Practical Network Coding in Wireless Networks

Qunfeng Dong*, Jianming Wu
University of Wisconsin
Madison, WI, USA
{qunfeng, jianming}@cs.wisc.edu

Wenjun Hu, Jon Crowcroft
University of Cambridge
Cambridge, UK
{wh214, jac22}@cam.ac.uk

ABSTRACT

Network coding is seen as a promising technique to improve network throughput. In this paper, we study two important problems in localized network coding in wireless networks, which only requires each node to know about and coordinate with one-hop neighbors. In particular, we first establish a condition that is both necessary and sufficient for useful coding to be possible. We show this condition is much weaker than expected, and hence allows a variety of coding schemes to suit different network conditions and application preferences. Based on the understanding we establish, we are able to design a robust coding technique called *loop coding* that can improve network throughput and TCP throughput simultaneously.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*

General Terms

Analysis, Design, Performance, Reliability

Keywords

Network Coding, TCP, Throughput, Wireless Networks

1. INTRODUCTION

To improve network throughput, the idea of *network coding* [1] has been proposed for forwarding nodes to mix the bits in forwarded packets. In this work, we focus on one specific type of network coding in wireless networks, where we XOR packets for unicast flows. This type of network coding has recently received a lot of practical interest for its ease of implementation and the importance of unicast communication.

The basic idea of network coding can be illustrated using the Alice-and-Bob scenario [4] in Figure 1, where Alice wants to send packet P_1 to Bob and Bob wants to send packet P_2 to Alice. They rely on a *relay* in the middle to exchange packets. In the terminology of network coding, a non-encoded original packet (such as P_1 and P_2) is referred to as a *native packet*. *Network coding is about what packet(s) should the relay transmit, in order for the native packets to be obtained by their intended receiver(s).*

*Qunfeng Dong was supported in part by NSF grants CNS-0639434, CNS-0627589, CNS-0627102, and CNS-0520152.

MobiCom'07, September 9–14, 2007, Montréal, Québec, Canada.
ACM 978-1-59593-681-3/07/0009.

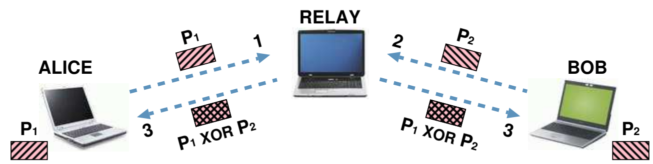

Figure 1: The basic Alice-and-Bob scenario.

Encode at relay The relay XORs P_1 and P_2 together (with padding if necessary) and broadcasts $P_1 \oplus P_2$, which we refer to as an *XOR packet*.

Decode at receiver(s) Upon receiving $P_1 \oplus P_2$, Alice can decode P_2 by $P_2 = P_1 \oplus (P_1 \oplus P_2)$. Similarly, Bob can decode P_1 by $P_1 = (P_1 \oplus P_2) \oplus P_2$.

Thus, in order to relay P_1 and P_2 to their intended receiver, the relay only needs to transmit one packet (i.e., the XOR packet) instead of two (i.e., P_1 and P_2).

Later on, Katti *et al.* proposed COPE [2] to exploit overhearing in wireless networks. In COPE, each node overhears all native packets transmitted by its neighbors. For example, in Figure 2 node E overhears packet P_1, which is not addressed to itself. COPE allows the relay to XOR a set of k native packets $P_1, P_2, \cdots, P_k$ into one XOR packet if the receiver of each native packet P_i has the $k-1$ native packets other than P_i. After receiving the XOR packet, the receiver of each native packet P_i will be able to decode P_i by $P_i = (P_1 \oplus P_2 \oplus \cdots \oplus P_k) \oplus P_1 \oplus P_2 \oplus \cdots \oplus P_{i-1} \oplus P_{i+1} \cdots \oplus P_k$.

COPE is applicable in more cases. However, it is still often the case that useful network coding is actually possible while COPE is oblivious of it. This can be illustrated using the example in Figure 2, where COPE is not applicable. Nevertheless, instead of forwarding P_1, P_3, and P_4 individually, the relay R actually only needs to broadcast two XOR packets: $P_1 \oplus P_3$ and $P_3 \oplus P_4$. Once C, D, and E have correctly received these XOR packets, they can decode the native packets addressed to themselves.

In light of this initial observation, we are interested to study two important problems.

2. APPLICABILITY

We have illustrated in Figure 2 that COPE has limited applicability. But after all, COPE is just one practice of network coding. To really understand the applicability of network coding in general, it remains to establish the necessary condition and sufficient condition for useful network coding to be possible. In this section, we establish such the following unique condition that is both necessary and sufficient, and demonstrate during its proof how to design a useful network coding scheme, if this condition is satisfied.

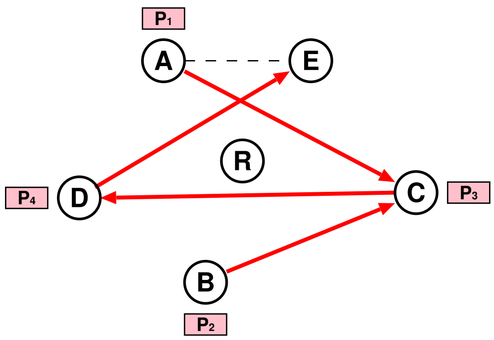

	sender	receiver	key set
P_1	A	C	$\{P_3\}$
P_2	B	C	$\{P_3\}$
P_3	C	D	$\{P_4\}$
P_4	D	E	$\{P_1\}$

Figure 2: The relay R can hear and be heard by all. Two surrounding nodes can hear each other if they are connected by a dashed line.

DEFINITION 1 (THE APPLICABILITY CONDITION). *There exists some non-empty subset $P' \subseteq P$ of native packets such that for any native packet $P_i \in P'$, $K_i \cap P' \neq \phi$.*

It will become clear in Section 3 that the understanding we acquire in this section is fundamental to further study.

2.1 Definitions

In this work, we study localized network coding, where each node only needs to know about and coordinate with one-hop neighbors. It has become the theme of recent research [4, 2] because it is more practical and scalable, is the building block to construct network-wide solutions, and is incrementally deployable.

In this localized problem setting, there is a relay R surrounded by a set of m communicating nodes, denoted by $N = \{N_1, N_2, \cdots, N_m\}$. Among these m nodes, there are a set $P = \{P_1, P_2, \cdots, P_n\}$ of n native packets, each being sent by some node N_i to some other node N_j via the relay R. We refer to N_i as the *sender* and N_j as the *receiver* of that native packet, respectively. Without loss of generality, we also assume native packets are of the same size. For each native packet P_i, we refer to the subset of native packets available to the receiver of P_i (before the relay transmits any packet) as the *key set* of P_i, denoted by $K_i \subseteq P$. Native packets in the key set of P_i are either overheard or transmitted by the receiver of P_i. As we only need to consider native packets that have to be forwarded by the relay, that implies $P_i \notin K_i$. Figure 2 illustrates an example.

Relevant to each network coding problem are the sender, receiver, and key set of each native packet. Given this information, a network coding scheme (or coding scheme for short) basically defines a set of XOR packets to be transmitted by the relay. If an XOR packet is obtained by XORing k native packets together, we call it a *k-ary XOR packet* that *contains* those k native packets. For ease of discussion, we consider each native packet individually forwarded by the relay to be a *unary* XOR packet. If a coding scheme does not define any l-ary XOR packet such that $l > k$, we call this coding scheme a *k-ary coding scheme*. A coding scheme is considered *feasible*, if correct reception of its defined set X of XOR packets enables each native packet to be decoded by its receiver (using its key set). If the number of XOR packets defined by a feasible coding scheme is less than the number of native packets to be forwarded by the relay, we call the coding scheme a *useful* one.

2.2 The applicability condition is sufficient

If the applicability condition is satisfied, there must exist a useful coding scheme, *regardless of the network topology (i.e., overhearing relationship among nodes) and flow configuration.* In fact, it is true even if we are restricted to use binary coding schemes only, i.e., coding scheme where each XOR packet is obtained by XORing at most two native packets together. We prove this by considering the coding/decoding process in a graph theoretic setting.

DEFINITION 2 (CODING GRAPH). *Given a binary coding scheme which defines a set X of XOR packets, its* coding graph *$G_X = (P, E)$ is defined as follows.*

- *The vertex set P is exactly the set P of native packets.*
- *Two vertices share an edge if they are XORed together.*

LEMMA 1. *Using a binary coding scheme which defines a set X of XOR packets, a native packet P_i can be decoded by its receiver if there exists some native packet $P_j \in K_i$ such that P_j and P_i are connected in G_X.*

PROOF. Suppose the edges in the path from P_j to P_i are (P_j, P_{i_1}), $(P_{i_1}, P_{i_2}), \cdots, (P_{i_k}, P_i)$ in order. The decoding process at the receiver of P_i is essentially a hop-by-hop "walk" from P_j to P_i along the path as follows.
Using P_j and $P_j \oplus P_{i_1}$, we can reach P_{i_1};
Using P_{i_1} and $P_{i_1} \oplus P_{i_2}$, we can reach P_{i_2};
......
Using P_{i_k} and $P_{i_k} \oplus P_i$, we can reach P_i. □

LEMMA 2. *If the applicability condition is satisfied, we must be able to define a useful binary coding schemes.*

PROOF. Let $P' = \{P_{i_1}, P_{i_2}, \cdots, P_{i_k}\}$ be a non-empty subset of k native packets that satisfies the applicability condition. The follow binary coding scheme is a useful one.

- Transmit each native packet $P_i \notin P'$ separately.
- For the k native packets in P', transmit $k-1$ binary XOR packets $P_{i_1} \oplus P_{i_2}, P_{i_2} \oplus P_{i_3}, \cdots, P_{i_{k-1}} \oplus P_{i_k}$.

Because in the corresponding coding graph G_X, the k native packets in P' are connected into a chain. By the definition of P', for any native packet $P_i \in P'$, there exists some other native packet $P_j \in K_i$ that is also in P'. It follows by Lemma 1 that P_i can be decoded by its receiver. □

Compared with the applicability condition, the condition stated in the following corollary is more intuitively clear. This condition is actually stronger than the applicability condition. Specifically, Corollary 1 requires $P' = P$. Indeed, the applicability condition as a sufficient condition is very weak.

COROLLARY 1. *We can devise a useful coding scheme if each receiver has transmitted or overheard one native packet.*

Revisiting COPE Now we are better informed to understand the limitation of COPE. In COPE, each native packet has to be decoded using a single XOR packet. As we have pointed out in the proof of Lemma 1, the decoding process of any native packet P_i is essentially a hop-by-hop "walk" from some other native packet $P_j \in K_i$ to P_i, along the path connecting them in the coding graph. As each XOR packet corresponds to an edge in the coding graph[1], COPE essentially looks as far as only one hop away. The applicability of COPE is thus limited.

[1]By saying this, we actually cheat a bit here since edges in the coding graph we have defined represent binary XOR packets only. Nevertheless, we believe the intuition is there and the anology is hopefully not twisted.

2.3 The applicability condition is necessary

We now prove the applicability condition is necessary, for which we first prove the following auxiliary lemma.

LEMMA 3. *Given an arbitrary feasible coding scheme that defines a set X^0 of XOR packets, we can always transform it into a new feasible coding scheme which defines a new set X of XOR packets, such that:*

I. Let Q denote the set of unary XOR packets in X and $P' = P - Q$. Every native packet P_i with an empty intersection $K_i \cap P'$ is contained in Q as a unary XOR packet.

II. The XOR packets in $X - Q$ contain and only contain native packets in P'.

III. $|X| \leq |X^0|$.

PROOF. We start with $X = X^0$ and iteratively modify X to satisfy the stated properties.

Property II: Given $P' \cap Q = \phi$ and the fact that every native packet in P' can be decoded by its receiver, we know every native packet in P' must be contained in some XOR packet(s) in $X - Q$. On the other hand, notice that native packets in Q are available to the receivers. Whenever an XOR packet X_i in $X - Q$ contains some native packet P_i in Q, we can safely remove P_i from X_i, by XORing X_i with P_i. Because whenever the original X_i is needed, receivers can always obtain it by XORing P_i back into X_i. Therefore, property II can always be satisfied.

Property I: To prove property I, let us assume there exists some violating native packet $P_i \in P'$ with an empty intersection $K_i \cap P'$. Since $K_i \cap P' = \phi$ and $Q \cap P' = \phi$, we know $(K_i \cup Q) \cap P' = \phi$. Therefore, the fact that P_i can be decoded by its receiver means there exists a subset $X' \subseteq X - Q$ of XOR packets such that

$$\bigoplus_{X_i \in X'} X_i = P_i.$$

Because by property II, the XOR packets in $X - Q$ only contain native packets in P', while the receiver of P_i only has native packets in $K_i \cup Q$ and hence has no native packets in P'. Because $(K_i \cup Q) \cap P' = \phi$. If for any $X' \subseteq X - Q$, $\bigoplus_{X_i \in X'} X_i$ either does not contain P_i at all or contains some native packet $P_j \in P'$ other than P_i, then the receiver of P_i will have no way to decode P_i.

Apparently, X' contains at least one XOR packet X_i that contains P_i. XORing both sides of the above equation with all XOR packets in X' except X_i gives us

$$X_i = P_i \oplus \left(\bigoplus_{X_j \in X' - \{X_i\}} X_j \right).$$

Thus, we can add P_i to Q and safely remove X_i from $X - Q$. Because whenever X_i is needed, receivers can always obtain X_i using the above formula. To preserve the fact that $P' = P - Q$, we remove P_i from P'. Then to preserve property II, we also remove P_i from every XOR packet in $X - Q$ that still contains P_i, which we have shown to be safe in the proof of property II. Thus far, we have safely removed a violating native packet P_i from P'.

Property III: The total number of XOR packets will not increase since we removed at least one packet from $X - Q$ and added only one packet to Q. Property III is thus satisfied.

If needed, we can repeat the above process to remove all such violating native packets from P'. In the worst case, $P' = \phi$ and $Q = P$ will satisfy all the properties. (Note that the applicability condition requires P' to be non-empty.) □

LEMMA 4. *There exists a useful network coding scheme only if the applicability condition is satisfied.*

PROOF. Assume there exists a useful network coding scheme, which defines a set X of XOR packets. Given this assumption, we show we must be able to find a *non-empty* subset $P' \subseteq P$ that satisfies the applicability condition. Initially, let $Q \subseteq X$ denote the set of unary packets in X and let $P' = P - Q$. We follow the techniques described in the proof of Lemma 3 to remove violating native packets from P'. In this case, keep removing violating native packets from P' is guaranteed to give us a non-empty subset P', which satisfies the stated properties in Lemma 3 and hence the applicability condition.

Since the coding scheme is a useful one, it is clear that $Q \subset P$. Moreover, $P' = P - Q$ initially contains at least two native packets. Because if P' contains only one native packet, we know by Lemma 3 that $X - Q$ contains at least one XOR packet while Q contains $n - 1$ packets, which contradicts the fact that the coding scheme is a useful one.

If P' does not satisfy the applicability condition, we keep removing violating native packets from P' until the applicability condition is satisfied, or P' contains exactly two native packets P_i and P_j. In the latter case, since the given coding scheme is a useful one and we never increase the number of XOR packets, we know $X - Q$ contains only one XOR packet. By property II, that XOR packet must be $P_i \oplus P_j$. By property I, P_i and P_j are not in Q. Therefore, P_i and P_j must have P_j and P_i in their key set, respectively, so that they can be decoded by their receivers. That means P' satisfies the applicability condition. □

3. RELIABILITY

Packet loss is very common in wireless networks. In [3], De Couto *et al.* have reported that most available wireless links exhibit significant link layer loss rates. Therefore, besides defining a minimum number of XOR packets (to be transmitted by the relay), another effective way to improve network throughput is to reduce packet loss rate and hence the number of packet retransmissions.

Moreover, reducing packet loss rate is of independent importance — TCP throughput significantly degrades even at modest packet loss rates. This is because the congestion control mechanism of TCP interprets packet loss as a signal of network congestion and halves TCP congestion window upon reception of every such signal. To effectively improve TCP throughput, merely minimizing the number of XOR packets (to be transmitted by the relay), which is the sole purpose of previously proposed network coding schemes such as COPE, is no longer enough. In fact, with the employment of overhearing (e.g. in COPE) in network coding, this packet loss problem becomes especially challenging. Because on one hand, XOR packets transmitted by the relay must be reliably delivered to intended receivers. On the other hand, receivers may also need to correctly overhear some native packets to be able to decode received XOR packets.

In the network protocol stack, network coding schemes operate between the wireless link layer and higher layer protocols (such as TCP). Therefore, it is possible to design some robust coding schemes to mask the underlying link layer loss rate from higher layer protocols. (To facilitate discussion, we refer to the probability that a native packet will not be correctly decoded by its receiver as the *post-coding loss rate*.) Given the dual importance of reducing packet loss rate on improving both network throughput and TCP throughput, an ideal solution would be to design a practical network coding scheme that can effectively reduce packet loss rate, without incurring additional communication cost. In Section 2, we have established the key understanding that in a coding scheme, a native packet can be decoded as long as in the coding graph it is connected with some native packet in its receiver's key set. Based on this understanding, we are now able to define such a class of robust coding schemes, called *loop coding*, meaning a coding scheme whose coding graph contains some loop(s).

For the example in Figure 3, where COPE and Alice-and-Bob are not applicable, we choose the loop coding scheme which defines three binary XOR packets: $P_1 \oplus P_2$, $P_2 \oplus P_3$, and $P_3 \oplus P_1$. Here,

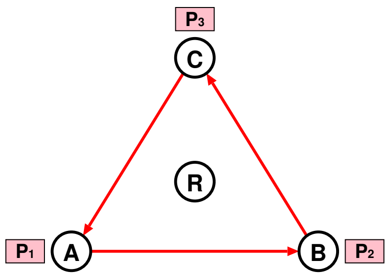

Figure 3: Node $A/B/C$ wants to send packet $P_1/P_2/P_3$ to node $B/C/A$, respectively. None of them is in the transmission range of each other. The relay can hear and be heard by everybody.

the loop coding scheme only requires the relay to transmit three XOR packets, the same as the number of native packets. However, the post-coding loss rate is effectively reduced. The intuitive explanation is that, for each native packet P_i in the loop, the loop contains two disjoint paths between P_i and the native packet $P_j \in K_i$ (i.e., P_j is sent by P_i's receiver). Thus, even if some link layer reason makes the receiver of P_i miss one XOR packet from the relay, which is one edge in the loop, it is still able to decode P_i since P_i and P_j are still connected in the residual coding graph.

To see how effectively the loop coding scheme can reduce post-coding loss rate, let us assume for simplicity but without loss of generality that all the wireless links have a link layer loss rate of p. If the relay forwards each native packet without coding, the post-coding loss rate is p. But using the loop coding scheme, P_1 can be decoded by its receiver B if node B either (1) receives $P_1 \oplus P_2$ or (2) receives both $P_2 \oplus P_3$ and $P_3 \oplus P_1$. Thus, the post-coding loss rate of P_1, namely the probability of failing both, is $p\left(1-(1-p)^2\right)$, which is less than the link layer loss rate p. For instance, $p\left(1-(1-p)^2\right) < 2\%$ when $p = 10\%$. The same analysis and calculation apply to P_2 and P_3 as well.

Compared with COPE, one may suspect that the loop coding scheme achieves such packet loss rates much lower than p because it transmits more XOR packets than COPE. This conjecture is not true. Even when COPE is applicable, no matter how many XOR packets COPE transmits, it can not avoid the probability of unsuccessful overhearing (except in the Alice-and-Bob scenario), which is already p. Consequently, the post-coding loss rate of COPE will never be lower than p, no matter how many (copies of) XOR packets it transmits. Without overhearing, COPE is only applicable in the Alice-and-Bob scenario. In contrast, our loop coding scheme is applicable in many other scenarios as well.

The above design of loop coding is also applicable to cases where the loop contains a different number of native packets. Furthermore, it is also applicable with overhearing. For instance, let us look at the example in Figure 4. Let us again consider the loop coding scheme which defines three binary XOR packets: $P_1 \oplus P_2$, $P_2 \oplus P_3$ and $P_3 \oplus P_1$. The post-coding loss rate of P_1 and P_2 is the same as in Figure 3. But to be able to decode P_3, its receiver

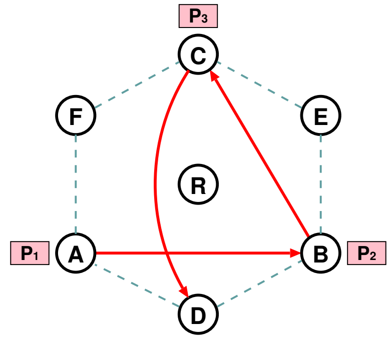

Figure 4: Node $A/B/C$ wants to send packet $P_1/P_2/P_3$ to node $B/C/D$, respectively. The relay R can hear and be heard by everybody. Two surrounding nodes can hear each other if they are connected by a dashed line.

D must overhear packet P_1. Then, it still needs to either receives $P_3 \oplus P_1$ or receives both $P_1 \oplus P_2$ and $P_2 \oplus P_3$. Therefore, the post-coding loss rate of P_3 is $p + (2p^2 - 3p^3 + p^4) < p + 2p^2$. For instance, when $p = 10\%$, the post-coding loss rate of P_1 and P_2 are both less than 2%, while the post-coding loss rate of P_3 is less than 12%. Their average post-coding loss rate is reduced from 10% to 5.3%.

To understand why loop coding is applicable in such different scenarios, we hereby introduce another auxiliary graph theoretic tool, called *key graph*.

DEFINITION 3 (KEY GRAPH). *Given a network coding problem, its* key graph *$G_K = (P, A)$ is defined as follows.*

- *The vertex set P is exactly the set P of native packets.*
- *Unlike coding graphs, the key graph is a directed graph. It contains a directed arc from P_i to P_j, if and only if P_j is in the key set of P_i.*

Given a directed cycle traversing k native packets in the key graph, we can simply define a loop coding scheme to encode these k native packets into a k-loop, as we did for the above examples. The examples in Figure 3 and 4 share the key graph in Figure 5 and hence the same loop coding scheme.

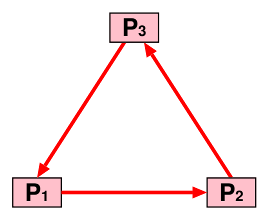

Figure 5: Key graph of the examples in Figure 3 and Figure 4.

4. REFERENCES

[1] R. Ahlswede, N. Cai, S.-Y. Li, and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4), 2000.

[2] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. XORs in the air: Practical wireless network coding. In *ACM SIGCOMM*, 2006.

[3] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *ACM MobiCom*, 2003.

[4] Y. Wu, P. A. Chou, and S. Y. Kung. Information exchange in wireless networks with network coding and physical-layer broadcast. Technical Report MSR-TR-2004-78, Microsoft Research, 2004.