words, the channel between any two nodes is a packet erasure channel with an erasure probability of $1 - P_{u,v}$, or, equivalently, the link has an ETX (expected transmission count) metric of $\frac{1}{P_{u,v}}$. These probabilities are considered to be independent both among the links and among successive transmissions on the same link. We emphasize that links may be asymmetric, i.e. we allow $P_{u,v} \neq P_{v,u}$ for any pair of nodes u, v, and that some of the channel probabilities may be set to zero (corresponding to pairs of nodes that do not have a link at all); thus, the model represents a general network topology and is not limited to a two-hop setting.

With the above model, we are interested to find a pair of paths, from A to B (the 'forward' path) and from B to A (the 'reverse' path), such that the total expected number of transmissions before the packets are successfully delivered to their destinations is minimized. This is not equivalent to finding the shortest path in each direction (with respect to the ETX metric), due to the possibility of network coding; namely, if a relay node is common to both of the paths, then a coding opportunity is created and both the forward and reverse packets can be handled with a single transmission. For brevity, we henceforth refer to the expected number of transmissions associated with a given pair of paths as the *cost* of those paths.

We now proceed to illustrate the considerations arising in the calculation of the cost of a pair of paths using the example of a network with two intermediate relay nodes. These will be subsequently applied towards the generic problem definition and path cost formula at the end of this section.

In a network where the endpoints A, B communicate via two relays R_1, R_2, there are four possible paths for packets from A to B (and, similarly, from B to A):

1) $A \to R_1 \to B$;
2) $A \to R_2 \to B$;
3) $A \to R_1 \to R_2 \to B$;
4) $A \to R_2 \to R_1 \to B$.

As the links can be asymmetric, the best paths in the opposite directions may differ, hence any of the 16 possible combinations of forward and reverse routes may be optimal in this network. Some of these combinations are shown in Figure 1. In the case of Figure 1a, the paths are completely disjoint and there are no network coding opportunities, while in other cases such as in Figure 1c, the forward and reverse paths include two nodes in common, either of which can perform the initial coding. We now compute the total expected number of transmissions, i.e. the cost of the paths, for each of the cases illustrated in Figure 1.

In the simplest case of Figure 1a, no network coding is possible, and the cost consists simply of the sum of ETX metrics of all the links involved:

$$\frac{1}{P_{A,R_1}} + \frac{1}{P_{R_1,B}} + \frac{1}{P_{B,R_2}} + \frac{1}{P_{R_2,A}}. \qquad (1)$$

In the case of Figure 1b, the forward and reverse paths have a common node of R_1; thus, A and B send their packets to R_1, and R_1 transmits their coded combination until successfully received by both A and B. Therefore, the number of transmissions saved by the use of network coding is equal to the expected number of transmissions until one of the destinations receives the packet successfully (further transmissions of the coded packet thereafter are equivalent to transmissions of the respective native packet). Thus, the total cost of the paths amounts to

$$\frac{1}{P_{A,R_1}} + \frac{1}{P_{R_1,B}} + \frac{1}{P_{B,R_1}} + \frac{1}{P_{R_1,A}} - \frac{1}{1-(1-P_{R_1,A})(1-P_{R_1,B})} \qquad (2)$$

The case of Figure 1c is similar, except that now there are two relay nodes that are common between the two paths, and the initial coding (combining) of the two packets can be performed in either of them. Following a similar reasoning as above, we obtain the cost associated with this scenario:

$$\frac{1}{P_{A,R_1}} + \frac{1}{P_{R_1,R_2}} + \frac{1}{P_{R_2,B}} + \frac{1}{P_{B,R_2}} + \frac{1}{P_{R_2,R_1}} + \frac{1}{P_{R_1,A}} - \begin{cases} \frac{1}{1-(1-P_{R_1,R_2})(1-P_{R_1,A})} & \text{if } R_1 \text{ is doing NC} \\ \frac{1}{1-(1-P_{R_2,B})(1-P_{R_2,R_1})} & \text{if } R_2 \text{ is doing NC} \end{cases} \qquad (3)$$

Using a similar reasoning, one can calculate the costs of all other path combinations, where the combined packet is

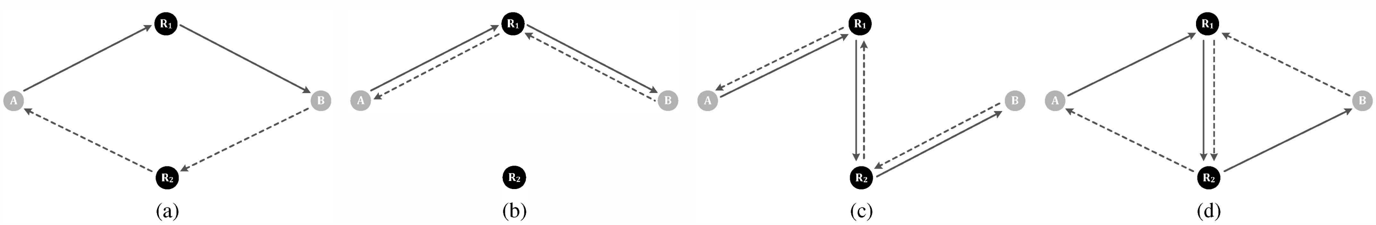

Fig. 1: Different path combinations in a two-relay network.

transmitted to two different next hop nodes. The exceptional case arises in Figure 1d, where the two paths share an overlapping section in the *same* direction. The gains from network coding in this case (which, to the best of our knowledge, has not been observed in existing literature) are twofold. First, the coded packet (which is generated by R_1, i.e. at the start of the shared section) only needs to be communicated over the overlapping section once rather than twice. In addition, where the paths eventually split (i.e. at R_2 in this example), the transmission of the coded packet generates a saving similar to that in Figures 1b–1c, i.e. equal to the expected number of transmissions until one of the destinations receives the coded packet. The total cost of both paths is therefore

$$\frac{1}{P_{A,R_1}} + \frac{1}{P_{R_2,B}} + \frac{1}{P_{B,R_1}} + \frac{1}{P_{R_2,A}} + \frac{1}{P_{R_1,R_2}} - \frac{1}{1-(1-P_{R_2,B})(1-P_{R_2,A})}, \quad (4)$$

emphasizing that $\frac{1}{P_{R_1,R_2}}$ is only counted once for both paths.

We now generalize the above observations to a network with any number of intermediate relays between the endpoints A and B. To that end, we divide the links forming the paths into the following three distinct categories, representing the different ways in which the cost is affected by network coding:

- The *independent* category consists of those portions of both paths where the transmissions to deliver the two packets are made separately, e.g. the links $A \to R_1$ and $B \to R_1$ in Figure 1d. Links in this category do not count for any savings due to network coding. We use $\mathcal{I}_a$ and $\mathcal{I}_b$ to denote the sets of independent links corresponding to the forward and reverse paths, respectively.
- The *shared* category consists of the links that are common to both paths (and in the same direction), which only need to be used once for both paths with the help of coded transmissions, e.g. the link $R_1 \to R_2$ in Figure 1d. The set of links belonging to this category is denoted by $\mathcal{S}$; this set may be empty if the two paths do not have any links in common.
- The *point-of-split* category consists of precisely the two links from the last node of the shared section (if any) to two different nodes, from where the remainders of the two paths are independent. In the example of Figure 1d, these are the links $R_2 \to A$ and $R_2 \to B$. At the point of split, the cost reduction due to network coding is the expected number of transmissions until one of the two links is successful for the first time. We denote the two links by $\mathcal{L}_a$ and $\mathcal{L}_b$, corresponding to the forward and reverse paths respectively. If the forward and reverse paths are entirely node-disjoint, this category is undefined.

We illustrate the categories further using a generic example in Figure 2, consisting of 5 relay nodes in addition to the endpoints. The shared portion of the paths consists of the link $R_2 \to R_3$, while the links $R_3 \to R_4$ and $R_3 \to R_5$ constitute the point-of-split. All the other links belong to the independent category; specifically, these are the links $A \to R_1$, $R_1 \to R_2$, and $R_5 \to B$ in the forward path, and $B \to R_2$ and $R_4 \to A$ in the reverse path. We emphasize that the links $R_5 \to B$ and $R_4 \to A$ are counted as independent even though they are used to carry the coded packet, rather than native ones. The reason is that their ETX metric still counts towards the total cost in the same way; in other words, no saving due to network coding can be attributed to these links, compared to the situation where network coding is not used at all and only native packets are transmitted throughout.

Using the above categories and the corresponding notations, we finally arrive at the cost formula for any given pair of paths:

$$\sum_{L_{u,v} \in \mathcal{I}_a} \frac{1}{P_{u,v}} + \sum_{L_{u,v} \in \mathcal{I}_b} \frac{1}{P_{u,v}} + \sum_{L_{u,v} \in \mathcal{S}} \frac{1}{P_{u,v}} + \frac{1}{P_{\mathcal{L}_a}} + \frac{1}{P_{\mathcal{L}_b}} - \frac{1}{1-(1-P_{\mathcal{L}_a})(1-P_{\mathcal{L}_b})} \quad (5)$$

(the last term in (5) does not apply if the paths are entirely node-disjoint; in that case, of course, the set $\mathcal{S}$ is empty as well). The remainder of the paper focuses on the problem of finding the pair of paths, and the node where network coding is performed, such that the cost given by (5) is minimized.

III. The Optimal Coding-Aware Routing Solution

In this section we present a solution algorithm for the above routing optimization problem. Without loss of generality, we henceforth disregard the possibility that network coding is not employed in the optimal solution at all. Indeed, testing for that possibility is straightforward; if network coding is not used, then the optimal solution must consist simply of the shortest path (with respect to the ETX metric) in each direction, and the resulting cost can be simply compared with the outcome of the coding-aware algorithm described henceforth.

To that end, we first define some further notation related to the categories defined in the previous section. Specifically, we attach the following labels to the nodes serving as boundaries between links in different categories, as follows:

- $N1$ is the node where the shared portion of the paths commences. It is also the node where network coding is initially performed.

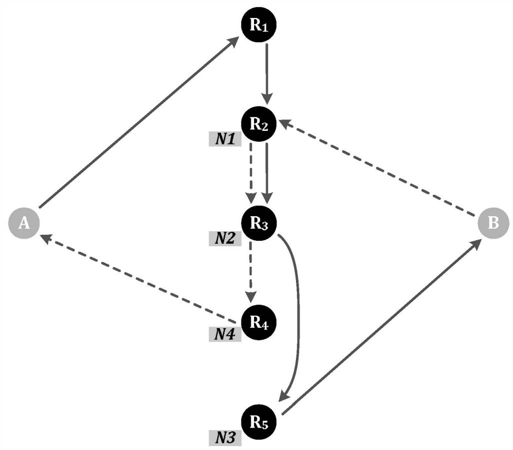

Fig. 2: Example of a network with 5 intermediate relays.

- $N2$ is the node where the shared portion of the paths ends, and is the transmitter of the point-of-split links.
- $N3, N4$ are the nodes at the receiving end of the point-of-split links, i.e. the direct successors of $N2$ on the forward and reverse paths, respectively.

Thus, links in the shared category form a path from the node labeled $N1$ to $N2$; the two links in the point-of-split are $N2 \to N3$ and $N2 \to N4$; and all the remaining links are independent. An example of these labels can be seen in Figure 2. We now clarify the following rules regarding the definitions of nodes $N1, N2, N3, N4$:

1) The source nodes A and B may not be labeled as $N1$ or $N2$; these labels can only refer to relay nodes (since, by definition, the endpoints do not transmit coded packets).
2) The labels $N1$ and $N2$ may refer to the same node, if the forward and reverse paths contain no shared links.
3) The nodes labeled $N3$ and $N4$ must be direct neighbors of $N2$, i.e. $P_{N2,N3} > 0$, $P_{N2,N4} > 0$. They must be distinct from $N1$, $N2$, and each other; however, node $N3$ may coincide with the destination B, and node $N4$ may be the destination A.

We now present the following essential structural property of the optimal routing solution.

Lemma 1. *In the optimal solution, the following path sections:*

- *between A and $N1$ on the forward path,*
- *between B and $N1$ on the reverse path,*
- *between $N1$ and $N2$ (the shared part of the paths),*
- *between $N3$ and B on the forward path, and*
- *between $N4$ and A on the reverse path,*

must, by themselves, be the shortest paths between the corresponding nodes (with respect to the ETX metric).

Proof: It can be seen that, if any of the above sections is not by itself the shortest path between the corresponding nodes, then replacing that section with a different set of links that do form the shortest path between those nodes (and making no changes to other parts of the paths) will lead to a reduction in the corresponding component in the cost expression (5), without changing the other components (in light of the independence assumption of link probabilities), in contradiction to the optimality. ■

input : The set of nodes $\mathcal{N} = \{A, B, R_1, R_2, \dots\}$ and a matrix $(P_{u,v})$ of successful transmission probabilities between every pair of nodes;
output: The lowest cost achievable by any pair of paths, and the corresponding tuple $\langle N1, N2, N3, N4\rangle$.

Define $d_{u,v} = \frac{1}{P_{u,v}}$ (the ETX metric) for every $u, v \in \mathcal{N}$;
Calculate the shortest paths between all nodes, such that $s_{u,v}$ is the shortest distance from u to v;
for $N1 \in \mathcal{N} \setminus \{A, B\}$ **do**
  for $N2 \in \mathcal{N} \setminus \{A, B\}$ **do**
    for $N3 \in \mathcal{N} \setminus \{N1, N2, A\}$ *s.t.* $P_{N2,N3} > 0$ **do**
      for $N4 \in \mathcal{N} \setminus \{N1, N2, N3, B\}$ *s.t.* $P_{N2,N4} > 0$ **do**
        $C_{ind} \leftarrow s_{A,N1} + s_{B,N1} + s_{N3,B} + s_{N4,A}$;
        $C_{shared} \leftarrow s_{N1,N2}$;
        $C_{split} \leftarrow d_{N2,N3} + d_{N2,N4} - \frac{1}{1-(1-P_{N2,N3})(1-P_{N2,N4})}$;
        $C_{total} \leftarrow C_{ind} + C_{shared} + C_{split}$;
        if $C_{total} < C^*$ **then**
          $N^* \leftarrow \langle N1, N2, N3, N4\rangle$;
          $C^* \leftarrow C_{total}$;
        end
      end
    end
  end
end
return C^*, N^*

Algorithm 1: The optimal routing algorithm for bidirectional flows with network coding.

A corollary of Lemma 1 is that, in the optimal solution, the shared portion of the two paths is contiguous, i.e. it is not possible for the paths to join, split, and then join again at a later point. Otherwise, if a pair of paths has two shared sections with a split in between, then it would take fewer transmissions to send the coded packet along just one of the split branches rather than unnecessarily duplicating it on both. Hence, our labeling definitions, which implicitly require that the entire portion between $N1$ and $N2$ is shared between the forward and reverse paths, indeed characterize the optimal paths without loss of generality.

With the help of Lemma 1 and the above corollary, it becomes evident that the problem of finding the lowest-cost pair of paths can be reduced to an enumeration of all the possible labelings $N1, N2, N3, N4$, calculating the shortest paths in all the respective section, and recording the labeling that leads to the lowest total cost. Algorithm 1 hereabove implements the above, listed in pseudo-code form. This algorithm has a running time of $O(N^4)$ (note that the initial calculation of shortest paths between all pairs of nodes can be accomplished in $O(N^3)$ time, using the Floyd-Warshall algorithm).

IV. Performance Evaluation

In this section we evaluate the performance of the optimal network coding paths as found by our algorithm, measured in

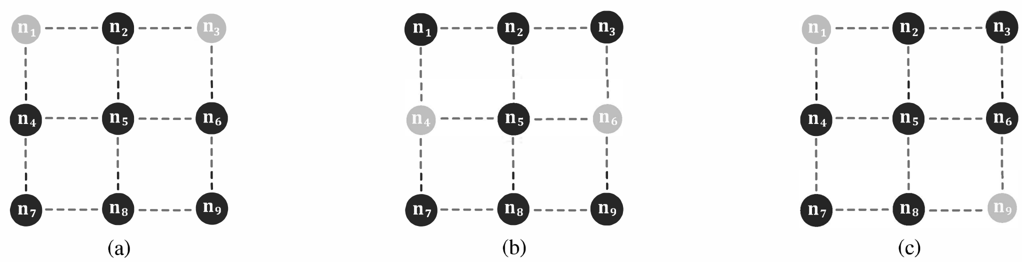

Fig. 3: A 3×3 network topology, with endpoint nodes at (a) opposite sides of first row; (b) opposite sides of middle row; (c) opposite sides of main diagonal.

terms of the expected number of transmissions (ETX) required to convey one packet in each direction. Our goal is to compare the optimal paths found by the algorithm with the prevailing network coding approach, which requires the packets to be decoded immediately at next hops of their respective flows. Note that, in terms of the notation in our algorithm, this requirement means that each of the nodes $N3$ and $N4$ must be on the path from the respective source node to $N1$ (and, therefore, have previously received the native packet in the opposite direction). In addition, to gain a further insight into the extent that allowing the two paths to contain a shared portion contributes to the improvement, we also compare the optimal ETX cost with the case where the forward and reverse paths may not contain shared links; in other words, when $N1$ and $N2$ are forced to be the same node.

We evaluate our algorithm and the above two variations using an underlying topology consisting of a mesh network of k^2 nodes, arranged as a square grid of k columns by k rows, such that each node is connected to its adjacent nodes in the same row and the same column. We consider four sets of scenarios, corresponding to $k = 3, 4, 5, 6$ respectively (Figure 3 illustrates the topology for $k = 3$). Within each set of scenarios, we consider three cases differing in the position of the flow endpoint nodes (A,B): (1) case 1 where the endpoint nodes are located at the two ends of the first row, (2) case 2 where the endpoint nodes are located at the two ends of the middle row (in order to minimize the edge effect); and (3) case 3 where the endpoints are located at the far two ends of the main diagonal of the grid. For example, in the set corresponding to $k = 3$, the endpoint nodes are set at $n1, n3$ in scenario 1, $n4, n6$ in scenario 2, and $n1, n9$ in scenario 3 (see Figures 3a, 3b and 3c). In all of the scenarios, the transmission success probability is set for each link independently to a random value drawn from a uniform distribution between 0 and 1. All the results reported below are obtained from averages where each of the scenarios is simulated for 1000 times; specifically, since the absolute ETX values vary over a significant range due to the random probability values of the individual links, the subsequent results (unless stated otherwise) present the average cost improvement of the optimal path in *relative* (percentage) terms.

We first compare the performance of the optimal paths to the prevailing approach in network coding literature, which requires combined packets to be immediately decodable at the next hop. Figure 4a depicts the average cost improvement gained by the optimal routing across the various scenarios. As can be seen, the average improvement due to allowing packets to be forwarded in coded form are highest in the 3×3 grid, particularly when the endpoint nodes are not located on the diagonal (nearly 35%). While the average improvement decreases with the size of the network, dropping to about 2.5% for the 6×6 networks with endpoints located diagonally, Figure 4b shows that the *frequency* of scenarios (i.e. the percentage out of all the random topologies tested) in which the optimal routing had some positive (non-zero) benefit is very high, remaining at more than 60% across all the sets.

In order to understand where the gains from forwarding of coded packets mostly occur, we turn our focus to the shared category of links in the optimal paths (i.e. the part between nodes $N1$ and $N2$ in the algorithm), and study the ETX cost reduction of the optimal paths compared to the variation where $N1$ and $N2$ are forced to be the same node (but without requiring that $N3$ and $N4$ must be able to decode the combined packet immediately). The average improvement in this case is shown in Figure 5a. We observe similar trends as in Figure 4a, and, furthermore, we note that, except for the 3×3 mesh, a large portion of the overall improvement of the optimal routing paths can be attributed to the transmissions over the shared section. For completeness, the frequency of the scenarios with a non-zero improvement due to allowing shared links is depicted in Figure 5b. Again, we observe that the frequency of such scenarios increases as the network grows, but, interestingly, is significantly lower across all sets of scenarios than in Figure 4b.

There are several important insights about the benefits of optimal coding-aware routing in asymmetric networks that can be made from the above observations. The most important conclusion is that, even though some reduction of ETX is achieved in many topologies simply by allowing the combined packet to be forwarded (from the node that performs the network coding initially) towards both destinations in coded form, the particularly significant gains are dominated by instances where the topology allows both of the respective paths to share one or more links in the same direction before splitting towards the two destinations. Furthermore, the position of the endpoints in a random asymmetric network is of paramount importance; for example, when the flow endpoints are positioned at the opposite sides of the diagonal, the shortest paths (with respect to ETX) in the opposite directions are least likely to contain overlapping links, and, consequently, the gain from allowing

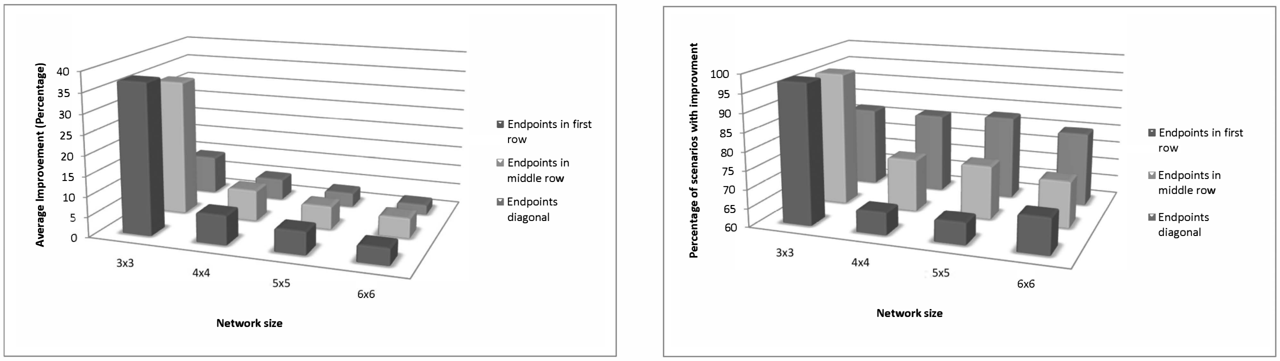

(a) Average ETX cost improvement (percentage)

(b) Frequency of instances with nonzero improvement (percentage)

Fig. 4: Improvement of total path cost due to forwarding of coded packets.

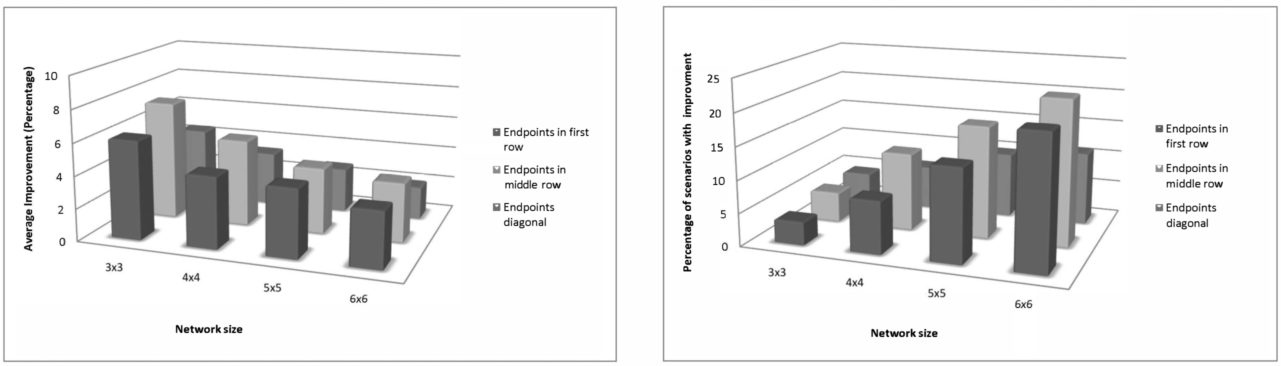

(a) Average ETX cost improvement (percentage)

(b) Frequency of instances with nonzero improvement (percentage)

Fig. 5: Improvement of total path cost due to links in shared category.

the forwarding of coded packets is the least pronounced, compared to other endpoint positions.

V. Conclusion

We have considered the problem of optimal routing of packets in a bidirectional flow between a pair of endpoints in the presence of wireless network coding, where the wireless links have positive and asymmetric packet erasure probabilities. We have shown that the optimal pair of paths, which account both for the link qualities and the coding opportunities arising where the paths intersect, can exhibit a number of different ways of intersecting, and may even have a common portion where both forward and reverse packets travel in the same direction. This effect is a direct result of allowing coded packets to be routed over several hops (and ultimately decoded at the endpoints of the flow), unlike the majority of related work on coding-aware routing where a coding opportunity may arise only when the packet can be decoded immediately at all the next-hop nodes. We have presented a polynomial-time algorithm to find the optimal paths so as to minimize the expected number of transmissions required to deliver one packet in each direction of the flow, and evaluated it over a range of random topologies, showing that the cost of the paths can be reduced significantly (up to 35% in some instances) by exploiting the ability to forward coded packets along overlapping links in both directions. The extensions of our solution so as to alleviate the assumption of independence among the link erasure probabilities, and combine it with opportunistic routing techniques (based on actual reception or overhearing of packets in real time) as opposed to static routing precomputed in advance [7], [8], remain the subject of ongoing work.

References

[1] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft. XORs in the air: Practical wireless network coding. *IEEE/ACM Trans. on Networking*, 16(3):497–510, June 2008.

[2] Y.E. Sagduyu and A. Ephremides. On joint MAC and network coding in wireless ad hoc networks. *IEEE Trans. on Information Theory*, 53(10):3697–3713, October 2007.

[3] B. Ni, N. Santhapuri, Z. Zhong, and S. Nelakuditi. Routing with opportunistically coded exchanges in wireless mesh networks. In *Proc. IEEE Workshop on Wireless Mesh Networks (WiMesh)*, pages 157–159, Reston, VA, September 2006.

[4] J. Le, J.C.S. Lui, and D.M. Chiu. DCAR: Distributed coding-aware routing protocol for wireless networks. *IEEE Trans. on Mobile Computing*, 9(4):596–608, April 2010.

[5] J. Sun, Y. Liu, H. Hu, and D. Yuan. On-demand coding-aware routing in wireless mesh networks. *The Journal of China Universities of Posts and Telecommunications*, 17(5):80–86, October 2010.

[6] S. Sengupta, S. Rayanchu, and S. Banerjee. Network coding-aware routing in wireless networks. *IEEE/ACM Trans. on Networking*, 18(4):1158–1170, August 2010.

[7] T. Mehmood and L. Libman. Towards optimal forwarding in wireless networks: Opportunistic routing meets network coding. In *Proc. IEEE Conference on Local Computer Networks (LCN)*, Zurich, Switzerland, October 2009.

[8] L. Libman, G.S. Paschos, L. Georgiadis, and X. Zhao. Throughput regions and optimal policies in wireless networks with opportunistic routing. In *Proc. International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Tsukuba, Japan, May 2013.