Efficient Universal Recovery in Broadcast Networks

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

***Abstract*—Consider a 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. In a previous paper we established necessary and sufficient conditions on the number of transmissions from each node required for universal recovery (in which each node recovers all k packets). However, these conditions are numerous and cumbersome. The present paper gives a series of relatively simple conditions for universal recovery that apply when the number of packets is large and the distribution of packets among the nodes is well behaved.**

Our first result, which applies to any fixed network topology, uses only simple cuts in the network to characterize a set of transmission strategies such that for any $\epsilon > 0$ these strategies require at most $k\epsilon$ transmissions above the minimum required for universal recovery. For certain topologies including nonsingular d-regular d-connected networks, we explicitly construct transmission strategies that achieve universal recovery while using at most N transmissions above the minimum even when the total number of required transmissions is very large. These explicit constructions essentially resolve the problem completely for many canonical networks (e.g. cliques, rings, grids on tori, etc.).

I. Introduction

CONSIDER a connected 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 the topology of the network and which packets each node already knows. The network seeks universal recovery (i.e. to disseminate the k packets to every node in the network). This paper identifies the minimum number of transmissions required for universal recovery and demonstrates how to efficiently design a transmission strategy and communication schedule that closely approach this minimum.

For the special case when the network is a clique, upper and lower bounds for this problem have been recently described in [11]. This paper provides feasibility results and asymptotic bounds on the minimum number of transmissions required for the general case. It also provides bounds (within a constant number of transmissions) on the minimum number of transmissions required for universal recovery in a general class of networks which includes many frequently considered canonical networks. In doing so, it significantly extends the recent work by the present authors [1].

The problem addressed in this paper is related to the index coding problem originally introduced by Birk and Kol in [3]. 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.

In general, nonlinear coding is required to minimize the required transmissions for the index coding problem. (See [2], [9].) As discussed above, the universal recovery 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.

This paper is organized as follows. Section II states the problem and introduces basic definitions and notation. Section III describes our three main results: (1) necessary and sufficient conditions describing all transmission strategies permitting universal recovery, (2) the characterization (in terms of simple cuts in the network) of a set of transmission strategies requiring at most $k\epsilon$ transmissions above the minimum number required for universal recovery, and (3) explicit transmission strategies for certain topologies (including many canonical networks) that achieve universal recovery while using at most N transmissions above the minimum number required. Section IV delivers the conclusions. The Appendix sketches the proofs of the results.

II. System Model

This paper considers a 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. V is the set of vertices $\{v_i : i \in \{1, \ldots, N\}\}$, each corresponding to a node. E is the set of edges connecting nodes. Although we assume that the edges in E are undirected, our work easily extends to directed graphs.

This paper seeks to determine the minimum amount of communication required to achieve universal recovery of k desired packets. Each node wishes to recover the same k desired packets, and each node begins with a (possibly empty) subset of the desired packets. Specifically, node i begins with the set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$. $\{\mathcal{P}_i\}_{i=1}^N$ satisfies $\bigcup_{i=1}^N \mathcal{P}_i = \{p_1, \ldots, p_k\}$. Each $p_j \in \mathbb{F}$, where $\mathbb{F}$ is some sufficiently large finite field (e.g. $\mathbb{F} = \text{GF}(2^m)$).

This paper considers sequences of networks indexed by the total number of packets k. In such a sequence, the network topology is fixed, but the distribution of packets depends on k. In order to emphasize this dependence on k, we will often refer to a sequence of packet distributions indexed by k as $\{\mathcal{P}_i(k)\}_{i=1}^N$. Naturally, these sets must satisfy $\left|\bigcup_{i=1}^N \mathcal{P}_i(k)\right| = k$.

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).

We say that a sequence of packet distributions is *well-behaved* if the limit of the packet distribution

$$\mathbb{P}_{\mathcal{S}} \triangleq \lim_{k\to\infty} \frac{1}{k} \left| \bigcup_{i\in\mathcal{S}} \mathcal{P}_i(k) \right| \tag{1}$$

exists for all subsets $\mathcal{S} \subseteq \{1,\ldots,N\}$. $\mathbb{P}_{\mathcal{S}}$ is the limit of the empirical probability that any node in $\mathcal{S}$ receives a particular packet. Let $\mathbb{P}^c_{\mathcal{S}} = 1 - \mathbb{P}_{\mathcal{S}}$. The condition (1) can be replaced with convergence in probability, as discussed in subsection III-D.

Let the set $\Gamma(i)$ be the neighbors of node i. For convenience, $i \in \Gamma(i)$. There exists an edge $e \in E$ connecting two vertices $v_i, v_j \in V$ iff $i \in \Gamma(j)$. 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 $\Gamma(\mathcal{S}) = \cup_{i\in\mathcal{S}}\Gamma(i)$. The neighbors of a set $\mathcal{S}$ not included in $\mathcal{S}$ itself (i.e., the boundary of $\mathcal{S}$) are denoted $\partial(\mathcal{S}) = \Gamma(\mathcal{S}) \cap \mathcal{S}^c$.

Throughout this paper, we adopt the conventional notation $[k] \triangleq \{1,\ldots,k\}$. An indexed set (or vector) $\{x_1,\ldots,x_k\}$ is referred to by the shorthand notation $\{x_i\}_{i=1}^k$. When the range of the index is apparent from context, we will sometimes abbreviate it further as $\{x_i\}$.

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. The implications of divisible packets were considered previously in [1]. In this paper, we maintain a combinatorial approach throughout and assume that transmissions must consist of an integer number of packets unless otherwise stated. However, divisible packets can be modeled by choosing an appropriate sequence of well-behaved packet distributions in which packets are replicated to produce the chunks of a subsequent packet distribution.

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 $j \in \Gamma(i)$. For the remainder of this section, let M be the minimum required number of such transmissions that permit 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 M 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 $M \geq 3$. 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 $M = 3$.

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 $M \geq 2$. 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 $M = 2$. 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.

III. Main Results

In this section, we state our main results. In subsection III-A, we define a set of conditions which are both necessary and sufficient for universal recovery. In subsection III-B, we derive asympotic bounds on the number of transmissions required for universal recovery by considering a slice of the region defined in subsection III-A. Subsection III-C specializes the results of subsection III-A for a wide class of networks. In particular, we demonstrate how to explicitly construct a transmission strategy which uses at most a constant number of transmissions more than the minimum number of transmissions required for universal recovery. Subsection III-D summarizes the results in the context of a probabilistic framework (i.e., when random packet distributions are considered). Proofs of all results are delayed until the Appendix.

A. Transmission Strategies Permitting Universal Recovery

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 :

$$\sum_{j=1}^{r} \sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i\in\mathcal{S}_r} \mathcal{P}_i^c \right|. \tag{2}$$

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 given by (2) 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 $j \in [r]$. For convenience, we refer to this type of sequence of sets as an *appropiate sequence of sets*. The inequalities defined by all appropriate

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

In this paper, we are interested in the minimum number of transmissions required to permit universal recovery in a network. Since the number of transmissions required is finite, there exists some finite r_0 for which a vector $\{b_i^j\} \in \mathcal{R}_{r_0}$ defines a transmission strategy requiring the minimum number of transmissions.

For any feasible vector $\{b_i^j\}$ (i.e. one that satisfies $\{b_i^j\} \in \mathcal{R}_r$ for some 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 [12]. Therefore, the difficulty in describing a transmission strategy which achieves (or approaches) the minimum number of transmissions lies solely in solving a minimization problem over $\mathcal{R}_r$.

Remark 1: Theorem 1 (and the subsequent results) can be readily extended to include the case where only a subset of the nodes wishes to reconstruct the original k packets and the other nodes serve only as helpers or relays. However, in order to keep the arguments and notation simple, we restrict our attention to the case where universal recovery is the objective.

B. Asymptotic Results for General Networks

This section considers what happens asymptotically as the number of packets k becomes large. Of particular interest is the case where the sequence of packet distributions $\{\mathcal{P}_i(k)\}_{i=1}^N$ is *well-behaved* as defined in (1).

Theorem 1 characterizes all transmission strategies by means of a very complicated sequence of cuts. By considering the asymptotic (large k) case, Theorem 2 simplifies this result to yield a more meaningful region which still characterizes feasible transmission strategies of interest. Using only simple cuts in the network, Theorem 2 characterizes a sufficient condition permitting universal recovery in the asymptotic regime under a per-round constraint on the number of transmissions by each node. In other words, for any $\epsilon > 0$, we can find transmission strategies requiring at most $k\epsilon$ more transmissions than the minimum number allowing universal recovery. This asymptotic result is analogous to a recent result of Mohajer et. al. [13] for the networks under consideration.

Theorem 2: For a fixed network topology, $\tau \triangleq k/r$ fixed, and a well-behaved sequence of packet distributions $\{\mathcal{P}_i(k)\}_{i=1}^N$, if $\{b_i\}$ satisfies:

$$\sum_{i \in \partial(\mathcal{S})} b_i > \tau \mathbb{P}_{\mathcal{S}}^c, \qquad \forall\ \emptyset \subsetneq \mathcal{S} \subsetneq [N] \tag{3}$$

then any vector $\{b_i^j\}$ satisfying $b_i^j \geq b_i$ for all $i \in [N], j \in [r]$ permits universal recovery for all sufficiently large k.

The interpretation of Theorem 2 is as follows: if node i is allowed to make at least b_i transmissions per communication round, then universal recovery is possible if the b_i's flowing into any set $\mathcal{S}$ from its bordering nodes $\partial(\mathcal{S})$ represent a total information flow larger than the information $\mathcal{S}$ is missing. Theorem 1 yields a corresponding converse.

Theorem 3: Universal recovery is not possible if node i is allowed to make at most b_i transmissions per communication round and there is some set $\emptyset \subsetneq \mathcal{S} \subsetneq [N]$ for which

$$\sum_{i \in \partial(\mathcal{S})} b_i < \frac{1}{r} \left| \bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c \right|.$$

Equation (3) indicates a possible asymptotic gap of $r\epsilon_r$ transmissions between the constraints of Theorem 2 and a truly optimal transmission strategy computed by optimizing over the region $\mathcal{R}_{r_0}$ directly. However, a key constraint in the proof of Theorem 2 is that the number of communication rounds r grows linearly with k as $k \to \infty$. This yields a potential gap of $k\epsilon_k$ transmissions where $\epsilon_k = \epsilon_r/\tau$.

From a combinatorial point of view, it is perhaps more interesting to approximate the minimum number of transmissions required within a constant factor (not depending on k or r). The next section accomplishes this for a wide class of networks by constructing transmission strategies that require at most N more transmissions than the required minimum.

C. Bounding Excess Transmissions by the Number of Nodes

Analysis in [1] analytically computed the minimum number of transmissions required for universal recovery in clique networks with a special type of packet distribution. This section presents tight bounds on the minimum number of transmissions for a wide class of networks by considering the asymptotic scenario where k is large.

Vector $\{b_i^j\}$ is "within ℓ transmissions of optimal" if the minimum number of transmissions required for universal recovery is M and $\sum_{i,j} b_i^j \leq M + \ell$. In this section, we show for a wide class of networks that we can construct vectors $\{b_i^j\}$ that are within N transmissions of optimal. Recall that N is the number of nodes in the network. Thus, we can bound the minimum number of transmissions by the network parameter N, instead of by a factor linear in k as in the previous section.

Since we plan to determine (or approximate) the minimum required number of transmissions analytically, it is necessary to introduce some notational machinery based on the network topology and the limiting packet distribution $\{\mathbb{P}_{\mathcal{S}}\}_{\mathcal{S} \subseteq [N]}$. In what follows, we shall always assume that we are working with a well-behaved sequence of packet distributions with limiting distribution $\{\mathbb{P}_{\mathcal{S}}\}_{\mathcal{S} \subseteq [N]}$.

Define A to be the adjacency matrix of the graph $\mathcal{G}_{\mathcal{T}}$ (i.e. $a_{i,j} = 1$ if $(i,j) \in E$ and $a_{i,j} = 0$ otherwise). If A is nonsingular, we say that the network is nonsingular.

Let $\vec{\mathcal{P}}^c(k) = [|\mathcal{P}_1^c(k)|, \ldots, |\mathcal{P}_N^c(k)|]^T$. Assuming A is nonsingular, define $\delta_i(k) = \left[A^{-1}\vec{\mathcal{P}}^c(k)\right]_i$, where $[\vec{x}]_i$ denotes the i^{th} coordinate of the vector $\vec{x}$. The following theorem gives sufficient conditions that, when met, allow us to explicitly construct a transmission strategy $\{b_i^j\}$ from the quantities k, r, $\delta_i(k)$.

Theorem 4: If $\delta_i(k) \geq 0$ and there exists some k_0 so that

for all $k \geq k_0$:

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

for all appropriate sequences of sets with $|\mathcal{S}_r| \geq 2$, then choosing b_i^j so that $b_i^j \geq \lfloor \frac{1}{r}\delta_i(k) \rfloor$ and $\sum_{j=1}^{r} b_i^j = \lceil \delta_i(k) \rceil$ yields a vector $\{b_i^j\}$ that is within N transmissions of optimal.

Remark 2: If packets are divisible, then we can always choose $\{b_i^j\}$ so that $\sum_{j=1}^{r} b_i^j = \delta_i(k)$. Thus, the optimum can be achieved by packet splitting. See [1] for a detailed discussion regarding divisible packets.

The main application of Theorem 4 is to networks which have some structure allowing the δ_i's to be easily computed. In this paper, we are particularly interested in the case of d-regular networks (i.e. every node has d neighbors) and packet distributions where $\mathbb{P}_i = \rho$ for all $i \in [N]$ (i.e. each node has approximately $\rho \times k$ packets). The choice of d-regularity includes many canonical networks (e.g. cliques, rings, grids on tori, Cayley graphs, etc.), and the choice of the uniform-type packet distribution is inspired by those real-world applications where the networks are approximately homogeneous.

Theorem 5 below requires the network to be d-connected and nonsingular[1]. A d-connected network is a network in which at least d nodes must be removed in order to disconnect the network. Clearly the connectivity of a d-regular network is at most d since one can remove the nodes in $\partial(i)$ to disconnect i from the rest of the network. It turns out (see [14]) that almost every large random d-regular network is d-connected, therefore our choice to enforce the d-connectivity criterion serves to eliminate certain pathological realizations of networks which would rarely (if ever) appear in practice.

Theorem 5: All nonsingular d-regular, d-connected networks with $\rho > 0$ fixed,

$$\mathbb{P}_i = \rho < \mathbb{P}_{\mathcal{S}}, \quad \forall i \in [N], \ \forall \mathcal{S} : |\mathcal{S}| > 1, \tag{4}$$

and

$$(N - |\mathcal{S}|) \cdot \mathbb{P}_i^c > d \cdot \mathbb{P}_{\mathcal{S}}^c, \quad \forall \mathcal{S} : |\mathcal{S}| > N - d \tag{5}$$

satisfy the conditions of Theorem 4. Accordingly, one can analytically compute a vector $\{b_i^j\}$ that is within N transmissions of optimum. Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d}\sum_{i \in [N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i \in [N] \\ j \in [r]}} b_i^j \leq \frac{1}{d}\sum_{i \in [N]} |\mathcal{P}_i^c(k)| + N. \tag{6}$$

In certain situations, some conditions required by Theorem 5 are not necessary. We state these as corollaries.

Corollary 1: If $\mathbb{P}_i < \frac{1}{d+1}$, then condition (5) in Theorem 5 can be omitted.

Corollary 2: If $\mathbb{P}_{\mathcal{S}} = 1$ whenever $|\mathcal{S}| \geq N - d + 1$, then condition (5) in Theorem 5 can be omitted.

[1]Costello and Vu have recently conjectured that, for $d > 2$, almost every large d-regular network is nonsingular (See [15], [16].).

Corollary 3: If the packet distribution satisfies

$$\mathbb{P}_{\mathcal{S}} = \frac{1 - (1-q)^{|\mathcal{S}|}}{1 - (1-q)^N}, \quad \forall \mathcal{S} \subseteq [N],$$

for any $0 < q < 1$, then conditions (4-5) in Theorem 5 can be omitted.

Whenever the conditions of Theorem 5 or its corollaries are met, we can explicitly construct a transmission strategy $\{b_i^j\}$ that is within N transmissions of optimal. Thus, we have an analytical expression for transmission strategies that approximate the performance of the optimum strategy computable by solving an Integer Linear Program over the region $\mathcal{R}_{r_0}$.

D. Random Distributions of Packets

Instead of considering well-behaved sequences of packet distributions as defined in (1), we can instead consider random distributions of packets which satisfy $\frac{1}{k}\left|\bigcup_{i \in \mathcal{S}} \mathcal{P}_i(k)\right| \to \mathbb{P}_{\mathcal{S}}$ in probability for all subsets $\mathcal{S} \subseteq [N]$. Then, by standard arguments, it can be shown that all the previous results hold with arbitrarily high probability as $k \to \infty$.

One particularly important random distribution that we consider is when each node has a packet with probability q independent from other nodes. Necessarily, we must condition on the fact that each packet is available to at least one node, yielding:

$$\mathbb{P}_{\mathcal{S}} = \frac{1 - (1-q)^{|\mathcal{S}|}}{1 - (1-q)^N}, \quad \forall \mathcal{S} \subseteq [N]. \tag{7}$$

Abusing terminology slightly, we call this the *independent distribution*.

Due to the importance of this interpretation, we summarize some of the previous results in this probabilistic context. Note that regardless of the random distribution considered, we always require that $\mathbb{P}_{[N]} = 1$ (i.e., all N nodes collectively have all the packets).

Theorem 6: Suppose that each of k total packets is distributed in a nonsingular, d-regular, d-connected network according to an i.i.d. process, and that the marginal probability of a node having a packet is ρ for all nodes. If any one of the following conditions are true:

1) $\rho < \frac{1}{d+1}$.
2) The probability of any subset of nodes of size greater than $N - d$ missing a packet is zero.
3) Packets are distributed according to the independent distribution (7).

then the transmission strategy $\{b_i^j\}$ obtained by choosing b_i^j so that $b_i^j \geq \lfloor \frac{1}{r}\delta_i(k) \rfloor$ and $\sum_{j=1}^{r} b_i^j = \lceil \delta_i(k) \rceil$ yields a vector $\{b_i^j\}$ that is within N transmissions of optimal and is feasible with arbitrarily high probability (as $k \to \infty$). Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d}\sum_{i \in [N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i \in [N] \\ j \in [r]}} b_i^j \leq \frac{1}{d}\sum_{i \in [N]} |\mathcal{P}_i^c(k)| + N.$$

IV. Concluding Remarks

This paper studies connected broadcast networks of N nodes that all wish to recover k desired packets originally dispersed among the nodes. The key results give a series of relatively simple conditions for universal recovery that apply when the number of packets is large and the distribution of packets among the nodes is well-behaved.

For any fixed network topology, it suffices to consider only simple cuts in the network to characterize a set of transmission strategies that requires at most $k\epsilon$ transmissions above the minimum required for universal recovery (ϵ can be arbitrarily small). For certain topologies including non-singular d-regular d-connected networks, this paper shows how to explicitly construct transmission strategies that achieve universal recovery while using at most N transmissions above the minimum number required, thus essentially resolving the universal recovery problem completely for many canonical networks (e.g. cliques, rings, grids on tori, etc.).

Appendix

This appendix contains all the proofs of the results given in the paper.

Proof of Theorem 1 : 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 [4].

First, fix the number of communication rounds r large enough to permit universal recovery. For a network $\mathcal{T}$, construct the network-coding graph $\mathcal{G}_{\mathcal{T}}^{NC} = (V_{NC}, E_{NC})$ as follows. The vertex set, V_{NC} is defined as:

$$V_{NC} = \{s, u_1, \ldots, u_k\} \cup \bigcup_{j=0}^{r} \{v_1^j, \ldots, v_N^j\} \cup \bigcup_{j=1}^{r} \{w_1^j, \ldots, w_N^j\}.$$

The edge set, E_{NC}, consists of directed edges and is constructed as follows:

- For each $i \in [k]$, there is an edge of unit capacity[2] from s to u_i.
- If $p_i \in \mathcal{P}_j$, then there is an edge of infinite capacity from u_i to v_j^0.
- For each $j \in [r]$ and each $i \in [N]$, there is an edge of infinite capacity from v_i^{j-1} to v_i^j.
- For each $j \in [r]$ and each $i \in [N]$, there is an edge of capacity b_i^j from v_i^{j-1} to w_i^j.
- For each $j \in [r]$ and each $i \in [N]$, there is an edge of infinite capacity from w_i^j to $v_{i'}^j$ iff $i' \in \Gamma(i)$.

The interpretation of this graph is as follows: the vertex u_i is introduced to represent packet p_i, the vertex v_i^j represents node i after the j^{th} round of communication, and the vertex w_i^j represents the broadcast of node i during the j^{th} round of communication. If the b_i^j's are chosen such that the graph $\mathcal{G}_{\mathcal{T}}^{NC}$ admits a network coding solution which supports a multicast of k units from s to $\{v_1^r, \ldots, v_N^r\}$, then this network coding solution also solves the universal recovery problem for the network $\mathcal{T}$ when node i is allowed to make at most b_i^j transmissions during the j^{th} round of communication.

We now formally prove the equivalence of the network coding problem on $\mathcal{G}_{\mathcal{T}}^{NC}$ and the universal recovery problem defined by $\mathcal{T}$.

Suppose a set of encoding functions $\{f_i^j\}$ and a set decoding functions $\{\phi_i\}$ describe a transmission strategy which solves the universal recovery problem for a network $\mathcal{T}$ in r rounds of communication. Let b_i^j be the number of transmissions made by node i during the j^{th} round of communication, and let $\mathcal{I}_i^j$ be all the information known to node i prior to the j^{th} round of communication (e.g. $\mathcal{I}_i^1 = \mathcal{P}_i$). The function f_i^j is the encoding function for user i during the j^{th} round of communication (i.e. $f_i^j(\mathcal{I}_i^j) \in \mathbb{F}^{b_i^j}$), and the decoding functions satisfy:

$$\phi_i\left(\mathcal{I}_i^r, \cup_{i' \in \Gamma(i)} \{f_{i'}^r(\mathcal{I}_{i'}^r)\}\right) = \{p_1, \ldots, p_k\}.$$

Note that, given the encoding functions and the $\mathcal{P}_i$'s, the $\mathcal{I}_i^j$'s can be defined recursively as:

$$\mathcal{I}_i^{j+1} = \mathcal{I}_i^j \cup \bigcup_{i' \in \Gamma(i)} \{f_{i'}^j(\mathcal{I}_{i'}^j)\}.$$

The functions $\{f_i^j\}$ and $\{\phi_i\}$ can be used to generate a network coding solution which supports k units of flow from s to $\{v_1^r, \ldots, v_N^r\}$ on $\mathcal{G}_{\mathcal{T}}^{NC}$ as follows:

For each vertex $v \in V_{NC}$, let $\text{IN}(v)$ be whatever v receives on its incoming edges. Let g_v be the encoding function at vertex v, and $g_v(e, \text{IN}(v))$ be the encoded message which vertex v sends along e (e is an outgoing edge from v).

If e is an edge of infinite capacity emanating from v, let $g_v(e, \text{IN}(v)) = \text{IN}(v)$.

Let s send p_i along edge (s, u_i). At this point, we have $\text{IN}(v_i^0) = \mathcal{P}_i = \mathcal{I}_i^1$. For each $i \in [N]$, let $g_{v_i^0}((v_i^0, w_i^1), \text{IN}(v_i^0)) = f_i^1(\mathcal{I}_i^1)$. By a simple inductive argument, defining the encoding functions $g_{v_i^j}((v_i^j, w_i^{j+1}), \text{IN}(v_i^j))$ to be equal to f_i^{j+1} yields the result that $\text{IN}(v_i^r) = \left(\mathcal{I}_i^r, \cup_{i' \in \Gamma(i)} \{f_{i'}^r(\mathcal{I}_{i'}^r)\}\right)$. Hence, the decoding function ϕ_i can be used at v_i^r to allow error-free reconstruction of the k-unit flow.

The equivalence argument is completed by showing that a network coding solution which supports a k-unit multicast flow from s to $\{v_1^r, \ldots, v_N^r\}$ on $\mathcal{G}_{\mathcal{T}}^{NC}$ also solves the universal recovery problem on $\mathcal{T}$. This is argued in a similar manner as above, and is therefore omitted.

Since we have shown that the universal recovery problem on $\mathcal{T}$ is equivalent to a network coding problem on $\mathcal{G}_{\mathcal{T}}^{NC}$, the celebrated max-flow min-cut result of Ahlswede et. al [4] is applicable. In particular, a fixed vector $\{b_i^j\}$ admits a solution to the universal recovery problem where node i makes at most b_i^j transmissions during the j^{th} round of communication if and only if any cut separating s from some v_i^r in $\mathcal{G}_{\mathcal{T}}^{NC}$ has capacity at least k.

What remains to be shown is that the inequalities in (2) are satisfied if and only if any cut separating s from some v_i^r in $\mathcal{G}_{\mathcal{T}}^{NC}$ has capacity at least k.

[2] An edge of unit capacity can carry one field element $z \in \mathbb{F}$ per unit time.

Suppose we have a cut $(\mathcal{S}, \mathcal{S}^c)$ satisfying $s \in \mathcal{S}^c$ and $v_i^r \in \mathcal{S}$ for some $i \in [N]$. We will modify the cut $(\mathcal{S}, \mathcal{S}^c)$ to produce a new cut $(\mathcal{S}', \mathcal{S}'^c)$ with capacity less than or equal to the capacity of the original cut $(\mathcal{S}, \mathcal{S}^c)$.

Define the set $\mathcal{S}_0 \subseteq [N]$ as follows: $i \in \mathcal{S}_0$ iff $v_i^r \in \mathcal{S}$ (by definition of $\mathcal{S}$, we have that $\mathcal{S}_0 \neq \emptyset$).

Initially, let $\mathcal{S}' = \mathcal{S}$. Modify the cut $(\mathcal{S}', \mathcal{S}'^c)$ as follows:

M1) If $i \in \Gamma(\mathcal{S}_0)$, then place w_i^r into $\mathcal{S}'$.
M2) If $i \notin \Gamma(\mathcal{S}_0)$, then place w_i^r into $\mathcal{S}'^c$.

Modifications M1 and M2 are justified (respectively) by J1 and J2:

J1) If $i \in \Gamma(\mathcal{S}_0)$, then there exists an edge of infinite capacity from w_i^r to some $v_{i'}^r \in \mathcal{S}$. Thus, moving w_i^r to $\mathcal{S}'$ (if necessary) does not increase the capacity of the cut.
J2) If $i \notin \Gamma(\mathcal{S}_0)$, then there are no edges from w_i^r to $\mathcal{S}$, hence we can move w_i^r into $\mathcal{S}'^c$ (if necessary) without increasing the capacity of the cut.

Modifications M1 and M2 guarantee that $w_i^r \in \mathcal{S}'$ iff $i \in \Gamma(\mathcal{S}_0)$. Thus, assume that $(\mathcal{S}', \mathcal{S}'^c)$ satisfies this condition and further modify the cut as follows:

M3) If $i \in \mathcal{S}_0$, then place v_i^{r-1} into $\mathcal{S}'$.
M4) If $i \notin \Gamma(\mathcal{S}_0)$, then place v_i^{r-1} into $\mathcal{S}'^c$.

Modifications M3 and M4 are justified (respectively) by J3 and J4:

J3) If $i \in \mathcal{S}_0$, then there exists an edge of infinite capacity from v_i^{r-1} to $v_i^r \in \mathcal{S}$. Thus, moving v_i^{r-1} to $\mathcal{S}'$ (if necessary) does not increase the capacity of the cut.
J4) If $i \notin \Gamma(\mathcal{S}_0)$, then there are no edges from v_i^{r-1} to $\mathcal{S}'$ (since $w_i^r \notin \mathcal{S}'$ by assumption), hence we can move v_i^{r-1} into $\mathcal{S}'^c$ (if necessary) without increasing the capacity of the cut.

At this point, define the set $\mathcal{S}_1 \subseteq [N]$ as follows: $i \in \mathcal{S}_1$ iff $v_i^{r-1} \in \mathcal{S}'$. Note that the modifications of $\mathcal{S}'$ guarantee that $\mathcal{S}_1$ satisfies $\mathcal{S}_0 \subseteq \mathcal{S}_1 \subseteq \Gamma(\mathcal{S}_0)$.

This procedure can be repeated for each layer of the graph resulting in a sequence of sets $\emptyset \subsetneq \mathcal{S}_0 \subseteq \cdots \subseteq \mathcal{S}_r \subseteq [N]$ satisfying $\mathcal{S}_j \subseteq \Gamma(\mathcal{S}_{j-1})$ for each $j \in [r]$.

We now perform a final modification of the cut $(\mathcal{S}', \mathcal{S}'^c)$:

M5) If $p_j \in \cup_{i \in \mathcal{S}_r} \mathcal{P}_i$, then place u_j into $\mathcal{S}'$.
M6) If $p_j \notin \cup_{i \in \mathcal{S}_r} \mathcal{P}_i$, then place u_j into $\mathcal{S}'^c$.

Modifications M5 and M6 are justified (respectively) by J5 and J6:

J5) If $p_j \in \cup_{i \in \mathcal{S}_r} \mathcal{P}_i$, then there is an edge of infinite capacity from u_j to $\mathcal{S}'$ and moving u_j into $\mathcal{S}'$ (if necessary) does not increase the capacity of the cut.
J6) If $p_j \notin \cup_{i \in \mathcal{S}_r} \mathcal{P}_i$, then there are no edges from u_j to $\mathcal{S}'$, hence moving u_j (if necessary) into $\mathcal{S}'^c$ cannot increase the capacity of the cut.

A quick calculation shows that the modified cut $(\mathcal{S}', \mathcal{S}'^c)$ has capacity greater than or equal to k iff:

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{r+1-j} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|. \tag{8}$$

Since every modification of the cut either preserved or reduced the capacity of the cut, the original cut $(\mathcal{S}, \mathcal{S}^c)$ also has capacity greater than or equal to k if the above inequality is satisfied.

By the equivalence of the universal recovery problem on a network $\mathcal{T}$ to the network coding problem on $\mathcal{G}_{\mathcal{T}}^{NC}$ and the max-flow min-cut theorem for network information flow, if a transmission scheme permits universal recovery in $\mathcal{T}$, then the associated b_i^j's must satisfy the constraints of the form given by (8). Conversely, for any set of b_i^j's which satisfy the constraints of the form given by (8), there exists a transmission scheme using exactly those numbers of transmissions which permits universal recovery.

Remark 3: Since $\left| \bigcap_{i \in [N]} \mathcal{P}_i^c \right| = 0$, constraints where $\mathcal{S}_r = [N]$ are trivially satisfied. Therefore, we can restrict our attention to sequences of sets where $\mathcal{S}_r \subsetneq [N]$.

■

Proof of Theorem 2: Fix $\{b_i\}$ and assume that

$$\sum_{i \in \mathcal{S}^c \cap \Gamma(\mathcal{S})} b_i \geq \tau \left(\mathbb{P}_{\mathcal{S}}^c + \epsilon \right)$$

for some sufficiently small $\epsilon > 0$ and all subsets $\mathcal{S} \subseteq [N]$. Take k sufficiently large so that $\mathbb{P}_{\mathcal{S}}^c \geq \frac{1}{k} \left| \cap_{i \in \mathcal{S}} \mathcal{P}_i^c(k) \right| - \frac{\epsilon}{2}$ for all $\mathcal{S}$ and also so that $k\frac{\epsilon}{2} \geq \sum_{i=1}^{N} b_i$. Then, if $b_i^j \geq b_i$, the following string of inequalities together with Theorem 1 proves that $\{b_i^j\}$ permits universal recovery.

$$\begin{aligned}
\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^j &\geq \sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i \\
&= \sum_{j=1}^{r} \sum_{i \in \mathcal{S}_{j-1}^c \cap \Gamma(\mathcal{S}_{j-1})} b_i - \sum_{i \in \mathcal{S}_0^c \cap \mathcal{S}_r} b_i \\
&\geq \tau \sum_{j=1}^{r} \left(\mathbb{P}_{\mathcal{S}_{j-1}}^c + \epsilon \right) - \sum_{i=1}^{N} b_i \\
&\geq r\tau \left(\mathbb{P}_{\mathcal{S}_r}^c + \epsilon \right) - \sum_{i=1}^{N} b_i \\
&\geq k \left(\frac{1}{k} \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c(k) \right| + \frac{\epsilon}{2} \right) - \sum_{i=1}^{N} b_i \\
&\geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c(k) \right|.
\end{aligned}$$

■

Proof of Theorem 3: Take $\mathcal{S}_i = \mathcal{S}$ for $i \in [r] \cup \{0\}$. This

yields the following string of inequalities:

$$\begin{aligned}\sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})} b_i^j &= \sum_{j=1}^{r}\sum_{i\in\mathcal{S}^c\cap\Gamma(\mathcal{S})} b_i^j \\ &\le \sum_{j=1}^{r}\sum_{i\in\partial(\mathcal{S})} b_i \\ &< \left|\bigcap_{i\in\mathcal{S}_r}\mathcal{P}_i^c\right|.\end{aligned}$$

Thus, universal recovery is not possible by Theorem 1. ∎

Proof of Theorem 4: Suppose $\sum_{j=1}^{r}\hat{b}_i^j = \delta_i(k)$, then $\sum_{i,j}\hat{b}_i^j$ is a lower bound on the minimum number of transmissions required because

$$\sum_{j=1}^{r}\sum_{i\in\partial(i)}\hat{b}_i^j = |\mathcal{P}_i^c(k)| \quad \text{for all } i.$$

In other words, each node can receive at most the number of packets it is missing. Therefore, if $\sum_{j=1}^{r} b_i^j = \lceil\delta_i(k)\rceil$, and $\{b_i^j\}$ is feasible, then it describes a transmission strategy must be within N transmissions of optimal since $\sum_{i,j} b_i^j \le \sum_{i,j}\hat{b}_i^j + N$. Thus, it just remains to be shown that $\{b_i^j\}$ is feasible when the conditions of the theorem are met.

Note that we can always choose b_i^j so that $b_i^j \ge \lfloor\frac{1}{r}\delta_i(k)\rfloor$ and $\sum_{j=1}^{r} b_i^j = \lceil\delta_i(k)\rceil$. By this definition of $\{b_i^j\}$, we have that $\{b_i^j\}$ satisfies (2) when $|\mathcal{S}_r| = 1$, therefore, in what follows we only consider $|\mathcal{S}_r| \ge 2$.

$$\begin{aligned}\sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})} b_i^j &\ge \sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})}\left\lfloor\frac{1}{r}\delta_i(k)\right\rfloor \\ &\ge \frac{1}{r}\sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})}\delta_i(k) - rN \\ &\ge \left|\bigcap_{i\in\mathcal{S}_r}\mathcal{P}_i^c(k)\right|.\end{aligned}$$

Therefore, $\{b_i^j\} \in \mathcal{R}_r$, and permits universal recovery by Theorem 1. ∎

Proof of Theorem 5: Since the packet distributions are well behaved, for any $\epsilon > 0$ we have that $\frac{k}{d}(1-\rho-\epsilon) \le \delta_i(k) \le \frac{k}{d}(1-\rho+\epsilon)$ for all i when k is sufficiently large. By the hypothesis that the network is d-connected, $|\partial(\mathcal{S})| \ge d$ whenever $|\mathcal{S}| \le N-d$ and $|\partial(\mathcal{S}_1)| \ge N - |\mathcal{S}_2|$ whenever $|\mathcal{S}_2| \ge N-d$ and $\mathcal{S}_1 \subseteq \mathcal{S}_2$. We consider the cases where $2 \le |\mathcal{S}_r| \le N-d$ and $N-d < |\mathcal{S}_r| \le N-1$ separately.

First, note the following inequality:

$$\begin{aligned}\sum_{j=1}^{r}\sum_{i\in\mathcal{S}_j^c\cap\Gamma(\mathcal{S}_{j-1})}\delta_i(k) &= \sum_{j=1}^{r}\sum_{i\in\partial(\mathcal{S}_{j-1})}\delta_i(k) - \sum_{i\in\mathcal{S}_r\cap\mathcal{S}_0^c}\delta_i(k) \\ &\ge \frac{k}{d}(1-\rho)\left(\sum_{j=1}^{r}|\partial(\mathcal{S}_{j-1})| - N\right) \\ &\quad - k\nu(\epsilon),\end{aligned}$$

where $\lim_{\epsilon\to 0}\nu(\epsilon) = 0$. Noting that $(1-\rho) = \mathbb{P}_i^c$, in order to verify the conditions of Theorem 4, it is sufficient to check that

$$\frac{1}{d}\mathbb{P}_i^c\left(\sum_{j=1}^{r}|\partial(\mathcal{S}_{j-1})| - N\right) > r\cdot\mathbb{P}_{\mathcal{S}_r}^c.$$

Or equivalently, that the following holds:

$$\frac{1}{rd}\sum_{j=1}^{r}|\partial(\mathcal{S}_{j-1})| > \frac{\mathbb{P}_{\mathcal{S}_r}^c}{\mathbb{P}_i^c} + \frac{N}{rd}.$$

Assume that $1 < |\mathcal{S}_r| \le N-d$, then we have the following:

$$\begin{aligned}\frac{1}{rd}\sum_{j=1}^{r}|\partial(\mathcal{S}_{j-1})| &\ge \frac{1}{rd}\sum_{j=1}^{r} d \\ &= 1 \\ &> \frac{\mathbb{P}_{\mathcal{S}_r}^c}{\mathbb{P}_i^c} + \frac{N}{rd}.\end{aligned}$$

Where the last inequality holds if r is taken sufficiently large (how large depends only on N, d, and $\max_{\mathcal{S}:|\mathcal{S}|>1}\mathbb{P}_{\mathcal{S}}^c$).

Next, assume that $|\mathcal{S}_r| > N-d$, then by hypothesis, we have:

$$\begin{aligned}\frac{1}{rd}\sum_{j=1}^{r}|\partial(\mathcal{S}_{j-1})| &\ge \frac{N-|\mathcal{S}_r|}{d} \\ &> \frac{\mathbb{P}_{\mathcal{S}_r}^c}{\mathbb{P}_i^c} + \frac{N}{rd}.\end{aligned}$$

Where the last inequality again holds for sufficiently large r (depending only on N, d, and $\max_{\mathcal{S}:|\mathcal{S}|>N-d}\mathbb{P}_{\mathcal{S}}^c$).

The bounds on $\sum_{i,j} b_i^j$ are a result of direct computation. ∎

Proof of Corollary 1: Note that $\mathbb{P}_i < \frac{1}{d+1}$ implies that $\mathbb{P}_i^c > \frac{d}{d+1}$. Then, note that we have the following string of inequalities:

$$\begin{aligned}d\cdot\mathbb{P}_{\mathcal{S}}^c &\le d\cdot\mathbb{P}_{\mathcal{S}^c} \\ &\le d\cdot\sum_{i\in\mathcal{S}^c}\mathbb{P}_i \\ &= d\cdot(N-|\mathcal{S}|)\cdot\mathbb{P}_i \\ &< \frac{d}{d+1}\cdot(N-|\mathcal{S}|) \\ &< \mathbb{P}_i^c\cdot(N-|\mathcal{S}|).\end{aligned}$$

Thus, $d\cdot\mathbb{P}_{\mathcal{S}}^c < (N-|\mathcal{S}|)\cdot\mathbb{P}_i^c$ and the conditions of Theorem 5 are satisfied. ∎

Proof of Corollary 2: This trivially implies that $\mathbb{P}_{\mathcal{S}}^c = 0$. Thus, $d\cdot\mathbb{P}_{\mathcal{S}}^c = 0 < (N-|\mathcal{S}|)\cdot\mathbb{P}_i^c$ and the conditions of Theorem 5 are satisfied. ∎

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

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

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

$$\delta = \min_{\ell \in \{2,\ldots,N-1\}} \left[\frac{N-\ell}{N-1} - \frac{(1-q)^\ell - (1-q)^N}{1-q-(1-q)^N} \right]$$

completes the proof. ∎

Proof of Corollary 3: Under the assumption that

$$\mathbb{P}_{\mathcal{S}} = \frac{1-(1-q)^{|\mathcal{S}|}}{1-(1-q)^N}, \quad \forall \mathcal{S} \subseteq [N],$$

we have the following by Lemma 1:

$$\begin{aligned} \frac{N-|\mathcal{S}|}{d} &\geq \frac{N-|\mathcal{S}|}{N-1} \\ &> \frac{(1-q)^{|\mathcal{S}|} - (1-q)^N}{(1-q)-(1-q)^N} \\ &= \frac{\mathbb{P}^c_{\mathcal{S}}}{\mathbb{P}^c_i}. \end{aligned}$$

Thus, $d \cdot \mathbb{P}^c_{\mathcal{S}} < (N-|\mathcal{S}|) \cdot \mathbb{P}^c_i$ and the conditions of Theorem 5 are satisfied. ∎

Proof of Theorem 6: The theorem is an immediate consequence of Corollaries 1-3 and the weak law of large numbers. ∎

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.

REFERENCES

[1] T. Courtade, B. Xie, and R. Wesel, "Optimal Exchange of Packets for Universal Recovery in Broadcast Networks," To appear at MILCOM 2010, San Jose, CA, October 31 - November 3, 2010. Preprint available online at http://www.ee.ucla.edu/∼csl/files/publications.html.

[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] 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.

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

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

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

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

[12] 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.

[13] S. Mohajer, C. Tian, and S. Diggavi, "On Source Transmission over Deterministic Relay Networks," IEEE Information Theory Workshop, Cairo, Egypt, January 6-8, 2010.

[14] B. Bollobás, Random Graphs, 2nd edition, Cambridge University Press (2001).

[15] Costello, K.; Vu, Van, The rank of random graphs. Random Structures Algorithms 33 (2008), no. 3, 269-285.

[16] Costello, K., Vu, Van, On the rank of random sparse matrices. Combin. Probab. Comput. 19 (2010), no. 3, 321-342.