

Multiuser broadcast erasure channel with feedback – capacity and algorithms

Marios Gatzianas, Leonidas Georgiadis, Leandros Tassiulas

► **To cite this version:**

Marios Gatzianas, Leonidas Georgiadis, Leandros Tassiulas. Multiuser broadcast erasure channel with feedback – capacity and algorithms. NET-COOP 2010 - 4th Workshop on Network Control and Optimization, Nov 2010, Gent, Belgium. 2010. <inria-00596234>

HAL Id: inria-00596234

<https://hal.inria.fr/inria-00596234>

Submitted on 26 May 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multiuser broadcast erasure channel with feedback — capacity and algorithms

Marios Gatzianas^{*§}, Leonidas Georgiadis^{*§} and Leandros Tassioulas^{†§}

^{*}Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, Thessaloniki, 54 124, Greece.

[†]Computer Engineering and Telecommunications Department, University of Thessaly, Volos, 38 221, Greece.

[§]Center for Research and Technology Hellas, Thessaloniki, 60 361, Greece.

Emails: {mgkatzia, leonid}@auth.gr, leandros@uth.gr

Abstract—We consider the N -user broadcast erasure channel where feedback from the users is fed back to the transmitter in the form of ACK messages. We first provide a generic outer bound to the capacity of this system; we then propose a coding algorithm, compute its throughput region and show that it achieves capacity under certain conditions on channel statistics, assuming that instantaneous feedback is known to all users. Removing this assumption results in a rate region which differs from the outer bound by a factor $O(N/L)$, where L is packet length. For the case of non-symmetric channels, we present a modification of the previous algorithm whose achievable region is identical to the outer bound for $N \leq 3$, when instant feedback is known to all users, and differs from the bound by $O(N/L)$ when each user knows only its own ACK. The proposed algorithms do not require any prior knowledge of channel statistics.

Index Terms—Broadcast erasure channels, feedback-based coding, capacity achieving algorithms.

I. INTRODUCTION

Broadcast channels have been extensively studied by the information theory community since their introduction in [1]. Although their capacity remains unknown in the general case, special cases have been solved, including the important category of “degraded” channels [2]. Another class of channels that has received significant attention is erasure channels, where either the receiver receives the input symbol unaltered or the input symbol is erased (equivalently, dropped) at the receiver. The latter class is usually employed as a model for lossy packet networks.

Combining the above classes, a broadcast erasure channel (BEC) is a suitable abstraction for wireless communications modeling since it captures the broadcast nature of the medium as well as the potential for packet loss (due to fading, packet collision etc). Since this channel is not always degraded, the computation of its feedback capacity region is an open problem. Numerous variations of this channel, under different assumptions, have been studied, a brief summary of which follows.

For multicast traffic, an outer bound to the capacity region of erasure channels is derived in [3], in the form of a suitably defined minimum cut, and it is proved that the bound can be achieved by linear coding at intermediate nodes. The broadcast nature is captured by requiring each node to transmit the

same signal on all its outgoing links, while it is assumed that the destinations have complete knowledge of any erasures that occurred on *all* source-destination paths. In a sense, [3] is the “wireless” counterpart to the classical network coding paradigm of [4], since it carries all results of [4] (which were based on the assumption of error-free channels) into the wireless regime.

The concept of combining packets for efficient transmission based on receiver feedback is also used in [5], where broadcast traffic is assumed and a rate-optimal, zero-delay, offline algorithm is presented for 3 users. Online heuristics that attempt to minimize the decoding delay are also provided. Reference [6] expands on this work by presenting an online algorithm that solves at each slot a (NP-hard) set packing problem in order to decide which packets to combine. This algorithm also aims in minimizing delay.

Multiple unicast flows, which are traditionally difficult to handle within the network coding framework, are studied in [7] for a network where each source is connected to a relay as well as to all destinations, other than its own, and all connections are modeled as BECs. A capacity outer bound is presented for an arbitrary number of users N and is shown to be achievable for $N = 3$ and almost achievable for $N = 4, 5$. The capacity-achieving algorithm operates in two stages with the relay having knowledge of the destination message side information at the end of the first stage but not afterward (i.e. once the second stage starts, the relay does not receive feedback from the destinations).

A similar setting is studied in [8], where ACK-based packet combining is proposed and emphasis is placed on the overhead and complexity requirements of the proposed scheme. An actual implementation of the use of packet XORing in an intermediate layer between the IP and 802.11 MAC layers is presented and evaluated in [9], while [10] proposes a replacement for the 802.11 retransmission scheme based on exploiting knowledge of previously received packets.

This paper expands upon earlier work in [11] (which studied the case $N = 2$) and differs from the aforementioned works in that, although it also uses the idea of packet mixing (similar to the network coding sense), it provides explicit performance guarantees. Specifically, an outer bound to the feedback capacity region for multiple unicast flows (one for each user) is

[‡]This work was supported by the EU project N-CRAVE, FP7-215252.

computed and two online algorithms are presented that achieve this bound for the following settings, respectively: arbitrary N with channel statistics that satisfy a specific order relation (this includes the cases of symmetric and spatially independent channels with fairness constraints), and $N = 3$ with arbitrary channel statistics.

The algorithms do not require any knowledge of channel parameters (such as erasure probabilities) or future events so that they can be applied to any BEC. They use receiver feedback to combine packets intended for different users into a single packet which is then transmitted. The combining scheme (i.e. choosing which packets to combine and how) relies on a set of virtual queues, maintained in the transmitter, which are updated based on per-slot available receiver ACK/NACKs. This queue-based coding concept has also been used in [12], albeit for broadcast traffic with stochastic arrivals where the stability region of the proposed algorithm becomes asymptotically optimal as the erasure probability goes to 0, whereas we consider systems with an arbitrarily fixed number of packets per unicast session where the capacity is achieved for arbitrary values of erasure probability.

We were recently informed that C. Wang has independently studied in [13] the same problem as appears here and proposed coding algorithms that achieve capacity under the same conditions as ours. Although the two works share common ideas (namely, introducing degraded channels to derive capacity outer bounds and performing packet coding based on receiver feedback), the proposed algorithms (including the procedures for handling overhead), as well as the methodology used for deriving their throughput regions, are quite different.

The paper is structured as follows. Section II describes the exact model under investigation and provides the necessary definitions in order to derive the capacity outer bound in Section III. The first coding algorithm is presented in Section IV, which also contains a discussion of the intuition behind the algorithm, its correctness and optimal performance under certain conditions. The incorporation of overhead and the corresponding reduction in the achievable region are also examined. A modification of the algorithm that achieves capacity for 3 users under arbitrary channel conditions is presented in Section V, while Section VI concludes the paper. Due to space restrictions, the proofs of all stated results are omitted and presented in [14] instead.

II. SYSTEM MODEL AND DEFINITIONS

Consider a time slotted system where messages (packets) of length L bits are transmitted in each slot. We normalize to unity the actual time required to transmit a single bit so that the time interval $[(l-1)L, lL)$, for $l = 1, 2, \dots$, corresponds to slot l . The system consists of a single transmitter and a set $\mathcal{N} \triangleq \{1, 2, \dots, N\}$ of receivers, while there exists at the transmitter a distinct set \mathcal{K}_i of unicast packets destined for each user i . The channel is modeled as memoryless¹ broadcast erasure

¹the memoryless property assumption is necessary for keeping the analysis at a tractable level. In fact, Lemmas 2, 3, which form the starting point of this work, depend crucially on this assumption.

so that each broadcast packet is either received unaltered by a user or is dropped (i.e. the user does not receive it), in which case an erasure occurs for the user. This is equivalent to considering that the user receives the special symbol E , which is distinct from any transmitted symbol. Hence, each user knows whether an erasure has occurred or not by examining its received symbol.

Define $Z_{i,l} \triangleq \mathbb{I}[\text{user } i \text{ receives } E \text{ in slot } l]$, where $\mathbb{I}[\cdot]$ denotes an indicator function, and consider the random vector $\mathbf{Z}_l = (Z_{1,l}, Z_{2,l}, \dots, Z_{N,l})$. The sequence $\{\mathbf{Z}_l\}_{l=1}^{\infty}$ is assumed to consist of iid vectors (we denote with $\mathbf{Z} = (Z_1, \dots, Z_N)$ the random vector with distribution equal to that of \mathbf{Z}_l), although, for a fixed slot, arbitrary correlation between erasures in different users is allowed. For any index set $\mathcal{I} \subseteq \mathcal{N}$, we define the probability that an erasure occurs to all users in \mathcal{I} as $\Pr(Z_i = 1, \forall i \in \mathcal{I}) \triangleq \varepsilon_{\mathcal{I}}$, where, by convention, it holds $\varepsilon_{\emptyset} = 1$. For simplicity, we write ε_i instead of $\varepsilon_{\{i\}}$ and assume $\varepsilon_i < 1$ to avoid trivial cases.

According to the introduced notation, when the transmitter, at the beginning of slot l , broadcasts symbol X_l , each user i receives symbol $Y_{i,l}$ given by $Y_{i,l} = Z_{i,l}E + (1 - Z_{i,l})X_l$, where we denote $\mathbf{Y}_l \triangleq (Y_{i,l})_{i \in \mathcal{N}}$. At the end of each slot l , all users inform the transmitter whether the symbol was received or not, which is equivalent to each user i sending the value of $Z_{i,l}$ through an error-free control channel. In information-theoretic terms [15], the broadcast channel is described by the input alphabet \mathcal{X} , the output alphabets $\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_N$ for users $1, 2, \dots, N$, respectively, and the probability transition function $p(\mathbf{Y}_l | X_l)$. Due to the memoryless property, the transition probability function is independent of l , so that it can be written as $p(\mathbf{Y} | X)$. In the rest of the paper, we set $\mathcal{X} = \mathbb{F}_q$, with \mathbb{F}_q a suitable field of size q , so that, by definition of erasure channel, it holds $\mathcal{Y}_i = \mathcal{X} \cup \{E\}$ for all $i \in \mathcal{N}$.

We use the standard definitions of channel codes with feedback, probability of erroneous decoding and achievable rates from [15, Chap. 15], which we omit due to space restrictions. The following definition, introduced in [2], will also be useful in deriving the BEC capacity outer bound.

Definition 1: A broadcast, not necessarily erasure, channel $(\mathcal{X}, (\mathcal{Y}_i)_{i \in \mathcal{N}}, p(\mathbf{Y} | X))$ with receiver set \mathcal{N} is physically degraded if there exists a permutation π on \mathcal{N} such that the sequence $X \rightarrow Y_{\pi(1)} \rightarrow \dots \rightarrow Y_{\pi(N)}$ forms a Markov chain. A generalization to N users of the 2-user proof in [16] provides the following result.

Lemma 1: Feedback does not increase the capacity region of a physically degraded broadcast channel. We now have all necessary tools to compute the actual capacity outer bound.

III. CAPACITY OUTER BOUND

The derivation of the capacity outer bound is based on a method similar to the approaches in [17]–[19]. We initially state a general result on the capacity of broadcast erasure channels *without feedback* [20].

Lemma 2: The capacity region (measured in information bits per transmitted symbol) of a broadcast erasure channel

with receiver set \mathcal{N} and no feedback is

$$C_{nf} = \left\{ \mathbf{R} \geq \mathbf{0} : \sum_{i \in \mathcal{N}} \frac{R_i}{1 - \varepsilon_i} \leq L \right\}, \quad (1)$$

which implies that a simple timesharing, between the users, algorithm achieves capacity.

We denote with C the channel under consideration and, for an arbitrary permutation π on \mathcal{N} , introduce a new, hypothetical, broadcast channel \hat{C}_π with the same input/output alphabets as C and an erasure indicator function of $\hat{Z}_{\pi(i),l} = \prod_{j=1}^i Z_{\pi(j),l}$. In other words, a user $\pi(i)$ in \hat{C}_π erases a symbol if and only if all users $\pi(j)$, with $j \leq i$, erase the symbol in channel C . This occurs with probability $\hat{\varepsilon}_{\pi(i)} \triangleq \varepsilon_{\cup_{j=1}^i \{\pi(j)\}}$. The following two results are proved in [14].

Lemma 3: Channel \hat{C}_π is physically degraded.

Lemma 4: Denote with \mathcal{C}_f , $\hat{\mathcal{C}}_{\pi,f}$ the feedback capacity regions of channels C , \hat{C}_π , respectively. It holds $\mathcal{C}_f \subseteq \hat{\mathcal{C}}_{\pi,f}$.

Notice that Lemma 4 already provides an outer bound to \mathcal{C}_f . In order to derive this bound, we note that the previous results imply that the feedback capacity region of the physically degraded channel \hat{C}_π is identical, due to Lemma 1, to the capacity region of \hat{C}_π without feedback. The latter is described, in general form, in Lemma 2 whence we deduce the following result.

Lemma 5: The feedback capacity region of \hat{C}_π is given by

$$\hat{\mathcal{C}}_{\pi,f} = \left\{ \mathbf{R} \geq \mathbf{0} : \sum_{i \in \mathcal{N}} \frac{R_{\pi(i)}}{1 - \hat{\varepsilon}_{\pi(i)}} \leq L \right\}. \quad (2)$$

The above analysis was based on a particular permutation π . Considering all $N!$ permutations on \mathcal{N} provides a tighter general outer bound.

Theorem 1: Denote with \mathcal{P} the set of all possible permutations on \mathcal{N} . It then holds $\mathcal{C}_f \subseteq \mathcal{C}^{out} \triangleq \bigcap_{\pi \in \mathcal{P}} \hat{\mathcal{C}}_{\pi,f}$.

IV. CODING ALGORITHM

In this section, we present a coding algorithm named CODE1, show its correctness, characterize its throughput region and provide conditions under which it achieves capacity. It turns out that CODE1 achieves capacity for the following special cases, among others:

- symmetric channels, i.e. channels which satisfy the condition $\varepsilon_{\mathcal{I}} = \varepsilon_{\mathcal{J}}$ whenever $|\mathcal{I}| = |\mathcal{J}|$, for any $\mathcal{I}, \mathcal{J} \subseteq \mathcal{N}$.
- spatially independent channels (i.e. $\varepsilon_{\mathcal{I}} = \prod_{i \in \mathcal{I}} \varepsilon_i$ for all $\mathcal{I} \subseteq \mathcal{N}$) with one-sided fairness constraints. The latter notion, which appears in [13], is defined as follows: a rate $\mathbf{R} = (R_1, \dots, R_N)$ is one-sided fair iff it holds $\varepsilon_i R_i \geq \varepsilon_j R_j$ for all $i < j$. Denoting with \mathcal{R}_{fair} the set of one-sided fair rates, we are searching for algorithms that achieve a rate region of $\mathcal{C}^{out} \cap \mathcal{R}_{fair}$.

It should be mentioned that the proposed algorithms in [13] also achieve capacity for both of the above settings.

In the following, we assume that each user knows the size $|\mathcal{K}_i|$ of all sessions and instant feedback is available to all users. The first assumption can be easily satisfied in practice while the second one will be removed in a later section.

Before the algorithm's description, a brief discussion of its underlying rationale will be useful. Since each user i must decode exactly $|\mathcal{K}_i|$ packets, one way of achieving this is by sending linear combinations, over the field \mathbb{F}_q , of appropriate packets so that user i eventually receives $|\mathcal{K}_i|$ linearly independent combinations of the packets in \mathcal{K}_i . Specifically, all packets in \mathcal{K}_i are viewed as elements of \mathbb{F}_q , while each transmitted symbol (or packet) s has the form $s = \sum_{p \in \cup_{i \in \mathcal{N}} \mathcal{K}_i} a_s(p)p$, where $a_s(p)$ are suitable coefficients in \mathbb{F}_q . If the symbol s can also be written as

$$s = \sum_{p \in \mathcal{K}_i} b_s(p)p + c_s, \quad (3)$$

where $\mathbf{b}_s \triangleq (b_s(p), p \in \mathcal{K}_i)$, c_s are known to user i , then s is considered to be a "token" for i . Additionally, if s is received by i and the \mathbf{b}_s coefficients of s , along with the $\mathbf{b}_{s'}$ coefficients of all previously received (by i) tokens s' , form a linearly independent set of vectors over \mathbb{F}_q , then s is considered to be an "innovative token" for i . In words, an innovative token for i is any packet s that allows i to effectively construct a new equation (with the packets in \mathcal{K}_i as unknowns, since \mathbf{b}_s , c_s are known), that is linearly independent w.r.t. all previously constructed equations by i . Hence, each user i must receive $|\mathcal{K}_i|$ innovative tokens in order to decode its packets. Note that it is quite possible, and actually very desirable, for the same packet to be a token (better yet, an innovative token) for multiple users.

In order to avoid inefficiency and, hopefully, achieve the outer bound of Section III, it is crucial that, under certain circumstances, a symbol (i.e. a linear combination of packets) that is erased by some users, but is received by at least one other user, is stored in an appropriate queue so that it can be combined in the future with other erased symbols to provide tokens for multiple users (and thus compensate for the loss). The crux of the algorithm is in the careful bookkeeping required to handle these cases.

A. Description of algorithm CODE1

For the reader's convenience, algorithm CODE1 is succinctly presented in pseudocode in Fig. 1. Specifically, the transmitter maintains a network of virtual queues $Q_{\mathcal{S}}$, indexed by the non-empty subsets \mathcal{S} of \mathcal{N} (see Fig. 2 for an illustration for 4 users). The queues are initialized with the unicast packets as follows: $Q_{\mathcal{S}} = \mathcal{K}_i$ if $\mathcal{S} = \{i\}$, otherwise $Q_{\mathcal{S}} = \emptyset$. Additionally, with each queue $Q_{\mathcal{S}}$, indices $T_{\mathcal{S}}^i$ are maintained for all $i \in \mathcal{S}$ and are initialized as: $T_{\mathcal{S}}^i = |\mathcal{K}_i|$ if $\mathcal{S} = \{i\}$, otherwise $T_{\mathcal{S}}^i = 0$.

It will become apparent from the algorithm's description that index $T_{\mathcal{S}}^i$ represents the number of *innovative* tokens (i.e. packets of the form in (3)) that user i must receive successfully from $Q_{\mathcal{S}}$ in order to decode its packets² (due to the performed initialization, this statement is trivially true

²the transmitted combination of packets from $Q_{\mathcal{S}}$ can never become a token for any user $i \in \mathcal{N} - \mathcal{S}$, so that the transmitter does not need to maintain indices for them.

Algorithm CODE1

```

1: initialize  $Q_S, T_S^i$  for all  $S \subseteq \mathcal{N}$  and  $i \in S$ ;
2: for  $phase \leftarrow 1, \dots, N$  do
3:   for all  $Q_S \in \mathcal{Q}_n$  do  $\triangleright$  arbitrary order in  $\mathcal{Q}_n$ 
4:     while ( $T_S^i > 0$  for at least one  $i \in S$ ) do
5:       compute suitable coefficients ( $a_s(p), p \in Q_S$ );
6:       transmit packet  $s = \sum_{p \in Q_S} a_s(p)p$ ;
7:       apply ACTFB1 based on feedback for  $s$ ;
8:     end while
9:   end for
10: end for

```

Fig. 1. Pseudocode for algorithm CODE1.

for all S with $|\mathcal{S}| = 1$). These indices are dynamically updated during the algorithm's execution based on the received feedback, as will be explained soon. Finally, each receiver $i \in \mathcal{N}$ maintains its own set of queues R_S^i , for all non-empty $S \subseteq \mathcal{N}$ with $i \in S$, where it stores the innovative tokens it receives from Q_S . We assume for now that all users know which queue the packet they receive comes from. All queues R_S^i are initially empty.

Denote with \mathcal{Q}_n the set of all queues Q_S with $|\mathcal{S}| = n$. The algorithm operates in N phases so that in phase n , with $1 \leq n \leq N$, only transmissions of linear combinations of packets in one of the queues in \mathcal{Q}_n occur. Specifically, at phase n , the transmitter orders the set \mathcal{Q}_n according to a predetermined rule, known to all users (say, according to lexicographic order, which corresponds to the top-to-bottom ordering shown in Fig. 2). The transmitter then examines the first (according to this order) queue Q_S and transmits a symbol (or packet) s that is a linear combination of all packets in Q_S , i.e. $s = \sum_{p \in Q_S} a_s(p)p$. We slightly abuse parlance and say that “ s is transmitted from Q_S ”, although it is clear that s is not actually stored in Q_S . The coefficients $a_s(p) \in \mathbb{F}_q$ can be produced either via a pseudo-random number generator or through structured codes. The exact generation method for $a_s(p)$ is unimportant as long as the following requirements are met:

- the generation procedure is known to all users, so that they can always reproduce the values of $a_s(p)$ even when they don't receive packet s . This implies that the receivers must also know the size of all queues $Q_S, S \subseteq \mathcal{N}$, at all times.
- the set of coefficient vectors ($a_s(p) : p \in Q_S$), for all packets (i.e. linear combinations) s transmitted from Q_S , is a linearly independent set of vectors over \mathbb{F}_q .

If the coefficients $a_s(p)$ are randomly generated, the generation procedure can be augmented, for robustness purposes, with a subroutine that checks whether the linear independence condition is indeed true and (if false) creates a new set of coefficients until the condition becomes true. This augmented algorithm implies *zero* probability of erroneous decoding, albeit at increased computational cost (due to the additional




Fig. 2. Transmitter virtual queues required for 4 users and some possible index transitions.

checking routine). If the computational cost becomes excessive, the second condition can be relaxed so that it is satisfied with probability arbitrarily close to 1, for sufficiently large field size q , which leads to a non-zero (but arbitrarily small) probability of error.

Depending on the received feedback for the packet s transmitted from queue Q_S , the following steps, collectively referred to as ACTFB1, are taken (all 4 cases must be examined)

- 1) if no user in \mathcal{N} receives s , it is retransmitted.
- 2) for each user $i \in \mathcal{S}$ that receives s and satisfies $T_S^i > 0$, s is added to queue R_S^i and T_S^i is decreased by 1.
- 3) if s has been erased by at least one user $i \in \mathcal{S}$ and has been received by *all* users in some set \mathcal{G} , with $\emptyset \neq \mathcal{G} \subseteq \mathcal{N} - \mathcal{S}$, the following 2 steps are performed
 - packet s is added to queue $Q_{\mathcal{S} \cup \mathcal{G}}$ (no packets are removed from Q_S).
 - for each user $i \in \mathcal{S}$ that erased s and satisfies $T_S^i > 0$, T_S^i is reduced by 1 and $T_{\mathcal{S} \cup \mathcal{G}}^i$ is increased by 1.
- 4) if the set \mathcal{G} of users that receive s is a subset of \mathcal{S} such that $T_S^i = 0$ for all $i \in \mathcal{G}$, s is retransmitted.

Fig. 2 presents the allowable index transitions from queues $Q_{\{1\}}, Q_{\{1,3\}}$ that occur in step 3 of ACTFB1 (the other transitions are not shown to avoid graphical clutter; dashed lines correspond to step 2 of ACTFB1). Transmission of linear combinations of packets from Q_S continues for as long as there exists at least one $i \in \mathcal{S}$ with $T_S^i > 0$. When it holds $T_S^i = 0$ for all $i \in \mathcal{S}$, the transmitter moves to the next queue $Q_{S'}$ in the ordering of \mathcal{Q}_n and repeats the above procedure until it has visited all queues in \mathcal{Q}_n . When this occurs, phase n is complete and the algorithm moves to phase $n+1$. CODE1 terminates at the end of phase N .

B. Properties and correctness of CODE1

The second statement in the following Lemma, which is rigorously proved in [14] although it can be intuitively

ascertained through induction on $|\mathcal{S}|$, is the crucial property of CODE1 and follows from its construction.

Lemma 6: Any packet s that is stored in queue $Q_{\mathcal{S}}$ with $|\mathcal{S}| \geq 2$ is a linear combination of all packets in queue $Q_{\mathcal{I}_s}$, for some non-empty $\mathcal{I}_s \subset \mathcal{S}$, that has been received by *exactly* all users in $\mathcal{S} - \mathcal{I}_s$. Hence, any packet in queue $Q_{\mathcal{S}}$ is a token for all $i \in \mathcal{S}$ (and only these $i \in \mathcal{S}$), and there exist coefficients $a_s(p)$ such that the linear combination of all packets in $Q_{\mathcal{S}}$ is an innovative token for all $i \in \mathcal{S}$ with $T_{\mathcal{S}}^i > 0$.

The above Lemma gives a very intuitive explanation to the algorithm's operation. Specifically, step 2 of ACTFB1 is equivalent to saying that whenever user i receives a useful token (meaning that $T_{\mathcal{S}}^i > 0$ so that there remain innovative tokens to receive) from $Q_{\mathcal{S}}$, this (innovative) token should be added to $R_{\mathcal{S}}^i$. If this is not the case and there exist users, comprising set $\mathcal{G} \subseteq \mathcal{N} - \mathcal{S}$, who receive this packet (step 3), then the packet has become a token for users in $\mathcal{S} \cup \mathcal{G}$ and should be placed in queue $Q_{\mathcal{S} \cup \mathcal{G}}$. This allows the token to be simultaneously received by multiple users in the future and thus compensate for the current loss. Additionally, since user i can now recover this token more efficiently from $Q_{\mathcal{S} \cup \mathcal{G}}$ instead of $Q_{\mathcal{S}}$, the indices $T_{\mathcal{S}}^i, T_{\mathcal{S} \cup \mathcal{G}}^i$ should be modified accordingly to account for the token transition. Step 4 of ACTFB1 merely states that the packet is retransmitted when it is only received by users who have already recovered all tokens intended for them.

Finally, since for any slot t that some $T_{\mathcal{S}}^i$ is reduced by 1, either some other $T_{\mathcal{S} \cup \mathcal{G}}^i$ is increased by 1 or (exclusive or) some packet is added to queue $R_{\mathcal{S}}^i$, it follows that the quantity $\sum_{\mathcal{S}: i \in \mathcal{S}} |R_{\mathcal{S}}^i(t)| + \sum_{\mathcal{S}: i \in \mathcal{S}} T_{\mathcal{S}}^i(t)$ is constant during the execution of CODE1 and, due to the performed initialization, is equal to $|\mathcal{K}_i|$ for all $i \in \mathcal{N}$. Since the algorithm terminates when it holds $T_{\mathcal{S}}^i = 0$ for all non-empty $\mathcal{S} \subseteq \mathcal{N}$ and all $i \in \mathcal{S}$, we conclude that at the end of the terminating slot t_f it holds $\sum_{\mathcal{S}: i \in \mathcal{S}} |R_{\mathcal{S}}^i(t_f)| = |\mathcal{K}_i|$ for all $i \in \mathcal{N}$. Hence, each user has recovered $|\mathcal{K}_i|$ tokens which, by choosing a sufficiently large field size q (which in turn implies a sufficiently large L), can be made linearly independent with probability 1 (or arbitrarily close to 1, depending on whether or not the independence checking routine is used during the coefficient generation procedure). Thus, all users can decode their packets with either zero or a vanishing probability of error and CODE1 operates correctly. Notice that this result holds for arbitrary BECs, so that, in principle, CODE1 is universally applicable. In addition, no prior knowledge of channel parameters is required for its execution.

C. Performance of CODE1

The complete analysis of the performance of CODE1 is quite lengthy with full details being given in [14]. We present here the starting point of the analysis along with the main results. We assume that $\min_{i \in \mathcal{N}} |\mathcal{K}_i|$ is sufficiently large to invoke the strong law of large numbers, and define the events $E_{\mathcal{S}} \triangleq \{Z_i = 1, \forall i \in \mathcal{S}\}$ and $R_{\mathcal{G}} \triangleq \{Z_i = 0, \forall i \in \mathcal{G}\}$, which

imply (c stands for set complement and \uplus for disjoint union)

$$R_{\mathcal{G}}^c = \biguplus_{\mathcal{H} \neq \emptyset: \mathcal{H} \subseteq \mathcal{G}} (E_{\mathcal{H}} \cap R_{\mathcal{G} - \mathcal{H}}). \quad (4)$$

For completeness, we define $E_{\emptyset} = R_{\emptyset} = \Omega$ (the sample space). Denoting with $p_{\mathcal{S}, \mathcal{G}} \triangleq \Pr(E_{\mathcal{S}} \cap R_{\mathcal{G}})$ the probability that a packet is received *exactly* by all users in \mathcal{G} , and combining the identity $E_{\mathcal{S}} = (E_{\mathcal{S}} \cap R_{\mathcal{G}}) \uplus (E_{\mathcal{S}} \cap R_{\mathcal{G}}^c)$ with (4) yields the following recursion for $p_{\mathcal{S}, \mathcal{G}}$ in terms of $\varepsilon_{\mathcal{S}}$

$$p_{\mathcal{S}, \mathcal{G}} = \varepsilon_{\mathcal{S}} - \sum_{\mathcal{H} \neq \emptyset: \mathcal{H} \subseteq \mathcal{G}} p_{\mathcal{S} \cup \mathcal{H}, \mathcal{G} - \mathcal{H}}, \quad (5)$$

which is solved in closed form in [14].

If T^{**} is the number of slots (including retransmissions due to steps 1,4 of ACTFB1) required for all users to fully decode their packets under CODE1, the achieved rate for user i (in information symbols per transmission) is $R_i = |\mathcal{K}_i|/T^{**}$; therefore, we only need to compute T^{**} . Denoting with $T_{\mathcal{S}}^*$ the number of slots required for the processing of queue $Q_{\mathcal{S}}$ under CODE1, it holds

$$T^{**} = \sum_{\emptyset \neq \mathcal{S} \subseteq \mathcal{N}} T_{\mathcal{S}}^*, \quad (6)$$

$$T_{\mathcal{S}}^* = \max_{i \in \mathcal{S}} \left(\frac{k_{\mathcal{S}}^i}{1 - \varepsilon_{\mathcal{N} - (\mathcal{S} - \{i\})}} \right),$$

with $k_{\mathcal{S}}^i$ being the value of index $T_{\mathcal{S}}^i$ just before processing of $Q_{\mathcal{S}}$ begins. Hence, the throughput of CODE1 is essentially known once $k_{\mathcal{S}}^i$ is computed.

Step 3 of ACTFB1 also implies that $k_{\mathcal{S}}^i = \sum_{\emptyset \neq \mathcal{I} \subseteq \mathcal{S}} k_{\mathcal{I}, \mathcal{S}}^i$, where $k_{\mathcal{I}, \mathcal{S}}^i$ is the number of linear combinations of packets in $Q_{\mathcal{I}}$ (with $i \in \mathcal{I}$) that are erased by i and received by $\mathcal{S} - \mathcal{I}$ (so that step 3 of ACTFB1 is applicable). The previous relation, which can be interpreted as a principle of tokens conservation, leads to the following recursion

$$k_{\mathcal{S}}^i = \sum_{\substack{\emptyset \neq \mathcal{I} \subseteq \mathcal{S} \\ i \in \mathcal{I}}} \frac{k_{\mathcal{I}}^i}{1 - \varepsilon_{\mathcal{N} - (\mathcal{I} - \{i\})}} p_{\mathcal{N} - (\mathcal{S} - \{i\}), \mathcal{S} - \mathcal{I}}, \quad (7)$$

for all \mathcal{S} with $|\mathcal{S}| \geq 2$.

It turns out that the last recursion has an explicit solution [14], from which T^{**} (and therefore R_i) can be computed through (6) to produce the two following results:

Theorem 2: Denote $f_{\mathcal{S}}^i \triangleq \sum_{\mathcal{H} \subseteq \mathcal{S} - \{i\}} \frac{(-1)^{|\mathcal{S}| - |\mathcal{H}| - 1}}{1 - \varepsilon_{\mathcal{N} - \mathcal{H}}}$. For arbitrary channel statistics, the throughput region of CODE1 is

$$\mathcal{R}_{\text{CODE1}} = \left\{ \mathbf{R} : \sum_{\mathcal{S} \subseteq \mathcal{N}} \max_{i \in \mathcal{S}} (R_i f_{\mathcal{S}}^i) \leq L \right\}.$$

Theorem 3: If the nodes in \mathcal{N} can be rearranged via a renaming (i.e. a permutation) of indices so that, for the renamed nodes $i \in \mathcal{N}$ and for all $\mathbf{R} \in \mathcal{C}^{out}$, it holds $\arg \max_{i \in \mathcal{S}} (f_{\mathcal{S}}^i R_i) = \min\{i \in \mathcal{S}\}$, for all $\mathcal{S} \subseteq \mathcal{N}$, then CODE1 achieves \mathcal{C}^{out} . Specifically, it holds

$$\mathcal{R}_{\text{CODE1}} = \mathcal{C}^{out} = \left\{ \mathbf{R} \geq \mathbf{0} : \sum_{i=1}^N \frac{R_i}{1 - \varepsilon_{\{1, \dots, i\}}} \leq L \right\}.$$

It is shown in [14] that the conditions of Theorem 3 hold for symmetric channels as well as spatially independent channels with fairness constraints, so that CODE1 achieves the capacity in both settings.

D. Taking the overhead into account

The previous analysis rests on the assumption that complete feedback is available to all users. To remove this assumption (so that each user need only know its own feedback), the feedback information must be conveyed to the receivers by the transmitter at the expense of channel capacity (i.e. the incorporation of overhead) and increased complexity at the receivers. The following procedure is proposed, under the assumption that the transmitter and the receivers use the same random number generator, with the same seed, to generate random packet coefficients.

A single overhead bit h , initially set to 0, is reserved in each packet of length L . The transmitter executes CODE1 normally, taking the received feedback into account according to ACTFB1. For each transmitted linear combination s (including retransmissions due to steps 1, 4 of ACTFB1), the transmitter also creates an N -bit group (f_1, \dots, f_N) , where f_i is 1 or 0, depending on whether or not user i received s , respectively. Since, by Theorem 2, the number of slots required by CODE1 to process all queues is $T^{**} = \sum_{S \subseteq \mathcal{N}} \max_{i \in S} (f_S^i | \mathcal{K}_i|)$, an equal number of N -bit groups is created. Meanwhile, each user stores the packets it receives in a single queue since, at this point, it can do nothing more without additional information on the other users' feedback.

When CODE1 terminates (i.e. phase N is complete), the transmitter splits the entire feedback log into packets of length L (so that a total of NT^{**}/L packets is required) and broadcasts random linear combinations of these packets until all users receive NT^{**}/L linearly independent combinations (this procedure, which was also used in [11], is more efficient than repeatedly transmitting each feedback packet unencoded until all users receive it). All packets containing feedback log information are marked with $h = 1$, so that the receivers can distinguish them from pure information packets. Once a user has received the necessary linearly independent combinations, it can decode its feedback packets and accurately recreate the entire feedback log, so that it can essentially "replay" the execution of CODE1 (provided that the order in which the queues $Q_S \in \mathcal{Q}_n$ are visited is known a priori).

The following result now follows from the fact that, for sufficiently large T^{**} , the number of slots required for all users to receive NT^{**}/L linearly independent combinations of the feedback packets is $NT^{**}/(L(1 - \varepsilon_{max}))$, where $\varepsilon_{max} = \max_{i \in \mathcal{N}} \varepsilon_i$.

Theorem 4: Under the overhead accounting scheme described above, CODE1 achieves the following rate region for arbitrary channels

$$\mathcal{R}_{\text{CODE1}}^{\text{over}} = \left\{ \mathbf{R} : \sum_{S \subseteq \mathcal{N}} \max_{i \in S} (R_i f_S^i) \leq \frac{L-1}{1 + \frac{N}{L(1-\varepsilon_{max})}} \right\}. \quad (8)$$

$\mathcal{R}_{\text{CODE1}}^{\text{over}}$ approximates $\mathcal{R}_{\text{CODE1}}$ very closely as $L \gg N/(1 - \varepsilon_{max})$, so that the overhead-induced throughput loss is minimal. An example, for $N = 50$ and $\varepsilon_{max} = 0.5$ (the latter represents very poor channel conditions; ε_{max} is usually much smaller), a length of $L = 8000$ bits leads to a throughput loss of 1.25% w.r.t. $\mathcal{R}_{\text{CODE1}}$.

It should be noted that the above overhead accounting scheme may lead to a feedback log that grows arbitrarily large (with a correspondingly arbitrarily small probability), since it includes retransmissions due to steps 1, 4 of ACTFB1. This scheme can be modified so that the size of the feedback log is deterministically bounded by a quantity that is independent of channel statistics. More details are provided in [14].

V. THE 3-RECEIVER CASE FOR ARBITRARY CHANNELS

Although CODE1 achieves the capacity outer bound of Theorem 1 for certain classes of channels, for sufficiently large L , this is not always true for arbitrary channels, i.e. there exist rates $\mathbf{R} \in \mathcal{C}^{\text{out}}$ that are *not* achievable by CODE1. This is easily verified in the following scenario: consider the case of equal rates, i.e. $R_i = R$ for all $i \in \{1, 2, 3\}$ (which implies that $|\mathcal{K}_i| = K$ for all i), and assume that it holds $\varepsilon_1 = \varepsilon_2 = \varepsilon_3$ and $\varepsilon_{\{1,2\}} > \varepsilon_{\{1,3\}} > \varepsilon_{\{2,3\}}$. Considering all possible permutations on $\{1, 2, 3\}$ and applying Theorem 1 yields the following bound

$$\mathcal{C}_{eq}^{\text{out}} = \left\{ R\mathbf{1} : R \left(\frac{1}{1 - \varepsilon_1} + \frac{1}{1 - \varepsilon_{\{1,2\}}} + \frac{1}{1 - \varepsilon_{\{1,2,3\}}} \right) \leq L \right\}. \quad (9)$$

The number of slots \tilde{T}^{**} required for the application of CODE1 in this setting is computed in [14] as

$$\tilde{T}^{**} = K \cdot \max \left[\frac{1}{1 - \varepsilon_1} + \frac{1}{1 - \varepsilon_{\{1,2\}}} + \frac{1}{1 - \varepsilon_{\{1,2,3\}}}, \frac{1 - \varepsilon_{\{2,3\}}}{(1 - \varepsilon_2)(1 - \varepsilon_{\{1,3\}})} + \frac{1}{1 - \varepsilon_{\{1,2\}}} + \frac{1}{1 - \varepsilon_{\{1,2,3\}}}, (\cdot) \right]. \quad (10)$$

The third term appearing in (10) is written as (\cdot) since it does not influence the fact that the second term is strictly larger than the first (since it holds $\varepsilon_1 = \varepsilon_2$ and $1 - \varepsilon_{\{2,3\}} > 1 - \varepsilon_{\{1,3\}}$). This implies that the rate (in information bits per transmitted symbol) $R = KL/\tilde{T}^{**}$ achieved by CODE1 is strictly smaller than the bound in (9), which demonstrates the suboptimality of CODE1.

A more intuitive explanation for the suboptimal performance of CODE1 under asymmetric channels for the 3-receiver case can also be given by the following argument. Assume that in phase 2 of CODE1, the order in which the queues are visited is $\{1, 2\}, \{1, 3\}, \{2, 3\}$. When the transmitter sends linear combinations of packets from $Q_{\{1,2\}}$, it is quite possible that the indices $T_{\{1,2\}}^1, T_{\{1,2\}}^2$ do not become zero simultaneously. Say it happens that $T_{\{1,2\}}^1 = 0$ and $T_{\{1,2\}}^2 > 0$. By construction, CODE1 will continue to transmit linear combinations from $Q_{\{1,2\}}$ until $T_{\{1,2\}}^2$ also becomes 0. However, this creates a source of inefficiency, as implied by step 4 of ACTFB1.

Specifically, if a transmitted packet s is only received by 1, step 4 of ACTFB1 will force s to be retransmitted until either 2 or 3 receive it, in a sense “wasting” this slot. We claim that there exists potential for improvement at this point, by combining the packets in $Q_{\{1,2\}}$ with the packets in $Q_{\{1,2,3\}}$. A linear combination of packets in these queues creates a token for both 1 and 2. Hence, even if the packet is received only by 1, the slot is not wasted, since 1 recovers an innovative token (provided that $T_{\{1,2,3\}}^1 > 0$). Unfortunately, the previous reasoning implies that the rule of always combining packets from a single queue must be discarded if the objective is to achieve capacity. For $N > 3$, it is not even clear what structure a capacity achieving algorithm should have. However, for $N = 3$, we present the following algorithm, named CODE2, which achieves capacity for arbitrary channels.

CODE2 operates in phases as follows. Phase 1 of CODE2 is identical to phase 1 of CODE1, with the transmitter acting according to the rules in ACTFB1 (note that step 4 of ACTFB1 cannot occur in this phase of CODE2). In phase 2 of CODE2, the transmitter orders the queues Q_S in \mathcal{Q}_2 according to an arbitrary rule and transmits linear combinations from Q_S until *at least one* user $i \in \mathcal{S}$ recovers all innovative tokens from Q_S (i.e. $T_S^i = 0$). When this occurs, the transmitter moves to the next queue in \mathcal{Q}_2 . Again, the rules in ACTFB1 are applied. When all queues in \mathcal{Q}_2 have been visited, each $Q_S \in \mathcal{Q}_2$ has at most one surviving index (meaning some $i \in \mathcal{S}$ with $T_S^i > 0$). For convenience, we denote this epoch with t_s and define the survival number $su(i)$ of index $i \in \{1, 2, 3\}$ as $su(i) \triangleq |\{\mathcal{S} : |\mathcal{S}| = 2, T_S^i(t_s) > 0\}|$, where $T_S^i(t_s)$ is the value of the index at time t_s . In words, $su(i)$ is equal to the number of queues in \mathcal{Q}_2 which contain unrecovered innovative tokens for user i at time t_s . By definition, it holds $0 \leq su(i) \leq 2$ for all $i \in \{1, 2, 3\}$. The transmitter now distinguishes cases as follows

- if it holds $su(i) = 0$ for all $i \in \{1, 2, 3\}$, CODE2 reverts to CODE1, starting at phase 3.
- if it holds $su(i) = 1$ for all $i \in \{1, 2, 3\}$, CODE2 reverts to CODE1, starting at phase 2. It can be shown [14] that, for sufficiently large $|\mathcal{K}_i|_{i=1}^3$, the probability of this event is arbitrarily small, so that the capacity region is unaffected by any actions taken henceforth.
- otherwise, there exists at least one user i^* such that $su(i^*) = 0$. In fact, simple enumeration reveals that all possible configurations for $su(i)$ fall in exactly one of the following 4 categories:
 - 1) there exist distinct users $i^*, j^*, k^* \in \{1, 2, 3\}$ such that $su(i^*) = 0, su(j^*) = 1, su(k^*) = 2$.
 - 2) there exist distinct users $i^*, j^*, k^* \in \{1, 2, 3\}$ such that $su(i^*) = 0, su(j^*) = su(k^*) = 1$.
 - 3) there exist distinct users $i^*, j^*, k^* \in \{1, 2, 3\}$ such that $su(i^*) = su(j^*) = 0$ and $su(k^*) = 2$.
 - 4) there exist distinct users $i^*, j^*, k^* \in \{1, 2, 3\}$ such that $su(i^*) = su(j^*) = 0$ and $su(k^*) = 1$.

To provide some concrete examples, Fig. 3 contains 4 possible configurations (each belonging, from left to




Fig. 3. Possible states of innovative token indices for the queues in \mathcal{Q}_2 at epoch t_s .

right, to one of the above categories), where circles are used to denote surviving indices. The values (i^*, j^*, k^*) for each configuration are $(3, 2, 1), (2, 1, 3), (3, 2, 1), (3, 2, 1)$, respectively.

The transmitter now constructs the set $\mathcal{Q}_{su} = \{Q_{\{i^*, j\}} : su(i^*) = 0, T_{\{i^*, j\}}^j > 0\}$ consisting of all queues in \mathcal{Q}_2 that contain a surviving index j and an index i^* with $su(i^*) = 0$. Relative order within \mathcal{Q}_{su} is unimportant. A subphase, called 2.1, is now initiated, in which the following actions are performed:

- the transmitter visits each queue $Q_{\{i^*, j\}}$ in \mathcal{Q}_{su} and transmits a packet s which is a linear combination of all packets in queues $Q_{\{i^*, j\}}$ and $Q_{\{1,2,3\}}$. Depending on the received feedback, the following actions, collectively referred to as ACTFB2, are taken
 - 1) if j receives s , $T_{\{i^*, j\}}^j$ is decreased by 1.
 - 2) if i^* receives s and it holds $T_{\{1,2,3\}}^{i^*} > 0$, $T_{\{1,2,3\}}^{i^*}$ is decreased by 1.
 - 3) if j erases s and $k \in \{1, 2, 3\} - \{i^*, j\}$ receives it, s is added to $Q_{\{1,2,3\}}$, $T_{\{i^*, j\}}^j$ is decreased by 1 and $T_{\{1,2,3\}}^j$ is increased by 1.
 - 4) if s is erased by all users or is received only by i^* when it holds $T_{\{1,2,3\}}^{i^*} = 0$, s is retransmitted.

Notice that, apart from step 2) in the above list, ACTFB2 is similar to ACTFB1. The above procedure is repeated until it holds $T_{\{i^*, j\}}^j = 0$, at which point the next queue in \mathcal{Q}_{su} is visited. The above procedure is repeated until all queues in \mathcal{Q}_{su} have been visited.

- once all queues in \mathcal{Q}_{su} have been processed, the transmitter computes the new values of $su(i)$ for $i \in \{1, 2, 3\}$ and constructs \mathcal{Q}_{su} from scratch. If $\mathcal{Q}_{su} = \emptyset$, CODE2 reverts to CODE1 starting at phase 3, otherwise it repeats the above procedure verbatim for the new \mathcal{Q}_{su} . It is easy to verify that at most 2 iterations of this procedure will be performed until it holds $\mathcal{Q}_{su} = \emptyset$.

As a final comment, step 4 of ACTFB2 is similar to step 4 of ACTFB1 so one could argue that CODE2 still performs inefficiently. However, by construction of \mathcal{Q}_{su} , it is easy to verify that if, during the combination of $Q_{\{i^*, j\}} \in \mathcal{Q}_{su}$ with $Q_{\{1,2,3\}}$, $T_{\{1,2,3\}}^{i^*}$ becomes 0 before $T_{\{i^*, j\}}^j$ does, then i^* has recovered all innovative tokens (i.e. it holds $T_S^{i^*} = 0$ for all $\mathcal{S} \subseteq \mathcal{N}$). Hence, i^* cannot gain any more innovative tokens

by combining $Q_{\{i^*,j\}}$ with $Q_{\{1,2,3\}}$ and no efficiency is lost.

To provide a concrete justification for the last statement, consider the application of subphase 2.1 to the leftmost configuration in Fig. 3. It holds $\mathcal{Q}_{su} = \{Q_{\{1,3\}}, Q_{\{2,3\}}\}$ and the transmitter starts combining $Q_{\{1,3\}}$ with $Q_{\{1,2,3\}}$ until $T_{\{2,3\}}^2$ becomes 0. If it happens that $T_{\{1,2,3\}}^3$ becomes 0 before $T_{\{2,3\}}^2$, then 3 has indeed recovered all innovative tokens so that, even if step 4 occurs, no efficiency gain is possible. The same conclusion is reached by examining the 3 other categories shown in Fig. 3. Hence, at the end of subphase 2.1, it holds $T_S^i = 0$ for all $i \in \mathcal{S}$ with $|\mathcal{S}| = 2$ and CODE2 reverts to CODE1 starting at phase 3. Reference [14] contains the proof of the following important result, which ensures the correctness of CODE2 (i.e. guarantees that each user i will receive $|\mathcal{K}_i|$ innovative packets)

Lemma 7: Assume that, at the beginning of subphase 2.1, it holds $T_{\{i,j\}}^i = 0$, $T_{\{i,j\}}^j > 0$. During subphase 2.1, the coefficients $a_s(p)$ of a linear combination s of all packets p in queues $Q_{\{i,j\}}$, $Q_{\{1,2,3\}}$ can be selected such that s is an innovative token for j, i , as long as it holds $T_{\{i,j\}}^j > 0$, $T_{\{1,2,3\}}^i > 0$, respectively.

The analysis of the performance of CODE2 is relatively straightforward (essentially being a repetition of the analysis of CODE1, with a careful calculation of the number of indices moved during the combination of the queues in \mathcal{Q}_2 with $Q_{\{1,2,3\}}$) but lengthy so we only present the final result [14].

Theorem 5: CODE2 achieves the capacity outer bound of C^{out} , assuming complete feedback is known to all users.

The assumption of complete feedback known to all users can be removed by overhead mechanisms essentially identical to the one described in Section IV-D, with a similar reduction in the achievable region. This issue will not be pursued any further.

VI. CONCLUSIONS

This paper presented 2 coding algorithms, CODE2 and CODE1, that achieve the feedback capacity of N -user broadcast erasure channels with multiple unicast sessions for the following cases 1) arbitrary channel statistics, for $N \leq 3$, and 2) arbitrary N and channel statistics that satisfy the general conditions of Theorem 3 (this includes symmetric and one-sided fair spatially independent channels), respectively. The main characteristic of the algorithms is the introduction of virtual queues to store packets, depending on received feedback, and the appropriate mixing of the packets to allow for simultaneous reception of innovative packets by multiple users, while none of them requires knowledge of channel statistics. Since only an outer bound to the capacity region is known for $N \geq 4$ and arbitrary channels, future research may involve the search for capacity achieving algorithms for $N \geq 4$. It is expected that such algorithms cannot be

constructed through minor modifications of CODE1 and may possibly require complete knowledge of channel statistics. If this is the case, adaptive algorithms that essentially “learn” the relevant statistics may be pursued. Suboptimal algorithms with guaranteed performance bounds in the spirit of [12] may also be of interest.

REFERENCES

- [1] T. Cover, “Broadcast channels,” *IEEE Trans. Inf. Theory*, vol. 18, no. 1, pp. 2–14, January 1972.
- [2] P. Bergmans, “Random coding theorem for broadcast channels with degraded components,” *IEEE Trans. Inf. Theory*, vol. 19, no. 2, pp. 197–207, March 1973.
- [3] A. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, “Capacity of wireless erasure networks,” *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 789–804, March 2006.
- [4] R. Ahlswede, C. Ning, S. Li, and R. Yeung, “Network information flow,” *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
- [5] L. Keller, E. Drinea, and C. Fragouli, “Online broadcasting with network coding,” in *Proc. 4th Workshop on Network Coding, Theory and Applications*, 2008.
- [6] P. Sadeghi, D. Traskov, and R. Koetter, “Adaptive network coding for broadcast channels,” in *Proc. 5th Workshop on Network Coding, Theory and Applications*, June 2009, pp. 80–86.
- [7] C. Wang, “On the capacity of wireless 1-hop intersession network coding — a broadcast packet erasure channel approach,” in *Proc. International Symposium on Information Theory (ISIT)*, June 2010, pp. 1893–1897.
- [8] P. Larsson and N. Johansson, “Multi-user ARQ,” in *Proc. Vehicular Technology Conference*, May 2006, pp. 2052–2057.
- [9] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, “XORs in the air: practical wireless network coding,” *IEEE/ACM Trans. Netw.*, vol. 16, no. 3, pp. 497–510, June 2008.
- [10] E. Rozner, A. Iyer, Y. Mehta, L. Qiu, and M. Jafry, “ER: efficient retransmission scheme for wireless LANs,” in *Proc. ACM CoNEXT*, December 2007.
- [11] L. Georgiadis and L. Tassiulas, “Broadcast erasure channel with feedback — capacity and algorithms,” in *Proc. 5th Workshop on Network Coding Theory and Applications*, June 2009, pp. 54–61.
- [12] Y. Sagduyu and A. Ephremides, “On broadcast stability of queue-based dynamic network coding over erasure channels,” *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5463–5478, December 2009.
- [13] C.-C. Wang, “Capacity of 1-to- K broadcast packet erasure channels with channel output feedback,” October 2010, presented in 48th Allerton Conference, submitted to *IEEE Trans. Inform. Theory*. [Online]. Available: <http://arxiv.org/abs/1010.2436v1>
- [14] M. Gatzianas, L. Georgiadis, and L. Tassiulas, “Multiuser broadcast erasure channel with feedback — capacity and algorithms.” [Online]. Available: http://users.auth.gr/~leonid/public/TechReports/tecreport_bec.pdf
- [15] T. Cover and J. Thomas, *Elements of information theory*, 2nd ed. John Wiley, 2006.
- [16] A. E. Gamal, “The feedback capacity of degraded broadcast channels,” *IEEE Trans. Inf. Theory*, vol. 24, no. 3, pp. 379–381, May 1978.
- [17] L. Ozarow and S. Leung-Yan-Cheong, “An achievable region and outer bound for the gaussian broadcast channel with feedback,” *IEEE Trans. Inf. Theory*, vol. 30, no. 4, pp. 667–671, July 1984.
- [18] S. Vishwanath, G. Kramer, S. Shamai, S. Jafar, and A. Goldsmith, “Capacity bounds for gaussian vector broadcast channels,” in *DIMACS Workshop on Signal Processing for Wireless Transmission*, October 2002, pp. 107–122.
- [19] R. Liu and H. Poor, “Secrecy capacity region of a multiple-antenna gaussian broadcast channel with conditional messages,” *IEEE Trans. Inf. Theory*, vol. 55, no. 3, pp. 1235–1249, March 2009.
- [20] A. Dana and B. Hassibi, “The capacity region of multiple input erasure broadcast channels,” in *Proc. International Symposium on Information Theory (ISIT)*, September 2005, pp. 2315–2319.