

NOVA University of Newcastle Research Online

nova.newcastle.edu.au

Ong, Lawrence; Ho, Chin Keong; Lim, Fabian "The single-uniprior index-coding problem: the single-sender case and the multi-sender extension". Originally published in IEEE Transactions on Information Theory Vol. 62, Issue 6, p. 3165-3182

Available from:

http://dx.doi.org/10.1109/TIT.2016.2555950

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Accessed from: http://hdl.handle.net/1959.13/1323828

Lawrence Ong, Chin Keong Ho, Fabian Lim

Abstract—Index coding studies multiterminal source-coding problems where a set of receivers are required to decode multiple (possibly different) messages from a common broadcast, and they each know some messages a priori. In this paper, at the receiver end, we consider a special setting where each receiver knows only one message a priori, and each message is known to only one receiver. At the broadcasting end, we consider a generalized setting where there could be multiple senders, and each sender knows a subset of the messages. The senders collaborate to transmit an index code. This work looks at minimizing the number of total coded bits the senders are required to transmit. When there is only one sender, we propose a pruning algorithm to find a lower bound on the optimal (i.e., the shortest) index codelength, and show that it is achievable by linear index codes. When there are two or more senders, we propose an appending technique to be used in conjunction with the pruning technique to give a lower bound on the optimal index codelength; we also derive an upper bound based on cyclic codes. While the two bounds do not match in general, for the special case where no two distinct senders know any message in common, the bounds match, giving the optimal index codelength. The results are expressed in terms of strongly connected components in directed graphs that represent the index-coding problems.

I. INTRODUCTION

We investigate a broadcast problem over noiseless channels with receiver side information, also known as index coding [1]. In the classical setup, *one* sender encodes a collection of messages, and broadcasts the codeword to multiple receivers. Each receiver knows some messages a priori, and is to decode a set of messages it wants from the single codeword broadcast by the sender. The aim is to find the shortest codeword that the sender needs to broadcast to ensure that all receivers can decode the messages they want. The index-coding problem remains open to date; only a small number of special cases have been solved [1]–[6].

In this paper, we consider a class of the index-coding problem, which we refer to as *single-uniprior*, where each receiver knows only one message a priori, but may request multiple messages, and each message is known to only one receiver. When there is only *one* sender, we completely solve the single-uniprior index-coding problem. We show that linear index codes are optimal for this class—although linear index

Lawrence Ong is the recipient of an Australian Research Council Future Fellowship (FT140100219).

codes are suboptimal in general [7], e.g., when two receivers know the same message.

1

We then extend the single-uniprior problem with a single sender to multiple senders, where each sender knows a subset of the messages. We derive lower and upper bounds on the shortest index codelength, and identify cases where the bounds match. In particular, the bounds match when no two senders know any message in common.

A. Motivation of the multi-sender single-uniprior index-coding problem

The single-uniprior index-coding problem formulation is motivated by satellite communications [8]–[10], where multiple clients exchange messages through a satellite, which acts as a relay. Having no direct communication links, the clients first send their messages to the satellite via an *uplink* channel. The satellite decodes the messages, re-encodes, and broadcasts to the clients via a *downlink* channel. Here, the downlink corresponds to a single-uniprior index-coding problem, where each client is to obtain messages of other clients, and it knows only its own message a priori.

When multiple satellites are used, the downlink corresponds to a multi-sender index-coding problem. In this case, different senders know different messages a priori, due to, for instance, decoding errors when getting the messages from the clients on the uplink, line-of-sight only to some clients, or limited storage of the senders. To the best of our knowledge, this is the first paper that investigates the multi-sender index-coding problem.

II. CHANNEL MODEL AND NOTATION

An *instance* of the multi-sender index-coding problem (including the single-sender as a special case) consists of the following:

- *m* independent messages, denoted by an ordered¹ set $\mathcal{M} = \{x_1, x_2, \dots, x_m\}.$
 - Each message x_i has q_i binary bits, denoted by $x_i = [x_i[1], x_i[2], \ldots, x_i[q_i]]$, where each bit $x_i[\cdot]$ is drawn from an independent and uniform distribution over $\{0, 1\}$.
- S senders: each sender s ∈ {1, 2, ..., S} knows a subset of the messages, denoted by an ordered¹ subset M_s ⊆ M.
- *n* receivers: each receiver $r \in \{1, 2, \dots, n\}$
 - knows an ordered¹ subset of the messages a priori, $\mathcal{K}_r \subset \mathcal{M}$, and

¹The elements are ordered in increasing indices.

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Part of the material in this paper was presented at the IEEE International Symposium on Information Theory, Istanbul, Turkey, July 7–12, 2013, and at the IEEE International Conference on Communications, Ottawa, Canada, June 10–15, 2012.

This work was done when Fabian Lim was with Massachusetts Institute of Technology under the support of NSF grant ECCS-1128226.

Senders



Fig. 1. The multi-sender multicast single-uniprior index coding: We consider distributed-like settings where each sender s is limited to know a subset \mathcal{M}_s of the message set \mathcal{M} . Also, we consider a multicast setup where each receiver r knows a single unique message x_r a priori, and wants a subset of messages $\mathcal{W}_r \subset \mathcal{M}$. Each sender s broadcasts c_s to all receivers noiselessly. The aim is to find the minimum aggregated codelength $|c_1| + |c_2| + |c_3|$.

- wants/requests an ordered¹ subset of the messages, $\mathcal{W}_r \subseteq \mathcal{M}.$

Figure 1 depicts an example of a multi-sender single-uniprior index-coding instance, where $\mathcal{K}_i = \{x_i\}$.

Without loss of generality, we assume that $\mathcal{W}_r \cap \mathcal{K}_r = \emptyset$ for any r, and that $\bigcup_{s=1}^{S} \mathcal{M}_s = \mathcal{M}$, i.e., each message is available at some sender(s).

We define a multi-sender index code for the above setup:

Definition 1 (Multi-Sender Index Code): An index code for an index-coding instance consists of

1) an encoding function for each sender $s \in \{1, 2, \dots, S\}$,

$$\mathbf{E}_s: \{0,1\}^{\sum_{i:x_i \in \mathcal{M}_s} q_i} \mapsto \{0,1\}^{\ell_s},$$

2) a decoding function for each receiver $r \in \{1, 2, ..., n\}$,

$$\mathbb{D}_r: \{0,1\}^{\sum_{s=1}^S \ell_s + \sum_{i:x_i \in \mathcal{K}_r} q_i} \mapsto \{0,1\}^{\sum_{i:x_i \in \mathcal{W}_r} q_i},$$

such that $\mathcal{W}_r = D_r(E_1(\mathcal{M}_1), E_2(\mathcal{M}_2), \dots, E_S(\mathcal{M}_S), \mathcal{K}_r)$ for each receiver r.

This means each sender s encodes its known messages into an ℓ_s -bit sub-codeword. The sub-codewords of all senders are given to all receivers. The total number of transmitted bits is $\tilde{\ell} \triangleq \sum_{s=1}^{S} \ell_s$, which is the index codelength. We seek the optimal (i.e., the minimum) index codelength, denoted by ℓ^* , as well as index codes with the optimal length. If there is only one sender, S = 1, we simplify the notation by denoting the optimal index codelength as ℓ^* .

The case S = 1 reduces to the single-sender index-coding problem studied in many works [1]-[7], [11]-[24]. Clearly, for the same receiver setting $\{\mathcal{W}_r, \mathcal{K}_r\}_{r=1}^n$, a multi-sender index code for any S > 1 case will also be a code for the singlesender S = 1 case. This is because, by definition, the only sender in the single-user case, who knows all the messages, can transmit any multi-sender index code. But the converse is not true. So, techniques for the single-sender case do not straightforwardly apply to the multi-sender case.

For both the single-sender and the multi-sender indexcoding problems, we assume that each sender knows which messages (not the message content) each receiver knows, \mathcal{K}_r , and requests, W_r ; and which messages each sender has, M_s . This can be achieved by having the nodes to communicate (the message-availability and message-request information) prior to index coding. Under such scenarios, index coding is useful if the messages are sufficiently long such that the overhead incurred during the prior communication is negligible.² In some communication scenarios, side information is strategically placed at the receivers (for example, data pre-fetching or caching [25]), and the messages in the senders' storage are populated according to their geographical locations. In such cases, the senders can be informed of all message indices in $\mathcal{K}_r, \mathcal{W}_r$, and \mathcal{M}_s .

If we set the number of senders to one, S = 1, and the messages to be of equal length, $q_i = q$ for all i, the optimal normalized index codelength-over all index codes and all message lengths-is commonly known as the optimal broadcast rate [4]:

$$\beta \triangleq \inf_{q} \frac{\ell^*}{q} = \lim_{q \to \infty} \frac{\ell^*}{q} \text{ codeword length per message length.}$$
(1)

The second equality is due to the sub-additivity property of ℓ^* [12].

A. Terminology

We will represent index-coding instances by graphs, and derive bounds on ℓ^* by executing graph operations. We will use common graph terminology: A directed path of length L from vertex v_0 to vertex v_L consists of the following (L+1)distinct vertices (except possibly the first and last) and L arcs: $\{v_0, (v_0 \to v_1), v_1, (v_1 \to v_2), \dots, (v_{L-1} \to v_L), v_L\}$. A cycle is a path with the same first and last vertices. A strongly connected component (SCC) of a graph is a maximal subgraph such that for any vertex pair (i, j) in the subgraph, there is a directed path in the subgraph from i to j, and another from j to i. A trivial SCC is an SCC with only one vertex. A *leaf* vertex in a graph has no outgoing arcs. Similarly, a leaf non-trivial SCC is an SCC with two or more vertices, in which there is no arc from any vertex in the SCC to any vertex outside the SCC. In other words, for each vertex in the SCC, its out-neighbors are all vertices in the SCC. A vertex jis a *predecessor* of vertex *i* if and only if there is a directed path from j to i. A path on an undirected graph is similarly defined with arcs replaced by edges. A subgraph is *connected* if there is a path within the subgraph from any vertex to any other vertex. A tree is a connected undirected subgraph with no cycle.

III. INDEX-CODING CLASSIFICATION AND GRAPHICAL REPRESENTATION

A single-sender instance is fully described by the message lengths $\{q_i\}_{i=1}^m$ and the receiver setting $\{\mathcal{W}_r, \mathcal{K}_r\}_{r=1}^n$. For the multi-sender extension, we further require the sender setting $\{\mathcal{M}_s\}_{s=1}^S.$

²For simplicity, in the multi-sender section of this paper, we assume that each message consists of a single bit. However, all proofs are also valid even when each message consist of q bits, for any positive integer q.

A. Classification of index-coding instances based on the receiver setting

We first consider the single-sender setup, and categorize the instances as follows:

1) Information flow: We classify different types of information flow from the sender to the receivers. We say that an index-coding instance is *unicast* if

$$\mathcal{W}_i \cap \mathcal{W}_j = \emptyset, \quad \forall i \neq j,$$
 (2)

meaning that each message can be requested by at most one receiver. In addition, we say that the instance is *single-unicast* if, in addition to (2), we also have that $|W_i| = 1$ for all *i*, meaning each receiver requests exactly one message—but the message may consist of multiple bits.

Note that any unicast instance can be recast as a singleunicast instance [13]. Suppose that a receiver i wants two messages $W_i = \{x_1, x_2\}$, and knows \mathcal{K}_i . As far as index codes are concerned, we can replace receiver i by two receivers i'and j', both knowing the same set \mathcal{K}_i , and each wanting a single message, i.e., x_1 and x_2 respectively. We can further split the receivers in a similar manner so that each message contains only a single bit.

Unicast instances were investigated by Bar-Yossef et al. [1], Lubetzky and Stav [7], Shanmugam et al. [18], [19], Arbabjolfaei et al. [3], [20], [23], Yu and Neely [6], Wong et al. [26], Ebrahimi and Siavoshani [21], and Ong et al. [5], [22], [24].

2) *Side information:* We next classify different types of side information at the receivers. We say that an index coding instance is *uniprior* if

$$\mathcal{K}_i \cap \mathcal{K}_j = \emptyset, \quad \forall i \neq j,$$
 (3)

meaning that each message is known a priori to at most one receiver. In addition, we say that the instance is *single-uniprior* if, in addition to (3), we also have that $|\mathcal{K}_i| = 1$ for all *i*, meaning that each receiver knows exactly one unique message a priori.

Unlike unicast instances, single-uniprior instances do not subsume all uniprior instances. A receiver who knows $\{x_1, x_2\}$ is not equivalent to two receivers each knowing only one of them.

Unicast uniprior instances were investigated by Neely et al. [2].

With the above terminology, we call index-coding instances with no restrictions on W_i and \mathcal{K}_i multicast multiprior instances. Multicast multiprior instances were investigated by Blasiak et al. [4], Alon et al. [12], Shanmugam et al. [20], Tehrani et al. [16], and Neely et al. [2].

B. Graphical representation of the receiver setting

Existing works on index coding focused on the single-sender case, and many used the following graphical representations to capture the receivers' side information and requests.



(a) Side-information graph G_1 (b) Side-information graph G_2 Fig. 2. Incorrect side-information graphs for the uniprior instance in Example 1

1) Unicast multiprior (or simply unicast): As mentioned above, any unicast instance can be converted into an equivalent single-unicast instance. Without loss of generality, we can assume exactly n receivers and n messages, as message not requested by any receivers can be removed. Single-unicast index-coding instances are commonly represented by *sideinformation graphs* [1] with n vertices $\{1, 2, ..., n\}$, where an arc exists from vertex i to vertex j if and only if receiver iknows the message requested by receiver j.

2) Unicast uniprior: The class of index-coding instances with unicast and uniprior was investigated by Neely et al. [2], where each message bit is (i) known to only one receiver, and (ii) requested by only one receiver. Neely et al. represented this class by a weighted compressed graph, where each vertex represents a receiver, and an arc of weight q exists from vertex i to j if and only if receiver j wants q messages known to receiver i.

Although a unicast instance can be recast into a singleunicast instance, not all unicast uniprior instances can be recast into single-unicast uniprior instances. This is because if we split a receiver, there will be more than one receivers knowing the same message (i.e., the instance is no longer uniprior).

3) Multicast multiprior: This is the most general class of the index-coding instances, where there is no restriction on W_i and \mathcal{K}_i . To completely capture the information in both W_i and \mathcal{K}_i , Neely et al. [2] proposed a bipartite-graph representation, where there are *n* receiver vertices, $\{r_1, r_2, \ldots, r_n\}$, and *m* messages vertices, $\{x_1, x_2, \ldots, x_m\}$. An arc from r_i to x_j exists if and only if receiver *i* wants x_j , and arc from x_i to r_j exists if and only if receiver *j* knows x_i . Blasiak et al. [4] and Alon et al. [12] used hypergraphs to represent multicast multiprior instances.

4) Multicast single-uniprior (or simply single-uniprior): This is the class of index-coding instances considered in this paper. Without loss of generality, we consider n receivers and n messages. This subsumes the case where a message x_j is not known to any receiver. To see this, add a dummy receiver n+1who knows x_j but does not want any message. Clearly, adding this receiver does not change the problem.

Side-information graphs used for single-unicast instances cannot capture all single-uniprior instances. We illustrate this using the example below:

Example 1: Consider the following uniprior instance: $\mathcal{K}_1 = \{x_2\}, \mathcal{K}_2 = \{x_1\}, \mathcal{K}_3 = \{x_4\}, \mathcal{K}_4 = \{x_3\}, \mathcal{W}_1 = \{x_3\}, \mathcal{W}_2 = \{x_4\}, \mathcal{W}_3 = \{x_1, x_2\}, \mathcal{W}_4 = \{x_1, x_2\}, \text{ where each message } x_i \text{ contains one bit. We will see later (Corollary 1) that the optimal index codelength for this instance is three bits. If we apply the rule for constructing side-information graphs that says "<math>i \rightarrow j$ exists if receiver i knows the message requested by receiver j", we get graph \mathcal{G}_1 shown in Figure 2. For a unicast instance represented by \mathcal{G}_1 , the optimal codelength is two bits (using clique cover [11]). Thus, \mathcal{G}_1 incorrectly represents the uniprior instance described here. However, since receivers 3 and 4 request multiple messages, we can obtain an equivalent instance by splitting receiver 3 to two receivers (say, 3 and 5), each knowing the same message but requesting a different message: $\mathcal{K}_3 =$ $\{x_4\}, \mathcal{K}_5 = \{x_4\}, \mathcal{W}_3 = \{x_1\}, \mathcal{W}_5 = \{x_2\}.$ We similarly split receiver 4: $\mathcal{K}_4 = \{x_3\}, \mathcal{K}_6 = \{x_3\}, \mathcal{W}_4 = \{x_1\}, \mathcal{W}_6 = \{x_2\}.$ Constructing a side-information graph for this equivalent instance, we get \mathcal{G}_2 in Figure 2. But for a unicast instance represented by \mathcal{G}_2 , the optimal codelength is four bits (again, using clique cover). Thus \mathcal{G}_2 again incorrectly represents the uniprior instance in this example. The reason is that sideinformation graphs are not able to capture multiple receivers requesting a same message.

Instead, we propose *information-flow graphs* with *n* weighted vertices. Without loss of generality, let $\mathcal{K}_i = \{x_i\}$ for all receiver $i \in \{1, 2, ..., n\}$. An arc exists from vertex *i* to vertex *j* if and only if node *j* wants the message known to receiver *i*, i.e., $x_i \in W_j$. The weight of vertex *i* is q_i , which is the number of bits in message x_i . Although a uniprior instance can be represented by either a hypergraph or a bipartite graph, using the information-flow graph results in a smaller (compare to the bipartite graph) and simpler (a directed graph instead of a hyper graph) graph.

C. Graphical representation of the sender setting

When there are multiple senders, we need to represent $\{\mathcal{M}_s\}$ in addition to $\{q_i, \mathcal{K}_r, \mathcal{M}_r\}$. We propose to represent $\{\mathcal{M}_s\}$ using an *n*-vertex undirected *message* graph, denoted by \mathcal{U} . An edge between vertices *i* and *j* on \mathcal{U} (denoted by (i, j)) exists if and only if messages x_i and x_j are known to the same sender, i.e., $i, j \in \mathcal{M}_s$ for some *s*.

In this paper, we will consider multi-sender single-uniprior instances each described by an information-flow graph \mathcal{G} and a message graph \mathcal{U} , both defined on the same set of vertices $\mathcal{V} = \{1, 2, ..., n\}$. We denote by $\tilde{\ell}^*(\mathcal{G}, \mathcal{U})$ the optimal index codelength for the instance represented by $(\mathcal{G}, \mathcal{U})$.

Note that an edge (i, j) in \mathcal{U} does not indicate which sender(s) owns both the messages (i.e. the set of *s* such that $i, j \in \mathcal{M}_s$). This ambiguity will not affect the techniques developed in this paper. In fact, for some cases, this is sufficient to derive the optimal index codelength. However, we will point out in the conclusion the existence of an instance where resolving this ambiguity can further improve our results.

IV. EXISTING RESULTS AND OUR CONTRIBUTIONS

In this section, we will survey existing results on ℓ^* for the single-sender problem. We will then review related works on multiterminal networks with side information that have settings similar to that in the multi-sender index-coding problem. Lastly, we present the main results of this paper.

A. Existing results for single-sender index coding

Recall that any unicast instance can be first converted to an equivalent single-unicast instance with binary messages and then represented by a side-information graph. For binary messages, Bar-Yossef et al. [1] found the optimal index codelength ℓ^* for all acyclic side-information graphs. Suppose that a side-information graph has cycle(s) and is *symmetrical* (meaning that for every arc from vertex *i* to vertex *j*, there is another arc from vertex *j* to vertex *i*). For this special case, the directed side-information graph can be converted to an undirected graph where an edge between vertices *i* and *j* represent the arcs in both directions. Bar-Yossef et al. solved all unicast index-coding instances with the following undirected side-information graphs: (i) perfect graphs, (ii) odd holes with five or more vertices, and (iii) odd anti-holes with five or more vertices.

Arbabjolfaei et al. [3] found ℓ^* for all unicast instances up to five receivers. They used random-coding arguments to prove achievability. Later, Ong [5] showed that binary linear codes are sufficient to achieve ℓ^* for all unicast instances up to five receivers. Blasiak et al. [4] and Thapa et al. [24] found ℓ^* for certain special classes of graphs.

Also for unicast instances, using the bipartite-graph representation, Yu and Neely [6] showed that if the graph is planar,³ then ℓ^* can be found using linear programming (solutions not in closed form). This class of unicast instances subsume all unicast uniprior instances up to four receivers as special cases. Neely et al. [2] found ℓ^* for unicast uniprior instances if the weighted compressed graph contains only arc-disjoint cycles.

For the most general multicast multiprior case, Neely et al. [2] found ℓ^* for all acyclic bipartite graphs.

It has been shown [14] that (i) the general multicast multiprior index-coding problem is NP-hard, and (ii) the multicast (non-unicast) index-coding problem is even NP-hard to approximate.

For all the above-mentioned classes of index-coding instances where ℓ^* has been found, *linear* index codes are optimal. However, for some single-unicast instances, Lubetzky and Stav [7] showed that non-linear index codes can outperform linear codes.

B. Related works on multiterminal networks with side information

In the literature, a problem setup related to the multi-sender index-coding problem is multiterminal networks where a set of clients exchange messages directly through a shared noiseless medium. The aim is to minimize the total transmission *cost* (e.g., total number of transmissions) while satisfying all clients' requirements. This problem can be viewed as a special case of the multi-sender index-coding problem where for each client, there is a sender having the same messages.

Ozgul and Sprintson [27] considered this setup where each client requests *all the messages it does not know*. They used multiple rounds of random linear coding to achieve the optimal transmission cost with high probability.

Hou, Hsu, and Sprintson [28] considered the same setup where each client requests *one message (i.e. unicast)*, and it is *selfish* in the sense that it tries to minimize its own

 $^{{}^{3}}A$ graph is planar if it can be drawn on a two-dimensional plane in such a way that its edges intersect only at their endpoints.



Fig. 3. Summary of results in this paper and existing results. Note that any unicast instance can be recast as a single-unicast instance. The relative volumes are not indicative of the number of instances in the classes.

transmissions. They proposed a game-theoretic algorithm (via bidding) for transmission, and numerically showed that their algorithm is close to optimal.

Ji, Caire, and Molisch [29] also considered this setup where the side information of the clients (which admits a fixed number of bits) can be chosen, but the requests (also with a fixed number of bits) are random variables. They obtained an upper bound and a lower bound (on the optimal *rate*, i.e., total transmitted bits per side-information bits per client) within a multiplicative gap, by carefully designing the client's side information, such that all requests can be fulfilled by their transmission scheme.

Unlike the setup considered by Ji et al., the setup in this paper does not allow one to design the side information of the receivers. Also, unlike the setups by Ozgul and Sprintson, and Hou et al., we consider all possible message requests by the receivers.

C. Our contributions

1) Single-sender: In this paper, we solve all single-uniprior (including unicast and multicast) index-coding problems, and show that the solution can be found in polynomial time in the worst case.

More specifically, we characterize the minimum codelength ℓ^* of any single-uniprior index-coding instance in terms of the number of *leaf non-trivial strongly connected components* (SCC) of the corresponding information-flow graph \mathcal{G} . To this end, we design a *pruning algorithm* that removes arcs from \mathcal{G} to "destroy" all leaf non-trivial SCCs, to get a graph through which we derive a lower bound on ℓ^* . By carefully pruning \mathcal{G} , we show that the lower bound is achievable using *linear*

cyclic codes. Hence, we incidentally show that linear index codes are optimal for all single-uniprior instances. This is in stark contrast to the single-unicast instances where non-linear codes can outperform linear codes [7].

For the rest of this paper, when we say leaf SCCs, we simply mean leaf non-trivial SCCs.

2) Multi-sender: For multi-sender single-uniprior instances, we construct upper and lower bounds on the optimal index codelength based on the information-flow graph \mathcal{G} and the message graph \mathcal{U} . For the lower bound, we construct an algorithm that not only prunes \mathcal{G} (using the aforementioned pruning algorithm), but also adds arcs, edges, and vertices, and derive a lower bound on $\tilde{\ell}^*$ through the resultant graph. For the upper bound, we first define a special type of *tree*, and we count the number of such trees that we can fit into the graph. We then propose a coding scheme for each tree and some leaf SCCs to obtain a linear index code.

We will show that our upper and lower bounds match for a class of multi-sender single-uniprior instances. This class includes as a special case the scenario where no two senders know the same message, i.e., $\mathcal{M}_s \cap \mathcal{M}_t = \emptyset$ for all $s \neq t$.

3) Summary of results and techniques: Table I shows the sections in which we derive different bounds and the techniques involved. Figure 3 summarizes our results in relation to other results.

V. THE SINGLE-SENDER SINGLE-UNIPRIOR CASE

In this section we focus on the single-sender single-uniprior case. We will propose a *pruning* algorithm to obtain a lower bound on ℓ^* , and then construct a linear index code that achieves this lower bound.

For the single-uniprior case, we work on the informationflow graph, denoted by $\mathcal{G} = \{\mathcal{V}, \mathcal{A}\}$, where $\mathcal{V} = \{1, 2, ..., n\}$ is the set of vertices, and \mathcal{A} is the set of arcs. This models an index-coding instance where there are *n* receivers, and each receiver *i* knows x_i a priori and requests $\{x_j : (j \to i) \in \mathcal{A}\}$. For the remaining of this paper, we refer to the message x_i , known to receiver *i*, simply as the message of vertex *i*. Let $\ell^*(\mathcal{G})$ denote the optimal index codelength for a single-sender uniprior instance represented by \mathcal{G} .

A. Lower bound

1) Useful lemmas: We first establish a few lemmas, which we will use to establish our lower bound.

Lemma 1: Each receiver $i \in \mathcal{V}$ must be able to decode the messages of its predecessors in \mathcal{G} .

Proof: An index code always allows receiver i to decode x_j if an arc $(j \rightarrow i)$ exists. In the single-uniprior instance, x_j is all the side information receiver j has. Therefore, receiver i must also be able to decode all the messages that receiver j is able to decode, i.e., receiver i can decode all $\mathcal{W}_j = \{x_k : (k \rightarrow j) \in \mathcal{A}\}$. Further chaining of this argument proves that receiver i must be able to decode the messages of the predecessors of vertex i.

Lemma 2: Every receiver can decode the messages of all predecessors of any leaf vertex, even without utilizing its prior message.

Proof: Consider any leaf vertex i, its prior message x_i is not requested by any receiver. So, we can arbitrarily set $x_i = 0$ without affecting the decoding of any receiver, or the codelength. With this, any receiver (even one without any side information) knows $x_i = 0$, and, by Lemma 1, it must be able to decode all predecessors of the leaf vertex i. Since the choice of the leaf vertex was arbitrary, we conclude the proof.

2) A simple lower bound: Define the set of leaf vertices in \mathcal{G} as $\mathcal{L}(\mathcal{G})$. We have the following from Lemma 2:

Lemma 3: For any graph \mathcal{G} , we have

$$\ell^*(\mathcal{G}) \ge \sum_{i: i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G})} q_i \qquad (4)$$

Proof: Denote the messages of all predecessors of all leaf vertices by $\mathbf{X} \triangleq \{X_i : i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G})\},^4$ and the index code by \mathbf{C} . It follows from Lemma 2 that any receiver can decode \mathbf{X} from \mathbf{C} , i.e., $H(\mathbf{X}|\mathbf{C}) = 0$, or $H(\mathbf{X}, \mathbf{C}) = H(\mathbf{C})$, where $H(\cdot)$ is the entropy function. It follows that $\ell^*(\mathcal{G}) \ge H(\mathbf{C}) = H(\mathbf{X}, \mathbf{C}) \ge H(\mathbf{X}) \stackrel{(a)}{=} \sum_{i:i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G})} H(X_i) \stackrel{(b)}{=} Q$, where Q is the right-hand side of (4). Equality (a) is derived because the messages are independent, and (b) is derived because each message bit is uniformly distributed.

The lower bound in Lemma 3 is not useful when the graph \mathcal{G} contains not many predecessors of leaf vertices. For example, the right-hand side is zero if \mathcal{G} is a cycle. In light of this, we "process" the graph to obtain a better lower bound in the following subsection.

3) A better lower bound: We observe the following:

Lemma 4: Let $\mathcal{G} = \{\mathcal{V}, \mathcal{A}\}$ and $\mathcal{G}' = \{\mathcal{V}, \mathcal{A}'\}$, where $\mathcal{A}' \subseteq \mathcal{A}$. We have that

$$\ell^*(\mathcal{G}) \ge \ell^*(\mathcal{G}'). \tag{5}$$

Proof: Removing arcs in an information-flow graph reduces decoding requirements, while maintaining the prior messages of each receiver. This means any index code for \mathcal{G} is also an index code for \mathcal{G}' .

Combining Lemmas 3 and 4 gives the following lower bound: Lemma 5: Given any $\mathcal{G} = \{\mathcal{V}, \mathcal{A}\}$, the optimal index codelength of a single-sender single-uniprior index-coding instance represented by \mathcal{G} is lower bounded as

$$\ell^*(\mathcal{G}) \ge \max_{\mathcal{G}' = \{\mathcal{V}, \mathcal{A}'\} \text{ s.t. } \mathcal{A}' \subseteq \mathcal{A}} \sum_{i: i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G}')} q_i$$
(6)

The lower bound in Lemma 5 involves maximizing $\sum q_i$ over all arc-removed subgraphs \mathcal{G}' . We now present a way to optimally remove arcs, to attain the right-hand side of (6). To this end, we propose an algorithm to remove certain arcs to obtain a lower bound. Then, in the next section, we will show that this lower bound is indeed achievable, thereby proving that the bound obtained using this algorithm is tight, i.e., that it attains the right-hand side of (6).

We first define the following:

Definition 2: A *leaf SCC* is a *non-trivial*⁵ SCC that has no outgoing arc (i.e., from a vertex in the SCC to a vertex outside the SCC).

We now present a lower bound, which is essentially Lemma 5 evaluated with a specific G':

Theorem 1 (Single-sender: lower bound): The optimal index codelength of a single-sender single-uniprior index-coding instance is lower bounded as

$$\ell^{*}(\mathcal{G}) \geq \sum_{k \in \mathcal{V}} q_{k} - \sum_{i \in \mathcal{L}(\mathcal{G})} q_{i} - \sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_{a}, \quad (7)$$

where $\mathbb{V} \triangleq \{\mathcal{V}_{leafSCC(1)}, \mathcal{V}_{leafSCC(2)}, \dots\}$ is the set of all leaf SCCs in \mathcal{G} .

We defer the proof of Theorem 1 to Section V-A4, after introducing a pruning algorithm.

Remark 1: The lower bound here is more general than the result in our previously-published conference paper [15], for which the messages are restricted to be the same size.

4) Proving Theorem 1 using a pruning algorithm: We start with the following definitions:

Definition 3: A vertex is said to be *grounded* if it is either a leaf vertex, or a predecessor of some leaf vertex.

Definition 4: A graph is said to be grounded if every vertex in the graph is grounded.

Lemma 6: The two statements below are equivalent for any directed graph \mathcal{G} :

1) \mathcal{G} has no leaf SCC.

2) \mathcal{G} is grounded.

Proof: See Appendix A.

⁵Recall that a non-trivial SCC has two or more vertices.

We now propose an arc-removing algorithm, called the *pruning algorithm*, which we will use to prove Theorem 1.

⁴We use upper-case letters to denote random variables.

Algorithm 1: The Pruning Algorithm

 $\begin{array}{l} \textbf{input} : \text{A directed graph } \mathcal{G} = \{\mathcal{V}, \mathcal{A}\} \text{ with vertex} \\ & \text{weights } \{q_i\}_{i \in \mathcal{V}} \\ \textbf{output : A grounded directed graph} \\ \textbf{foreach Leaf SCC, say } \mathcal{G}_{\text{leafSCC}} = \{\mathcal{V}_{\text{leafSCC}}, \mathcal{A}_{\text{leafSCC}}\} \subseteq \mathcal{G} \\ \textbf{do} \\ & \text{Arbitrarily select a vertex with the least weight, i.e.,} \\ & \text{any } i \in \operatorname{argmin}_{a \in \mathcal{V}_{\text{leafSCC}}} q_a; \\ & \text{Remove all outgoing arcs from vertex } i, \text{ i.e., all} \\ \end{array}$

 $(i \to k) \in \mathcal{A};$

i:

Executing the pruning algorithm gives the following result: Lemma 7: Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{A}\}$. Let the resultant graph after running the pruning algorithm be $\mathcal{G}^{p} = \{\mathcal{V}^{p}, \mathcal{A}^{p}\}$. We have

$$\mathcal{V}^{\mathbf{p}} = \mathcal{V},\tag{8}$$

 $\mathcal{A}^{\mathsf{p}} \subseteq \mathcal{A},\tag{9}$

$$\sum_{i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G}^{\mathsf{p}})} q_i$$
$$= \sum_{k \in \mathcal{V}} q_k - \sum_{i \in \mathcal{L}(\mathcal{G})} q_i - \sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a.$$
(10)

Proof: Since the algorithm only removes arcs, we have (8) and (9).

We now show that \mathcal{G}^p has no leaf SCC. Consider a leaf SCC in \mathcal{G} for which vertex *i* is selected and all its outgoing arcs are removed. By definition, every vertex in the SCC has a path to *i*. After the **foreach** iteration, *i* is made a leaf vertex, and all vertices in the leaf SCC are therefore grounded. Also, as all SCCs are vertex disjoint, this iteration *destroys* one leaf SCC and does not create any new leaf SCC. When the algorithm terminates, \mathcal{G}^p has no leaf SCC, and it follows from Lemma 6 that \mathcal{G}^p is grounded.

Since \mathcal{G}^p is grounded, each non-leaf vertex is a predecessor of some leaf vertex. This means

$$\sum_{i: i \text{ is a predecessor of some vertex in } \mathcal{L}(\mathcal{G}^{p})} q_{i} = \sum_{i \in \mathcal{V} \setminus \mathcal{L}(\mathcal{G}^{p})} q_{i}$$
$$= \sum_{k \in \mathcal{V}} q_{k} - \sum_{i \in \mathcal{L}(\mathcal{G}^{p})} q_{i}.$$
(11)

Recall that $\mathbb{V} \triangleq \{\mathcal{V}_{\text{leafSCC}(1)}, \mathcal{V}_{\text{leafSCC}(2)}, \dots\}$ is the set of all leaf SCCs in \mathcal{G} . These leaf SCCs are vertex-disjoint by definition. In Algorithm 1, one vertex (one with the least weight) in each leaf SCC in \mathcal{G} is made a leaf vertex. So,

$$\sum_{i \in \mathcal{L}(\mathcal{G}^{\mathfrak{p}})} q_i = \sum_{i \in \mathcal{L}(\mathcal{G})} q_i + \sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a.$$
(12)

Substituting (12) into (11), we have (10).

Proof of Theorem 1: Combining Lemmas 5 and 7, we have Theorem 1.

Remark 2: (Complexity of the pruning algorithm) All SCCs in a graph can be found in linear time in $|\mathcal{V}| + |\mathcal{A}|$ (see,

for example, the Kosaraju-Sharir algorithm [30] or Tarjan's algorithm [31]). To check if an SCC is a leaf SCC, we check whether all out-neighbors of each vertex in the SCC is in the same SCC. This takes at most $|\mathcal{A}||\mathcal{V}|$ checks. Finding the vertex with the minimum message length in a leaf SCC can be done by a sorting algorithm with a worst-case complexity of $|\mathcal{V}| \log |\mathcal{V}|$. So, the pruning algorithm runs in $O(|\mathcal{V}|^3)$ in the worst case.

The main idea behind our lower bound here is that removing some requests by a receiver (i.e., reducing the sets $\{\mathcal{W}_r\}$) cannot increase the optimal index codelength. This corresponds to removing arcs in the information-flow graph. A similarbut different—concept was used to derive a lower bound for the unicast index-coding instances (which are represented by side-information graphs), namely, the maximum-acylic-inducedsubgraph (MAIS) lower bound [1], [2]. The main idea behind the MAIS bound is that removing messages (i.e., reducing \mathcal{M} and adjusting $\{\mathcal{K}_r, \mathcal{W}_r\}$ accordingly) from the system cannot increase the optimal index codelength. This corresponds to a vertex-induced subgraph. While the MAIS lower bound is in general not tight for single-sender unicast index-coding instances [4], we will see that our proposed lower bound based on removing arcs is tight for all single-sender uniprior indexcoding instances.

B. Upper bound (achievability): cyclic codes

In this subsection, we construct linear index codes that achieve the lower bound given in Theorem 1. Although the lower bound is calculated based on the pruned graph \mathcal{G}^{p} , for achievability, we need to design index codes for the original graph \mathcal{G} . It turns out that by coding on leaf SCCs in \mathcal{G} , we can construct index codes that attain the lower bound.

It is easier to discuss the lower bound, i.e., the right-hand side of (7), in terms of a trivial index code and "savings". To see this, we compare it with a trivial index code that simply transmits $\sum_{k \in \mathcal{V}} q_k - \sum_{i \in \mathcal{L}(\mathcal{G})} q_i$ bits, which are the uncoded messages of the non-leaf vertices (since, by definition, the messages of the leaf vertices are not requested by any receiver). We will further show that an extra savings of $\sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a$ bits is possible—a consequence of the leaf SCCs. This corresponds to a savings, over the trivial index code, of $\min_{a \in \mathcal{V}_{\text{leafSCC}}} q_a$ bits for each leaf SCC $\mathcal{V}_{\text{leafSCC}} \in \mathbb{V}$.

We will use *cyclic codes*, defined below, to realize the required savings:

Definition 5 (Cyclic codes): Consider n q-bit messages, $\{x_1, x_2, \ldots, x_n\}$. A cyclic code is a length q(n-1)-bit code, constructed as follows:

$$x_1 \oplus x_2, x_2 \oplus x_3, \ldots, x_{n-1} \oplus x_n,$$

where \oplus bit-wise XOR.

Theorem 2 (Single-sender: upper bound): The optimal index codelength of a single-sender single-uniprior index-coding instance is upper bounded as

$$\ell^*(\mathcal{G}) \le \sum_{k \in \mathcal{V}} q_k - \sum_{i \in \mathcal{L}(\mathcal{G})} q_i - \sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a.$$
(13)

Proof of Theorem 2: For each non-leaf vertex i that does not belong to any leaf SCC, we send x_i uncoded.

For vertices in leaf SCCs, we construct the following code: Consider a leaf SCC in \mathcal{G} , and let its vertices be $\mathcal{V}_{\text{leafSCC}} = \{1, 2, \dots, v\}$. For now, let us first assume all messages $\{x_i\}_{i=1}^{v}$ in the leaf SCC are of equal length, say q bits. Then construct the following cyclic code:

$$x_1 \oplus x_2, x_2 \oplus x_3, \ldots, x_{v-1} \oplus x_v.$$
 (14)

Note that for any $i \in \mathcal{V}_{\text{leafSCC}}$, receiver i

- can decode all messages in V_{leafSCC} from (14) and its side information x_i ,
- can also decode the messages of all non-leaf vertices not in any leaf SCC (as these messages are sent uncoded), and
- does not request any messages in other leaf SCCs (due to the definition of leaf SCC).

Repeating the above cyclic code for each leaf SCC, we satisfy the decoding requirements for all receivers in all leaf SCCs. Also, any receiver not in any leaf SCC can also decode its requested messages, as these messages are sent uncoded. Note that, by definition, messages in any leaf SCC are not requested by any receiver outside the leaf SCC.

To generalize the above to non-equal-length messages of q_i bits for each vertex *i*, simply form (14) by substituting each x_i with its first q_{\min} bits, where q_{\min} is the shortest message length in the leaf SCC, i.e., $q_{\min} = \min_{\substack{1 \le a \le v \\ 1 \le q_{\min}}} q_a$. Then the total number of coded bits in (14) is $(v-1)q_{\min}$ bits. We send the remaining bits $\sum_{i=1}^{v} q_i - vq_{\min}$ uncoded. Hence, the total savings per leaf SCC is q_{\min} .

Repeating this for each leaf SCC, we save a total of $\sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a \text{ bits for the entire graph } \mathcal{G}.$

Remark 3: The cyclic code (14) used in the proof of Theorem 2 can be derived from the partial-clique-cover scheme [11] (originally proposed for unicast index coding). Consider (i) v messages $\{x_i\}_{i=1}^v = \mathcal{M}$ each taking values from a finite field \mathbb{F} ; (*ii*) a set of receivers $\{1, 2, \dots, n\}$, where each receiver $i \in \{1, 2, ..., n\}$ knows a priori at least k messages in \mathcal{M} , i.e., $|\mathcal{K}_i \cap \mathcal{M}| > k$. The partial-clique-cover scheme gives a codeword based on minimum-distance-separable codes with the following properties: (a) it consists of (v - k) finite-field symbols; (b) it simultaneously lets each receiver decode all messages in M. For a leaf SCC in a single-uniprior indexcoding problem, we have k = 1. By setting $\mathbb{F} = \{0, 1\}$, and appropriately constructing a minimum-distance-separable code, one can obtain (14) for q = 1 (a codeword of (v - 1) bits), which allows all receivers in the SCC to decode all messages in the SCC. The multiple-bit version of (14) can be obtained by concatenating q > 1 copies of the binary codewords.

C. The optimal index codelength

Combining Theorems 1 and 2, we have the following:

Theorem 3: The optimal index codelength of a single-sender single-uniprior index-coding instance is

$$\ell^*(\mathcal{G}) = \sum_{k \in \mathcal{V}} q_k - \sum_{i \in \mathcal{L}(\mathcal{G})} q_i - \sum_{\mathcal{V}_{\text{leafSCC}(j)} \in \mathbb{V}} \min_{a \in \mathcal{V}_{\text{leafSCC}(j)}} q_a.$$
(15)



Fig. 4. An example of a single-uniprior index-coding instance: (a) Graph \mathcal{G} is an information-flow graph representing the problem instance. (b) Graph \mathcal{G}' is the resultant graph after running the pruning algorithm on \mathcal{G}

Furthermore, linear codes can achieve the optimal index codelength.

Example 2: We illustrate how to obtain the optimal codelength and an optimal code for an index-coding instance with one single-bit message (i.e., $x_1 = x_1[1]$), four two-bit messages (i.e., $x_i = [x_i[1], x_i[2]]$ for $i \in \{2, 3, 4, 5\}$), and five receivers. Each receiver i knows $\mathcal{K}_i = \{x_i\}$, and the receivers each request the following messages: $W_1 = \{x_2, x_3\}, W_2 =$ $\{x_1, x_3, x_4\}, \mathcal{W}_3 = \{x_1, x_2\}, \mathcal{W}_4 = \emptyset, \mathcal{W}_5 = \{x_4\}.$ This problem can be represented by a graph G depicted in Figure 4. It contains one leaf SCC with vertices $V_{\text{leafSCC}} = \{1, 2, 3\}$, in which message x_1 has the shortest length of one bit. There is only one leaf vertex in \mathcal{G} , which is node 5. Theorem 1 gives a lower bound $\ell^*(\mathcal{G}) \geq 9 - 2 - 1 = 6$. To construct an optimal index code, we form a two-bit cyclic code $(x_1[1] \oplus x_2[1], x_2[1] \oplus x_3[1])$ for the leaf SCC in \mathcal{G} , and four uncoded bits for the remaining messages in the leaf SCC and the messages of the other the non-leaf vertices $(x_2[2], x_3[2], x_4[1], x_4[2])$. Here, $\ell^*(\mathcal{G}) = 6$.

For the special cases where all messages are binary, Theorem 3 simplifies to the following:

Corollary 1: The optimal index codelength of a single-sender single-uniprior index-coding instance with *binary messages* is

$$\mathcal{L}^{*}(\mathcal{G}) = n - |\mathcal{L}(\mathcal{G})| - N_{\text{leaf}}(\mathcal{G}),$$
 (16)

where $|\mathcal{L}(\mathcal{G})|$ is the number of leaf vertices in \mathcal{G} and $N_{\text{leaf}}(\mathcal{G})$ is the total number of leaf SCCs in \mathcal{G} .

VI. THE MULTI-SENDER CASE: A LOWER BOUND

In this section, we will modify the pruning algorithm to include information from \mathcal{U} , and further propose an *appending* step to modify the graph \mathcal{G} , based on which we derive a lower bound for the multi-sender case.

For simplicity and clarity, we will only consider the case where all messages x_i are binary. This means all vertices in the information-flow graph \mathcal{G} all have weight one. One can extend the ideas developed in this section to the more general case of unequal message sizes.

A. The single-sender bound may be loose for multi-sender

Denote a multi-sender instance by $(\mathcal{G}, \mathcal{U})$. The optimal codelength $\tilde{\ell}^*(\mathcal{G}, \mathcal{U})$ clearly is lower bounded by $\ell^*(\mathcal{G})$ for the single-sender case, because a single sender who knows all the messages can also send any multi-sender index code. We also know from Lemma 3 that, for a graph \mathcal{G} with binary messages, $\ell^*(\mathcal{G})$ is lower bounded by the number of vertices

that are each a predecessor of some leaf vertex, denoted by $V_{\text{predec}}(\mathcal{G})$. This gives

$$\ell^*(\mathcal{G}, \mathcal{U}) \ge \ell^*(\mathcal{G}) \ge V_{\text{predec}}(\mathcal{G}).$$
 (17)

For the single-sender problem, we used the pruning algorithm to prune all leaf SCCs to get a grounded graph \mathcal{G}^{p} . As pruning does not increase the optimal codelength (see Lemma 4), we have the following single-sender bound for multi-sender instances:

$$\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) \ge \ell^*(\mathcal{G}) \ge \ell^*(\mathcal{G}^p) \ge V_{\text{predec}}(\mathcal{G}^p) = V_{\text{out}}(\mathcal{G}^p), \quad (18)$$

where $V_{out}(\mathcal{G})$ denotes the total number of non-leaf vertices in \mathcal{G} . The equality above is derived as \mathcal{G}^{p} is grounded.

However, the single-sender lower bound may be loose due to the sender constraints specified by the undirected graph \mathcal{U} (see its definition in Section III-C). Consider Example 2 again, and suppose that there are five senders, where each sender $i \in \{1, 2, \dots, 5\}$ has only x_i . Under this constraint, it is not possible to transmit $x_1[1] \oplus x_2[1]$, as the messages belong to different senders. The minimum codelength for this case is seven bits—sending $[x_1, x_2, x_3, x_4]$ uncoded. On the other hand, if some sender has $\{x_1, x_2\}$ and another has $\{x_2, x_3\}$, then the six-bit lower bound is tight using the coding scheme mentioned in the example.

B. A tighter bound by appending

We will propose a new algorithm that gives a better lower bound. Our new algorithm produces a resultant graph, say \mathcal{G}^{\dagger} , with a possibly higher number of $V_{\text{predec}}(\mathcal{G}^{\dagger}) > V_{\text{predec}}(\mathcal{G}^{p})$, where \mathcal{G}^{p} is the resultant graph after running Algorithm 1. To obtain the desired \mathcal{G}^{\dagger} , we propose a new technique that appends some leaf SCCs, by adding an outgoing arc from each of these leaf SCC to another vertex or a newly introduced grounded vertex. In the new algorithm, we iterate on the leaf SCCs (either pruning or appending) until we get a grounded information-flow graph \mathcal{G}^{\dagger} .⁶ From the resultant \mathcal{G}^{\dagger} , we get

$$\tilde{\ell}^*(\mathcal{G}^{\dagger}, \mathcal{U}^{\dagger}) \ge \ell^*(\mathcal{G}^{\dagger}) \ge V_{\text{predec}}(\mathcal{G}^{\dagger}) = V_{\text{out}}(\mathcal{G}^{\dagger}), \qquad (19)$$

for all \mathcal{U}^{\dagger} , where the inequalities follow from (17), and the equality is derived because \mathcal{G}^{\dagger} is grounded.

To link this lower bound to the original graph $(\mathcal{G}, \mathcal{U})$, we need to make sure that the optimal index codelength cannot increase after each pruning/appending step. If this is satisfied, we have the lower bound

$$\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) \ge \tilde{\ell}^*(\mathcal{G}^{\dagger}, \mathcal{U}^{\dagger})$$
(20a)

$$\geq V_{\text{out}}(\mathcal{G}^{\intercal}).$$
 (20b)

We aim to maximize the number of non-leaf vertices in the resultant graph (which is grounded).

We will see that pruning reduces the number of non-leaf vertices by one, while appending does not change the number of non-leaf vertices. So, potentially $V_{\text{out}}(\mathcal{G}^{\dagger}) > V_{\text{out}}(\mathcal{G}^{p})$, which means the lower bound (20b) is potentially tighter than (18). We have seen in the single-sender case that (20a) is always true if we prune all leaf SCCs. However, appending in general increases decoding requirements, which may cause the optimal index codelength to increase. The main challenge here is to append certain types of leaf SCCs such that the optimal index codelength cannot increase, i.e., (20a) holds. To this end, we will classify leaf SCCs based on \mathcal{U} .

C. Classifying leaf SCCs to reflect the sender setting

We now give some intuition on how we use \mathcal{U} to classify the leaf SCCs. Recall that \mathcal{U} is an undirected graph where an edge exists between vertices i and j, if and only if some sender has both x_i and x_j . Connectivity in \mathcal{U} restricts the construction of index codes, which allows us to append some leaf SCCs. For example, if there is no path between two vertices in \mathcal{U} , say vertices 1 and 2, then it turns out that we can partition any index codeword into two sub-codewords, where one sub-codeword is a function of only messages from a set, say $\mathcal{M}' \subset \mathcal{M}$, containing x_1 , and the other sub-codeword a function of the only messages from the set $\mathcal{M} \setminus \mathcal{M}'$, which contains x_2 . This, we will show, leads to the result that any receiver can decode all messages in a leaf SCC that contains two vertices not connected in \mathcal{U} . As a consequence, we can append this type of leaf SCC, while satisfying (20a), by adding an arc from the leaf SCC to a dummy vertex.

We will now formally classify different types of leaf SCCs based on their connectivity in \mathcal{U} . We first define *neighboring vertices* in \mathcal{U} . For a vertex set $\mathcal{V}_{S} \subseteq \mathcal{V}$, we say that a vertex $i \notin \mathcal{V}_{S}$ is a *neighbor* of \mathcal{V}_{S} if and only if there is an edge (i, v)in \mathcal{U} between *i* and some $v \in \mathcal{V}_{S}$.

- 1) Message-connected leaf SCC: A leaf SCC in \mathcal{G} is said to be message-connected if and only if there always exists a path⁷ in \mathcal{U} between any two vertices in the SCC, where the path consists of vertices only in the SCC.
- 2) Message-disconnected leaf SCC: A leaf SCC is message-disconnected if and only if there are two vertices in the SCC with no path in \mathcal{U} between them (even if the path can contain vertices outside the SCC).
- 3) Semi-message-connected leaf SCC A leaf SCC that is neither message-connected nor message-disconnected is semi-message-connected, referred also as semi leaf SCC for short. Here, we can always find a vertex pair, where all paths between them must contain some vertex outside the SCC. We further classify semi leaf SCCs:
 - a) Degenerated leaf SCC: A semi leaf SCC, with vertex set \mathcal{V}_{S} , is said to be *degenerated* if and only if we can find two vertex sets

 - $\mathcal{V}_{S}^{inside} \subset \mathcal{V}_{S}$, which is inside the leaf SCC, and $\mathcal{V}_{S}^{outside} \subseteq \mathcal{V} \setminus \mathcal{V}_{S}$, which is outside the leaf SCC, such that
 - there is no edge in \mathcal{U} between $\mathcal{V}_S^{\text{inside}}$ and $\mathcal{V}_S \setminus \mathcal{V}_S^{\text{inside}}$ (the complement of $\mathcal{V}_S^{\text{inside}}$ in the leaf SCC),

⁶A lower bound based on a non-grounded resultant graph is strictly suboptimal, because the graph must contain at least one leaf SCC, and by pruning the leaf SCC, we can increase $V_{\text{predec}}(\cdot)$.

⁷Recall that we use *path* to denote an undirected path (of edges) in \mathcal{U} , or a directed path (of arcs) in G.



Fig. 5. Classification of leaf SCCs: By definition, leaf SCCs are determined by \mathcal{G} (where arcs are drawn with arrows), but their various types are determined also in accordance with \mathcal{U} (where edges are drawn with solid lines). This graph illustrates concurrently three leaf SCC types: (i) message-connected, where there is a path between any two vertices through only vertices in the SCC; (ii) message-disconnected, containing at least two vertices that cannot be connected by any path; and (iii) semi-message-connected, where some vertices must be connected by a path with vertices outside the SCC. The semi leaf SCC here (with vertex set \mathcal{V}_S) is degenerated because we can find two vertex sets, \mathcal{V}_S^{inside} in the leaf SCC and $\mathcal{V}_S^{outside}$ outside the SCC, such that (a) there is no edge between \mathcal{V}_S^{inside} and $\mathcal{V}_S \setminus \mathcal{V}_S^{inside}$, (b) $\mathcal{V}_S^{outside}$ contains at most one non-leaf vertex, and (c) all neighbors of \mathcal{V}_S^{inside} are predecessors of $\mathcal{V}_S^{outside}$.

- there is at most one non-leaf vertex in $\mathcal{V}^{\text{outside}}_{S},$ and
- every neighbor of V_S^{inside} is either in V_S^{outside}, or a predecessor (with respect to the directed graph G) of some vertex in V_S^{outside}.
- b) Non-degenerated leaf SCCs: A semi leaf SCC that is not degenerated is said to be non-degenerated.

Any leaf SCC must belong to one of these four types. Figure 5 shows examples of three types of leaf SCCs. As cryptic as the definition of a degenerated leaf SCC may seem, it has been carefully crafted to allow us to append it.

D. Which leaf SCCs to append and which to prune?

Recall that our lower bound is given by $\hat{\ell}^*(\mathcal{G}, \mathcal{U}) \geq V_{\text{out}}(\mathcal{G}^{\dagger})$, and we want the resultant grounded information-flow graph \mathcal{G}^{\dagger} to have as many non-leaf vertices as possible. As mentioned before, we would prefer appending to pruning, as the latter always reduces the number of non-leaf vertices, but the former does not. However, for message-connected and nondegenerated leaf SCCs, we are unable to show that appending will not increase the optimal index codelength, as required by (20a). In fact, appending a message-connected leaf SCC can strictly increase the optimal index codelength, as shown by the following example:

Example 3: Consider a single-uniprior index-coding instance, whose information-flow graph is given by



and a sender has both the messages. Here, the graph itself is a message-connected leaf SCC. The optimal codelength is one, attained by $x_1 \oplus x_2$. The single-sender lower bound $\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) \geq 1$ is tight. If we append (we will formally define the appending step in the next section) this leaf SCC, by creating a dummy vertex 3, and an arc $(2 \rightarrow 3)$, the resultant grounded graph \mathcal{G}^{\dagger} is

and we get $\ell^*(\mathcal{G}^{\dagger}, \mathcal{U}^{\dagger}) \geq V_{\text{predec}}(\mathcal{G}^{\dagger}) = 2$. For this example, (20a) is not true, and $V_{\text{predec}}(\mathcal{G}^{\dagger})$ is not a lower bound on $\ell^*(\mathcal{G}, \mathcal{U})$.

So, we will

- prune message-connected and non-degenerated leaf SCCs, and
- append message-disconnected and degenerated leaf SCCs.

In the following three subsections, we will prove that appending message-disconnected and degenerated leaf SCCs indeed guarantees (20a). We will also derive the number of non-leaf vertices after each pruning/appending step, and show that the number of leaf SCCs that remain in the graph cannot increase (so that our algorithm always terminates).

We will use the following notation. Let \mathcal{G} and \mathcal{U} be the information-flow graph and the message graph (respectively) before a pruning/appending step, and \mathcal{G}' and \mathcal{U}' be the respective graphs after the step.

1) Appending a message-disconnected leaf SCC:

Definition 6 (Appending a message-disconnected leaf SCC): To append a message-disconnected leaf SCC (with a vertex set V_S), we add

- a dummy vertex n + 1 in both \mathcal{G} and \mathcal{U} where $x_{n+1} = 0$, which is known to all receivers and is hence never transmitted by the sender, and
- a dummy arc in G from an arbitrarily chosen vertex in the leaf SCC to the dummy vertex, i.e., choose some v ∈ V_S and add (v → n + 1) in G.

Proposition 1: After appending a message-disconnected leaf SCC, we have

$$N_{\text{leaf}}(\mathcal{G}') = N_{\text{leaf}}(\mathcal{G}) - 1, \qquad (21)$$

$$\tilde{\ell}^*(\mathcal{G}',\mathcal{U}') = \tilde{\ell}^*(\mathcal{G},\mathcal{U}),\tag{22}$$

$$V_{\rm out}(\mathcal{G}') = V_{\rm out}(\mathcal{G}). \tag{23}$$

Proof: See Appendix B.

2) Appending a degenerated leaf SCC: Recall that a semi leaf SCC (with vertex set V_S) is degenerated if and only if we can find two vertex sets, V_S^{inside} in the SCC and V_S^{outside} outside the SCC, as per the definition given in Section VI-C. We append a degenerated leaf SCC as follows:

Definition 7 (Appending a degenerated leaf SCC): To append a degenerated leaf SCC, first pick any vertex in $V_{\rm S}^{\rm inside}$, say $v_{\rm inside}$. We then add an arc according to the following two possibilities:

- Case 1: If $\mathcal{V}_{S}^{\text{outside}}$ contains exactly one non-leaf vertex, add an arc from v_{inside} to the non-leaf vertex.
- Case 2: Otherwise (all vertices in $V_{\rm S}^{\rm outside}$ are leaf vertices), add an arc from $v_{\rm inside}$ to an arbitrarily chosen vertex in $V_{\rm S}^{\rm outside}$.

Proposition 2: After appending a degenerated leaf SCC, we have

$$N_{\rm leaf}(\mathcal{G}') = \int N_{\rm leaf}(\mathcal{G}), \text{ or } (24)$$

$$\tilde{\ell}^*(\mathcal{C}'\mathcal{U}') - \tilde{\ell}^*(\mathcal{C}\mathcal{U}) - 1, \qquad (25)$$

$$V_{\text{out}}(\mathcal{G}') = V_{\text{out}}(\mathcal{G}).$$
(27)



Fig. 6. An algorithm that recursively prune/append leaf SCCs to produce a grounded information-flow graph

3) Pruning a leaf SCC: The pruning step in the multi-sender case is the same as the single-sender case. Since we assume that all the messages are binary, we can choose any vertex from which we prune its outgoing arcs. Formally, we have the following:

Definition 8 (Pruning a leaf SCC): To prune a leaf SCC, with vertex set V_S , we arbitrarily select a vertex $v \in V_S$, and remove all outgoing arcs from v.

Proposition 3: After pruning a leaf SCC, we have

$$N_{\text{leaf}}(\mathcal{G}') = N_{\text{leaf}}(\mathcal{G}) - 1, \qquad (28)$$

$$\tilde{\ell}^*(\mathcal{G}',\mathcal{U}') \le \tilde{\ell}^*(\mathcal{G},\mathcal{U}),\tag{29}$$

$$V_{\text{out}}(\mathcal{G}') = V_{\text{out}}(\mathcal{G}) - 1.$$
(30)

Proof: The pruning step grounds all the vertices in the leaf SCC, and does not change the connectivity of other vertices outside the leaf SCC. This gives (28).

As removing arcs reduces decoding requirements, any index code for $(\mathcal{G}, \mathcal{U})$ also satisfies the decoding requirement for $(\mathcal{G}', \mathcal{U}')$. So, we have (29).

As we remove all outgoing arc from a non-leaf vertex, we have (30).

E. Constructing an algorithm by pruning and appending

We are now ready to assemble the pruning and appending steps to form an algorithm that returns a grounded information-flow graph \mathcal{G}^{\dagger} .

From the previous section, we know that each step increases neither the optimal index-codelength nor the number of leaf SCCs. In fact, we have the following lemma:



Fig. 7. An example showing the importance of the pruning/appending sequence. Arcs in \mathcal{G} are marked with arrows, and edges in \mathcal{U} with solid lines.

Lemma 8: We can always get a grounded information-flow graph after a finite number of pruning and appending steps on all the leaf SCCs.

Proof: Consider an algorithm that recursively (i) appends a message-disconnected or a degenerated leaf SCC, and (ii) prunes a message-connected or a non-degenerated leaf SCC. We now show that the algorithm always terminates after a finite number of steps.

From Propositions 1, 2, and 3, the number of leaf SCCs always (i) decreases by one, or (ii) remains the same (In this case, the leaf SCC assimilates one or more vertices that do not belong to any leaf SCC. See Appendix C for proof.) after each pruning/appending step. The pruning/appending step that results in case (i) can be executed at most n/2 times as there are at most n/2 leaf SCCs in a graph. The pruning/appending step that results in case (ii) can be executed at most n-2 times, as this step requires a leaf SCC to start with (which must contain at least two vertices), and there are at most n-2 vertices that does not belong to any leaf SCC.

So, the algorithm will terminate after at most 3n/2 - 2 appending/pruning steps.

Figure 6 depicts our algorithm to find a lower bound on $\tilde{\ell}^*(\mathcal{G},\mathcal{U})$.

F. Obtaining a lower bound

The algorithm depicted in Figure 6 returns possibly different lower bounds, $V_{out}(\mathcal{G}^{\dagger})$, depending on which leaf SCC is selected in each iteration. We give an example to show this, after proving the following lemma:

Lemma 9: A non-degenerated leaf SCC may become degenerated after appending/pruning other leaf SCCs.

Proof: Although no new leaf SCC is created after each appending/pruning step, more vertices may be grounded, and some non-degenerated leaf SCCs may become degenerated as the degeneration conditions are satisfied with more grounded vertices.

Example 4: Consider the index-coding instance depicted in Figure 7. The directed subgraph induced by vertices $\{1, 2\}$ is a message-connected leaf SCC, that by $\{3, 4\}$ non-degenerated, and that by $\{5, 6\}$ degenerated. Suppose that we first prune the non-degenerated leaf SCC $\{3, 4\}$, followed by pruning the message-connected $\{1, 2\}$, and then appending the degenerated $\{5, 6\}$. This pruning/appending sequence gives a resultant grounded graph with $V_{\text{out}}(\mathcal{G}^{\dagger}) = 4$. However, if we append $\{5, 6\}$ and prune $\{1, 2\}$ first, $\{3, 4\}$ will be made degenerated and will be appended instead. This gives a different \mathcal{G}^{\dagger} with a higher $V_{\text{out}}(\mathcal{G}^{\dagger}) = 5$.



Fig. 8. An heuristic algorithm (which processes non-degenerated leaf SCCs last) that recursively prune/append leaf SCCs to produce a grounded information-flow graph

To obtain the tightest lower bound using the algorithm in Figure 6, we perform the following combinatorial optimization:

Theorem 4 (Multi-sender: lower-bound): The optimal indexcodelength for a multi-sender index-coding instance represented by $(\mathcal{G}, \mathcal{U})$ is lower bounded as

$$\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) \ge \max V_{\text{out}}(\mathcal{G}^\dagger),$$
(31)

where \mathcal{G}^{\dagger} is the resultant graph after running the algorithm depicted in Figure 6, and the maximization is taken over all possible sequences of pruning/appending the leaf SCCs in the algorithm.

G. Another lower bound using a heuristic algorithm

The lower bound obtained by the combinatorial optimization in Theorem 4 has some implementation issues:

- For a large number of leaf SCCs, the optimization runs in factorial time in the number of leaf SCCs.
- It provides little insights to when the algorithm is optimal (i.e., giving a tight lower bound), and how good the algorithm is (whether we can bound the gap between the optimal index codelength and the lower bound).

In light of these issues, we design a heuristic algorithm. We first prove the following lemma:

Lemma 10: The status of a leaf SCC being messageconnected, message-disconnected, or degenerated is not affected by the appending/pruning of *other* leaf SCCs.

Proof: See Appendix D.

Lemma 10 suggests that it does not matter in what order we prune or append message-connected, message-disconnected, or degenerated leaf SCCs. At the end, all message-connected



leaf SCCs will be pruned, and all message-disconnected and degenerated leaf SCCs appended.⁸

It matters, however, in which order we process *non-degenerated* leaf SCCs. As pointed out in the previous section, having more grounded vertices increases the chances of a non-degenerated leaf SCC become degenerated for which we can append (instead of pruning it in the event that it stays non-degenerated). We will design a heuristic algorithm that leaves the *pruning of the non-degenerated leaf SCCs to the last*—we ground the rest as far as possible. This heuristic algorithm is shown in Figure 8. Note that this heuristic algorithm is a one-pass algorithm which does not iterate over operating different sequences of leaf SCCs.

While the heuristic algorithm in Figure 8 produces a grounded \mathcal{G}^{\dagger} that we require, we re-write the algorithm in a slightly different way, as Algorithm 2, which leads to a lower bound (see Theorem 5) in a form similar to the upper bound to be derived in Theorem 6 later. Having similar forms, we will then identify classes of $(\mathcal{G}, \mathcal{U})$ for which the bounds match.

We now describe the idea behind Algorithm 2. Our lower bound, $V_{out}(\mathcal{G}^{\dagger})$, relates to the number of non-leaf vertices in the original graph \mathcal{G} via (23), (27), and (30). We see that we only need to count the number of pruning steps. So, we design Algorithm 2 such a way that we

• identify and count each pruning step,

 8 Here, we do not optimize how to append a degenerated leaf SCC, when there are two or more ways to append it. So, we may get to a suboptimal solution.

- separate the pruning of message-connected leaf SCCs in the original graph \mathcal{G} out in the Initialization phase (so that we can express the lower bound as a function of \mathcal{G}),
- process non-degenerated leaf SCCs last (after processing all other types of leaf SCCs).

The Initialization and Iteration phases are the same except that we only prune non-degenerated leaf SCCs in the latter, after pruning and counting all message-connected leaf SCCs in \mathcal{G} .

Note that any leaf SCC being operated is destroyed except for degenerated leaf SCCs (see Appendix C). In this case, the degenerated leaf SCC may assimilate other vertices to form a larger leaf SCC, which can be of any type, depending on how the new leaf SCC is connected in \mathcal{U} . The type of this newly-formed leaf SCC is to be identified during runtime.

After running Algorithm 2, we obtain the following lower bound:

Theorem 5 (Multi-sender: lower bound): The optimal index codelength for a multi-sender index-coding instance represented by $(\mathcal{G}, \mathcal{U})$ is lower bounded as

$$\hat{\ell}^*(\mathcal{G},\mathcal{U}) \ge V_{\text{out}}(\mathcal{G}^\dagger)$$
 (32a)

$$= V_{\text{out}}(\mathcal{G}) - \#(\text{pruning steps in Initialization})$$

$$-$$
 #(pruning steps in all Iterations) (32b)

$$= V_{\text{out}}(\mathcal{G}) - (N_{\text{msg-conn}}(\mathcal{G}, \mathcal{U}) + N_{\text{iteration}}), \quad (32c)$$

where $N_{\text{msg-conn}}(\mathcal{G}, \mathcal{U})$ is the number of message-connected leaf SCCs in \mathcal{G} , and $N_{\text{iteration}}$ is the number of times Iteration repeated in Algorithm 2.

Proof: The lower bound is obtained by running Algorithm 2. Equation (32a) follows from (20b); (32b) follows from (23), (27), and (30); (32c) follows as one and only one leaf SCC is pruned during each Iteration.

Remark 4: From (32c), we note that the lower bound depends on the number of times Iteration is executed. In each iteration, the choice of which non-degenerated leaf SCC to prune will affect the lower bound. As a rule of thumb, we prune a non-degenerated leaf SCC that will degenerate other nondegenerated leaf SCCs, so as to reduce the number of Iterations.

VII. THE MULTI-SENDER CASE: AN UPPER BOUND

We now derive an upper bound, via a coding scheme, by constructing special trees in the message graph \mathcal{U} , referred to as *connecting trees*. The vertex set of a connecting tree, denoted by \mathcal{V}^{T} , has the following properties:

- Each vertex in V^T has one or more outgoing arcs in G, and only to other vertices in V^T. This means V^T has no outgoing arc to V \ V^T,
- 2) No vertex in \mathcal{V}^{T} belongs to any message-connected leaf SCCs or another connecting tree.

Let $N_{\text{tree}}(\mathcal{G}, \mathcal{U})$ denote the number of connecting trees that can be found for a given instance $(\mathcal{G}, \mathcal{U})$. We will propose a coding scheme that achieves the following index codelength:

Theorem 6 (Multi-sender: upper bound): The optimal index codelength for a multi-sender index-coding instance represented by $(\mathcal{G}, \mathcal{U})$ is upper bounded as

$$\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) \le V_{\text{out}}(\mathcal{G}) - (N_{\text{msg-conn}}(\mathcal{G}, \mathcal{U}) + N_{\text{tree}}(\mathcal{G}, \mathcal{U})).$$
(33)



Fig. 9. An index-coding instance represented by \mathcal{G} (with arcs) and \mathcal{U} (with edges). There are two leaf SCCs in the graph: (i) message-connected, and (ii) degenerated. Here, we can form a connecting tree using the hollow vertices.

The upper bound is optimized by finding the maximum number of connecting trees.

For example, consider the graphs in Figure 9. We can form one connecting tree using the hollow vertices.

Proof of Theorem 6: To simplify notation, we will show that there exists an index code of length $V_{\text{out}} - (N_{\text{msg-conn}} + N_{\text{tree}})$, where we have dropped the arguments as they are clear from the context. Let a set of connecting trees be $\{\mathcal{T}_t = (\mathcal{V}_t^{\text{T}}, \mathcal{E}_t^{\text{T}}):$ $t \in \{1, 2, \ldots, N_{\text{tree}}\}\}$, and all the message-connected leaf SCCs in \mathcal{G} be $\{\mathcal{C}_c = (\mathcal{V}_c^{\text{C}}, \mathcal{A}_c^{\text{C}}): c \in \{1, 2, \ldots, N_{\text{msg-conn}}\}\}$. Furthermore, let the remaining vertices in \mathcal{G} be $\mathcal{V}' = \mathcal{V} \setminus \{\bigcup_{t=1}^{N_{\text{tree}}} \mathcal{V}_t^{\text{T}} \cup \bigcup_{c=1}^{N_{\text{msg-conn}}} \mathcal{V}_c^{\text{C}}\}$. Denote by \mathcal{V}_{out} the set of all nonleaf vertices in \mathcal{V}' . By definition, all $\mathcal{V}_t^{\text{T}}, \mathcal{V}_c^{\text{C}}$, and \mathcal{V}' are disjoint. Our coding scheme is as follows:

1) For each connecting tree $(\mathcal{V}_t^{\mathrm{T}}, \mathcal{E}_t^{\mathrm{T}})$, we transmit all $\{x_i \oplus x_i \in (i, i) \in \mathcal{E}^{\mathrm{T}}\}$ i.e. we transmit the XOP of the

- $x_j : (i, j) \in \mathcal{E}_t^{\mathrm{T}}$, i.e., we transmit the XOR of the associated message pair for each edge. Note that we transmit $|\mathcal{V}_t^{\mathrm{T}}| 1$ bits.
- 2) For each message-connected leaf SCC $(\mathcal{V}_c^{C}, \mathcal{A}_c^{C})$ (which is edge-connected by definition), we first obtain a spanning tree in \mathcal{U} , denoted by $\mathcal{T}_c^{ST} = (\mathcal{V}_c^{C}, \mathcal{E}_c^{ST})$, where $\mathcal{E}_c^{ST} \subseteq \mathcal{E}$. We then transmit all $\{x_i \oplus x_j : (i, j) \in \mathcal{E}_c^{ST}\}$. Note that we transmit $|\mathcal{V}_c^{C}| - 1$ bits.
- For the rest of the non-leaf vertices, we transmit {x_i : i ∈ V'_{out}}, i.e., we transmit the message bits uncoded.

Each vertex in the connecting trees and the message-connected SCCs has at least one outgoing arc. Hence, the coding scheme generates an index code of length $V_{\text{out}} - (N_{\text{msg-conn}} + N_{\text{tree}})$.

We can easily verify that the index code can be transmitted, as each message pair to be XORed is associated with an edge, i.e., both the message bits belong to some sender.

Finally, we show that each receiver is able to obtain its requested messages. Recall that each receiver *i* needs to decode all messages in $\{x_j : (j \to i) \in \mathcal{A}\}$. Now, each receiver *i* must belong to one—and only one—of the following groups:

- (Connecting tree) i ∈ V_t^T: Knowing x_i, receiver i can decode all {x_j : j ∈ V_t^T} from {x_j ⊕ x_k : (j, k) ∈ E_t^T} by traversing the tree (which is connected by definition). It can also decode the messages {x_k : k ∈ V'_{out}}, sent uncoded. Since all connecting trees and message-connected leaf SCCs have no outgoing arcs, each incoming arc to i must be from either V_t^T \ {i} or V'_{out}. So, receiver i is able to decode all its requested messages.
- 2) (Message-connected leaf SCC) $i \in \mathcal{V}_c^{\mathbb{C}}$: Using the same argument as that for the connecting trees, we can show that receiver *i* can decode all its requested messages.



Fig. 10. We run Algorithm 2 on the graph in Figure 9. Subfigures (a) and (b) show the initialization step where the message-connected leaf SCC is pruned (by removing the dashed arc in (a)) and the degenerated leaf SCC is appended (by adding the dashed arc in (b)), respectively. Subfigure (c) shows the final graphs after the algorithm terminates.

3) (The remaining vertices) i ∈ V': Using the argument in point 1, all incoming arcs to vertex i must come from V'_{out} \ {i}. Since we sent {x_j : j ∈ V'_{out}} uncoded, receiver i can decode all its requested messages.

VIII. SPECIAL CASES WHERE THE BOUNDS ARE TIGHT AND AN EXAMPLE

Combining Theorems 5 and 6, we conclude $N_{\text{iteration}} \ge N_{\text{tree}}(\mathcal{G}, \mathcal{U})$, and thus the optimal index codelength is found within $N_{\text{iteration}} - N_{\text{tree}}(\mathcal{G}, \mathcal{U})$ bits. Recall that $N_{\text{iteration}}$ is the number of times Iteration repeated in Algorithm 2, and it is upper bounded by the number of leaf SCCs in \mathcal{G} that are not message-connected. In the following special cases, we have $N_{\text{iteration}} = N_{\text{tree}}(\mathcal{G}, \mathcal{U})$, and the lower bound is tight.

Corollary 2: If no leaf SCC remains after running the initialization of Algorithm 2, then

$$\tilde{\ell}^*(\mathcal{G}, \mathcal{U}) = V_{\text{out}}(\mathcal{G}) - N_{\text{msg-conn}}(\mathcal{G}, \mathcal{U}).$$
(34)

Proof: If Algorithm 2 terminates after the initialization step, we have $N_{\text{iteration}} = 0$, which implies that $N_{\text{tree}}(\mathcal{G}, \mathcal{U}) = N_{\text{iteration}} = 0$.

Corollary 3: If each bit x_i in the message set \mathcal{M} is known to only one sender (i.e., the *n* sender constraint sets \mathcal{M}_s partition \mathcal{M}), then the optimal index codelength is given by (34).

Proof: If messages x_i and x_j belong to some sender s (i.e. $x_i, x_j \in \mathcal{M}_s$), then there exists an edge (i, j) in the message graph \mathcal{U} . Otherwise, if the messages x_i, x_j belong to different senders, it is impossible to have a *path* between i and j. This means we have only message-connected or -disconnected leaf SCCs, i.e., there is no semi leaf SCC. Thus, $N_{\text{iteration}} = 0$.

Corollary 3 includes the result of the single-sender instance [15] as a special case.

Note that the scenario in Corollary 3, where the senders' message sets do not overlap, cannot be trivially solved by splitting the multi-sender instance into multiple single-sender instances, where each instance consists of the messages known to one sender. We illustrate this with an example in Appendix E.



Fig. 11. An example where the optimal index codelength $\ell^* = 4$ is achievable by XOR of three bits

A. An example

We now illustrate the lower bound and the achievability for the graph in Figure 9. For lower bound, we run Algorithm 2 to obtain the resultant graph shown in Figure 10(c). Specifically, starting from Figure 9, note that the furthest-left four vertices form a message-connected leaf SCC. We prune this leaf SCC by removing all outdoing arcs from an arbitrarily selected vertex. This results in Figure 10(a), where we remove the outgoing arc (dashed arc) from *a*. Now, the two middle vertices forms a degenerated leaf SCC. We append this leaf SCC by adding an arc (indicated by the dashed arc in Figure 10(b)) from v_{inside} to $v_{outside}$. This completes the initialization step.

Entering the iteration stage, we note that the right three vertices in Figure 10(b) now form a message-connected leaf SCC. We prune this leaf SCC by removing all outgoing arc(s) from the arbitrarily selected vertex v_{outside} . The algorithm terminates here, after destroying all leaf SCCs. From Theorem 5, we have the lower bound $\tilde{\ell}^* \geq 7 - (1+1) = 5$.

For achievability, recall that we can form a connecting tree using the three hollow vertices in Figure 9. So, Theorem 6 gives the upper bound $\tilde{\ell}^* \leq 7 - (1+1) = 5$.

In this example, we have a scenario where $N_{\text{iteration}} = N_{\text{tree}}(\mathcal{G}, \mathcal{U}) = 1$, and hence both the lower and upper bounds are tight, giving the optimal index codelength of five bits. This example illustrates that the upper and lower bounds can coincide even if $N_{\text{iteration}} > 0$ (cf. Corollary 2).

IX. DISCUSSIONS AND FUTURE WORK

We now show that the pairwise linear coding proposed in Section VII can be suboptimal. Consider an index-coding instance with six messages and six receivers, with $\mathcal{A} = \{(1 \leftrightarrow 2), (3 \leftrightarrow 4), (5 \leftrightarrow 6)\}$ where $(i \leftrightarrow j) \triangleq \{(i \rightarrow j), (j \rightarrow i)\}$, and four senders having the following messages (x_1, x_3, x_5) , $(x_3, x_5, x_2), (x_5, x_2, x_4)$, and (x_2, x_4, x_6) respectively.

The graphical representation is depicted in Figure 11. In this example, there is no message-connected, message-disconnected, or degenerated leaf SCC in \mathcal{G} . Running Algorithm 2, we get the lower bound $\tilde{\ell}^* \geq 4$.

We can also show that the largest $N_{\text{tree}} = 1$, i.e., the pairwise linear coding (Theorem 6) can only achieve 5 bits. However, if each sender sends the XOR of its three message bits, the 4-bit lower bound is achievable.

This example illustrates that—in stark contrast to the single-sender case—the pairwise coding scheme described in Section VII is not always optimal. It also shows a disadvantage of using \mathcal{U} : it cannot differentiate a sender having (x_1, x_2, x_3) from three senders having $(x_1, x_2), (x_2, x_3)$, and (x_1, x_3) respectively. For future work, we will investigate a more general coding scheme and a more informative graphical representation.

APPENDIX A A GRAPH CONTAINS NO LEAF SCC IF AND ONLY IF IT IS GROUNDED

Every directed graph with no leaf SCC is grounded: Given any graph \mathcal{G} , we form a supergraph \mathcal{G}^s by replacing each (leaf or non-leaf) SCC with at least two vertices by a special vertex, referred to as a supernode. First, \mathcal{G}^s cannot contain any directed cycle. Otherwise, all supernodes and vertices in the cycle form an SCC, and it would have been collapsed into a supernode. Since \mathcal{G}^s is acyclic, a path from any node/supernode must terminate somewhere. Now, if the \mathcal{G} has no leaf SCC, then \mathcal{G}^s has no leaf supernode, and any path must terminate at a regular (i.e., non-supernode) leaf vertex.

Every grounded directed graph has no leaf SCC: This is easily seen by the fact that all vertices in the leaf SCC cannot be grounded, because these vertices can only reach vertices within the same leaf SCC, all of which are non-leaf.

APPENDIX B Proof of Proposition 1

The proof of Proposition 1 requires the following lemma. We defer the proof of this lemma to after that of the proposition.

Lemma 11: From any index code, any receiver is able to decode the messages of any message-disconnected leaf SCC.

Proof of Proposition 1: We first show that (21) holds. The appending step adds a dummy leaf vertex and an arc from the message-connected leaf SCC to the dummy vertex. This grounds all vertices in the leaf SCC, and the leaf SCC is no longer a leaf SCC. The connectivity of all other vertices outside the leaf SCC is unchanged. So, this appending step neither creates any new leaf SCC nor destroys any other existing leaf SCC. Hence the number of leaf SCCs decreases by one after this step. We have (21).

Next, we show that (22) holds. Adding the arc $(v \to n+1)$ now additionally requires the dummy receiver to decode messages in the appended leaf SCC, so $\tilde{\ell}^*(\mathcal{G}', \mathcal{U}') \geq \tilde{\ell}^*(\mathcal{G}, \mathcal{U})$. But as the appended leaf SCC is message-disconnected, Lemma 11 above says that its messages can be decoded without using any side-information. Hence, the dummy receiver can decode its required message defined in \mathcal{G}' even using the index code for the original graph $(\mathcal{G}, \mathcal{U})$. We have (22).

Finally we show (23). A new arc is added from $v \in V_S$ to a newly added dummy vertex n + 1. Vertex v already had at least one outgoing arc before the appending step, and n + 1 is a leaf vertex. So, we have (23).

It still remains to prove Lemma 11.

Proof of Lemma 11: Let \mathcal{V}_{S} be the vertex set of a messagedisconnected leaf SCC. By definition, there are two vertices $a, b \in \mathcal{V}_{S}$ that are not connected by any path in \mathcal{U} . Let \mathcal{V}_{a} be a set of all vertices connected to a, and $\mathcal{V}_{b} \triangleq \mathcal{V} \setminus \mathcal{V}_{a}$. This means $a \in \mathcal{V}_{a}, b \in \mathcal{V}_{b}$, and there is no edge across \mathcal{V}_{a} and \mathcal{V}_{b} .

The lack of edge-connectivity (in \mathcal{U}) between \mathcal{V}_a and \mathcal{V}_b implies that any index code can be partitioned into two parts (or subcodes), $\mathbf{c} = (\mathbf{c}_a, \mathbf{c}_b)$, such that every bit in \mathbf{c}_a depends on only $\{x_i : i \in \mathcal{V}_a\}$ and not on $\{x_i : j \in \mathcal{V}_b\}$, and vice versa. This means the random variables X_a and (C_b, X_b) are independent, even given C_a .⁹ In other words,

$$X_a - \boldsymbol{C}_a - (\boldsymbol{C}_b, X_b) \tag{35}$$

forms a Markov chain.

_

Intuitively, the parts c_a (which contains x_a) and c_b do not have any message bit in common. Then it must be true that receiver a decodes x_b using only the part c_b , i.e., without using its prior message x_a , which appears only in the other separate part c_a (its prior messages x_a can help receiver a to decode messages only in c_a). Hence, if receiver a can decode x_b , so can any receiver—even one without any prior message.

We now prove this formally. Since a and b belong to an SCC, receiver a must decode x_b from the index code c and its prior information x_a (see Lemma 1). This means

$$H(\boldsymbol{C}_{a}, \boldsymbol{C}_{b}, \boldsymbol{X}_{a}, \boldsymbol{X}_{b}) = H(\boldsymbol{C}_{a}, \boldsymbol{C}_{b}, \boldsymbol{X}_{a})$$
(36a)

$$= H(\boldsymbol{C}_{a}, \boldsymbol{C}_{b}) + H(\boldsymbol{X}_{a}|\boldsymbol{C}_{a}, \boldsymbol{C}_{b}, \boldsymbol{X}_{b})$$

$$= H(\boldsymbol{C}_{a}, \boldsymbol{C}_{b}) + H(\boldsymbol{X}_{a}|\boldsymbol{C}_{a}, \boldsymbol{C}_{b})$$
(36b)

$$H(\boldsymbol{C}_a, \boldsymbol{C}_b, X_b) + H(X_a | \boldsymbol{C}_a)$$

$$=H(\boldsymbol{C}_{a},\boldsymbol{C}_{b})+H(X_{a}|\boldsymbol{C}_{a})$$
(36c)

$$\Rightarrow \quad H(\boldsymbol{C}_a, \boldsymbol{C}_b, X_b) = H(\boldsymbol{C}_a, \boldsymbol{C}_b) \tag{36d}$$

Here, (36c) is due to the Markov chain (35). Equation (36d) means that any receiver can obtain x_b from the index code alone. Consequently, any receiver can also decode the messages of all predecessors of *b* (using the argument in the proof of Lemma 1), which include all messages in the message-disconnected leaf SCC. We have Lemma 11.

APPENDIX C PROOF OF PROPOSITION 2

Let the appended arc be $(v_{\text{inside}} \rightarrow v_{\text{outside}})$. Unlike appending a message-disconnected leaf SCC, here, we change the connectivity of some vertex outside the leaf SCC, i.e., v_{outside} , by creating an incoming arc to it.

The following lemma will be used to prove Proposition 2. *Lemma 12:* Consider the appending step in Definition 7 for a degenerated leaf SCC with some chosen $\mathcal{V}_{S}^{\text{inside}}$ and v_{outside} . From any index code for $(\mathcal{G}, \mathcal{U})$ (the problem instance *before* the appending step), receiver v_{outside} must be able to decode all messages in $\mathcal{V}_{S}^{\text{inside}}$.

Proof: By Lemmas 1 and 2, receiver v_{outside} can decode the messages of all neighbors of $\mathcal{V}_{\text{S}}^{\text{inside}}$, denoted as $\mathcal{N}(\mathcal{V}_{\text{S}}^{\text{inside}})$, as each vertex in $\mathcal{N}(\mathcal{V}_{\text{S}}^{\text{inside}})$ is either grounded, a predecessor of v_{outside} , v_{outside} itself, or a combination of these conditions.

Choose any $a \in \mathcal{V}_{S} \setminus \mathcal{V}_{S}^{\text{inside}}$ and $b \in \mathcal{V}_{S}^{\text{inside}}$. receiver a must decode x_{b} , because there is a path from b to a by the definition of SCC. We will show that if receiver a can decode x_{b} , so can receiver v_{outside} . Consequently, receiver v_{outside} can decode all messages of $\mathcal{V}_{S}^{\text{inside}}$.

There is no edge across $\mathcal{V}_{S}^{\text{inside}}$ and $\mathcal{V}_{S} \setminus \mathcal{V}_{S}^{\text{inside}}$ by definition, meaning that any index codebit cannot be a function of the messages from both the sets. So, we can partition any index code for $(\mathcal{G}, \mathcal{U})$ into $\boldsymbol{c} = (\boldsymbol{c}_1, \boldsymbol{c}_2)$, where \boldsymbol{c}_1 does not contain

⁹Recall that we use upper-case letters to denote the corresponding random variables.

any message of $\mathcal{V}_{S}^{\text{inside}}$, and c_2 contains only the messages of $\mathcal{V}_{S}^{\text{inside}}$ and $\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})$.

Intuitively, any advantage in decoding x_b that receiver a has over receiver v_{outside} is due to knowing x_a , but x_a can help receiver a only in decoding the messages in c_1 , which receiver acan then use to decode the messages in c_2 (which contain x_b). As receiver v_{outside} is able to decode the prior messages of $\mathcal{N}(\mathcal{V}_{\text{S}}^{\text{inside}})$, which contains all the overlap of messages in c_1 and c_2 , receiver v_{outside} is as capable as receiver a in decoding x_b .

We now prove this formally. Recall that C_2 is a function of only messages $(X_{\mathcal{V}_{S}^{\text{inside}}}, X_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})})$, where $X_{\mathcal{S}} \triangleq \{X_i : i \in S\}$ for some set S; C_1 can be written as a function of $(X_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, X_{\mathcal{V}^{\text{remaining}}})$, where $\mathcal{V}^{\text{remaining}} = \mathcal{V} \setminus (\mathcal{V}_{S}^{\text{inside}} \cup \mathcal{N}(\mathcal{V}_{S}^{\text{inside}}))$. Clearly, we have

$$I(\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}}; \boldsymbol{X}_{\mathcal{V}_{\text{S}}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{\text{S}}^{\text{inside}})}) = 0$$
(37a)

$$\Rightarrow H(\mathbf{X}_{\mathcal{V}^{\text{remaining}}} | \mathbf{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}) - H(\mathbf{X}_{\mathcal{V}^{\text{remaining}}} | \mathbf{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \mathbf{X}_{\mathcal{V}_{S}^{\text{inside}}}) = 0$$
(37b)

$$\Rightarrow H(\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{C}_{2}) - H(\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}}, \boldsymbol{C}_{2}) \leq 0$$
(37c)

$$\Rightarrow I(\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}}; \boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{C}_{2}) = 0$$
(37d)

$$\Rightarrow H(\boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{C}_{2}) - H(\boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{X}_{\mathcal{V}^{\text{remaining}}}, \boldsymbol{C}_{2}) = 0 \quad (37e) \Rightarrow H(\boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{C}_{2}, \boldsymbol{C}_{1})$$

$$-H(\boldsymbol{X}_{\mathcal{V}_{S}^{\text{inside}}}|\boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}, \boldsymbol{X}_{\mathcal{V}^{\text{remaining}}}, \boldsymbol{C}_{2}, \boldsymbol{C}_{1}) \leq 0$$
(37f)

$$\Rightarrow I(\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}}; \boldsymbol{X}_{\mathcal{V}_{\text{S}}^{\text{inside}}} | \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{\text{S}}^{\text{inside}})}, \boldsymbol{C}_{2}, \boldsymbol{C}_{1}) = 0.$$
(37g)

Here,

(37a) is due to the independence of the messages,

(37c) is derived because conditioning cannot increase entropy, and C_2 is a function of only messages $(X_{\mathcal{V}_{S}^{\text{inside}}}, X_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})})$, (37d) and (37g) follow from the non-negativity of conditional mutual information,

(37f) is derived because conditioning cannot increase entropy, and C_1 is a function of $(X_{\mathcal{N}(\mathcal{V}_s^{inside})}, X_{\mathcal{V}^{remaining}})$.

Equation (37g) implies that

$$\boldsymbol{X}_{\mathcal{V}^{\text{remaining}}} - (\boldsymbol{X}_{\mathcal{N}(\mathcal{V}^{\text{inside}}_{\text{S}})}, \boldsymbol{C}_{2}, \boldsymbol{C}_{1}) - \boldsymbol{X}_{\mathcal{V}^{\text{inside}}_{\text{S}}}$$
 (38)

forms a Markov chain. Since $a \in \mathcal{V}^{\text{remaining}}$ and $b \in \mathcal{V}^{\text{inside}}_{s}$,

$$X_a - (\boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{\mathsf{S}}^{\mathsf{inside}})}, \boldsymbol{C}_2, \boldsymbol{C}_1) - X_b$$
(39)

also forms a Markov chain.

Since a and b are both in the leaf SCC, receiver a must be able to decode x_b from the index code and its prior information x_a , i.e.,

$$H(X_b|\boldsymbol{C}_1, \boldsymbol{C}_2, X_a) = 0 \quad (40a)$$

$$\Rightarrow \quad H(X_b | \boldsymbol{C}_1, \boldsymbol{C}_2, X_a, X_{v_{\text{outside}}}, \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_{\mathsf{S}}^{\text{inside}})}) = 0 \qquad (40b)$$

$$\Rightarrow \qquad H(X_b | \boldsymbol{C}_1, \boldsymbol{C}_2, X_{v_{\text{outside}}}, \boldsymbol{X}_{\mathcal{N}(\mathcal{V}_s^{\text{inside}})}) = 0 \qquad (40c)$$

$$\Rightarrow \qquad H(X_b | \boldsymbol{C}_1, \boldsymbol{C}_2, X_{v_{\text{outside}}}) = 0. \quad (40d)$$

Here,

(40b) is derived because conditioning cannot increase entropy,

(40c) is due to the Markov chain (39),

(40d) is due to $X_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}$ being a deterministic function of $(C_1, C_2, X_{v_{\text{outside}}})$ because receiver v_{outside} can decode $X_{\mathcal{N}(\mathcal{V}_{S}^{\text{inside}})}$ (which are the messages of the predecessors of vertex v_{outside}).

Equation (40d) implies that knowing $x_{v_{\text{outside}}}$, receiver v_{outside} can decode x_b . Using the argument in the proof of Lemma 1, receiver v_{outside} must be able to decode all messages in the leaf SCC.

Proof of Proposition 2: First, we show that either (24) or (25) must hold.

Consider Case 2 in Definition 7. Similar to the proof of Proposition 1, adding the arc causes the original leaf SCC V_S to become non-leaf. Furthermore, all other leaf SCCs do not change their status, since v_{outside} is a leaf vertex and cannot belong to any leaf SCC. So, we get (25).

For Case 1 in Definition 7, there are three possibilities: v_{outside} is

- Grounded: This case is similar to Case 2 above. The above argument holds even if v_{outside} is not a leaf vertex but is grounded. Thus (25) holds.
- Not grounded and has no directed path to v_{inside} : First, we show there is neither any destruction of existing SCCs, nor creation of any new one. This follows because no cycle is created by the addition of the arc ($v_{inside} \rightarrow v_{outside}$), since $v_{outside}$ has no directed path back to v_{inside} . Next, we need to consider if the appended arc may change some SCCs from leaf to non-leaf or vice versa. For the leaf SCC V_S , the appended arc makes it become non-leaf, since this is an outgoing arc which does not form a cycle. For SCCs other than V_S , the appended arc is an incoming arc, which will not change the SCCs' being leaf or non-leaf. So, we have (25) for this particular case.
- Not grounded and has a directed path to v_{inside} : Here, the appended arc creates a cycle between v_{inside} (in V_S) and $v_{outside}$ (outside V_S). This causes the original leaf SCC V_S to become a larger SCC by assimilating some external nodes (that include $v_{outside}$). If we could show that none of the assimilated external nodes belongs to any leaf SCC (meaning that no other leaf SCC is affected by this appending step), then clearly either (24) or (25) holds depending on whether the enlarged SCC have outgoing arcs or not. Indeed none of the assimilated external nodes belong to any leaf SCC, because they each have a directed path to v_{inside} . Thus, we conclude the argument for this case.

To show (26), we follow the same line of argument as in the proof of Proposition 1. Let $(\mathcal{G}, \mathcal{U})$ and $(\mathcal{G}', \mathcal{U}')$ denote the graphs before and after the appending step, respectively. First, adding an arc increases the decoding requirements and result in $\tilde{\ell}^*(\mathcal{G}', \mathcal{U}') \geq \tilde{\ell}^*(\mathcal{G}, \mathcal{U})$. On the other hand, Lemma 12 says that receiver v_{outside} can decode all messages of $\mathcal{V}_{\text{S}}^{\text{inside}}$ using any index code for $(\mathcal{G}, \mathcal{U})$. This implies that any index code for $(\mathcal{G}, \mathcal{U})$ is also an index code for $(\mathcal{G}', \mathcal{U}')$. So, $\tilde{\ell}^*(\mathcal{G}, \mathcal{U})$ is achievable for the problem instance $(\mathcal{G}', \mathcal{U}')$. So, (26) must hold.

Finally, as v_{inside} already has at least one outgoing arc before the appending step (since it belongs to an SCC), this step does not change the number of vertices with outgoing arcs. We have (27).

APPENDIX D

Proof of Lemma 10

From Sections VI-D1 to VI-D3 and Appendices B to C, we see that, besides the leaf SCC being operated on, all other leaf SCCs remain as leaf SCCs. In addition, pruning and appending neither remove nor add any edge in \mathcal{U} , and thus all other message-connected and message-disconnected leaf SCCs remain as leaf SCCs of the same type. For a degenerated leaf SCC not being operated on, it remains degenerated because we have the following after the pruning/appending operation on another leaf SCC:

- There is no addition or removal of edges in U, and so *V*^{inside} remains unchanged, and the leaf SCC will neither become message-connected nor message-disconnected.
- The number of non-leaf vertices in $\mathcal{V}^{outside}_S$ can only decrease, as all leaf vertices remain leaf vertices, and some non-leaf vertices may become leaf vertices (we might add more leaf vertices into $\mathcal{V}^{outside}_S$, see the argument below); and
- Each neighbor of $\mathcal{V}_{S}^{\text{inside}}$ can still be in $\mathcal{V}_{S}^{\text{outside}}$ or be a predecessor of some vertex in $\mathcal{V}_{S}^{\text{outside}}$. This is because appending other leaf SCCs will not change the neighbor set. If we prune some other leaf SCC by removing all outgoing arcs from a vertex, say v, that happens to be in the path from a neighbor of $\mathcal{V}_{S}^{\text{inside}}$ to $\mathcal{V}_{S}^{\text{outside}}$, we just add v (which has been made a leaf vertex) to $\mathcal{V}_{S}^{\text{outside}}$.

APPENDIX E

A MULTI-SENDER INSTANCE CANNOT BE TRIVIALLY SPLIT INTO MULTIPLE SINGLE-SENDER INSTANCES

Consider the following information-flow graph:



Suppose that there are two senders, one having (x_1, x_2) , and the other (x_3, x_4) . We have a message-disconnected leaf SCC, and the optimal index-codelength is $\tilde{\ell}^* = 4$.

Consider splitting this index-coding instance into two singlesender instances according to the senders' messages. The first instance consists of the first sender with messages (x_1, x_2) . Since there are only two messages, we cannot have vertices 3 and 4 in the information-flow graph, i.e., we have



This first single-sender instance has an optimal index codelength of $\ell^* = 1$, and an optimal code of $x_1 \oplus x_2$.

The second instance can be similarly defined: a sender having (x_3, x_4) and the information-flow graph is as follows:



This second single-sender instance has an optimal index codelength of $\ell^* = 1$, and an optimal code of $x_3 \oplus x_4$. Combining these two instances gives a total codelength of $\tilde{\ell}^* = 2$ with a code $(x_1 \oplus x_2, x_3 \oplus x_4)$. However, this is not an index code for the original instance as receiver 2 cannot decode x_4 . In conclusion, we cannot trivially split a multi-sender uniprior instance into separate single-sender uniprior instances. Having said that, one can do so for *unicast* instances.

REFERENCES

- Z. Bar-Yossef, Y. Birk, T. S. Jayram, and T. Kol, "Index coding with side information," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 1479–1494, Mar. 2011.
- [2] M. J. Neely, A. S. Tehrani, and Z. Zhang, "Dynamic index coding for wireless broadcast networks," in *Proc. 31st IEEE Conf. Comput. Commun.* (*INFOCOM*), Orlando, USA, Mar. 25–30 2012, pp. 316–324.
- [3] F. Arbabjolfaei, B. Bandemer, Y.-H. Kim, E. Şaşoğlu, and L. Wang, "On the capacity region for index coding," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, July 7–12 2013, pp. 962–966.
- [4] A. Blasiak, R. Kleinberg, and E. Lubetzky, "Broadcasting with side information: Bounding and approximating the broadcast rate," *IEEE Trans. Inf. Theory*, vol. 59, no. 9, pp. 292–298, Sept. 2013.
- [5] L. Ong, "Linear codes are optimal for index-coding instances with five or fewer receivers," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, USA, June 29–July 4 2014, pp. 491–495.
- [6] H. Yu and M. J. Neely, "Duality codes and the integrality gap bound for index coding," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 7256–7268, Nov. 2014.
- [7] E. Lubetzky and U. Stav, "Nonlinear index coding outperforming the linear optimum," *IEEE Trans. Inf. Theory*, vol. 55, no. 8, pp. 3544–3551, Aug. 2009.
- [8] A. D. Wyner, J. K. Wolf, and F. M. J. Willems, "Communicating via a processing broadcast satellite," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1243–1249, June 2002.
- [9] D. Gündüz, A. Yener, A. Goldsmith, and H. V. Poor, "The multi-way relay channel," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Seoul, Korea, June 28–July 3 2009, pp. 339–343.
- [10] L. Ong, S. J. Johnson, and C. M. Kellett, "The capacity region of multiway relay channels over finite fields with full data exchange," *IEEE Trans. Inf. Theory*, vol. 57, no. 5, pp. 3016–3031, May 2011.
- [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, June 2006.
- [12] N. Alon, A. Hassidim, E. Lubetzky, U. Stav, and A. Weinstein, "Broadcasting with side information," in *Proc. IEEE Symp. Found. Comput. Sci. (FOCS)*, Philadelphia, USA, Oct. 25–28 2008, pp. 823–832.
 [13] S. El Rouayheb, A. Sprintson, and C. Georghiades, "On the index coding
- [13] S. El Rouayheb, A. Sprintson, and C. Georghiades, "On the index coding problem and its relation to network coding and matroid theory," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3187–3195, July 2010.
- [14] M. A. R. Chaudhry, Z. Asad, A. Sprintson, and M. Langberg, "On the complementary index coding problem," in *Proc. IEEE Int. Symp. on Inf. Theory (ISIT)*, St Petersburg, Russia, July 31–Aug. 5 2011, pp. 244–248.
- [15] L. Ong and C. K. Ho, "Optimal index codes for a class of multicast networks with receiver side information," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Ottawa, Canada, June 10–15 2012, pp. 2223–2228.
- [16] A. S. Tehrani, A. G. Dimakis, and M. J. Neely, "Bipartite index coding," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, USA, July 1–6 2012, pp. 2256–2260.
- [17] S. H. Dau, V. Skachek, and Y. M. Chee, "On the security of index coding with side information," *IEEE Trans. Inf. Theory*, vol. 58, no. 6, pp. 3975–3988, June 2012.
- [18] K. Shanmugam, A. G. Dimakis, and M. Langberg, "Local graph coloring and index coding," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, July 7–12 2013, pp. 1152–1156.
- [19] K. Shanmugam and A. G. Dimakis, "Bounding multiple unicasts through index coding and locally repairable codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, USA, June 29–July 4 2014, pp. 296–300.

- [20] F. Arbabjolfaei and Y.-H. Kim, "Local time sharing for index coding," in Proc. IEEE Int. Symp. Inf. Theory (ISIT), Honolulu, USA, June 29–July 4 2014, pp. 286–290.
- [21] J. B. Ebrahimi and M. J. Siavoshani, "On index coding and graph homomorphism," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Hobart, Australia, Nov. 2–5 2014, pp. 541–545.
- [22] L. Ong, "A new class of index coding instances where linear coding is optimal," in *Proc. Int. Symp. on Netw. Coding (NetCod)*, Aalborg, Denmark, June 27–28 2014.
- [23] F. Arbabjolfaei and Y.-H. Kim, "Structural properties of index coding capacity using fractional graph theory," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Hong Kong, China, June 14–19 2015, pp. 1034–1038.
- [24] C. Thapa, L. Ong, and S. J. Johnson, "A new index coding scheme exploiting interlinked cycles," in *Proc. IEEE Int. Symp. Inf. Theory* (*ISIT*), Hong Kong, China, June 14–19 2015, pp. 1024–1028.
- [25] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," IEEE Trans. Inf. Theory, vol. 60, no. 5, pp. 2856–2867, Nov. 2014.
- [26] M. F. Wong, M. Langberg, and M. Effros, "On a capacity equivalence between network and index coding and the edge removal problem," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, July 7–12 2013, pp. 972–976.
- [27] D. Ozgul and A. Sprintson, "An algorithm for cooperative data exchange with cost criterion," in *Inf. Theory Appl. Workshop (ITA)*, La Jolla, USA, Feb. 6-11 2011.
- [28] I.-H. Hou, Y.-P. Hsu, and A. Sprintson, "Truthful and non-monetary mechanism for direct ata exchange," in *Proc. 51st Allerton Conf. Commun. Control Comput. (Allerton Conf.)*, Monticello, USA, Oct. 2–4 2013, pp. 406–412.
- [29] M. Ji, G. Caire, and A. F. Molisch, "Fundamental limits of distributed caching in D2D wireless networks," in *Proc. IEEE Inf. Theory Workshop* (*ITW*), Seville, Spain, Sept. 9–13 2013.
- [30] M. Sharir, "A strong connectivity algorithm and its applications to data flow analysis," *Comput. Math. Appl.*, vol. 1, no. 7, pp. 67–72, 1981.
- [31] R. Tarjan, "Depth-first search and linear graph algorithms," Siam J. Comput., vol. 2, no. 1, pp. 146–160, June 1972.

Lawrence Ong (S'05–M'10) received the BEng degree (1st Hons) in electrical engineering from the National University of Singapore (NUS), Singapore, in 2001. He subsequently received the MPhil degree from the University of Cambridge, UK, in 2004 and the PhD degree from NUS in 2008. He was with MobileOne, Singapore, as a system engineer from 2001 to 2002. He was a research fellow at NUS, from 2007 to 2008. From 2008 to 2012, he was a postdoctoral researcher at the University of Newcastle, Australia. He was awarded a Discovery Early Career Researcher Award in 2012 and a Future Fellowship in 2014, both by the Australian Research Council. He is currently a Future Fellow at the University of Newcastle

Chin Keong Ho (S'05-M'07) received the B. Eng. (First-Class Hons., Minor in Business Admin.), and M. Eng degrees from the Department of Electrical Engineering, National University of Singapore in 1999 and 2001, respectively. He obtained his Ph.D. degree in 2009 at the Eindhoven University of Technology, The Netherlands, where he concurrently conducted research work in Philips Research. Since August 2000, he has been with Institute for Infocomm Research (I²R), A*STAR, Singapore. He is Deputy Department Head of Cognitive Communications Technology Department and Lab Head of Energy-Aware Communications Lab. His research interest includes green wireless communications with focus on energy-efficient solutions and with energy harvesting constraints; cooperative and adaptive wireless communications; and implementation aspects of multi-carrier and multiantenna communications. He has developed key technologies and led teams in algorithmic development for industrial WiFi projects, and in exploratory work for energy harvesting and wireless power to support wireless communications. His work in unified study of wireless power and wireless communications received the IEEE Marconi Prize Paper Award in 2015.

Fabian Lim received the B.Eng and M.Eng degrees from the National University of Singapore in 2003 and 2006, respectively, and the Ph.D. degree from the University of Hawaii, Manoa in 2010, all in electrical engineering.

He is currently a Research Staff Member in IBM Research Collaboratory, Singapore. Previously from Nov 2014 to March 2016 he was with SK Hynix Memory Solutions Inc, San Jose. From June 2013 to Nov 2014, he was with LSI Corporation, San Jose. He was co-awarded an NSF grant in support of his post-doctorate research, during the period Sept 2010 to May 2013 at Massachusetts Institute of Technology, Cambridge. He has held short-term visiting research positions at Harvard University in 2004 and 2005. From Oct 2005 to May 2006, he was a staff member in the Data Storage Institute in Singapore. From May 2008 to July 2008, he as an intern at Hitachi Global Storage Technologies, San Jose. In March 2009, he was a visitor at the Research Center for Information Security, Japan.