

Coded Cooperative Data Exchange in Multihop Networks

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

Abstract—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 exchanges coded packets with its neighbors. This paper provides necessary and sufficient conditions that characterize the set of all transmission strategies that permit every node to ultimately learn (recover) all k packets. When the network satisfies certain regularity conditions and packets are randomly distributed, this paper provides tight concentration results on the number of transmissions required to achieve universal recovery. For the case of a fully connected network, a polynomial-time algorithm for computing an optimal transmission strategy is derived. An application to secrecy generation is discussed.

Index Terms—Coded cooperative data exchange, universal recovery, network coding.

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 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. We refer to this as the *Coded Cooperative Data Exchange* (CCDE) problem, or just the Cooperative Data Exchange problem.

This work is motivated in part by emerging issues in distributed data storage. Consider the problem of backing up data on servers in a large data center. One commonly employed method to protect data from corruption is replication. Using this method, large quantities of data are replicated in several locations so as to protect from various sources of corruption (*e.g.*, equipment failure, power outages, natural disasters, *etc.*). As the quantity of information in large data

centers continues to increase, the number of file transfers required to complete a periodic replication task is becoming an increasingly important consideration due to time, equipment, cost, and energy constraints. The results contained in this paper address these issues by characterizing the minimum number of multicast transmissions that must take place in order to successfully replicate the data collectively held by the servers at all locations.

This model also has natural applications in the context of 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 transmitted packets due to interference, obstructions, and other signal integrity issues. In order to recover the transmission, the nodes are free to communicate with their nearby 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? The results in the present paper provide insight into how a cooperative protocol at the application-layer (so that communication links appear noiseless and error-free) can be optimized.

Beyond the examples mentioned above, the results presented herein can also be applied to practical secrecy generation amongst a collection of nodes. We consider this application in detail in Section IV.

A. Related Work

Distributed data exchange problems have received a great deal of attention over the past several years. The powerful techniques afforded by network coding [4], [5] have paved the way for cooperative communications at the packet-level.

The coded cooperative data exchange problem (also called the universal recovery problem in [1]–[3]) was originally introduced by El Rouayheb *et al.* in [6], [7] for a fully connected network (*i.e.*, a single-hop network). For this special case, a randomized algorithm for finding an optimal transmission strategy was given in [8], and the first deterministic algorithm was recently given in [9]. In the concluding remarks of [9], the weighted universal recovery problem (in which the objective is to minimize the weighted sum of transmissions by nodes) was posed as an open problem. However, this was solved using a variant of the same algorithm in [10], and independently by the present authors using the submodular optimization algorithm presented herein (originally appearing in [3]).

Manuscript received March 15, 2012; revised May 31, 2013; accepted July 5, 2013. Date of publication November 14, 2013; date of current version January 15, 2014. This work was supported in part by Rockwell Collins under Contract 4502769987 and in part by the NSF Center for Science of Information under Grant CCF-0939370. This paper was presented at the 2010 Military Communications Conference [1] and the 2010–2011 Allerton Conference on Communication, Control, and Computing [2], [3].

T. A. Courtade was with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095 USA. He is now with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: courtade@stanford.edu).

R. D. Wesel is with the Electrical Engineering Department, University of California, Los Angeles, CA 90095 USA (e-mail: wesel@ee.ucla.edu).

Communicated by M. Franceschetti, Associate Editor for Communications. Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2013.2290993

The coded cooperative data exchange problem is related to the index coding problem originally introduced by Birk and Kol in [11]. 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.

Also closely related to the cooperative data exchange problem is the extensive work on gossip algorithms for information dissemination (cf. [12, Chapter 3] and the references therein). Gossip algorithms are defined as a class of decentralized network communication protocols which can be employed to disseminate data within a network. Gossip algorithms for data dissemination assume each node initially begins with a unique piece of data and that nodes communicate in a unicast fashion. In contrast, the coded cooperative data exchange problem allows for nodes to initially possess common data and permits multicast transmissions. Moreover, typical results in the literature on gossip algorithms show that a gossip protocol can disseminate data in order-optimal time, whereas in the coded cooperative data exchange problem, we seek to characterize the *exact* number of transmissions which must take place in order to successfully disseminate data. Despite these fundamental differences, we would like to point out that the work on gossip algorithms that is most closely related to the present paper is [13]. In [13], the authors show that a gossip algorithm employing random linear network coding can attain order optimal dissemination time. Similarly, in the present paper, we show that linear network coding is sufficient to optimally solve the coded cooperative data exchange problem.

This paper applies principles of cooperative data exchange to generate secrecy in the presence of an eavesdropper. In this context, the secrecy generation problem was originally studied in [14]. In [14], Csiszar and Narayan gave single-letter characterizations of the secret-key and private-key capacities for a network of nodes connected by an error-free broadcast channel. While general and powerful, these results left two practical issues as open questions. First, (as with many information-theoretic investigations) the results require the nodes to observe arbitrarily long sequences of i.i.d. source symbols, which is generally not practical. Second, no constructive algorithm is provided in [14] which achieves the respective secrecy capacities. More recent work in [15], [16] addressed the latter point. Also related, though to a lesser extent, is the body of work on secret-sharing schemes. We refer the reader to the survey [17] and the references therein for further details.

B. Our Contributions

In this paper, we provide necessary and sufficient conditions for achieving universal recovery¹ in arbitrarily connected multihop networks. We specialize these necessary and sufficient conditions to obtain precise results in the case where the underlying network topology satisfies some modest regularity conditions.

¹In this paper, we use the term *universal recovery* to refer to the ultimate condition where every node has successfully recovered all packets.

For the case of a fully connected network, we provide an algorithm based on submodular optimization which solves the cooperative data exchange problem. This algorithm is unique from the others previously appearing in the literature (cf. [8]–[10]) in that it exploits submodularity². As a corollary, we provide exact concentration results when packets are randomly distributed in a network.

In this same vein, we also obtain tight concentration results and approximate solutions when the underlying network is d -regular and packets are distributed randomly.

Furthermore, if packets are divisible (allowing transmissions to consist of partial packets), we prove that the traditional cut-set bounds can be achieved for any network topology. In the case of d -regular and fully connected networks, we show that splitting packets does not typically provide any significant benefits.

Finally, for the application to secrecy generation, we leverage the results of [14] in the context of the cooperative data exchange problem for a fully connected network. In doing so, we provide an polynomial-time algorithm that achieves the secrecy capacity without requiring any quantities to grow asymptotically large.

C. Organization

This paper is organized as follows. Section II formally introduces the problem and provides basic definitions and notation. Section III presents our main results. Section IV discusses the application of our results to secrecy generation by a collection of nodes in the presence of an eavesdropper. Section V contains the relevant proofs. Section VI delivers the conclusions and discusses directions for future work.

II. SYSTEM MODEL AND DEFINITIONS

Before we formally introduce the problem, we establish some notation. Let $\mathbb{N} = 0, 1, 2, \dots$ denote the set of natural numbers. For two sets A and B , the relation $A \subset B$ implies that A is a proper subset of B (i.e., $A \subseteq B$ and $A \neq B$). If it is clear from context that A is a subset of a ground set B , the shorthand notation A^c will be used to denote $B \setminus A$. For a set A , the corresponding power set is denoted $2^A := \{B : B \subseteq A\}$. We use the notation $[m]$ to denote the set $\{1, \dots, m\}$.

This paper considers a network of n nodes. The network must be connected, but it need not be fully connected (i.e., it need not be a complete graph). A graph $\mathcal{G} = (V, E)$ describes the specific connections in the network, where V is the set of vertices $\{v_i : i \in \{1, \dots, n\}\}$ (each corresponding to a node) and E is the set of edges connecting nodes. We assume that the edges in E are undirected.

Each node wishes to recover the same k desired packets, and each node begins with a (possibly empty) subset of the desired packets. Formally, let $P_i \subseteq \{p_1, \dots, p_k\}$ be the (indexed) set of packets originally available at node i , and $\{P_i\}_{i=1}^n$ satisfies $\bigcup_{i=1}^n P_i = \{p_1, \dots, p_k\}$. Each $p_j \in \mathbb{F}$, where \mathbb{F} is some

²We note that [18], which appeared around the same time as [3], proposed an algorithm exploiting submodularity for computing optimal rate allocations in the closely related *communication for omniscience problem* (see [14]). We comment on the relationship in Section III-D and Appendix A.




Fig. 1. For the given graph, a set of vertices S and its neighborhood $\Gamma(S)$ are depicted. The set $\partial(S)$ (i.e., the boundary of S) consists of the four vertices in $\Gamma(S)$ which are not in S .

finite field³ (e.g. $\mathbb{F} = \text{GF}(2^m)$). For our purposes, it suffices to assume $|\mathbb{F}| \geq 2n$. The set of packets initially missing at node i is denoted $P_i^c := \{p_1, \dots, p_k\} \setminus P_i$.

Throughout this paper, we assume that each packet $p_i \in \{p_1, \dots, p_k\}$ is equally likely to be any element of \mathbb{F} . Moreover, we assume that packets are independent of one another. Thus, no correlation between different packets or prior knowledge about unknown packets can be exploited.

To simplify notation, we will refer to a given problem instance (i.e., a graph and corresponding sets of packets available at each node) as a network $\mathcal{T} = \{\mathcal{G}, P_1, \dots, P_n\}$. When no ambiguity is present, we will refer to a network by \mathcal{T} and omit the implicit dependence on the parameters $\{\mathcal{G}, P_1, \dots, P_n\}$.

Let the set $\Gamma(i)$ be the neighborhood 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, we put $i \in \Gamma(i)$. Node i sends (possibly coded) packets to its neighbors $\Gamma(i)$ over discrete, memoryless, and interference-free channels. In other words, if node i transmits a message, then every node in $\Gamma(i)$ receives that message. If S is a set of nodes, then we define $\Gamma(S) = \cup_{i \in S} \Gamma(i)$. In a similar manner, we define $\partial(S) = \Gamma(S) \setminus S$ to be the boundary of the vertices in S . An example of sets S , $\Gamma(S)$, and $\partial(S)$ is given in Figure 1.

This paper seeks to determine the minimum number of transmissions required to achieve universal recovery (when every node has learned all k packets). With the exception of the results given in Section III-D, we focus exclusively on the case where packets are deemed indivisible, motivating the following definitions:

Definition 1: A transmission by node i consists of sending a packet (i.e., some $z \in \mathbb{F}$) to all nodes $j \in \Gamma(i)$.

Definition 2: Given a network \mathcal{T} , the minimum number of transmissions required to achieve universal recovery is denoted $M^*(\mathcal{T})$.

To clarify these definitions, we briefly consider two examples:

Example 1 (Line Network): Suppose \mathcal{T} is a network of nodes connected along a line as follows: $V = \{v_1, v_2, v_3\}$,




Fig. 2. An illustration of the transmission strategy employed in Example 1. During the first time instant, Nodes 1 and 3 broadcast packets p_1 and p_2 , respectively. During the second time instant, Node 2 broadcasts the XOR of packets p_1 and p_2 . This strategy requires three transmissions and achieves universal recovery.

$E = \{(v_1, v_2), (v_2, v_3)\}$, $P_1 = \{p_1\}$, $P_2 = \emptyset$, and $P_3 = \{p_2\}$. Note that each node must transmit at least once in order for all nodes to recover $\{p_1, p_2\}$, hence $M^*(\mathcal{T}) \geq 3$. 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, respectively) 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 node to recover $\{p_1, p_2\}$. Hence $M^*(\mathcal{T}) = 3$.

Example 1 demonstrates a transmission schedule that uses two rounds of communication. The transmissions 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. P_i and previously received transmissions from neighboring nodes). In other words, the transmissions are causal. The transmission strategy employed in Example 1 is illustrated in Figure 2.

Example 2 (Fully Connected Network): Suppose \mathcal{T} is a 3-node fully connected network in which \mathcal{G} is a complete graph on 3 vertices, and $P_i = \{p_1, p_2, p_3\} \setminus p_i$. Clearly one transmission is not sufficient, thus $M^*(\mathcal{T}) \geq 2$. It can be seen that two transmissions suffice: let node 1 transmit p_2 which lets node 2 have $P_2 \cup \{p_2\} = \{p_1, p_2, p_3\}$. Now, node 2 transmits $p_1 \oplus p_3$, allowing nodes 1 and 3 to each recover all three packets. Thus $M^*(\mathcal{T}) = 2$. 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.

In the above examples, we notice that the transmission strategies are partially characterized by a schedule dictating which nodes transmit during which round of communication. We formalize this notion with the following definition:

Definition 3 (Transmission Schedule): A set of integers $\{b_i^j : i \in [n], j \in [r], b_i^j \in \mathbb{N}\}$ is called a transmission schedule for r rounds of communication if node i makes exactly b_i^j transmissions during communication round j .

When the parameters n and r are clear from context, a transmission schedule will be denoted by the shorthand notation $\{b_i^j\}$. Note that b_i^j can exceed one, implying that multiple transmissions per communication round are permitted. However, any transmission schedule with multiple

³We restrict packets to be an element of a field \mathbb{F} for notational simplicity. In practice, packets are typically comprised of a vector of field elements, and arithmetic in the vector space is performed using the operations inherited from \mathbb{F} .

transmissions per round can easily be transformed to a transmission schedule with one or fewer transmissions per round by simply increasing r .

Although finding a transmission schedule that permits universal recovery is relatively easy (e.g., each node transmits all packets in their possession at each time instant), finding one that permits universal recovery with $M^*(\mathcal{T})$ transmissions can be extremely difficult. This is demonstrated by the following example:

Example 3 (Optimal CCDE is NP-Hard.): Suppose \mathcal{T} is a network with $k = 1$ corresponding to a bipartite graph with left and right vertex sets V_L and V_R respectively. Let $P_i = \{p_1\}$ for each $i \in V_L$, and let $P_i = \emptyset$ for each $i \in V_R$. In this case, $M^*(\mathcal{T})$ is given by the minimum number of sets in $\{\Gamma(i)\}_{i \in V_L}$ which cover all vertices in V_R . Thus, finding $M^*(\mathcal{T})$ is at least as hard as the Minimum Set Cover problem, which is NP-complete [19].

Recently, Gonen and Langberg have studied the hardness of the cooperative coded data exchange problem in greater detail. Since we do not address this issue beyond the example above, we refer the reader to their work [20] for a detailed treatment.

Several of our results are stated in the context of *randomly distributed packets*. Assume $0 < q < 1$ is given. Our model is essentially that each packet is available independently at each node with probability q . However, we must condition on the event that each packet is available to at least one node. Thus, when packets are randomly distributed, the underlying probability measure is given by

$$\Pr \left[p_i \in \bigcup_{j \in S} P_j \right] = \frac{1 - (1 - q)^{|S|}}{1 - (1 - q)^n} \quad (1)$$

for all $i \in [k]$ and all nonempty $S \subseteq V = [n]$.

Finally, we introduce one more definition which links the network topology with the number of communication rounds, r .

Definition 4: For a graph $\mathcal{G} = (V, E)$ on n vertices, define $\mathcal{S}^{(r)}(\mathcal{G}) \subset (2^V)^{r+1}$ as follows: $(S_0, S_1, \dots, S_r) \in \mathcal{S}^{(r)}(\mathcal{G})$ if and only if the sets $\{S_i\}_{i=0}^r$ satisfy the following two conditions:

$$\begin{aligned} \emptyset \subset S_i \subset V & \quad \text{for each } 0 \leq i \leq r, \text{ and} \\ S_{i-1} \subseteq S_i \subseteq \Gamma(S_{i-1}) & \quad \text{for each } 1 \leq i \leq r. \end{aligned}$$

In words, any element in $\mathcal{S}^{(r)}(\mathcal{G})$ is a nested sequence of subsets of vertices of \mathcal{G} . Moreover, the constraint that each set in the sequence is contained in its predecessor's neighborhood implies that the sets cannot expand too quickly relative to the topology of \mathcal{G} . Note that, for any graph \mathcal{G} , the set $\mathcal{S}^{(r)}(\mathcal{G}) \subset (2^V)^{r+1}$ is unique.

To make the definition of $\mathcal{S}^{(r)}(\mathcal{G})$ more concrete, we have illustrated a sequence $(S_0, S_1, S_2) \in \mathcal{S}^{(2)}(\mathcal{G})$ for a particular choice of graph \mathcal{G} in Figure 3.

III. MAIN RESULTS

In this section, we present our main results. Proofs are delayed until Section V.




Fig. 3. An example of a sequence $(S_0, S_1, S_2) \in \mathcal{S}^{(2)}(\mathcal{G})$ for a particular choice of graph \mathcal{G} .

A. Necessary and Sufficient Conditions for Universal Recovery

First, we provide necessary and sufficient conditions for achieving universal recovery in a network \mathcal{T} . It turns out that these conditions are characterized by a particular set of transmission schedules $\mathcal{R}_r(\mathcal{T})$ which we define as follows:

Definition 5: For a network $\mathcal{T} = \{\mathcal{G}, P_1, \dots, P_n\}$, define the region $\mathcal{R}_r(\mathcal{T}) \subseteq \mathbb{N}^{n \times r}$ to be the set of all transmission schedules $\{b_i^j\}$ satisfying:

$$\sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in S_r} P_i^c \right| \quad (2)$$

for each $(S_0, \dots, S_r) \in \mathcal{S}^{(r)}(\mathcal{G})$.

For a given network \mathcal{T} , the region $\mathcal{R}_r(\mathcal{T})$ characterizes a subset of transmission schedules. While the definition of $\mathcal{R}_r(\mathcal{T})$ may initially appear somewhat heavy-handed, its importance is revealed through the following theorem.

Theorem 1: For a network \mathcal{T} , a transmission schedule $\{b_i^j\}$ permits universal recovery in r rounds of communication if and only if $\{b_i^j\} \in \mathcal{R}_r(\mathcal{T})$.

Theorem 1 reveals that the set of transmission schedules permitting universal recovery is characterized precisely by the region $\mathcal{R}_r(\mathcal{T})$. In fact, given a transmission schedule in $\mathcal{R}_r(\mathcal{T})$, a corresponding coding scheme that achieves universal recovery can be computed in polynomial time using the algorithm in [21] applied to the network coding graph discussed in the proof of Theorem 1. Alternatively, one could employ random linear network coding over a sufficiently large field size [22]. If transmissions are made in a manner consistent with a schedule in $\mathcal{R}_r(\mathcal{T})$, universal recovery will be achieved with high probability.

Thus, the problem of achieving universal recovery with the minimum number of transmissions reduces to solving a combinatorial optimization problem over $\mathcal{R}_r(\mathcal{T})$. As this problem was shown to be NP-hard in Example 3, we do not attempt to solve it in its most general form. Instead, we apply Theorem 1 to obtain surprisingly simple characterizations for several cases of interest.

Before proceeding, we provide a few examples to shed some light on the region $\mathcal{R}_r(\mathcal{T})$. First, we show how the traditional cut-set bounds can be recovered from Theorem 1.

Example 4 (Cut-Set Bounds): Considering the constraint defining $\mathcal{R}_r(\mathcal{T})$ in which the nested subsets that form $\mathcal{S}^{(r)}(\mathcal{G})$ are all identical. That is, $(S, S, \dots, S) \in \mathcal{S}^{(r)}(\mathcal{G})$ for some nonempty $S \subset V$. We see that any transmission schedule $\{b_i^j\} \in \mathcal{R}_r(\mathcal{T})$ must satisfy the familiar cut-set bounds:

$$\sum_{j=1}^r \sum_{i \in \partial(S)} b_i^j \geq \left| \bigcap_{i \in S} P_i^c \right|. \quad (3)$$

In words, the total number of packets that flow into the set of nodes S must be greater than or equal to the number of packets that the nodes in S are collectively missing.

As a second example, we apply Theorem 1 to the line network \mathcal{T} considered in Example 1 to find the set of permissible transmission schedules for $r = 2$ rounds of communication. Although it is a tedious application of Theorem 1, we will see that it supplies the desired result.

Example 5: Let \mathcal{T} be the line network considered in Example 1, and let \mathcal{G} be the corresponding graph. Three different choices of $(S_0, S_1, S_2) \in \mathcal{S}^{(2)}(\mathcal{G})$ are given below with the corresponding constraints induced by $\mathcal{R}_2(\mathcal{T})$.

S_0	S_1	S_2	induced constraint
{1}	{1}	{1, 2}	$b_2^1 \geq 1$
{1}	{1, 2}	{1, 2}	$b_3^1 \geq 1$
{3}	{2, 3}	{2, 3}	$b_1^2 \geq 1$

In words, the above constraints imply that nodes 1 and 3 must each transmit once in the first communication round, and node 2 must transmit once in the second communication round. Arguing along the lines of Example 1, these conditions are both necessary and sufficient for a transmission schedule to permit universal recovery in $r = 2$ rounds of communication. Indeed, the reader can check that any $(S_0, S_1, S_2) \in \mathcal{S}^{(2)}(\mathcal{G})$ not listed in the table above induces a constraint via $\mathcal{R}_2(\mathcal{T})$ that is a linear combination of those given above.

B. Fully Connected Networks

When \mathcal{T} is a fully connected network, the graph \mathcal{G} is a complete graph on n vertices. This is perhaps one of the most practically important cases to consider. For example, in a wired computer network, clients can multicast their messages to all other terminals which are cooperatively exchanging data. In wireless networks, broadcast is a natural transmission mode. Indeed, there are protocols tailored specifically to wireless networks which support reliable network-wide broadcast capabilities (cf. [23]–[26]). It is fortunate then, that the cooperative data exchange problem can be solved in polynomial time for fully connected networks:

Theorem 2: For a fully connected network \mathcal{T} , a transmission schedule requiring only $M^*(\mathcal{T})$ transmissions can be computed in polynomial time. Necessary and sufficient conditions for universal recovery in this case are given by the cut-set constraints (3). Moreover, a single round of communication is sufficient to achieve universal recovery with $M^*(\mathcal{T})$ transmissions.

For the fully connected network in Example 2, we remarked that only one round of transmission was required.

Theorem 2 states that this property extends to any fully connected network. Moreover, an optimal transmission schedule requiring $M^*(\mathcal{T})$ transmissions can be computed using the algorithm derived in Appendix A.

Although Theorem 2 applies to arbitrary sets of packets P_1, \dots, P_n , it is insightful to consider the case where packets are randomly distributed in the network according to (1). In this case, the minimum number of transmissions required for universal recovery converges in probability to a simple function of the (random) sets P_1, \dots, P_n .

Theorem 3: If \mathcal{T} is a fully connected network and packets are randomly distributed, then

$$M^*(\mathcal{T}) = \left\lceil \frac{1}{n-1} \sum_{i=1}^n |P_i^c| \right\rceil. \quad (4)$$

with probability approaching 1 as the number of packets $k \rightarrow \infty$.

We remark that, if packets are randomly distributed according to (1), then $M^*(\mathcal{T})$ is a random variable. Hence, Theorem 3 shows that this random variable agrees with $\left\lceil \frac{1}{n-1} \sum_{i=1}^n |P_i^c| \right\rceil$ with high probability in k .

At first glance, it appears that the result of Theorem 3 is independent of the parameter q , which dictates how random packets are distributed in (1). However, q is implicitly present in the result since it parameterizes the distribution of the random variables $|P_i^c|$.

C. Networks Which Are d -Regular

Given that precise results can be obtained for fully connected networks, it is natural to ask whether these results can be extended to a larger class of networks which includes fully connected networks as a special case. In this section, we partially answer this question in the affirmative. To this end, we define d -regular networks.

Definition 6 (d -Regular Networks): A network \mathcal{T} is said to be d -regular if $\partial(i) = d$ for each $i \in V$ and $\partial(S) \geq d$ for each nonempty $S \subset V$ with $|S| \leq n - d$. In other words, a network \mathcal{T} is d -regular if the associated graph \mathcal{G} is d -regular and d -vertex-connected.

Immediately, we see that the class of d -regular networks includes fully connected networks as a special case with $d = n - 1$. Further, the class of d -regular networks includes many frequently studied network topologies (e.g., cycles, grids on tori, etc.).

Unfortunately, the deterministic algorithm of Theorem 2 does not appear to extend to d -regular networks. However, a slightly weaker concentration result similar to Theorem 3 can be obtained when packets are randomly distributed. Before stating this result, consider the following Linear Program (LP) with variable vector $x \in \mathbb{R}^n$ defined for a network \mathcal{T} :

$$\text{minimize } \sum_{i=1}^n x_i \quad (5)$$

$$\text{subject to: } \sum_{i \in \partial(j)} x_i \geq |P_j^c| \quad \text{for each } j \in V. \quad (6)$$

Let $M_{LP}(\mathcal{T})$ denote the optimal value of this LP. Interpreting x_i as $\sum_j b_i^j$, the constraints in the LP are a subset of the cut-set constraints of (3) which are a subset of the necessary constraints for universal recovery given in Theorem 1. Furthermore, the integer constraints on the x_i 's are relaxed. Thus $M_{LP}(\mathcal{T})$ certainly bounds $M^*(\mathcal{T})$ from below. Surprisingly, if \mathcal{T} is a d -regular network and the packets are randomly distributed, $M^*(\mathcal{T})$ is very close to this lower bound with high probability:

Theorem 4: If \mathcal{T} is a d -regular network and the packets are randomly distributed, then

$$M^*(\mathcal{T}) < M_{LP}(\mathcal{T}) + n \quad (7)$$

with probability approaching 1 as the number of packets $k \rightarrow \infty$.

We make two important observations. First, $M^*(\mathcal{T}) \in [M_{LP}(\mathcal{T}), M_{LP}(\mathcal{T}) + n]$ with high probability. Hence, even though the number of packets k may be extremely large, $M^*(\mathcal{T})$ can be estimated accurately by computing the solution to LP (5-6) for a given instantiation of the packet distribution $\{P_i\}_{i=1}^n$. Second, as k grows large, $M^*(\mathcal{T})$ is dominated by the local topology of \mathcal{T} . This is readily seen since the constraints defining $M_{LP}(\mathcal{T})$ correspond only to nodes' immediate neighborhoods. The importance of the local neighborhood was also seen in [27] where network coding capacity for certain random networks is shown to concentrate around the expected number of nearest neighbors of the source and the terminals.

D. Large (Divisible) Packets

We now return to general networks with arbitrarily distributed packets. However, we now consider the case where packets are "large" and can be divided into several smaller pieces (e.g., packets actually correspond to large files). To formalize this, assume that each packet can be partitioned into t chunks of equal size, and transmissions can consist of a single chunk (as opposed to an entire packet). In this case, we say the packets are t -divisible. To illustrate this point more clearly, we return to Example 2, this time considering 2-divisible packets.

Example 6 (2-Divisible Packets): Let \mathcal{T} be the network of Example 2 and split each packet into two halves: $p_i \rightarrow (p_i^{(1)}, p_i^{(2)})$. Denote this new network \mathcal{T}' with corresponding sets of packets:

$$P'_i = \{p_1^{(1)}, p_1^{(2)}, p_2^{(1)}, p_2^{(2)}, p_3^{(1)}, p_3^{(2)}\} \setminus \{p_i^{(1)}, p_i^{(2)}\}. \quad (8)$$

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 $p_1^{(2)} \oplus p_2^{(1)}$. It is readily verified from (3) that 3 chunk-transmissions are required to permit universal recovery. Thus, $M^*(\mathcal{T}') = 3$. Hence, if we were allowed to split the packets of Example 2 into two halves, it would suffice to transmit 3 chunks. Normalizing the number of transmissions by the number of chunks per packet, we say that universal recovery can be achieved with 1.5 packet transmissions.

Motivated by this example, define $M_t^*(\mathcal{T})$ to be the minimum number of (normalized) packet-transmissions required to achieve universal recovery in the network \mathcal{T} when packets are t -divisible. For the network \mathcal{T} in Example 2, we saw above that $M_2^*(\mathcal{T}) = 1.5$.

It turns out, if packets are t -divisible and t is large, the cut-set bounds (3) are "nearly sufficient" for achieving universal recovery. To see this, let $M_{\text{cut-set}}(\mathcal{T})$ be the optimal value of the LP:

$$\text{minimize } \sum_{i=1}^n x_i \quad (9)$$

$$\text{subject to: } \sum_{i \in \partial(S)} x_i \geq \left| \bigcap_{i \in S} P_i^c \right| \quad \text{for all } \emptyset \subsetneq S \subset V. \quad (10)$$

Clearly $M_{\text{cut-set}}(\mathcal{T}) \leq M_t^*(\mathcal{T})$ for any network \mathcal{T} with t -divisible packets because the LP producing $M_{\text{cut-set}}(\mathcal{T})$ relaxes the integer constraints and is constrained only by (3) rather than the full set of constraints given in Theorem 1. However, there exist transmission schedules which can approach this lower bound. Stated more precisely:

Theorem 5: For any network \mathcal{T} , the minimum number of (normalized) packet-transmissions required to achieve universal recovery with t -divisible packets satisfies

$$\lim_{t \rightarrow \infty} M_t^*(\mathcal{T}) = M_{\text{cut-set}}(\mathcal{T}). \quad (11)$$

Precisely how large t is required to be in order to approach $M_{\text{cut-set}}(\mathcal{T})$ within a specified tolerance is not clear for general networks. Since the inequalities (10) have integer coefficients, Cramer's rule and Hadamard's inequality imply that $t \leq n^{n/2}$ is sufficient. However, we should anticipate a much better estimate in general since an immediate consequence of Theorem 3 is that $t = n-1$ is sufficient to achieve $M_{\text{cut-set}}(\mathcal{T})$ with high probability when packets are randomly distributed in a fully connected network.

We note that [18] presents an algorithm exploiting submodularity that solves LP (9) in the special case where \mathcal{T} is a fully connected network. This is in contrast to the algorithm we present in Appendix VI (originally appearing in [3]), which solves the ILP version of this problem corresponding to indivisible packets.

Finally, we remark that it is a simple exercise to construct examples where the cut-set bounds alone are not sufficient to characterize transmission schedules permitting universal recovery when packets are not divisible (e.g., a 4-node line network with packets p_1 and p_2 at the left-most and right-most nodes, respectively). Thus, t -divisibility of packets provides the additional degrees of freedom necessary to approach the cut-set bounds more closely.

E. Remarks

One interesting consequence of our results is that splitting packets does not significantly reduce the required number of packet-transmissions for many scenarios. Indeed, at most one transmission can be saved if the network is fully connected (under any distribution of packets). If the network is d -regular, we can expect to save fewer than n transmissions

if packets are randomly distributed (in fact, at most one transmission per node). It seems possible that this result could be strengthened to include arbitrary distributions of packets in d -regular networks (as opposed to randomly distributed packets), but a proof has not been found.

The limited value of dividing packets has practical ramifications since there is usually some additional communication overhead associated with dividing packets (*e.g.* additional headers, *etc.* for each transmitted chunk are required). Thus, if the packets are very large, say each packet is a video file, our results imply that entire coded packets can be transmitted without significant loss, avoiding any additional overhead incurred by dividing packets.

IV. AN APPLICATION: SECRECY GENERATION

In this section, we consider the setup of the cooperative data exchange problem for a fully connected network \mathcal{T} , but we consider a different goal. In particular, we wish to generate a secret-key among the nodes that cannot be derived by an eavesdropper privy to all of the transmissions among nodes. Also, like the nodes themselves, the eavesdropper is assumed to know the indices of the packets initially available to each node and the communication protocol which the nodes follow. In particular, the only information unknown to the eavesdropper is the contents of the nodes' packets. The goal is for the nodes to generate the maximum amount of "secrecy" that cannot be determined by the eavesdropper.

The theory behind secrecy generation among multiple terminals was originally established in [14] for a very general class of problems in the usual information-theoretic, asymptotic setting. Our results should be interpreted as a practical application of the theory originally developed in [14]. Indeed, our proofs essentially follow those in [14], but have been streamlined and adapted to deal with the non-asymptotic and combinatorial scenario under consideration. The aim of the present section is to show how secrecy can be generated in a practical scenario. By *practical*, we mean that it is possible to generate the maximum amount of secrecy (as defined in [14]) among nodes in a fully connected network $\mathcal{T} = \{\mathcal{G}, P_1, \dots, P_n\}$ using a polynomial time algorithm. Moreover, this is possible in the non-asymptotic regime (*i.e.*, there are no ϵ 's and we don't require the number of packets or nodes to grow arbitrarily large), and it is possible to generate perfect secrecy instead of ϵ -secrecy without any sacrifice.

A. Practical Secrecy Results

In this subsection, we state two results on secrecy generation in fully connected networks \mathcal{T} . Proofs are again postponed until Section V. We begin with some definitions⁴. Let \mathbf{F} denote the set of all transmissions (all of which are available to the eavesdropper by definition). A function K of the packets $\{p_1, \dots, p_k\}$ in the network is called a secret key (SK) if K

is recoverable by all nodes after observing \mathbf{F} , and it satisfies the (perfect) secrecy condition

$$I(K; \mathbf{F}) = 0, \quad (12)$$

and the uniformity condition

$$\Pr(K = key) = \frac{1}{|\mathcal{K}|} \quad \text{for all } key \in \mathcal{K}, \quad (13)$$

where \mathcal{K} is the alphabet of possible keys.

We define $C_{SK}(P_1, \dots, P_n)$ to be the secret-key capacity for a particular distribution of packets. We will drop the notational dependence on P_1, \dots, P_n where it doesn't cause confusion. By this we mean that a secret-key K can be generated if and only if $\mathcal{K} = \mathbb{F}^{C_{SK}}$. In other words, the nodes can generate at most C_{SK} packets worth of secret-key. Implicit in this definition is the requirement that C_{SK} is an integer. That is, we are interested in the maximum *integer* number of packets of secret key which can be generated. Our first result of this section is the following analogue of [14, Theorem 1]:

Theorem 6: The secret-key capacity is given by: $C_{SK}(P_1, \dots, P_n) = k - M^*(\mathcal{T})$.

Next, consider the related problem where a subset $D \subset V$ of nodes is compromised, and assume all nodes are aware of which nodes are compromised (*i.e.*, D is known to all nodes). In this problem, the eavesdropper has access to \mathbf{F} and P_i for $i \in D$. In this case, the secret-key should also be kept hidden from the nodes in D (or else the eavesdropper could also recover it). Thus, for a subset of nodes D , let $P_D = \bigcup_{i \in D} P_i$, and call K a private-key (PK) if it is a secret-key which is only recoverable by the nodes in $V \setminus D$, and also satisfies the stronger secrecy condition:

$$I(K; \mathbf{F}, P_D) = 0. \quad (14)$$

Similar to above, define $C_{PK}(P_1, \dots, P_n, D)$ to be the private-key capacity for a particular distribution of packets and subset of nodes D . Again, we mean that a private-key K can be generated if and only if $\mathcal{K} = \mathbb{F}^{C_{PK}}$. In other words, the nodes in $V \setminus D$ can generate at most C_{PK} packets worth of private-key. As with C_{SK} , we require that C_{PK} is an integer. That is, we are interested in the maximum integer number of packets of private key which can be generated. Note that, since P_D is known to the eavesdropper, each node $i \in D$ can transmit its respective set of packets P_i without any loss of secrecy capacity.

Define a new network $\mathcal{T}_D = \{\mathcal{G}_D, \{P_i^{(D)}\}_{i \in V \setminus D}\}$ as follows. Let \mathcal{G}_D be the complete graph on $V \setminus D$, and let $P_i^{(D)} = P_i \setminus P_D$ for each $i \in V \setminus D$. Thus, \mathcal{T}_D is a fully connected network with $n - |D|$ nodes and $k - |P_D|$ packets. Our second result of this section is the following analogue of [14, Theorem 2]:

Theorem 7: The private-key capacity is given by:

$$C_{PK}(P_1, \dots, P_n, D) = (k - |P_D|) - M^*(\mathcal{T}_D). \quad (15)$$

The basic idea for private-key generation is that the users in $V \setminus D$ should generate a secret-key from $\{p_1, \dots, p_k\} \setminus P_D$.

By the definitions of the SK and PK capacities, Theorem 2 implies that it is possible to compute these capacities efficiently (*i.e.*, in polynomial time). Moreover, as we will see in the achievability proofs, these capacities can be achieved

⁴We attempt to follow the notation of [14] where appropriate.

by performing coded cooperative data exchange amongst the nodes. Thus, the algorithm developed in Appendix A combined with the algorithm in [21] can be employed to solve the secrecy generation problem we consider in polynomial time.

Remark 1: If one drops the integrality requirement for C_{SK} and C_{PK} , this amounts to determining the maximum (possibly fractional) number of packets of secret/private key that can be generated. In this case, we simply replace $M^*(\mathcal{T})$ with $M_{LP}(\mathcal{T})$ in Theorem 6 and $M^*(\mathcal{T}_D)$ with $M_{LP}(\mathcal{T}_D)$ in Theorem 7. These fractional secret- and private-key capacities can be approached as the field size $|\mathbb{F}|$ grows large. The proofs are nearly identical to those of Theorems 6 and 7 (splitting packets where appropriate), and are therefore omitted.

We conclude this subsection with an example to illustrate the results.

Example 7: Consider again the network of Example 2 and assume $\mathbb{F} = \{0, 1\}$ (i.e., each packet is a single bit). The secret-key capacity for this network is 1 bit. After Nodes 1-3 perform universal recovery as described in Example 2, an eavesdropper would observe p_2 and the parity $p_1 \oplus p_3$. A perfect secret-key is $K = p_1$ (we could alternatively use $K = p_3$). If any of the nodes are compromised by the eavesdropper, the private-key capacity is 0.

We remark that the secret-key in the above example can in fact be attained by all nodes using only one transmission (i.e., universal recovery is not a prerequisite for secret-key generation). However, it remains true that only one bit of secrecy can be generated.

V. PROOFS OF MAIN RESULTS

A. Necessary and Sufficient Conditions for Universal Recovery

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

$$V_{NC} = \{s, u_1, \dots, u_k\} \cup \bigcup_{j=0}^r \{v_1^j, \dots, v_n^j\} \cup \bigcup_{j=1}^r \{w_1^j, \dots, w_n^j\}.$$

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 capacity⁵ from s to u_i .
- If $p_i \in P_j$, 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 $i' \in \Gamma(i)$.

⁵An edge of unit capacity can carry one field element $z \in \mathbb{F}$ per unit time.




Fig. 4. The graph \mathcal{G}^{NC} corresponding to the line network of Example 1. Edges represented by broken lines have infinite capacity. Edges with finite capacities are labeled with the corresponding capacity value.

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. In other words, \mathcal{G}^{NC} is the “time-expanded graph” corresponding to the universal recovery problem (the construction of directed acyclic networks via time-expansion originated in [4] in the context of network coding). If the b_i^j 's are chosen such that the graph \mathcal{G}^{NC} admits a network coding solution which supports a multicast of k units from s to $\{v_1^r, \dots, v_n^r\}$, 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. The graph \mathcal{G}^{NC} corresponding to the line network of Example 1 is given in Figure 4.

We now formally prove the equivalence of the network coding problem on \mathcal{G}^{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. $\mathcal{I}_i^1 = P_i$). 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:

$$\phi_i(\mathcal{I}_i^r, \cup_{i' \in \Gamma(i)} \{f_{i'}^r(\mathcal{I}_{i'}^r)\}) = \{p_1, \dots, p_k\}. \quad (16)$$

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

$$\mathcal{I}_i^{j+1} = \mathcal{I}_i^j \cup \bigcup_{i' \in \Gamma(i)} \{f_{i'}^j(\mathcal{I}_{i'}^j)\}. \quad (17)$$

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 $\{v_1^r, \dots, v_n^r\}$ on \mathcal{G}^{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 $\text{IN}(v_i^0) = P_i = \mathcal{I}_i^1$. 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 $\text{IN}(v_i^r) = (\mathcal{I}_i^r, \cup_{i' \in \Gamma(i)} \{f_{i'}^r(\mathcal{I}_{i'}^r)\})$. 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 $\{v_1^r, \dots, v_n^r\}$ on \mathcal{G}^{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}^{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}^{NC} has capacity at least k .

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

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

Define the set $S_0 \subseteq [n]$ as follows: $i \in S_0$ iff $v_i^r \in S$ (by definition of S , we have that $S_0 \neq \emptyset$).

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

- M1) If $i \in \Gamma(S_0)$, then place w_i^r into S' .
- M2) If $i \notin \Gamma(S_0)$, then place w_i^r into S'^c .

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

- J1) If $i \in \Gamma(S_0)$, then there exists an edge of infinite capacity from w_i^r to some $v_{i'}^r \in S$. Thus, moving w_i^r to S' (if necessary) does not increase the capacity of the cut.
- J2) If $i \notin \Gamma(S_0)$, then there are no edges from w_i^r to S , hence we can move w_i^r into S'^c (if necessary) without increasing the capacity of the cut.

Modifications M1 and M2 guarantee that $w_i^r \in S'$ iff $i \in \Gamma(S_0)$. Thus, assume that (S', S'^c) satisfies this condition and further modify the cut as follows:

- M3) If $i \in S_0$, then place v_i^{r-1} into S' .
- M4) If $i \notin \Gamma(S_0)$, then place v_i^{r-1} into S'^c .

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

- J3) If $i \in S_0$, then there exists an edge of infinite capacity from v_i^{r-1} to $v_i^r \in S$. Thus, moving v_i^{r-1} to S' (if necessary) does not increase the capacity of the cut.
- J4) If $i \notin \Gamma(S_0)$, then there are no edges from v_i^{r-1} to S' (since $w_i^r \notin S'$ by assumption), hence we can move v_i^{r-1} into S'^c (if necessary) without increasing the capacity of the cut.

At this point, define the set $S_1 \subseteq [n]$ as follows: $i \in S_1$ iff $v_i^{r-1} \in S'$. Note that the modifications of S' guarantee that S_1 satisfies $S_0 \subseteq S_1 \subseteq \Gamma(S_0)$.

This procedure can be repeated for each layer of the graph resulting in a sequence of sets $\emptyset \subsetneq S_0 \subseteq \dots \subseteq S_r \subseteq [n]$ satisfying $S_j \subseteq \Gamma(S_{j-1})$ for each $j \in [r]$.

We now perform a final modification of the cut (S', S'^c) :

- M5) If $p_j \in \cup_{i \in S_r} P_i$, then place u_j into S' .
- M6) If $p_j \notin \cup_{i \in S_r} P_i$, then place u_j into S'^c .

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

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

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

$$\sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i^{r+1-j} \geq \left| \bigcap_{i \in S_r} P_i^c \right|. \quad (18)$$

Since every modification of the cut either preserved or reduced the capacity of the cut, the original cut (S, S^c) also has capacity greater than or equal to k if the above inequality is satisfied. In Figure 5, we illustrate a cut (S, S^c) and its modified minimal cut (S', S'^c) for the graph \mathcal{G}^{NC} corresponding to the line network of Example 1.

By the equivalence of the universal recovery problem on a network \mathcal{T} to the network coding problem on \mathcal{G}^{NC} and the max-flow min-cut theorem for network information flow, if a transmission strategy solves the universal recovery problem on \mathcal{T} , then the associated b_i^j 's must satisfy the constraints




Fig. 5. The graph \mathcal{G}^{NC} corresponding to the line network of Example 1 with original cut (S, S^c) and the corresponding modified minimal cut (S', S'^c) . In this case, $S_0 = S_1 = S_2 = \{1\}$. Upon substitution into (18), this choice of S_0, S_1, S_2 yields the inequality $b_2^1 + b_2^2 \geq 1$.

of the form given by (18). Conversely, for any set of b_i^j 's which satisfy the constraints of the form given by (18), there exists a transmission strategy using exactly those numbers of transmissions which solves the universal recovery problem for \mathcal{T} . Thus the constraints of (18), and hence the inequalities defining $\mathcal{R}_r(\mathcal{T})$, are satisfied if and only if any cut separating s from some v_i^r in \mathcal{G}^{NC} has capacity at least k .

Remark 2: Since $|\bigcap_{i \in [n]} P_i^c| = 0$, constraints where $S_r = [n]$ are trivially satisfied. Therefore, we can restrict our attention to sequences of sets where $S_r \subsetneq [n]$.

Remark 3: Our problem formulation does not require linear encoding schemes, however the proof above reveals that linear encoding schemes are sufficient to achieve universal recovery with the minimum number of transmissions. ■

B. Fully Connected Networks

Having established necessary and sufficient conditions for universal recovery in Theorem 1, we specialize these

conditions to the setting of a fully connected network to establish Theorems 2 and 3.

Proof of Theorem 2: In the case where \mathcal{T} is a fully connected network, we have that $S_j^c \cap \Gamma(S_{j-1}) = S_j^c$ for any nonempty $S \subset V$. Therefore, the constraints defining $\mathcal{R}_r(\mathcal{T})$ become:

$$\sum_{j=1}^r \sum_{i \in S_j^c} b_i^{r+1-j} \geq \left| \bigcap_{i \in S_r} P_i^c \right|. \quad (19)$$

Now, suppose a transmission schedule $\{b_i^j\} \in \mathcal{R}_r(\mathcal{T})$ and consider the modified transmission schedule $\{\tilde{b}_i^j\}$ defined by: $\tilde{b}_i^r = \sum_{j=1}^r b_i^j$ and $\tilde{b}_i^j = 0$ for $j < r$. By construction, $S_{j+1}^c \subseteq S_j^c$ in the constraints defining $\mathcal{R}_r(\mathcal{T})$. Therefore, using the definition of $\{\tilde{b}_i^j\}$, we have:

$$\sum_{i \in S_j^c} \tilde{b}_i^r \geq \sum_{j=1}^r \sum_{i \in S_j^c} b_i^{r+1-j} \geq \left| \bigcap_{i \in S_r} P_i^c \right|. \quad (20)$$

Thus the modified transmission schedule is also in $\mathcal{R}_r(\mathcal{T})$. Since $|\bigcap_{i \in S_1} P_i^c| \geq |\bigcap_{i \in S_r} P_i^c|$, when \mathcal{T} is a fully connected network, it is sufficient to consider constraints of the form:

$$\sum_{i \in S^c} b_i^1 \geq \left| \bigcap_{i \in S} P_i^c \right| \quad \text{for all nonempty } S \subset V. \quad (21)$$

This proves the following: that the cut-set constraints are necessary and sufficient for universal recovery when \mathcal{T} is a fully connected network, and that a single round of communication is sufficient to achieve universal recovery with $M^*(\mathcal{T})$ transmissions (since (21) does not depend on r).

With these results established, an optimal transmission schedule can be obtained by solving the following integer linear program:

$$\begin{aligned} & \text{minimize} \quad \sum_{i=1}^n b_i \\ & \text{subject to:} \quad \sum_{i \in S^c} b_i \geq \left| \bigcap_{i \in S} P_i^c \right| \quad \text{for all nonempty } S \subset V. \end{aligned} \quad (22)$$

In order to accomplish this, we identify $B_i \leftarrow P_i^c$ and set $w_i = 1$ for $i \in [n]$ and apply the submodular algorithm presented in Appendix A. We remark briefly that the algorithm presented in Appendix A is sufficiently general to solve the ILP above with the objective (22) replaced by $\sum_{i=1}^n w_i b_i$, where $\{w_i\}_{i=1}^n$ are nonnegative reals. In other words, we can easily compute transmission schedules which (i) permit universal recovery, and (ii) minimize a *weighted* sum of nodes' transmissions. ■

Now we consider fully connected networks in which packets are randomly distributed according to (1), which is parameterized by q . The proof of Theorem 3 requires the following technical lemma:

Lemma 1: Let $0 < q < 1$ be fixed. For each $\ell \in \{2, \dots, n-1\}$, define

$$\delta_\ell = \binom{n-\ell}{n-1} \times \left(\frac{1-q-(1-q)^n}{(1-q)^\ell - (1-q)^n} \right) - 1. \quad (23)$$

Then, for each $\ell \in \{2, \dots, n-1\}$:

$$\frac{n-\ell}{n-1} = \frac{(1-q)^\ell - (1-q)^n}{1-q - (1-q)^n} (1 + \delta_\ell) \quad (24)$$

with $0 < \delta_2 < \delta_3 \dots < \delta_{n-1}$.

Proof: Applying Jensen's inequality to the strictly convex function $f(x) = (1-q)^x$ using the convex combination $\ell = \theta \cdot 1 + (1-\theta) \cdot n$ yields:

$$\frac{n-\ell}{n-1} > \frac{(1-q)^\ell - (1-q)^n}{1-q - (1-q)^n} \quad (25)$$

for $\ell \in \{2, \dots, n-1\}$. Thus it is readily verified that (24) holds with $\delta_\ell > 0$ by simply substituting in the definition of δ_ℓ . It is a simple calculus exercise to verify that δ_ℓ is strictly increasing in ℓ for $\ell \in \{2, \dots, n-1\}$. Indeed, let the parameter ℓ vary continuously in the interval $(1, n)$, and look at the derivatives

$$\begin{aligned} & \frac{\partial}{\partial \ell} \left\{ \frac{n-\ell}{(1-q)^\ell - (1-q)^n} \right\} \\ &= (1-q)^\ell \left[\log \left(\frac{(1-q)^\ell}{(1-q)^n} \right) - 1 \right] + (1-q)^n \quad (26) \end{aligned}$$

$$\begin{aligned} & \frac{\partial^2}{\partial \ell^2} \left\{ \frac{n-\ell}{(1-q)^\ell - (1-q)^n} \right\} \\ &= (1-q)^\ell \log(1-q) \left[\log \left(\frac{(1-q)^\ell}{(1-q)^n} \right) + 1 \right] < 0. \quad (27) \end{aligned}$$

Hence, we can conclude that (23) is concave in ℓ over the interval $(1, n)$, with a maximum being attained when $\ell \rightarrow n$. ■

The proof of Theorem 3 proceeds by showing that the necessary and sufficient conditions of Theorem 1 admit a simply-structured solution when packets are randomly distributed according to (1).

Proof of Theorem 3: We begin by showing that the LP

$$\text{minimize } \sum_{i=1}^n b_i \quad (28)$$

$$\text{subject to: } \sum_{i \in S^c} b_i \geq \left| \bigcap_{i \in S} P_i^c \right| \quad \text{for all } \emptyset \subsetneq S \subset V. \quad (29)$$

has an optimal value of $\frac{1}{n-1} \sum_{i=1}^n |P_i^c|$ with high probability as $k \rightarrow \infty$, over random instantiations of the P_i 's (recall that packets are randomly distributed according to (1)). To this end, note that the inequalities

$$\sum_{\substack{i=1 \\ i \neq j}}^n b_i \geq |P_j^c| \quad \text{for } 1 \leq j \leq n. \quad (30)$$

are a subset of the inequality constraints (29). Summing both sides of (30) over $1 \leq j \leq n$ reveals that any feasible vector $b \in \mathbb{R}^n$ for LP (28)-(29) must satisfy:

$$\sum_{i=1}^n b_i \geq \frac{1}{n-1} \sum_{i=1}^n |P_i^c|. \quad (31)$$

This establishes a lower bound on the optimal value of the LP. We now identify a solution that is feasible with probability

approaching 1 as $k \rightarrow \infty$ while achieving the lower bound of (31) with equality. To begin note that

$$\tilde{b}_j = \frac{1}{n-1} \sum_{i=1}^n |P_i^c| - |P_j^c| \quad (32)$$

is a solution to the system of linear equations given by (30) and achieves (31) with equality. Now, we prove that $(\tilde{b}_1, \dots, \tilde{b}_n)$ is a feasible solution to LP (28) with high probability. To be specific, we must verify that

$$\sum_{i \in S^c} \tilde{b}_i \geq \left| \bigcap_{i \in S} P_i^c \right| \quad (33)$$

holds with high probability for all subsets $S \subset V$ satisfying $2 \leq |S| \leq n-1$ (the case $|S| = 1$ is satisfied by the definition of $\{\tilde{b}_i\}_{i=1}^n$). Substitution of (32) into (33) along with simple algebra yields that the following equivalent conditions must hold:

$$\left(\frac{n-|S|}{n-1} \right) \sum_{i=1}^n \frac{1}{k} |P_i^c| - \sum_{i \in S^c} \frac{1}{k} |P_i^c| \geq \frac{1}{k} \left| \bigcap_{i \in S} P_i^c \right|. \quad (34)$$

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

$$\zeta_{|S|} \triangleq \Pr(X_j^S = 1) = \frac{(1-q)^{|S|} - (1-q)^n}{1 - (1-q)^n}. \quad (35)$$

For $\ell = 2, 3, \dots, n-1$, define δ_ℓ as in Lemma 1. Now, define

$$v_1 \triangleq \frac{\delta_2}{2n + \delta_2(2n-1)}, \quad (36)$$

and define v_ℓ to satisfy:

$$1 + v_\ell = (1 + \delta_\ell) (1 - v_1(2n-1)) \quad \text{for } \ell = 2, \dots, n-1.$$

By definition, $v_1 > 0$, but also observe that $v_\ell > 0$ for $\ell \geq 2$ since Lemma 1 implies $\delta_\ell \geq \delta_2$, and therefore

$$1 + v_\ell = (1 + \delta_\ell) (1 - v_1(2n-1)) \quad (37)$$

$$\geq (1 + \delta_2) (1 - v_1(2n-1)) \quad (38)$$

$$= 1 + v_1. \quad (39)$$

Now, for notational convenience, define the function $\zeta(n, q)$ of (n, q) as follows:

$$\zeta(n, q) = 2 \times \left(\min_{1 \leq \ell \leq [n-1]} \left\{ (v_\ell \zeta_\ell)^2 \right\} \right). \quad (40)$$

Note that Hoeffding's inequality yields the following concentration result:

$$\Pr \left(\left| \frac{1}{k} \left| \bigcap_{i \in S} P_i^c \right| - \zeta_{|S|} \right| > v_{|S|} \zeta_{|S|} \right) \leq 2 \exp \left[-2k (v_{|S|} \zeta_{|S|})^2 \right] \quad (41)$$

$$\leq 2 \exp \left[-k \cdot \zeta(n, q) \right]. \quad (42)$$

Therefore, by the union of events bound, we have

$$\Pr \left(\bigcup_{S: S \subset [n], S \neq \emptyset} \left\{ \left| \frac{1}{k} \left| \bigcap_{i \in S} P_i^c \right| - \zeta_{|S|} \right| > \nu_{|S|} \zeta_{|S|} \right\} \right) \leq 2^{n+1} \exp[-k \cdot \zeta(n, q)]. \quad (43)$$

Thus, with probability greater than $1 - 2^{n+1} \exp[-k \cdot \zeta(n, q)]$, the following inequalities hold for $S \subset V$ with $2 \leq |S| \leq n-1$:

$$\left(\frac{n - |S|}{n - 1} \right) \sum_{i=1}^n \frac{1}{k} |P_i^c| - \sum_{i \in S^c} \frac{1}{k} |P_i^c| \quad (44)$$

$$\geq \left(\frac{n - |S|}{n - 1} \right) n \zeta_1 (1 - \nu_1) - \left(\frac{n - |S|}{n - 1} \right) (n - 1) \zeta_1 (1 + \nu_1) \quad (45)$$

$$= \left(\frac{n - |S|}{n - 1} \right) \left(\frac{(1 - q) - (1 - q)^n}{1 - (1 - q)^n} \right) (1 - \nu_1 (2n - 1)) \quad (46)$$

$$= \left(\frac{(1 - q)^{|S|} - (1 - q)^n}{1 - (1 - q)^n} \right) (1 + \delta_{|S|}) (1 - \nu_1 (2n - 1)) \quad (47)$$

$$= \left(\frac{(1 - q)^{|S|} - (1 - q)^n}{1 - (1 - q)^n} \right) (1 + \nu_{|S|}) \quad (48)$$

$$= \zeta_{|S|} (1 + \nu_{|S|}) \quad (49)$$

$$\geq \frac{1}{k} \left| \bigcap_{i \in S} P_i^c \right|, \quad (50)$$

where

- (45) follows from (43).
- (46) follows from elementary algebra and the definition of ζ_1 .
- (47)-(49) follow from the definitions of $\delta_{|S|}$, $\nu_{|S|}$ and $\zeta_{|S|}$, respectively.
- (50) follows from (43).

This proves that (34) holds, and therefore $(\tilde{b}_1, \dots, \tilde{b}_n)$ is a feasible solution to LP (28) with probability greater than $1 - 2^{n+1} \exp[-k \cdot \zeta(n, q)]$.

Finally, Corollary 1 in Appendix A states that the optimal values of the ILP:

$$\min \left\{ \sum_{i=1}^n b_i : \sum_{i \in S^c} b_i \geq \left| \bigcap_{i \in S} P_i^c \right|, S \subset V, S \neq \emptyset, b_i \in \mathbb{Z} \right\}$$

and the corresponding LP relaxation:

$$\min \left\{ \sum_{i=1}^n b_i : \sum_{i \in S^c} b_i \geq \left| \bigcap_{i \in S} P_i^c \right|, S \subset V, S \neq \emptyset, b_i \in \mathbb{R} \right\}$$

differ by less than 1. Combining this fact with Theorem 2 completes the proof. \blacksquare

Remark 4: Observe that we have actually proven a stronger concentration result than was needed to prove Theorem 3. Indeed, we have shown that if packets are randomly distributed according to (1), then

$$M^*(T) = \left[\frac{1}{n-1} \sum_{i=1}^n |P_i^c| \right] \quad (51)$$

with probability greater than $1 - 2^{n+1} \exp[-k \cdot \zeta(n, q)]$, where $\zeta(n, q)$ is an explicit function of n and q .

C. d -Regular Networks

Before we begin, let us define $\|\cdot\|_2$, $\|\cdot\|_\infty$ to be the usual ℓ_2 - and ℓ_∞ -norms, respectively. Also, let $\mathbf{1}$ denote the column vector of all ones. In order to simplify the proof of Theorem 4, we require two technical lemmas:

Lemma 2: Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with nonnegative entries and all column sums equal to d . Let \bar{x}_y be the vector of minimum Euclidean norm which minimizes $\|Ax_y - y\|_2$. There exists an optimal solution x^* to the linear program

$$\text{minimize } \mathbf{1}^T x \quad (52)$$

$$\text{subject to: } Ax \geq y \quad (53)$$

which satisfies

$$\|x^* - \bar{x}_y\|_\infty \leq c_A \|A\bar{x}_y - y\|_2, \quad (54)$$

where c_A is a constant depending only on A .

Proof: See Appendix B. \blacksquare

Lemma 3: Assume packets are randomly distributed in a d -regular network \mathcal{T} according to (1). For any $\epsilon > 0$, there exists an optimal solution x^* to LP (5-6) which satisfies

$$\left\| x^* - \frac{1}{d} \mathbb{E}[|P_i^c|] \mathbf{1} \right\|_\infty < \epsilon k \quad (55)$$

with probability approaching 1 as $k \rightarrow \infty$, where \mathbb{E} indicates expectation.

Before proceeding with the proof of Lemma 3, we recall that the random packet model (1) implies that $\mathbb{E}[|P_j^c|] = \mathbb{E}[|P_1^c|]$ for all $j \in [n]$. Hence, there is nothing noteworthy about the appearance of P_1^c in the statement of the lemma, any P_j^c would suffice.

Proof: Let $\vec{P} = (|P_1^c|, \dots, |P_n^c|)^T$ and let A be the adjacency matrix of \mathcal{G} (i.e., $a_{i,j} = 1$ if $(i, j) \in E$ and 0 otherwise). Observe that A is symmetric and $A\mathbf{1} = d\mathbf{1}$, where $\mathbf{1}$ denotes a column vector of 1's. With this notation, LP (5) can be rewritten as:

$$\text{minimize } \mathbf{1}^T x \quad (56)$$

$$\text{subject to: } Ax \geq \vec{P}, \quad (57)$$

where " $a \geq b$ " for vectors $a, b \in \mathbb{R}^n$ means that $a_i \geq b_i$ for $i = 1, \dots, n$.

Let A^+ denote the Moore-Penrose pseudoinverse of A . Observe that the linear least squares solution $\bar{x}_{LS} \triangleq \arg \min_x \|Ax - \vec{P}\|_2$ is given by:

$$\bar{x}_{LS} = A^+ \vec{P} \quad (58)$$

$$= A^+ \mathbb{E} \vec{P} + A^+ (\vec{P} - \mathbb{E} \vec{P}) \quad (59)$$

$$= \frac{1}{d} \mathbb{E} \vec{P} + A^+ (\vec{P} - \mathbb{E} \vec{P}). \quad (60)$$

For the last step above, note that $\mathbb{E} \vec{P}$ is an eigenvector of A with eigenvalue d so $\mathbb{E} \vec{P}$ will also be an eigenvector of A^+ with eigenvalue $\frac{1}{d}$. Hence,

$$\|\bar{x}_{LS} - \frac{1}{d} \mathbb{E} \vec{P}\|_2 = \|A^+ (\vec{P} - \mathbb{E} \vec{P})\|_2 \quad (61)$$

$$\leq \|A^+\|_2 \|\vec{P} - \mathbb{E} \vec{P}\|_2, \quad (62)$$

where $\|A^+\|_2$ is the operator norm of A^+ induced by the 2-norm on vectors (also known as the spectral norm of A^+). Thus, for any vector y , we have

$$\|y - \frac{1}{d}\mathbb{E}\vec{P}\|_\infty \leq \|y - \bar{x}_{LS}\|_\infty + \|\bar{x}_{LS} - \frac{1}{d}\mathbb{E}\vec{P}\|_\infty \quad (63)$$

$$\leq \|y - \bar{x}_{LS}\|_\infty + \|\bar{x}_{LS} - \frac{1}{d}\mathbb{E}\vec{P}\|_2 \quad (64)$$

$$\leq \|y - \bar{x}_{LS}\|_\infty + \|A^+\|_2 \|\vec{P} - \mathbb{E}\vec{P}\|_2, \quad (65)$$

where (63) follows from the triangle inequality, (64) follows from the definitions of the ℓ_2 - and ℓ_∞ -norms, and (65) follows from (62).

Now, Lemma 2 guarantees the existence of an optimal solution x^* to LP (56) (and consequently LP (5)) which satisfies $\|x^* - \bar{x}_{LS}\|_\infty \leq c_A \|A\bar{x}_{LS} - \vec{P}\|_2$, where c_A is a constant depending only on A . Hence, letting $y = x^*$ in (65), we have:

$$\begin{aligned} \|x^* - \frac{1}{d}\mathbb{E}\vec{P}\|_\infty &\leq \|x^* - \bar{x}_{LS}\|_\infty + \|A^+\|_2 \|\vec{P} - \mathbb{E}\vec{P}\|_2 \\ &\leq c_A \|A\bar{x}_{LS} - \vec{P}\|_2 + \|A^+\|_2 \|\vec{P} - \mathbb{E}\vec{P}\|_2 \\ &\leq c_A \frac{1}{d} \|A\mathbb{E}\vec{P} - \vec{P}\|_2 + \|A^+\|_2 \|\vec{P} - \mathbb{E}\vec{P}\|_2 \\ &= c_A \|\mathbb{E}\vec{P} - \vec{P}\|_2 + \|A^+\|_2 \|\vec{P} - \mathbb{E}\vec{P}\|_2, \end{aligned}$$

where the third inequality follows by definition of \bar{x}_{LS} . By the weak law of large numbers, $\|\vec{P} - \mathbb{E}\vec{P}\|_2 \leq \epsilon k$ with probability tending to 1 as $k \rightarrow \infty$ for any $\epsilon > 0$. Noting that $\mathbb{E}\vec{P} = \mathbb{E}[\mathbb{1}[P_1^c]]\mathbb{1}$ completes the proof. ■

Now, we are in a position to prove Theorem 4. Roughly speaking, the proof of Theorem 4 is similar to that of 3. In particular, we show that if packets are randomly distributed according to (1), then a simply-structured transmission schedule satisfies the conditions of Theorem 1. However, care must be taken to handle the more general network topologies.

Proof of Theorem 4: We begin with some observations and definitions:

- First, recall that our model for randomly distributed packets (1) implies that

$$\mathbb{E} \left[\prod_{i \in S} P_i^c \right] = k \frac{(1-q)^{|S|} - (1-q)^n}{1 - (1-q)^n} \quad (66)$$

for all nonempty $S \subset V$.

- Applying the identity (66), we have

$$\mathbb{E} [\mathbb{1}[P_1^c]] \geq (1+c_q) \mathbb{E} \left[\prod_{i \in S} P_i^c \right] \quad (67)$$

for all $S \subset V$, $|S| \geq 2$, where $c_q > 0$ is defined by

$$c_q \triangleq \frac{q(1-q)}{2((1-q)^2 - (1-q)^n)}. \quad (68)$$

- Next, Lemma 1 implies the existence of a constant $\delta_q > 0$ such that for any $S \subset V$ with $2 \leq |S| \leq n-1$:

$$\frac{n-|S|}{n-1} \geq \frac{(1-q)^{|S|} - (1-q)^n}{(1-q) - (1-q)^n} + \delta_q \quad (69)$$

$$= \frac{\mathbb{E} [\mathbb{1}[\prod_{i \in S} P_i^c]]}{\mathbb{E} [\mathbb{1}[P_1^c]]} + \delta_q, \quad (70)$$

where the equality follows from (66).

- Arguing in a manner similar to the discussion surrounding (43), it is readily verified that Hoeffding's inequality⁶ implies

$$\left(1 + \frac{\min\{\delta_q, c_q\}}{4}\right) \mathbb{E} \left[\prod_{i \in S} P_i^c \right] \geq \left| \prod_{i \in S} P_i^c \right| \quad (71)$$

with probability approaching 1 as $k \rightarrow \infty$.

- Finally, for the proof below, we will take the number of communication rounds sufficiently large to satisfy

$$r \geq \max \left\{ \frac{2n}{d\delta_q}, \frac{2n(1+c_q)}{dc_q} \right\}. \quad (72)$$

Fix $\epsilon > 0$. Lemma 3 guarantees that there exists an optimal solution x^* to LP (5) satisfying

$$\left\| x^* - \frac{1}{d} \mathbb{E}[\mathbb{1}[P_1^c]] \mathbb{1} \right\|_\infty < \epsilon k \quad (73)$$

with probability tending to 1 in k . Now, it is always possible to construct a transmission schedule $\{b_i^j\}$ which satisfies $\sum_j b_i^j = \lceil x_i^* \rceil$ and $\lfloor \frac{1}{r} x_i^* \rfloor \leq b_i^j \leq \lceil \frac{1}{r} x_i^* \rceil$ for each i, j . Observe that $\sum_{i,j} b_i^j < n + \sum_i x_i^*$. Thus, proving that $\{b_i^j\} \in \mathcal{R}_r(\mathcal{T})$ with high probability will prove the theorem.

Since the network is d -regular, $|\partial(S)| \geq d$ whenever $|S| \leq n-d$, and $|\partial(S_1)| \geq n-|S_2|$ whenever $|S_2| \geq n-d$ and $S_1 \subseteq S_2$. To see that the latter point is true, observe the following. If $|S_1| \leq n-d$, then $|\partial(S_1)| \geq d$ and hence $|S_2| \geq n-d$ implies that $|\partial(S_1)| \geq n-|S_2|$. On the other hand, if $|S_1| \geq n-d$, then $|\partial(S_1)| = n-|S_1| \geq n-|S_2|$, where the inequality follows since $S_1 \subseteq S_2$.

We consider the cases where $2 \leq |S_r| \leq n-d$ and $n-d < |S_r| \leq n-1$ separately. The case where $|S_r| = 1$ coincides precisely with the constraints (6), and hence is satisfied by definition of $\{b_i^j\}$.

Considering the case where $2 \leq |S_r| \leq n-d$, we have the following string of inequalities:

$$\begin{aligned} &\sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i^{(r+1-j)} \\ &\geq \sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} \left\lfloor \frac{1}{r} x_i^* \right\rfloor \end{aligned} \quad (74)$$

$$\geq \left(\frac{1}{r} \sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} x_i^* \right) - nr \quad (75)$$

$$\geq \left(\frac{1}{r} \sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} \frac{1}{d} \mathbb{E}[\mathbb{1}[P_1^c]] \right) - nk\epsilon - nr \quad (76)$$

$$\begin{aligned} &= \left(\frac{1}{r} \sum_{j=1}^r \sum_{i \in \partial(S_{j-1})} \frac{1}{d} \mathbb{E}[\mathbb{1}[P_1^c]] \right) \\ &\quad - \left(\frac{1}{r} \sum_{i \in S_r \cap S_0^c} \frac{1}{d} \mathbb{E}[\mathbb{1}[P_1^c]] \right) - nk\epsilon - nr \end{aligned} \quad (77)$$

⁶In fact, the weak law of large numbers is sufficient here, but Hoeffding's inequality can be used to prove exponential concentration if desired.

$$\geq \frac{1}{rd} \mathbb{E}[\mathbb{I}[P_1^c]] \left(\sum_{j=1}^r |\partial(S_{j-1})| - n \right) - nr(k\epsilon + 1) \quad (78)$$

$$\geq \frac{1+c_q}{rd} \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] \left(\sum_{j=1}^r |\partial(S_{j-1})| - n \right) - nr(k\epsilon + 1) \quad (79)$$

$$\geq \frac{1+c_q}{rd} \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] (rd - n) - nr(k\epsilon + 1) \quad (80)$$

$$\geq \left(1 + \frac{c_q}{2}\right) \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] - nr(k\epsilon + 1) \quad (81)$$

$$\geq \left(1 + \frac{c_q}{4}\right) \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] \quad (82)$$

$$\geq \left| \bigcap_{i \in S_r} P_i^c \right|. \quad (83)$$

The above string of inequalities holds with probability tending to 1 as $k \rightarrow \infty$. They can be justified as follows:

- (74) follows by definition of $\{b_i^j\}$.
- (75) follows since $\lfloor \frac{1}{r} x_i^* \rfloor \geq \frac{1}{r} x_i^* - 1$ and $|S_j^c \cap \Gamma(S_{j-1})| \leq n$.
- (76) follows from (73) and the fact that $|S_j^c \cap \Gamma(S_{j-1})| \leq n$.
- (77) follows from writing $\cup_{j=1}^r S_j^c \cap \Gamma(S_{j-1})$ as $(\cup_{j=1}^r \partial(S_{j-1})) \setminus (S_r \cap S_0^c)$ and expanding the sum.
- (78) is true since $|S_0^c \cap S_r| \leq n$.
- (79) follows from (67).
- (80) follows from $|\partial(S_{j-1})| \geq d$ by d regularity and the assumption that $2 \leq |S_r| \leq n - d$.
- (81) follows from our choice of r given in (72). To see this, recall that we have chosen $r \geq \frac{2n(1+c_q)}{dc_q}$, implying the inequality

$$\frac{1+c_q}{rd} (rd - n) \geq 1 + \frac{c_q}{2}. \quad (84)$$

- (82) follows since $\frac{c_q}{4} \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] \geq nr(k\epsilon + 1)$ for ϵ sufficiently small and k sufficiently large. Indeed, this is equivalent to

$$\begin{aligned} \frac{c_q}{4} \mathbb{E} \left[\frac{1}{k} \mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] &= \frac{c_q}{4} \times \frac{(1-q)^{|S_r|} - (1-q)^n}{1 - (1-q)^n} \\ &\geq nr \left(\epsilon + \frac{1}{k} \right). \end{aligned} \quad (85)$$

Since the left hand side of (85) is strictly positive and independent of k, ϵ , the claim follows.

- (83) follows from (71).

Next, consider the case where $n - d \leq |S_r| \leq n - 1$. Starting from (78), we obtain:

$$\begin{aligned} &\sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i^{(r+1-j)} \\ &\geq \frac{1}{rd} \mathbb{E}[\mathbb{I}[P_1^c]] \left(\sum_{j=1}^r |\partial(S_{j-1})| - n \right) - nr(k\epsilon + 1) \end{aligned} \quad (86)$$

$$\geq \frac{1}{rd} \mathbb{E}[\mathbb{I}[P_1^c]] \left(r(n - |S_r|) - n \right) - nr(k\epsilon + 1) \quad (87)$$

$$= \mathbb{E}[\mathbb{I}[P_1^c]] \left(\frac{n - |S_r|}{d} - \frac{n}{rd} \right) - nr(k\epsilon + 1) \quad (88)$$

$$\geq \mathbb{E}[\mathbb{I}[P_1^c]] \left(\frac{n - |S_r|}{n-1} - \frac{n}{rd} \right) - nr(k\epsilon + 1) \quad (89)$$

$$\geq \mathbb{E}[\mathbb{I}[P_1^c]] \left(\frac{\mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right]}{\mathbb{E}[\mathbb{I}[P_1^c]]} + \delta_q - \frac{n}{rd} \right) - nr(k\epsilon + 1) \quad (90)$$

$$\geq \mathbb{E}[\mathbb{I}[P_1^c]] \left(\frac{\mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right]}{\mathbb{E}[\mathbb{I}[P_1^c]]} + \frac{\delta_q}{2} \right) - nr(k\epsilon + 1) \quad (91)$$

$$\geq \left(1 + \frac{\delta_q}{4}\right) \mathbb{E} \left[\mathbb{I} \left[\bigcap_{i \in S_r} P_i^c \right] \right] \quad (92)$$

$$\geq \left| \bigcap_{i \in S_r} P_i^c \right|. \quad (93)$$

The above string of inequalities holds with probability tending to 1 as $k \rightarrow \infty$. They can be justified as follows:

- (86) is simply (78) repeated for convenience.
- (87) follows since $n - d \leq |S_r| \leq n - 1$ and hence d -regularity implies that $|\partial(S_{j-1})| \geq (n - |S_r|)$.
- (89) follows since $d \leq n - 1$.
- (90) follows from (70).
- (91) follows from our definition of r given in (72). To see this, recall that we have chosen $r \geq \frac{2n}{d\delta_q}$, implying the inequality $\delta_q - \frac{n}{rd} \geq \frac{\delta_q}{2}$.
- (92) follows since $\frac{\delta_q}{4} \mathbb{E}[\mathbb{I}[P_1^c]] \geq nr(k\epsilon + 1)$ for ϵ sufficiently small and k sufficiently large. The argument justifying this is similar to (85).
- (93) follows from (71).

Thus, we conclude that, for ϵ sufficiently small, the transmission schedule $\{b_i^j\}$ satisfies each of the inequalities defining $\mathcal{R}_r(\mathcal{T})$ with probability tending to 1. Since the number of such inequalities is finite, an application of the union bound completes the proof that $\{b_i^j\} \in \mathcal{R}_r(\mathcal{T})$ with probability tending to 1 as $k \rightarrow \infty$. ■

Remark 5: As was done in the proof of Theorem 3, it is possible to prove an exponential concentration result for Theorem 4. However, in favor of a simpler presentation, we omit the details and content ourselves with the given asymptotic result.

D. Divisible Packets

The proof of Theorem 5 is relatively straightforward. In particular, given a transmission schedule that satisfies the cutset constraints, we create a corresponding transmission schedule that satisfies the necessary and sufficient conditions of Theorem 1 adapted to t -divisible packets.

Proof of Theorem 5: Fix any $\epsilon > 0$ and let x^* be an optimal solution to LP (9). Put $b_i = x_i^* + \epsilon$. Note that b_i is nonnegative. This follows by considering the set $S \setminus \{i\}$ in the inequality constraint (10), which implies $x_i^* \geq 0$.

Now, take an integer $r \geq \epsilon^{-1}n \max_{1 \leq i \leq n} b_i$. If packets are t -divisible, we can find a transmission schedule $\{b_i^j\}$ such that $\frac{1}{r}b_i \leq b_i^j \leq \frac{1}{r}b_i + \frac{1}{r}$ for all $i \in [n], j \in [r]$.

Thus, for any $(S_0, \dots, S_r) \in \mathcal{S}^{(r)}(\mathcal{G})$ we have the following string of inequalities:

$$\begin{aligned} & \sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i^{(r+1-j)} \\ & \geq \frac{1}{r} \sum_{j=1}^r \sum_{i \in S_j^c \cap \Gamma(S_{j-1})} b_i \end{aligned} \quad (94)$$

$$= \frac{1}{r} \sum_{j=1}^r \sum_{i \in \partial(S_{j-1})} b_i - \frac{1}{r} \sum_{i \in S_0^c \cap S_r} b_i \quad (95)$$

$$\begin{aligned} & = \frac{1}{r} \sum_{j=1}^r \sum_{i \in \partial(S_{j-1})} x_i^* + \frac{\epsilon}{r} \sum_{j=1}^r |\partial(S_{j-1})| \\ & \quad - \frac{1}{r} \sum_{i \in S_0^c \cap S_r} b_i \end{aligned} \quad (96)$$

$$\geq \frac{1}{r} \sum_{j=1}^r \sum_{i \in \partial(S_{j-1})} x_i^* + \epsilon - \frac{n}{r} \max_{1 \leq i \leq n} b_i \quad (97)$$

$$\geq \frac{1}{r} \sum_{j=1}^r \left| \bigcap_{i \in S_{j-1}} P_j^c \right| + \epsilon - \frac{n}{r} \max_{1 \leq i \leq n} b_i \quad (98)$$

$$\geq \left| \bigcap_{i \in S_r} P_j^c \right|, \quad (99)$$

where

- (94) follows since $b_i^j \geq \frac{1}{r}b_i$.
- (95) follows from writing $\cup_{j=1}^r S_j^c \cap \Gamma(S_{j-1})$ as $(\cup_{j=1}^r \partial(S_{j-1})) \setminus (S_r \cap S_0^c)$ and expanding the sum.
- (96) follows from our definition $b_i = x_i^* + \epsilon$.
- (97) follows since $|\partial(S_{j-1})| \geq 1$ and $|S_0^c \cap S_r| \leq n$.
- (98) follows since x^* is an optimal solution to LP (9), and therefore must satisfy the claimed inequality.
- (99) follows from our choice or $r \geq \epsilon^{-1}n \max_{1 \leq i \leq n} b_i$.

Hence, Theorem 1 implies that the transmission schedule $\{b_i^j\}$ is sufficient to achieve universal recovery. Noting that

$$\sum_{i,j} b_i^j \leq \sum_{i=1}^n b_i + \frac{nr}{t} \leq \sum_{i=1}^n x_i^* + n \left(\frac{r}{t} + \epsilon \right) \quad (100)$$

completes the proof of the theorem. \blacksquare

E. Secrecy Generation

In this subsection, we prove Theorems 6 and 7. We again remark that our proofs essentially follow those in [14], but have been adapted for the problem at hand. For notational convenience, define $P = \{p_1, \dots, p_k\}$. We will require the following lemma.

Lemma 4: Given a packet distribution P_1, \dots, P_n , let K be a secret-key achievable with communication \mathbf{F} . Then the

following holds:

$$H(K|\mathbf{F}) = H(P) - \sum_{i=1}^n x_i. \quad (101)$$

for some vector $x = (x_1, \dots, x_n)$ which is feasible for the following ILP:

$$\text{minimize } \sum_{i=1}^n x_i \quad (102)$$

$$\text{subject to: } \sum_{i \in S} x_i \geq \left| \bigcap_{i \in S^c} P_i^c \right| \text{ for all } \emptyset \subsetneq S \subset V. \quad (103)$$

Moreover, if K is a PK (with respect to a set D) and each node $i \in D$ transmits its respective set of packets P_i , then

$$H(K|\mathbf{F}) = H(P|P_D) - \sum_{i \in V \setminus D} x_i. \quad (104)$$

for some vector $x = (x_1, \dots, x_n)$ which is feasible for the ILP:

$$\text{minimize } \sum_{i \in V \setminus D} x_i \quad (105)$$

$$\text{subject to: } \sum_{i \in S} x_i \geq \left| \bigcap_{i \in S^c} P_i^c \right| \quad \forall \emptyset \subsetneq S \subset V \setminus D. \quad (106)$$

Remark 6: We remark that (103) and (106) are necessary and sufficient conditions for achieving universal recovery in the networks \mathcal{T} and \mathcal{T}_D considered in Theorems 6 and 7, respectively. Thus, the optimal values of ILPs (102) and (105) are equal to $M^*(\mathcal{T})$ and $M^*(\mathcal{T}_D)$, respectively.

Proof of Lemma 4: We assume throughout that all entropies are with respect to the base- $|\mathbb{F}|$ logarithm (*i.e.*, information is measured in packets). For this and the following proofs, let $\mathbf{F} = (F_1, \dots, F_n)$ and $F_{[1,i]} = (F_1, \dots, F_i)$, where F_i denotes the transmissions made by node i . For simplicity, our proof does not take into account interactive communication, but can be modified to do so. Allowing interactive communication does not change the results. See [14] for details.

Since K and \mathbf{F} are functions of P :

$$H(P) = H(\mathbf{F}, K, P_1, \dots, P_n) \quad (107)$$

$$\begin{aligned} & = \sum_{i=1}^n H(F_i | F_{[1,i-1]}) + H(K|\mathbf{F}) \\ & \quad + \sum_{i=1}^n H(P_i | \mathbf{F}, K, P_{[1,i-1]}). \end{aligned} \quad (108)$$

Set $x_i = H(F_i | F_{[1,i-1]}) + H(P_i | \mathbf{F}, K, P_{[1,i-1]})$. Then, the substituting x_i into the above equation yields:

$$H(K|\mathbf{F}) = H(P) - \sum_{i=1}^n x_i. \quad (109)$$

Note that $x = (x_1, \dots, x_n)$ is a feasible vector for the linear relaxation of ILP (102), since:

$$\left| \bigcap_{i \in S^c} P_i^c \right| = H(P_S | P_{S^c}) \quad (110)$$

$$= H(\mathbf{F}, K, P_S | P_{S^c}) \quad (111)$$

$$= H(K | \mathbf{F}, P_{S^c}) + \sum_{i=1}^n H(F_i | F_{[1, i-1]}, P_{S^c}) \\ + \sum_{i \in S} H(P_i | \mathbf{F}, K, P_{[1, i-1]}, P_{S^c \cap [i+1, n]}) \quad (112)$$

$$\leq \sum_{i \in S} H(F_i | F_{[1, i-1]}) + \sum_{i \in S} H(P_i | \mathbf{F}, K, P_{[1, i-1]}) \quad (113)$$

$$= \sum_{i \in S} x_i. \quad (114)$$

In the above inequality, we used the fact that conditioning reduces entropy, the fact that K is a function of (\mathbf{F}, P_{S^c}) for any $S \neq V$, and the fact that F_i is a function of P_i (by the assumption that communication is not interactive). Now, since C_{SK} is an integer, (109) implies that $\sum_{i=1}^n x_i$ is an integer. Thus, Corollary 1 in Appendix A (which states that the optimal values of ILP (102) and its linear relaxation differ by strictly less than 1), implies the existence of an integral vector $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$ which is feasible for ILP (102) and satisfies

$$\sum_{i=1}^n \tilde{x}_i = \sum_{i=1}^n x_i. \quad (115)$$

This proves the first part of the lemma.

To prove the second part of the lemma, we can assume $D = \{1, \dots, \ell\}$. The assumption that each node i in D transmits all of the packets in P_i implies $F_i = P_i$. Thus, for $i \in D$ we have $x_i = H(P_i | P_{[1, i-1]})$. Repeating the above argument, we obtain

$$H(K | \mathbf{F}) = H(P) - H(P_D) - \sum_{i \in V \setminus D} x_i \quad (116)$$

$$= H(P | P_D) - \sum_{i \in V \setminus D} x_i. \quad (117)$$

Exchanging $x = (x_1, \dots, x_n)$ for an integral vector $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$ as above completes the proof of the lemma. ■

Proof of Theorem 6: Converse Part. Suppose K is a secret-key achievable with communication \mathbf{F} . Then, by definition of a SK and Lemma 4 we have

$$C_{SK} = H(K) = H(K | \mathbf{F}) \quad (118)$$

$$= H(P) - \sum_{i=1}^n x_i \quad (119)$$

$$\leq H(P) - M^*(\mathcal{T}) \quad (120)$$

$$= k - M^*(\mathcal{T}). \quad (121)$$

Achievability Part. By definition, universal recovery can be achieved with $M^*(\mathcal{T})$ transmissions. Moreover, the communication \mathbf{F} can be generated as a linear function of P (see the proof of Theorem 1 and [21]). Denote this linear transformation by $\mathbf{F} = \mathcal{L}P$. Note that \mathcal{L} only depends on the

indices of the packets available to each node, not the values of the packets themselves (see [21]). Let $\mathcal{P}_{\mathbf{F}} = \{P' : \mathcal{L}P' = \mathbf{F}\}$ be the set of all packet distributions which generate \mathbf{F} .

By our assumption that the packets are i.i.d. uniform from \mathbb{F} , each $P' \in \mathcal{P}_{\mathbf{F}}$ is equally likely given \mathbf{F} was observed. Since \mathbf{F} has dimension $M^*(\mathcal{T})$, $|\mathcal{P}_{\mathbf{F}}| = \mathbb{F}^{k - M^*(\mathcal{T})}$. Thus, we can set $\mathcal{K} = \mathbb{F}^{k - M^*(\mathcal{T})}$ and label each $P' \in \mathcal{P}_{\mathbf{F}}$ with a unique element in \mathcal{K} . The label for the actual P (which is reconstructed by all nodes after observing \mathbf{F}) is the secret-key. Thus, $C_{SK} \geq k - M^*(\mathcal{T})$.

We remark that this labeling can be accomplished in polynomial time by an appropriate linear transformation mapping P to K . ■

Proof of Theorem 7: Converse Part. Suppose K is a private-key. Then, by definition of a PK and Lemma 4,

$$C_{PK} = H(K) = H(K | \mathbf{F}) \quad (122)$$

$$= H(P | P_D) - \sum_{i \in V \setminus D} x_i \quad (123)$$

$$\leq H(P | P_D) - M^*(\mathcal{T}_D) \quad (124)$$

$$= (k - |P_D|) - M^*(\mathcal{T}_D). \quad (125)$$

Achievability Part. Let each node $i \in D$ transmit P_i so that we can update $P_j \leftarrow P_j \cup P_D$ for each $j \in V \setminus D$. Now, consider the universal recovery problem for only the nodes in $V \setminus D$. $M^*(\mathcal{T}_D)$ is the minimum number of transmissions required among the nodes in $V \setminus D$ so that each node in $V \setminus D$ recovers P . At this point, the achievability proof proceeds identically to the SK case. ■

VI. CONCLUDING REMARKS

In this paper, we derive necessary and sufficient conditions for achieving universal recovery in an arbitrarily connected network. For the case when the network is fully connected, we provide a polynomial-time algorithm based on submodular optimization which solves the cooperative problem. This algorithm and its derivation yield tight concentration results for the case when packets are randomly distributed. Moreover, concentration results are provided when the network is d -regular and packets are distributed randomly. If packets are divisible, we prove that the traditional cut-set bounds are achievable. As a consequence of this and the concentration results, we show that splitting packets does not typically provide a significant benefit when the network is d -regular. Finally, we discuss an application to secrecy generation in the presence of an eavesdropper. We demonstrate that our submodular algorithm can be used to generate the maximum amount of secrecy in a practical manner.

It is conceivable that the coded cooperative data exchange problem can be solved (or approximated) in polynomial time if the network is d -regular, but packets aren't necessarily randomly distributed. This is one possible direction for future work.

APPENDIX A

AN EFFICIENTLY SOLVABLE INTEGER LINEAR PROGRAM

In this appendix, we introduce a special ILP and provide an polynomial-time algorithm for solving it. This algorithm

can be used to solve the cooperative data exchange problem when the underlying graph is fully-connected. We begin by introducing some notation⁷.

Let $E = \{1, \dots, n\}$ be a finite set with n elements. We denote the family of all subsets of E by 2^E . We frequently use the compact notation $E \setminus U$ and $U + i$ to denote the sets $E \cap U^c$ and $U \cup \{i\}$ respectively. For a vector $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, define the corresponding functional $x : 2^E \rightarrow \mathbb{R}$ as:

$$x(U) := \sum_{i \in U} x_i, \quad \text{for } U \subseteq E. \quad (126)$$

Throughout this section, we let $\mathcal{F} = 2^E - \{\emptyset, E\}$ denote the family of nonempty proper subsets of E . Let $\mathcal{B} = \{B_1, \dots, B_n\}$. No special structure is assumed for the B_i 's except that they are finite.

With the above notation established, we consider the following Integer Linear Program (ILP) in this section:

$$\min \left\{ \sum_{i \in E} w_i x_i : x(U) \geq \left| \bigcap_{i \in E \setminus U} B_i \right|, \quad \forall U \in \mathcal{F}, x_i \in \mathbb{Z} \right\}. \quad (127)$$

It is clear that any algorithm that solves this ILP also solves ILP (22) by putting $B_i \leftarrow P_i^c$ and $w = \mathbb{1}$.

A. Submodular Optimization

Our algorithm for solving ILP (127) relies heavily on submodular function optimization. To this end, we give a very brief introduction to submodular functions here.

A function $g : 2^E \rightarrow \mathbb{R}$ is said to be submodular if, for all $X, Y \in 2^E$,

$$g(X) + g(Y) \geq g(X \cap Y) + g(X \cup Y). \quad (128)$$

Over the past three decades, submodular function optimization has received a significant amount of attention. Notably, several polynomial time algorithms have been developed for solving the Submodular Function Minimization (SFM) problem

$$\min \{g(U) : U \subseteq E\}. \quad (129)$$

We refer the reader to [28]–[30] for a comprehensive overview of SFM and known algorithms. As we will demonstrate, we can solve ILP (127) via an algorithm that iteratively calls a SFM routine. The most notable feature of SFM algorithms is their ability to solve problems with exponentially many constraints in polynomial time. One of the key drawbacks of SFM is that the problem formulation is very specific. Namely, SFM routines typically require the function g to be submodular on *all* subsets of the set E .

B. The Algorithm

We begin by developing an algorithm to solve an equality constrained version of ILP (127). We will remark on the

⁷We attempt to keep the notation generic in order to emphasize that the results in this appendix are not restricted to the context of the cooperative data exchange problem.

general case at the conclusion of this section. To this end, let M be a positive integer and consider the following ILP:

$$\text{minimize } w^T x \quad (130)$$

$$\text{subject to: } x(U) \geq \left| \bigcap_{i \in E \setminus U} B_i \right| \quad \text{for all } U \in \mathcal{F}, \text{ and} \quad (131)$$

$$x(E) = M. \quad (132)$$

Remark 7: We assume $w_i \geq 0$, else in the case without the equality constraint we could allow the corresponding $x_i \rightarrow +\infty$ and the problem is unbounded from below.

Algorithm A.1: SOLVEILP(\mathcal{B}, E, M, w)

comment: Define $f : 2^E \rightarrow \mathbb{R}$ as in eq. (134).

$x \leftarrow \text{COMPUTE POTENTIALX}(f, M, w)$

if CHECKFEASIBLE(f, x)

then return (x)

else return (Problem Infeasible)

Theorem 8: Algorithm A.1 solves the equality constrained ILP (130) in polynomial time. If feasible, Algorithm A.1 returns an optimal x . If infeasible, Algorithm A.1 returns “Problem Infeasible”.

Proof: The proof is accomplished in three steps:

- 1) First, we show that if our algorithm returns an x , it is feasible.
- 2) Second, we prove that if a returned x is feasible, it is also optimal.
- 3) Finally, we show that if our algorithm does not return an x , then the problem is infeasible.

Each step is given its own subsection, and our proof roughly follows the development in [30] on crossing families. ■

Algorithm A.1 relies on three basic subroutines given below:

Algorithm A.2: COMPUTE POTENTIALX(f, M, w)

comment: If feasible, returns x satisfying (131) and (132) that minimizes $w^T x$.

comment: Order elements of E

so that $w_1 \geq w_2 \geq \dots \geq w_n$.

for $i \leftarrow n$ **to** 2

do $\left\{ \begin{array}{l} \text{comment: Define } f_i(U) := f(U + i) \\ \text{for } U \subseteq \{i, \dots, n\}. \\ x_i \leftarrow \text{SFM}(f_i, \{i, \dots, n\}) \end{array} \right.$

$x_1 \leftarrow M - \sum_{i=2}^n x_i$

return (x)

Algorithm A.3: CHECKFEASIBLE(f, x)

comment: Check if $x(U) \leq f(U)$ for all $U \in \mathcal{F}$ with $1 \in U$.

comment: Define $f_1(U) := f(U + 1)$ for $U \subseteq E$.

if SFM(f_1, E) < 0

then return (false)

else return (true)

Algorithm A.4: SFM(f, V)

comment: Minimize submodular function f over groundset V . See [28] for details.
 $v \leftarrow \min \{f(U) : U \subseteq V\}$
return (v)

C. Feasibility of a Returned x

In this section, we prove that if Algorithm A.1 returns a vector x , it must be feasible. We begin with some definitions.

Definition 7: A pair of sets $X, Y \subset E$ is called **crossing** if $X \cap Y \neq \emptyset$ and $X \cup Y \neq E$.

Definition 8: A function $g : 2^E \rightarrow \mathbb{R}$ is **crossing submodular** if

$$g(X) + g(Y) \geq g(X \cap Y) + g(X \cup Y) \quad (133)$$

for X, Y crossing.

We remark that minimization of crossing submodular functions is well established, however it involves a lengthy reduction to a standard submodular optimization problem. However, the crossing family \mathcal{F} admits a straightforward algorithm, which is what we provide in Algorithm A.1. We refer the reader to [30] for complete details on the general case.

For M a positive integer, define

$$f(U) := M - \left| \bigcap_{i \in U} B_i \right| - x(U), \quad \text{for } U \in \mathcal{F}. \quad (134)$$

Lemma 5: The function f is crossing submodular on \mathcal{F} .

Proof: For $X, Y \in \mathcal{F}$ crossing:

$$\begin{aligned} f(X) + f(Y) &= M - \left| \bigcap_{i \in X} B_i \right| - x(X) + M - \left| \bigcap_{i \in Y} B_i \right| - x(Y) \quad (135) \\ &= M - \left| \bigcap_{i \in X} B_i \right| - x(X \cap Y) \\ &\quad + M - \left| \bigcap_{i \in Y} B_i \right| - x(X \cup Y) \quad (136) \end{aligned}$$

$$\begin{aligned} &\geq M - \left| \bigcap_{i \in X \cap Y} B_i \right| - x(X \cap Y) \\ &\quad + M - \left| \bigcap_{i \in X \cup Y} B_i \right| - x(X \cup Y) \quad (137) \\ &= f(X \cap Y) + f(X \cup Y). \quad (138) \end{aligned}$$

Observe that, with f defined as above, the constraints of ILP (130) can be equivalently written as:

$$f(U) = M - \left| \bigcap_{i \in U} B_i \right| - x(U) \geq 0 \quad \text{for all } U \in \mathcal{F}, \quad (139)$$

$$x(E) = M. \quad (140)$$

Without loss of generality, assume the elements of E are ordered lexicographically so that $w_1 \geq w_2 \geq \dots \geq w_n$.

At iteration i in Algorithm A.2, $x_j = 0$ for all $j \leq i$. Thus, setting

$$x_i \leftarrow \min_{U \subseteq \{i, \dots, n\}} \{f_i(U)\} \quad (141)$$

$$= \min_{U \subseteq \{i, \dots, n\}; i \in U} \{f(U)\} \quad (142)$$

$$= \min_{U \subseteq \{i, \dots, n\}; i \in U} \left\{ M - \left| \bigcap_{i \in U} B_i \right| - x(U) \right\} \quad (143)$$

and noting that the returned x satisfies $x(E) = M$, rearranging (143) guarantees that

$$x(E \setminus U) \geq \left| \bigcap_{i \in U} B_i \right|, \quad \text{for all } U \subseteq \{i, \dots, n\}, i \in U \quad (144)$$

as desired. Iterating through $i \in \{2, \dots, n\}$ guarantees (144) holds for $2 \leq i \leq n$.

Remark 8: In the feasibility check routine (Algorithm A.3), we must be able to evaluate $f_1(E)$. The reader can verify that putting $f(E) = 0$ preserves submodularity.

Now, in order for the feasibility check to return **true**, we must have

$$\min_{U \subseteq E} \{f_1(U)\} = \min_{U \subseteq E: 1 \in U} \{f(U)\} \quad (145)$$

$$= \min_{U \subseteq E: 1 \in U} \left\{ M - \left| \bigcap_{i \in U} B_i \right| - x(U) \right\} \quad (146)$$

$$\geq 0, \quad (147)$$

implying that

$$x(E \setminus U) \geq \left| \bigcap_{i \in U} B_i \right|, \quad \text{for all } U \subseteq E, 1 \in U. \quad (148)$$

Combining (144) and (148) and noting that $x(E) = M$ proves that x is indeed feasible. Moreover, x is integral as desired.

D. Optimality of a Returned x

In this section, we prove that if Algorithm A.1 returns a feasible x , then it is also optimal. First, we require two more definitions and a lemma.

Definition 9: A constraint of the form (139) corresponding to U is said to be **tight** for U if

$$f(U) = M - \left| \bigcap_{i \in U} B_i \right| - x(U) = 0. \quad (149)$$

Lemma 6: If x is feasible, X, Y are crossing, and their corresponding constraints are tight, then the constraints corresponding to $X \cap Y$ and $X \cup Y$ are also tight.

Proof: Since the constraints corresponding to X and Y are tight, we have

$$0 = f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y) \geq 0. \quad (150)$$

The first inequality is due to submodularity and the last inequality holds since x is feasible. This implies the result. ■

Definition 10: A family of sets \mathcal{L} is **laminar** if $X, Y \in \mathcal{L}$ implies either $X \cap Y = \emptyset$, $X \subset Y$, or $Y \subset X$.

At iteration k ($1 < k \leq n$) of Algorithm A.2, let U_k be the set where (143) achieves its minimum. Note that $k \in U_k \subseteq$

$\{k, \dots, n\}$. By construction, the constraint corresponding to U_k is tight. Also, the constraint $x(E) = M$ is tight. From the U_k 's and E we can construct a laminar family as follows: if $U_j \cap U_k \neq \emptyset$ for $j < k$, then replace U_j with $\tilde{U}_j \leftarrow U_k \cup U_j$. By Lemma 6, the constraints corresponding to the sets in the newly constructed laminar family are tight. Call this family \mathcal{L} . For each $i \in E$, there is a unique smallest set in \mathcal{L} containing i . Denote this set L_i . Since $k \in U_k \subseteq \{k, \dots, n\}$, $L_i \neq L_j$ for $i \neq j$. Note that $L_1 = E$ and $L_i \subset L_j$ only if $j < i$.

For each $L_i \in \mathcal{L}$ there is a unique smallest set L_j such that $L_i \subset L_j$. We call L_j the least upper bound on L_i .

Now, consider the dual linear program to (130):

$$\text{maximize } - \sum_{U \in \mathcal{F}} \pi_U \left(M - \left| \bigcap_{i \in U} B_i \right| \right) - \pi_E M \quad (151)$$

$$\text{subject to: } \sum_{U \in \mathcal{F}: i \in U} \pi_U + \pi_E + w_i = 0, \quad \text{for } 1 \leq i \leq n$$

$$\pi_U \geq 0 \quad \text{for } U \in \mathcal{F}, \text{ and } \pi_E \text{ free.} \quad (152)$$

For each $L_i \in \mathcal{L}$, let the corresponding dual variable $\pi_{L_i} = w_j - w_i$, where L_j is the least upper bound on L_i . By construction, $\pi_{L_i} \geq 0$ since it was assumed that $w_1 \geq \dots \geq w_n$. Finally, let $\pi_E = -w_1$ and $\pi_U = 0$ for $U \notin \mathcal{L}$.

Now, observe that:

$$\sum_{U \in \mathcal{F}: i \in U} \pi_U + \pi_E + w_i = 0 \quad (153)$$

as desired for each i . Thus, π is dual feasible. Finally, note that $\pi_U > 0$ only if $U \in \mathcal{L}$. However, the primal constraints corresponding to the sets in \mathcal{L} are tight. Thus, (x, π) form a primal-dual feasible pair satisfying complementary slackness conditions, and are therefore optimal.

E. No Returned $x = \text{Infeasibility}$

Finally, we prove that if the feasibility check returns **false**, then ILP (130) is infeasible. Note by construction that the vector x passed to the feasibility check satisfies

$$M - \left| \bigcap_{i \in U} B_i \right| - x(U) \geq 0 \quad \text{for all nonempty } U \subseteq \{2, \dots, n\},$$

and $x(E) = M$. Again, let U_k be the set where (143) achieves its minimum and let \mathcal{L} be the laminar family generated by these U_k 's and E exactly as before. Again, the constraints corresponding to the sets in \mathcal{L} are tight (this can be verified in a manner identical to the proof of Lemma 6). Now, since x failed the feasibility check, there exists some exceptional set T with $1 \in T$ for which

$$M - \left| \bigcap_{i \in T} B_i \right| - x(T) < 0. \quad (154)$$

Generate a set L_T as follows: Initialize $L_T \leftarrow T$. For each $L_i \in \mathcal{L}$, $L_i \neq E$, if $L_T \cap L_i \neq \emptyset$, update $L_T \leftarrow L_T \cup L_i$. Now, we can add L_T to family \mathcal{L} while preserving the laminar property. We pause to make two observations:

- 1) By an argument similar to the proof of Lemma 6, we have that

$$M - \left| \bigcap_{i \in L_T} B_i \right| - x(L_T) < 0. \quad (155)$$

- 2) The sets in \mathcal{L} whose least upper bound is E form a partition of E . We note that L_T is a nonempty class of this partition. Call this partition $\mathcal{P}_{\mathcal{L}}$.

Again consider the dual constraints, however, let $w_i = 0$ (this does not affect feasibility). For each $L \in \mathcal{P}_{\mathcal{L}}$ define the associated dual variable $\pi_L = \alpha$, and let $\pi_E = -\alpha$. All other dual variables are set to zero. It is easy to check that this π is dual feasible. Now, the dual objective function becomes:

$$\begin{aligned} & - \sum_{U \in \mathcal{F}} \pi_U \left(M - \left| \bigcap_{i \in U} B_i \right| \right) - \pi_E M \\ &= -\alpha \sum_{L \in \mathcal{P}_{\mathcal{L}}} \left(M - \left| \bigcap_{i \in L} B_i \right| - x(L) + x(L) \right) + \alpha M \\ &= -\alpha \left(M - \left| \bigcap_{i \in L_T} B_i \right| - x(L_T) \right) - \alpha x(E) + \alpha M \\ &= -\alpha \left(M - \left| \bigcap_{i \in L_T} B_i \right| - x(L_T) \right) \\ &\rightarrow +\infty \text{ as } \alpha \rightarrow \infty. \end{aligned}$$

Thus, the dual is unbounded and therefore the primal problem must be infeasible.

As an immediate corollary we obtain the following:

Corollary 1: The optimal values of the ILP:

$$\min \{x(E) : x(U) \geq |\bigcap_{i \in E \setminus U} B_i|, U \in \mathcal{F}, x_i \in \mathbb{Z}\}$$

and the corresponding LP relaxation:

$$\min \{x(E) : x(U) \geq |\bigcap_{i \in E \setminus U} B_i|, U \in \mathcal{F}, x_i \in \mathbb{R}\}$$

differ by less than 1.

Proof: Algorithm A.1 is guaranteed to return an optimal x if the intersection of the polytope and the hyperplane $x(E) = M$ is nonempty. Thus, if M^* is the minimum such M , then the optimal value of the LP must be greater than $M^* - 1$. ■

F. Solving the General ILP

Finally, we remark on how to solve the general case of the ILP without the equality constraint given in (127). First, we state a simple convexity result.

Lemma 7: Let $p_w^*(M)$ denote the optimal value of ILP (130) when the equality constraint is $x(E) = M$. We claim that $p_w^*(M)$ is a convex function of M .

Proof: Let M_1 and M_2 be integers and let $\theta \in [0, 1]$ be such that $M_\theta = \theta M_1 + (1 - \theta) M_2$ is an integer. Let $x^{(1)}$ and be $x^{(2)}$ optimal vectors that attain $p_w^*(M_1)$ and $p_w^*(M_2)$ respectively. Let $x^{(\theta)} = \theta x^{(1)} + (1 - \theta) x^{(2)}$. By convexity, $x^{(\theta)}$ is feasible, though not necessarily integer. However, by




Fig. 6. Empirical computation times for Algorithm A.1, where each data point is averaged over 10 trials. For the broken line, the multiplicative constant α and exponent β were chosen to minimize the MSE $\sum_{i=1}^n |\log(\alpha n^\beta) - \log(\hat{m}_n)|^2$, where \hat{m}_n is the sample mean of the computation times for $|E| = n$.

the results from above, optimality is always attained by an integral vector. Thus, it follows that:

$$\begin{aligned} \theta p_w^*(M_1) + (1 - \theta)p_w^*(M_2) &= \theta w^T x^{(1)} + (1 - \theta)w^T x^{(2)} \\ &= w^T x^{(\theta)} \end{aligned} \quad (156)$$

$$\geq p_w^*(M_\theta). \quad (157)$$

Noting that $p_w^*(M)$ is convex in M , we can perform bisection on M to solve the ILP in the general case. For our purposes, it suffices to have relatively loose upper and lower bounds on M since the complexity only grows logarithmically in the difference. A simple lower bound on M is given by $M \geq \max_i |B_i|$.

G. Complexity

Our aim in this paper is not to give a detailed complexity analysis of our algorithm. This is due to the fact that the complexity is dominated by the SFM over the set E in Algorithm A.3. Therefore, the complexity of Algorithm A.1 is dictated by the complexity of the SFM solver employed. Indeed, Algorithm A.1 sequentially performs SFM over groundsets of size $1, 2, \dots, n-1$, respectively. Hence the total complexity of Algorithm A.1 is $O(n \cdot \text{SFM}(n))$, where $\text{SFM}(m)$ is the complexity of performing SFM over a groundset of size m . It is well-known that $\text{SFM}(m)$ is polynomial in m , for instance the survey [28] describes an Ellipsoid Algorithm-based approach which runs in $O(m^5 \cdot \gamma + m^7)$ time, where γ is the complexity of evaluating the submodular function under consideration. We remark that the algorithm presented in [18] solves the LP relaxation of (127) in $O(n^3 \cdot \text{SFM}(n))$ time.

Despite the fact that SFM has somewhat large worst-case polynomial, we have performed a series of numerical experiments to demonstrate that Algorithm A.1 performs quite well in practice. In our implementation, we ran the Fujishige-Wolfe (FW) algorithm for SFM [31] based largely on a Matlab routine by A. Krause [32]. While the FW algorithm has not

been proven to run in polynomial time, it has been shown to work quite well in practice [31] (similar to the Simplex algorithm for solving Linear Programs). Whether or not FW has worst-case polynomial complexity is an open problem to date. As mentioned above, there are SFM algorithms that run in strongly polynomial time which could be used if a particular application requires polynomially bounded worst-case complexity [28].

In our series of experiments, we chose $B_i \subset F$ randomly, where $|F| = 50$. We let $n = |E|$ range from 10 to 190 in increments of 10. For each value of n , we ran 10 experiments. The average computation time is shown in Figure 6, with error bars indicating one standard deviation. We consistently observed that the computations run in approximately $O(n^{1.85})$ time. Due to the iterative nature of the SFM algorithm, we anticipate that the computation time could be significantly reduced by implementing the algorithm in C/C++ instead of Matlab. However, the $O(n^{1.85})$ trend should remain the same. Regardless, we are able to solve the ILP problems under consideration with an astonishing 2^{190} constraints in approximately one minute.

APPENDIX B PROOF OF LEMMA 2

For convenience, we repeat the statement of Lemma 2.

Lemma 2: Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with nonnegative entries and all column sums equal to d . Let \bar{x}_y be the vector of minimum Euclidean norm which minimizes $\|Ax_y - y\|_2$. There exists an optimal solution x^* to the linear program

$$\text{minimize } \mathbb{1}^T x \quad (158)$$

$$\text{subject to: } Ax \geq y \quad (159)$$

which satisfies

$$\|x^* - \bar{x}_y\|_\infty \leq c_A \|A\bar{x}_y - y\|_2, \quad (160)$$

where c_A is a constant depending only on A .

Before moving to the proof of the lemma, we comment briefly on the interpretation. If the matrix A in the statement of Lemma 2 is nonsingular, then the solution x^* to the linear program (158) is unique and equal to \bar{x}_y . Indeed, this can be seen by premultiplying both sides of (159) by $\mathbb{1}^T$, implying $\mathbb{1}^T x \geq \frac{1}{d} \mathbb{1}^T y$ for all feasible x . This lower bound is achieved only if $Ax = y$. On the other hand, if A is singular, there are many optimal solutions to the linear program (158). More precisely, if x^* is one optimal solution for the linear program (158), then so is $x^* + z$ for any z such that $Az = 0$. Thus, Lemma 2 guarantees that we can find an optimal solution x^* which is “close” to the least squares solution \bar{x}_y .

Proof of Lemma 2: To begin the proof, we make a few definitions. Let λ be the absolute value of the nonzero eigenvalue of A with smallest modulus (at least one exists since d is an eigenvalue). Define $\mathcal{N}(A)$ to be the nullspace of A , and let $\mathcal{N}^\perp(A)$ denote its orthogonal complement. Finally, let A^+ denote the Moore-Penrose pseudoinverse of A .

Fix $\bar{x}_y \in \mathbb{R}^n$, and note that x^* is an optimal solution to LP (158) if and only if $x^* - \bar{x}_y$ is an optimal solution to the

linear program

$$\text{minimize } \mathbf{1}^T(x + \bar{x}_y) \quad (161)$$

$$\text{subject to: } A(x + \bar{x}_y) \geq y \quad (162)$$

with variable $x \in \mathbb{R}^n$. With this in mind, put $\bar{x}_y = A^+y$ and define $b = y - A\bar{x}_y$. By definition of the pseudoinverse, \bar{x}_y is the vector of minimum Euclidean norm which minimizes $\|Ax_y - y\|_2$. Moreover, $b \in \mathcal{N}(A)$.

Thus, in order to prove the lemma, it suffices to show the existence of an optimal solution x^* to the linear program

$$\text{minimize } \mathbf{1}^T x \quad (163)$$

$$\text{subject to: } Ax \geq b \quad (164)$$

which also satisfies the additional constraints

$$|x_i| \leq c_A \|b\|_2 \quad \text{for } i = 1, \dots, n, \quad (165)$$

where c_A is a constant depending only on A .

Claim 1: There exists an optimal solution x^* to Linear Program (163) which satisfies

$$x_i^* \leq (d\lambda)^{-1}n\|b\|_\infty \quad \text{for } i = 1, \dots, n. \quad (166)$$

The proof relies heavily on duality. The reader is directed to [33] or any other standard text for details.

To prove the claim, consider LP (163). By premultiplying the inequality constraint by $d^{-1}\mathbf{1}^T$ on both sides, we see that $\mathbf{1}^T x \geq d^{-1}\mathbf{1}^T b > -\infty$. Hence, the objective is bounded from below, which implies that strong duality holds. Thus, let \tilde{z} be an optimal solution to the dual LP of (163):

$$\text{maximize } b^T \tilde{z} \quad (167)$$

$$\text{subject to: } A\tilde{z} = \mathbf{1} + y \quad (168)$$

$$z \geq 0 \quad (169)$$

with dual variable $z \in \mathbb{R}^n$.

Next, consider the dual LP of (163) with the additional inequality constraints corresponding to (166):

$$\text{maximize } b^T z - (d\lambda)^{-1}n\|b\|_\infty \mathbf{1}^T y \quad (170)$$

$$\text{subject to: } Az = \mathbf{1} + y \quad (171)$$

$$z \geq 0 \quad (172)$$

$$y \geq 0 \quad (173)$$

with dual variables $z \in \mathbb{R}^n$ and $y \in \mathbb{R}^n$. Equivalently, by setting $z = \tilde{z} + \Delta z$ and observing that $y = A\Delta z$, we can write the dual LP (170) as

$$\text{maximize } b^T \tilde{z} + b^T \Delta z - (d\lambda)^{-1}n\|b\|_\infty \mathbf{1}^T A\Delta z \quad (174)$$

$$\text{subject to: } A\Delta z \geq 0 \quad (175)$$

$$\tilde{z} + \Delta z \geq 0 \quad (176)$$

with dual variables $\Delta z \in \mathbb{R}^n$. We prove the claim by showing that the dual LPs (167) and (174) have the same optimal value. Since strong duality holds, the corresponding primal problems must also have the same optimal value.

Without loss of generality, we can uniquely decompose $\Delta z = \Delta z_1 + \Delta z_2$ where $\Delta z_1 \in \mathcal{N}(A)$ and $\Delta z_2 \in \mathcal{N}^\perp(A)$.

Since $b \in \mathcal{N}(A)$, we have $b^T \Delta z_2 = 0$ and we can rewrite (174) yet again as

$$\text{maximize } b^T \tilde{z} + b^T \Delta z_1 - (d\lambda)^{-1}n\|b\|_\infty \mathbf{1}^T A\Delta z_2 \quad (177)$$

$$\text{subject to: } A\Delta z_2 \geq 0 \quad (178)$$

$$\tilde{z} + \Delta z_1 + \Delta z_2 \geq 0 \quad (179)$$

$$\Delta z_1 \in \mathcal{N}(A), \Delta z_2 \in \mathcal{N}^\perp(A). \quad (180)$$

By definition of λ , for any unit vector $u \in \mathcal{N}^\perp(A)$ with $\|u\|_2 = 1$ we have $\|Au\|_2 \geq \lambda$. Using this and the fact that $A\Delta z_2 \geq 0$ for all feasible Δz_2 , we have the following inequality:

$$\mathbf{1}^T A\Delta z_2 = \|A\Delta z_2\|_1 \geq \|A\Delta z_2\|_2 \geq \lambda \|\Delta z_2\|_2. \quad (181)$$

Thus, the objective (177) can be upper bounded as follows:

$$\begin{aligned} b^T \tilde{z} + b^T \Delta z_1 - (d\lambda)^{-1}n\|b\|_\infty \mathbf{1}^T A\Delta z_2 \\ \leq b^T \tilde{z} + b^T \Delta z_1 - d^{-1}n\|b\|_\infty \|\Delta z_2\|_2. \end{aligned} \quad (182)$$

Next, we obtain an upper bound on $b^T \Delta z_1$. To this end, observe that constraint (179) implies that $\tilde{z} + \Delta z_1 \geq -\mathbf{1} \|\Delta z_2\|_\infty$. Motivated by this, consider the following ϵ -perturbed LP:

$$\text{minimize } -b^T v \quad (183)$$

$$\text{subject to: } \tilde{z} + v \geq -\epsilon \mathbf{1}$$

$$v \in \mathcal{N}(A).$$

with variable v . Let $p^*(\epsilon)$ denote the optimal value of the ϵ -perturbed problem. First observe that $p^*(0) = 0$. To see this, note that if $\tilde{z} + v \geq 0$, then $b^T v \leq 0$, else we would contradict the optimality of \tilde{z} since $z = \tilde{z} + v$ is a feasible solution to the dual LP (167) in this case. Now, weak duality implies

$$-b^T v \geq p^*(\epsilon) \geq p^*(0) - \epsilon \mathbf{1}^T w^*, \quad (184)$$

where w^* corresponds to an optimal solution to the dual LP of the unperturbed primal LP (183), given by:

$$\text{maximize } -\tilde{z}^T (Aw - b) \quad (185)$$

$$\text{subject to: } Aw \geq b. \quad (186)$$

Hence, (184) implies that

$$b^T \Delta z_1 \leq \|\Delta z_2\|_\infty \mathbf{1}^T w^* \quad (187)$$

if $\Delta z_1, \Delta z_2$ are feasible for LP (177).

By definition of \tilde{z} , $\tilde{z}^T A = \mathbf{1}^T$, and hence a vector w^* is optimal for (185) if and only if it also optimizes:

$$\text{minimize } \mathbf{1}^T w \quad (188)$$

$$\text{subject to: } Aw \geq b. \quad (189)$$

Combining this with (187), we have

$$b^T \Delta z_1 \leq \|\Delta z_2\|_\infty \mathbf{1}^T w^* \leq \|\Delta z_2\|_\infty \mathbf{1}^T w \quad (190)$$

for any vector w satisfying $Aw \geq b$. Trivially, $w = d^{-1}\|b\|_\infty \mathbf{1}$ satisfies this, and hence we obtain:

$$b^T \Delta z_1 \leq d^{-1}n\|b\|_\infty \|\Delta z_2\|_\infty. \quad (191)$$

Finally, we substitute this into (182) and see that

$$\begin{aligned} b^T z &\leq b^T \tilde{z} + d^{-1} n \|b\|_\infty \|\Delta z_2\|_\infty - d^{-1} n \|b\|_\infty \|\Delta z_2\|_2 \\ &\leq b^T \tilde{z} + d^{-1} n \|b\|_\infty \|\Delta z_2\|_2 - d^{-1} n \|b\|_\infty \|\Delta z_2\|_2 \\ &\leq b^T \tilde{z} \end{aligned}$$

for all vectors z which are feasible for the dual LP (170). This completes the proof of Claim 1.

Claim 2: There exists an optimal solution x^* to Linear Program (163) which satisfies

$$|x_i| \leq c_A \|b\|_2 \quad \text{for } i = 1, \dots, n \quad (192)$$

for some constant c_A depending only on A .

First note that $\|b\|_\infty \leq \|b\|_2$ for any $b \in \mathbb{R}^n$, hence it suffices to prove the claim for the infinity norm. Claim 1 shows that each of the x_i 's can be upper bounded by $(d\lambda)^{-1} n \|b\|_\infty$ without affecting the optimal value of LP (163). To see the lower bound, let a_j^T be a row of A with entry $a_{ji} \geq d/n$ in the i^{th} coordinate (at least one exists for each i since the columns of A sum to d). Now, the inequality constraint $Ax \geq b$ combined with the upper bound on each x_i implies:

$$a_{ji} x_i + (d - a_{ji}) \lambda^{-1} n \|b\|_\infty \geq a_j^T x \geq b_j \geq -\|b\|_\infty. \quad (193)$$

Since $a_{ji} \geq d/n$, (193) implies:

$$x_i \geq -\lambda^{-1} n(n-1) \|b\|_\infty. \quad (194)$$

Hence, we can take $c_A = \lambda^{-1} n \times \max\{n-1, d^{-1}\}$. This proves Claim 2, and, by our earlier remarks, proves the lemma. ■

ACKNOWLEDGMENT

The authors would like to thank Kent Benson, Alex Sprintson, Pavan Datta, and Chi-Wen Su for the helpful conversations and suggestions which led to this paper. The authors also wish to thank one of the anonymous reviewers who provided many constructive comments.

REFERENCES

- [1] T. Courtadé, B. Xie, and R. Wesel, "Optimal exchange of packets for universal recovery in broadcast networks," in *Proc. Military Commun. Conf.*, Nov. 2010, pp. 2250–2255.
- [2] T. Courtadé and R. Wesel, "Efficient universal recovery in broadcast networks," in *Proc. 48th Annu. Allerton Conf. Commun., Control, Comput.*, Oct. 2010, pp. 1542–1549.
- [3] T. Courtadé and R. Wesel, "Weighted universal recovery, practical secrecy, and an efficient algorithm for solving both," in *Proc. 49th Annu. Allerton Conf. Commun., Control, Comput.*, Oct. 2011, pp. 1349–1357.
- [4] R. Ahlswede, N. Cai, S. Yen Robert Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.
- [5] S.-Y. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inf. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.
- [6] S. El Rouayheb, M. Chaudhry, and A. Sprintson, "On the minimum number of transmissions in single-hop wireless coding networks," in *Proc. IEEE ITW*, Sep. 2007, pp. 120–125.
- [7] S. El Rouayheb, A. Sprintson, and P. Sadeghi, "On coding for cooperative data exchange," in *Proc. IEEE ITW*, Jan. 2010, pp. 1–5.
- [8] A. Sprintson, P. Sadeghi, G. Booker, and S. El Rouayheb, "A randomized algorithm and performance bounds for coded cooperative data exchange," in *Proc. IEEE ISIT*, Jun. 2010, pp. 1888–1892.
- [9] A. Sprintson, P. Sadeghi, G. Booker, and S. El Rouayheb, "Deterministic algorithm for coded cooperative data exchange," in *ICST QShine*. New York, NY, USA: Springer-Verlag, Nov. 2010.
- [10] D. Ozgul and A. Sprintson, "An algorithm for cooperative data exchange with cost criterion," in *Proc. ITA Workshop*, Feb. 2011, pp. 1–4.
- [11] Y. Birk and T. Kol, "Coding on demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2825–2830, Jun. 2006.
- [12] D. Shah, "Gossip algorithms," *Found. Trends Netw.*, vol. 3, no. 1, pp. 1–125, 2008.
- [13] S. Deb, M. Medard, and C. Choute, "Algebraic gossip: A network coding approach to optimal multiple rumor mongering," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2486–2507, Jun. 2006.
- [14] I. Csiszar and P. Narayan, "Secrecy capacities for multiple terminals," *IEEE Trans. Inf. Theory*, vol. 50, no. 12, pp. 3047–3061, Dec. 2004.
- [15] C. Ye and P. Narayan, "Secret key and private key constructions for simple multiterminal source models," in *Proc. IEEE ISIT*, Sep. 2005, pp. 2133–2137.
- [16] C. Ye and A. Reznik, "A simple secret key construction system for broadcasting model," in *Proc. 44th Annu. CISS*, Mar. 2010, pp. 1–6.
- [17] A. Beimel, *Secret-Sharing Schemes: A Survey*, Y. Chee, Z. Guo, S. Ling, F. Shao, Y. Tang, H. Wang, and C. Xing, Eds. New York, NY, USA: Springer-Verlag, 2011.
- [18] N. Milosavljevic, S. Pawar, S. El Rouayheb, M. Gastpar, and K. Ramchandran, "Deterministic algorithm for the cooperative data exchange problem," in *Proc. IEEE Symp. Inf. Theory*, Aug. 2011, pp. 410–414.
- [19] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, Eds. New York, NY, USA: Plenum Press, 1972, pp. 85–103.
- [20] M. Gonen and M. Langberg, "Coded cooperative data exchange problem for general topologies," in *Proc. IEEE ISIT*, Jul. 2012, pp. 2606–2610.
- [21] S. Jaggi, P. Sanders, P. Chou, M. Effros, S. Egner, K. Jain, et al., "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1973–1982, Jun. 2005.
- [22] T. Ho, M. Medard, R. Koetter, D. Karger, M. Effros, J. Shi, et al., "A random linear network coding approach to multicast," *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.
- [23] T. Halford, K. Chugg, and A. Polydoros, "Barrage relay networks: System and protocol design," in *Proc. 21st IEEE Int. Symp. PIMRC*, Sep. 2010, pp. 1133–1138.
- [24] T. Halford and K. Chugg, "Barrage relay networks," in *Proc. ITA Workshop*, Feb. 2010, pp. 1–8.
- [25] T. Halford and G. Hwang, "Barrage relay networks for unmanned ground systems," in *Proc. Military Commun. Conf.*, Nov. 2010, pp. 1274–1280.
- [26] A. Blair, T. Brown, K. Chugg, T. Halford, and M. Johnson, "Barrage relay networks for cooperative transport in tactical MANETS," in *Proc. Military Commun. Conf.*, Nov. 2008, pp. 1–7.
- [27] A. Ramamoorthy, J. Shi, and R. D. Wesel, "On the capacity of network coding for random networks," *IEEE Trans. Inf. Theory*, vol. 51, no. 8, pp. 2878–2885, Aug. 2005.
- [28] S. T. McCormick, *Submodular Function Minimization*. In *Handbooks in Operations Research and Management Science*, vol. 12, New York, NY, USA: Elsevier, 2005.
- [29] S. Fujishige, *Submodular Functions and Optimization*, 2nd ed. New York, NY, USA: Elsevier, 2010.
- [30] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. New York, NY, USA: Springer-Verlag, 2003.
- [31] S. Fujishige, T. Hayashi, and S. Isotani, "The minimum-norm-point algorithm applied to submodular function minimization and linear programming," Ph.D. dissertation, Res. Inst. Math. Sci., Kyoto Univ., Kyoto, Japan, 2006.
- [32] A. Krause and S. Sonnenburg, "SFO: A toolbox for submodular function optimization, the," *J. Mach. Learn. Res.*, vol. 11, pp. 1141–1144, Jan. 2010.
- [33] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K: Cambridge University Press, 2004.

Thomas A. Courtadé (S'06-M'13) received the B.S. degree in Electrical Engineering from Michigan Technological University in 2007, and the M.S. and Ph.D. degrees in Electrical Engineering from UCLA in 2008 and 2012, respectively. In 2012, he was awarded the inaugural Postdoctoral Research Fellowship through the NSF Center for Science of Information. He currently holds this position, and resides at Stanford University. His honors include a Distinguished Ph.D. Dissertation award and an Excellence in Teaching award from the UCLA Department of Electrical Engineering and a Best Student Paper Award at the 2012 International Symposium on Information Theory.

Richard D. Wesel (S'91–M'96–SM'01) is a Professor with the UCLA Electrical Engineering Department and is the Associate Dean for Academic and Student Affairs for the UCLA Henry Samueli School of Engineering and Applied Science. He joined UCLA in 1996 after receiving his Ph.D. in electrical engineering from Stanford. His B.S. and M.S. degrees in electrical engineering are from MIT. His research is in the area of communication theory with particular interest in channel coding. He has received the National Science Foundation (NSF) CAREER Award, an Okawa Foundation award for research in information theory and telecommunications, and the Excellence in Teaching Award from the Henry Samueli School of Engineering and Applied Science.