Decentralized Pliable Index Coding

Tang Liu and Daniela Tuninetti
University of Illinois at Chicago, Chicago, IL 60607 USA,
Email: tliu44, danielat@uic.edu

***Abstract*—This paper introduces the *decentralized* Pliable Index CODing (PICOD) problem: a variant of the Index Coding (IC) problem, where a central transmitter serves *pliable* users with message side information; here, pliable refers to the fact that a user is satisfied by decoding *any* t messages that are not in its side information set. In the decentralized PICOD, a central transmitter with knowledge of all messages is not present, and instead users share among themselves massages that can only depend on their local side information set. This paper characterizes the capacity of two classes of decentralized *complete–*S PICOD(t) problems with m messages (where the set $S \subset [m]$ contains the sizes of the side information sets, and the number of users is $n = \sum_{s\in S}\binom{m}{s}$, with no two users having the same side information set): (i) the *consecutive case* $S = [s_{\min} : s_{\max}]$ for some $0 \leqslant s_{\min} \leqslant s_{\max} \leqslant m - t$, and (ii) the *complement-consecutive case* $S = [0 : m - t]\backslash[s_{\min} : s_{\max}]$, for some $0 < s_{\min} \leqslant s_{\max} < m - t$. Interestingly, the optimal code-length for the decentralized PICOD in those cases is the same as for the classical (centralized) PICOD counterpart, except when the problem is no longer pliable, that is, it reduces to an IC problem where every user needs to decode all messages not in its side information set. Although the optimal code-length may be the same in both centralized and decentralized settings, the actual optimal codes are not. For the decentralized PICOD, sparse Maximum Distance Separable (MDS) codes and vector linear index codes are used (as opposed to scalar linear codes).**

I. Introduction

A. Motivation

Index coding (IC), first proposed when considering satellite communication [1], is a simple model to study the impact of message side information at the receivers in broadcast communication networks. The IC consists of one transmitter with m independent messages to be delivered to n users through an error-free broadcast link. Each user has some messages as side information available to it and needs to reliably decode some messages that are not in its side information set; the desired messages for each user are predetermined. In IC, one asks what is the minimum number of transmissions (i.e., minimum code-length) such that every user is able to decode its desired messages successfully. In this paper we are interested in the *decentralized pliable index coding problem*, which is motivated by two variants of IC: Pliable Index CODing (PICOD), and decentralized IC.

The *PICOD* problem is motivated by the flexibility in choosing the desired messages for the users in some practical scenarios, such as online advertisement systems. Firstly proposed in [2], in the PICOD(t) there is a single transmitter, with m message, and n users, with message side information, which are connected via an error-free rate-limited broadcast channel, as in IC. Different from IC, in the PICOD(t) the desired messages at the users are not pre-determined and each user is satisfied whenever it can decode *any* t messages not in its side information set. This provides the transmitter more encoding opportunities, as it now encodes based on its own choice of desired messages for the users. The goal in the PICOD(t) is to find the assignment of desired messages for the users that leads to the smallest possible code-length.

The *decentralized IC* is motivated by peer-to-peer and ad-hoc network, where a central controller / transmitter does not exist and instead communication occurs among peers / users. The decentralized IC can be seen as a special case of the distributed IC [7]. In the distributed IC with m messages, there are $2^m - 1$ servers; each sender has knowledge of a unique subset of the message set (and can thus only encode based on its local knowledge) and is connected to the users through a separate error-free rate-limited link. The decentralized IC is thus a distributed IC where there are as many servers as users, and each server has the same message knowledge as one of the users. The goal for the decentralized IC is to determine the shortest code-length such that all users can decode.

The *decentralized PICOD* proposed in this paper is a combination of the (centralized) PICOD and the decentralized IC, namely, a central transmitter with knowledge of all messages is not present, and instead users share among themselves massages that can only depend on their local side information set. The decentralized PICOD problem is motivated by coded cooperative data exchang [3], [9], distributed storage [5], and distributed computation [12].

B. Past Work

Several achievable schemes have been proposed for PICOD, based on scalar linear codes; the results of [2], [11] show an exponential code-length reduction for PICOD compared to IC.

For converse results, the optimal code-length under the restriction that the transmitter can only use linear schemes was shown in [2] for the *oblivious* PICOD(t), where the transmitter only knows the size of the side information at the users. In [6], we used techniques based on combinatorial design to prove tight converse bounds for some *complete–*S PICOD(t) (see next for a formal definition) problems that generalize of the oblivious class, and showed the information theoretic optimality of scalar linear codes.

The multi-sender IC has been studied in [10], where the focus was on the "single uniprior" case (where users have only one single message as side information). The general multi-sender IC, or distributed IC, was investigated in [7], where converse bounds (leveraging the submodularity of entropy)

and achievable bounds (based on composite IC coding) were proposed; those bounds were numerically verified to match for the case of symmetric rate and symmetric link capacities in all settings with no more than four messages; the use of those bounds in general settings is however problematic because the number of variables involved is exponential in the number of servers (thus double exponential in the number of messages).

C. Contributions

In this paper we derive tight information theoretic converse bounds (i.e., no restrictions on the class of codes used by the users) for two classes of decentralized PICOD(t) problems, namely: (i) the *complement-consecutive complete-S* PICOD(t), and (ii) the *consecutive complete-S* PICOD(t). The complete–S PICOD(t), where S is a subset of $[0:m-t]$ (where m is the number of messages at the transmitter and t the number of messages to be decoded by each user), is a system where all side information sets / users with size indexed by S are present. We say that S is *consecutive* if $S=[s_{\min}:s_{\max}]$ for some $0\leqslant s_{\min}\leqslant s_{\max}\leqslant m-t$, and *complement-consecutive* if $S=[0:m-t]\backslash[s_{\min}:s_{\max}]$ for some $0<s_{\min}\leqslant s_{\max}<m-t$. We characterized the optimal code-length in those cases in [6] for the centralized PICOD(t) case. Here, we examine the decentralized case.

Trivially, a centralized PICOD(t) has optimal block-length no larger than that of the corresponding decentralized problem (because a centralized transmitter can mimic any decentralized transmission scheme). In this work, we thus start by analyzing the decentralized version of those PICOD(t) problems whose optimal code-length we characterized in [6], and use our tight past result as a "trivial centralized converse bound." *Surprisingly, we show that such a "trivial centralized converse bound" is tight whenever the decentralized PICOD(t) remains indeed pliable.* More precisely, we show that by using vector linear codes (in contrast to the simple linear scalar schemes that are optimal in the corresponding centralized setting [6]) we can achieve the "trivial centralized converse bound" except for the case where the problem parameters are such that every user must decode all the messages that are not in its side information set, that is, the problem becomes a special case of the "communications for omniscience" problem [4].

D. Paper Organization

Section II introduces the system model and definitions; Section III lists our main contributions; Section IV, resp. V, provides the proof for consecutive, resp. complement-consecutive, complete–S PICOD(t). Section VI concludes the paper.

E. Notation

Throughout the paper we use capital letters to denote sets, calligraphic letters for family of sets, and lower case letters for elements in a set. For integers $1\leqslant a_1\leqslant a_2$ we let $[a_1:a_2]:=\{a_1,a_1+1,\ldots,a_2\}$, and $[a_2]:=[1:a_2]$. A capital letter as a subscript denotes set of elements whose indices are in the set, i.e., $W_A:=\{w_a:w_a\in W,a\in A\}$. For two sets A and B, $A\backslash B$ is the set that consists all the elements that are in A but not in B.

II. System Model

A decentralized PICOD(t) system consists of: (i) $n\in\mathbb{N}$ users and no central transmitter. The user set is denoted as $U:=\{u_1,u_2,\ldots,u_n\}$. (ii) $m\in\mathbb{N}$ independent and uniformly distributed binary messages of $\kappa\in\mathbb{N}$ bits each. The message set is denoted as $W:=\{w_1,w_2,\ldots,w_m\}$. (iii) User u_i knows the messages indexed by its side information set $A_i\subset[m]$, $i\in[n]$. The collection of all side information sets is denoted as $\mathcal{A}:=\{A_1,A_2,\ldots,A_n\}$, which is assumed globally known at all users. Note that for a decentralized PICOD problem to have a solution, one must have $\cup_{i=1}^n A_i\supset A_j,\forall j\in[n]$, that is, for every user there must be an unknown message that is in the side information set of some other users. (iv) An error-free broadcast link is shared among all users and allows one user to transmit while all the remaining users receive. (v) The codeword $x^{\kappa\ell}:=(x^{\kappa\ell_1},x^{\kappa\ell_2},\ldots,x^{\kappa\ell_n})$ is eventually received by all users, where $\ell:=\sum_{j\in[n]}\ell_j$ and

$$x^{\kappa\ell_j}:=\mathsf{ENC}_j(W_{A_j},\mathcal{A}),\ \forall j\in[n],$$

is the encoding function at user u_j. (vi) The decoding function for user u_j is

$$\{\widehat{w}_1^{(j)},\ldots,\widehat{w}_t^{(j)}\}:=\mathsf{DEC}_j(W_{A_j},x^{\ell\kappa}),\ \forall j\in[n].$$

(vii) A code is *valid* if and only if every user can successfully decode at least t messages not in its side information set, i.e., the decoding functions $\{\mathsf{DEC}_j,\forall j\in[n]\}$ are such that

$$\Pr[\exists\{d_{j,1},\ldots,d_{j,t}\}\cap A_j=\varnothing:\ \{\widehat{w}_1^{(j)},\ldots,\widehat{w}_t^{(j)}\}\neq\{w_{d_{j,1}},\ldots,w_{d_{j,t}}\}]\leqslant\epsilon,$$

for some $\epsilon\in(0,1)$. For a valid code, $\{\widehat{w}_1^{(j)},\ldots,\widehat{w}_t^{(j)}\}=\{w_{d_{j,1}},\ldots,w_{d_{j,t}}\}$ is called the *desired message set* for user $u_j,\ j\in[n]$. The indices of the desired messages are denoted as $D_j:=\{d_{j,1},\ldots,d_{j,t}\}$ where $D_j\cap A_j=\varnothing,\forall j\in[n]$. The choice of desired messages for the users is denoted as $\mathcal{D}=\{D_1,D_2,\ldots D_n\}$. (viii) The goal is to find a valid code with minimum length, that is, to determine

$$\ell^\star:=\min\{\ell:\exists\text{ a valid }x^{\kappa\ell}\text{ for some }\kappa\}.$$

In the following we shall focus on the decentralized *complete–S* PICOD(t), for a given set $S\subseteq[0:m-t]$. In this complete–S system, there are $n:=\sum_{s\in S}\binom{m}{s}$ users, where no two users have the same side information set, i.e., all possible users with distinct side information sets that are subsets of size s of the message set, for all $s\in S$, are present in the system. In particular, we focus on the *consecutive complete-S* PICOD(t) and the *complement-consecutive complete-S* PICOD(t), where we say that S is *consecutive* if $S=[s_{\min}:s_{\max}]$ for some $0\leqslant s_{\min}\leqslant s_{\max}\leqslant m-t$ (i.e., S contains consecutive integers, from $s_{\min}$ to $s_{\max}$), and *complement-consecutive* if $S=[0:m-t]\backslash[s_{\min}:s_{\max}]$ for some $0<s_{\min}\leqslant s_{\max}<m-t$ (note that the set S includes elements 0 and $m-t$). We characterized the optimal centralized code-length in those two cases in [6] and we examine here the decentralized version. Note that $S=\{0\}$ is not considered since it violates the condition $\cup_{i=1}^n A_i\neq A_j,\forall j\in[n]$.

III. Main Results

The main contributions of this paper are as follows.

Theorem 1 (consecutive)**.** *For the decentralized complete–S PICOD(t) with m messages and $S = [s_{\min} : s_{\max}]$ for some $0 \leqslant s_{\min} \leqslant s_{\max} \leqslant m - t$, the optimal code-length is*

$$\ell^\star = \begin{cases} \frac{\binom{m}{m-t}}{\binom{m}{m-t}-1} t, & s_{\max} = s_{\min} = m - t, \\ \min\{s_{\max} + t, m - s_{\min}\}, & \textit{otherwise.} \end{cases} \tag{1}$$

Theorem 2 (complement-consecutive)**.** *For the decentralized complete–S PICOD(t) with m messages and $S = [0 : m - t] \backslash [s_{\min} : s_{\max}] = [0 : s_{\min} - 1] \cup [s_{\max} + 1 : m - t]$ for some $0 < s_{\min} \leqslant s_{\max} < m - t$, the optimal code-length is*

$$\ell^\star = \min\{m, |S| + 2t - 2\}. \tag{2}$$

Proof details are in Sections IV and V. Note:

1) Theorem 2 says that, for the same parameters of $(m, t, s_{\min}, s_{\max})$, the centralized and the decentralized settings have the same optimal code-length; similarly for Theorem 1, except for the case $s_{\max} = s_{\min} = m - t$.
2) Having the same optimal code-length does not necessarily imply that the same code is optimal in both cases. In [6], we showed that for the centralized setting simple *scalar linear* codes are optimal; in particular, the central transmitter either sends $\ell^\star$ distinct messages one by one, or $\ell^\star$ random linear combinations of all the messages. Clearly, the former strategy can be implemented in a decentralized setting, but not the latter. In this case we use *vector linear* codes; in particular, our achievable scheme uses *sparse Maximum Distance Separable (MDS) codes*.
3) When necessary to distinguish the optimal code length of the centralized and decentralized settings, we shall use the notation $\ell^{\star,\text{cen}}$ and $\ell^{\star,\text{dec}}$, respectively. Note that $\ell^{\star,\text{dec}} = \ell^\star$ where $\ell^\star$ was defined in Section II.
   Among all PICOD cases studied in this work, the only case where the decentralized optimal code-length $\ell^{\star,\text{dec}}$ is strictly larger than the corresponding centralized optimal code-length $\ell^{\star,\text{cen}}$ is when $s_{\min} = s_{\max} = m - t$. This is the only case in centralized PICOD(t) where $\ell^{\star,\text{cen}} = t$. Since $\ell^{\star,\text{dec}} > t$ for all decentralized PICOD(t) (as what is sent by a user is not useful for that user), our results show that for the consecutive and complete consecutive complete–S PICOD(t), $\ell^{\star,\text{dec}} \neq \ell^{\star,\text{cen}}$ if only if $\ell^{\star,\text{cen}} = t$; interestingly, this is also the case where the PICOD problem "looses its pliability," that is, it reduces to an IC problem where every user needs to decode all messages not in its side information set.
4) Theorems 1 and 2 can be extended to all the centralized complete–S PICOD(t) that we have been solved in [6], which are not reported here because of space limitations. In those cases too we obtained $\ell^{\star,\text{dec}} = \ell^{\star,\text{cen}}$ whenever $\ell^{\star,\text{cen}} \neq t$; and $\ell^{\star,\text{dec}} = \frac{n}{n-1} t$ if $\ell^{\star,\text{cen}} = t$. An intriguing question is whether this holds true for all complete–S PICOD(t), even those not solved by the technique in [6]. Answering this question is part of ongoing work.
5) A similar proof technique can show that for the decentralized PICOD(1) where the network topology hypergraph is a circular-arc, the optimal code-length is: $\ell^\star = 2$ if 1-factor does not exists; or $\ell^\star = \frac{p}{p-1}$, where p is the maximum size of the 1-factor of the network topology hypergraph – see definitions in [6]. This serves a tight bound for the decentralized PICOD beyond the complete–S case. Finding tight bounds for the general decentralized PICOD(t) is one direction for future work.

IV. Proof for Theorem 1

We split the proof into sub cases. For $\ell^{\star,\text{cen}} = \min\{s_{\max} + t, m - s_{\min}\} < t$ for the centralized consecutive complete–S PICOD(t), in which case $\ell^{\star,\text{cen}} = \ell^{\star,\text{dec}} = \ell^\star$, we study separately the cases $s_{\max} + t \leqslant m - s_{\min}$ (Section IV-A) and $t < m - s_{\min} < s_{\max} + t$ (Section IV-B). The case $\ell^{\star,\text{cen}} = t$ is studied in Section IV-C, in which case $\ell^{\star,\text{cen}} < \ell^{\star,\text{dec}} = \ell^\star = \frac{\binom{m}{m-t}}{\binom{m}{m-t}-1} t$ and is only possible for $s_{\min} = s_{\max} = m - t$.

A. Case $s_{\max} + t \leqslant m - s_{\min}$

We send $s_{\max} + t$ messages, one at a time. This can be done in a decentralized setting since each message is in the side information set of at least one user. Therefore, such a user can transmit the message to the rest of the users in one channel use. This achievable scheme is optimal since $s_{\max} + t$ is the optimal code-length for the corresponding centralized setting. We thus conclude $\ell^\star = s_{\max} + t$ for $s_{\max} + t \leqslant m - s_{\min}$.

B. Case $t < m - s_{\min} < s_{\max} + t$

We show that in this case a decentralized scheme with $m - s_{\min}$ transmissions can satisfy all users; being $m - s_{\min}$ the optimal code-length for the corresponding centralized setting, such a scheme is thus optimal. In the centralized case, the optimal code involves $m - s_{\min}$ linearly independent linear combinations of all the messages, or alternatively an MDS code; this is not a possible decentralized scheme because at most $s_{\max} + t$ messages can be used to produce a valid code (assuming that a user has side information set of size $s_{\max}$ and sends after having decoded t messages).

Next, when describing achievable schemes, instead of working with messages and codewords in bits (as done in the description of the channel model in Section II), we represent each message of κ bits as one symbols in the finite field $\mathbb{F}_{2^\kappa}$. With an abuse of notation, we also let x^ℓ denote the codeword of length ℓ symbols from the finite field $\mathbb{F}_{2^\kappa}$, and where each symbol corresponds to a transmission by a user. A linear code for the decentralized system is thus $x^\ell = \mathbf{G} w^m$, where $\mathbf{G}$ is the code generator matrix of size $\ell \times m$ and w^m is the vector of length m containing all the messages.

For a valid optimal decentralized linear code, we look for a matrix $\mathbf{G} = [\mathbf{C}, \mathbf{0}]$, where $\mathbf{0}$ is zero matrix of size $\ell^\star \times (m - s_{\max} - t)$ with $\ell^\star = m - s_{\min}$, and where $\mathbf{C}$ is a matrix of size $\ell^\star \times (s_{\max} + t)$ that satisfies two conditions:

1) [C1] each row has at most $s_{\max}$ non-zero elements, and
2) [C2] any submatrix of p columns, with $t \leqslant p \leqslant \ell^\star$, has rank p / is full rank.

The reason for these conditions is as follows. Each row of $\mathbf{G}$ is the encoding vector used by a user; C1 is because a user knows at most $s_{\max}$ messages (in its side information set). C2 is for successful decoding at the users; once the contribution of the messages in the side information set has been subtracted off from the code, each user sees a subset of the remaining messages encoded by a full rank submatrix of p columns; the range of p is because each user must decode at least t messages, thus $t \leqslant p$, and at most all messages in the code that are not in the side information, thus $p \leqslant \ell^\star$.

Note that condition C2 is equivalent to require that all $\ell^\star \times \ell^\star$ submatrices of $\mathbf{G}$ are full rank. This is because any submatrix obtained by taking a subset of columns of a full rank square matrix is full rank. Therefore, instead of having to consider all possible sets of p columns in condition C2, we only look at submatrix of size $\ell^\star \times \ell^\star$, which is the so-called MDS-property of a linear code of dimension $\ell^\star$. We show that the desired matrix $\mathbf{G}$ exists as a *spare MDS code* generator matrix for sufficiently large κ, that is, for a sufficiently large field size.

We now introduce the "zero pattern" matrix for the spare MDS code generator matrix. The "zero pattern" matrix $\mathbf{Z} \in \{0,1\}^{(m-s_{\min})\times(s_{\max}+t)}$ of $\mathbf{C}$ is a matrix whose entry is 0 if the corresponding entry in $\mathbf{C}$ is 0, and 1 otherwise. Consider the following $\mathbf{Z} = [z_{ij}]$

$$z_{ij} = \begin{cases} 1, & \text{for } 0 \leqslant i+j \pmod{(s_{\max}+t)} \leqslant s_{\max}-1, \\ 0, & \text{otherwise.} \end{cases}$$

Let $Z_i := \{j \in [s_{\max}+t] : z_{ij} = 0\}$ be the set of the zero entries in the ith row, $|Z_i| = t, \forall i \in [m - s_{\min}]$. Since $s_{\max}+t > m - s_{\min}$, we have $Z_i \neq Z_j$, $i \neq j$. Therefore, all Z_i are different "shifted" version of Z_1. In $\cap_{i\in P} Z_i$ there are $|P|-1$ "shifts", which reduce the size of the intersection by at least $|P|-1$. We then have the inequality

$$|P| + |\cap_{i\in P} Z_i| \leqslant |P| + t - (|P|-1) = t+1 \leqslant \ell^\star,$$

which is known as the "MDS condition"(which is sufficient for the existence of an MDS generator matrix over some finite field [8]). Therefore, there exists a matrix $\mathbf{C}$ that satisfies conditions C1 and C2 with the specified "zero pattern" $\mathbf{Z}$. From [8], a finite field of size $m - s_{\min} + s_{\max} + t - 1$ suffices. Since $\mathbf{G}$ satisfies condition C1, this code thus can be generated in a distributed way.

After receiving the codewords of length $\ell^\star = m - s_{\min}$, user u_i subtracts off the messages in its side information set A_i and is left with a linear code for the messages $W_{[s_{\max}+t]\setminus A_i}$. Condition C2 guarantees that all users can decode at least t messages that are not in their side information. Therefore all users can be satisfied by this code of length $m - s_{\min}$.

C. Case $s_{\min} = s_{\max} = m - t$

Let $s := s_{\min} = s_{\max} = m - t$. This is the case where the "trivial centralized converse bound" $\ell^{\star,\text{cen}} = \min\{m - s, s + t\} = t \leqslant \ell^\star$ is not tight, and for which we want to show $\ell^\star = \frac{t\binom{m}{s}}{\binom{m}{s}-1} > t = \ell^{\star,\text{cen}}$. In this case, the decentralized PICOD(t) becomes an actual multicast decentralized IC problem, we must show both achievability and converse.

1) Converse: An intuitive explanation for the converse proof is as follows. The $n := \binom{m}{s}$ users in the system are symmetric, i.e., by relabeling the messages we can swap any pair of users. Therefore all users have the same "chance" $1/n$ to be the one who sends part of the overall codeword x^ℓ. In the decentralized setting, the part of x^ℓ sent by a user is generated based on its own side information set, and such a transmission cannot benefit the transmitting user. Therefore, at most a fraction $\frac{n-1}{n}$ of x^ℓ can be useful for each user. Since each transmission can convey at most one message, in order to let each user decode at least t messages, the total number of transmissions satisfies $\frac{n-1}{n}\ell \geqslant t$.

We next provide the formal proof for the converse. Let $\ell_i\kappa$ be the number of bits sent by user $u_i, i \in [n]$, and $x^{\kappa\ell} := (x^{\kappa\ell_1}, x^{\kappa\ell_2}, \ldots, x^{\kappa\ell_n})$ be the overall codeword used for decoding by the users, with $\ell := \sum_{i\in[n]} \ell_i$. With an abuse of notation, let $x^{(\ell-\ell_i)\kappa}$ indicate the bits in the transmitted codeword $x^{\ell\kappa}$ that were not sent by user $u_i, i \in [n]$.

By Fano's inequality, with $\lim_{\kappa\to\infty} \epsilon_\kappa = 0$, we have

$$\begin{aligned} \ell\kappa\epsilon_\kappa &\geqslant H(W_{D_i}|x^{\ell\kappa}, W_{A_i}) = H(W_{D_i}|x^{(\ell-\ell_i)\kappa}, W_{A_i}) \\ &= H(W_{D_i}|W_{A_i}) - I(W_{D_i}; x^{(\ell-\ell_i)\kappa}|W_{A_i}) \\ &= H(W_{D_i}) - I(W_{D_i}; x^{(\ell-\ell_i)\kappa}|W_{A_i}). \end{aligned}$$

Therefore, for $\forall i \in [n]$, we have

$$\begin{aligned} (\ell-\ell_i)\kappa &\geqslant H(x^{(\ell-\ell_i)\kappa}) \geqslant H(x^{(\ell-\ell_i)\kappa}|W_{D_i}) \\ &\geqslant I(W_{D_i}; x^{(\ell-\ell_i)\kappa}|W_{A_i}) \\ &\geqslant H(W_{D_i}) - \ell\kappa\epsilon_\kappa \geqslant t\kappa - \ell\kappa\epsilon_\kappa, \end{aligned}$$

and therefore, for large enough κ, by summing the above inequalities we obtain the converse bound

$$\ell \geqslant \frac{nt}{n-1} = \frac{\binom{m}{s}}{\binom{m}{s}-1} t. \tag{3}$$

2) Achievability: The achievability involves message splitting and random linear coding, i.e., we use a vector linear code, in contrast to the scalar linear code used in Section IV-B.

We split each message into f sub-messages, $w_i = [w_{i,1}, w_{i,2}, \ldots, w_{i,f}], i \in [m]$. The size of the sub-message is κ/f bits, which is assumed to be an integer. The parameter f will be appropriately chosen later. Each sub-message is thus on the finite field $\mathbb{F}_{2^{\kappa/f}}$. Each user uses $\ell' = \frac{f\ell}{n}$ sub-timeslots (as the messages are split into f pieces, the time slots are split into f pieces as well) to transmit. In each sub-timeslot the user transmits a linear combination of all the sub-messages it has in its side information set, i.e., at sub-timeslot h, user u_i transmits $\sum_{g\in A_i, j\in[f]} a_{gj}(h) w_{g,j}$, where the coefficients $a_{gj}(h)$ are on $\mathbb{F}_{2^{\kappa/f}}$. The linear code has generator matrix $\mathbf{G}$ consisting of $a_{gj}(h)$ for $g \in [m], j \in [f], h \in [f\ell]$, and of size $n\ell' \times mf$. Each row of $\mathbf{G}$ has at most sf nonzero entries.

For each user, among all $n\ell'$ sub-timeslots, only $(n-1)\ell'$ are useful for its decoding since the other ℓ' sub-timeslots are transmitted by itself. Therefore, we choose ℓ' and f such that

$$(n-1)\ell' = (m-s)f, \; n = \binom{m}{s}, \; \ell' = \frac{f\ell}{n}.$$

For each user, the submatrix of $\mathbf{G}$ corresponding to what all other users have sent needs to be a full rank square matrix of size $(n-1)\ell' \times (m-s)f$ so that each user can successfully decode, i.e., every submatrix of $\mathbf{G}$ formed by $(m-s)f$ columns is full rank. Similarly to the proof in Section IV-B, the "MDS condition" on its zero-pattern matrix is as follows

$$|P| + |\cap_{i\in P} Z_i| \leqslant n\ell' + |P| - \ell' - \frac{|P|-\ell'}{\ell'} f \leqslant n\ell'.$$

Therefore the proposed code generator matrix $\mathbf{G}$ exists for some large enough κ. By this scheme each user decodes all the $(m-s)f$ sub-messages that are not in its side information.

The total number of transmissions by this scheme is

$$\ell = \frac{\ell'}{f} n = \frac{1}{f} \frac{f(m-s)}{n-1} n = \frac{nt}{n-1}, \tag{4}$$

which coincides with the converse bound in 3. The achievability scheme is information theoretically optimal.

Note that a scalar linear code can only achieve an integer number of transmissions. Thus our vector linear code strictly outperforms scalar linear codes in this case.

V. Proof for Theorem 2

Also for this decentralized complement-consecutive complete–S PICOD(t), where $S = [0:m-1]\backslash[s_{\min}:s_{\max}]$ for some $0 < s_{\min} \leqslant s_{\max} < m-t$, we need to show a decentralized achievable scheme that meet the "trivial centralized converse bound."

In the centralized case, the achievable scheme consists of two scalar linear codes: one to serve all the users with side information of size in $[0:s_{\min}-1]$, and the other to serve all the users with side information of size in $[s_{\max}+1:m-t]$. Also for the decentralized scheme, we separate the users into these two groups: $U_1 = \{u_i : |A_i| \in [0:s_{\min}-1]\}$ and $U_2 = \{u_i : |A_i| \in [s_{\max}+1:m-t]\}$. The analysis of the achievability scheme is divided into two parts: $s_{\min}-1+t < s_{\max}+1 = m-t$ and the rest.

A. Case $s_{\min}-1+t < s_{\max}+1 = m-t$

In this case the decentralized scheme is different from the centralized one. This is because U_2 in this case represents a consecutive complete–S case discussed in Section IV-C, where the centralized converse bound is not tight. Therefore, we can not treat the problem of serving the users in U_1 and U_2 as two independent subproblems, as the scheme does in centralized case. The achievability scheme takes two steps:

- Step 1: Send messages $W_{[s_{\min}-1+t]}$ one by one. All users in U_1 are satisfied. $s_{\min}-1+t \geqslant t$ messages are sent in this step. Since all users in U_2 have side information sets of size $s_{\max}+1 = m-t$, there exists at least one user in U_2 that has been satisfied in the first step.
- Step 2: The user in U_2 that was satisfied in Step 1 has the knowledge of all messages and can thus act as the centralized transmitter of the centralized PICOD(t) [6], sending t linearly independent linear combinations of all messages. Since all users in U_2 have t messages not in the side information, by having t linear independent linear combinations of all messages, all users in U_2 are satisfied.

It thus takes $s_{\min}-1+t+t = |S|+2t-2$ number of transmissions to satisfy all users.

B. Other Case

The achievable scheme in Section IV-A satisfies the users in U_1 by using $s_{\min}-1+t$ transmissions. The achievable scheme in Section IV-B satisfies the users in U_2 by using $m-(s_{\max}+1)$ transmissions. Therefore, the total number of transmissions is $s_{\min}-1+t+m-s_{\max}-1 = |S|+2t-2$.

Note that $\ell = m$ is a trivially achievable number of transmissions for the decentralized setting as well, we conclude for decentralized complement-consecutive complete–S PICOD(t) the optimal number of transmissions is $\ell^\star = \min\{m, |S|+2t-2\}$, which is the same as the centralized setting.

VI. Conclusion

In this paper we introduced and found the capacity of some decentralized complete–S PICOD(t) problems. For most cases we found that the optimal code-length for the decentralized setting is the same as for the centralized one. Among the cases we have explored, we found that when all users request all the messages that are not in their side information set then the decentralized PICOD has a strictly larger optimal code-length then the centralized one. Whether there are other cases where the centralized and decentralized settings have the same code-length, and if there is a fundamental connection between lack of "pliability" and different code-lengths between centralized and decentralized settings, is part of ongoing work.

References

[1] Y. Birk and T. Kol, "Informed-source coding-on-demand (ISCOD) over broadcast channels," *Proc. IEEE 17th INFOCOM*, pp. 1257–1264, 1998.

[2] S. Brahma and C. Fragouli, "Pliable index coding," *IEEE Trans. Information Theory*, vol. 61, no. 11, pp. 6192–6203, Nov 2015.

[3] T. A. Courtade and R. D. Wesel, "Coded cooperative data exchange in multihop networks," *IEEE Trans. Information Theory*, vol. 60, no. 2, pp. 1136 – 1158, Feb. 2014.

[4] N. Ding, C. Chan, Q. Zhou, R. A. Kennedy, and P. Sadeghi, "Determining optimal rates for communication for omniscience," *IEEE Trans. Information Theory*, vol. 64, no. 3, pp. 1919 – 1944, March 2018.

[5] M. Ji, G. Caire, and A. F. Molisch, "Fundamental limits of caching in wireless D2D networks," *IEEE Trans. Information Theory*, 2016.

[6] T. Liu and D. Tuninetti, "Tight information theoretic converse results for some pliable index coding problems," *arXiv:1810.02451*, 2018.

[7] Y. Liu, P. Sadeghi, F. Arbabjolfaei, and Y.-H. Kim, "Capacity theorems for distributed index coding," *arXiv:1801.09063*, 2018.

[8] S. Lovett, "MDS matrices over small fields: A proof of the gm-mds conjecture," *arXiv:1803.02523v2*, 2018.

[9] N. Milosavljevic, S. Pawar, S. E. Rouayheb, M. Gastpar, and K. Ramchandran, "Efficient algorithms for the data exchange problem," *IEEE Trans. Information Theory*, vol. 62, no. 4, pp. 1878 – 1896, April 2016.

[10] L. Ong, C. K. Ho, and F. Lim, "The single-uniprior index-coding problem: The single-sender case and the multi-sender extension," *IEEE Trans. Information Theory*, vol. 62, no. 6, pp. 3165 – 3182, June 2016.

[11] L. Song and C. Fragouli, "A polynomial-time algorithm for pliable index coding," *IEEE Trans. on Information Theory*, vol. 64, no. 2, pp. 979 – 999, Feb 2018.

[12] L. Song, S. R. Srinivasavaradhan, and C. Fragouli, "The benefit of being flexible in distributed computation," *arXiv:1705.08464v2*.