Minimum cost content distribution using network coding: Replication vs. coding at the source nodes

Shurui Huang and Aditya Ramamoorthy
Department of Electrical and Computer Engineering
Iowa State University
Ames, Iowa 50011
Email: {hshurui, adityar}@iastate.edu

Muriel Médard
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
Email: {medard}@mit.edu

***Abstract*—Consider a large file that needs to be multicast over a network to a given set of terminals. Storing the file at a single server may result in server overload. Accordingly, there are distributed storage solutions that operate by dividing the file into pieces and placing copies of the pieces (replication) or coded versions of the pieces (coding) at multiple source nodes. Suppose that the cost of a given network coding based solution to this problem is defined as the sum of the storage cost and the cost of the flows required to support the multicast. In this work, we consider a network with a set of source nodes that can either contain subsets or coded versions of the pieces of the file and are interested in finding the storage capacities and flows at minimum cost. We provide succinct formulations of the corresponding optimization problems by using information measures. In particular, we show that when there are two source nodes, there is no loss in considering subset sources. For three source nodes, we derive a tight upper bound on the cost gap between the two cases. Algorithms for determining the content of the source nodes are also provided.**

I. Introduction

Large scale content distribution over the Internet is a topic of great interest in recent years. The dominant mode of content distribution is the client-server model, where a given client requests a central server for the file, which then proceeds to service the request. A single server, however is likely to be overwhelmed when a large number of users request for a file at the same time and the websites are often replicated by the use of mirrors. One can also consider the usage of coding for replicating the content, e.g., if one uses erasure codes such as Reed-Solomon codes or fountain codes, then it turns out that obtaining a certain number of coded packets from each of mirrors will suffice. Peer-to-peer networks have also been proposed for content distribution in a distributed manner.

The technique of network coding has also been considered for content distribution in networks. Network coding allows us to use the network resources more efficiently in the case of multicast. Under network coding based multicast, the problem of allocating resources such as rates and flows in the network can be solved in polynomial time. Moreover, one can arrive at distributed solutions to these problems in an easier manner.

In this work, we consider the following problem. Suppose that there is a large file, that can be broken into small pieces, that needs to be transmitted to a given set of clients over a network using network coding. The network has a designated set of nodes (source nodes) that have storage space. Each unit of storage space and each unit of flow over a certain edge has a known linear cost. We want to determine the optimal storage capacities and flow patterns such that this can be done with minimum cost. Within this problem setting, we distinguish two different cases: (i) *Subset sources case:* Each source node only contains a subset of the pieces of the file. (ii) *Coded sources case:* Each source node can contain arbitrary functions of the pieces of the file. In the subset sources case, we consider a file represented as (a, b, c, d) in Figure 1, where each component has unit-entropy, and a network where each edge has capacity 3. The cost of transmitting is 1 per unit rate over any edge, the cost of storage at the sources is 1 per unit storage. As shown in the figure, the case of partial replication when the source nodes contain dependent information has lower cost compared to the cases when the source nodes contain independent information or identical information (full replication). The case of subset sources, is interesting for multiple reasons. For example, it may be the case that a given terminal is only interested in a part of the original file. In this case, if one places coded pieces of the original file at the source nodes, then the terminal may need to obtain a large number of coded pieces before it can recover the part that it is interested in. In the extreme case, if coding is performed across all the pieces of the file, then the terminal will need to recover all the sources before it can recover the part it is interested in. Note however, that in this work we do not explicitly consider scenarios where a given terminal requires part of the file. From a theoretical perspective as well, it is interesting to examine how much loss one incurs by not allowing coding at the sources.

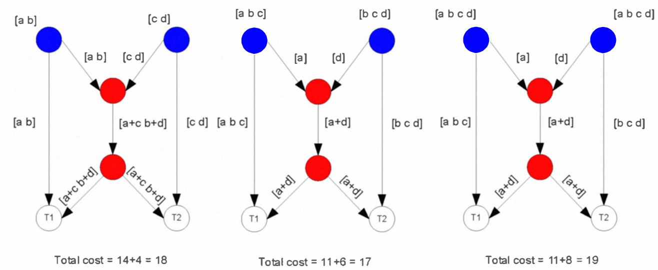

Fig. 1. The cost comparison of three schemes when a document $[a\ b\ c\ d]$ needs to be transmitted to two terminals through two sources.

This research was supported in part by NSF grant CNS-0721453.

Optimization issues in network coding have been examined in the past. The work of [1], proposed linear programming formulations for minimum cost flow allocation network coding based multicast. Lee et al. [2] constructed minimum cost subgraphs for the multicast of two correlated sources. It also proposed the problem of optimizing the correlation structure of sources and their placement. However, a solution was not presented there. Efficient algorithms for jointly allocating flows and rates were proposed for the multicast of a large number of correlated sources in [3]. The work of Jiang [4], considered a formulation that is similar to ours. It shows that under network coding, the problem of minimizing the joint transmission and storage cost can be formulated as a linear program. Furthermore, it considers a special class of networks called generalized tree networks and shows that there is no cost difference whether one considers subset sources or coded sources. In contrast, in this work we consider general networks. The work of Bhattad et al. [5] proposed a problem formulation for cost minimization when some nodes are only allowed routing instead of network coding. Our problem formulation can be considered as a specific instance of it. However, our formulation is much simpler than [5] and allows us to compare the cost of subset sources vs. coded sources. In addition, we recover stronger results in the case when there are only two or three sources. Note also our solution approach is quite different and uses the concept of information measures.

Our main contributions are: (1) we provide a precise formulation of the different optimization problems by leveraging the properties of the information measure (I-measure) introduced in [6]. (2) The usage of the properties of information measure allows us to conclude that when there are two source nodes, there is no loss in considering subset sources. Furthermore, in the case of three source nodes, we derive an upper bound on the cost gap that is shown to be tight. Finally, we formulate a gap LP to determine the cost gap for general cases.

This paper is organized as follows. In Section II, we give the main theory that defines the subset source problem in the content of information measures. In section III, we formulate the problems in both the subset case and the coded case, and give the source construction procedure. In Section IV, we discuss the cost gap between the two cases. In Section V, we provide the simulation results based on an example network and random networks. Section VI concludes the paper. Owing to space limitations, we omit many of the proofs in this paper. These can be found in [7].

II. Preliminaries

In this section we develop some key results, that will be used throughout the paper. In particular, we shall deal extensively with the I-measure introduced in [6]. We refer the reader to [6] for the required background in this area.

Let $\mathcal{N}_S = \{1, 2, \cdots, n\}$ and consider n random variables $X_1, X_2, \cdots, X_n$. Let $\widetilde{X}_i$ be a set corresponding to X_i and let $\widetilde{X}_V = \cup_{i\in V}\widetilde{X}_i$. We denote the set of nonempty atoms of $\mathcal{F}_n$ by $\mathcal{A}$, where $\mathcal{F}_n$ is the field generated by the sets $\widetilde{X}_1, \widetilde{X}_2, \cdots, \widetilde{X}_n$. Similarly, X_V denotes the collection of random variables $(X_i, i \in V)$, where $V \subseteq \mathcal{N}_S$. Construct the signed measure $\mu^*(\widetilde{X}_V) = H(X_V)$, $\forall V \subseteq \mathcal{N}_S$.

Theorem 1: (1) Suppose that there exists a set of $2^n - 1$ nonnegative values, one corresponding to each atom of $\mathcal{F}_n$, i.e, $\alpha(A) \geq 0, \forall A \in \mathcal{A}$. Then, we can define a set of independent random variables, $W_A, A \in \mathcal{A}$ and construct random variables $X_j = (W_A : A \in \mathcal{A}, A \subset \widetilde{X}_j)$, such that the measures of the nonempty atoms of the field generated by $\widetilde{X}_1, \widetilde{X}_2, \cdots, \widetilde{X}_n$ correspond to the values of α, i.e., $\mu^*(A) = \alpha(A), \forall A \in \mathcal{A}$.
(2) Conversely, let $Z_i, i \in \{1, \ldots, m\}$ be a collection of independent random variables. Suppose that a set of random variables $X_i, i = 1, \ldots, n$ is such that $X_i = Z_{V_i}$, where $V_i \subseteq \{1, \ldots, m\}$. Then the set of atoms of the field generated by $\widetilde{X}_1, \widetilde{X}_2, \cdots, \widetilde{X}_n$, have non-negative measures.

proof: (1) Independent random variables $W_A, A \in \mathcal{A}$, such that $H(W_A) = \alpha(A)$ can be constructed [6]. Set $X_i = (W_A : A \in \mathcal{A}, A \subset \widetilde{X}_i)$. Then we can check the consistency of the measures using the properties of the signed measure.

(2) It can be shown than all the measures are nonnegative using induction. Without loss of generality, we analyze $\mu^*(\widetilde{X}_1 \cap \cdots \cap \widetilde{X}_l \cap_{k:k\in K} \widetilde{X}_k^c)$ where $K \subseteq \mathcal{N}_S \setminus \{1, 2, \cdots, l\}$.

When $l = 1$, the measure corresponds to conditional entropy, $\forall K \subseteq \mathcal{N}_S \setminus \{1\}$, $\mu^*(\widetilde{X}_1 \cap_{k:k\in K} \widetilde{X}_k^c) = H(X_1|X_K) \geq 0$.

When $l = 2$, we have, $\forall K \subseteq \mathcal{N}_S \setminus \{1, 2\}$

$$\mu^*(\widetilde{X}_1 \cap \widetilde{X}_2 \cap_{k:k\in K} \widetilde{X}_k^c) = I(X_1; X_2|X_K)$$
$$= \sum_{i\in V_1\cap V_2 \cap_{k:k\in K} V_k^c} H(Z_i) \geq 0$$

We can then prove $\mu^*(\widetilde{X}_1 \cap \cdots \cap \widetilde{X}_l \cap_{k:k\in K} \widetilde{X}_k^c) \geq 0$ by induction, where $K \subseteq \mathcal{N}_S \setminus \{1, 2, \cdots, l\}$, $\forall l \leq n$. The details of the proof are skipped. In a similar manner it is easy to see that all atom measures are non-negative. ■

We emphasize that in general, atoms can be negative [6].

III. Problem Formulation

We now present the precise problem formulations for the subset sources case and the coded sources case. Suppose that we are given a directed graph $G = (V, E, C)$ that represents the network, V denotes the set of vertices, E denotes the set of edges, and C_{ij} denotes the capacity constraint for edge $(i, j) \in E$. There is a set of source nodes $S \subset V$ (numbered $1, \ldots, n$) and terminal nodes $T \subset V$. We assume that the original source, that has a certain entropy can be represented as the collection of equal entropy independent sources $\{OS_j\}_{j=1}^{Q}$, where Q is a sufficiently large integer. This assumption is equivalent to assuming that a file can be split into arbitrarily small pieces. Let X_i represent the source at the i^{th} source node. Suppose that each edge (i, j) incurs a linear cost $f_{ij}z_{ij}$ for a flow of value z_{ij} over it, and each source incurs a linear cost $d_i H(X_i)$ for the information X_i.

A. Subset Sources Case

In this case each source $X_i, i = 1, \ldots, n$ is constrained to be a subset of the pieces of the original source. We leverage Theorem 1 from the previous section that tells us that in this case that $\mu^*(A) \geq 0$ for all $A \in \mathcal{A}$.

We construct an augmented graph $G^* = (V^*, E^*, C^*)$ as follows. We append a virtual super node s^* to G, and connect s^* and each source node i with virtual edges, such that its capacity is infinity and its cost is d_i.

Let $x_{ij}^{(t)}$, $t \in T$ represent the flow variable over G^* corresponding to the terminal t along edge (i,j) and let z_{ij} represent $\max_{t \in T} x_{ij}^{(t)}, \forall (i,j \in E)$. We pose the problem as one of first recovering all the sources, X_i, $i \in S$ at each terminal and then the original source. Note that since these sources are correlated, this formulation is equivalent to the Slepian-Wolf problem over a network [3]. We introduce the variable $R_i^{(t)}, t \in T$ that represents the rate from source i to terminal t. Thus $R^{(t)} = (R_1^{(t)}, R_2^{(t)}, \cdots, R_n^{(t)})$ represents the rate vector for terminal t. In order for terminal t to recover the sources, the rate vector $R^{(t)}$ needs to lie within the Slepian-Wolf region of the sources, which is defined as follows.

$$\mathcal{R}_{\mathcal{SW}} = \{(R_1, \cdots, R_n) : \forall U \subseteq S, \sum_{i \in U} R_i \geq H(X_U | X_{S \backslash U})\}$$

Moreover, the rates also need to be in the capacity region such that the network has enough capacity to support them for each terminal. From Theorem 1, we have the subset constraints $\mu^*(A) \geq 0, \forall A \in \mathcal{A}$. The optimization problem is defined as SUBSET-MIN-COST. The formulation is as follows.

$$\text{minimize } \textstyle\sum_{(i,j) \in E} f_{ij} z_{ij} + \sum_{i \in S} d_i z_{s^* i}$$

$$\text{subject to } 0 \leq x_{ij}^{(t)} \leq z_{ij} \leq c_{ij}^*, \ \forall (i,j) \in E^*, \forall t \in T \quad (1)$$

$$\sum_{\{j | (i,j) \in E^*\}} x_{ij}^{(t)} - \sum_{\{j | (j,i) \in E^*\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \ i \in V^*, \ t \in T$$

$$x_{s^* i}^{(t)} \geq R_i^{(t)}, \forall i \in S, t \in T \quad (2)$$

$$R^{(t)} \in \mathcal{R}_{\mathcal{SW}}, \forall t \in T \quad (3)$$

$$\mu^*(A) \geq 0, \forall A \in \mathcal{A} \quad (4)$$

$$z_{s^* i} = H(X_i), \forall i \in S \quad (5)$$

$$H(X_1, X_2, \cdots, X_n) = \sum_{A \in \mathcal{A}} \mu^*(A) \quad (6)$$

where
$$\sigma_i^{(t)} = \begin{cases} H(X_1, X_2, \cdots, X_n) & \text{if } i = s^* \\ -H(X_1, X_2, \cdots, X_n) & \text{if } i = t \\ 0 & \text{otherwise} \end{cases}$$

Though not expressed explicitly, there should be $H(X_i) = \sum_{A: A \subset \tilde{X}_i} \mu^*(A)$ and $H(X_U | X_{S \backslash U}) = \sum_{A: A \not\subset \tilde{X}_{S \backslash U}} \mu^*(A)$, so that the measures and the entropies are consistent.

This formulation will be useful when we compare the costs of the coded and subset cases. An alternative formulation that has less variables and constraints can be found in [7].

B. Solution explanation and construction

Assume that we solve the above problem and obtain the values of all the atoms $\mu^*(A), \forall A \in \mathcal{A}$. These will in general be fractional. We now outline the algorithm that decides the content of each source node. We use the assumption that the original source can be represented as a collection of independent equal-entropy random variables $\{OS_i\}_{i=1}^{Q}$, for large enough Q at this point. Suppose that $H(OS_1) = \beta$. In turn, we can conclude that there exist integers $\alpha_A, \forall A \in \mathcal{A}$, such that $\alpha_A \times \beta = \mu^*(A), \forall A \in \mathcal{A}$ and that $\sum_{A \in \mathcal{A}} \alpha_A = Q$. Consider an ordering of the atoms, denoted as $A_1, A_2, \cdots, A_{2^n - 1}$. The atom random variables can then be assigned as follows: For each A_i, assign $W_{A_i} = (OS_{\sum_{j<i} \alpha_{A_j} + 1}, OS_{\sum_{j<i} \alpha_{A_j} + 2}, \ldots, OS_{\sum_{j \leq i} \alpha_{A_j}})$. It is clear that the resultant atom random variables are independent and that $H(W_A) = \mu^*(A), \forall A \in \mathcal{A}$. Then $X_i = (W_A : A \subset \widetilde{X}_i)$.

The assumption on the original source is essentially equivalent to saying that a large file can be broken into arbitrarily small pieces. To see this assume that each edge in the network has a capacity of 1000 bits/sec. At this time-scale, suppose that we treat each edge as unit-capacity. If the smallest unit of a file is a single bit, then we can consider it to be consisting of sources of individual entropy equal to 10^{-3}.

C. Coded source network

Given the same network, if we allow coded information stored at the sources, using the augmented graph G^*, the storage at the sources can be viewed as the transmission along the edges connecting the virtual source and real sources. Then the problem becomes the standard minimum cost multicast with network coding problem (CODED-MIN-COST) where the variables are only the flows z_{ij} and $x_{ij}^{(t)}$.

$$\text{minimize } \textstyle\sum_{(i,j) \in E} f_{ij} z_{ij} + \sum_{i \in S} d_i z_{s^* i}$$

$$\text{subject to } 0 \leq x_{ij}^{(t)} \leq z_{ij} \leq c_{ij}^*, \ (i,j) \in E^*, t \in T$$

$$\sum_{\{j | (i,j) \in E^*\}} x_{ij}^{(t)} - \sum_{\{j | (j,i) \in E^*\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \ i \in V^*, \ t \in T$$

Assume we have the solution for CODED-MIN-COST, we can use the random coding scheme introduced by [8] to construct the sources and the flow of each edge.

IV. Cost comparison between the coded case and the subset case

For given instances of the problem, we can certainly compute the cost gap by solving the corresponding optimization problems SUBSET-MIN-COST or the alternative formulation in [7] and CODED-MIN-COST. Because the subset case is a special case of the coded case, we define the cost gap as the difference between the optimums of the subset case and the coded case. In this section, we first formulate an optimization problem similar to SUBSET-MIN-COST. The main difference is that we consider the source node can contain any arbitrary functions of the pieces of the original source. Accordingly, we require the atoms to satisfy the information inequalities [6] that consist of Shannon type inequalities and non-Shannon type inequalities when $n \geq 4$ [9]. Because there are infinitely many non-Shannon type inequalities, it is impossible to list all the information inequalities when the source number exceeds 4. However, if we remove the non-Shannon type inequalities from the constraints, the optimal value of the coded case will not increase. In turn, this means that the gap computed by comparing these optimal values will still be a valid upper bound for the gap between the subset and coded cases.

Following this, we can find an upper bound on the cost gap as the solution to another optimization problem. In the general case, of n sources, even this optimization has constraints that are exponential in n. However, this formulation still has advantages. We are able to prove that there is a closed form upper bound in the case of three sources, which can be shown to be tight, i.e., there exist instances such that the cost gap is met with equality.

A. General case

We now present the new coded formulation ATOM-CODED-MIN-COST using the augmented graph G^*.

minimize $\sum_{(i,j)\in E} f_{ij} z_{ij} + \sum_{i\in S} d_i z_{s^* i}$

subject to $0 \le x_{ij}^{(t)} \le z_{ij} \le c_{ij}^*, \forall (i,j) \in E^*, t \in T$

$$\sum_{\{j|(i,j)\in E^*\}} x_{ij}^{(t)} - \sum_{\{j|(j,i)\in E^*\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \forall i \in V^*,\ t \in T \quad (7)$$

$$x_{s^* i}^{(t)} \ge R_i^{(t)}, \forall i \in S, t \in T \quad (8)$$

$$R^{(t)} \in \mathcal{R}_{\mathcal{SW}}, \forall t \in T \quad (9)$$

$$H(X_i | X_{S\setminus\{i\}}) \ge 0, \forall i \in S \quad (10)$$

$$I(X_i; X_j | X_K) \ge 0, \forall i \in S, j \in S, i \ne j, K \subseteq S\setminus\{i,j\} \quad (11)$$

$$z_{s^* i} = H(X_i), \forall i \in S \quad (12)$$

$$H(X_1, X_2 \cdots, X_n) = \sum_{A\in\mathcal{A}} \mu^*(A) \quad (13)$$

The formulation is the same as SUBSET-MIN-COST except that we remove (4), and add constraints (10) and (11) that are elemental inequalities, which guarantee that all Shannon type inequalities are satisfied [6]. The elemental inequalities can be represented in the form of atoms, $\forall K \subseteq S \setminus \{i,j\}$:

$$H(X_i|X_{S\setminus\{i\}}) = \mu^*(A), A \not\subseteq \widetilde{X}_{S\setminus\{i\}}$$
$$I(X_i;X_j|X_K) = \sum_{A\in\mathcal{A}: A\subset\widetilde{X}_i, A\subset\widetilde{X}_j, A\not\subseteq\widetilde{X}_K} \mu^*(A)$$

Now, suppose that we know the optimal value of the above optimization problem, i.e., the flows $x_{ij,1}^{(t)}, z_{ij,1}^{(t)}, t \in T, (i,j) \in E^*$, the measure of the atoms $\mu^*(A)_1, \forall A \in \mathcal{A}$, and the corresponding conditional entropies $H_1(X_U|X_{S\setminus U}), \forall U \subseteq S$. It can be shown that we can construct a feasible solution for SUBSET-MIN-COST such that the flows over E^* are the same as $x_{ij,1}^{(t)}$(and $z_{ij,1}^{(t)}$), $t \in T, (i,j) \in E$, then we can arrive at an upper bound for the gap. This is done below. Let $\mu^*(A), \forall A \in \mathcal{A}$ denote the variables for the atom measures for the subset case. We have the gap LP,

$$\min \sum_{A\in\mathcal{A}} \Big(\sum_{\{i\in S: A\subset\widetilde{X}_i\}} d_i\Big)\mu^*(A) - \sum_{A\in\mathcal{A}} \Big(\sum_{\{i\in S: A\subset\widetilde{X}_i\}} d_i\Big)\mu^*(A)_1$$

$$\text{subject to} \sum_{A: A\not\subseteq\widetilde{X}_{S\setminus U}} \mu^*(A) \le H_1(X_U|X_{S\setminus U}), \forall U \subset S \quad (14)$$

$$\mu^*(A) \ge 0, \forall A \in \mathcal{A}$$

$$\sum_{A: A\in\mathcal{A}} \mu^*(A) = H(X_1, X_2, \cdots, X_n)$$

where $H_1(X_U|X_{S\setminus U}) = \sum_{A: A\not\subseteq\widetilde{X}_{S\setminus U}} \mu^*(A)_1, \forall U \subset S$. In the SUBSET-MIN-COST, we assign $x_{ij}^{(t)} = x_{ij,1}^{(t)}, (i,j) \in E^*$, $z_{ij}^{(t)} = z_{ij,1}^{(t)}, (i,j) \in E$ and $z_{s^* i} = \sum_{A: A\subset\widetilde{X}_i} \mu^*(A), \forall i \in S$.

To see that this is feasible, note that

$$\begin{aligned} z_{s^* i} &= \sum_{A: A\subset\widetilde{X}_i} \mu^*(A) = H(X_i) \\ &= H(X_1, \cdots, X_n) - H(X_1, \cdots, X_{i-1}, X_{i+1}, \cdots, X_n | X_i) \\ &\overset{(a)}{\ge} H(X_1, \cdots, X_n) - H_1(X_1, \cdots, X_{i-1}, X_{i+1}, \cdots, X_n | X_i) \\ &= H_1(X_i) = z_{s^* i,1} \ge x_{s^* i,1}^{(t)} = x_{s^* i}^{(t)} \end{aligned}$$

Then constraint (1) is satisfied.

$$\sum_{i: i\in U} x_{s^* i}^{(t)} = \sum_{i: i\in U} x_{s^* i,1}^{(t)} \ge H_1(X_U|X_{S\setminus U}) \overset{(b)}{\ge} H(X_U|X_{S\setminus U})$$

Then constraints (2) and (3) are satisfied.

Both (a) and (b) come from constraint (14). The difference in the costs is only due to the different storage costs, since the flow costs are exactly the same.

B. Three sources case

The case of three sources is special because, (i) Non-Shannon type inequalities do not exist for three random variables. This implies that we can find three random variables using the atom measures solution of ATOM-CODED-MIN-COST. (ii) Moreover, there is at most one atom, $\mu^*(\widetilde{X}_1 \cap \widetilde{X}_2 \cap \widetilde{X}_3)$ that can be negative. Let the atom measures found by solving ATOM-CODED-MIN-COST be denoted by the variables shown in Figure 2(a).

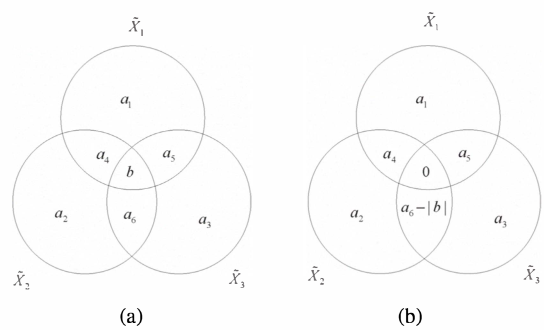

Fig. 2. (a) The coded case, (b) The corresponding subset case.

Claim 1: Consider random variables X_1, X_2 and X_3 with $H(X_1, X_2, X_3) = h$. Then, $b \ge -\frac{h}{2}$.

proof: The proof is omitted. ■

Using this we can obtain the following lemma:

Lemma 1: Supose that we have three source nodes. Let the joint entropy of the original source be h and let f_{opt2} represent the optimal value of SUBSET-MIN-COST and f_{opt1}, the optimal value of CODED-MIN-COST. Then, $f_{opt2} - f_{opt1} \le (\min_{i\in S}(d_i))h/2$.

proof: Without loss of generality, assume that $\min_{i\in S}(d_i) = d_1$. Suppose that in the optimal solution for ATOM-CODED-MIN-COST, $b \le 0$. We construct a feasible solution for SUBSET-MIN-COST by keeping the flow values the same, but changing the atom values suitably. Let $a_i', i = 1, \ldots, 6, b'$

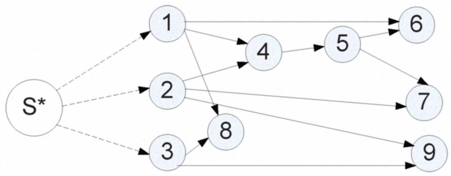

Fig. 3. Network with source nodes at 1, 2 and 3; terminals at 6, 7, 8 and 9. Append a virtual source S^* connecting real sources.

denote the atom values for the subset case. Consider the following assignment,

$$a'_i = a_i, i = 1, \dots, 5.\ a'_6 = a_6 - |b|, \text{ and } b' = 0.$$

which is shown pictorially in Figure 2. We can check constraint (14) to see that the solution is feasible for the gap LP for three sources. By further checking the KKT condition, we conclude that the solution is the optimal solution for the gap LP. The flows do not change over transforming the coded case to the subset case. The only cost increased is $d_1 \times |b| \leq d_1 \times h/2$. In the results section, we will show an instance of a network where this upper bound is tight. ■

Finally we note that when there are only two source nodes, there is no cost difference between the subset case and the coded case, since for two random variables, all atoms have to be nonnegative. We state this as a lemma below.

Lemma 2: Suppose that we have two source nodes. Let f_{opt2} represent the optimal value of SUBSET-MIN-COST and f_{opt1}, the optimal value of CODED-MIN-COST. Then, $f_{opt2} = f_{opt1}$.

V. Results

In this section we present an example of a network with three sources where our upper bound derived in Section IV-B is tight. We also performed several experiments with randomly generated graphs to study whether the difference in cost between the two cases occurs very frequently.

Consider the network in Figure 3 with three sources nodes, 1, 2 and 3 and four terminal nodes, 6, 7, 8, and 9. The entropy of the original source = $H(X_1, X_2, X_3) = 2$ and all edges are unit-capacity. The costs are such that $f_{ij} = 1, \forall (i, j) \in E$ and $d_1 = d_2 = 2, d_3 = 1$.

The optimal cost in the subset sources case is 17. The corresponding atom values are: $a_4 = 0.5809$, $a_5 = 0.6367$, $a_6 = 0.7824$ and other atoms have measures 0. In this case we have $H(X_1) = 1.22, H(X_2) = 1.36$ and $H(X_3) = 1.42$.

In the coded sources case, the optimal value is 16, with $H(X_1) = H(X_2) = H(X_3) = 1$. In this case the gap between the optimal values is precisely = $\frac{2}{2} \times 1 = 1$, i.e., the upper bound derived in the previous section is met with equality.

We generated several directed graphs at random with $|V| = 87$, $|E| = 322$. The linear cost of each edge was fixed to an integer in $\{1, 2, 3, 4, 5, 6, 29, 31\}$. We ran 5000 experiments with fixed parameters $(|S|, |T|, h)$, where $|S|$ - number of source nodes, $|T|$ - number of terminal nodes and h - entropy of the original source. The locations of the source and terminal nodes were chosen randomly. The capacity of each edge was chosen at random from the set $\{1, 2, 3, 4, 5\}$. In many cases the network did not have enough capacity to support the recovery at the terminals. These instances were discarded.

TABLE I
Comparisons of two schemes in 5000 random directed graphs

$(\lvert S\rvert, \lvert T\rvert, h)$	(3, 3, 3)	(4, 4, 4)	(5, 5, 5)	(4, 5, 5)	(5, 4, 5)	(4, 4, 5)
Equal	3893	2855	1609	1577	2025	1954
Nonequal	1	3	10	9	6	8

The results are shown in Table I. The "Equal" row corresponds to the number of instances when both the coded and subset cases have the same cost, and "Nonequal" corresponds to the number of instances where the coded case has a lower cost. Note that in most cases, the two cases have the exact same cost. We also evaluated the gap LP using random graphs. Note that the gap LP is only an upper bound since it is derived assuming that the flow patterns do not change between the coded case and the subset case. When $(|S|, |T|, h) = (4, 3, 4)$, among 5000 experiments, 3269 instances could support both cases. Out of these, there were 481 instances where the upper bound determined by the gap LP was not tight.

VI. Conclusion

In this work, we considered network coding based content distribution, under the assumption that the content can be considered as a collection of independent equal entropy sources. Given a network with a specified set of source nodes, we minimize the joint cost of transmission and storage for the subset sources case and the coded source case. We provided succinct formulations of the corresponding optimization problems by using the properties of information measures. In particular, we show that when there are two source nodes, there is no loss in considering subset sources. For three source nodes, we derive a tight upper bound on the cost gap between the two cases. A gap LP for estimating the cost gap for a given instance was provided. Finally, we also provided algorithms for determining the content of the source nodes.

Our results indicate that when the number of source nodes is small, in many cases constraining the source nodes to only contain subsets of the content does not incur a loss.

References

[1] D. S. Lun, N. Ratnakar, M. Médard, R. Koetter, D. R. Karger, T. Ho, E. Ahmed, and F. Zhao, "Minimum-Cost Multicast over Coded Packet Networks," *IEEE Trans. on Info. Th.*, vol. 52, pp. 2608–2623, June 2006.

[2] A. Lee, M. Médard, K. Z. Haigh, S. Gowan, and P. Rubel, "Minmum-Cost Subgraphs for Joint Distributed Source and Network Coding," in *the Third Workshop on Network Coding, Theory, and Applications*, Jan. 2007.

[3] A. Ramamoorthy, "Minimum cost distributed source coding over a network," in *IEEE Intl. Symposium on Info. Th.*, 2007, pp. 1761–1765.

[4] A. Jiang, "Network Coding for Joint Storage and Transmission with Minimum Cost," in *IEEE Intl. Symposium on Info. Th.*, 2006, pp. 1359–1363.

[5] K. Bhattad, N. Ratnakar, R. Koetter, and K. Narayanan, "Minimal Network Coding for Multicast," in *IEEE Intl. Symposium on Info. Th.*, 2005, pp. 1730–1734.

[6] R. Yeung, *Information Theory and Network Coding.* Springer, 2008.

[7] S. Huang, in *MS Thesis, Iowa State Univeristy.* [Online]. Available: http://www.ece.iastate.edu/~hshurui/thesisdraft.pdf

[8] T. Ho, M. Médard, J. Shi, M. Effros, and D. R. Karger, "On Randomized Network Coding," in *41st Allerton Conference on Communication, Control, and Computing*, 2003.

[9] Z. Zhang and R. W. Yeung, "A non-Shannon-type Condtional Inequality of Information Quantities," *IEEE Trans. on Info. Th.*, vol. 43, pp. 1982–1986, 1997.