Optimal Exchange of Packets for Universal Recovery in Broadcast Networks

Thomas A. Courtade, *Student Member, IEEE,* Bike Xie, *Student Member, IEEE,*
and Richard D. Wesel, *Senior Member, IEEE*

***Abstract*—Consider an arbitrarily connected broadcast network of N nodes that all wish to recover k desired packets. Each node begins with a subset of the desired packets and broadcasts messages to its neighbors. For the case where nodes must transmit an integer number of packets, this paper provides necessary and sufficient conditions which characterize the set of all transmission schemes that permit universal recovery (in which every node learns all k packets). By relaxing the integer transmission constraint, this paper gives a computable lower-bound on the amount of information required to be broadcast to achieve universal recovery. Furthermore, a network-coding-based scheme (computable in polynomial time) can always achieve this lower bound if packet splitting is permitted. In this way, packet splitting can provide a significant reduction in the amount of communication required for universal recovery. For cliques with N nodes, this paper shows that splitting the packet into $N-1$ chunks allows the lower bound to be achieved with high probability.**

I. Introduction

CONSIDER an arbitrary network of N nodes that all wish to recover k desired packets. Each node begins with a subset of the desired packets and broadcasts messages to its neighbors over discrete, memoryless, and interference-free channels. Furthermore, every node knows which packets are already known by each node and knows the topology of the network. How many transmissions are required to disseminate the k packets to every node in the network? How should this be accomplished? These are the essential questions addressed.

This model has several applications to tactical networks, and we give one of them here. Consider a scenario in which an aircraft flies over a group of nodes on the ground and tries to deliver a video stream. Each ground node might only receive a subset of the packets due to interference, obstructions, and other signal integrity issues. In order to recover the transmission, the nodes are free to communicate with their neighbors, but would like to minimize the number of transmissions in order to conserve battery power (or avoid detection, etc.). How should the nodes share information, and what is the minimum number of transmissions required so that the entire network can recover the video stream?

For the special case when the network is a clique, upper and lower bounds for this problem have been recently described in [14]. Our paper provides results for the general case, describes the performance improvement facilitated by packet splitting, and derives precise results on the advantage of packet splitting for the special case of a clique.

The problem addressed in this paper is related to the index coding problem originally introduced by Birk and Kol in [4]. Specifically, generalizing the index coding problem to permit each node to be a transmitter (instead of having a single server) and further generalizing so that the network need not be a single hop network leads to a class of problems that includes our problem as a special case in which each node desires to receive all packets.

One significant result in index coding is that nonlinear index coding outperforms the best linear index code in certain cases (See [12], [2].). As discussed above, our problem is distinct from the index coding problem, and it turns out that linear encoding does achieve the minimum number of transmissions required for universal recovery and at least sometimes this solution is computable in polynomial time.

This paper is organized as follows. Section II defines the problem and introduces basic definitions and notation. Section III describes our main results: (1) necessary and sufficient conditions describing all transmission strategies allowing universal recovery, and (2) achievability of a (polynomial-time) computable lower bound using packet-splitting. These results for general networks are specialized and strengthened for the special case when the network is a clique. Section IV sketches the proofs for these results. Section V delivers the conclusions and discusses directions for future work.

II. System Model

This paper considers an arbitrary network $\mathcal{T}$ of N nodes. The network must be connected, but it need not be fully connected (a clique). A graph $\mathcal{G}_{\mathcal{T}} = (V, E)$ describes the specific connections in the network, where V is the set of vertices $\{v_i : i \in \{1, \ldots, N\}\}$ (each corresponding to a node) and E is the set of edges connecting nodes. We assume that the edges in E are undirected, but our work easily extends to directed graphs. We frequently use the notation $[m]$ to denote the set $\{1, \ldots, m\}$ throughout.

Each node wishes to recover the same k desired packets, and each node begins with a (possibly empty) subset of the desired packets. $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$ is the set of packets originally available at node i, and $\{\mathcal{P}_i\}_{i=1}^N$ satisfies $\bigcup_{i=1}^N \mathcal{P}_i = \{p_1, \ldots, p_k\}$. Each $p_j \in \mathbb{F}$, where $\mathbb{F}$ is some finite field (e.g. $\mathbb{F} = \mathrm{GF}(2^m)$).

The authors are with the Electrical Engineering Department, University of California, Los Angeles, CA, 90095 USA (email: tacourta@ee.ucla.edu; wesel@ee.ucla.edu).

This research was supported by Rockwell Collins through contract #4502769987.

Let the set $\Gamma(i)$ be the neighbors of node i. There exists an edge $e \in E$ connecting two vertices $v_i, v_j \in V$ iff $i \in \Gamma(j)$. For convenience, $i \in \Gamma(i)$. Node i broadcasts messages to its neighbors $\Gamma(i)$ over discrete, memoryless, and interference-free channels. If $\mathcal{S}$ is a set of nodes, then we define $\Gamma(\mathcal{S}) = \cup_{i\in\mathcal{S}}\Gamma(i)$.

This paper seeks to determine the minimum amount of communication required to achieve universal recovery of the k packets. An important consideration is whether packets are considered indivisible (so that the smallest unit of transmission is a packet) or packets may be split into chunks so that a fraction of a packet may be transmitted. We will consider both cases and ultimately show the benefit of packet splitting.

A. Indivisible Packets

When packets are deemed indivisible, a single *transmission* by user i consists of sending a packet (some $z \in \mathbb{F}$) to all nodes $j \in \Gamma(i)$. Following [2], let $B(\mathcal{T})$ be the minimum required number of such transmissions that allow universal recovery. Throughout this paper, we assume error-free broadcast channels and orthogonal multiple-access channels (i.e. there is no interference from simultaneous transmissions). Thus, the task of determining $B(\mathcal{T})$ is combinatorial.

Example 1 (Line Network): Suppose $\mathcal{T}$ is a network of nodes connected along a line as follows: $V = \{v_1, v_2, v_3\}$, $E = \{(v_1, v_2), (v_2, v_3)\}$, $\mathcal{P}_1 = \{p_1\}$, $\mathcal{P}_2 = \emptyset$, and $\mathcal{P}_3 = \{p_2\}$. Note that each node must transmit at least once in order for all nodes to recover $\{p_1, p_2\}$, hence $B(\mathcal{T}) \geq 3$. Suppose node 1 transmits p_1 and node 3 transmits p_2. Then (upon receipt of p_1 and p_2 from nodes 1 and 3) node 2 transmits $p_1 \oplus p_2$, where $\oplus$ indicates addition in the finite field $\mathbb{F}$. This strategy requires 3 transmissions and allows each user to recover $\{p_1, p_2\}$. Hence $B(\mathcal{T}) = 3$.

Example 1 demonstrates a transmission strategy that uses two *rounds* of communication. The broadcasts by node i in a particular round of communication can depend only on the information available to node i prior to that round (i.e. $\mathcal{P}_i$ and previously received transmissions from neighboring nodes). In other words, the broadcasts are causal.

Example 2 (Clique): Suppose $\mathcal{T}$ is a network of nodes that are fully connected so as to form a clique as follows: $\mathcal{P}_i = \{p_1, p_2, p_3\}\backslash p_i$, and $\mathcal{G}_\mathcal{T}$ is a clique of size 3. Clearly one transmission is not sufficient, thus $B(\mathcal{T}) \geq 2$. It can be seen that two transmissions suffice: let node 1 transmit p_2 which lets node 2 have $\mathcal{P}_2 \cup p_2 = \{p_1, p_2, p_3\}$. Now, node 2 transmits $p_1 \oplus p_3$, allowing nodes 1 and 3 to each recover all three packets. Thus $B(\mathcal{T}) = 2$. Since each transmission was only a function of the packets originally available at the corresponding node, this transmission strategy can be accomplished in a single round of communication.

B. Divisible Packets

Suppose that we can split each packet into t equally sized chunks (i.e. $p_i \in \mathbb{F}^t$). The resulting problem has the same general form as the original problem except that the original k packets are replaced by kt smaller packets (chunks) and the t chunks that form one of the original packets are either all present or all absent from any particular node. Using the notation of [2], we denote the new network of kt packets as the product of t instances of the original network $\mathcal{T} \times \cdots \times \mathcal{T}$ (t times) or simply $\mathcal{T}^t$.

Note that $B(\mathcal{T}^t)$ for the network $\mathcal{T}^t$ created by splitting each packet into t chunks will be the number of "chunk" transmissions. Each chunk transmission has length only $1/t$ of an original packet transmission. To normalize so that we can properly evaluate the benefit of packet splitting, define $B_t(\mathcal{T}) = \frac{B(\mathcal{T}^t)}{t}$ (as in [2]).

Example 3 (Clique with Packet Splitting): Let $\mathcal{T}$ be the network of Example 2 and consider the network $\mathcal{T}^2 = \mathcal{T} \times \mathcal{T}$. Using superscripts to identify the chunks of a given packet, we can express each $\mathcal{P}_i$ as follows:

$$\mathcal{P}_i = \{p_1^{(1)}, p_1^{(2)} p_2^{(1)}, p_2^{(2)}, p_3^{(1)}, p_3^{(2)}\}\backslash\{p_i^{(1)}, p_i^{(2)}\}.$$

Three chunk transmissions allow universal recovery as follows: Node 1 transmits $p_2^{(2)} \oplus p_3^{(2)}$. Node 2 transmits $p_1^{(1)} \oplus p_3^{(1)}$. Node 3 transmits $p_1^{(2)} \oplus p_2^{(1)}$. Thus, $B(\mathcal{T}^2) = 3$ and $B_2(\mathcal{T}) = 1.5$. Hence, if we were allowed to split the packets of Example 2 into two symbols each (e.g. $p_i \rightarrow (p_i^{(1)}, p_i^{(2)})$), it would suffice to transmit 3 chunks (i.e. 1.5 packets instead of 2 packets).

Note that the benefit of packet splitting is not monotonic. This can be observed by considering $B_3(\mathcal{T})$ for the network $\mathcal{T}$ defined in Example 2.

III. Main Results

In this section, we state our main results. In subsection III-A, we give results that are applicable to an arbitrary broadcast network. In subsection III-B, we give several results that are specialized to clique networks.

A. Arbitrarily Connected Networks

Let b_i^j be the number of transmissions from node i during round j. In this way, the total number of packet transmissions summing over all rounds is $\sum_{i=1}^{N}\sum_{j=1}^{r} b_i^j$. Also, let $\{b_i^j\}$ denote the set of b_i^j values for $i \in [N]$ and $j \in [r]$.

Define the region $\mathcal{R}_r \subset \mathbb{Z}_+^{N\times r}$ as follows:

$\{b_i^j\} \in \mathcal{R}_r$ if and only if:

$\forall\ \emptyset \subsetneq \mathcal{S}_0 \subseteq \cdots \subseteq \mathcal{S}_r \subsetneq [N]$ satisfying $\mathcal{S}_j \subseteq \Gamma(\mathcal{S}_{j-1})$

for each $j \in [r]$, the following inequalities hold :

$$\sum_{j=1}^{r} \sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left|\bigcap_{i\in\mathcal{S}_r} \mathcal{P}_i^c\right|. \tag{1}$$

Theorem 1: For a fixed number of communication rounds r, a transmission strategy in which node i makes exactly b_i^j transmissions during the j^{th} round of communication permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Thus, the family of inequalities defining $\mathcal{R}_r$ are necessary and sufficient for universal recovery. These inequalities utilize sequences of sets of the form $\emptyset \subsetneq \mathcal{S}_0 \subseteq \cdots \subseteq \mathcal{S}_r \subsetneq [N]$ satisfying $\mathcal{S}_j \subseteq \Gamma(\mathcal{S}_{j-1})$ for each $j \in [r]$. The inequalities

can be thought of as "generalized cutset bounds" which govern the information flow in the multihop broadcast networks under consideration.

Remark 1: As the proof of Theorem 1 (given in Section IV) demonstrates, for any feasible vector $\{b_i^j\}$ (i.e. $\{b_i^j\} \in \mathcal{R}_r$), a corresponding transmission strategy in which user i makes exactly b_i^j transmissions during round j can be computed in polynomial time using the algorithm described in [15]. Therefore, the difficulty in solving the broadcast problem lies solely in finding the optimal solution to the ILP defined by minimizing $\sum_{i=1}^{N}\sum_{j=1}^{r} b_i^j$ subject to $\{b_i^j\} \in \mathcal{R}_r$.

We give two brief examples to demonstrate the utility of Theorem 1.

Example 4 (Reduction to Cutset Bounds): In their most general form, the constraints defining $\mathcal{R}_r$ may be difficult to interpret, however they reduce to the familiar cutset bounds when $\mathcal{S}_0 = \cdots = \mathcal{S}_r = \mathcal{S}$ for some nonempty $\mathcal{S} \subsetneq [N]$:

$$\sum_{i \in \mathcal{S}^c \cap \Gamma(\mathcal{S})} \sum_{j=1}^{r} b_i^j \geq \left| \bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c \right|.$$

In words, the total number of packets that flow into the set of nodes $\mathcal{S}$ must be greater than or equal to the number of packets that no node in $\mathcal{S}$ has.

Example 5 (Line Network Revisited): In the 3-node line network of Example 1, the reader can verify that $\mathcal{R}_2$ reduces to:

$$b_1^1 \geq 1, \quad b_3^1 \geq 1, \quad b_2^2 \geq 1$$

after eliminating redundant constraints.

From this example (in particular since $b_2^2 \geq 1$), it is clear that multiple rounds may be necessary. The delay created by information propagating through a multihop network is captured by the constraints defining $\mathcal{R}_r$.

While the above examples are relatively simple, some more interesting applications of this theorem to clique networks are given in subsection III-B.

Although Theorem 1 defines an ILP which can be solved to obtain $B(\mathcal{T})$, solving an ILP is NP-hard in general. Thus, it is only practical to compute $B(\mathcal{T})$ for relatively small networks. However, a lower bound for $B(\mathcal{T})$ can be efficiently computed by solving a corresponding Linear Program (LP) relaxation. In particular, let $\mathcal{R}_r^{LP} \subset \mathbb{R}^{N \times r}$ be defined in exactly the same way as $\mathcal{R}_r$ with the integer constraint on b_i^j replaced by $b_i^j \geq 0$. This yields the following LP[1]:

$$\min_{\left\{\{b_i^j\} : \{b_i^j\} \in \mathcal{R}_r^{LP}\right\}} \sum_{i=1}^{N}\sum_{j=1}^{r} b_i^j. \tag{2}$$

Define $B_{LP}(\mathcal{T})$ to be the optimal value of LP (2) for a network $\mathcal{T}$ and r sufficiently large. Because the LP is a relaxation of the original ILP, $B_{LP}(\mathcal{T}) \leq B(\mathcal{T})$. However, it turns out that $B_{LP}(\mathcal{T})$ is achievable with packet splitting as stated in the following theorem:

Theorem 2: For any network $\mathcal{T}$, we have:

$$\lim_{t \to \infty} B_t(\mathcal{T}) = B_{LP}(\mathcal{T}).$$

Moreover, there exists a finite $t_0 \in \mathbb{N}$ such that $B_{t_0}(\mathcal{T}) = B_{LP}(\mathcal{T})$.

In other words, Theorem 2 states that one can achieve the LP lower bound (which is efficiently computable) by packet-splitting (i.e. splitting the packets into smaller sub-packets). This was observed in Example 3, where $B_2(\mathcal{T}) = B_{LP}(\mathcal{T}) = 1.5$ and $B(\mathcal{T}) = 2$.

A natural question is whether or not t_0 can be determined for specific networks of interest. This is explored for clique networks in the following subsection.

B. Clique Networks

Theorems 1 and 2 can be applied to specific networks or topologies of interest to obtain results that can be practical and powerful. In this paper, due to space limitations, we focus our attention on the special case of clique networks as an important example. Extensions to a much broader class of networks can be found in the follow-on paper [16] by the present authors.

Theorem 3: For a given network $\mathcal{T}$, if $\mathcal{G}_{\mathcal{T}}$ is a clique of size N, then $B(\mathcal{T})$ is the optimal value of the following Integer Linear Program (ILP):

$$\begin{aligned} \textbf{minimize} \quad & \sum_{i=1}^{N} b_i \\ \text{subject to:} \quad & \sum_{i \in \mathcal{S}} b_i \geq \left| \bigcap_{i \in \mathcal{S}^c} \mathcal{P}_i^c \right|, \ \forall \emptyset \subsetneq \mathcal{S} \subsetneq [N]. \\ & (b_i \text{ is integral.}) \end{aligned} \tag{3}$$

Moreover, if $\{b_i^*\}$ is an optimal solution to the above ILP, then there exists an optimal transmission strategy in which each node i makes exactly b_i^* transmissions.

The constraints in the ILP of Theorem 3 are equivalent to $\{b_i\} \in \mathcal{R}_1$ (i.e., the traditional cutset bounds), and hence one round of communication is sufficient for achieving the minimum number of transmissions when the broadcast network is a clique. This implies that node i only needs to transmit functions of $\mathcal{P}_i$. Adaptation based on previous transmissions is not necessary.

In many real-world scenarios, packets are available at each node according to some random process. For the following theorem, consider a model in which packets are distributed randomly in a network. Specifically, let each packet be available at each node (independently) with probability q.

In this context universal recovery means that all nodes recover all of the packets that at least one node had originally. In other words, we disregard the packets that aren't available at any node. This induces a probability measure on the set of possible networks where $k = |\cup_{i=1}^{N} \mathcal{P}_i|$ is fixed. In the sequel, probabilities are with respect to this measure.

[1] As posed, LP (2) has an exponential number of constraints. However, it can be recast in terms of N different max-flow problems, each with a polynomial number of constraints and variables. Thus, solving LP (2) has polynomial complexity. This argument is omitted due to space constraints.

Theorem 4: With probability arbitrarily close to 1 as $k \to \infty$, the following relationship holds:

$$B_{N-1}(\mathcal{T}) = B_{LP}(\mathcal{T}) = \frac{1}{N-1}\sum_{i=1}^{N}|\mathcal{P}_i^c|,$$

and an optimal $\{b_i^*\}$ is:

$$b_i^* = \frac{1}{N-1}\sum_{j=1}^{N}|\mathcal{P}_j^c| - |\mathcal{P}_i^c| \quad \text{for } i \in [N].$$

In other words, splitting the packets into at most $N-1$ chunks will achieve the LP lower bound with high probability. Moreover, with high probability, we have the remarkable result that $\sum_{j\neq i} b_j^* = |\mathcal{P}_i^c|$ for all i (i.e., each node receives *exactly* the number of packets it needed, and no more).

Theorem 4 is stronger than Theorem 2 in the sense that it provides an exact value of t_0. However its strength is weakened because it holds with high probability instead of with probability one. Below, we conjecture that further packet-splitting is never required in a clique network.

Conjecture 1: Splitting packets into at most $N-1$ chunks is sufficient to achieve the LP lower bound on the number of broadcasts required in a clique network.

Remark 2: Clearly the conjecture holds for $N = 2$. It can be verified (somewhat tediously) that the conjecture also holds for $N = 3$ by manipulating the constraints of Theorem 3. Monte Carlo experiments for $N \leq 12$ and a variety of k values support this conjecture.

IV. Proofs

This section contains the proofs of the results presented in section III.

A. Arbitrary Networks

Proof of Theorem 1 : Due to space limitations, the complete proof of Theorem 1 has been omitted and can be found in the follow-on paper [16] by the present authors. This proof is accomplished by reducing the problem at hand to an instance of a single-source network coding problem and invoking the Max-Flow Min-Cut Theorem for network information flow [6]. Then, it can be shown that it is both necessary and sufficient to consider cuts which are characterized by the set of inequalities described by (1). ∎

Definition 1: A rational LP is a linear program in which the constraint coefficients and cost vector coefficients are rational numbers.

Lemma 1: If the triple (A, b, c) defines a rational LP that is feasible and bounded from below, then there exists a vector x^* with rational coefficients which minimizes $c^T x$ subject to $Ax \leq b$.

Proof: If the LP is feasible and bounded from below, then the optimum is achieved at some extreme point x^* which can be obtained by solving a rational system of equations, and hence is rational itself. ∎

Definition 2: A sequence $\{a_n\}_{n=1}^{\infty}$ is subadditive if $a_{m+n} \leq a_m + a_n$ for all $m, n \in \mathbb{N}$.

Lemma 2: $B_t(\mathcal{T})$ is subadditive.

Proof: Let $t = t_1 + t_2$. Note that $B(\mathcal{T}^t) \leq B(\mathcal{T}^{t_1}) + B(\mathcal{T}^{t_2})$ since $B(\mathcal{T}^{t_1}) + B(\mathcal{T}^{t_2})$ broadcasts are achievable for the network $\mathcal{T}^t$ by concatenating the optimal broadcast strategies for networks $\mathcal{T}^{t_1}$ and $\mathcal{T}^{t_2}$. Thus: $B(\mathcal{T}^t) = \frac{B(\mathcal{T}^t)}{t} \leq \frac{B(\mathcal{T}^{t_1})+B(\mathcal{T}^{t_2})}{t_1+t_2} \leq B_{t_1}(\mathcal{T}) + B_{t_2}(\mathcal{T})$. ∎

Lemma 3 (Fekete's Lemma [17]): If $\{a_n\}_{n=1}^{\infty}$ is a subadditive sequence of nonnegative terms, then $\lim_{n\to\infty}\frac{a_n}{n}$ exists and is equal to $\inf\{\frac{a_n}{n} : n \in \mathbb{N}\}$.

Proof of Theorem 2 : By Lemma 1, there is a set of rational numbers $\{b_i^{*j}\}_{i,j}$ which minimizes LP (2). Express each b_i^{*j} in its rational form as $b_i^{*j} = p_{i,j}/q_{i,j}$, where $p_{i,j}, q_{i,j} \in \mathbb{N}$. Let $t_0 = \text{LCM}(\{q_{i,j}\}_{i,j})$. Then, each $b_i^{*j} = p'_{i,j}/t_0$, where $p'_{i,j} \in \mathbb{N}$. If we consider the network $\mathcal{T}^{t_0}$, then $|\cap_{i\in\mathcal{S}}\mathcal{P}_i^c| \mapsto t_0|\cap_{i\in\mathcal{S}}\mathcal{P}_i^c|$ and defines the corresponding region $\mathcal{R}_r^{t_0}$ for $\mathcal{T}^{t_0}$. It follows that the set of integers $\{p'_{i,j}\}_{i,j}$ minimizes the following ILP, which, by definition, has optimal value equal to $B_{t_0}(\mathcal{T})$:

$$\min_{\left\{\{b_i^j\}:\{b_i^j\}\in\mathcal{R}_r^{t_0}\right\}} \frac{1}{t_0}\sum_{i=1}^{N}\sum_{j=1}^{r} b_i^j.$$

Thus, $B_{LP}(\mathcal{T}) = B_{t_0}(\mathcal{T})$ for some finite $t_0 \in \mathbb{N}$. By Lemmas 2 and 3, this establishes that $\lim_{t\to\infty} B_t(\mathcal{T}) = B_{LP}(\mathcal{T})$. ∎

B. Clique Networks

Proof of Theorem 3: In the case where the broadcast network is a clique of size N, we have that $\mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1}) = \mathcal{S}_j^c$ for any $\emptyset \subsetneq \mathcal{S}_{j-1} \subseteq \mathcal{S}_j \subsetneq [N]$. Therefore, the constraints defining $\mathcal{R}_r$ in Theorem 1 become:

$$\sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c} b_i^{r+1-j} \geq \left|\bigcap_{i\in\mathcal{S}_r}\mathcal{P}_i^c\right|. \tag{4}$$

Now, suppose some $\{b_i^j\}_{i,j}$ satisfy the conditions of (4) for any appropriate increasing sequence of sets $\{\mathcal{S}_j\}_{j=1}^r$, and consider a modified set $\{\tilde{b}_i^j\}_{i,j}$ defined as: $\tilde{b}_i^r = \sum_{j=1}^r b_i^j$ and $\tilde{b}_i^j = 0$ for $j \in [r-1]$. By construction, $\mathcal{S}_{j+1}^c \subseteq \mathcal{S}_j^c$, therefore, using the definition of $\{\tilde{b}_i^j\}_{i,j}$, we have:

$$\sum_{i\in\mathcal{S}_1^c}\tilde{b}_i^r \geq \sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c} b_i^{r+1-j} \geq \left|\bigcap_{i\in\mathcal{S}_r}\mathcal{P}_i^c\right|.$$

Since $\left|\bigcap_{i\in\mathcal{S}_1}\mathcal{P}_i^c\right| \geq \left|\bigcap_{i\in\mathcal{S}_r}\mathcal{P}_i^c\right|$, when the broadcast network is a clique of size N, it is sufficient to consider constraints of the form:

$$\sum_{i\in\mathcal{S}^c} b_i \geq \left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right|, \quad \forall\emptyset \subsetneq \mathcal{S} \subsetneq [N]. \tag{5}$$

∎

In the model where packets are randomly distributed, quantities such as $|\mathcal{P}_i|$ and $|\cup_{i\in\mathcal{S}}\mathcal{P}_i|$ are random variables. We continue to adopt the notation that $|\cup_{i=1}^N\mathcal{P}_i| = k$. Consider a

single packet p_j. The following observation will be useful in what follows:

$$\Pr\left(p_j \in \bigcap_{i\in\mathcal{S}} \mathcal{P}_i^c\right) = \frac{(1-q)^{|\mathcal{S}|} - (1-q)^N}{1-(1-q)^N}. \quad (6)$$

This observation follows by conditioning on the fact that packets not available to any node are excluded. The following lemma will also be useful:

Lemma 4: If $0 < q < 1$ is fixed, then there exists some $\delta > 0$ such that the following inequality holds for all $\ell \in \{2, \ldots, N-1\}$:

$$\frac{N-\ell}{N-1} \geq \frac{(1-q)^\ell - (1-q)^N}{1-q-(1-q)^N} + \delta.$$

Proof: We write $\ell = \theta \cdot 1 + (1-\theta)\cdot N$, where $\theta = \frac{N-\ell}{N-1}$. Since $(1-q)^x$ is strictly convex in x, by Jensen's inequality we have:

$$\frac{(1-q)^\ell - (1-q)^N}{1-q-(1-q)^N} < \frac{N-\ell}{N-1}.$$

Taking δ to be the minimum gap in the above inequality for the values $\ell \in \{2, \ldots, N-1\}$ completes the proof. ∎

Proof of Theorem 4 : The proof of the theorem is carried out in two steps: verifying feasibility of the solution and verifying optimality of the solution. This is accomplished in Lemmas 5 and 6. ∎

First, consider the following set of inequalites which are a subset of the inequalities described by (5):

$$\sum_{\substack{i=1\\ i\neq j}}^{N} b_i \geq |\mathcal{P}_i^c| \quad \text{for } 1 \leq j \leq N. \quad (7)$$

These form a set of N linearly independent inequalities, hence the vector $\{\tilde{b}_i\}_{i=1}^N$ which satisfies each inequality in (7) with equality is uniquely defined. In fact, some algebra reveals that each $\tilde{b}_j$ can be expressed as:

$$\tilde{b}_j = \frac{1}{N-1}\sum_{i=1}^{N} |\mathcal{P}_i^c| - |\mathcal{P}_j^c|. \quad (8)$$

Lemma 5 (Feasibility of $\{\tilde{b}_i\}_{i=1}^N$): As $k \to \infty$, $\{\tilde{b}_i\}_{i=1}^N$ is feasible with probability approaching 1.

Proof: We must verify that:

$$\sum_{i\in\mathcal{S}^c} \tilde{b}_i \geq \left|\bigcap_{i\in\mathcal{S}} \mathcal{P}_i^c\right| \quad (9)$$

holds with high probability for all subsets $\mathcal{S} \subseteq [N]$ satisfying $2 \leq |\mathcal{S}| \leq N-1$ (the case $|\mathcal{S}| = 1$ is satisfied by the definition of $\{\tilde{b}_i\}_{i=1}^N$). Substitution of (8) into (9) along with some algebra yields that the following equivalent conditions must hold:

$$\left(\frac{N-|\mathcal{S}|}{N-1}\right)\sum_{i=1}^{N}\frac{1}{k}|\mathcal{P}_i^c| - \sum_{i\in\mathcal{S}^c}\frac{1}{k}|\mathcal{P}_i^c| \geq \frac{1}{k}\left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right|$$

Now, note that for any $\mathcal{S}$, $\left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right|$ is a random variable which can be expressed as $\left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right| = \sum_{j=1}^k X_j^{\mathcal{S}}$, where $X_j^{\mathcal{S}}$ is an indicator random variable taking the value 1 if $p_j \in \bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c$ and 0 otherwise. From (6) we have:

$$\Pr\left(X_j^{\mathcal{S}} = 1\right) = \frac{(1-q)^{|\mathcal{S}|} - (1-q)^N}{1-(1-q)^N}.$$

By the Law of Large Numbers, for any $\eta > 0$:

$$\Pr\left(\left|\frac{1}{k}\left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right| - \frac{(1-q)^{|\mathcal{S}|}-(1-q)^N}{1-(1-q)^N}\right| > \eta\right) < \epsilon_k,$$

where $\epsilon_k \to 0$ as $k \to \infty$. Thus, by the union bound, Lemma 4, and taking η sufficiently small, the following string of inequalities holds with arbitrarily high probability as $k \to \infty$:

$$\begin{aligned}
&\left(\frac{N-|\mathcal{S}|}{N-1}\right)\sum_{i=1}^{N}\frac{1}{k}|\mathcal{P}_i^c| - \sum_{i\in\mathcal{S}^c}\frac{1}{k}|\mathcal{P}_i^c| \\
&\geq \left(\frac{N-|\mathcal{S}|}{N-1}\right)\left(\frac{(1-q)-(1-q)^N}{1-(1-q)^N} - (2N-1)\eta\right) \\
&\geq \frac{(1-q)^{\mathcal{S}}-(1-q)^N}{1-(1-q)^N} + \eta \\
&\geq \frac{1}{k}\left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right|.
\end{aligned}$$

∎

Lemma 6 (Optimality of $\{\tilde{b}_i\}_{i=1}^N$): If $\mathcal{N}$ is a clique, then the following inequality holds:

$$B_{LP}(\mathcal{N}) \geq \frac{1}{N-1}\sum_{i=1}^{N}|\mathcal{P}_i^c| = \sum_{i=1}^{N}\tilde{b}_i. \quad (10)$$

Proof: The equality in (10) follows by definition of $\{\tilde{b}_i\}_{i=1}^N$.

Consider the dual of the LP relaxation of (3). The dual LP can be expressed as:

$$\begin{aligned}
\textbf{maximize} \quad & \sum_{\mathcal{S}} \lambda_{\mathcal{S}} \left|\bigcap_{i\in\mathcal{S}}\mathcal{P}_i^c\right| \quad (11)\\
\text{subject to:} \quad & \\
& \sum_{\mathcal{S}} \lambda_{\mathcal{S}} I(i \notin \mathcal{S}) \leq 1, \quad \forall i \in [N] \\
& \lambda_{\mathcal{S}} \geq 0.
\end{aligned}$$

In the above formulation, there is a dual variable $\lambda_{\mathcal{S}}$ for each nonempty subset $\mathcal{S} \subsetneq [N]$, and $I(\cdot)$ is the usual indicator function.

Consider the following choice of dual variables:

$$\lambda_{\mathcal{S}} = \begin{cases} \frac{1}{N-1} & \text{if } \mathcal{S} = \{j\} \text{ for some } j \\ 0 & \text{otherwise.} \end{cases}$$

A quick check shows that $\sum_{\mathcal{S}} \lambda_{\mathcal{S}} I(i \notin \mathcal{S}) = 1$ for all i, hence this choice of dual variables corresponds to a dual

feasible point. Further, the objective function evaluated at this point is:

$$\sum_{\mathcal{S}} \lambda_{\mathcal{S}} |\bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c| = \frac{1}{N-1} \sum_{i=1}^{N} |\mathcal{P}_i^c|.$$

Since the dual objective function evaluated at any feasible point provides a lower bound on the optimal value in the primal problem, the lemma is proved.

■

V. Conclusions

This paper studies an arbitrarily connected broadcast network of N nodes that all wish to recover k desired packets originally dispersed among the nodes. For the case where nodes must transmit an integer number of packets, we derive necessary and sufficient conditions which characterize the set of transmission strategies permitting universal recovery.

Relaxing the integer transmission constraint provides a computable lower-bound on the amount of information required to be broadcast to achieve universal recovery. Furthermore, a network-coding-based scheme (computable in polynomial time) can always achieve this lower bound if packet splitting is permitted. In this way, packet splitting can provide a significant reduction in the amount of communication required for universal recovery.

For cliques with N nodes, this paper shows that splitting the packet into $N-1$ chunks allows the lower bound to be achieved with high probability. In fact, this high probability result includes explicit construction for an optimal transmission strategy. This solution is completely efficient in the sense that every node receives an amount of information that exactly equals its number of missing packets.

Acknowledgment

This research was supported by Rockwell Collins through contract #4502769987. The authors would like to thank Kent Benson of Rockwell Collins for proposing the investigation that led to this paper. Conversations with Professor Benny Sudakov led us in the direction of Theorem 4. We also enjoyed numerous useful conversations with Pavan Datta and Chi-Wen Su.

References

[1] S. Boyd and L. Vandenberghe, Convex Optimization. Cambridge University Press, 2004.

[2] N. Alon, A. Hassidim, E. Lubetzky, U. Stav and A. Weinstein, Broadcasting with side information, Proc. of the 49th IEEE FOCS (2008), 823-832.

[3] Z. Bar-Yossef, Y. Birk, T.S. Jayram and T. Kol, Index coding with side information, Proc. of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 197-206.

[4] Y. Birk and T. Kol, Coding-on-demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients, IEEE Transactions on Information Theory 52 (2006), 2825-2830.

[5] H. S. Witsenhausen, The zero-error side information problem and chromatic numbers, IEEE Transactions on Information Theory, 22(5) (1976), 592-593.

[6] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung, "Network information flow", IEEE Transactions on Information Theory, July 2000.

[7] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding", IEEE Transactions on Information Theory, Feb. 2003.

[8] C. Fragouli, J. Widmer, and J.-Y. L. Boudec, "A network coding approach to energy efficient broadcasting: From theory to practice", in IEEE INFOCOM, Barcelona, Spain, Apr. 2006.

[9] C. Fragouli, J. Widmer, and J.-Y. L. Boudec, "Efficient Broadcasting Using Network Coding," IEEE/ACM Transactions on Networking, Vol. 16, No. 2, April 2008, 450-463.

[10] Y. Wu, J. Padhye, R. Chandra, V. Padmanabhan, and P. A. Chou. The local mixing problem. In Proc. Information Theory and Applications Workshop, San Diego, Feb. 2006.

[11] S. El Rouayheb, M.A.R. Chaudhry, and A. Sprintson. On the minimum number of transmissions in single-hop wireless coding networks. In IEEE Information Theory Workshop (Lake Tahoe), 2007.

[12] E. Lubetzky and U. Stav. Non-linear index coding outperforming the linear optimum. In Proc. of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 161167, 2007.

[13] S. El Rouayheb, A. Sprintson, and C. N. Georghiades. On the relation between the index coding and the network coding problems. Proc. of IEEE International Symposium on Information Theory (ISIT08), 2008.

[14] S El Rouayheb, A Sprintson, P Sadeghi, "On Coding for Cooperative Data Exchange" - Arxiv preprint arXiv:1002.1465, 2010.

[15] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial Time Algorithms for Multicast Network Code Construction, IEEE Transactions on Information Theory, vol. 51, no. 6, pp. 19731982, 2005.

[16] T. A. Courtade and R. D. Wesel, "On the Minimum Number of Transmissions Required for Universal Recovery in Broadcast Networks," Forty-Eighth Annual Allerton Conference on Communication, Control, and Computing: Monticello, IL, Sept. Oct. 1, 2010.

[17] Fekete, M. "Uber die Verteilung der Wurzeln bei gewissen algebraischen Gleichungen mit. ganzzahligen Koeffizienten." Mathematische Zeitschrift 17 (1923), pp. 228-249.