RESEARCH

A modular framework for implementing joint wireless network coding and scheduling algorithms

Ioannis Broustis[1], Georgios S Paschos[1,2*], Dimitris Syrivelis[1], Leonidas Georgiadis[1,3] and Leandros Tassiulas[1,4]

Abstract

Intersession network coding (NC) can provide significant performance benefits via mixing packets at wireless routers; these benefits are especially pronounced when NC is applied in conjunction with intelligent scheduling. NC, however, imposes certain processing operations, such as encoding, decoding, copying, and storage. When not utilized carefully, all these operations can induce tremendous processing overheads in practical settings. Our testbed measurements suggest that such processing operations can severely degrade the router throughput, especially at high bit rates. Motivated by this, we design network coding framework for rate adaptive wireless links (NCRAWL). The design of NCRAWL facilitates low overhead NC operations, thereby effectively approaching the theoretically expected throughput benefits of joint NC and scheduling. We implement and evaluate NCRAWL on a wireless testbed. Our experiments demonstrate that NCRAWL meets the theoretical predicted throughput gain while requiring much less CPU processing, compared to related frameworks.

1 Introduction

Intersession network coding (NC) enables the local processing and mixing of independent traffic flows. Combining such flows at wireless routers can increase the available capacity [1-3]. However, such increase is evident only when (a) routers (which perform the encoding operations) are able to quickly identify efficient coding opportunities that increase the NC gain, (b) packet decoders are able to correctly decipher the encoded packets and acknowledge the decoded packets that they receive in diverse channel conditions, and (c) the overheads imposed due to the inclusion of additional packet headers as well as packet processing operations [4] are kept minimal. Moreover, while NC can increase the router throughput in random-access networks [5], prior studies have shown that when the packets are scheduled carelessly, NC may not offer significant benefits [2,6]. With multi-rate links and when decisions are made based on statistical information, scheduling is necessary to avoid packet losses. All

these factors should be taken into account when designing and developing practical, efficient NC algorithms and systems.

1.1 Implementing efficient wireless NC in practice is a challenging task

Although intersession NC can theoretically offer unprecedented wireless router capacity benefits, realizing these benefits in practical systems is by no means an easy task. Prior implemented efforts towards practical NC [3,7-9] have demonstrated throughput benefits at low transmission rates but have also discovered a series of complexity issues arising in such implementations. Our testbed measurements suggest that it is quite challenging to map the theoretically expected benefits offered by wireless NC in practical multi-rate deployments. This is due to two reasons, which motivate our study; we explain them below.

1.1.1 The overheads incurred by NC may be excessive

With NC, the routers need to be aware of the packets that have been successfully overheard by each neighbor, in order to further decide which packets to encode together and when. One method that has been proposed for addressing this requirement is by enforcing every

*Correspondence: gpasxos@mit.edu
[1]Center for Research and Technology Hellas, Volos, Greece
[2]LIDS, Massachusetts Institute of Technology, Cambridge, MA, USA
Full list of author information is available at the end of the article

neighbor into explicitly acknowledging overheard packets (e.g., see [3,7]). However, unless clever implementation techniques are employed, this approach may not improve the performance; this is because significant additional packet processing needs to be performed at routers, which may intrusively increase the already imposed processing overhead. Thus, although the wireless channel may be conducive to the use of high bit rates, routers may be incapable of transmitting as many packets in order to meet those rates. Our experiments with various NC implementations suggest that with overhead-agnostic design choices, the theoretical benefits offered by NC cannot be mapped in practice due to the excessive processing and network overheads that are imposed.

1.1.2 Scheduling for NC has not been exploited in practice
Previous studies have demonstrated the benefits of jointly applying NC and link scheduling [2,6]. However, prior practical implementations (1) have not incorporated any such techniques and (2) have not been designed to host scheduling algorithms in an efficient manner. This necessitates the design of a broad, although lightweight framework, which can facilitate the efficient coexistence of wireless NC and scheduling.

1.2 Our contributions
Towards addressing the above two issues, we design and implement network coding framework for rate adaptive wireless links (NCRAWL). NCRAWL has been optimized at each stage of the NC operations to utilize the router resources efficiently. It is a modular tool which can host the implementation of intersession NC schemes that are either stand-alone or tightly integrated with scheduling algorithms. The NCRAWL modules manage all the NC operations, such as encoding, decoding, storage, and routing in a lightweight manner, allowing for overhead-limited network operations. Note that to the best of our knowledge, our framework is the first to facilitate the practical coexistence of NC and scheduling. Theoretically shown throughput benefits can be easily assessed on NCRAWL and adapted for operating on real networks with limited effort. We implement NCRAWL on *Click* as a Linux kernel module [10]; we evaluate our framework on a wireless testbed via measurements, under various indoor and outdoor topological and traffic settings. Our experiments demonstrate that NCRAWL offers significant throughput improvements even at high bit rate regimes, where previously implemented schemes are unable to operate due to excessive computational and networking overheads.

1.3 The scope of our work
Our focus is not on proposing optimal scheduling policies for NC but on developing an accurate, lightweight and easy-to-deploy experimental tool that can host NC and scheduling schemes. Our system design and implementation decisions could be embedded into previously proposed platforms, such as COPE [3], ER [7], and CLONE [8]. However, we have chosen to implement NCRAWL from scratch, instead of trying to modify such platforms, only because it has been much easier to implement certain design choices in a particular, modular manner. Moreover, in this paper, we focus on local NC, where an encoded packet does not traverse more than one hop. This has direct applications in WLANs, where clients exchange data via their associated access point, thereby enabling high-bandwidth applications such as online gaming and video streaming. Note that NCRAWL can also be applied on top of multi-hop topologies with long routes, where native packets are potentially encoded/decoded multiple times at intermediate routers along their traversed routes.

The rest of the paper is structured as follows. In Section 2, we discuss relevant previous studies. In Section 3, we provide a high level overview of the considered NC scheme. In Section 4, we present the modular design and implementation of NCRAWL. In Section 5, we demonstrate the strengths of our design through a scheduling-driven case study. In Section 6, we assess the performance of NCRAWL via extensive testbed measurements. Finally, our conclusions form Section 7.

2 Related work
In this section, we discuss previous related NC studies and differentiate our work.

2.1 Experimental studies on wireless coding
Katti et al. [3] propose COPE, the seminal implementation of wireless NC. With COPE, routers are fully aware of packets that have been overheard by every neighbor. For this, each node is required to inform the router about overheard packets. Experiments with COPE on a wireless testbed show that even with very simple encoding operations, intersession NC can provide significant capacity gains. Rozner et al. in [7] present ER, a scheme that adopts the design of COPE and employs NC to perform efficient packet retransmissions. With ER, packets that need to be retransmitted are coded together, such that one retransmission can recover multiple packet losses. Kim et al. [11] extend the design of COPE to include NC-aware bit rate control and clever selection of nodes that acknowledge the reception of encoded packets. Rayanchu et al. [8] propose CLONE, a suit of algorithms for NC that take into account losses on wireless links. However, [7,8,11] all follow COPE's logic regarding the dissemination of information about which packets have been stored as keys. Moreover, these studies do not make online decisions about whether to enable coding or not, based on the link quality. MORE [9] is a routing protocol, where

routers perform random mixing of packets before forwarding them. MIXIT [12] encodes symbols rather than packets. Similarly to MORE, batches of packets are coded together. However, since a packet is a sequence of symbols, Intermediate relays use hints from the PHY layer in order to infer which symbols within a packet are correctly received with high probability.

The above experimental approaches differ from ours in two main ways: (1) We consider the use of stochastic information for overhearing instead of blindly acknowledging each particular packet. This allows for efficient implementation and avoids computationally expensive packet processing operations. Note that our design choices can also be applied to most of the above approaches. (2) NCRAWL is a 'tool' for studying problems regarding joint NC and scheduling with feedback, and can potentially be intertwined with an optimal algorithm to provide the best solution within the class of implicit ACK schemes. To the best of our knowledge, our work is the first to provide a coherent, lightweight framework to practically assess joint NC and scheduling schemes.

2.2 Analytical and simulation NC studies
In the recent literature, there exists a series of astonishing theoretical results regarding NC in wireless networks. Lun et al.[13] show that the problem of minimizing the communication cost can be formulated as a linear program and solved in a distributed manner. Chaporkar and Proutiere [2] show that unless appropriate scheduling is applied, NC may lead to performance degradation. We support this claim by identifying an additional example when implicit ACKs are used. Liu and Xue [14] consider NC for two-way relaying in a three-node network. They analytically characterize the achievable rate regions for the traditional Alice-Relay-Bob topology, and they find the theoretically optimal end-to-end sum rates. Scheuermann et al. [6] propose noCoCo, a deterministic packet scheduling scheme for NC within two-way multi-hop traffic flows. Their scheme involves per-hop packet scheduling, NC, and congestion control. Seferoglu and Markopoulou [15] provide an understanding of the interplay between application data rate control and NC. Finally, Vieira et al. [16] provide observations on how the combination of NC and bit rate diversity affects the performance of practical broadcasting protocols. They show that it is possible for multi-rate link layer broadcasts and NC to jointly increase the network throughput in multicast applications.

In recent theoretical work [17], it is shown that for the two-user case, maximum throughput can be achieved without explicit ACKs, i.e., by utilizing scheduling algorithms that guess, code, and then correct the transmissions using feedback. In particular, this scheme has no loss of optimality if either the downlink channel rates are equal

for the two receivers or if the overhearing probability at the slow receiver is 1 (perfect overhearing). Moreover, whenever the overhearing probabilities are high enough (>0.6), the loss of throughput in comparison to the explicit ACKs is very small (<5%). This motivates further our work, which is in line with such schedulers. In particular, our framework enables the implementation of such algorithms for reasons of performance analysis and prototyping. We believe that the experimental evaluation is very important since the complexity of some of the newly proposed approaches may degrade the performance and result in smaller benefits compared to that of the theoretically predicted ones.

3 Network coding scheme
We study topologies that include two-hop flows crossing a central node, as shown in Figures 1 and 2. The central node, called the *relay* (or router), is connected to all other nodes, called *neighbors*. Links between neighbors may exist as well. Each link connecting nodes i and j is characterized by two channel rates (r_{ij} and r_{ji}) and two probabilities (q_{ij} and q_{ji}) which correspond to the packet delivery ratios in each direction.

In the uplink part (the first hop of these flows), *native* packets are transmitted without NC towards the relay. In the downlink part, the relay selects a number of native packets, applies the XOR operator, and transmits the constructed *encoded* packet. If a receiver recognizes its address in the header, it attempts to perform decoding in order to obtain the intended native packet. To achieve this, the receiver should be in possession of all the other native packets that comprise the encoded packet in order to perform decoding. These packets are known to the receiver either because they have been generated by it (in case of symmetric flows) or because they have been obtained by means of opportunistic listening, i.e., overhearing of the uplink transmissions [3]. Packets that are successfully obtained are acknowledged to the relay via a separate mechanism that operates in a higher layer (feedback packets are generated and sent to the relay), see Subsection 4.3.

In order to experiment on high gain topologies, we consider two particular classes of topologies which we name as $\frac{x}{2}$ wheel and x wheel; these names are inspired by [3]. Such topologies are ideal for performing wireless NC, since they allow a maximum number of employed two-hop flows to be mixed due to the existence of appropriate overhearing links. For example, in Figure 1 (right), the relay may encode six packets and send them in one transmission, while the decoding is guaranteed due to the existence of appropriate links (black arrows). A constructive way to build an $\frac{x}{2}$-wheel topology is the following: Split the neighbor nodes in two equal sets (x must be even), the source and the destination set, and select a

matching of these two sets which correspond to $\frac{x}{2}$ flows (pairs of nodes). Then, allow all nodes (the relay and the neighbors) to be connected except those that are paired in the above selected matching. Finally, create $\frac{x}{2}$ more flows by inverting the roles of source and destination. If only the initial $\frac{x}{2}$ flows are enabled, we refer to the topology as '$\frac{x}{2}$-wheel setting'; otherwise, if all x flows are enabled, we refer to it as 'x-wheel setting'. See Figures 1 and 2 for examples of such settings, wherein NC can be applied.

Evidently, the two-wheel setting corresponds to the well-known Alice-Relay-Bob topology. Throughout this paper, we also use $\frac{4}{2}$-wheel settings (referred to as 'half-cross') and four-wheel ones (referred to as 'cross'), as well as a six-wheel setting, wherein we activate the flows serially, one after the other (again see Figures 1 and 2 above). We should point out that NCRAWL supports any random subgraph of the N-wheel topology with any possible set of flows activated on top of it. In addition, it supports settings with any possible combination of link qualities and/or channel rates. The wheel topology is the most general topology to be considered around a single node. Any actual network topology can be reduced to a subgraph of a wheel topology if the nodes, links, and flows irrelevant to NC on any given node are removed. Since NCRAWL runs on all nodes in the network, our scheme works with any arbitrary network topology, providing the maximum possible local NC gain subject to some constraints we explicitly mention below. As we discuss later, the encoding opportunities at each node are automatically discovered by the combination of NCRAWL and Srcr [18]. Hence, NCRAWL operates under any assumed graph providing throughput gains opportunistically. We study the wheel setting, where the maximum such gains arises, in order to showcase that NCRAWL can achieve it in many cases despite the overheads and complexity induced by NC.

We have incorporated the following features in order to demonstrate the practicality of NCRAWL in a tractable manner:
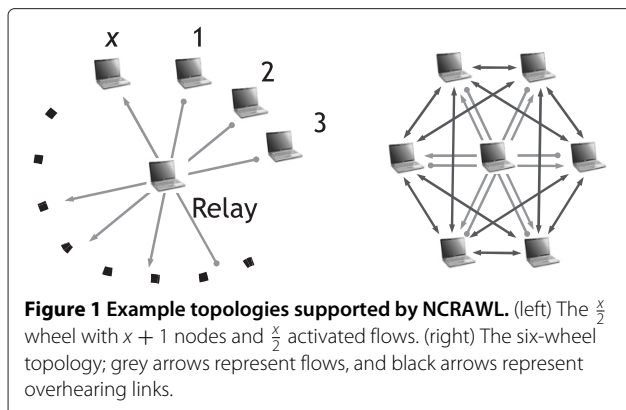

Figure 1 Example topologies supported by NCRAWL. (left) The $\frac{x}{2}$ wheel with $x+1$ nodes and $\frac{x}{2}$ activated flows. (right) The six-wheel topology; grey arrows represent flows, and black arrows represent overhearing links.
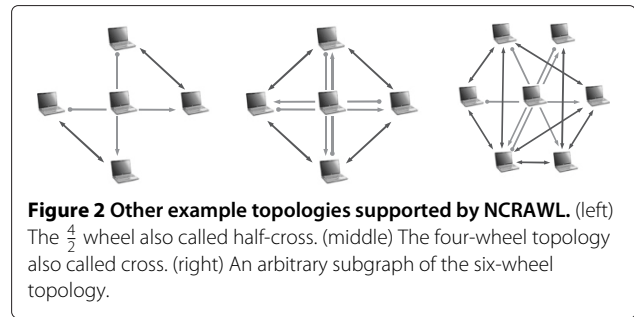

Figure 2 Other example topologies supported by NCRAWL. (left) The $\frac{4}{2}$ wheel also called half-cross. (middle) The four-wheel topology also called cross. (right) An arbitrary subgraph of the six-wheel topology.

- The XOR (binary field) operator is used instead of more general linear coding, similarly to [3].
- We enforce the decoding of encoded packets at the next hop, since this is practically the most possible case in today's wireless access deployments. Multihop flows can still benefit from our scheme by cascading the two-hop coding operations.
- Only native packets are stored as keys.
- We consider the use of *implicit ACKs* of overheard packets, as we explain below.
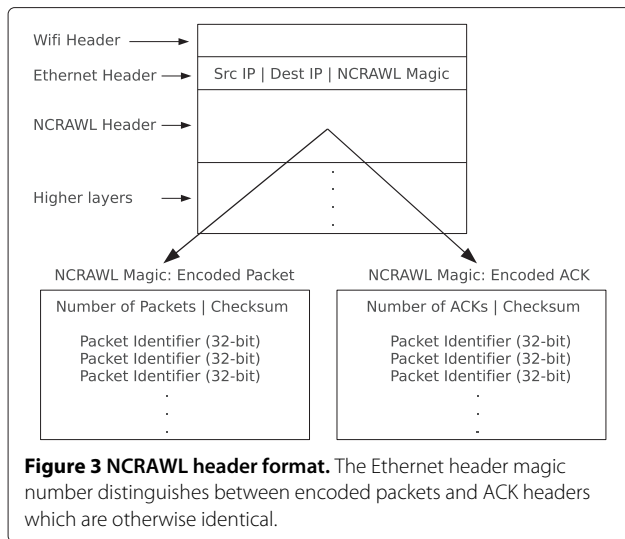- We use equal packet size (maximum MAC PDU); if a packet has smaller number of bits, we pad with zeros.

The decision on imposing these features is justified by the necessity to keep the NC scheme simple, practical, implementable, and efficient in terms of processing overhead. Clearly, the modular nature of NCRAWL allows the development of more complex features, taking into account potential application and policy/administration requirements.

4 Architectural blueprint

In this section, we describe the modular design and implementation of NCRAWL.

4.1 Employing Click as the basis of our framework

NCRAWL has been developed in the Click modular router framework [10]. Click can be used to develop primarily OSI layer 3 packet processors, which can be directly deployed inside the standard Linux network stack. A Click processor is mainly comprised of (a) processing stages which are called *elements* and (b) an element interconnection configuration that indicates the processing flow. Execution in Click is event-driven, with four different types of asynchronous events, namely the *incoming packet* event, the *ready-to-forward packet* event, the *timer expire* event, and the *external read or write* events to Click memory. Since all Click events are asynchronous, a Click packet processor leverages internal queues to temporarily store incoming packets, since they are to be forwarded asynchronously.

Figure 3 NCRAWL header format. The Ethernet header magic number distinguishes between encoded packets and ACK headers which are otherwise identical.

In what follows, we describe the NCRAWL system design and implementation. We also present the NCRAWL interface that can be used to develop new algorithms and deploy and manage experiments on wireless testbeds.

4.2 Design preliminaries

NCRAWL realizes an OSI layer-2.5 protocol that lies immediately under the routing layer. More specifically, it can be considered as an extension to the Click modular router implementation of the Srcr protocol [19], which is the heart of the MIT Roofnet wireless network. NCRAWL operates below the routing layer; this

simplifies the format of the NCRAWL packet headers, which are now used only to encapsulate encoded packets, and carry the corresponding ACKs for successful decoding back to the sender. The NCRAWL header formats are depicted in Figure 3. Network-wide unique 32-bit packet identifiers are used by applying the *sdbm* hashing algorithm [20] on data tuples that comprised packet source IP, the IP header sequence number, and the respective offset.

4.3 NCRAWL system description

Next, we discuss the modular design of NCRAWL in detail. Our platform is based on a Click network packet processor that includes the Srcr routing protocol implementation for wireless networks [19]. We have included two additional processing stages: the NCRAWL *decoder* and the NCRAWL *encoder*. We have developed these stages as individual Click elements, and we have placed them before the beginning and after the end of the Srcr processing flow, respectively, as depicted in Figure 4.

4.3.1 NCRAWL acknowledgments (implicit ACKS)

As opposed to Katti et al. [3], who uses explicit ACKs, where each successful overhearing of a packet is acknowledged to the router[a], NCRAWL uses ACK packets only for the verification of correct decoding. Following the transmission of a coded packet, the router expects layer-3 ACK messages from the intended receivers of the coded packet. This takes place for two reasons:

- The router obtains MAC-level ACK only from one receiver; thus without the layer-3 ACK, the router

Figure 4 NCRAWL framework system components and their interaction.
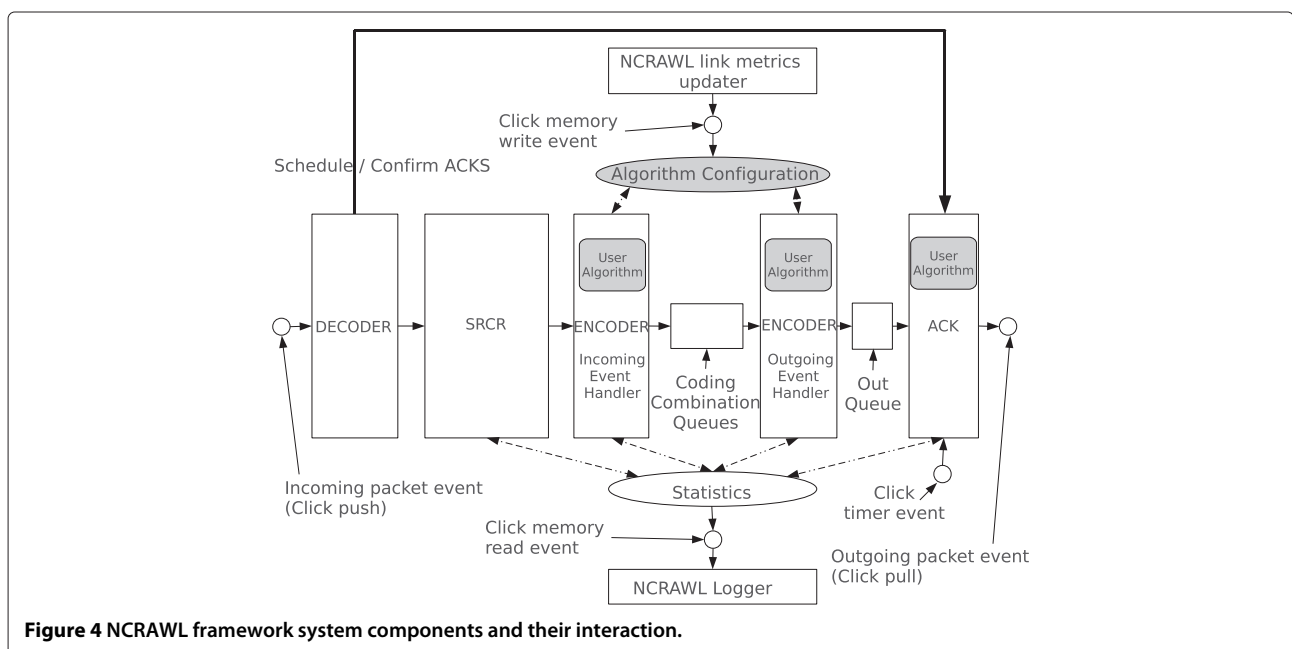
cannot be notified of a loss that has occurred at the other intended receivers.

- Since the overhearing events are unknown to the router, it is impossible to know if the decoding was successful (i.e., if the receivers have the necessary keys in their buffers).

Note that a negative ACK (NACK) architecture is possible if the broadcast 802.11 mode is used. If NACKs are used, the overhead of ACK messages can be reduced significantly in practical scenarios (95% reduction is possible). Even with the ACK architecture though, ACK messaging is reduced due to the lack of unnecessary overhearing ACKs.

The router holds a buffer for the packets to be acknowledged. When an ACK arrives, the buffer is inspected, and the acknowledged packet is removed. The buffer has a periodic timer, and each packet is stored in the buffer for at least a specific period. When the timer expires, packets found to have been stored for more than that period are analyzed, and the corresponding non-acknowledged native packets are sent back to the system queues with the information of 'decoding failure'. Note that a decoding failure can be attributed either to the MAC failure or to the inability of the intended receiver to decode (lack of proper keys). Note also that a timer-expire event triggers the transmission of ACKs in separate packets when there is not enough outgoing traffic to piggyback them (Figure 4).

4.3.2   The packet decoder
The main tasks of the decoder module are the following:

- To use the available (from overhearing or ownership) *key* packets in order to decode the received encoded packet.
- To schedule the transmission of layer-3 ACKs for the correctly retrieved native packets derived by the decoding operation.
- To determine any potential pending ACKs, as well as to verify any received ACKs.
- To tag and store all the correctly overheard data packets as potential keys; as discussed above, these will be potentially used in the near future to decode the received encoded packets. Moreover, the key repository is used for packet retransmissions, in case an expected ACK never arrives.

The decoder resides at the packet receiving side of the system and is invoked by the corresponding packet arrival event.

4.3.3   The NC packet encoder
The NCRAWL encoder element resides at the sending side of the system and is more complicated, since

it maintains and manages the processor packet queues. A part of the element handles incoming packet events; another part deals with outgoing packet events, and there is also a code that gets invoked upon timer expiry as well as read and write Click configuration events (Figure 4). It is this element that exports the framework API which can be used to develop NC algorithms. Specifically, the main assigned tasks for this module are the following:

- To process and place incoming native packets into specifically maintained queues. Our system supports a plurality of queueing operations, which can be configured as per the requirements of the NC algorithm under development.
- To identify and combine packets together towards forming encoded packets. The selection of the appropriate packet set follows the directions of the NC algorithm under consideration, supported by NCRAWL.
- To piggyback any ACKs (through the use of scheduled, upcoming data packet transmissions) that have been scheduled by the decoder element.
- To generate potentially expected ACK tokens for each of the native packets comprising an encoded packet.

4.3.4   Maintaining up-to-date topological information
The *link metrics updater* is responsible for collecting information about the transmission rates and PDR values of the links with neighbors. This information is gathered and passed to the rest of the system via the Click memory write event mechanism. Furthermore, the software procedures that set and apply the encoding combination policies are invoked as needed.

Gathering link quality information The NCRAWL updater relies on the Srcr routing protocol component [18], which maintains link connectivity information and performs periodic link measurements. Srcr sends probe packets at all rates to determine the PDR for each link and chooses the highest rate that performs well. PDR information is then used by Srcr to calculate the ETX or ETT metric [21,22], which provides information about the entire routing paths (not just one-hop links). This information is kept in the Srcr link table and is accessible by the NCRAWL components. The Srcr measurement period can be set as desired (the default value [18], also used in our work, is 3 s).

Collecting information about neighbor node links Based on the information stored in the Srcr link table, the link updater maintains its own so called neighbor table (NT), which includes information for its neighbors. Initially, the NT is empty. The updater periodically reads the Srcr link table and updates NT as needed. The

NT contents are updated whenever (1) a new neighbor appears, (2) an existing neighbor disappears, or (3) a certain link quality changes. In such cases, the NCRAWL updater broadcasts a packet with the new NT contents and sets a timer. When such a NT packet is received (overheard) by other nodes, their corresponding updater utility replies by broadcasting its NT, provided it has not done so recently. Specifically, in order to avoid excessive network overheads due to such messaging, we leverage a NT *reply suppression threshold*; while the value of this threshold is tunable (and can be configured based on network and application dynamics), we set it to be equal to the Srcr broadcasting interval period. The NT packets are used by the NCRAWL updater to maintain the so called received NT table (RNTT), which complements NT by storing information about the link quality, as measured by neighbors.

Feeding NC algorithms with updated topological information Each time the updater modifies the contents of the NT or RNTT (i.e., each time it proactively sends or receives a NT packet which leads to an update of the RNTT), a timer is set. Upon timer expiry, the new link qualities are passed to the main NCRAWL system, where they will potentially drive adaptive NC decisions, based on the NC designer's needs. This timeout is (generously) set to 1 s, providing ample time for any NT reply packets to arrive. The timeout value can be set based on the dynamics of the network on which NCRAWL is applied. Note here that overheads are kept low with NCRAWL, since the updater employs its own threads of execution to perform these information maintenance tasks; the current implementation uses two threads, but these remain suspended most of the time, making this component quite lightweight in terms of CPU occupancy. Note that the implementation of the NCRAWL updater can be trivially adjusted to cooperate with other link information gathering protocols as well, i.e., it is not tied to Srcr.

4.3.5 NCRAWL logger
The *read* events are used by another utility, the NCRAWL *logger*, which gathers various statistics that are generated online by both the encoder and decoder.

Utilizing resources effectively Efficient resource utilization is an inherent property of NCRAWL. The repository that stores copies of packets uses a FIFO queue as the main indexing mechanism and can host up to a user-defined quantity. After the storage limit is reached, the oldest packet is removed in order for a new one to get stored. The same packets are also indexed in a hash table based on their network-wide unique identifiers. The hash table is used to quickly retrieve packets used either as keys for decoding or for re-sending them in case an expected ACK token expires. The same indexing approach has been used for the ACKs and expected ACK tokens as well.

4.4 Implementing NC algorithms
NCRAWL exports an application programming interface (API) that can be used to implement scheduling algorithms for intersession NC. This API is a library of functions that can be used to carry out NCRAWL common tasks and mandatory function extensions to the handling code of each event. Points of extensibility and/or programmability are denoted in Figure 4 with shadowed boxes.

4.4.1 Implementing packet handling operations
Regarding the incoming packet event, the designer should account for placing arriving packets into proper queues. In particular, each flow is associated with a queue, and the scheduler checks all available controls, i.e., activates a set of queues by combining one packet from each queue. It maintains a list with the expected score (reward) of all controls and selects one of those controls at each time instance. The controls that activate only one queue correspond to the case of transmitting native packets. It is always possible to deactivate NC by imposing the use of only those latter controls. With this, the developer may implement logic that disables NCRAWL when needed. Furthermore, after the placement of the incoming packet, the developer may (a) invoke the function that chooses the next encoded packet queue combination according to the NC algorithm, (b) retrieve the packets from the respective queues, (c) encode them and schedule the encoded packet for transmission by placing it into the outgoing queue. This operation may be repeated until the outgoing queue contains a user-defined number of packets. Apart from the queue combination retrieval function, which needs to be implemented by the algorithm developer, the rest of the required functionality is already seamlessly provided by NCRAWL.

4.4.2 Implementing the core NC logic
The main algorithm implementation takes place in the context of the Click memory write event, generated by the NCRAWL updater. The latter provides the user with a table of links with entries denoted by the corresponding source and destination IP pairs. Each entry holds the link direction, the PDR, and the transmission rate. This information can be used by the NC algorithm to decide upon valid NC combinations by selecting the queues to activate together. Since this part of the code runs periodically, the developer is encouraged to implement any complex algorithm steps here and thoroughly index the NC available combinations. With such an approach, the overhead of choosing the most beneficial packet combination during the incoming or outgoing packet events is minimized.

4.4.3 Sending data and ACK packets

The outgoing packet event checks the size of the outgoing packet queue. If this is below the defined (configurable) threshold, the functions that choose and encode the combinations are called. The developer may also add logic for the handling of ACKs. By default, NCRAWL resends packets that have not been acknowledged by directly placing them on the outgoing queue. Note that NCRAWL allows the development of algorithms that deal with the failed packets in a different way (see Section 5 for an example). Since NCRAWL groups ACKs that belong to the same encoded packet, the scheduling algorithm knows which packets have been decoded successfully at which destination and, hence, may further decide whether a packet should be resent directly or potentially reconsidered in encoding combinations.

4.4.4 System monitoring

Our framework allows for user-defined timer events. Statistics for incoming/outgoing packet activity as well as queue lengths are all logged using counters. The NCRAWL logger periodically retrieves statistics and notifies the user at runtime about the flow stability and the corresponding queue lengths. The latter are also available for use in the NC algorithm if needed. Finally, the developer may also implement additional debugging support for inspecting the algorithm configuration at runtime using Click's read handler.

4.5 Deploying NCRAWL experiments

We have integrated NCRAWL deployment scripts with the OMF framework [23] for wireless testbeds as an optional feature. OMF is a control, management, and measurement framework that provides users with tools to easily describe, execute, and collect testbed results. NCRAWL leverages the following OMF utilities (see [23] for details):

- Gridservices. This is a set of web services that are used by OMF to fetch information and perform actions remotely on the nodes.
- Nodehandler. This component resides on the central server that interacts with the user for the experiment submission. Moreover, it provides the necessary applications for node system image loading, experiment execution, image saving, and node status check.
- Nodeagent. This component obtains instructions for the experiment deployment, arriving from the nodehandler.

NCRAWL extensions have been written for both nodehandler and nodeagent. The former performs transfers of the Click executable along with user defined parameters. The latter retrieves local node information (e.g., the network interface name and MAC address) and then parametrizes a generic NCRAWL deployment script. Finally, the nodehandler is notified if the deployment was successful. With OMF, the NCRAWL experiments can be deployed with minimal user effort.

5 Case study

This section demonstrates by example how the NCRAWL framework allows for easy implementation of scheduling algorithms. We also discuss a case study which we use for the assessment of NCRAWL in the next section.

At each WiFi transmission opportunity, the scheduler selects a number of packets, encodes them, and sends the encoded packet to the MAC layer. The problem is of scheduling nature: *which are the best packets to select considering the queue backlogs, the transmission rates, and the overhearing probabilities?* Note that selecting only one packet corresponds to transmission without NC.

A well-known family of optimal algorithms for scheduling is the maximum weight (maxweight) algorithms applied in the stability theory of stochastic networks and input queues switches (see for example [24,25]). In these algorithms, *control actions* are chosen to maximize a reward that depends on the product of link rates and queue lengths. The application of such algorithms towards solving the joint NC and scheduling problem with arbitrary rates is then promising [26].

For the case of the implicit ACKing scheme used by our framework, one observes that the actual rate of service is random. The randomness comes from the decoding; the scheduler (at the relay node) does not know *a priori* whether packets needed for decoding are missing or not by some of the receivers. Nevertheless, a NCRAWL equipped relay knows the probabilities of overhearing and is then in position to determine the probability of decoding and, thus, the *expected service rate*. Let C represent a selected subset of receivers (neighbors of the relay). If the scheduling decision is C, then the reflected action is to combine the first packet from each of the queues with packets destined to nodes C and send to the MAC layer the coded packet. The expected service rate for receiver $i \in C$ with such an action will be

$$\mu_i(C) = \min_{j \in C}\{r_j\} \prod_{\substack{j \in C \\ j \neq i}} q_{ji}, \qquad (1)$$

where $\min_{j \in C}\{r_j\}$ determines the actual transmission rate with this action; with this action, q_{ji} is the probability that a transmission of a packet destined to receiver j is overheard at receiver i, and $\prod_{\substack{j \in C \\ j \neq i}} q_{ji}$ determines the probability of decoding at receiver i. Here, we have assumed that the overhearing events are independent. This is a realistic

assumption since the time and space for each overhearing event are different. Collisions and Rayleigh fading may be the causes for this randomness. Below, we discuss some example algorithms that can be implemented and assessed with NCRAWL.

5.1 Algorithm 1

Let Q_i be the queue backlog at the relay for flow i, and $\mathcal{C}$ the set of all possible control actions. Algorithm 1 searches set $\mathcal{C}$ for the action that maximizes the *total* expected service rate.

Algorithm 1 MaxWeight Algorithm without feedback

Input: $Q_i, \mu_i(C)$

Output: C^*

$C^* := \arg\max_{C \in \mathcal{C}}\{\sum_{i \in C} Q_i \mu_i(C)\};$

Algorithm 1 is throughput-optimal under the condition that the random decoding event distribution does not change during transmissions, and thus, the expected service rate computed above is accurate. This happens when (1) the probabilities are 0 or 1, as in Alice-Relay-Bob topology (and any other symmetric flow setting), or when (2) upon a decoding failure, we reschedule the uplink transmission for the failed flow. The latter may arise, e.g., in a TCP scenario. In the general case, however, whenever a particular encoded packet is not correctly decoded, extra feedback information is obtained and the expected service rate changes. If, for example, $P_1 \oplus P_2$ is not decoded by both receivers, the relay knows that these two packets are not overheard by receivers 2 and 1, respectively, and they should not be coded together again. The impact of feedback clearly biases the probabilities of decoding. Therefore, in the general case, Algorithm 1 is not optimal; also, when the overhearing probabilities are quite small, it might perform quite badly.

5.2 Algorithm 2: the case of two flows

Utilizing feedback information for making scheduling decisions is a notoriously difficult problem. In [27,28], an *optimal* scheduling algorithm is discussed, which however requires a significant number of calculations, since it operates on a virtual network with $O(4^N)$ virtual queues, where N is the number of physical neighbors. Here, we discuss a simple myopic scheduling algorithm that is not optimal but is able to effectively utilize feedback and outperform Algorithm 1. For tractability, we discuss the case for two flows; note that it can be extended to the case of arbitrary number of flows and pairwise coding.

In order to cope with feedback, we add two more knowledge states. Apart from *unknown* (newly arrived) packets,

behavior of which is captured by known probabilities, we have a state for *good* packets (successfully overheard) and one for *bad* packets (those not overheard), where the scheduler has deterministic knowledge of the decoding event. Given this, the system maintains the queues Q_i^s, where $i \in \{1, 2\}$ signifies the flow and $s \in \{u, g, b\}$ signifies the state. The new set of controls $\mathcal{C}_2$ contains all controls that activate one or two queues with the constraint that no two queues from the same flow can be activated.

Algorithm 2 Myopic Algorithm with feedback

Input: $Q_i^s, \mu_i^s(C)$

Output: C^*

At feedback time

- For each packet that was not correctly decoded, define whether it is good or bad using the decoding event of the paired packet.
- Bad packets are directly sent to the MAC layer for transmission without coding.
- Good packets are sent to the corresponding queue at the good state.

At decision time

- $C^* := \arg\max_{C \in \mathcal{C}_2}\{\sum_{i \in C} Q_i^s \mu_i^s(C)\};$

The packets are initially injected in the queues corresponding to the unknown state. Once a packet is not decoded properly, the relay classifies it as either good or bad based on feedback information. If it is bad, it is retransmitted without encoding. If it is deemed as good, it is transferred to the corresponding queue at the good state (Q_1^g or Q_2^g depending on the flow it belongs to). When calculating the average service rates, the packets at the good state have the probability of overhearing equal to one. Apart from these alterations, Algorithm 5.2 works in the same way as Algorithm 1.

5.3 Algorithm 3: fixed threshold policy

This algorithm operates only with implicit ACKs and makes decisions based on principles used in the COPE framework. In this sense, it emulates COPE in its probabilistic mode. An important difference is that instead of calculating average service rates, Algorithm 3 (namely δ-*fixed threshold policy* or δ-FTP) simply compares the overhearing probabilities $q_{i,j}$ with a fixed threshold $\delta \in [0, 1]$ and sets them to 1 if they exceed the threshold or zero otherwise. The algorithm selects at each decision instance the control that maximizes the number of transmitted packets.

5.4 Implementing the three algorithms in NCRAWL

Next, we demonstrate how to implement the three above algorithms on NCRAWL. For all three cases, we configure NCRAWL at each node to maintain one queue per flow for incoming packets.

5.4.1 Implementing Algorithm 1

We first describe how one may organize queues in an efficient manner. Subsequently, we show how to utilize the queue information to apply NC.

- Organizing packet queues. To begin with, we dedicate one vector per control C which contains the identity of the involved queues (e.g., the flow it belongs to and/or the state) and the identities of the packets enqueued at the involved queues. The formed vectors are stored in a double linked list. Each vector is assigned a weight (see the argument of the maximization in the above algorithms); the higher this weight, the higher the preference of the encoder to use the combination. This weight is recalculated every time the backlog size of a member queue changes. The linked list is formed such that the head of the list contains always the current maximum weight. For the sake of low-processing overhead, vectors are also directly indexed by their member queues; with this, the weight update process is fast. Vectors as well as their linked list are all constructed during the NCRAWL updater write event.

- Applying NC operations. Given the construction of the control list, the encoder event examines the head of the list and further: (a) retrieves packets from their respective queues, (b) updates the vector weights (since the respective backlogs are decremented), and (c) sets the vector with the highest weighted combination as the head of the list. The latter is actually a process with slowly scaling complexity with the number of vector combinations, since each updated vector weight is just compared against the weight of the current head and only takes its place if it is higher. Retrieved packets are subsequently combined using the NCRAWL *encode* library call, and the resulting encoded packet is scheduled for transmission.

5.4.2 Implementation considerations for Algorithm 2

This algorithm is similar to Algorithm 1; however, it involves an additional acknowledgment scheme logic. Therefore, for each flow, NCRAWL now maintains two queues: (a) one with new incoming packets and (b) one with packets that have been successfully logged as keys by fellow nodes but have not reached their ultimate destinations [b]. Algorithm 5.2 exploits the NCRAWL acknowledgment scheme facility; this process groups the packet acknowledgment tokens, which have been created for outgoing packets combined together in the same encoded packet. This information is provided by NCRAWL to the developer. Algorithm 5.2 directly sends packets that have not yet reached their destinations; such packets are not reconsidered for encoding. However, the algorithm considers favorable queues and 'unknown' queues for the same flow separately, when forming vectors. Note that the vectors formed with this algorithm scale intrusively compared to the simple maxweight algorithm described previously. Throughout our measurements, we only consider the scenario of two flows and, thus, avoid the arising complexity. This issue is expected to be resolved in the future using the NCRAWL framework.

5.4.3 Algorithm 3 in NCRAWL

For the implementation of the third algorithm, we simply need to create vectors (i.e., controls or queue combinations) for which the decoding probability is nonzero, according to the user-defined threshold δ and the channel quality. As soon as the packets are available in all queues that constitute a vector, they are combined and transmitted at once, without considering or updating the queue backlogs. This algorithm selects controls that mix the largest possible number of packets each time.

Note that as opposed to COPE [3], NCRAWL does not use any time-threshold policy towards increasing the backlog size of the incoming packet queues before deciding to send outgoing packets. With NCRAWL, queue backlogs will increase when the relay's outgoing packet rate is smaller than that of the incoming packet rate. In such cases, NC proves to be a panacea for the router stability; if the NC algorithmic operations are supported by a lightweight implementation, the router capacity can be truly increased, as our measurements suggest.

6 Evaluating our framework

In this section, we evaluate NCRAWL in conjunction with scheduling algorithms (NCRAWL + alg1, NCRAWL + alg2, and δ-FTP) described in Section 5, in terms of both throughput and resource utilization. We begin by describing the wireless testbed infrastructure and the configurations that we used to deploy experiments. Next, we quantify the CPU overheard that is introduced by each NCRAWL processing stage, under maximum traffic loads, and we compare total CPU utilization to (1) the public COPE implementation that uses an explicit acknowledgment scheme and (2) legacy IEEE 802.11b-g. Following, we demonstrate that NCRAWL can support theoretical gains even when coding opportunities lead to more than two-packet combinations. Finally, we deploy experiments that demonstrate how the proposed algorithms perform in cases with variable link qualities and different rates.
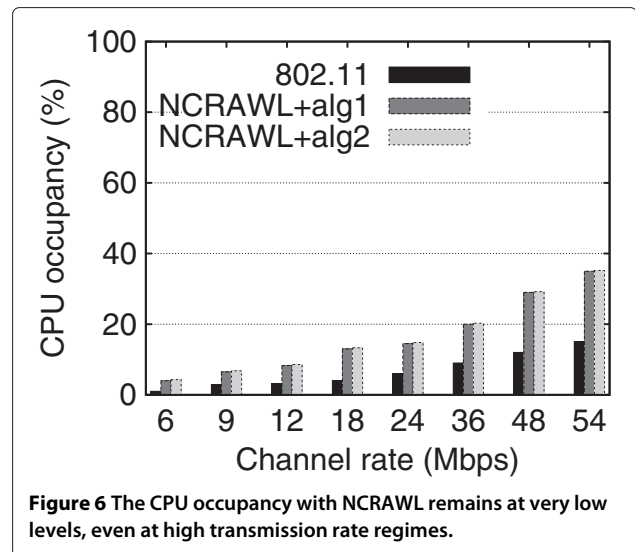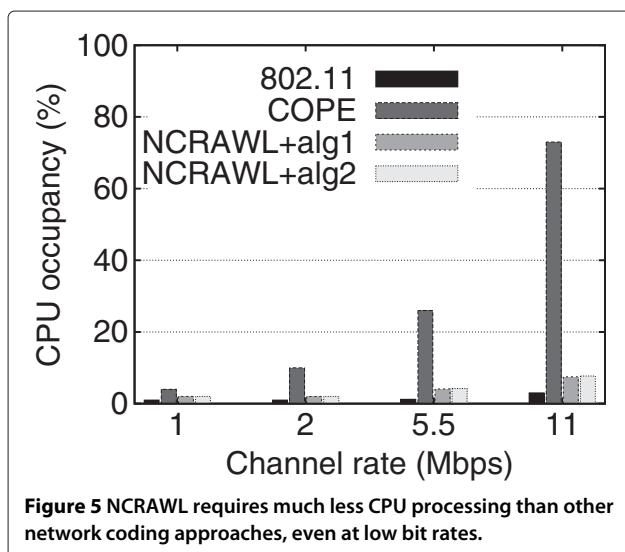
6.1 Experimental setup

Our testbed comprised 20 ORBIT nodes, deployed both indoors and outdoors, in the Department of Computer Engineering and Telecommunications of UTH. Each node consists of one 1 GHz i386 processor, 512 MB of RAM, two Ethernet ports, and two miniPCI slots which are used to host two AR5212 Atheros 802.11a/b/g WiFi cards. UDP and TCP traffic is generated using the *iperf* software tool [29]. For CPU occupancy measurements, we appropriately instrument NCRAWL with the Linux *getrusage* system call, which accurately estimates CPU usage time. We place several getrusage calls at the borders of each processing stage, we record the average usage time of each stage, and compare it to the whole NCRAWL system usage time. We have repeatedly performed all of our experiments late at night in order to avoid interference from collocated networks.

6.2 CPU occupancy measurements

In order to measure the efficiency of our framework in terms of CPU occupancy, we compare it to the case of running COPE [3], as well as the legacy IEEE 802.11 protocol.

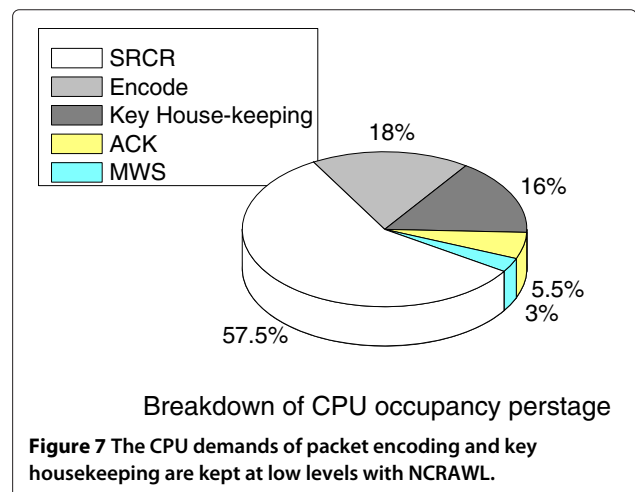#### 6.2.1 NCRAWL is more CPU friendly than previous approaches

We invoke the Alice-Relay-Bob setting (see Section 3), and we inject fully saturated traffic in both flows. We compare NCRAWL + alg1, NCRAWL + alg2, COPE, and the plain 802.11 for the case of 802.11b; Figure 5 depicts the results. Note that COPE can support at most the IEEE 802.11b rate set as discussed in Section 1. For the sake of a fair comparison here, we use this mode of operation for NCRAWL as well. We observe that NCRAWL makes use of the CPU resources in a very efficient manner: it
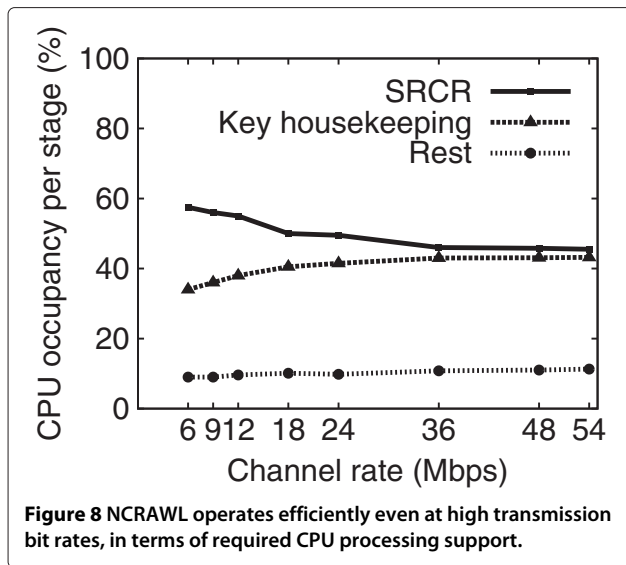

Figure 6 The CPU occupancy with NCRAWL remains at very low levels, even at high transmission rate regimes.

reduces the CPU utilization by at least two and as much as seven times compared to COPE (we have validated these observations for the case of ER [7]as well, which is based on COPE). Furthermore, we test NCRAWL for the case of 802.11g. Our measurements(Figure 6) suggest that NCRAWL does not need to occupy more than 37% of the CPU resources for NC operations at 54 Mbps, with fully saturated UDP traffic. This implies that the design of NCRAWL includes low additional overhead functions (as opposed to legacy 802.11).

6.2.2 Evaluating individual operations of NCRAWL

Next, we deploy getrusage calls and measure the breakdown of CPU occupancy per processing stage (Figure 7). The most CPU-intensive operation is the Srcr stage (it contains legacy IEEE 802.11 operations as well). The most computationally heavy pieces of NCRAWL are the encode

Figure 5 NCRAWL requires much less CPU processing than other network coding approaches, even at low bit rates.

Breakdown of CPU occupancy perstage

Figure 7 The CPU demands of packet encoding and key housekeeping are kept at low levels with NCRAWL.

Figure 8 NCRAWL operates efficiently even at high transmission bit rates, in terms of required CPU processing support.

Figure 10 The NCRAWL gain is equal to theoretically predicted one (case with 802.11b).

stage and the key housekeeping [c]. Note here that these two lie at the heart of any NC system and in a way represent unavoidable costs. It should also be noted that the processing stage of the scheduler remains at very low values, and there is a certain percentage dedicated to dealing with ACKs. Furthermore, as depicted in Figure 8, by increasing the channel rate (and thus the number of packets into the system per unit time), the coding stage increases in complexity disproportionally with the Srcr stage. This implies that the coding complexity increases faster than Srcr as the rate increases. Nevertheless, for high channel rates, the differences are reduced. This suggests that NCRAWL could potentially operate efficiently at much higher channel rates, such as with 802.11n systems. We plan to test NCRAWL on MIMO networks in our future work.
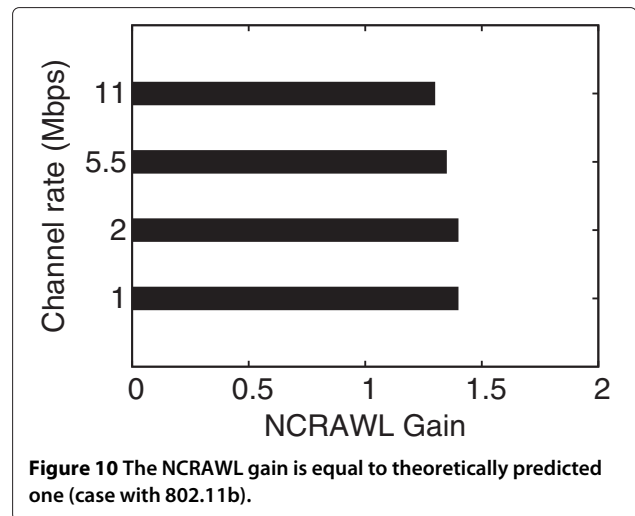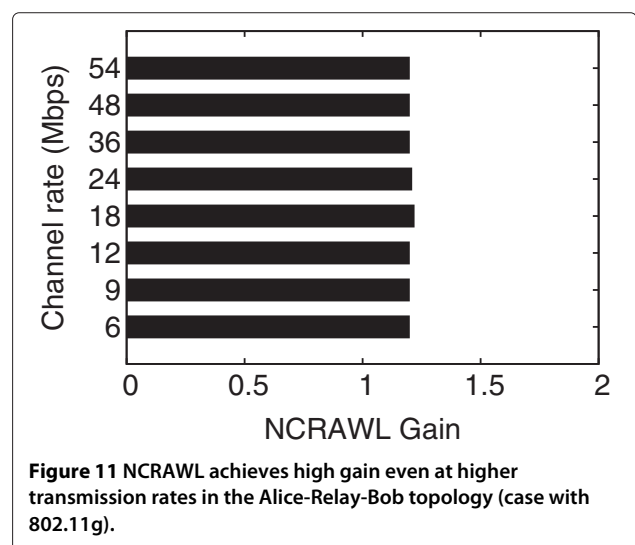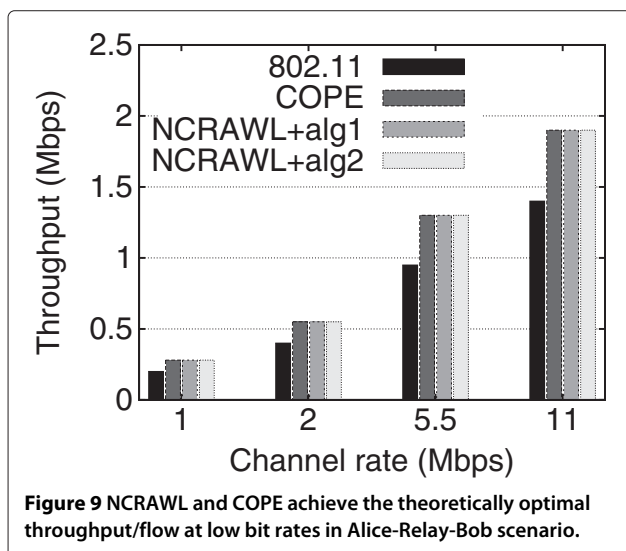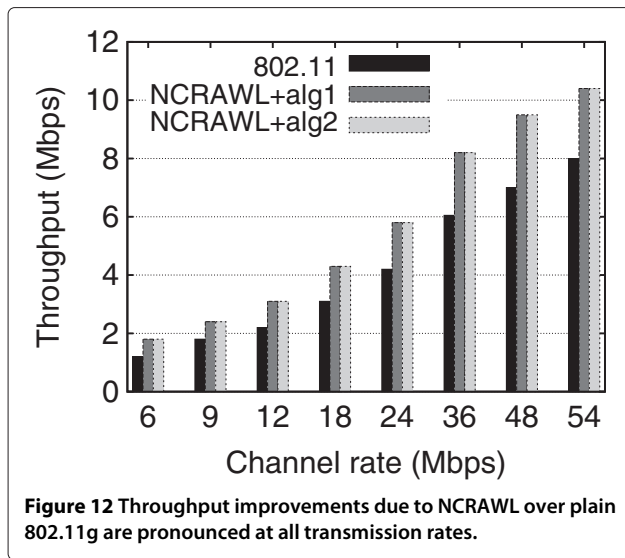
6.3 Throughput measurements with UDP

Next, we assess the ability of NCRAWL to approach the theoretically expected benefits of NC.
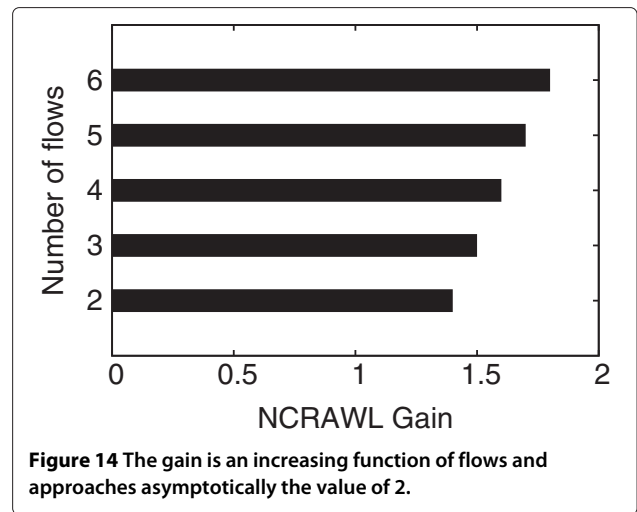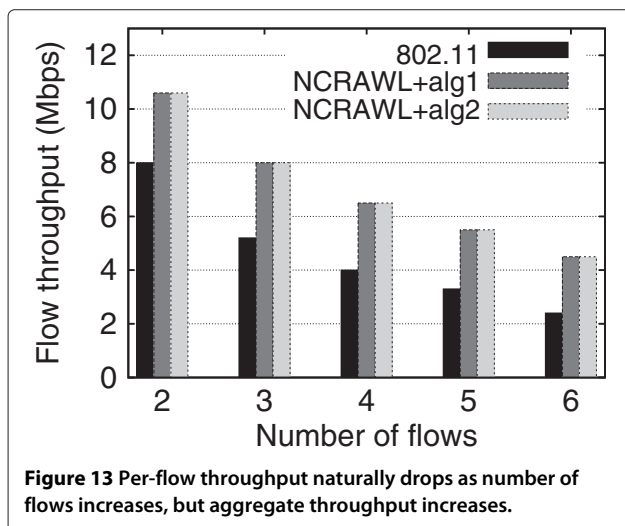
6.3.1 Experiments with the simple Alice-Relay-Bob topology

We calculate and measure the maximum throughput for both symmetric flows, such that the system remains stable (i.e.. the queues do not rise more than a large permissible number). Figures 9, 10, 11, and 12 show the results. Note that since the receivers always have the proper keys (these are the keys from their own transmitted packets [2]), decoding is always possible and thus Algorithm 1, Algorithm 5.2, and COPE are optimal in this setting. In each case, a gain in throughput of $\frac{4}{3}$ is identified, which matches the theoretical gain for this topology. Our measurements suggest that COPE achieves the theoretical throughput for small rates, but it fails to do so in higher rates. Note that

Figure 9 NCRAWL and COPE achieve the theoretically optimal throughput/flow at low bit rates in Alice-Relay-Bob scenario.

Figure 11 NCRAWL achieves high gain even at higher transmission rates in the Alice-Relay-Bob topology (case with 802.11g).

Figure 12 Throughput improvements due to NCRAWL over plain 802.11g are pronounced at all transmission rates.

Figure 14 The gain is an increasing function of flows and approaches asymptotically the value of 2.

the public COPE code was initially available for 802.11b only; while we carefully modified COPE to operate at 802.11g rates, we observed that such modifications lead to a very unstable system when rates higher than 18 Mbps are used. A closer look at certain individual components of the COPE implementation revealed that the reason for this instability is the excessive overhead induced by the NC system operations (as discussed earlier). For this reason, we do not explicitly compare COPE here at these high rates. Nevertheless, from these measurements, one can realize that COPE cannot provide benefits at rates higher than 18 Mbps due to the tremendous CPU processing overheads that its design incurs. In contrast, NCRAWL manages to reach the theoretical gain at high channel rates (e.g., at 54 Mbps), as shown in Figures 10 and 12.

Figure 13 Per-flow throughput naturally drops as number of flows increases, but aggregate throughput increases.
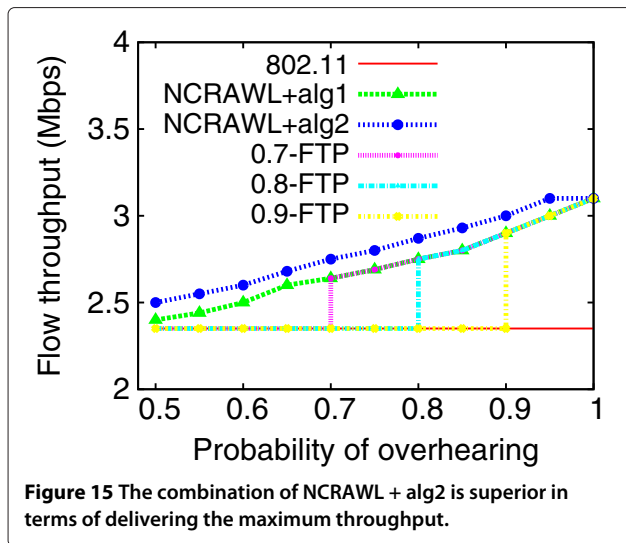
6.3.2 The case for wheel topologies

Furthermore, we scale the number of flows (see Figures 13 and 14); the topology is an $\frac{x}{2}$ wheel. The theoretical gain in this case is $\frac{2x}{x+1}$, where x is the number of flows combined at the downlink. Our measurements support the theoretically predicted gain at the channel rate of 54 Mbps. We observe that the per-flow throughput naturally drops as the number of flows increases, but the aggregate throughput increases. The gain (Figure 14) is an increasing function of flows and approaches asymptotically 2; note that this is perfectly aligned with the findings in [2] as well. Note also that in $\frac{x}{2}$-wheel topologies, piggybacking is not available since there is no return flow from the receivers. NCRAWL is able to select the appropriate ACKing method, and the results show that the overhead incurred is negligible.
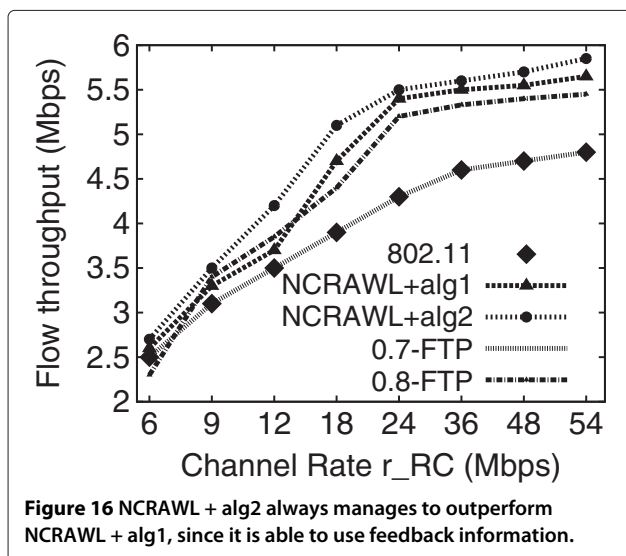
6.3.3 Experiments with cross topologies

We now present two more cases of interest that can appear in realistic environments. We setup various *cross* topologies with nodes in different locations across our testbed; we activate the flows Alice-Relay-Chloe and Bob-Relay-David. The arrivals are again chosen in a symmetric way, i.e., the arrival rate of the one flow is equal to the other.

- In the first case (Figure 15), David overhears Alice's uplink transmissions with probability 1, and Chloe hears Bob with probability *q*. The rates of all links are equally set to 12 Mbps (the channel rate is not important in this experiment). We measure the highest throughput that guarantees queue stability while varying the probability *q* by considering different node locations. We compare NCRAWL + alg1, NCRAWL + alg2, and IEEE 802.11g as well as δ-FTP for $\delta = \{0.7, 0.8, 0.9\}$

Figure 15 The combination of NCRAWL + alg2 is superior in terms of delivering the maximum throughput.

(see Section 5 for description). The results demonstrate the superiority of NCRAWL + alg2, which is able to deliver the maximum throughput in each case. Evidently, our framework in combination with the proposed scheduling algorithms is able to effectively handle the several link quality conditions.

- In the second case (Figure 16), the overhearing probability from Bob to Chloe is set to $q = 0.7$. All channel rates are set to 24 Mbps with the exception of the link Relay-Chloe which is varied. Our measurements demonstrate the inefficiency of policies oblivious to rates like the δ-FTP. In this case, the choice of a small value for δ is penalized when the Relay-Chloe link is slow enough. Instead, NCRAWL + alg2 is able to handle in an effective

Figure 16 NCRAWL + alg2 always manages to outperform NCRAWL + alg1, since it is able to use feedback information.

way the several rate and link conditions and deliver important throughput gains. From Figures 15 and 16, we also observe that given that overhearing links are not perfect in terms of PDR, NCRAWL + alg2 always outperforms NCRAWL + alg1, since it is able to use feedback information.

6.4 Performance with TCP traffic

Finally, we assess the efficacy of NCRAWL in scenarios with TCP traffic. In [3], experiments with TCP have demonstrated a loss in efficiency due to packet losses and reordering. First, throughout our experiments with the Alice-Relay-Bob topology, where no losses or delays are incurred, the throughput is reduced due to the additional TCP overheads. We observe that when the 54-Mbps rate is used, the per-flow throughput rate is 7 Mbps for plain 802.11 and 8.5 Mbps for NCRAWL + alg1. A slight loss in NC gain is observed; this is the result of mixing TCP ACKs with data packets. The same gain is obtained for all the other available bit rates.

Furthermore, we perform experiments with *half-cross* topologies, where flows are unidirectional (from Alice to Chloe and from Bob to Dave), with probabilities of overhearing $q_{AD} = q_{BC} = 0.7$ and several channel rates. In this case, NCRAWL + alg1 achieves a slightly lower throughput than IEEE 802.11. This is due to the fact that some packets are not correctly decoded at the destination; therefore; they arrive delayed and out of order. This causes abrupt reactions from TCP and leads to throughput reduction. When adding the reordering module of COPE [3], the packets arrive always in order; however, this module increases the delay for each packet. This in turn is interpreted by TCP as congestion; it ends up in TCP window increments and, thereby, decreases performance. NCRAWL is not optimized to cooperate with TCP at this point; thus, it faces the common problems of TCP in wireless networks. Improving this component is the main goal of our future work.

7 Conclusions

We design and develop NCRAWL. Our framework is an extended, generic NC framework that can be used to quickly develop networking systems in order to evaluate intersession NC and/or scheduling algorithms, entirely based on the implicit (probabilistic) acknowledgement that a packet can get decoded at the destination. The design of NCRAWL involves all the common processing steps that are always needed to implement such algorithms; these steps have been abstracted such that designers need to simply focus only on the implementation of their algorithms. Our measurements demonstrate that NCRAWL is a powerful NC development system. It offers significant throughput benefits even at high channel rates.

Endnotes

[a] Recall that the role of the router in an *ad hoc* setting might be played by virtually all nodes in the network

[b] For example, this could be due to the fact that the destination failed to decode properly a previously sent encoded packet.

[c] Key housekeeping refers to operations over a hash structure maintained for packet identification.

Author details
[1]Center for Research and Technology Hellas, Volos, Greece. [2]LIDS, Massachusetts Institute of Technology, Cambridge, MA, USA. [3]Aristotle University of Thessaloniki, Thessaloniki, Greece. [4]University of Thessaly, Volos, Greece.

References

1. R Ahlswede, N Cai, S Y, R Li, RW Yeung, Network information flow. IEEE Trans. Inform. Theory **46**(4), 1204–1216 (2000)
2. P Chaporkar, A Proutiere. Adaptive network coding and scheduling for maximizing throughput in wireless networks, in *Paper presented at the 13th ACM MobiCom* (Montreal, Canada, 9–14 September 2007
3. S Katti, H Rahul, W Hu, D Katabi, M Medard, J Crowcroft, XORs in the air: practical wireless network coding. IEEE/ACM Trans. Netw. **16**(3), 497–510 (2008)
4. JW Kim, YU Lho, YJ Kim, KB Kim, SW Lee. A memory copy reduction scheme for networked multimedia service in Linux kernel, in *Paper presented at EurAsia-ICT* Shiraz, Iran, 29–31 October 2002
5. IEEE 802.11 WG part 11b, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Higher Speed PHY Layer Extension in the 2.4 GHz Band, (1999)
6. B Scheuermann, W Hu, J Crowcroft. Near-optimal co-ordinated coding in wireless multihop networks, in *Paper presented at ACM CoNEXT* Columbia University New York USA 10–13 December 2007
7. E Rozner, A Iyer, Y Mehta, L Qiu, M Jafry. E R: Efficient retransmission scheme for wireless LANs, in *Paper presented at ACM CONEXT* Columbia University New York USA 10–13 December 2007
8. S Rayanchu, S Sen, J Wu, S Banerjee, S Sengupta. Loss-aware network coding for unicast wireless sessions: design, implementation, and performance evaluation, in *Paper presented at ACM SIGMETRICS* Annapolis, MD USA 2–6 June 2008
9. S Chachulski, M Jennings, S Katti, D Katabi. Trading structure for randomness in wireless opportunistic routing, in *Paper presented at ACM SIGCOMM* Kyoto Japan 27–31 August 2007
10. UCLA Engineering, Click modular router, (University of California 2013). http://read.cs.ucla.edu/click/. Accessed (10 September 2011)
11. TS Kim, S Vural, I Broustis, D Syrivelis, SV Krishnamurthy, TF La Porta. A framework for joint network coding and transmission rate control in wireless networks, in *Paper presented at IEEE INFOCOM* San Diego USA 14–19 March 2010
12. S Katti, D Katabi, H Balakrishnan, M Medard. Symbol-level network coding for wireless mesh networks, in *Paper presented at ACM SIGCOMM* Seattle USA 17–22 August 2008
13. DS Lun, N Ratnakar, M Mudard, R Koetter, DR Karger, T Ho, E Ahmed, F Zhao, Minimum-cost multicast over coded packet networks. IEEE Trans. Inform. Theory **52**(6), 2608–2623 (2006)
14. CH Liu, F Xue. Network coding for two-way relaying: rate region, sum rate and opportunistic scheduling, in *Paper presented at IEEE ICC* Beijing, China, 19–23 May 2008
15. H Seferoglu, A Markopoulou. Distributed rate control for video streaming over wireless networks with intersession network coding, in *Paper presented at seventeenth international packet video workshop* Seattle, Washington, 11–12 May 2009
16. LFM Vieira, A Misra, M Gerla. Performance of network coding in multi-rate wireless environments for multicast applications, in *Paper presented at IEEE MILCOM* Orlando, FL, USA ,29–31 October 2007
17. GS Paschos, C Fragiadakis, L Georgiadis, L Tassiulas. Wireless network coding with partial overhearing information, in *Paper presented at IEEE INFOCOM* Turin, Italy, 14–19 April 2013
18. J Bicket, D Aguayo, S Biswas, R Morris. Architecture and evaluation of an unplanned 802.11b mesh network, in *Paper presented at ACM MobiCom* Cologne, Germany, 28 August–2 September 2005
19. MIT CSAIL MIT Roofnet, (Massachusetts Insitute of Technology, 2013). http://pdos.csail.mit.edu/roofnet. Accessed (10 September 2011)
20. A Partow, SDBM hash function. (Arash Partow, 2013). http://www.partow.net/programming/hashfunctions/. Accessed (15 June 2009)
21. DSJ De Couto, D Aguayo, J Bicket, R Morris. A high throughput path metric for multiHop wireless routing, in *Paper presented at ACM MobiCom* San Diego, California, 14–19 September 2003
22. R Draves, J Padhye, B Zill. Routing in multi-radio, multi-hop wireless mesh networks, in *Paper presented at ACM MobiCom* Philadelphia, PA, USA, 26 September–1 October 2004
23. D Raychaudhuri, I Seskar, M Ott, S Ganu, K Ramachandran, H Kremo, R Siracusa, H Liu, M Singh, Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. IEEE WCNC. **3**, 1664–1669 (2005)
24. L Tassiulas, A Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. IEEE T. Automat. Contr. **37**(12), 1936–1948 (1992)
25. N Mckeown, A Mekkittikul, V Anantharam, J Walrand, Achieving 100% throughput in an input-queued switch. IEEE Trans. Commun. **47**(8), 1260–1267 (1999)
26. A Eryilmaz, DS Lun. Control for inter-session network coding., in *Paper presented at workshop on network coding (NetCod), theory and applications* Chinese University of Hong Kong, 3–4 January 2007
27. GS Paschos, L Georgiadis, L Tassiulas. Optimal scheduling of pairwise XORs under statistical overhearing and feedback, in *Paper presented at the 7th international workshop on resource allocation and cooperation in wireless networks–in conjunction with WiOpt 2011* Princeton NJ, 13 May 2011
28. GS Paschos, L Georgiadis, L Tassiulas, Scheduling with pairwise XORing of packets under statistical overhearing information and feedback. Queueing Syst. **72**(3–4), 361–395 (2012)
29. SourceForge, Iperf bandwidth measurement tool (NLANR/DAST, 2013). http://iperf.sourceforge.net/. Accessed (15 June 2009)