

XORs in The Air: Practical Wireless Network Coding

Sachin Katti[†] Hariharan Rahul[†] Wenjun Hu^{*} Dina Katabi[†] Muriel Médard[†] Jon Crowcroft^{*}
[†]MIT CSAIL ^{*}Univ. of Cambridge

ABSTRACT

This paper proposes COPE, a new architecture for wireless mesh networks. In addition to forwarding packets, routers mix (i.e., code) packets from different sources to increase the information content of each transmission. We show that intelligently mixing packets increases network throughput. Our design is rooted in the theory of network coding. Prior work on network coding is mainly theoretical and focuses on multicast traffic. This paper aims to bridge theory with practice; it addresses the common case of unicast traffic, dynamic and potentially bursty flows, and practical issues facing the integration of network coding in the current network stack. We evaluate our design on a 20-node wireless network, and discuss the results of the first testbed deployment of wireless network coding. The results show that COPE largely increases network throughput. The gains vary from a few percent to several folds depending on the traffic pattern, congestion level, and transport protocol.

Categories and Subject Descriptors

C.2.2 [Computer Systems Organization]: Computer-Communications Networks

General Terms

Algorithms, Design, Performance, Theory

Keywords

Network Coding, Wireless Networks

1. INTRODUCTION

Wireless networks are indispensable; they provide the means for mobility, city-wide Internet connectivity, distributed sensing, and outdoor computing. Current wireless implementations, however, suffer from a severe throughput limitation and do not scale to dense large networks.

This paper presents COPE, a new forwarding architecture that substantially improves the throughput of wireless networks. COPE inserts a coding shim between the IP and MAC layers, which identifies coding opportunities and benefits from them by forwarding multiple packets in a single transmission.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06, September 11–15, 2006, Pisa, Italy.
Copyright 2006 ACM 1-59593-308-5/06/0009 ...\$5.00.

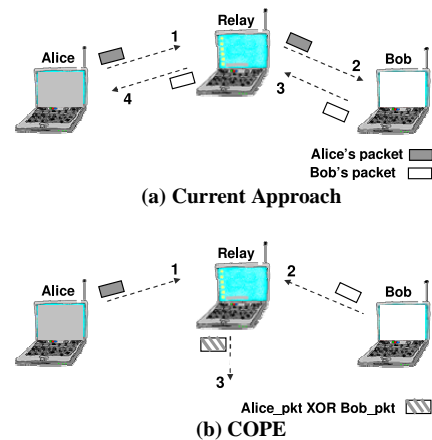


Figure 1—A simple example of how COPE increases the throughput. It allows Alice and Bob to exchange a pair of packets using 3 transmissions instead of 4 (numbers on arrows show the order of transmission).

To give the reader a feel for how COPE works, we start with a fairly simple example. Consider the scenario in Fig. 1, where Alice and Bob want to exchange a pair of packets via a router. In current approaches, Alice sends her packet to the router, which forwards it to Bob, and Bob sends his packet to the router, which forwards it to Alice. This process requires 4 transmissions. Now consider a network coding approach. Alice and Bob send their respective packets to the router, which XORs the two packets and broadcasts the XOR-ed version. Alice and Bob can obtain each other's packet by XOR-ing again with their own packet. This process takes 3 transmissions instead of 4. Saved transmissions can be used to send new data, increasing the wireless throughput.

In fact, COPE leads to larger bandwidth savings than are apparent from this example. COPE exploits the shared nature of the wireless medium which, for free, broadcasts each packet in a small neighborhood around its path. Each node stores the overheard packets for a short time. It also tells its neighbors which packets it has heard by annotating the packets it sends. When a node transmits a packet, it uses its knowledge of what its neighbors have heard to perform *opportunistic coding*; the node XORs multiple packets and transmits them as a single packet if each intended next hop has enough information to decode the encoded packet. This extends COPE beyond two flows that traverse the same nodes in reverse order (as in the Alice-and-Bob example), and allows it to XOR more than a pair of packets.

COPE's design is based on two key principles.

- COPE disposes of the point-to-point abstraction and embraces the broadcast nature of the wireless channel. Network designers typically abstract the wireless channel as a point-to-point link,

then adapt forwarding and routing techniques designed for wired networks for wireless. In contrast, COPE exploits the broadcast property of radios instead of hiding it under an artificial abstraction.

- *COPE employs network coding.* Our work is rooted in the theory of network coding, which allows the routers to mix the information content in the packets before forwarding them. Prior work on network coding is mainly theoretical and focuses on multicast traffic [2, 26, 20, 24, 17, 18, 27, 11]. Ours bridges theory with practice; it addresses the common case of unicast traffic, dynamic and potentially bursty flows, and practical issues facing the integration of network coding in the current network stack.

We have introduced the idea of opportunistic wireless network coding in an invited paper in Allerton 2005 [23]. This paper departs from our previous paper and all prior work on network coding in three main ways:

- (1) This paper presents the first system architecture for wireless network coding. It articulates a full-fledged design that integrates seamlessly into the current network stack, works with both TCP and UDP flows, and runs real applications.
- (2) The paper also implements the design in the Linux kernel and the Roofnet platform [1]. The implementation is deployed on a 20-node wireless testbed, creating the first deployment of network coding in a wireless network.
- (3) The paper studies the performance of COPE, and reveals its interactions with the wireless channel, routing, and higher layer protocols. Our findings can be summarized as follows:

- Network coding does have practical benefits, and can substantially improve wireless throughput.
- When the wireless medium is congested and the traffic consists of many random UDP flows, COPE increases the throughput of our testbed by 3-4x.
- If the traffic does not exercise congestion control (e.g., UDP), COPE's throughput improvement may substantially exceed the expected theoretical coding gain. This additional gain occurs because coding makes a router's queue smaller, reducing the probability that a congested downstream router will drop packets that have already consumed network resources.
- For a mesh network connected to the Internet via an access point, the throughput improvement observed with COPE varies depending on the ratio between total download and upload traffic traversing the access point, and ranges from 5% to 70%.
- Hidden terminals create a high collision rate that cannot be masked even with the maximum number of 802.11 retransmissions. In these environments, TCP does not send enough to utilize the medium, and thus does not create coding opportunities. With no hidden terminals, TCP's throughput increases by an average of 38% in our testbed.

2. BACKGROUND AND RELATED WORK

The idea underlying network coding is usually illustrated using the famous butterfly example [2]. Consider the network in Fig. 2, where source S_1 wants to deliver the stream of messages a_i to both R_1 and R_2 , and source S_2 wants to send the stream of messages b_i to the same two receivers. Assume all links have a capacity of one message per unit of time. If routers only forward the messages they receive, the middle link will be a bottleneck, which for every time unit, can either deliver a_i to R_1 or b_i to R_2 . In contrast, if the router feeding the middle link XORs the two messages and sends $a_i \oplus b_i$ (or any linear combination of a_i and b_i), as shown in the figure, both receivers obtain two messages in every time unit. Thus, network




Figure 2—A simple scenario showing how network coding improves throughput. All links have a capacity of one message per unit of time. By sending the XOR of a_i and b_i on the middle link, we can deliver two messages per unit of time to both receivers.

Term	Definition
Native Packet	A non-encoded packet
Encoded or XOR-ed Packet	A packet that is the XOR of multiple native packets
Next hops of an Encoded Packet	The set of next hops for the native packets XOR-ed to generate the encoded packet
Packet Id	A 32-bit hash of the packet's IP source address and IP sequence number
Output Queue	A FIFO queue at each node, where it keeps the packets it needs to forward
Packet Pool	A buffer where a node stores all packets heard in the past T seconds
Coding Gain	The ratio of the number of transmissions required by the current non-coding approach, to the number of transmissions used by COPE to deliver the same set of packets.
Coding+MAC Gain	The expected throughput gain with COPE when an 802.11 MAC is used, and all nodes are backlogged.

Table 1—Definitions of terms used in this paper.

coding, i.e., allowing the routers to mix the bits in forwarded messages, can increase network throughput.

Work on network coding started with a pioneering paper by Ahlswede et al. [2], who showed that having the routers mix information in different messages allows the communication to achieve multicast capacity. This was soon followed by the work of Li et al., who showed that, for multicast traffic (e.g., the butterfly scenario), linear codes are sufficient to achieve the maximum capacity bounds [26]. Koetter and Médard [24] presented polynomial time algorithms for encoding and decoding, and Ho et al. extended these results to random codes [17]. Some recent work studied wireless network coding [11, 31]. In particular, Lun et al. studied network coding in the presence of omni-directional antennae and showed that the problem of minimizing the communication cost can be formulated as a linear program and solved in a distributed manner [28]. All of this work is primarily theoretical and assumes multicast traffic. A few papers study specific unicast topologies showing that, for the studied scenario, network coding results in better throughput than pure forwarding [39, 16, 37]. This paper aims to bridge the gap between the theory of network coding and practical network design and provide an operational protocol for general unicast traffic.

Finally, a rich body of systems research has tackled the problem of improving the throughput of wireless networks. The proposed solutions range from designing better routing metrics [10, 5, 12] to tweaking the TCP protocol [33], and include improved routing and MAC protocols [6, 22, 15]. Our work builds on these foundations but adopts a fundamentally different approach; it explores the utility of network coding in improving the throughput of wireless networks.

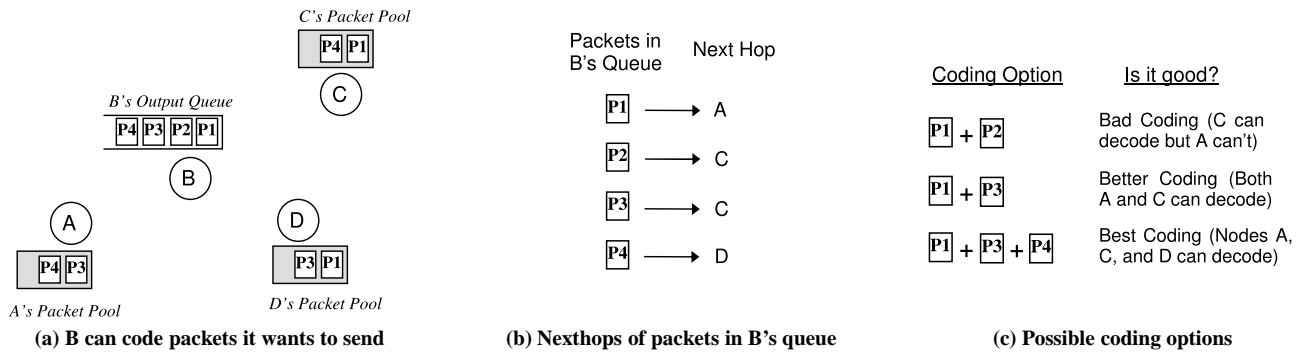


Figure 3—Example of Opportunistic Coding; Node B has 4 packets in its queue, whose next hops are listed in (b). Each neighbor of B has stored some packets as depicted in (a). Node B can make a number of coding decisions (as shown in (c)), but should select the last one because it maximizes the number of packets delivered in a single transmission.

3. COPE OVERVIEW

We introduce COPE, a new forwarding architecture for wireless mesh networks. It inserts a coding layer between the IP and MAC layers, which detects coding opportunities and exploits them to forward multiple packets in a single transmission. Before delving into details, we refer the reader to Table 1, which defines the terms used in the rest of the paper.

COPE incorporates three main techniques:

(a) Opportunistic Listening: Wireless is a broadcast medium, creating many opportunities for nodes to overhear packets when they are equipped with omni-directional antennae. COPE sets the nodes in promiscuous mode, makes them snoop on all communications over the wireless medium and store the overheard packets for a limited period T (the default is $T = 0.5s$).

In addition, each node broadcasts *reception reports* to tell its neighbors which packets it has stored. Reception reports are sent by annotating the data packets the node transmits. A node that has no data packets to transmit periodically sends the reception reports in special control packets.

(b) Opportunistic Coding: The key question is what packets to code together to maximize throughput. A node may have multiple options, but it should aim to *maximize the number of native packets delivered in a single transmission, while ensuring that each intended next hop has enough information to decode its native packet.*

The above is best illustrated with an example. In Fig. 3(a), node B has 4 packets in its output queue $p_1, p_2, p_3,$ and p_4 . Its neighbors have overheard some of these packets. The table in Fig 3(b) shows the next hop of each packet in B 's queue. When the MAC permits B to transmit, B takes packet p_1 from the head of the queue. Assuming that B knows which packets each neighbor has, it has a few coding options as shown in Fig. 3(c). It could send $p_1 \oplus p_2$. Since node C has p_1 in store, it could XOR p_1 with $p_1 \oplus p_2$ to obtain the native packet sent to it, i.e., p_2 . However, node A does not have p_2 , and so cannot decode the XOR-ed packet. Thus, sending $p_1 \oplus p_2$ would be a bad coding decision for B , because only one neighbor can benefit from this transmission. The second option in Fig. 3(c) shows a better coding decision for B . Sending $p_1 \oplus p_3$ would allow both neighbors C and A to decode and obtain their intended packets from a single transmission. Yet the best coding decision for B would be to send $p_1 \oplus p_3 \oplus p_4$, which would allow all three neighbors to receive their respective packets all at once.

The above example emphasizes an important coding issue. Packets from multiple unicast flows may get encoded together at some intermediate hop. But their paths may diverge at the next hop, at which point they need to be decoded. If not, unneeded data will

be forwarded to areas where there is no interested receiver, wasting much capacity. The coding algorithm should ensure that all next hops of an encoded packet can decode their corresponding native packets. This can be achieved using the following simple rule:

To transmit n packets, p_1, \dots, p_n , to n next hops, r_1, \dots, r_n , a node can XOR the n packets together only if each next-hop r_i has all $n - 1$ packets p_j for $j \neq i$.

This rule ensures that each next hop can decode the XOR-ed version to extract its native packet. Whenever a node has a chance to transmit a packet, it chooses the largest n that satisfies the above rule to maximize the benefit of coding.

(c) Learning Neighbor State: But how does a node know what packets its neighbors have? As explained earlier, each node announces to its neighbors the packets it stores in reception reports. However, at times of severe congestion, reception reports may get lost in collisions, while at times of light traffic, they may arrive too late, after the node has already made a suboptimal coding decision. Therefore, a node cannot rely solely on reception reports, and may need to guess whether a neighbor has a particular packet.

To guess intelligently, we leverage the routing computation. Wireless routing protocols compute the delivery probability between every pair of nodes and use it to identify good paths. For e.g., the ETX metric [10] periodically computes the delivery probabilities and assigns each link a weight equal to $1/(\text{delivery probability})$. These weights are broadcast to all nodes in the network and used by a link-state routing protocol to compute shortest paths. We leverage these probabilities for guessing. In the absence of deterministic information, COPE estimates the probability that a particular neighbor has a packet as the delivery probability of the link between the packet's previous hop and the neighbor.

Occasionally, a node may make an incorrect guess, which causes the coded packet to be undecodable at some next hop. In this case, the relevant native packet is retransmitted, potentially encoded with a new set of native packets.

4. UNDERSTANDING COPE'S GAINS

How beneficial is COPE? Its throughput improvement depends on the existence of coding opportunities, which themselves depend on the traffic patterns. This section provides some insight into the expected throughput increase and the factors affecting it.

4.1 Coding Gain

We defined the *coding gain* as the ratio of the number of transmissions required by the current non-coding approach, to the minimum

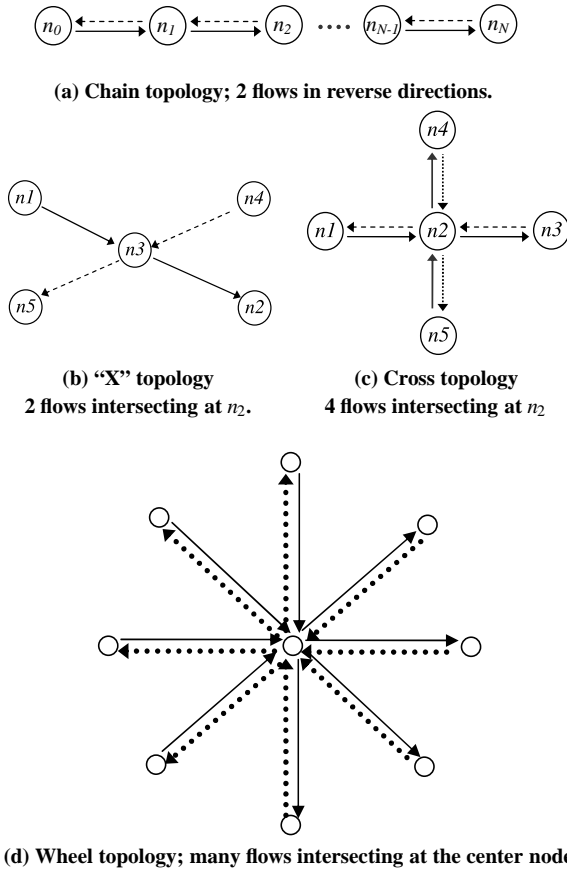


Figure 4—Simple topologies to understand COPE's Coding and Coding+MAC Gains.

number of transmissions used by COPE to deliver the same set of packets. By definition, this number is greater than or equal to 1.

In the Alice-and-Bob experiment, as described in §1, COPE reduces the number of transmissions from 4 to 3, thus producing a coding gain of $\frac{4}{3} = 1.33$.

But what is the maximum achievable coding gain, i.e., what is the theoretical capacity of a wireless network that employs COPE? The capacity of general network coding for unicast traffic is still an open question for arbitrary graphs [38, 16]. However, we analyze certain basic topologies that reveal some of the factors affecting COPE's coding gain. Our analysis assumes identical nodes, omnidirectional radios, perfect hearing within some radius, and the signal is not heard at all outside this radius, and if a pair of nodes can hear each other the routing will pick the direct link. Additionally, we assume that the flows are infinite and we only consider the steady state.

THEOREM 4.1. *In the absence of opportunistic listening, COPE's maximum coding gain is 2, and it is achievable.*

We prove the theorem by showing that the coding gain of the chain in Fig. 4(a) tends to 2 as the number of intermediate nodes increases. The complete proof is in Appendix A.

While we do not know the maximum gain for COPE with opportunistic listening, there do exist topologies where opportunistic listening adds to the power of COPE. For example, consider the "X"-topology shown in Fig. 4(b). This is the analogy of the Alice-and-Bob topology, but the two flows travel along link-disjoint paths. COPE without opportunistic listening cannot achieve any gains on

this topology. But with opportunistic listening and guessing, the middle node can combine packets traversing in opposite directions, for a coding gain of $\frac{4}{3} = 1.33$. This result is important, because in a real wireless network, there might be only a small number of flows traversing the reverse path of each other à la Alice-and-Bob, but one would expect many flows to intersect at a relay, and thus can be coded together using opportunistic listening and guessing.

The "X" and Alice-and-Bob examples can be combined to further improve the benefits of coding, as in the cross topology of Fig. 4(c). Without coding, 8 transmissions are necessary for each flow to send one packet to its destination. However, assuming perfect overhearing (n_4 and n_5 can overhear n_1 and n_3 , and vice versa), n_2 can XOR 4 packets in each transmission, thus reducing the number of transmissions from 8 to 5, producing a coding gain of $\frac{8}{5} = 1.6$.

We observe that while this section has focused on theoretical bounds, the gains in practice tend to be lower due to the availability of coding opportunities, packet header overheads, medium losses, etc. However, it is important to note that COPE increases the actual information rate of the medium far above the bit rate, and hence its benefits are sustained even when the medium is fully utilized. This contrasts with other approaches to improving wireless throughput, such as opportunistic routing [6], which utilize the medium better when it is not fully congested, but do not increase its capacity.

4.2 Coding+MAC Gain

When we ran experiments with COPE, we were surprised to see that the throughput improvement sometimes greatly exceeded the coding gain for the corresponding topology. It turns out that the interaction between coding and the MAC produces a beneficial side effect that we call the Coding+MAC gain.

The Coding+MAC gain is best explained using the Alice-and-Bob scenario. Because it tries to be fair, the MAC divides the bandwidth equally between the 3 contending nodes: Alice, Bob, and the router. Without coding, however, the router needs to transmit twice as many packets as Alice or Bob. The mismatch between the traffic the router receives from the edge nodes and its MAC-allocated draining rate makes the router a bottleneck; half the packets transmitted by the edge nodes are dropped at the router's queue. COPE allows the bottleneck router to XOR pairs of packets and drain them twice as fast, doubling the throughput of this network. Thus, the Coding+MAC gain of the Alice-and-Bob topology is 2.

The Coding+MAC gain assumes all nodes continuously have some traffic to send (i.e., backlogged), but are limited by their MAC-allocated bandwidth. It computes the throughput gain with COPE under such conditions. For topologies with a single bottleneck, like the Alice-and-Bob's, the Coding+MAC gain is the ratio of the bottleneck's draining rate with COPE to its draining rate without COPE.

Similarly, for the "X" and cross topologies, the Coding+MAC gain is higher than the coding gain. For the "X", the Coding+MAC gain is 2 since the bottleneck node is able to drain twice as many packets, given its MAC allocated rate. For the cross topology, the Coding+MAC gain is even higher at 4. The bottleneck is able to send 4 packets out in each transmission, hence it is able to drain four times as many packets compared to no coding. This begs the question: what is the maximum Coding+MAC gain? The maximum possible Coding+MAC gains with and without opportunistic listening are properties of the topology and the flows that exist in a network. Here we prove some upper bounds on Coding+MAC gains.

THEOREM 4.2. *In the absence of opportunistic listening, COPE's maximum Coding+MAC gain is 2, and it is achievable.*

Topology	Coding Gain	Coding+MAC Gain
Alice-and-Bob	1.33	2
“X”	1.33	2
Cross	1.6	4
Infinite Chain	2	2
Infinite Wheel	2	∞

Table 2—Theoretical gains for a few basic topologies.

The proof is in Appendix B.

THEOREM 4.3. *In the presence of opportunistic listening, COPE’s maximum Coding+MAC gain is unbounded.*

The proof, detailed in Appendix C, uses the wheel topology in Fig. 4(d). Assuming N edge nodes, with COPE the bottleneck node, in the center of the wheel, XORs N packets together, and consequently drains its queue N times faster than without COPE. As the number of edge nodes increases, i.e., $N \rightarrow \infty$, the gain becomes infinite. While the previous example is clearly artificial, it does illustrate the potential of COPE with opportunistic listening to produce a several-fold improvement in throughput, as in §7.3.

Table 2 lists the gains for a few basic topologies.

5. MAKING IT WORK

In order to integrate COPE effectively within the current network stack, we need to address some important system issues.

5.1 Packet Coding Algorithm

To build the coding scheme, we have to make a few design decisions. First, we design our coding scheme around the principle of *never delaying packets*. When the wireless channel is available, the node takes the packet at the head of its output queue, checks which other packets in the queue may be encoded with this packet, XORs those packets together, and broadcasts the XOR-ed version. If there are no encoding opportunities, our node does not wait for the arrival of a matching codable packet. COPE therefore lets the node opportunistically overload each transmission with additional information when possible, but does not wait for additional codable packets to arrive.

Second, COPE *gives preference to XOR-ing packets of similar lengths*, because XOR-ing small packets with larger ones reduces bandwidth savings. Empirical studies show that the packet-size distribution in the Internet is bimodal with peaks at 40 and 1500 bytes [29]. We can therefore limit the overhead of searching for packets with the right sizes by distinguishing between small and large packets. We might still have to XOR packets of different sizes. In this case, the shorter packets are padded with zeroes. The receiving node can easily remove the padding by checking the packet-size field in the IP header of each native packet.

Third, notice that COPE will *never code together packets headed to the same nexthop*, since the nexthop will not be able to decode them. Hence, while coding, we only need to consider packets headed to different nexthops. COPE therefore maintains two virtual queues per neighbor; one for small packets and another for large packets (The default setting uses a threshold of 100 bytes). When a new packet is added to the output queue, an entry is added to the appropriate virtual queue based on the packet’s nexthop and size.

Searching for appropriate packets to code is efficient due to the maintenance of virtual queues. When making coding decisions, COPE first dequeues the packet at the head of the FIFO output queue, and determines if it is a small or a large packet. Depending on the size, it looks at the appropriate virtual queues. For example, if the packet dequeued is a small packet, COPE first looks at the virtual queues for small packets. COPE looks only at the heads of

the virtual queues to limit packet reordering. After exhausting the virtual queues of a particular size, the algorithm then looks at the heads of virtual queues for packets of the other size. Thus for finding appropriate packets to code COPE has to look at $2M$ packets in the worst case, where M is the number of neighbors of a node.

Another concern is *packet reordering*. We would like to limit reordering packets from the same flow because TCP mistakes it as a congestion signal. Thus, we always consider packets according to their order in the output queue. Still, reordering may occur because we prefer to code packets of the same size. In practice, this reordering is quite limited because most data packets in a TCP flow are large enough to be queued in the large-packet queue, and thus be considered in order. We will see in §5.4, however, that reordering might arise from other reasons, particularly the need to retransmit a packet that has been lost due to a mistake in guessing what a neighbor can decode. Thus, we choose to deal with any reordering that might happen inside the network at the receiver. COPE has a module that puts TCP packets in order before delivering them to the transport layer as explained in §5.5.

Finally, we want to ensure that each neighbor to whom a packet is headed has a high probability of decoding its native packet. Thus, for each packet in its output queue, our relay node estimates the probability that each of its neighbors has already heard the packet. Sometimes the node can be certain about the answer, for example, when the neighbor is the previous hop of the packet, or when the reception reports from the neighbor state so. When neither of the above is true, the node leverages the delivery probabilities computed by the routing protocol; it estimates the probability the neighbor has the packet as the delivery probability between the packet’s previous hop and that neighbor. The node then uses this estimate to ensure that encoded packets are decodable by all of their nexthops with high probability.

In particular, suppose the node encodes n packets together. Let the probability that a nexthop has heard packet i be P_i . Then, the probability, P_D , that it can decode its native packet is equal to the probability that it has heard all of the $n - 1$ native packets XOR-ed with its own, i.e.,

$$P_D = P_1 \times P_2 \times \dots \times P_{n-1}.$$

Consider an intermediate step while searching for coding candidates. We have already decided to XOR $n - 1$ packets together, and are considering XOR-ing the n^{th} packet with them. The coding algorithm now checks that, for each of the n nexthops, the decoding probability P_D , after XOR-ing the n^{th} packet with the rest stays greater than a threshold G (the default value $G = 0.8$). If the above conditions are met, each nexthop can decode its packet with at least probability G . Finally, we note that for fairness we iterate over the set of neighbors according to a random permutation.

Formally, each node maintains the following data structures.

- Each node has a FIFO queue of packets to be forwarded, which we call *the output queue*.
- For each neighbor, the node maintains two *per-neighbor virtual queues*, one for small packets (e.g., smaller than 100 bytes), and the other for large packets. The virtual queues for a neighbor A contain pointers to the packets in the output queue whose nexthop is A .
- Additionally, the node keeps a hash table, *packet info*, that is keyed on packet-id. For each packet in the output queue, the table indicates the probability of each neighbor having that packet.

Whenever the MAC signals a sending opportunity, the node executes the procedure illustrated in Alg. 1.

1 Coding Procedure

```
Pick packet  $p$  at the head of the output queue.
Natives =  $\{p\}$ 
Nexthops =  $\{nexthop(p)\}$ 
if  $size(p) > 100$  bytes then
  which_queue = 1
else
  which_queue = 0
end if
for Neighbor  $i = 1$  to  $M$  do
  Pick packet  $p_i$ , the head of virtual queue  $Q(i, which\_queue)$ 
  if  $\forall n \in Nexthops \cup \{i\}, Pr[n \text{ can decode } p \oplus p_i] \geq G$  then
     $p = p \oplus p_i$ 
    Natives = Natives  $\cup \{p_i\}$ 
    Nexthops = Nexthops  $\cup \{i\}$ 
  end if
end for
which_queue = !which_queue
for Neighbor  $i = 1$  to  $M$  do
  Pick packet  $p_i$ , the head of virtual queue  $Q(i, which\_queue)$ 
  if  $\forall n \in Nexthops \cup \{i\}, Pr[n \text{ can decode } p \oplus p_i] \geq G$  then
     $p = p \oplus p_i$ 
    Natives = Natives  $\cup \{p_i\}$ 
    Nexthops = Nexthops  $\cup \{i\}$ 
  end if
end for
return  $p$ 
```

5.2 Packet Decoding

Packet decoding is simple. Each node maintains a *Packet Pool*, in which it keeps a copy of each native packet it has received or sent out. The packets are stored in a hash table keyed on packet id (see Table 1), and the table is garbage collected every few seconds. When a node receives an encoded packet consisting of n native packets, the node goes through the ids of the native packets one by one, and retrieves the corresponding packet from its packet pool if possible. Ultimately, it XORs the $n - 1$ packets with the received encoded packet to retrieve the native packet meant for it.

5.3 Pseudo-broadcast

The 802.11 MAC has two modes: unicast and broadcast. Since COPE broadcasts encoded packets to their next hops, the natural approach would be to use broadcast. Unfortunately, this does not work because of two reasons: poor reliability and lack of backoff.

Specifically, in the 802.11 unicast mode, packets are immediately ack-ed by their intended nexthops. The 802.11 protocol ensures reliability by retransmitting the packet at the MAC layer for a fixed number of times until a synchronous ack is received. Lack of an ack is interpreted as a collision signal, to which the sender reacts by backing off exponentially, thereby allowing multiple nodes to share the medium.

In contrast, 802.11 broadcast lacks both reliability and backoff. A broadcast packet has many intended receivers, and it is unclear who should ack. In the absence of the acks, the broadcast mode offers no retransmissions and consequently very low reliability. Additionally, a broadcast source cannot detect collisions, and thus does not back off. If multiple backlogged nodes share the broadcast channel, and each of them continues sending at the highest rate, the resulting throughput is therefore very poor, due to high collision rates.

Our solution is *pseudo-broadcast*, which piggybacks on 802.11 unicast and benefits from its reliability and backoff mechanism. Pseudo-broadcast unicasts packets that are meant for broadcast. The link-layer destination field is set to the MAC address of one of the intended recipients. An XOR-header is added after the link-layer header, listing all nexthops of the packet. Since all nodes are set in the promiscuous mode, they can overhear packets not addressed to

them. When a node receives a packet with a MAC address different from its own, it checks the XOR-header to see if it is a nexthop. If so, it processes the packet further, else it stores the packet in a buffer as an opportunistically received packet. As all packets are sent using 802.11 unicast, the MAC can detect collisions and backoff properly.

Pseudo-broadcast is also more reliable than simple broadcast. The packet is retransmitted multiple times until its designated MAC receiver receives the packet and acks it, or the number of retries is exceeded. A desirable side effect of these retransmissions is that nodes that are promiscuously listening to this packet have more opportunities to hear it. Pseudo-broadcast, however, does not completely solve the reliability problem, which we address in the next section.

5.4 Hop-by-hop ACKs and Retransmissions

(a) **Why hop-by-hop acks?** Encoded packets require all nexthops to acknowledge the receipt of the associated native packet for two reasons. First, encoded packets are headed to multiple nexthops, but the sender gets synchronous MAC-layer acks only from the nexthop that is set as the link layer destination of the packet (as explained in the previous section). There is still a probability of loss to the other nexthops from whom it does not get synchronous acks. Second, COPE may optimistically guess that a nexthop has enough information to decode an XOR-ed packet, when it actually does not.

The standard solution to wireless losses is to mask error-induced drops by recovering lost packets locally through acknowledgments and retransmissions [3, 19]. COPE too addresses this problem using local retransmissions; the sender expects the nexthops of an XOR-ed packet to decode the XOR-ed packet, obtain their native packet, and ack it. If any of the native packets is not ack-ed within a certain interval, the packet is retransmitted, potentially encoded with another set of native packets.

(b) **Asynchronous Acks and Retransmissions:** How should we implement these hop-by-hop acks? For non-coded packets, we simply leverage the 802.11 synchronous acks. Unfortunately, extending this synchronous ack approach to coded packets is highly inefficient, as the overhead incurred from sending each ack in its own packet with the necessary IP and WiFi headers would be excessive. Thus, in COPE encoded packets are ack-ed asynchronously.

When a node sends an encoded packet, it schedules a retransmission event for each of the native packets in the encoded packet. If any of these packets is not ack-ed within T_a seconds, the packet is inserted at the head of the output queue and retransmitted. (T_a is slightly larger than the round trip time of a single link.) Retransmitted packets may get encoded with other packets according to the scheme in §5.1.

A nexthop that receives an encoded packet decodes it to obtain its native packet, and immediately schedule an ack event. Before transmitting a packet, the node checks its pending ack events and incorporates the pending acks in the COPE header. If the node has no data packets to transmit, it sends the acks in periodic control packets—the same control packets used to send reception reports.

5.5 Preventing TCP Packet Reordering

Asynchronous acks can cause packet reordering, which may be confused by TCP as a sign of congestion. Thus, COPE has an *ordering agent*, which ensures that TCP packets are delivered in order. The agent ignores all packets whose final IP destinations differ from the current node, as well as non-TCP packets. These packets are immediately passed to the next processing stage. For each TCP flow ending at the host, the agent maintains a packet buffer and records the last TCP sequence number passed on to the network stack. Incoming packets that do not produce a hole in the TCP sequence

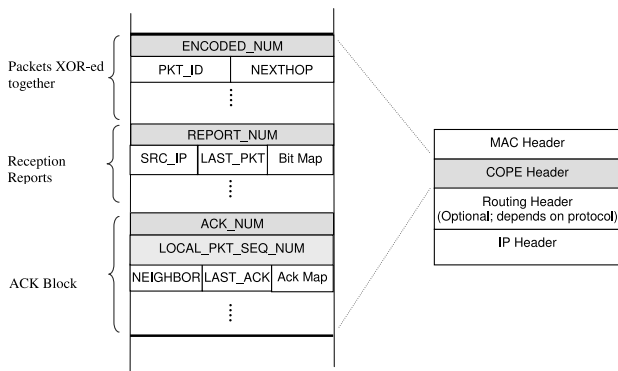


Figure 5—COPE Header. The first block identifies the native packets XOR-ed and their nexthops. The second block contains reception reports. Each report identifies a source, the last IP sequence number received from that source, and a bit-map of most recent packets seen from that source. The third block contains asynchronous acks. Each entry identifies a neighbor, an end point for the ACK map, and a bit-map of ack-ed packets.

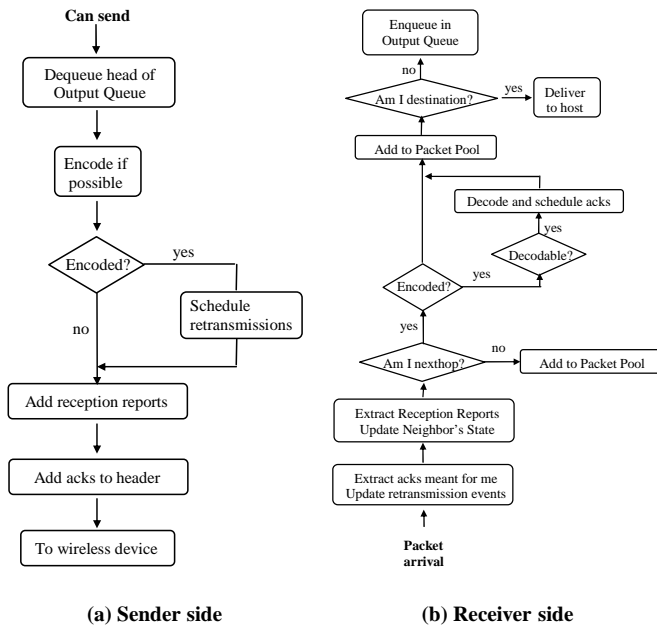


Figure 6—Flow chart for our COPE Implementation.

stream are immediately dispatched to the transport layer, after updating the sequence number state. Otherwise, they are withheld in the buffer till the gap in the sequence numbers is filled, or until a timer expires.

6. IMPLEMENTATION DETAILS

COPE adds special packet headers and alters the control flow of the router to code and decode packets. This section describes both parts.

6.1 Packet Format

COPE inserts a variable-length coding header in each packet, as shown in Fig. 5. If the routing protocol has its own header (e.g., Srcr [5]), COPE's header sits between the routing and the MAC headers. Otherwise, it sits between the MAC and IP headers. Only the shaded fields in Fig. 5 are required in every COPE header. The coding header contains the following 3 blocks.

(a) Ids of the coded native packets: The first block records metadata to enable packet decoding. It starts with ENCODED_NUM, the number of native packets XOR-ed together. For each native packet, the header lists its ID, which is a 32-bit hash of the packet's source IP address and IP sequence number. This is followed by the MAC address of the native packet's NextHop. When a node hears an XOR-ed packet, it checks the list of NextHops to determine whether it is an intended recipient for any of the native packets XOR-ed together, in which case it decodes the packet, and processes it further.

(b) Reception reports: Reception reports constitute the second block in the XOR header, as shown in Fig. 5. The block starts with the number of the reports in the packet, REPORT_NUM. This is followed by the IP sequence number of the last packet heard from that source LAST_PKT, and a bit-map of recently heard packets. For example, a report of the form {128.0.1.9, 50, 10000001} means that the last packet this node has heard from source 128.0.1.9 is packet 50, and it has also heard packets 42 and 49 from that source but none in between. The above representation for reception reports has two advantages: compactness and effectiveness. In particular, the bit-map allows the nodes to report each packet multiple times with minimal overhead. This guards against reception reports being dropped at high congestion.

(c) Expressing asynchronous acks compactly and robustly: To ensure ack delivery with minimum overhead, we use cumulative acks. Since they implicitly repeat ack information, cumulative acks are robust against packet drops. Each node maintains a per-neighbor 16-bit counter, called Neighbor_Seqno_Counter. Whenever the node sends a packet to that neighbor, the counter is incremented and its value is assigned to the packet as a local sequence number, LOCAL_PKT_SEQ_NUM. The two neighbors use this sequence number to identify the packet. Now, a node can use cumulative acks on a per-neighbor basis. Each coded packet contains an ack header as shown in Fig. 5. The ack block starts with the number of ack entries, followed by the packet local sequence number. Each ack entry starts with a neighbor MAC address. This is followed by a pointer to tell the neighbor where the cumulative acks stop, and a bit-map indicating previously received and missing packets. For example, an entry of {A, 50, 01111111} acks packet 50, as well as the sequence 43-49, from neighbor A. It also shows that packet 42 is still missing. Note that though we use cumulative acks, we do not guarantee reliability at link layer. In particular, each node retransmits a lost packet a few times (default is 2), and then gives up.

6.2 Control Flow

Fig. 6 abstracts the architecture of COPE. On the sending side, (shown in Fig. 6(a)), whenever the MAC signals an opportunity to send, the node takes the packet at the head of its output queue and hands it to the coding module (§5.1). If the node can encode multiple native packets in a single XOR-ed version, it has to schedule asynchronous retransmissions. Either way, before the packet can leave the node, pending reception reports and acks are added.

On the receiving side, (shown in Fig. 6(b)), when a packet arrives, the node extracts any acks sent by this neighbor to the node. It also extracts all reception reports and updates its view of what packets its neighbor stores. Further processing depends on whether the packet is intended for the node. If the node is not a nexthop for the packet, the packet is stored in the Packet Pool. If the node is a nexthop, it then checks if the packet is encoded. If it is, the node tries to decode by XOR-ing the encoded packet with the native packets it stores in its Packet Pool. After decoding it acks this reception to the previous




Figure 7—Node locations for one floor of the testbed.

hop and stores the decoded packet in the Packet Pool. The node now checks if it is the ultimate destination of the packet, if so it hands the packet off to the higher layers of the network stack. If the node is an intermediate hop, it pushes the packet to the output queue. If the received packet is not encoded, the packet is simply stored in the Packet Pool and processed in the same fashion as a decoded packet.

7. EXPERIMENTAL RESULTS

This section uses measurements from a 20-node wireless testbed to study both the performance of COPE and the interaction of network coding with the wireless channel and higher-layer protocols. Our experiments reveal the following findings.

- When the wireless medium is congested and the traffic consists of many random UDP flows, COPE delivers a 3-4 \times increase in the throughput of our wireless testbed.
- When the traffic does not exercise congestion control (e.g., UDP), COPE's throughput improvement substantially exceeds the expected coding gain and agrees with the Coding+MAC gain.
- For a mesh network connected to the Internet via a gateway, the throughput improvement observed with COPE varies depending on the ratio of download traffic to upload traffic at the the gateway, and ranges from 5% to 70%.
- Hidden terminals create a high loss rate that cannot be masked even with the maximum number of 802.11 retransmissions. In these environments, TCP does not send enough to utilize the medium and does not create coding opportunities. In environments with no hidden terminals, TCP's throughput improvement with COPE agrees with the expected coding gain.

7.1 Testbed

(a) Characteristics: We have a 20-node wireless testbed that spans two floors in our building connected via an open lounge. The nodes of the testbed are distributed in several offices, passages, and lounges. Fig. 7 shows the locations of the nodes on one of the floors. Paths between nodes are between 1 and 6 hops in length, and the loss rates of links on these paths range between 0 and 30%. The experiments described in this paper run on 802.11a with a bit-rate of 6Mb/s. Running the testbed on 802.11b is impractical because of a high level of interference from the local wireless networks.

(b) Software: Nodes in the testbed run Linux. COPE is implemented using the Click toolkit [25]. Our implementation runs as a user space daemon, and sends and receives raw 802.11 frames from the wireless device using a libpcap-like interface. The implementation exports a network interface to the user that can be treated like any other network device (e.g., eth0). Applications interact with

the daemon as they would with a standard network device provided by the Linux kernel. No modifications to the applications are therefore necessary. The implementation is agnostic to upper and lower layer protocols, and can be used by various protocols including UDP and TCP.

(c) Routing: Our testbed nodes run the Srcr implementation [5], a state-of-the-art routing protocol for wireless mesh networks. The protocol uses Dijkstra's shortest path algorithm on a database of link weights based on the ETT metric [5]. Router output queue is bounded at 100 packets.

(d) Hardware: Each node in the testbed is a PC equipped with an 802.11 wireless card attached to an omni-directional antenna. The cards are based on the NETGEAR 2.4 & 5 GHz 802.11a/g chipset. They transmit at a power level of 15 dBm, and operate in the 802.11 ad hoc mode, with RTS/CTS disabled.

(e) Traffic Model: We use a utility program called `udpgen` [35] to generate UDP traffic, and `tcpcp` [34] to generate TCP traffic. We either use long-live flows, or many shorter flows that match empirical studies of Internet traffic [30, 9], i.e., they have Poisson arrivals, and a Pareto file size with the shape parameter set to 1.17.

7.2 Metrics

Our evaluation uses the following metrics.

- *Network Throughput:* the measured total end-to-end throughput, i.e., the sum of the throughput of all flows in the network as seen by their corresponding applications.
- *Throughput Gain:* the ratio of the measured network throughputs with and without COPE. We compute the throughput gain from two consecutive experiments, with coding turned on, then off.

7.3 COPE in gadget topologies

We would like to compare COPE's actual throughput gain with the theoretical gains described in §4, and study whether it is affected by higher layer protocols. We start by looking at a few toy topologies with good link quality (medium loss rate after MAC retries < 1%), and no hidden terminals.

7.3.1 Long-Lived TCP Flows

We run long-lived TCP flows over 3 toy topologies: Alice-and-Bob, the "X", and the cross topologies depicted in Figs. 1 and 4. Fig. 8 plots the CDFs of the TCP throughput gain measured over 40 different runs. For the Alice-and-Bob topology the gain, shown in Fig. 8(a), is close to the theoretical coding gain of 1.33. The difference of 5 – 8% is due to the overhead of COPE's headers, as well as asymmetry in the throughput of the two flows, which prevents the router from finding a codemate for every packet. Similarly, for the "X"-topology, the gain in Fig. 8(b) is comparable to the optimal coding gain of 1.33. Finally, Fig. 8(c) shows the throughput gain for the cross topology with TCP. The gains are slightly lower than the expected coding gain of 1.6 because of header overhead, imperfect overhearing, and a slight asymmetry in the throughputs of the four flows.

The above experimental results reveal that when the traffic exercises congestion control, the throughput gain corresponds to the coding gain, rather than the Coding+MAC gain. The congestion control protocol, built into TCP, naturally matches the input rate at the bottleneck to its draining rate. When multiple long-lived TCP flows get bottlenecked at the same router, the senders back off and prevent excessive drops, leaving only pure coding gains.

7.3.2 UDP Flows

We repeat the above experiments with UDP and evaluate the




Figure 8—CDF of throughput gains obtained with COPE, for long-lived TCP flows.




Figure 9—CDF of throughput gains obtained with COPE, for UDP flows.

throughput gains. Fig. 9 plots a CDF of the UDP gain with COPE for the Alice-and-Bob, the “X”, and the cross topologies. The figure shows that the median UDP throughput gains for the three topologies are 1.7, 1.65, and 3.5 respectively.

Interestingly, the UDP gains are much higher than the TCP gains; they reflect the Coding+MAC gains for these toy topologies. Recall from §4 that the coding gain arises purely from the reduction in the number of transmissions achieved with COPE. Additionally, coding compresses the bottleneck queues, preventing downstream congested routers from dropping packets that have already consumed bandwidth, and producing a Coding+MAC gain. In §4, we have shown that the theoretical Coding+MAC gains for the above toy topologies are 2, 2, and 4 respectively. These numbers are fairly close to the numbers we observe in actual measurements.

One may wonder why the measured throughput gains are smaller than the theoretical Coding+MAC gain bounds. The XOR headers add a small overhead of 5-8%. However, the difference is mainly due to imperfect overhearing and flow asymmetry. Specifically, the nodes do not overhear all transmitted packets. Further, some senders capture the wireless channel sending more traffic in a particular direction, which reduces coding opportunities and overall gain.

In practice, traffic is a combination of congestion-controlled and uncontrolled flows. Further, most TCP flows are short-lived and do not fully exercise congestion control during slow-start. Thus, one would expect COPE’s gains to be higher than those observed with long-lived TCP and lower than those observed with UDP. Indeed, we have run experiments for the Alice-and-Bob scenario with short-lived TCP flows with Poisson arrivals and Pareto transfer size. Depending on the flow inter-arrival times, the measured throughput gains vary between the coding gain and the Coding+MAC gain.

7.4 COPE in an Ad Hoc Network

How does COPE perform in a wireless mesh network? We have advocated a simple approach to wireless network coding where each node relies on its local information to detect coding opportunities, and when possible XORs the appropriate packets. However, it is unclear how often such opportunities arise in practice, and whether they can be detected using only local information. Thus, in this section, we run experiments on our 20-node testbed to gauge the throughput increase provided by COPE in an ad hoc network.

7.4.1 TCP

We start with TCP flows that arrive according to a Poisson pro-




Figure 10—End-to-end loss rate and average queue size at the bottlenecks for the TCP flows in the testbed. Loss rates are as high as 14% even after 15 MAC retries; TCP therefore performs poorly. The queues at the bottlenecks almost never build up resulting in very few coding opportunities and virtually no gains.

cess, pick sender and receiver randomly, and transfer files whose sizes follow the distribution measured on the Internet [9].

Surprisingly, in our testbed, TCP does not show any significant improvement with coding (the average gain is 2-3%). The culprit is TCP’s reaction to collision-related losses. There are a number of nodes sending packets to the bottleneck nodes, but they are not within carrier sense range of each other, resulting in the classic hidden terminals problem. This creates many collision-related losses that cannot be masked even with the maximum number of MAC retries. To demonstrate this point, we repeat the TCP experiments with varying number of MAC retransmissions with RTS/CTS enabled. Note that disabling RTS/CTS exacerbates the problem further. Fig. 10 plots the end-to-end loss rates for TCP flows as a function of the number of MAC retransmissions. *These experiments have COPE turned off.* Even after 15 MAC retries (the maximum possible) the TCP flows experience 14% loss. As a result, the TCP flows suffer timeouts and excessive back-off, and are unable to ramp up and utilize the medium efficiently. Fig. 10 plots the average queue sizes at the bottleneck nodes.¹ The bottleneck nodes never see enough traffic to make use of coding; most of their time is spent without any packets in their queues or just a single packet. Few coding opportunities arise, and hence the performance is the same with and without coding.

Collision-related losses are common in wireless networks and recent work has studied their debilitating effect on TCP [14, 8]. Making TCP work in such a setting would imply solving the collision problem; such a solution is beyond the scope of this paper.

Would TCP be able to do better with COPE if we eliminated collision-related losses? We test the above hypothesis by performing

¹The few nodes connecting the two floors are where the flows intersect; they are main bottlenecks in our testbed.

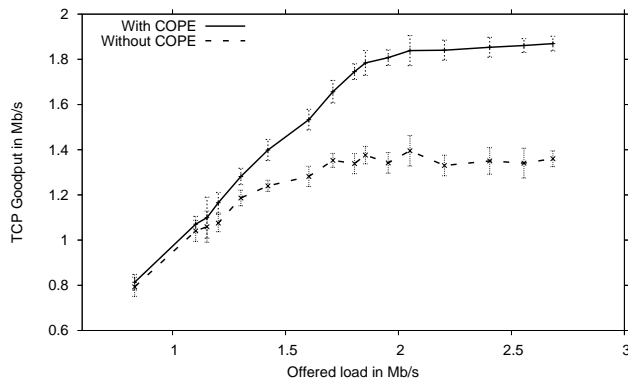


Figure 11—COPE provides 38% increase in TCP goodput when the testbed topology does not contain hidden terminals.

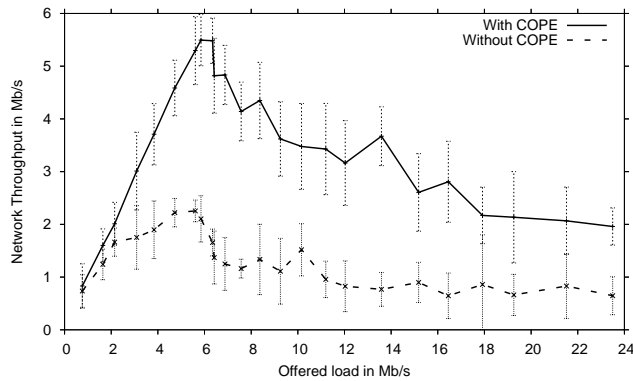


Figure 12—COPE can provide a several-fold (3-4x) increase in the throughput of wireless Ad hoc networks. Results are for UDP flows with randomly picked source-destination pairs, Poisson arrivals, and heavy-tail size distribution.

the following experiment. We compress the topology of the testbed by bringing the nodes closer together, so that they are within carrier sense range. We artificially impose the routing graph and inter-node loss rates of the original testbed. The intuition is that the nodes are now within carrier sense range and hence can avoid collisions. This will reduce the loss rates and enable TCP to make better use of the medium. We repeat the above experiment with increasing levels of congestion obtained by decreasing the inter-arrival times of the TCP flows. Fig. 11 plots the network TCP goodput with and without COPE as a function of the demand. For small demands, COPE offers a slight improvement since coding opportunities are scarce. As the demands increase, network congestion and coding opportunities increase, leading to higher goodput gains. As congestion increases beyond a certain level, the throughput levels off, reflecting the fact that the network has reached its capacity and cannot sustain additional load. At its peak, COPE provides 38% improvement over no coding. The medium loss rates after retransmissions are negligible. The TCP flows are therefore able to use the medium efficiently, providing coding opportunities which result in throughput gains.

7.4.2 UDP

We repeat the large scale testbed experiments with UDP. The flows again arrive according to a Poisson process, pick sender and receiver randomly, and transfer files whose sizes follow the distribution measured on the Internet [9]. We vary the arrival rates of the Poisson process to control the offered load. For each arrival rate, we run 10 trials, with coding on and then off (for a total of 500 experiments), and compute the network throughput in each case.

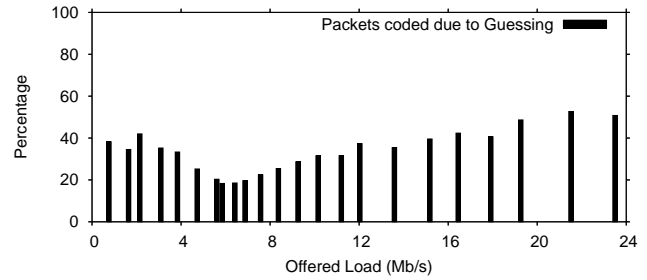


Figure 13—Percentage of packets coded in the testbed due to guessing, as a function of offered load, for the set of experiments in Fig. 12.




Figure 14—Distribution of number of packets coded together in the test bed at the peak point of Fig. 12.

Fig. 12 shows that COPE greatly improves the throughput of these wireless networks, by a factor of 3-4x on average. The figure plots the aggregate end-to-end throughput as a function of the demands, both with COPE and without. At low demands (below 2Mb/s), coding opportunities are scarce, and COPE performs similarly to no coding. As demands increase, both network congestion and the number of coding opportunities increase. In such dense networks, the performance without coding deteriorates because of the high level of contention and consequent packet loss due to collisions. In contrast, coding reduces the number of transmissions, alleviates congestion, and consequently yields higher throughput.

It is interesting to examine how much of the coding is due to guessing, as opposed to reception reports. Fig. 13 plots the percentage of packets that have been coded because of guessing for the experiments in Fig.12. It is calculated as follows: If n packets are coded together, and at most k packets could be coded using reception reports alone, then $n - k$ packets are considered to be coded due to guessing. The figure shows that the benefit of guessing varies with demands. At low demands, the bottleneck nodes have small queues, leading to a short packet wait time. This increases dependence on guessing because reception reports could arrive too late, after the packets have been forwarded. As demands increase, the queues at the bottlenecks increase, resulting in longer wait times, and consequently allowing more time for reception reports to arrive. Hence, the importance of guessing decreases. As demands surge even higher, the network becomes significantly congested, leading to high loss rates for reception reports. Hence, a higher percentage of the coding decisions is again made based on guessing.

Let us now examine in greater detail the peak point in Fig. 12, which occurs when demands reach 5.6 Mb/s. Fig. 14 shows the PDF of the number of native packets XOR-ed at the bottleneck nodes (i.e., the nodes that drop packets). The figure shows that, on average, nearly 3 packets are getting coded together. Due to the high coding gain, packets are drained much faster from the queues of the bottleneck nodes. The result is an average throughput gain of 3-4x.

7.5 COPE in a Mesh Access Network

There is growing interest in providing cheap Internet access using




Figure 15—COPE’s throughput gain as a function of the ratio of uplink to downlink traffic at in a congested mesh access network.

multi-hop wireless networks that connect to the rest of the Internet via one or more gateways/access points [1, 4, 32, 36]. We evaluate COPE in such a setting, where traffic is flowing to and from the closest gateway. We divide the nodes in the testbed into 4 sets. Each set communicates with the Internet via a specific node that plays the role of a gateway. We use UDP flows,² and control the experiments by changing the ratio of upload traffic to download traffic. Fig. 15 plots the throughput gains as a function of this ratio.

The throughput gain increases as the fraction of uplink traffic increases. When the amount of uplink traffic is small, gains are correspondingly modest; around 5 – 15%. As uplink traffic increases, gains increase to 70%. COPE’s throughput gain relies on coding opportunities, which depend on the diversity of the packets in the queue of the bottleneck node. For example, in the Alice-and-Bob topology, if only 10% of the packets in the bottleneck queue are from Alice and 90% from Bob, then coding can at best sneak Alice’s packets out on Bob’s packets. Hence, as the ratio of uplink traffic increases, the diversity of the queues at bottlenecks increases, more coding opportunities arise, and consequently higher throughput gains are obtained.

7.5.1 Fairness

The access network experiment above illuminates the effect fairness has on coding opportunities. An important source of unfairness in wireless networks is the comparative quality of the channels from the sources to the bottleneck, usually referred to as the *capture effect*. For example, in the Alice and Bob experiment, if the channel between Alice and the router is worse than that between Bob and the router, Alice might be unable to push the same amount of traffic as Bob. Although the 802.11 MAC should give a fair allocation to all contenders, the sender with the better channel (here Bob) usually captures the medium for long intervals. The routing protocol tries to discount the capture effect by always selecting the stronger links; but in practice, capture always happens to some degree.

We study the effect of capture on COPE by intentionally stressing the links in the Alice and Bob topology. We set it up such that both Alice and Bob are equidistant from the router, and compute the total network throughput. We then gradually move Alice’s node away from the router, and repeat the experiment and the measurements.

Fig. 16 shows the network throughput as a function of the ratio of Alice’s and Bob’s distance to the router. It also shows the percentage of coded packets and the *fairness index*, computed as the ratio of Alice’s throughput to Bob’s. As Alice moves further away, Bob increasingly captures the channel, reducing fairness, coding opportunities, and the aggregate network throughput. Interestingly, without coding, fairness and efficiency are conflicting goals; throughput increases if the node with the better channel captures the medium and

²As mentioned earlier, in the uncompressed testbed, TCP backs off excessively because of collision-based losses from hidden terminals, and does not send enough to fully utilize the medium.




Figure 16—Effect of unequal channel qualities on coding opportunities and throughput gain in the Alice-and-Bob topology. COPE aligns the fairness and efficiency objectives. Increased fairness increases coding opportunities and hence improves the aggregate throughput.

sends at full blast. Coding, however, aligns these two objectives; increasing fairness increases the overall throughput of the network.

8. DISCUSSION AND CONCLUSION

Finally, we would like to comment on the scope of COPE. The present design targets stationary wireless mesh networks, where the nodes are not resource-constrained. More generally, COPE can be used in multi-hop wireless networks that satisfy the following:

- *Memory*: COPE’s nodes need to store recently heard packets for future decoding. Only packets in flight are used in coding; there is no need to store packets that have already reached their destination. Consequently, the storage requirement should be slightly higher than a delay-bandwidth product. (For e.g., an 11 Mb/s network with a 50ms RTT has a delay-bandwidth product of 70 KB.)
- *Omni-directional antenna*: Opportunistic listening requires omni-directional antennas to exploit the wireless broadcast property.
- *Power requirements*: Our current design of COPE does not optimize power usage and assumes the nodes are not energy limited.

The ideas in COPE may be applicable beyond WiFi mesh networks. Note that COPE can conceptually work with a variety of MAC protocols including WiMax and TDMA. One may envision modifying COPE to address the needs of sensor networks. Such a modification would take into account that only a subset of the sensor nodes is awake at any point of time and can participate in opportunistic listening. Sensor nodes may also trade-off saved transmissions for reduced battery usage, rather than increased throughput. Additionally, COPE may be useful for cellular relays. Deploying cellular base stations is usually expensive. A cheap way to increase coverage is to deploy relay nodes that intervene between the mobile device and the base station [13], creating a multi-hop cellular backbone. COPE would allow cellular relays to use the bandwidth more efficiently. Indeed, after the publication of COPE, we have learned that Ericsson has independently proposed a design for cellular relays with a subset of COPE’s functionality, where the cellular relay XORs only duplex flows, as in the Alice-and-Bob scenario [13]. This scheme can be extended to make full usage of the ideas embedded in COPE.

Our community knows a few fundamental approaches that can improve wireless throughput, including more accurate congestion control, better routing, and efficient MAC protocols. We believe that COPE is an important step forward in our understanding of the potential of wireless networks because it presents a new orthogonal axis that can be manipulated to extract more throughput; namely, how to maximize the amount of data delivered in a single transmission. This is coding, which is an old theme, traditionally used at the physical and application layers. But COPE and a few other recent projects [7, 21] introduce coding to the networking community as a practical tool that can be integrated with forwarding, routing, and reliable delivery.

Acknowledgments

We thank Szymon Chachulski, John Bicket, Sanjit Biswas, and our shepherds, John Byers and Scott Shenker, for their insightful comments, and Leo for providing pleasurable distractions during many long nights of work. Katti and Katabi are supported by a NSF Career Award CNS-0448287, while Rahul is supported by an Intel gift. The opinions and findings in this paper are those of the authors and do not necessarily reflect the views of NSF or Intel.

9. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *ACM SIGCOMM, 2004*.
- [2] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung. Network Information Flow. In *IEEE Transactions on Information Theory*, 2000.
- [3] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A comparison of mechanisms for improving TCP performance over wireless links.
- [4] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. In *HotNets*, 2003.
- [5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *ACM MobiCom, 2005*.
- [6] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. *ACM SIGCOMM, 2005*.
- [7] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. More: Exploiting spatial diversity with network coding. In *MIT CSAIL Technical Report*, 2006.
- [8] C. cheng Chen, E. Seo, H. Kim, and H. Luo. Self-learning collision avoidance for wireless networks. In *Proceedings of IEEE INFOCOM, 2006*.
- [9] M. E. Crovella, M. S. Taqqu, and A. Bestavros. Heavy-tailed probability distributions in the World Wide Web. In R. J. Adler, R. E. Feldman, and M. S. Taqqu, editors, *A Practical Guide To Heavy Tails*, chapter 1, pages 3–26. Chapman and Hall, New York, 1998.
- [10] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *ACM MobiCom '03*, San Diego, California, September 2003.
- [11] S. Deb, M. Effros, T. Ho, D. R. Karger, R. Koetter, D. S. Lun, M. Médard, and N. Ratnakar. Network coding for wireless applications: A brief tutorial. In *IWWAN, 2005*.
- [12] R. Draves, J. Padhye, and B. Zill. Comparison of Routing Metrics for Multi-Hop Wireless Networks. In *Proceedings of ACM SIGCOMM, 2004*.
- [13] Definition and assessment of relay based cellular deployment concepts for future radio scenarios considering 1st protocol characteristics. Chapter 5. <https://www.ist-winner.org/DeliverableDocuments/D3.4.pdf>.
- [14] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla. The impact of multihop wireless channel on tcp throughput and loss. In *INFOCOM 2003*.
- [15] M. Heusse, F. Rousseau, R. Guillier, and A. Duda. Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans. In *ACM SIGCOMM, 2005*.
- [16] T. Ho and R. Koetter. Online incremental network coding for multiple unicasts. In *DIMACS Working Group on Network Coding, 2005*.
- [17] T. Ho, R. Koetter, M. Médard, D. Karger, and M. Effros. The Benefits of Coding over Routing in a Randomized Setting. In *ISIT, 2003*.
- [18] T. Ho, B. Leong, Médard, R. Koetter, Y. Chang, and M. Effros. The Utility of Network Coding in Dynamic Environments. In *IWWAN, 2004*.
- [19] p. a. IEEE 802.11 WG. Wireless lan medium access control (mac) and physical layer (phy) specifications. *Standard Specification, IEEE, 1999*.
- [20] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. *IEEE Transactions on Information Theory*, 2003.
- [21] A. Kamra, J. Feldman, V. Misra, and D. Rubenstein. Growth codes: Maximizing sensor network data persistence. In *ACM SIGCOMM, 2006*.
- [22] B. Karp. *Geographic Routing for Wireless Networks*. PhD thesis, Harvard University, 2000.
- [23] S. Katti, D. Katabi, W. Hu, H. S. Rahul, and M. Médard. The importance of being opportunistic: Practical network coding for wireless environments, 2005.
- [24] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, 2003.
- [25] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The click modular router. *ACM Transactions on Computer Systems*, August 2000.
- [26] S. R. Li, R. W. Yeung, and N. Cai. Linear network coding. In *IEEE Transactions on Information Theory*, 2003.
- [27] D. S. Lun, M. Médard, R. Koetter, and M. Effros. Further results on coding for reliable communication over packet networks. In *IEEE International Symposium on Information Theory (ISIT 05)*, 2005.
- [28] D. S. Lun, N. Ratnakar, R. Koetter, M. Médard, E. Ahmed, and H. Lee. Achieving Minimum-Cost Multicast: A Decentralized Approach Based on Network Coding. In *IEEE INFOCOM, 2005*.
- [29] Internet packet size distributions: Some observations. <http://netweb.usc.edu/rsinha/pkt-sizes/>.
- [30] V. Paxson and S. Floyd. Wide-area traffic: the failure of poisson modeling. In *IEEE/ACM Transactions on Networking*, 3(3):226–244, 1995.

- [31] A. Ramamoorthy, J. Shi, and R. Wesel. On the capacity of network coding for wireless networks. In *41st Annual Allerton Conference on Communication Control and Computing*, Oct. 2003.
- [32] Nokia rooftop wireless routing. White Paper.
- [33] P. Sinha, T. Nandagopal, N. Venkitaraman, R. Sivakumar, and V. Bharghavan. WTCP: A reliable transport protocol for wireless wide-area networks. *Wireless Networks*, 8(2-3):301–316, 2002.
- [34] tcp. <http://ftp.arl.mil/ftp/pub/tcp/>.
- [35] udpgen. <http://pdos.csail.mit.edu/click/ex/udpgen.html>.
- [36] Wireless Community Network List. <http://www.toaster.net/wireless/community.html>.
- [37] Y. Wu, P. A. Chou, and S. Y. Kung. Information Exchange in Wireless Networks with Network Coding and Physical-layer Broadcast. *MSR-TR-2004-78*.
- [38] Z. Li and B. Li. Network Coding in Undirected Networks. In *CISS 04, 2004*.
- [39] Z. Li and B. Li. Network coding: The case for multiple unicast sessions. In *Allerton Conference on Communications*, 2004.

APPENDIX

A. PROOF OF THEOREM 4.1

PROOF. We first prove the upper bound of 2. Note that if the intermediate node codes N native packets together, these packets have to be to N different next-hops, by the coding rule of §3(b). In the absence of opportunistic listening, the only neighbor that has a packet is the previous hop of that packet. Suppose the intermediate hop codes ≥ 2 packets from the same neighbor. All other neighbors must have $\leq N - 2$ packets in the encoded packet, which violates the coding rule. As a result, the intermediate hop can code at most one packet from a neighbor. Without opportunistic listening, this is the only native packet in the encoded packet that this neighbor has. Invoking the coding rule, this implies that the intermediate hop can code at most 2 packets together. This implies that the total number of transmissions in the network can at most be halved with coding, for a coding gain of 2.

Indeed, this gain is achievable in the chain of N links in Fig. 4(a). This topology is an extension of the Alice-and-Bob example where $N = 2$. The no-coding case requires a total of $2N$ transmissions to deliver a packet from Alice to Bob, and vice-versa. On the other hand, in the presence of coding, each of the $N - 1$ intermediate nodes on the path can transmit information simultaneously to neighbors on either side by coding the two packets traversing in opposite directions, for a total of $N + 1$ transmissions. The coding gain in this case is $\frac{2N}{N+1}$, which tends to 2 as the chain length grows. \square

B. PROOF OF THEOREM 4.2

PROOF. As proved above, in the absence of opportunistic listening, a node can code at most 2 packets together. Hence, a bottleneck node can drain its packets at most twice as fast, bounding the Coding+MAC gain at 2. This gain is achieved even in the simple Alice-and-Bob experiment as explained above (longer chains result in the same Coding+MAC gain). \square

C. PROOF OF THEOREM 4.3

PROOF. Consider the wheel topology with radius r in Fig. 4(d) with N nodes uniformly placed on the circumference, and one node at the center of the circle. Assume that when a node transmits, all other nodes in the circle overhear this transmission, except for the diametrically opposed node (i.e., the radio range is $2r - \epsilon$, where $\epsilon \approx 0$). Suppose now that there are flows between every pair of diametrically opposed nodes. Note that nodes on either end of a diameter cannot communicate directly, but can communicate using a two-hop route through the middle node. In fact, this route is the geographically shortest route between these nodes. In the absence of coding, a single flow requires 1 transmission from an edge node, and 1 transmission from the middle node. This adds to a total of 1 transmission per edge node, and N transmissions for the middle node, across all packets. Since the MAC gives each node only a $\frac{1}{N+1}$ share of the medium, the middle node is the bottleneck in the absence of coding. However, COPE with opportunistic listening allows the middle node to code all the N incoming packets and fulfill the needs of all flows with just one transmission, thereby matching its input and output rates. Hence, the Coding+MAC gain is N , which grows without bound with the number of nodes. \square