Pliable Index Coding via Conflict-Free Colorings of Hypergraphs

Prasad Krishnan, Rogers Mathew, Subrahmanyam Kalyanasundaram

Abstract

We present a hypergraph coloring based approach to pliable index coding (PICOD). We represent the given PICOD problem using a hypergraph consisting of m messages as vertices and the request-sets of the n clients as hyperedges. A *conflict-free coloring* of a hypergraph is an assignment of colors to its vertices so that each hyperedge contains a uniquely colored vertex. We show that various parameters arising out of conflict-free colorings (and some new variants) of the PICOD hypergraph result in new upper bounds for the optimal PICOD length. Using these new upper bounds, we show the existence of single-request PICOD schemes with length $O(\log^2 \Gamma)$, where Γ is the maximum number of hyperedges overlapping with any hyperedge. For the t-request PICOD scenario, we show the existence of PICOD schemes of length $\max(O(\log \Gamma \log m), O(t \log m))$, under some mild conditions on the graph parameters. These results improve upon earlier work in general. We also show that our achievable lengths in the t-request case are asymptotically optimal, up to a multiplicative factor of $\log t$. Our existence results are accompanied by randomized constructive algorithms, which have complexity polynomial in the parameters of the PICOD problem, in expectation or with high probability.

I. INTRODUCTION

The Index Coding problem introduced by Birk and Kol in [2] consists of a system with a server containing m messages and n receivers connected by a broadcast channel. Each receiver has a

Dr. Krishnan is with the Signal Processing & Communications Research Center, International Institute of Information Technology Hyderabad, India, email:prasad.krishnan@iiit.ac.in. Dr. Mathew and Dr. Kalyanasundaram are with the Department of Computer Science and Engineering, Indian Institute of Technology Hyderabad, email: {rogers,subruk}@iith.ac.in.

Acknowledgements: The first author acknowledges support from SERB-DST projects MTR/2017/000827 and CRG/2019/005572. The second author was supported by a grant from the Science and Engineering Research Board, Department of Science and Technology, Govt. of India (project number: MTR/2019/000550). The third author acknowledges SERB-DST projects MTR/2020/000497 and CRG/2022/009400 for supporting this research.

Parts of an earlier version of this work were presented virtually (please see [1]) at the IEEE International Symposium on Information Theory, 2021, held in Melbourne, Australia, during July 12-20, 2021.

subset of the messages at the server as *side-information* and demands a particular new message. The goal of the index coding problem is to design a transmission scheme (called an *index code*) at the server which uses less number of transmissions (the *length* of the index code) to serve all receivers. The index coding problem is a canonical problem in information theory and has been addressed by a variety of techniques, including graph theory [3], [4], linear programming [5], and interference alignment [6].

A variant of the index coding problem, called pliable index coding (PICOD), was introduced by Brahma and Fragouli in [7]. The pliable index coding problem relaxes the index coding setup, such that each receiver requests *any* message which is not present in its side-information (i.e., any message from its *request-set*). It was shown in [7] that finding the optimal length of a PICOD problem is NP-hard in general. However, the existence of a code with length $O(\min\{\log m(1+\log^+(\frac{n}{\log m})), m, n\})$ was proved using a probabilistic argument (where $\log^+(x) = \max\{0, \log(x)\}$). When $m = n^\delta$ for some constant $\delta > 0$, this means that $O(\log^2 n)$ is sufficient. Some algorithms for designing pliable index codes based on greedy and set-cover techniques were also presented and compared in [7]. In [8], a polynomial-time algorithm was presented for general PICOD problems which achieves a length $O(\log^2 n)$. Thus, unlike the index coding problem which has instances for which the required length can be $\Theta(n)$ (for instance, the directed n-cycle problem [3]), much fewer transmissions are sufficient in general for PICOD instances. Further, the work [8] also extended this algorithm to the case of t-request PICOD, where each receiver wants any t messages from its request-set (for some $t \geq 1$), or the entire request-set if the request-set has less than t messages. Specifically, an achievable scheme is shown in [8] for the t-request case, which achieves length $O(t \log n + \log^2 n)$. For several special classes of PICOD problems, distinguished by the structure of the side-information or request-sets of the receivers, achievability and converse results were presented in [9], [10]. Converse techniques were further developed in [11], [12], using which the optimal lengths of specific classes of PICOD problems were obtained. Other extensions of PICOD such as vector pliable index codes [8], multiple requests [7], [8], secure pliable index codes [13] and decentralized pliable index codes [14] have been studied recently. Pliable index coding has also been proposed for efficient data exchange in real-world applications, such as in the data shuffling phase of distributed computing [15].

In this work, we present a graph coloring approach to pliable index coding. A *conflict-free coloring* of a hypergraph is an assignment of labels to its vertices so that each hyperedge of the hypergraph contains at least one vertex which has a label distinct from others. Conflict-free colorings were introduced by Even et al. in [16], motivated by a problem of frequency assignment in wireless communications. Since then, it has been extensively studied in the context of general hypergraphs, hypergraphs induced by neighborhoods in graphs, hypergraphs induced by simple paths in a graph, hypergraphs that naturally arise in geometry, etc. See [17] for a survey on conflict-free colorings.

Any PICOD problem can be equivalently represented using a hypergraph with the vertices representing the messages (taking values in some finite field), and the request-sets as hyperedges. We show that conflict-free colorings (and its variants) of this hypergraph give rise to achievability schemes for PICOD. By leveraging this connection, we use relevant tools from the probabilistic method (recalled in Appendix A) to show the existence of pliable index coding schemes with small lengths (as functions of the parameters of the given problem). Specifically, the lengths of PICOD schemes we identify are directly proportional to $\log \Gamma$, where Γ is defined as the maximum number of other edges that any edge in the PICOD hypergraph intersects with. A small value of Γ signifies that any request-set has common messages with only a few other request-sets. This suggests that more coding opportunities are available, which means that PICOD schemes with small lengths would suffice. This is the intuition behind the proportionality relationship of the length of our PICOD schemes to Γ. In a single-server multi-client communication scenario (which can be modelled as a PICOD problem), a small value of Γ can arise when the symbols in the side-information sets (and therefore, the request-sets) of distinct clients are considerably distinct. For instance, consider a data shuffling phase of a distributed computing system, in which distinct data present at different clients have to be shuffled between. The PICOD problem that arises in this data shuffling setup (as presented in [15]) would naturally have a small value of Γ. Alternatively, the server itself could carefully choose a subset of clients, which have fairly disjoint side-information sets, to communicate with in a given round. Thus, the value of Γ in the PICOD setup arising from this choice of subset of clients would be small.

Apart from the existence results regarding the PICOD schemes, we also provide randomized algorithms for constructing these codes, which require time-complexity polynomial in the parameters of the PICOD problem in expectation or with high probability.

In Sections III and IV, we present results for single-request PICOD, while in Section V we present results for the t-request PICOD scenario. In Section VI, we extend some of our results to the k-vector PICOD scenario, in which each message is considered as a k-length vector. Our specific contributions and the organization of this paper are given below. In the following list of contributions, we denote by $\mathcal{H}$ be the PICOD hypergraph with m vertices and n hyperedges, where every hyperedge overlaps with at most Γ other hyperedges.

- We briefly review the PICOD problem setup in Section II and conflict-free colorings in Section III. In Section III-A, we show that the optimal length of a PICOD problem is at most the conflict-free chromatic number of the hypergraph corresponding to the PICOD problem (Theorem 2). The conflict-free chromatic number of a hypergraph is the smallest number of colors required to color the hypergraph using a conflict-free coloring.
- In Section III-B, we define the notion of *conflict-free collection of colorings* of hypergraphs. We call the corresponding chromatic number the *conflict-free covering number*. We show that the conflict-free covering number of the PICOD hypergraph is generally a tighter upper bound for the optimal PICOD length than the conflict-free chromatic number (Theorem 3 in Section III-B). By obtaining an upper bound on the conflict-free covering number (Lemma 6 in Section III-C), we show that $O(\log^2 \Gamma)$ transmissions suffice for scalar PICOD schemes (Theorem 4 in Section III-D). In Subsection III-D, we present a randomized algorithm to obtain a PICOD scheme of length $O(\log^2 \Gamma)$ for a given PICOD hypergraph. This algorithm has time complexity that is polynomial in the parameters of the problem with high probability. This new upper bound $O(\log^2 \Gamma)$ on the optimal PICOD length improves over known bounds [7], [8] for some parameter ranges. Further, the PICOD schemes obtained in Section III are finite-field oblivious (i.e., they do not demand any condition on the size of the finite field under consideration).
- In Section IV, we define a parameter called *local* conflict-free chromatic number and show a PICOD scheme whose length is upper bounded by this parameter (Theorem 5). We show that the local conflict-free chromatic number can be strictly smaller than the conflict-free chromatic number, thus improving the upper bound in Section III-A. Using conflict-free collection of colorings, in Section IV-A, we also generalize the covering number of Section III-B to its local variant and show the corresponding PICOD scheme (Theorem 6). The PICOD schemes obtained based on these local conflict-free parameters use Maximum Distance Separable (MDS) codes of length linear in the system parameters.
- In Section V, we present results for the t-request PICOD problem. We show that PICOD schemes for the t-request scenario can be obtained from t-strong conflict-free colorings of

the PICOD hypergraph, which have the property that each hyperedge E 'sees' $\min(t, |E|)$ colors exactly once. The corresponding chromatic number is called the *t-strong conflict-free chromatic number*. Analogous to conflict-free covering number, local conflict-free chromatic number and local conflict-free covering number, we define "t-strong" variants of each of these parameters. We show that each of these parameters bound the length of an optimal t-request PICOD scheme (Theorem 7) from above.

- In Sections V-A and V-B, we show that the t-strong local covering number and the t-strong covering number are both upper bounded by $\max(O(\log \Gamma \log m), O(t \log m))$, under some mild conditions on the parameters of the problem. These upper bounds are coupled with randomized algorithms to generate the respective colorings, which lead to corresponding PICOD schemes of length $\max(O(\log \Gamma \log m), O(t \log m))$, by leveraging Theorem 7. Our randomized constructions for these PICOD schemes have running time polynomial in the parameters of the problem, in expectation. These results are summarized in Theorem 8. In general, these achievability schemes give an improvement over the results of [8] where a PICOD scheme of length $O(t \log n + \log^2 n)$ was constructed. In Section V-C, we show that the lengths of the codes yielded by the above constructions are asymptotically tight up to a multiplicative factor of $\log t$. We do this by constructing a class of PICOD hypergraphs and showing a lower bound on the optimal PICOD length for the same (Lemma 13), by utilizing a result from [18].
- In Section VI, we give the correspondences between designing k-vector PICOD schemes and constructing *k-fold conflict-free colorings*, which is a generalization of conflict-free coloring and involves assigning k-sized sets of labels to the vertices of the PICOD hypergraph. We generalize the relationships between various bounds for the optimal PICOD length arising from conflict-free colorings to such k-fold colorings. The questions of existence and constructions of non-trivial k-fold colorings of specific lengths for general PICOD problems are not considered in this work, and remain open for future research.

Fig. 1 represents in a pictorial form some of the main results in the paper for the t-request PICOD scenario, for any $t \geq 1$, a section of which is applicable to the k-vector PICOD schemes also. Apart from the contributions listed above, we also show a number of minor results which bring out the relationships between the various studied in this work.

Notation: Let $[n] \triangleq \{1, \ldots, n\}$ for a positive integer n. For sets A, B we denote by $A \backslash B$ the set of elements in A but not in B. We abuse notation to denote $A \backslash \{b\}$ as $A \backslash b$. The set of k-sized

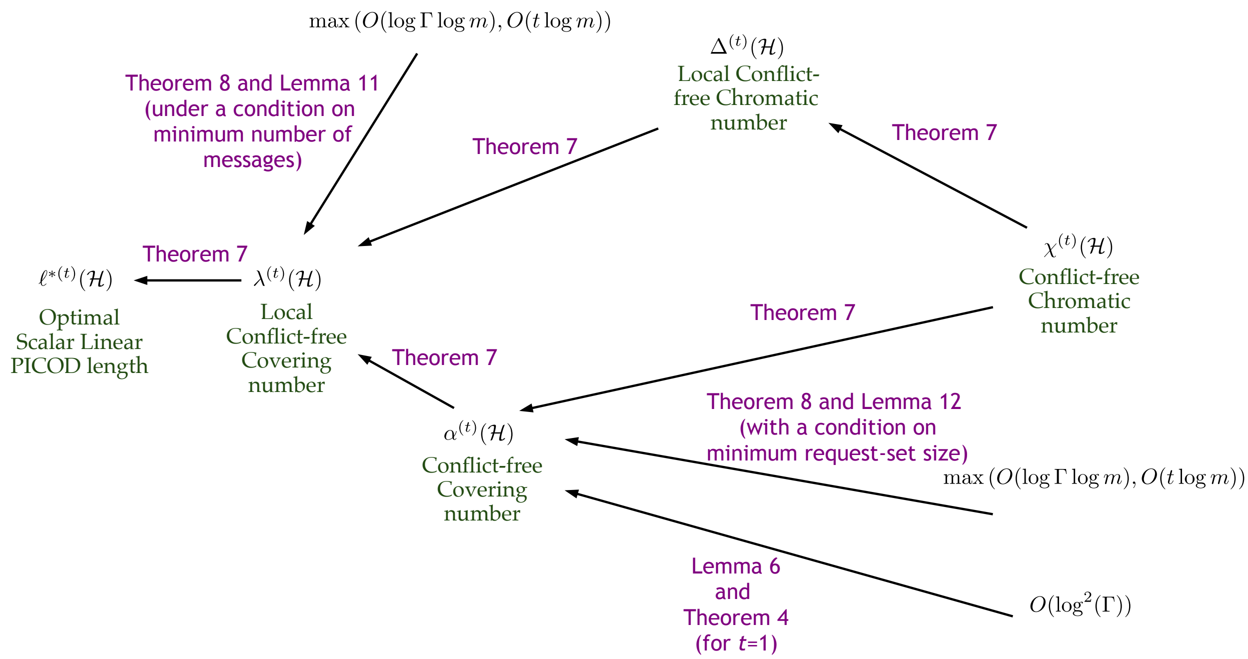

Figure 1: A map of the main quantities and results identified in this work for the t-request PICOD problem, for any $t \geq 1$. Each $\leftarrow$ arrow indicates an upper bound (a $\leq$ relationship). Note that Lemma 6 and Theorem 4 hold for $t = 1$ only, as indicated. Further, in Lemma 13, we show a class of hypergraphs for which the bound $\max(O(\log\Gamma\log m), O(t\log m))$ for $\ell^{*(t)}(\mathcal{H})$ is asymptotically tight, upto a multiplicative factor $\log t$.

subsets of any set A is denoted by $\binom{A}{k}$. The span of a set of vectors U is denoted by $\mathsf{span}(U)$. The dimension of a subspace W is denoted by $\mathsf{dim}(W)$. The base of the natural logarithm is denoted by e. Unless mentioned explicitly, all logarithms in the paper are to the base e. The empty is denoted by $\emptyset$. We denote the union of disjoint sets A and B as $A \uplus B$. A hypergraph $\mathcal{H}$ is a pair of sets $(V, \mathcal{E})$ where the set V is called the set of vertices of $\mathcal{H}$ (also denoted by $V(\mathcal{H})$) and $\mathcal{E}$ is a collection of subsets of V called the set of hyperedges (sometimes referred to as simply edges) of $\mathcal{H}$ (also denoted by $\mathcal{E}(\mathcal{H})$). Given a collection of hypergraphs $\mathcal{H}_p : p \in [P]$, we define their union as $\mathcal{H} = \cup_{p\in[P]}\mathcal{H}_p$ where $V(\mathcal{H}) = \cup_{p\in[P]}V(\mathcal{H}_p)$, and $\mathcal{E}(\mathcal{H}) = \cup_{p\in[P]}\mathcal{E}(\mathcal{H}_p)$. For asymptotics, we follow the standard Bachmann–Landau notation.

II. Pliable Index Coding Problem

We briefly review the pliable index coding problem, introduced in [7]. Consider a communication setup defined as follows. There are m messages denoted by $\{x_i : i \in [m]\}$ where x_i lies in some finite alphabet $\mathcal{A}$ (throughout, we impose a finite field structure on $\mathcal{A}$). These m messages are available at a server. Consider n receivers indexed by $[n]$. Assume that there is a noise-free broadcast channel between the server and the receivers. Each receiver r has some subset of messages available apriori, as *side-information*. The set of indices of the symbols available at client r is denoted S_r, and the indices of those that are not available are denoted using the set $I_r = [m] \setminus S_r$. We call $\{x_i : i \in I_r\}$ as the *request-set* of receiver r. The demand at receiver r is fulfilled if it receives any symbol in its request-set from the server. The messages indexed by $[m]$, the receivers indexed by $[n]$, and the request-sets $\mathfrak{I} \triangleq \{I_r : r \in [n]\}$ together define a $(n, m, \mathfrak{I})$-*pliable index coding problem (PICOD problem).* We assume that $|I_r| \geq 1, \forall r$, as any receiver with $|I_r| = 0$ can be removed from the problem description as it has all the symbols. Consider a hypergraph $\mathcal{H}$ with vertex set $V = [m]$ and edge set $\mathfrak{I} = \{I_r : r \in [n]\}$. This hypergraph represents the PICOD problem.

A *pliable index code* (PIC) consists of a collection of (a) an encoding function at the server which encodes the m messages to an ℓ-length codeword, denoted by $\phi : \mathcal{A}^m \rightarrow \mathcal{A}^\ell$ and (b) decoding functions $\{\psi_r : r \in [n]\}$ where $\psi_r : \mathcal{A}^\ell \times \mathcal{A}^{|S_r|} \rightarrow \mathcal{A}$ denotes the decoding function at client r such that

$$\psi_r \left(\phi(\{x_i : i \in [m]\}), \{x_i : i \in S_r\}\right) = x_d, \text{ for some } d \in I_r.$$

The quantity ℓ is called the *length* of the PIC. It is of interest to design pliable index codes of small length.

In this work, we assume $\mathcal{A} = \mathbb{F}^k$ for some finite field $\mathbb{F}$ and integer $k \geq 1$. We refer to these codes as k-*vector PICs*, while the $k = 1$ case is also called *scalar PIC*. We focus on linear PICs, i.e., those in which the encoding and decoding functions are linear. In that case, the encoder ϕ is represented by a $\ell \times mk$ matrix (denoted by G) such that $\phi(\{x_i : i \in [m]\}) = G\boldsymbol{x}^T$, where $\boldsymbol{x} = (x_{1,1}, \ldots, x_{1,k}, \ldots, x_{m,1}, \ldots, x_{m,k})$, in which $(x_{i,1}, \ldots, x_{i,k}) \in \mathbb{F}^k$ denotes the message x_i. For the PICOD problem given by hypergraph $\mathcal{H}$, the smallest ℓ for which there is a linear k-vector PIC (over some field $\mathbb{F}$) is denoted by $\ell_k^*(\mathcal{H})$.

The following definition and lemma (proved in [8]) describe when encoding $\boldsymbol{x}$ using the matrix G can lead to correct decoding at the receivers.

Definition 1. *For an $(n, m, \mathfrak{I})$-PICOD problem, a matrix G with mk columns indexed as $G_{i,j}$: $i \in [m], j \in [k]$, is said to satisfy receiver $r \in [n]$ with k symbols, if the following property (P) is satisfied by G.*

(P) *There exists some $d \in I_r$ such that* $\mathsf{dim}(\mathsf{span}(\{G_{d,j} : j \in [k]\})) = k$ *and*

$$\mathsf{span}(\{G_{d,j} : j \in [k]\}) \cap \mathsf{span}\left(\{G_{i,j} : \ \forall i \in I_r \backslash d, \ j \in [k]\}\right) = \{\mathbf{0}\}.$$

Lemma 1 (Lemmas 1 and 6 of [8])**.** *A matrix G with mk columns is the encoder of a k-vector PIC for an $(n, m, \mathfrak{I})$-PICOD problem if and only if the property (P) of Definition 1 is true for each receiver $r \in [n]$.*

Lemma 2 below is useful to prove achievability results for PICOD problems in this work.

Lemma 2. *For an $(n, m, \mathfrak{I})$-PICOD problem, let $\{G^p : p \in [P]\}$ denote a collection of matrices, where G^p is of size $L_p \times mk$, such that for each $r \in [n]$, there exists some matrix G^p which satisfies receiver r with k symbols. Then the matrix $G = \begin{bmatrix} G^1 \\ \vdots \\ G^P \end{bmatrix}$ of size $(\sum_{p \in [P]} L_p) \times mk$ is the encoder of a k-vector PIC for the given PICOD problem.*

Proof: For each $r \in [n]$, there exists some matrix G^p such that Property (P) holds for r (with respect to some $d \in I_r$). By simple linear algebraic arguments, we see that the matrix G too must satisfy property (P) for receiver r (with respect to $d \in I_r$), and hence satisfies r. Applying Lemma 1, the proof is complete. ■

So far, we have formally presented the model for the single-request PICOD problem, in which a single message in the request-set is requested by each client. In Sections III and IV, we present our results on scalar index codes for such single-request PICOD problems. In Section V, we focus on the *t-request PICOD problem*, in which each receiver r demands $\min(t, |I_r|)$ messages in its request-set, for some $t \geq 1$. We capture the feasibility of scalar $(k = 1)$ PICs for the t-request problem by the following lemma, following similar ideas as in Definition 1 and Lemma 1. The proof of the below lemma is similar to Lemma 1 and is omitted.

Lemma 3. *For a t-request $(n, m, \mathfrak{I})$-PICOD problem, a matrix G with m columns indexed as $G_i : i \in [m]$, is the encoder of a valid scalar $(k = 1)$ PIC, if and only if the following property (P') is satisfied by G for each receiver $r \in [n]$:*

(P') *There exists some* $D \subseteq I_r$ *such that* $\mathsf{dim}(\mathsf{span}(\{G_d : d \in D\})) = |D| = \min(t, |I_r|)$ *and*

$$\mathsf{span}(\{G_d : d \in D\}) \cap \mathsf{span}\left(\{G_i : \ \forall i \in I_r \backslash D\}\right) = \{\mathbf{0}\}.$$

It is then easy to show a result similar to Lemma 2, with the property (P') as in Lemma 3 (instead of property (P)) that give the structure of scalar PIC schemes constructed in this work for t-request PICOD problems. We omit the formal statement and proof in the interest of avoiding repetition. In Section VI, we extend some of our scalar PIC results to the k-vector PIC schemes.

III. Scalar PICs arising from Conflict-free Colorings

Firstly, we review the definition of conflict-free colorings of a hypergraph $\mathcal{H}$ and the associated chromatic number $\chi_{CF}(\mathcal{H})$ and discuss some existing results. In Subsection III-A, we show that a conflict-free coloring of the hypergraph $\mathcal{H}(V = [m], \mathfrak{I})$ leads to a scalar linear PIC scheme for the PICOD problem given by $\mathcal{H}$, with length equal to the number of colors. We then define in Subsection III-B the notion of conflict-free collection of colorings of $\mathcal{H}$. The smallest number of colors that is used in any such collection is called conflict-free covering number and denoted by $\alpha_{CF}(\mathcal{H})$. We show that such a collection also leads to achievable PIC schemes, yielding tighter upper bounds on $\ell_1^*(\mathcal{H})$ than $\chi_{CF}(\mathcal{H})$, in general. We also show a 'sandwich' property for $\alpha_{CF}(\mathcal{H})$ based on functions of $\chi_{CF}(\mathcal{H})$. Using the framework established in Subsection III-B, in Subsection III-C, we show that a PICOD problem $\mathcal{H}$ has a PIC scheme with length $O(\log^2 \Gamma)$, where Γ is an edge-overlap parameter associated with $\mathcal{H}$. This new upper bound is better than the $O(\log^2 n)$ bound shown in [8] for specific parameter ranges, when Γ is small and n is large. Following reasons provided in Section I, this can happen, for instance, in a single-server distributed data shuffling scenario across a large number of helper nodes, where the data storage across distinct clients are fairly distinct.

Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph. Let $C : V \rightarrow [L]$ be a coloring of V, where L is a positive integer. Consider a hyperedge $E \in \mathcal{E}$. We say C is a *conflict-free coloring for the hyperedge* E if there is a vertex $v \in E$ such that $C(v) \neq C(u), \ \forall u \in E \setminus \{v\}$. That is, in such a coloring, E contains a vertex whose color is distinct from that of every other vertex in E (we also use the phrasing that E 'sees' a color exactly once). We say C is a *conflict-free coloring of the hypergraph* $\mathcal{H}$ if C is a conflict-free coloring for every $E \in \mathcal{E}$. The *conflict-free chromatic*

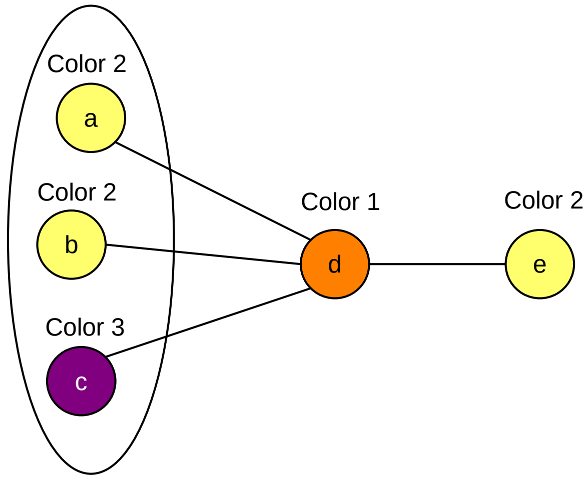

Figure 2: A hypergraph $\mathcal{H}$ with five nodes $\mathcal{V} = \{a, b, c, d, e\}$ and five edges $\mathcal{E} = \{\{d, a\}, \{d, b\}, \{d, c\}, \{d, e\}, \{a, b, c\}\}$ is shown. A conflict-free coloring is shown for $\mathcal{H}$. In the edge $\{a, b, c\}$, the vertex c has the distinct color. Each hyperedge 'sees' two colors locally.

number of $\mathcal{H}$, denoted by $\chi_{CF}(\mathcal{H})$, is the minimum L such that there is a conflict-free coloring $C : V \to [L]$ of $\mathcal{H}$.

Example 1. *Fig. 2 shows an example of a conflict-free coloring of a hypergraph $\mathcal{H}$, using three colors. Further, it is easy to see that any conflict-free coloring of this hypergraph requires at least three colors. Thus, $\chi_{CF}(\mathcal{H}) = 3$, for this example.* ■

The following theorem on conflict-free coloring on hypergraphs is due to Pach and Tardos [19], which we shall use to obtain one of our main results (Lemma 6 and Theorem 4) in Subsection III-C.

Theorem 1 (Theorem 1.2 in [19])**.** *For any positive integers t and Γ, the conflict-free chromatic number of any hypergraph in which each edge is of size at least $2t-1$ and each edge intersects with at most Γ others is $O(t\Gamma^{1/t} \log \Gamma)$. Further, there is a randomized polynomial time algorithm to find such a coloring.*

A. Relationship of PIC to Conflict-free Coloring

In this subsection, we show that a conflict-free coloring of the hypergraph $\mathcal{H}(V = [m], \mathfrak{I})$ gives a scalar PIC scheme for the PICOD problem given by $\mathcal{H}$. To do this, we define the following matrix associated with a conflict-free coloring of $\mathcal{H}$.

Definition 2 (Indicator Matrix associated with a coloring)**.** *Let $C : V \to [L]$ denote a coloring of $\mathcal{H}(V = [m], \mathfrak{I})$, where $C(i)$ denotes the color assigned to the vertex $i \in [m]$. Consider a standard basis of the L-dimensional vector space over $\mathbb{F}$, denoted by $\{e_1, \ldots, e_L\}$. Now consider the $L \times m$ matrix G (with columns indexed as $\{G_i : i \in [m]\}$) constructed as follows.*

- *For each $i \in [m]$, column G_i of G is fixed to be $e_{C(i)}$.*

We call G as the indicator matrix associated with the coloring C.

Using the indicator matrix associated with a conflict-free coloring of $\mathcal{H}$, we shall prove our first bound on $\ell_1^*(\mathcal{H})$.

Theorem 2. $\ell_1^*(\mathcal{H}) \leq \chi_{CF}(\mathcal{H})$.

Proof: Let $C : V \to [L]$ denote a conflict-free coloring of $\mathcal{H}$. We first show that there exists an L-length scalar linear PIC for the problem defined by $\mathcal{H}$. Let G denote the indicator matrix associated with the coloring C as defined in Definition 2. We show that G satisfies Lemma 1 and hence is a valid encoder for a linear PIC.

In any conflict-free coloring of $\mathcal{H}$, every edge I_r of $\mathcal{H}$ has a vertex d such that $C(d) \neq C(i), \forall i \in I_r \setminus d$. Then, clearly, $e_{C(d)} \neq e_{C(i)}$, for any $i \in I_r \backslash d$. This also means $\mathsf{span}(\{e_{C(d)}\}) \cap \mathsf{span}(\{e_{C(i)} : i \in I_r \backslash d\}) = \{\mathbf{0}\}$, as the vectors $\{e_1, \ldots, e_L\}$ are basis vectors. Further, $e_{C(d)}$ spans a one dimensional space. Thus, G satisfies every receiver r and is a valid encoder by Lemma 1. Note that the length of the code is exactly L. By definition of $\chi_{CF}(\mathcal{H})$, the proof is complete. ■

Example 2. *Consider the PICOD problem represented by the hypergraph $\mathcal{H}$ with vertex set $V = \{1, \ldots, 8\}$ and edge set*

$$\mathcal{E} = \{\{1,2,4,6\}, \{1,2,3,5\}, \{2,3,4,7\}, \{1,3,4,8\},$$
$$\{2,5,6,7\}, \{1,5,6,8\}, \{3,5,7,8\}, \{4,6,7,8\}\}.$$

Consider a coloring C which assigns color 1 to vertices $\{1,2,3,4\}$ and color 2 to vertices $\{5,6,7,8\}$. Note that this is a valid conflict-free coloring of $\mathcal{H}$. The indicator matrix associated with this coloring is given by

$$G = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

It can be checked that the above matrix satisfies the condition in Lemma 1 for the PICOD problem defined by $\mathcal{H}$. Thus, G is a valid PIC encoding matrix for $\mathcal{H}$. ■

B. Conflict-free coverings and PICOD

In the following discussion, we define a new parameter called the conflict-free covering number. Using this, we obtain a tighter bound (Theorem 3) on the optimal PIC length as compared to Theorem 2.

Definition 3 (*Conflict-free collection, conflict-free covering number*)**.** *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph. Let $\mathfrak{C} = \{C^1, \ldots, C^P\}$ where each $C^p : V \to [L_p]$ be colorings of the hypergraph $\mathcal{H}$. We say $\mathfrak{C}$ is a* conflict-free collection of colorings of $\mathcal{H}$, *if the following condition holds: For every $E \in \mathcal{E}$, there is $p \in [P]$ such that E sees some color exactly once under the coloring C^p. The quantity*

$$\alpha_{CF}(\mathcal{H}) \triangleq \min_{\mathfrak{C}} \sum_{p=1}^{P} L_p,$$

representing the minimum sum $\sum_{p=1}^{P} L_p$ over all possible collections $\mathfrak{C}$ (over all P) as defined above, is called the conflict-free covering number of $\mathcal{H}$.

Example 3. *Fig. 3 shows the hypergraph K_7 which is the complete graph on the vertex set $\{1,2,3,4,5,6,7\}$, containing all its 2-sized subsets as hyperedges. It requires 7 colors for a conflict-free coloring, thus $\chi_{CF}(K_7) = 7$. The three figures (b), (c) and (d) depict a collection of conflict-free colorings, with each figure corresponding to one coloring using 2 colors. It can be checked that each edge of K_7 is conflict-free in at least one of these colorings, and the total number of colors used in all the three colorings is 6. Thus $\alpha_{CF}(\mathcal{H}) \leq 6 < \chi_{CF}(\mathcal{H})$.* ■

In the following, we show that the parameter $\alpha_{CF}(\mathcal{H})$ is sandwiched between functions of $\chi_{CF}(\mathcal{H})$.

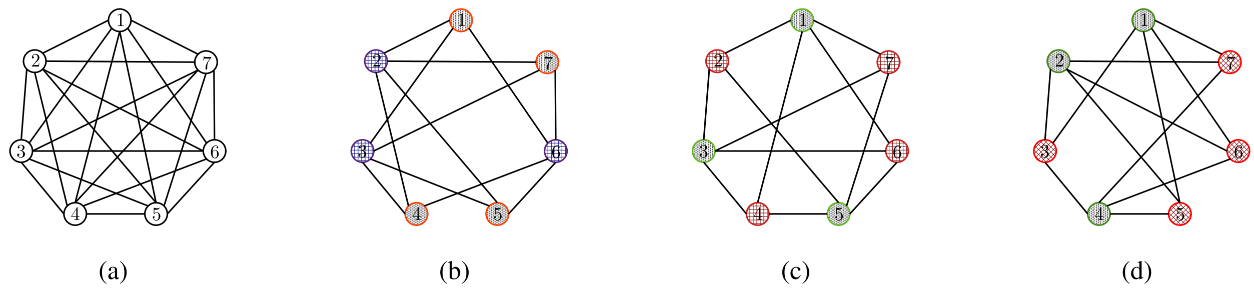

Figure 3: Figure (a) shows the hypergraph K_7 which is the complete graph on 7 vertices, $\{1, 2, 3, 4, 5, 6, 7\}$ containing all the 2-sized subsets as hyperedges. The three figures (b), (c) and (d) depict a collection of conflict-free colorings, with each figure corresponding to one coloring using 2 colors. Note that only those edges satisfied by the coloring are represented in (b), (c), and (d). In (b), the two color classes (a color class is a subset of vertices with the same color) are $\{1, 4, 5, 7\}$ and $\{2, 3, 6\}$. In (c), they are $\{2, 4, 6, 7\}$ and $\{1, 3, 5\}$ and in (d) they are $\{1, 2, 4\}$ and $\{3, 5, 6, 7\}$.

Lemma 4. *Let $\mathcal{H}$ be a hypergraph with $\chi_{CF}(\mathcal{H}) = \chi$ and let r be the smallest integer such that there exists a conflict-free coloring of $\mathcal{H}$ using χ colors in which the vertices in any hyperedge are colored with at most r colors. If $r = 1$, then $\chi = \alpha_{CF}(\mathcal{H}) = 1$ and if $r \geq 2$, then $\log_2(\chi) \leq \alpha_{CF}(\mathcal{H}) \leq \min(\chi, \frac{r^{r+2}}{r!} \log(\chi))$.*

Proof: If $r = 1$, it is easy to see that the lemma holds. So we assume $r \geq 2$ throughout.

Consider a conflict-free collection of $\mathcal{H}$ with P colorings with $\alpha_{CF}(\mathcal{H})$ total colors. Let $V = [m]$ be the set of vertices of $\mathcal{H}$. Consider the $\alpha_{CF}(\mathcal{H}) \times |V|$ matrix $G = \begin{bmatrix} G^1 \\ \vdots \\ G^P \end{bmatrix}$, where G^p represents the indicator matrix of the p^{th} coloring in the collection. Let $\{g_1, \ldots, g_L\}$ denote the set of all distinct columns of G. Thus, we must have that $\alpha_{CF}(\mathcal{H}) \geq \log_2(L)$. Now, consider a coloring C of the vertices of $\mathcal{H}$ with elements of $[L]$, where a vertex $i \in V$ gets label $\ell \in [L]$ if the i^{th} column of G is g_ℓ. By construction of G, C is a conflict-free coloring of $\mathcal{H}$ and thus $L \geq \chi$. Thus we see that $\alpha_{CF}(\mathcal{H}) \geq \log_2(\chi)$.

The upper bound $\alpha_{CF}(\mathcal{H}) \leq \chi$ follows, as any conflict-free coloring of $\mathcal{H}$ also gives a conflict-free collection containing only the same coloring. We now show the other upper bound. Let $C : V \to [\chi]$ be any conflict-free coloring with χ colors so that vertices in any hyperedge are colored with at most r colors. Firstly, for $P = \frac{r^{r+1} \log \chi}{r!}$, we show that there exists a $P \times \chi$ matrix with entries from $[r]$ such that any r columns of this matrix contain at least one row with

r distinct entries. Using this matrix, we construct a conflict-free collection of colorings with $\frac{r^{r+2}}{r!}\log(\chi)$ colors, which will complete the proof.

We now show the existence of this desired matrix. Let each entry of a random $P \times \chi$ matrix M be i.i.d and drawn uniformly at random from $[r]$. The probability that any particular r-subset $R \subset \chi$ of columns contains r distinct colors in the row i is given by $\frac{r!}{r^r}$. Thus, the probability q that at least one r-subset of columns of M does not contain distinct entries in any row, is given by

$$q \leq \binom{\chi}{r}\left(1-\frac{r!}{r^r}\right)^P \overset{(a)}{<} \chi^r \mathsf{e}^{-P\frac{r!}{r^r}} \leq 1,$$

where (a) follows as $\binom{\chi}{r} < \chi^r$ for $r \geq 2$, and since $1+x \leq \mathsf{e}^x, \forall x$. This means that there is at least one matrix M (say M_{cf}) of size $P \times \chi$ with entries from $[r]$, such that each r-subset of columns contains distinct entries in some row.

Now, using M_{cf}, we obtain a collection of P colorings of $\mathcal{H}$ in the following way. With respect to the p^{th} row of M_{cf}, we define a coloring $C^p : V \rightarrow [r]$ such that for each $i \in V$, $C^p(i)$ is equal to the $(p,j)^{th}$ entry of M_{cf}, where j is the color assigned to the vertex i in C. By the property of M_{cf} and the coloring C chosen, it can be verified that this collection $\{C^p : p \in [P]\}$ will be a conflict-free collection of colorings of χ. The collection uses r colors per coloring and thus totally there are rP colors being used. This completes the proof. ■

The following corollary to Lemma 4 shows that the conflict-free covering number can be strictly smaller than the conflict-free chromatic number in general. We consider the class of uniform hypergraphs (in which the size of each edge is the same) in this corollary, the proof of which follows easily from Lemma 4.

Corollary 1. *For constant r, consider the class of r-uniform hypergraphs on m vertices, $m \geq r$. For any hypergraph $\mathcal{H}$ in this class, $\alpha_{CF}(\mathcal{H}) = \Theta_r(\log_2 \chi_{CF}(\mathcal{H}))$.*

The following lemma shows a class of 2-uniform hypergraphs for which the separation between parameters $\alpha_{CF}(\mathcal{H})$ and $\chi_{CF}(\mathcal{H})$ can be quite large. Indeed, Fig. 3 gives an example hypergraph from this class and illustrates the lemma. We observe, however, that the minimum of the two upper bounds for $\alpha_{CF}(\mathcal{H})$ as in Lemma 4 is still (asymptotically) tight for this class of hypergraphs.

Lemma 5. *There exist a hypergraph $\mathcal{H}$ with n hyperedges for which $\alpha_{CF}(\mathcal{H}) = \Theta(\log_2 n)$ while $\chi_{CF}(\mathcal{H}) = \Theta(\sqrt{n})$.*

Proof: Consider the 2-uniform hypergraph with m vertices and all the 2-sized subsets of $[m]$ as hyperedges. Thus $n = \binom{m}{2}$. It is easy to see that any conflict-free coloring of this graph requires $m = \Theta(\sqrt{n})$ colors. Observe that in any conflict-free coloring of this hypergraph, each edge sees two colors. Using this fact in Lemma 4 completes the proof. ■

We now show that the optimal length of PIC for $\mathcal{H}$ is bounded by $\alpha_{CF}(\mathcal{H})$, thus improving the bound in Theorem 2.

Theorem 3. $\ell_1^*(\mathcal{H}) \leq \alpha_{CF}(\mathcal{H})$.

Proof: Let $\mathfrak{C} = \{C^p : p \in [P]\}$ be a conflict-free collection of colorings of $\mathcal{H}(V = [m], \mathfrak{I})$, where $C^p : V \to [L_p]$. We first show a PIC for $\mathcal{H}$ with length $\sum_{p\in[P]} L_p$. The proof then follows by definition of $\alpha_{CF}(\mathcal{H})$.

Let $G^p : p \in [P]$ denote the indicator matrices as defined in Definition 2 associated with the colorings $C^p : p \in [P]$ respectively. By definition of the conflict-free collection, for each $I_r \in \mathfrak{I}$ we have by arguments similar to the proof of Theorem 2, that there is some G^p which satisfies receiver r. Then, by Lemma 2, the matrix $G = \begin{bmatrix} G^1 \\ \vdots \\ G^P \end{bmatrix}$ is a valid encoder of a PIC of length $\sum_{p\in[P]} L_p$ for the given PICOD problem. ■

C. *Showing that $\alpha_{CF}(\mathcal{H}) = O(\log^2 \Gamma)$*

In the remainder of this section, our main goal is to show a new upper bound (Theorem 4) on $\alpha_{CF}(\mathcal{H})$ (and thus on the optimal linear scalar PIC length) based on a readily computable parameter associated with the hypergraph $\mathcal{H}$, and a randomized construction of such a PIC, which runs in polynomial time in the parameters of the problem, with high probability. Towards that end, we first show a new upper bound on $\alpha_{CF}(\mathcal{H})$. The following observation is needed to show our new upper bound.

Observation 1. *Let $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$. Then, $\alpha_{CF}(\mathcal{H}) \leq \alpha_{CF}(\mathcal{H}_1) + \alpha_{CF}(\mathcal{H}_2)$.*

We now show a new upper bound on $\alpha_{CF}(\mathcal{H})$.

Lemma 6. *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph where every hyperedge intersects with at most Γ other hyperedges. Then, $\alpha_{CF}(\mathcal{H}) = O(\log^2 \Gamma)$.*

Proof: Firstly, we consider the case when $\Gamma < \mathrm{e}$. Since Γ is an integer, it means that $\Gamma \leq 2$. Further, by definition, the quantity Γ bounds the maximum degree of any vertex (maximum number of edges any vertex is present in). From a result in [20] (Theorem 1 and Remark 1 in [20]), it then follows that $\alpha_{CF}(\mathcal{H}) \leq \chi_{CF}(\mathcal{H}) \leq \Gamma \leq 2$.

Now, we consider the case when $\Gamma > \mathrm{e}$. Let $\kappa := 2 \log \Gamma - 1$. Let $\mathcal{G} = (V, \mathcal{E}_G)$ be a hypergraph defined on the vertex set V with $\mathcal{E}_G = \{E \in \mathcal{E} \ : \ |E| \geq \kappa\}$. Applying Theorem 1 (with $t = \log \Gamma$), we have that $\chi_{CF}(\mathcal{G}) = O(\log^2 \Gamma)$. Thus, using Lemma 4 (which shows $\alpha_{CF}(\mathcal{G}) \leq \chi_{CF}(\mathcal{G})$), we have that $\alpha_{CF}(\mathcal{G}) = O(\log^2 \Gamma)$. Let $P := \lceil \log_2 \kappa \rceil$. For $0 \leq i \leq P$, let $\mathcal{H}_i = (V, \mathcal{E}_i)$, where $\mathcal{E}_i = \{E \in \mathcal{E} \ : \ \frac{k_i}{2} \leq |E| < k_i\}$ and $k_i = \frac{\kappa}{2^i}$. Clearly, $\mathcal{H} = \mathcal{G} \cup \mathcal{H}_0 \cup \mathcal{H}_1 \cup \cdots \cup \mathcal{H}_P$. Using Observation 1 repeatedly, we thus have that $\alpha_{CF}(\mathcal{H}) \leq \alpha_{CF}(\mathcal{G}) + \sum_{i=0}^{P} \alpha_{CF}(\mathcal{H}_i)$. Thus, to complete the proof, it is enough to show $\sum_{i=0}^{P} \alpha_{CF}(\mathcal{H}_i) = O(\log^2 \Gamma)$, which is what we do in the rest of the proof.

We first show that $\alpha_{CF}(\mathcal{H}_i) \leq 2(\lceil 5 k_i \log \Gamma \rceil)$, using the Lovász Local Lemma [21] (which we recall as Lemma 14 in Appendix A). Let $q_i = \frac{1}{k_i}$ and $t_i = \lceil 5 k_i \log \Gamma \rceil$. We do t_i rounds of coloring of the vertex set V of $\mathcal{H}_i$, using two new colors in each round. In any given round, we color each vertex $v \in V$ independently with probability q_i with the first color, and give it the second color with the remaining probability, i.e., $1 - q_i$. Consider some edge $E \in \mathcal{E}_i$. Let F_E denote the 'bad' event that none of the t_i colorings of V is a conflict-free coloring for E. The abovestated multi-round probabilistic coloring scheme will enable us to bound the probability of the events $F_E : E \in \mathcal{E}_i$, which will then allow us to apply the Local Lemma, thus proving the existence of a desired conflict-free collection of colorings for $\mathcal{H}_i$. We now provide the details. A conflict-free assignment of colors for E is surely realized by the coloring in a given round, if for any vertex $v \in E$, the vertex v gets the first color and all vertices in $E \setminus v$ get the second color, in that round. Thus, the probability that the coloring in a given round is a conflict-free coloring for E is at least $|E|q_i(1-q_i)^{|E|-1}$. Thus,

$$\begin{aligned}
\mathsf{Pr}[F_E] &\leq \left(1-|E|q_i(1-q_i)^{|E|-1}\right)^{t_i} \\
&\overset{(a)}{\leq} \frac{1}{\mathsf{e}^{t_i|E|q_i(1-q_i)^{|E|-1}}} \\
&\overset{(b)}{\leq} \frac{1}{\mathsf{e}^{t_i\frac{k_i}{2}\frac{1}{k_i}(1-\frac{1}{k_i})^{k_i-1}}} \\
&\overset{(c)}{\leq} \frac{1}{\mathsf{e}^{\frac{t_i}{2k_i}}} \\
&\leq \frac{1}{\Gamma^{2.5}} \quad (\text{since } t_i \geq 5k_i\log\Gamma),
\end{aligned}$$

where (a) holds by inequality $1+x \leq \mathsf{e}^x$, (b) holds as $q_i = \frac{1}{k_i}$ and $\frac{k_i}{2} \leq |E| < k_i$, (c) holds using the inequality $(1+x)^r \geq 1+rx$ for $x \geq -1, r \in \mathbb{R}\setminus(0,1)$. Thus, for each hyperedge E in $\mathcal{H}_i$, the probability of the bad event F_E is at most $\frac{1}{\Gamma^{2.5}}$. Observe that each such event F_E is independent of all the other events, but at most Γ events corresponding to those hyperedges intersecting with E. Now, we can invoke the Local Lemma (Lemma 14) with the following assignments:

- The events ($A_i : i \in [n]$, as in the Local Lemma (Lemma 14)) being $F_E : E \in \mathcal{E}_i$,
- $p := \frac{1}{\Gamma^{2.5}}$,
- $d := \Gamma$.

Observe that these values of p and d satisfy the condition $\mathsf{e}p(d+1) \leq 1$, as is desired in the Local Lemma. Thus, using the Local Lemma, we have that $\mathsf{Pr}[\cap_{E\in\mathcal{E}_i}\overline{F}_E] > 0$. Therefore, there exists a conflict-free collection containing t_i colorings, for $\mathcal{H}_i$, such that each coloring in the collection involves a distinct set of two colors. This proves that $\alpha_{CF}(\mathcal{H}_i) \leq 2t_i = 2(\lceil 5k_i\log\Gamma\rceil)$.

Using this fact, we have

$$\begin{aligned}
\sum_{i=0}^{P}\alpha_{CF}(\mathcal{H}_i) &\leq 2\sum_{i=0}^{P}\lceil 5k_i\log\Gamma\rceil \\
&\leq 2\sum_{i=0}^{P}(5k_i\log\Gamma+1) \leq 10\log\Gamma\left(\sum_{i=0}^{P}k_i\right)+2P+2 \\
&\leq 10\log\Gamma\left(\sum_{i\geq 0}\frac{\kappa}{2^i}\right)+2P+2 \\
&\leq 20\kappa\log\Gamma+2P+2 = O(\log^2\Gamma),
\end{aligned}$$

as $\kappa = O(\log\Gamma)$ and $P = O(\log(\log\Gamma))$. This completes the proof. ∎

D. *Constructing a PIC of length* $O(\log^2 \Gamma)$

Our result in Theorem 3 in conjunction with Lemma 6 implies the existence of a PIC scheme with length $O(\log^2 \Gamma)$. In this subsection, we show a randomized construction for a covering of the PICOD hypergraph with $O(\log^2 \Gamma)$ colors (and thus constructing a PIC scheme with length $O(\log^2 \Gamma)$). This construction runs in time that is polynomial in the parameters of the problem with high probability. Theorem 4 summarizes these observations and is the main result in this section. At the end of this subsection, we briefly remark that, using techniques available in the literature, our randomized algorithm can be derandomized, thus leading to a deterministic polynomial time algorithm for the PIC scheme with length $O(\log^2 \Gamma)$.

We now present randomized algorithms **CF-Coloring**($\mathcal{G} = (V, \mathcal{E}_G)$) and **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) which construct colorings for the subgraphs (in the proof of Lemma 6), $\mathcal{G}$ and $\mathcal{H}_i : i \in \{0, \ldots, \lceil \log \kappa \rceil\}$ (where $\kappa = (2 \log \Gamma - 1)$), respectively. We use the algorithmic version of the Local Lemma due to Moser and Tardos [22] (Theorem 12, recalled in Appendix A) to develop our randomized algorithms.

The first algorithm **CF-Coloring**($\mathcal{G} = (V, \mathcal{E}_G)$) gives a conflict-free coloring of the subgraph $\mathcal{G}$ (in the proof of Lemma 6), that uses $O(\log^2 \Gamma)$ colors. The second algorithm **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) returns a collection of conflict-free colorings for the subgraph $\mathcal{H}_i$ (in the proof of Lemma 6) that uses at most $2\left(\lceil 5k_i \log \Gamma \rceil\right)$ colors. Further, we show that these run in time that is polynomial in the system parameters, with high probability. Combining these, we show that a PIC scheme of length $O(\log^2 \Gamma)$ can be constructed in polynomial time with high probability.

CF-Coloring($\mathcal{G} = (V, \mathcal{E}_G)$)

Input: A hypergraph $\mathcal{G} = (V, \mathcal{E}_G)$, with $|V| = m$, such that (i) every hyperedge $E \in \mathcal{E}_G$ intersects with at most Γ other hyperedges $E' \in \mathcal{E}_G$, and (ii) for every hyperedge $E \in \mathcal{E}_G$, we have $|E| \geq \kappa$ where $\kappa = 2 \log \Gamma - 1$.

Output: A conflict-free coloring $c : V \to \{1, 2, \ldots, r\}$, where $r = O(\log^2 \Gamma)$.

Algorithm:

- Select κ vertices arbitrarily from each $E \in \mathcal{E}_G$, and discard the vertices that are not selected for any edge. Let $\mathcal{G}' = (V', \mathcal{E}'_G)$ be the resulting hypergraph induced by the

This algorithm is an algorithmic version of Theorem 1 (originally stated as Theorem 1.2 in [19], which the reader can refer to for a proof).

remaining vertices.

- For each $v \in V'$, do the following:
  - **Geometric assignment of colors:** Repeatedly toss a coin with Heads probability $p = 1/(60 \log \Gamma)$, till we get a Head. The color $c(v)$ is assigned to be the trial number in which the first Head was obtained.
- While $\exists E \in \mathcal{E}'_G$ such that the above coloring is not a conflict-free coloring of E
  - Choose an arbitrary $E \in \mathcal{E}'_G$ such that the above coloring is not a conflict-free coloring of E.
  - For each of the vertices $v \in E$, we reassign colors $c(v)$, using the above Geometric assignment of colors.
- For vertices $v \in V \setminus V'$, we assign $c(v)$ to be a new unused color.
- Output the coloring c.

Remark 1. *We now show that **CF-Coloring** runs in time that is polynomial in the problem parameters, with high probability.*

1) *Let us first consider the time taken for the geometric assignment of colors. This requires $1/p = O(\log \Gamma)$ trials in expectation, for each vertex. Therefore, it requires $O(m \log \Gamma)$ trials (in expectation) to color all the vertices in V'. Applying Markov's inequality, the probability of exceeding $O(mn(m+n)^2 \log \Gamma/\Gamma)$ trials is at most $\Gamma/((m+n)^2 n) < 1/(m+n)$ as $\Gamma \leq n$.*
2) *Checking the condition of the while loop, that is, if each hyperedge $E \in \mathcal{E}'_G$ is a conflict-free coloring, requires $O(nm)$ time. This is followed by a geometric reassignment of colors, which, as explained above takes $O(mn(m+n)^2 \log \Gamma/\Gamma)$ time. Thus, each run of the while loop takes at most $\max\{O(mn), O(mn(m+n)^2 \log \Gamma/\Gamma)\} = O(mn(m+n)^2 \log \Gamma/\Gamma)$ time, with probability at least $1 - \Gamma/((m+n)^2 n)$.*
3) *Now let us see how many times we need to run the while loop. Note that the **CF-Coloring** algorithm invokes the Algorithmic Local Lemma (Theorem 12) with the random variables being $c(v) : v \in V'$, and the (bad) events being $F_E : E \in \mathcal{E}'_G$ (that is, the event that E is not conflict-free colored). Since the probability of F_E is at most $1/\Gamma$ (as shown in proof of Lemma 6), as per Theorem 12, we need to run the while loop in **CF-Coloring** $O(n/\Gamma)$ number of times in expectation, in order to ensure that all edges $E \in \mathcal{E}'_G$ are conflict-free*

colored. Again, by Markov's inequality, the number of iterations of the while loop exceeds $O(n(m+n)/\Gamma)$ with probability at most $1/(m+n)$.

4) *Assume that the number of iterations of the while loop is at most $O(n(m+n)/\Gamma)$. The probability that a given run of the geometric reassignment exceeds $O(mn(m+n)^2 \log \Gamma/\Gamma)$ trials is at most $\Gamma/((m+n)^2 n)$. Taking union bound over the number of iterations of the while loop (which is at most $O(n(m+n)/\Gamma)$), the probability that any one of the geometric reassignments exceed $O(mn(m+n)^2 \log \Gamma/\Gamma)$ trials is at most $1/(m+n)$.*
5) *Using union bound over the bad probabilities given in the above points (1), (3) and (4) we can compute the probability that — (a) the initial geometric assignment of colors runs in at most $O(mn(m+n)^2 \log \Gamma/\Gamma)$ trials, (b) the while loop iterates at most $O(n(m+n)/\Gamma)$ times, and (c) each geometric reassignment takes at most $O(mn(m+n)^2 \log \Gamma/\Gamma)$ trials. This probability is at least $1-3/(m+n)$.*
6) *Combining the number of iterations of the while loop, with the time taken to run each iteration of the while loop, we get that the total time taken is at most $O(mn^2(m+n)^3 \log \Gamma/\Gamma^2)$, with probability at least $1-3/(m+n)$.*

■

CF-Covering($\mathcal{H}_i = (V, \mathcal{E}_i)$)

Input: A hypergraph $\mathcal{H}_i = (V, \mathcal{E}_i)$, with $|V| = m$, such that (i) every hyperedge $E \in \mathcal{E}_i$ intersects with at most Γ other hyperedges $E' \in \mathcal{E}_i$, and (ii) for every hyperedge $E \in \mathcal{E}_i$, we have $\frac{k_i}{2} \leq |E| < k_i$ where $k_i = \frac{2\log\Gamma - 1}{2^i}$.

Output: A conflict-free collection of colorings $\mathcal{C} = \{c_1, c_2, \dots, c_{t_i}\}$ of $\mathcal{H}_i$, where each $c_j : V \to \{1, 2\}$ and $t_i = \lceil 5k_i \log \Gamma \rceil$.

Algorithm:

- Let $q_i = \frac{1}{k_i}$.
- For $j = 1$ to t_i
  - For every $v \in V$, $c_j(v)$ is chosen independently with the following probabilities: $c_j(v) = 1$ with probability q_i and $c_j(v) = 2$ with probability $1 - q_i$.
- While $\exists E \in \mathcal{E}_i$ such that none of the above t_i colorings is a conflict-free coloring for E:

- Choose an arbitrary $E \in \mathcal{E}_i$ such that none of the t_i colorings is a conflict-free coloring for E.
- For each of the vertices $v \in E$, we reassign colors $c_j(v)$, for $j = 1$ to t_i, using the above probabilities (q_i and $1 - q_i$).

- Output $\mathcal{C} = \{c_1, c_2, \ldots, c_{t_i}\}$.

Remark 2. *The running time of Algorithm **CF-Covering** is dominated by the running time of the while loop. Within the while loop, testing for conflict-free property of all the hyper-edges requires $O(mt_in)$ time. The resampling step for the chosen hyperedge requires $O(mt_i)$ time. As per the algorithmic Local Lemma, we need to run the while loop $O(n/\Gamma)$ times in expectation (by similar calculations as in Remark 1). This yields an overall time complexity of $O(mt_in + mt_i) \cdot n/\Gamma = O(mt_in^2/\Gamma)$. Since $t_i = O(\log^2 \Gamma)$, we get an overall running time of $O(mn^2 \log^2 \Gamma/\Gamma)$ in expectation for the **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) procedure. By Markov's inequality, this also means that the probability of the overall running time for the **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) procedure being at most $O((m + n)mn^2 \log^2 \Gamma \log \kappa/\Gamma)$ is at least $1 - 1/((m + n) \log \kappa)$.* ■

Constructing a PIC scheme with length $O(\log^2 \Gamma)$: The following theorem summarizes the main contribution of this section.

Theorem 4. *For any $(n, m, \mathfrak{I})$-PICOD problem, let $\Gamma = \max_{r\in[n]} |\{r' \in [n]\backslash r : I_r \cap I_{r'} \neq \emptyset\}|$. Then there exists a binary linear scalar PIC for the given problem with length $O(\log^2 \Gamma)$. Thus $\ell_1^*(\mathcal{H}) = O(\log^2 \Gamma)$. Further, such a PIC can be constructed in a running time of $O(mn^2(m+n)^3)$, with probability at least $1 - 4/(m + n)$.*

Proof: The achievability follows from Lemma 6 in conjunction with Theorem 3. We now give the argument for the running time. Recall, from the proof of Lemma 6, that $\mathcal{H} = \mathcal{G} \cup \mathcal{H}_1 \cup \cdots \cup \mathcal{H}_{\lceil \log \kappa \rceil}$, where $\kappa = 2 \log \Gamma - 1$. The running time of **CF-Coloring**($\mathcal{G} = (V, \mathcal{E}_G)$) as given in Remark 1 is $O(mn^2(m + n)^3 \log \Gamma/\Gamma^2)$, with probability at least $1 - 3/(m + n)$. The running time of **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) for each $\mathcal{H}_i$ as given in Remark 2 is $O((m + n)mn^2 \log^2 \Gamma \log \kappa/\Gamma)$ with probability at least $1-1/((m+n) \log \kappa)$. Taking into account that we need to run **CF-Covering**($\mathcal{H}_i = (V, \mathcal{E}_i)$) for each $\mathcal{H}_i$, we get a total running time of all the calls to **CF-Covering** to be $O((m + n)mn^2 \log^2 \Gamma \log^2 \kappa/\Gamma)$ with probability at least $1 - 1/(m + n)$.

Combining the running time of **CF-Coloring** and all the calls to **CF-Covering**, we get a total running time of $O(mn^2(m+n)^3 \log\Gamma/\Gamma^2) + O((m+n)mn^2 \log^2\Gamma \log^2\kappa/\Gamma)$ with probability at least $1-4/(m+n)$. By considering the dominant terms, we get that the running time for obtaining all the colorings in the desired collection is upper bounded by $O(mn^2(m+n)^3 \log^2\Gamma \log^2\kappa/\Gamma) = O(mn^2(m+n)^3 \log^2\Gamma \log^2\log\Gamma/\Gamma)$ with probability at least $1-4/(m+n)$.

From the above constructed collection of colorings, and using similar arguments as in the proof of Theorem 3, we can obtain a PIC of length $O(\log^2\Gamma)$ for the given problem. To do this we have to stack the indicator matrices of the colorings in the collection. The indicator matrices of all the colorings in the collection can be done in time at most $O(m\log^2\Gamma)$ (as the total number of colors is at most $O(\log^2\Gamma)$ and the number of columns is m). There are $O(\log\log\Gamma)$ colorings, thus the stacking can be completed in time $O(\log\log\Gamma)$. Thus, the PIC can be obtained in total running time $O(mn^2(m+n)^3 \log^2\Gamma \log^2\log\Gamma/\Gamma) + O(m\log^2\Gamma) + O(\log\log\Gamma) = O(mn^2(m+n)^3)$, with probability at least $1-4/(m+n)$, where the last equality follows as $\log^2\Gamma\log^2\log\Gamma/\Gamma$ is at most a constant. ■

Remark 3. *We remark regarding deterministic algorithms for a conflict-free covering with $O(\log^2\Gamma)$ colors. To obtain a deterministic algorithm for any randomized algorithm generated as per the algorithmic Local Lemma (Theorem 12), we need to derandomize the Local Lemma. Derandomization of the Local Lemma has been studied [23], [24]. The details of such a derandomization procedure are sophisticated, hence we provide only pointers here. Applying an existing result from [24] (see Theorem 1.1 (1) in [24]), we can construct in polynomial time a conflict-free coloring of a hypergraph using $O(t\Gamma^{\frac{1+\epsilon}{t}}\log\Gamma)$ colors, where t and Γ are as defined in Theorem 1 and $\epsilon > 0$ is a constant. This suffices to get a deterministic polynomial time coloring algorithm for the hypergraph $\mathcal{G}$ in the proof of Lemma 6 using $O(\log^2\Gamma)$ colors. Following the same result from [24], one can get polynomial time algorithms for constructing conflict-free collection of colorings for hypergraphs $\mathcal{H}_i$ in the proof, such that the total number of colors used across all the colorings in such a collection is $O(k_i\log\Gamma)$. This enables us to get a PIC of length $O(\log^2\Gamma)$, in polynomial time.* ■

Comparison of the $O(\log^2\Gamma)$ upper bound on $\ell^*(\mathcal{H})$ with existing upper bounds: The original work of Brahma and Fragouli [7] showed the existence of an achievable scheme with length $O(\min\{\log m(1+\log^+(\frac{n}{\log m})), m, n\})$ (where $\log^+(x) = \max\{0, \log(x)\}$). For $m = n^\delta$ for some $\delta > 0$, this means the existence of a PIC with length $O(\log^2 n)$ is guaranteed. Further, in [8], an achievable scheme was presented for a PICOD problem with n receivers with length $O(\log^2 n)$. The algorithm in [8] was constructive, and had running time polynomial in the system parameters m, n. Our result, Theorem 4, gives an upper bound for the optimal PIC length based on the parameter Γ of the hypergraph. Given the set of vertices V and edges $\mathcal{E}$ of a hypergraph, the parameter Γ can be determined in $O(|V||\mathcal{E}|^2)$ time by a simple algorithm which runs through each edge, computing its intersection with all other edges. Further, the parameter $\Gamma \leq |\mathcal{E}|-1 = n-1$ always, but it could be much smaller in general, as suggested by the below example.

Example 4. *Consider the hypergraph $\mathcal{H} = (V, \mathcal{E})$, where $V = [m]$, $\mathcal{E} = \{\{i, i+1, i+2\} : i \in [m-2]\}$ for $m \geq 3$. Since every hyperedge overlaps with at most 3 other edges, we have $\Gamma = 3$. The result from [7] suggests the existence of a code of length $O(\log^2 m)$, where, by Lemma 6, we have a code of constant length (as m grows).* ■

IV. 'Local' Conflict-Free Chromatic Number and Pliable Index Coding

In this section, we define the *local* versions of the conflict-free chromatic number and covering number. This results in further refined upper bounds for $\ell_1^*(\mathcal{H})$. However, the PICOD schemes resulting from these local parameters may require a larger field size than the schemes in the previous sections, as they rely on the existence of MDS codes of appropriate lengths to exist.

We first define the local conflict-free chromatic number.

Definition 4 (Local Conflict-Free Chromatic Number)**.** *Given a hypergraph $\mathcal{H}(V, \mathcal{E})$, the local conflict-free chromatic number of $\mathcal{H}$ is given by*

$$\Delta(\mathcal{H}) = \min_{\substack{C: C \text{ is a CF} \\ \text{coloring of } \mathcal{H}}} \underbrace{\max_{E \in \mathcal{E}} |\{C(v) : v \in E\}|}_{\Delta_C(\mathcal{H})}.$$

For convenience, we define $\Delta_C(\mathcal{H}) = \max_{E \in \mathcal{E}} |\{C(v) : v \in E\}|$, where C is a conflict-free coloring of $\mathcal{H}$. Therefore, $\Delta(\mathcal{H}) = \min_C \Delta_C(\mathcal{H})$, where the minimum is over all such colorings C of $\mathcal{H}$.

Thus, the local chromatic number of $\mathcal{H}$ is the maximum number of colors that appears in any edge of $\mathcal{H}$, minimized across all conflict-free (CF) colorings.

Example 5. *For the hypergraph $\mathcal{H}$ in Fig. 2 (in Section III), $\chi_{CF}(\mathcal{H}) = 3$ and we see a coloring with three colors. However, only two colors appear in each edge, in the coloring shown. Further,* *by definition, in any conflict-free coloring of $\mathcal{H}$, every edge will indeed see at least two colors. Thus, $\Delta(\mathcal{H}) = 2$.* ■

The following observation is straightforward from the definition of $\Delta(\mathcal{H})$.

Observation 2. $\Delta(\mathcal{H}) \leq \chi_{CF}(\mathcal{H})$.

The following lemma shows that the gaps between $\Delta(\mathcal{H})$, $\alpha_{CF}(\mathcal{H})$, and $\chi_{CF}(\mathcal{H})$, can be arbitrarily large.

Lemma 7. *There exists a hypergraph $\mathcal{H}$ with n hyperedges for which $\Delta(\mathcal{H}) = 2$, while $\alpha_{CF}(\mathcal{H}) = \Theta(\log n)$ and $\chi_{CF}(\mathcal{H}) = \Theta(\sqrt{n})$.*

Proof: Consider the 2-uniform hypergraph with m vertices and all the 2-sized subsets of $[m]$ as hyperedges. We have already shown the values of the parameters $\alpha_{CF}(\mathcal{H})$ and $\chi_{CF}(\mathcal{H})$ in Lemma 5. To see that $\Delta(\mathcal{H}) = 2$, we see that in any conflict-free coloring of $\mathcal{H}$, each edge sees two distinct colors. ■

We now show that there is a scalar PIC of length $\Delta(\mathcal{H})$ for the PICOD problem given by $\mathcal{H}([m], \mathfrak{I})$, provided we are operating over a sufficiently large finite field. For K, N being positive integers such that $K \leq N$, let a linear code of dimension K and length N be referred to as an $[N, K]$ code. Recall that in the $K \times N$ generator matrix of a maximum distance separable (MDS) $[N, K]$ code, any K columns are linearly independent. Further, such MDS codes exist as long as the size of the finite field is at least N. Below, we define the MDS matrix associated with a given coloring C of the graph $\mathcal{H}([m], \mathfrak{I})$.

Definition 5 (MDS matrix associated with a CF coloring of $\mathcal{H}$)**.** *Let C be a conflict-free coloring of $\mathcal{H}([m], \mathfrak{I})$ that uses colors from $[L]$ (for some positive integer L). Let $\Delta_C = \max_{E \in \mathcal{E}} |\{C(v) : v \in E\}|$. Let G' denote the $\Delta_C \times L$ generator matrix of an $[L, \Delta_C]$ MDS code. We index the columns of G' by the set $[L]$, and denote the column indexed by $d \in [L]$ as G'_d. Consider the $\Delta_C \times m$ matrix G defined as follows (the columns of G are indexed as $G_i : i \in [m]$): For $i \in [m]$, we set $G_i = G'_{C(i)}$. We refer to G as the MDS matrix associated with the coloring C. Fig. 4 illustrates this construction.*

Using the matrix defined above, we show in the below theorem the achievability of length $\Delta(\mathcal{H})$.

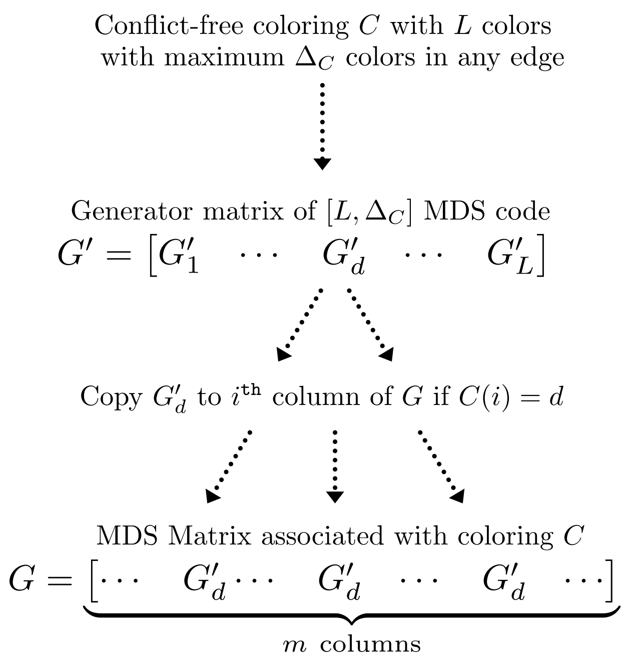

Figure 4: Illustration of the construction of the matrix G, which is the MDS matrix associated with the coloring C which consists of L colors and $\Delta_C(\mathcal{H}) = \max_{E\in\mathcal{E}} |\{C(v) : v \in E\}| = 2$, from the generator matrix of an $[L, \Delta_C]$ MDS code G'.

Theorem 5. *Let C be a conflict-free coloring of the PICOD hypergraph $\mathcal{H}$ using L colors. Then, there exists a PICOD of length $\Delta_C(\mathcal{H})$ for $\mathcal{H}$ over any field of size at least L. Thus,*

$$\ell_1^*(\mathcal{H}) \leq \Delta(\mathcal{H}).$$

Proof: Let C be the conflict-free coloring of $\mathcal{H}$, that uses colors from the set $[L]$. We now show that the MDS matrix G associated with the coloring C (of size $\Delta_C(\mathcal{H}) \times m$ as defined in Definition 5) satisfies the properties in Lemma 1, and hence is a valid PIC for $\mathcal{H}$. Note that this code would have length $\Delta_C(\mathcal{H})$, which we denote by Δ_C in this proof. By the definition of $\Delta(\mathcal{H})$, our proof would then be complete.

To see this, consider any $I_r \in \mathfrak{I}$. By definition of G and Δ_C, we have that

$$|\{G_i : i \in I_r\}| \leq \Delta_C. \tag{1}$$

Now, by the definition of conflict-free coloring C, at least one vertex $v \in I_r$ is such that $C(v) \neq C(v'), \forall v' \in I_r \backslash v$. Thus, by the definition of matrix G, we have the following.

We first note that the vector G_v appears exactly once in the collection $\{G_i : i \in I_r\}$. Since the columns of G are taken from the generator matrix G' of an $[L, \Delta_C]$ MDS code, any Δ_C distinct columns of G are linearly independent. Further, G_v is linearly independent of the space spanned by any collection of $(\Delta_C - 1)$ other columns of G'. By (1) and the above observations, the columns in $\{G_i : i \in I_r \backslash v\}$ lie in the subspace spanned by some $(\Delta_C - 1)$ columns of G' apart from G_v. Thus, we have that $\mathsf{span}(\{G_v\}) \cap \mathsf{span}(\{G_i : i \in I_r \backslash v\}) = \{\mathbf{0}\}$. Thus G satisfies receiver r by Definition 1. As r is arbitrary, by Lemma 1, G represents a valid PIC for $\mathcal{H}$. ■

Example 6. *We give an example of the code construction involved in proof of Theorem 5. Consider the* 2*-uniform hypergraph* $\mathcal{H}$ *of* $V = \{1, \ldots, 10\}$ *with all the* $\binom{10}{2}$ *hyperedges, and let* C *be a conflict-free coloring of* $\mathcal{H}$*. Any such coloring requires at least* 10 *colors, and there exists an edge with* 2 *colors. Thus, we have* $\Delta(\mathcal{H}) = 2$*. We define the encoder matrix as* G *as the encoder matrix of a* $[10, 2]$ *MDS code. It is easy to check that this satisfies all the receivers.*

■

A. Local conflict-free covering number and PICOD

In this subsection, we define a *local* version of the covering number arising due to conflict-free collections of $\mathcal{H}$ (as defined in Definition 3) and relate it to an achievable scheme for the PICOD problem. In Section V, we derive an upper bound on the local covering number, thus showing the existence of PIC schemes with length equal to this upper bound, for the more general scenario of t-request PICOD.

Definition 6 (Local Conflict-Free Covering Number)**.** *Given a hypergraph* $\mathcal{H}(V, \mathcal{E})$*, the local conflict-free covering number of* $\mathcal{H}$ *is given by*

$$\lambda(\mathcal{H}) = \min_{\substack{\mathfrak{C}:\mathfrak{C} \text{ is a CF} \\ \text{collection of } \mathcal{H}}} \overbrace{\sum_{C^p \in \mathfrak{C}} \underbrace{\left(\max_{E \in \mathcal{E}} |\{C^p(v) : v \in E\}| \right)}_{\Delta_{C^p}(\mathcal{H})}}^{\lambda_{\mathfrak{C}}(\mathcal{H})}.$$

For convenience, we define $\lambda_{\mathfrak{C}}(\mathcal{H}) = \sum_{C^p \in \mathfrak{C}} \Delta_{C^p}(\mathcal{H})$*, where* $\mathfrak{C}$ *is a conflict-free collection of colorings of* $\mathcal{H}$ *and* $\Delta_{C^p}(\mathcal{H})$ *is as in Definition 4. Therefore,* $\lambda(\mathcal{H}) = \min_{\mathfrak{C}} \lambda_{\mathfrak{C}}(\mathcal{H})$*, where the minimum is over all such collections* $\mathfrak{C}$ *of* $\mathcal{H}$*.*

Example 7. *Consider the graph* $\mathcal{H} = K_7$ *shown in Fig. 3a, for which we have a conflict-free collection* $\mathfrak{C}$ *consisting of the* $P = 3$ *colorings (denote them by* C^1, C^2 *and* C^3*) shown in Fig.*

3b-3d. In this collection, in each coloring, the maximum number of colors seen by each edge is 2*. Thus,* $\Delta_{C^p}(\mathcal{H}) = 2, \forall p \in \{1,2,3\}$*. Thus, for this collection, we have* $\lambda_{\mathfrak{C}}(\mathcal{H}) = \sum_{p=1}^{3} \Delta_{C^p}(\mathcal{H}) = 6$*. However, for the conflict-free collection* $\mathfrak{C}'$*, consisting of the single coloring* C *which assigns a distinct color for each of the* 7 *vertices, we have that* $\lambda_{\mathfrak{C}'}(\mathcal{H}) = \Delta_C(\mathcal{H}) = 2$*.* ■

The following theorem shows that there is an achievable PICOD scheme for the PICOD hypergraph $\mathcal{H}([m], \mathfrak{I})$ with length $\lambda(\mathcal{H})$, and summarizes the relationship of scalar PIC schemes with the various chromatic parameters considered in this work, for the single-request case.

Theorem 6. $\ell_1^*(\mathcal{H}) \leq \lambda(\mathcal{H}) \leq \min(\alpha(\mathcal{H}), \Delta(\mathcal{H})) \leq \chi(\mathcal{H})$.

Proof: We first show the first inequality. Let $\mathfrak{C} = \{C^1, \ldots, C^P\}$ be a conflict-free collection of colorings of $\mathcal{H}([m], \mathfrak{I})$. For convenience, in this proof, we use $\Delta(p)$ to denote $\Delta_{C^p}(\mathcal{H})$. We now show a valid PIC encoder matrix for $\mathcal{H}$ of size $\sum_{p \in [P]} \Delta(p)$.

Let the MDS matrix of size $\Delta(p) \times m$ associated with the coloring C^p be denoted by G^p. Consider the $\sum_{p \in [P]} \Delta(p) \times m$ matrix

$$G = \begin{bmatrix} G^1 \\ \vdots \\ G^P \end{bmatrix}.$$

We now show that the matrix G is a valid encoder of a PIC for $\mathcal{H}$. By definition of $\mathfrak{C}$, for any edge $I_r \in \mathfrak{I}$, there is some coloring C^p in which there exists some $v \in I_r$ such that v is colored by C^p and $C^p(v) \cap C^p(i) = \emptyset, \forall i \in I_r \setminus \{v\}$. By arguments similar to the proof of Theorem 5, we have that G^p satisfies receiver r. As r is arbitrary, by Lemma 2, G satisfies all receivers and is a valid PIC. Invoking the definition of $\lambda(\mathcal{H})$ completes the proof of the first inequality.

The inequality that $\lambda(\mathcal{H}) \leq \Delta(\mathcal{H})$ holds because any conflict-free coloring C of $\mathcal{H}$ generates a conflict-free collection of $\mathcal{H}$ containing only C. Further, by definitions of $\lambda(\mathcal{H})$ and $\alpha_{CF}(\mathcal{H})$, it holds that $\lambda(\mathcal{H}) \leq \alpha_{CF}(\mathcal{H})$. The final inequality follows from Observation 2 and Lemma 4. This completes the proof. ■

V. The t-Requests Case

In the previous sections, we considered the PICOD setting where each receiver demands one of the messages in its request-set. In the present section, we generalize our results to the scenario where each receiver r has to be sent any $\min(|I_r|, t)$-sized subset of messages indexed by its request-set I_r, where t is a positive integer. We shall call PICOD schemes which satisfy the above t-requests scenario as *t-request pliable index codes*, or t-request PICs. In the rest of the section, we shall denote the smallest length of any t-request PIC for the PICOD problem defined by hypergraph $\mathcal{H}$ as $\ell^{*(t)}(\mathcal{H})$. It was shown in [8] that, for any PICOD problem with n receivers, a t-request PIC with length $O(t \log n + \log^2 n)$ exists, and can be designed in polynomial time (in number of receivers n and the messages m). In the present section, we define the notion of t-strong conflict-free colorings and their related parameters. We prove the generalization of Theorem 6 to the t-request PICOD scenario. We then obtain upper bounds on some of the t-strong chromatic parameters, giving us upper bounds on $\ell^{*(t)}(\mathcal{H})$. We also show that these bounds are asymptotically tight, upto a multiplicative factor $\log t$.

The notion of strong conflict-free coloring of hypergraphs was introduced by Horev et al. in [25] as a conflict-free coloring in which any edge of the hypergraph 'sees' more than one distinct color. Formally, a *t-strong conflict-free coloring* of hypergraph $\mathcal{H} = (V, \mathcal{E})$ with L labels (or colors) is an assignment $C : V \to [L]$ such that the following holds.

- For any edge $E \in \mathcal{E}$, there exist $\min(t, |E|)$ vertices in E which get distinct labels, i.e., there exists $V_E \subseteq E$ such that (a) $|V_E| = \min(t, |E|)$, (b) $|\{C(v) : v \in V_E\}| = |V_E|$, and (c) $\{C(v) : v \in V_E\} \cap \{C(v') : v' \in E \setminus V_E\} = \emptyset$.

The minimum L such that a t-strong conflict-free coloring exists for $\mathcal{H}$ is then called the *t-strong conflict-free chromatic number* of $\mathcal{H}$, which we denote by $\chi^{(t)}(\mathcal{H})$. The notion of a *t-strong conflict-free collection* of colorings and the *t-strong conflict-free covering number* of $\mathcal{H}$ (denoted by $\alpha^{(t)}(\mathcal{H})$) are defined as in Definition 3, with the only difference being that we want the colorings $C^p : p \in [P]$ such that for each $E \in \mathcal{H}$, there exists some C^p such that E sees $\min(t, |E|)$ colors, each exactly once, under the coloring C^p. In other words, in a t-strong conflict-free collection, the colorings $C^p : p \in [P]$ are t-strong conflict-free colorings for subgraphs $\mathcal{H}_p : p \in [P]$ of $\mathcal{H}$ respectively, where $\cup_{p \in [P]} \mathcal{H}_p = \mathcal{H}$. Similarly, we can define the *t-strong local conflict-free chromatic number* as follows.

$$\Delta^{(t)}(\mathcal{H}) \triangleq \min\{\Delta_C(\mathcal{H}) : C \text{ is a } t\text{-strong conflict-free coloring for } \mathcal{H}\},$$

where $\Delta_C(\mathcal{H})$ is as in Definition 4. Finally, the *t-strong local conflict-free covering number* $\lambda^{(t)}(\mathcal{H})$ can be defined as

$$\lambda^{(t)}(\mathcal{H}) \triangleq \min\{\lambda_{\mathfrak{C}}(\mathcal{H}) : \mathfrak{C} \text{ is a } t\text{-strong conflict-free collection of } \mathcal{H}\},$$

where $\lambda_{\mathfrak{C}}(\mathcal{H})$ is as in Definition 6. The following results show the utility of t-strong conflict-free colorings to the t-request PICOD problems.

Lemma 8. *Let C be a t-strong conflict-free coloring for a PICOD hypergraph $\mathcal{H}$ that uses L colors. Then there exists a t-request PIC of length L for $\mathcal{H}$ over every field. Further, there is a t-request PIC for $\mathcal{H}$ of length $\Delta_C(\mathcal{H})$ over any field $\mathbb{F}$ with $|\mathbb{F}| \geq L$.*

Proof: The proof for the first part follows that of Theorem 2 (which will in turn invoke condition (P') in Lemma 3 for each receiver), while the proof for the second part follows that of Theorem 5 via MDS matrices as defined in Definition 5, which always exist for $|\mathbb{F}| \geq L$. ■

Lemma 9. *Let $\mathfrak{C} = \{C^p : p \in [P]\}$ be a t-strong conflict-free collection for a PICOD hypergraph $\mathcal{H}$, that uses $L_p : p \in [P]$ colors for the colorings $C^p : p \in [P]$ respectively. Then there is a t-request PIC of length $\sum_{p\in[P]} L_p$ for $\mathcal{H}$ over every field. Further, there is a t-request PIC of length $\lambda_{\mathfrak{C}}(\mathcal{H})$ for $\mathcal{H}$ over every field $\mathbb{F}$ such that $|\mathbb{F}| \geq \max_{p\in[P]} L_p$, where $\lambda_{\mathfrak{C}}(\mathcal{H})$ is as defined in Definition 6.*

Proof: The first part follows by arguments similar to Theorem 3, while the second part uses arguments as in Theorem 6. ■

We then have the following theorem, which summarizes the extensions of our results to the t-request PICOD scenario, which was also represented in Fig. 1.

Theorem 7. *For the PICOD problem defined by $\mathcal{H}$, we have,*

$$\ell^{*(t)}(\mathcal{H}) \leq \lambda^{(t)}(\mathcal{H}) \leq \min(\Delta^{(t)}(\mathcal{H}), \alpha^{(t)}(\mathcal{H})) \leq \chi^{(t)}(\mathcal{H}).$$

Proof: The claim that $\lambda^{(t)}(\mathcal{H})$ is an upper bound for $\ell^{*(t)}(\mathcal{H})$ follows from Lemma 9. The other inequalities follow essentially from the respective definitions, similar to the proof of Theorem 6. ■

A. Upper Bound for $\lambda^{(t)}(\mathcal{H})$

In this section, we show an upper bound for $\lambda^{(t)}(\mathcal{H})$ as $\max(O(\log \Gamma \log m), O(t \log m))$, under some condition on the number of messages. We then show that the same bound holds for $\alpha^{(t)}(\mathcal{H})$ also, under some restriction on the size of each edge. We first prove a partitioning lemma, which will prove useful later.

Lemma 10. *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph with $|V| = m$ and $|\mathcal{E}| = n$. It is given that every hyperedge in $\mathcal{H}$ intersects with at most Γ other hyperedges. Let ρ be a positive integer such that $\log(6(\Gamma+1)) \leq \rho < \frac{m}{12}$. Then, for some $r < \log_2 m$, there exist subsets $V_i \subseteq V : i \in [r]$, $\mathcal{E}_i \subseteq \mathcal{E} : i \in [r]$, and $\mathcal{E}' \subseteq \mathcal{E}$ such that*
(i) $(\uplus_i \mathcal{E}_i) \uplus \mathcal{E}' = \mathcal{E}$.
(ii) $\forall i \in [r]$, for each $E \in \mathcal{E}_i$, we have $6\rho < |E \cap V_i| \leq 36\rho$, and
(iii) for each $E \in \mathcal{E}'$, we have $|E| \leq 12\rho$.

Proof: Let $\mathcal{E}' = \{E \in \mathcal{E} \ : \ |E| \leq 12\rho\}$. Let r be the largest non-negative integer so that $12\rho < \frac{m}{2^{r-1}}$ (such an $r \geq 1$ exists since $12\rho < m$). Thus, $r < \log_2 m$. For each $i \in [r]$, let $m_i = \frac{m}{2^i}$ and let $\mathcal{E}_i = \{E \in \mathcal{E} \setminus \mathcal{E}' \ : \ m_i < |E| \leq m_{i-1}\}$. Consider an $i \in [r]$. Below, we explain how we construct V_i. Independently and uniformly at random select a vertex $v \in V$ into V_i with probability $\frac{12\rho}{m_i}$. Let X_E^i be a random variable that denotes $|E \cap V_i|$, for a hyperedge $E \in \mathcal{E}_i$. Let $\mu_E^i := E[X_E^i]$. Then, $\mu_E^i = \frac{12|E|\rho}{m_i}$. Since $m_i < |E| \leq m_{i-1}$, we have $12\rho < \mu_E^i \leq 24\rho$. Applying the Chernoff bound given in Theorem 13 (see Appendix A) with $\delta = 1/2$, we get $\mathsf{Pr}[|X_E^i - \mu_E^i| \geq \frac{1}{2}\mu_E^i] \leq 2\mathsf{e}^{-\frac{\mu_E^i}{12}} < 2\mathsf{e}^{-\frac{12\rho}{12}} \leq 2\mathsf{e}^{-\log(6(\Gamma+1))} = \frac{1}{3(\Gamma+1)}$.

Let A_E^i denote the bad event that $|X_E^i - \mu_E^i| \geq \frac{\mu_E^i}{2}$. We have shown that $\mathsf{Pr}[A_E^i] \leq \frac{1}{3(\Gamma+1)}$. Using the fact that A_E^i depends on at most Γ other events in $\{A_{\tilde{E}}^i : \tilde{E} \in \mathcal{E}_i\}$ (by definition. of Γ), and since it is true that $\mathsf{e} \cdot \frac{1}{3(\Gamma+1)} \cdot (\Gamma+1) \leq 1$, from Lemma 14 (the Local Lemma, given in Appendix A), we get $\mathsf{Pr}[\cap_{E \in \mathcal{E}_i}(\overline{A_E^i})] > 0$. Hence, there exists a V_i such that $\forall E \in \mathcal{E}_i$, $|X_E^i - \mu_E^i| < \frac{1}{2}\mu_E^i$. Since $12\rho < \mu_E^i \leq 24\rho$, this implies that there is a V_i such that $\forall E \in \mathcal{E}_i$, $6\rho < |E \cap V_i| \leq 36\rho$. This completes the proof. ∎

Using the algorithmic version of the Local Lemma (Theorem 12, Appendix A), we obtain the below randomized algorithm for constructing the set V_i, as in Lemma 10.

Partition($V, \mathcal{E}_i, \rho$)

Input: (i) A set V, (ii) $\mathcal{E}_i = \{E \subseteq V \ : \ m_i \leq |E| \leq m_{i-1}\}$, where $m_i = m/2^i$ (for some $i \leq r$ as in Lemma 10) and every set in $\mathcal{E}_i$ overlaps with at most Γ other sets in $\mathcal{E}_i$, and (iii) a positive integer ρ such that $\log(6(\Gamma+1)) \leq \rho < \frac{m}{12}$.
Output: A set $V_i \subseteq V$ such that every set $E \in \mathcal{E}_i$ satisfies the condition $6\rho < |E \cap V_i| \leq 36\rho$.
Algorithm:

- Each vertex $v \in V$ is independently and uniformly chosen to be in V_i with probability

$\frac{12\rho}{m_i}$.

- While $\exists E \in \mathcal{E}_i$ that does not satisfy the output condition
  - Choose an arbitrary $E \in \mathcal{E}_i$ that does not satisfy the output condition
  - Resample each vertex $v \in E$. That is, for each $v \in E$, independently and uniformly decide to include it in V_i with probability $\frac{12\rho}{m_i}$.
- Output V_i.

Remark 4. *In this remark, we argue that the sets $V_i : i \in [r]$ can be constructed in time polynomial in the graph parameters using Algorithm* ***Partition****. Let Y_v^i be the indicator random variable that vertex v is chosen into V_i in the above random process. We can apply Theorem 12 by considering the random variables $\mathcal{P} = \{Y_v^i : v \in V\}$, and the events A_E^i (as in the proof of Lemma 10) as the bad events. Since the probability of A_E^i is at most $\frac{1}{3(\Gamma+1)} \leq \frac{1}{\Gamma}$, we need at most $|\mathcal{E}_i|/\Gamma \leq n/\Gamma$ resamplings in expectation for satisfying the output condition $6\rho < |E \cap V_i| \leq 36\rho, \forall E \in \mathcal{E}_i$. Prior to each resampling, we would also need to test if all $E \in \mathcal{E}_i$ satisfies the output condition for the current choice of V_i. This takes $O(mn)$ time. As $r < \log_2 m$, we see that the time required to obtain all $V_i : i \in [r]$ is at most $O(mn^2 \log_2 m/\Gamma)$ in expectation.* ■

We now use Lemma 10 to show an upper bound on $\lambda^{(t)}(\mathcal{H})$, under a condition on the number of vertices.

Lemma 11. *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph with $|V| = m$ and $|\mathcal{E}| = n$. It is given that every hyperedge in $\mathcal{H}$ intersects with at most Γ other hyperedges. Then, for any positive integer t such that* $\max(\log(6(\Gamma+1)), t) < \frac{m}{12}$,

$$\lambda^{(t)}(\mathcal{H}) = \max(O(\log \Gamma \log m), O(t \log m)).$$

Proof: Let $\rho = \max(\log(6(\Gamma+1)), t)$. Then, from Lemma 10, we have $V_1, \ldots, V_r \subseteq V$ and $\mathcal{E}_1 \uplus \cdots \uplus \mathcal{E}_r \uplus \mathcal{E}' = \mathcal{E}$ with $r < \log_2 m$ such that (i) $\forall i \in [r]$, for each $E \in \mathcal{E}_i$, we have $6\rho < |E \cap V_i| \leq 36\rho$, and (ii) for each $E \in \mathcal{E}'$ we have $|E| \leq 12\rho$. For each $i \in [r]$, we define hypergraphs $\mathcal{H}_i = (V, \mathcal{E}_i)$. We define a t-strong conflict-free coloring c_i for $\mathcal{H}_i$ using $|V_i| + 1$ colors in which all the vertices in V_i get a distinct color from the first $|V_i|$ colors and all the vertices in $V \setminus V_i$ get the $(|V_i| + 1)^{th}$ color. In such a coloring, each hyperedge $E \in \mathcal{E}_i$ sees at least $6\rho + 1$ colors exactly once and at most $36\rho + 1$ colors in total. Finally, we define a t-strong conflict-free coloring c' for $\mathcal{H}'$ using $|V|$ colors that gives a distinct color to every vertex in V. In this coloring, every hyperedge $E \in \mathcal{E}'$ sees $|E|$ (which is $\leq 12\rho$) colors, each color exactly once. Thus, the collection $\mathfrak{C} = \{c_i : i \in [r]\} \cup \{c'\}$ is a conflict-free collection. It is easy to see that $\lambda_{\mathfrak{C}}(\mathcal{H}) = O(\rho r) = \max(O(\log \Gamma \log m, t \log m))$. By definition of $\lambda^{(t)}(\mathcal{H})$, the proof is complete. ■

Remark 5. *In this remark, we argue that we can obtain a construction of a conflict-free collection of colorings with properties as in Lemma 11, in expected running time that is polynomial in the system parameters. From the proof of Lemma 10, we clearly can construct the sets $\mathcal{E}_i \subseteq \mathcal{E} : i \in [r]$, and $\mathcal{E}' \subseteq \mathcal{E}$ satisfying conditions in Lemma 10 in $O(mn)$ time. Combining this observation with Remark 4, we obtain a randomized algorithm for the construction of a collection of colorings with properties as in Lemma 11, that runs in expected time at most $O(mn^2 \log_2 m/\Gamma)$.* ■

B. Upper Bound for $\alpha^{(t)}(\mathcal{H})$

We now prove that a similar upper bound as Lemma 11 holds for $\alpha^{(t)}(\mathcal{H})$ also, under a condition in the size of each edge of $\mathcal{H}$.

Lemma 12. *Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph with $|V| = m$ and $|\mathcal{E}| = n$. It is given that every hyperedge in $\mathcal{H}$ intersects with at most Γ other hyperedges. For t being any positive integer such that $|E| > 12 \max(\log(6(\Gamma + 1)), t), \forall E \in \mathcal{E}$, we have,*

$$\alpha^{(t)}(\mathcal{H}) = \max(O(\log \Gamma \log m), O(t \log m)).$$

Proof: Let $t_1 = \max(\log(6(\Gamma + 1)), t)$. Since $|E| > 12t_1$, it follows that $\max(\log(6(\Gamma + 1)), t_1) < \frac{m}{12}$. Applying Lemma 10 with $\rho = t_1$, we get $V_1, \ldots, V_r \subseteq V$ and $\mathcal{E}_1 \uplus \cdots \uplus \mathcal{E}_r = \mathcal{E}$ with $r < \log_2 m$ such that $\forall i \in [r]$, for each $E \in \mathcal{E}_i$), we have $6t_1 < |E \cap V_i| \leq 36t_1$. Note that $\mathcal{E}' = \emptyset$ as every hyperedge in $\mathcal{H}$ is of size greater than $12t_1$.

Consider an $i \in [r]$. We describe a t-strong conflict-free coloring c_i for the hyperedges in $\mathcal{E}_i$. For each vertex in V_i, assign a color that is chosen independently and uniformly at random from a set of $19\mathrm{e}^2 t_1$ colors. All the vertices in $V \setminus V_i$ are assigned the same color, a color different from the $19\mathrm{e}^2 t_1$ colors used to color the vertices in V_i. For each $E \in \mathcal{E}_i$, let $z_E^i \triangleq |E \cap V_i|$. We have $6t_1 < z_E^i \leq 36t_1$. Let B_E^i be the bad event that $E \cap V_i$ is colored with $\leq \lceil \frac{z_E^i + t}{2} \rceil \leq \frac{z_E^i + t + 1}{2}$ colors, where the last inequality holds since z_E^i and t are integers. Note that if B_E^i does not occur, then $E \cap V_i$ has some t colors that appear exactly once. Now we estimate the probability of B_E^i.

$$\begin{aligned}
\mathsf{Pr}[B_E^i] &\leq \binom{19\mathsf{e}^2 t_1}{(z_E^i + t + 1)/2} \left(\frac{(z_E^i + t + 1)/2}{19\mathsf{e}^2 t_1}\right)^{z_E^i} \\
&\leq \left(\frac{\mathsf{e} \cdot 19\mathsf{e}^2 t_1}{(z_E^i + t + 1)/2}\right)^{(z_E^i + t + 1)/2} \left(\frac{(z_E^i + t + 1)/2}{19\mathsf{e}^2 t_1}\right)^{z_E^i} \qquad \left(\text{since } \binom{n}{k} \leq \left(\frac{\mathsf{e}n}{k}\right)^k\right) \\
&= \mathsf{e}^{t+1} \left(\frac{z_E^i + t + 1}{38\mathsf{e} t_1}\right)^{(z_E^i - t - 1)/2} \\
&\leq \mathsf{e}^{t_1+1} \left(\frac{36t_1 + t_1 + 1}{38\mathsf{e} t_1}\right)^{(6t_1 + 1 - t_1 - 1)/2} \qquad (\text{since } t \leq t_1, \text{ and } 6t_1 < z_E^i \leq 36t_1) \\
&\leq \mathsf{e}^{t_1+1} \left(\frac{1}{\mathsf{e}}\right)^{5t_1/2} \leq \frac{\mathsf{e}}{\mathsf{e}^{1.5t_1}} \leq \frac{\mathsf{e}}{\mathsf{e}^{1.5\log(6(\Gamma+1))}} \\
&< \frac{1}{4(\Gamma + 1)} \quad .
\end{aligned}$$

Since $\mathsf{e} \cdot \frac{1}{4(\Gamma+1)} \cdot (\Gamma + 1) \leq 1$, by Lemma 14 (the Local Lemma in Appendix A, with $d = \Gamma$ and $p = \frac{1}{4(\Gamma+1)}$) , we get $\mathsf{Pr}[\cap_{E \in \mathcal{E}_i} (\overline{B_E^i})] > 0$. That is, $\forall E \in \mathcal{E}_i$, there is a t-strong conflict free coloring of $E \cap V_i$ with $19\mathsf{e}^2 t_1$ colors. Since such a color exists for each $i \in [r]$, we have that the conflict-free collection of colorings $\mathfrak{C} = \{c_i : i \in [r]\}$. The total number of colors in the collection is $O(r\rho) = O(rt_1)$. Using the fact that $t_1 = \max(\log(6(\Gamma+1)), t)$ and as $r < \log_2 m$, the proof is complete. ■

The following randomized algorithm generates the coloring with the properties according to Lemma 12 in time polynomial in system parameters, in the expectation.

t-Strong-CF-Covering($\mathcal{H} = (V, \mathcal{E}), t$)

Input: A hypergraph $\mathcal{H} = (V, \mathcal{E})$, with $|V| = m$, $|\mathcal{E}| = n$, and a positive integer t, such that (i) Every hyperedge $E \in \mathcal{E}$ intersects with at most Γ other hyperedges $E' \in \mathcal{E}$. (ii) Every $E \in \mathcal{E}$ satisfies $|E| > 12t_1$, where $t_1 = \max(\log(6(\Gamma + 1)), t)$.

Output: A conflict-free collection of colorings $\mathcal{C} = \{c_1, c_2, \ldots, c_r\}$ of $\mathcal{H}$, where each $c_i : V \to \{0\} \cup [19\mathsf{e}^2 t_1]$ and $r < \log_2 m$.

Algorithm:

- For $i = 1$ to $\log_2 m - 1$
  - $\mathcal{E}_i = \{E \in \mathcal{E} \,:\, m_i < |E| \leq m_{i-1}\}$, where $m_i = m/2^i$.

- $V_i = \text{Partition}(V, \mathcal{E}_i, t_1)$.
  - For every $v \in V \setminus V_i$, $c_i(v) = 0$.
  - For each vertex $v \in V_i$, let $c_i(v)$ be a color chosen independently and uniformly at random from $[19\mathrm{e}^2 t_1]$.
  - While $\exists E \in \mathcal{E}_i$ such that $E \cap V_i$ contains $\leq \lceil \frac{z_E^i + t}{2} \rceil$ distinct colors
    * Choose an arbitrary $E \in \mathcal{E}_i$ such that $E \cap V_i$ contains $\leq \lceil \frac{z_E^i + t}{2} \rceil$ distinct colors.
    * Recolor each vertex $v \in E$. That is, for each $v \in E$, independently and uniformly assign $c_i(v)$ from $[19\mathrm{e}^2 t_1]$.
- Output $\mathcal{C} = \{c_1, c_2, \ldots, c_r\}$.

Remark 6. *In this remark, we argue that the expected running time of Algorithm* t***-Strong-CF-Covering*** *is polynomial in the system parameters. For each* i*, we have a randomized algorithm for getting the* V_i *in Lemma 10 that runs in expected time at most* $O(mn^2/\Gamma)$ *(see Remark 4). We can again apply Theorem 12 by considering the random variables* $\mathcal{P} = \{c_i(v) : v \in V_i\}$ *and the events* B_E^i *(as in Lemma 12) as the bad events. So we need at most* $|\mathcal{E}_i|/\Gamma \leq n/\Gamma$ *recolorings in expectation. Prior to each recoloring, we would also need to test if any of the* B_E^i*'s occur for the current coloring. This takes* $O(mn)$ *time. So overall the time required is at most* $O(mn^2/\Gamma)$ *in expectation. Since* i *ranges from 1 to* $\log_2 m - 1$*, the total running time of Algorithm* t***-Strong-CF-Covering*** *is* $O(mn^2 \log m/\Gamma)$ *in expectation.* ■

The following theorem summarizes our results in the current and the previous subsection.

Theorem 8. *For a* t*-request PICOD problem defined by the hypergraph* $\mathcal{H}(V = [m], \mathfrak{I})$ *with* m *messages and the overlap parameter* Γ*, there exists a PIC that satisfies all clients and has length* $\ell = \max(O(\log \Gamma \log m), O(t \log m))$*, over any field of size at least* m*, if* $m > 12\max(t, \log(6(\Gamma + 1)))$*. Also, there is a binary (in other words, a finite-field oblivious) PIC satisfying all clients with the same length, if* $|I_r| > 12\max(t, \log(6(\Gamma+1)))$*, for each request-set* $I_r \in |\mathfrak{I}|$*. Further, these PIC schemes can be constructed in running time* $O(m^4 n^2 \log^2 m/\Gamma)$*, in expectation.*

Proof: The statement regarding the existence of the respective PICs follows directly from Lemma 11 (in conjunction with Theorem 7) and Lemma 12 (in conjunction with Theorem 7). Recall that we have a requirement of MDS generator matrices to obtain the PIC (with respect to the local chromatic parameter) from the collection of colorings in the proof of Theorem 7. Since for each coloring in any collection we need at most $|V| = m$ colors, we have that a field size larger than m is sufficient for the existence of all required MDS generator matrices in the proof of Theorem 7. Such fields exist, and thus the PIC using the collection of colorings in Lemma 11 (in conjunction with Theorem 7) can be constructed. On the other hand, to utilize the collection of colorings given by Lemma 12 (in conjunction with Theorem 7), we need only the indicator matrices of the colorings, which exist over any field. Thus, this construction too succeeds.

Now we come to the running time. From Remarks 5 and 6, the collections of colorings with properties as in Lemmas 11 and 12 can be constructed in time $O(mn^2 \log m/\Gamma)$, in expectation. From the collection of colorings satisfying properties in Lemma 11, we need to obtain the actual PIC scheme. To do this, we follow the techniques in Lemma 9, which require that the MDS matrices corresponding to the colorings in the collection need to be constructed and stacked. Following Definition 5, the MDS matrix corresponding to any coloring can be constructed in time $O(m)$, once we have a generator matrix for the suitable MDS code. It is easy to see that constructing such a generator matrix has at the most m^2 entries, and hence requires time at most $O(m^2)$ to construct. The time complexity for stacking the MDS matrices of the colorings in the collection is at most $m \log_2 m$, as the number of colorings in the collection as per Lemma 11 is at most $\log_2 m$ and the number of colors in any coloring is at most m. Thus, the total complexity of obtaining the PIC encoding matrix, after obtaining the collection of colorings, is $O(m^4 \log_2 m)$. Hence, the total time complexity for constructing a PIC scheme based on the collection of colorings in Lemma 11 is $O(m^4 n^2 \log^2 m/\Gamma)$. A similar argument can be given for the constructing of the PIC scheme using the collection of colorings in Lemma 12. This completes the proof. ■

Remark 7. *We now compare the results of Theorem 8 for the $t = 1$ case with that of Theorem 4. Theorem 4 shows the existence of a PIC with length $O(\log^2 \Gamma)$ for a t-request PICOD problem with $t = 1$, on a hypergraph with overlap parameter Γ. From Theorem 8, for $t = 1$ and for m, Γ such that $12 \log(6(\Gamma + 1)) < m$, there exists a PIC with length $O(\log \Gamma \log m)$, over a field of size at least m. Clearly, the upper bound from Theorem 8 is larger than the upper bound from Theorem 4, when $m > \Gamma$. This seems surprising, as the result from Theorem 8 uses the local covering parameter of a desirable conflict-free collection, which should ideally be smaller than the (non-local) covering parameter of another desirable conflict-free collection* *used in Theorem 4. The reason for this is essentially because of the difference in the hypergraph partitioning procedures used in the two results. In Theorem 4, the given PICOD hypergraph is partitioned (as per the proof of Lemma 6) into* $P+2$ *subgraphs, where* $P = O(\log(\log \Gamma))$. *We then prove in Lemma 6 that the first subgraph has a conflict-free coloring with* $O(\log^2 \Gamma)$ *colors, while a conflict-free collection with at most* $O(\log^2 \Gamma)$ *colors suffices for coloring all the other subgraphs. Thus, we need* $O(\log^2 \Gamma)$ *colors in total, to color the PICOD hypergraph using a conflict-free collection, as per Theorem 4. Whereas, in Theorem 8, the partitioning (as per Lemma 10, which is invoked through Lemma 11 in the proof of the theorem) results in* $O(\log m)$ *subgraphs, each of which can be then showed (through Lemma 11) to be conflict-free colored by a collection, such that the number of colors in each edge is* $O(\log \Gamma)$ *(for* $t = 1$*). Thus, we obtain a PIC of length* $O(\log \Gamma \log m)$, *using the notion of the local conflict-free covering parameter of this collection, in Theorem 8 for the* $t = 1$ *case. Thus, inspite of using the local covering parameter, the bound on the optimal length obtained in Theorem 8 is larger than that obtained from Theorem 4 (though the actual PIC schemes obtained from both results may have smaller lengths than the bounds we have proved). Reducing the number of subgraphs in the partitioning of Theorem 8 may be possible, invoking similar arguments as in Theorem 4. But it looks difficult to improve upon the* $O(\log^2 \Gamma)$ *bound (as we would continue to rely upon Theorem 1 as we did in Theorem 4). Indeed, proving a 'local' equivalent of Theorem 1, with a smaller bound on the 'local' number of colors in a conflict-free coloring (or collection), might help. We leave this as a possible direction for future research.* ■

Remark 8 (Comparison with the bound from [8])**.** *The work [8] presented a PICOD scheme having length* $O(t \log n + \log^2 n)$ *for the* t*-request problem with* m *messages and* n *clients. The proof of this bound was presented via a constructive argument and does not impose any restrictions on the parameters of the problem. Our bound in Theorem 8,* $\max(O(\log \Gamma \log m), O(t \log m))$, *was accompanied by a randomized polynomial-time algorithm. Our bound applies under some mild conditions on the parameters of the problem, as specified in Theorem 8. However, for the classes of PICOD problems which satisfy these conditions, our bound completely removes the dependency of the PIC length on the number of clients* n. *Moreover, for those classes of PICOD graphs satisfying the conditions of Theorem 8 with* $m < n$, *our bound improves upon the earlier bound, as* $\Gamma \leq n - 1$. *Further, we show in the next subsection that our bound is asymptotically optimal upto a multiplicative factor* $\log t$ *for a special class of graphs.* ■

C. A t-request instance that requires $\Omega(t \log m/\log t)$ length PICOD

In this subsection, we show that the bound in Theorem 8 is asymptotically tight, upto a multiplicative factor of $\log t$. Towards that end, we give the construction of a class of hypergraphs and show that this class of hypergraphs is the (asymptotically) tight example we are looking for. To do this, we need the following general statement on the lower bound for PICOD from [18].

Theorem 9. *[18] Consider a PICOD hypergraph $\mathcal{H} = (V, \mathcal{E})$ corresponding to the t-requests case. Suppose there exists a collection of subsets of $\mathcal{E}$, given by $\{\mathcal{E}_i \subseteq \mathcal{E} : i \in [r]\}$, such that the following condition holds for each $i \leq r-1$: For each $E \in \mathcal{E}_i$, and for any subset $T \subseteq E$ with $|T| = t$, there exists an edge $E' \in \mathcal{E}_{i+1}$, such that $E' \subset E$ and $T \cap E' = \emptyset$. Then, $\ell^{*(t)}(\mathcal{H}) \geq tr$.*

We now give our special class of hypergraphs. For suitable positive integers $r, t,$ and m such that $r = \lfloor (\log m/\log(12t)) \rfloor - 1$ and $t > \log m$, consider the hypergraph $\tilde{\mathcal{H}} = (V, \mathcal{E})$, where $V = [m]$, and $\mathcal{E} = \uplus_{i=1}^{r} \mathcal{E}_i$, where

$$\mathcal{E}_i = \left\{ \left\{ 1, 2, \ldots, \frac{m}{(12t)^{i-1}} \right\}, \left\{ \frac{m}{(12t)^{i-1}} + 1, \ldots, \frac{2m}{(12t)^{i-1}} \right\}, \ldots, \left\{ (m - \frac{m}{(12t)^{i-1}} + 1, \ldots, m \right\} \right\}.$$

The construction of this hypergraph class is illustrated in Fig. 5. We then have the following lemma, which is the main result in this subsection.

Lemma 13. *The hypergraph $\tilde{\mathcal{H}}$ satisfies $\ell^{*(t)}(\tilde{\mathcal{H}}) = O(t \log m)$. Further, it holds that*

$$\ell^{*(t)}(\tilde{\mathcal{H}}) = \Omega(t \log m/\log t).$$

Proof: Notice that $(12t)^{r+1} \leq m \leq (12t)^{r+2}$, by our choice of r. The total number of hyperedges is given by $|\mathcal{E}| = \sum_{i=1}^{r} (12t)^{i-1} = \frac{(12t)^{r-1}-1}{12t-1} < \frac{m}{11t}$. Hence we also have the overlapping parameter $\Gamma \leq |\mathcal{E}| < m/(11t)$.

Since $\Gamma < m/(11t)$, we have the following:

$$\log(6\Gamma + 1) < \log(6m/(11t) + 1) < \log m < t < m$$

Thus,we have that $\max(\log(6\Gamma+1), t) = t$. Since the hyperedges in $\mathcal{E}_r$ are of size $m/(12t)^{r-1} > 12t$, we satisfy all the conditions of Lemma 12, and of Lemma 11 as well. Thus, $\ell^{*(t)}(\tilde{\mathcal{H}}) = O(t \log m)$, by invoking Theorem 8. Finally, observe that the hypergraph $\tilde{\mathcal{H}}$ constructed satisfies the conditions in the Theorem 9. Thus, we get that $\ell^{*(t)}(\tilde{\mathcal{H}}) \geq tr = t\left(\lfloor (\log m/\log(12t)) \rfloor - 1\right) = \Omega(t \log m/\log t)$. This completes the proof. ■

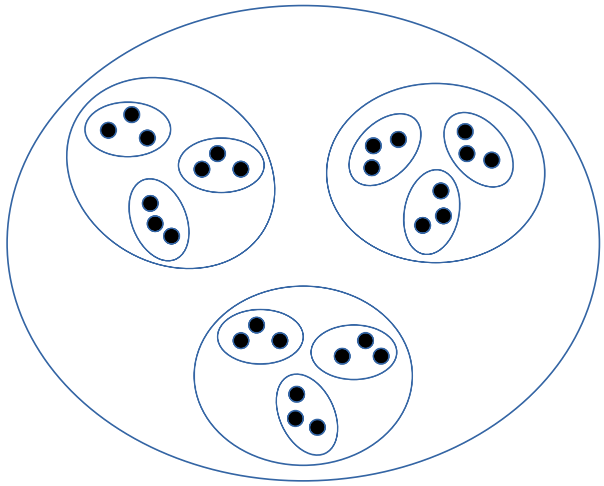

Figure 5: An illustration of how the class of hypergraphs in Subsection V-C is constructed. Each ellipse represents an edge. The largest edge therefore contains all the 27 vertices, and constitutes the collection $\mathcal{E}_1$. Within the largest edge, there are three edges in the figure, which form the sub-collection of edges $\mathcal{E}_2$. Similarly, the collection of nine inner edges forms the collection $\mathcal{E}_3$. Each dark dot represents a vertex. In our construction, there are $12t$ edges within each edge, rather than just three as shown in the figure here.

VI. Extension to k-vector PIC

In this section, we briefly show that the idea of using conflict-free colorings for constructing scalar PICs extends naturally to k-vector PICs as well. Towards this end, we define the notion of k-fold conflict-free coloring of a hypergraph, which generalizes the definition of a conflict-free coloring. To the best of our knowledge, this generalized notion is not available in literature.

Definition 7. *A k-fold coloring of a hypergraph $\mathcal{H} = (V, \mathcal{E})$ is an assignment of k-sized subsets of $[L]$ to the vertices V, given by $C : V \to \binom{[L]}{k}$. A k-fold coloring is conflict-free for edge $E \in \mathcal{E}$ if there exists some $v \in E$ such that $C(v) \cap C(v') = \emptyset$, for each $v' \in E \backslash v$. A coloring C is a k-fold conflict-free coloring for $\mathcal{H}$ if C is a k-fold conflict-free coloring for each edge in $\mathcal{E}$. We define the k-fold conflict-free chromatic number of $\mathcal{H}$ as the smallest L such that a k-fold conflict-free coloring of $\mathcal{H}$ exists as defined above, and denote it by $\chi_{k,CF}(\mathcal{H})$.*

Observe that $\chi_{1,CF}(\mathcal{H}) = \chi_{CF}(\mathcal{H})$. Fig. 6 gives an example of 1-fold and 2-fold conflict-free coloring. Clearly, $\chi_{k,CF}(\mathcal{H}) \leq k\chi_{CF}(\mathcal{H})$ as we can always obtain a k-fold conflict-free coloring

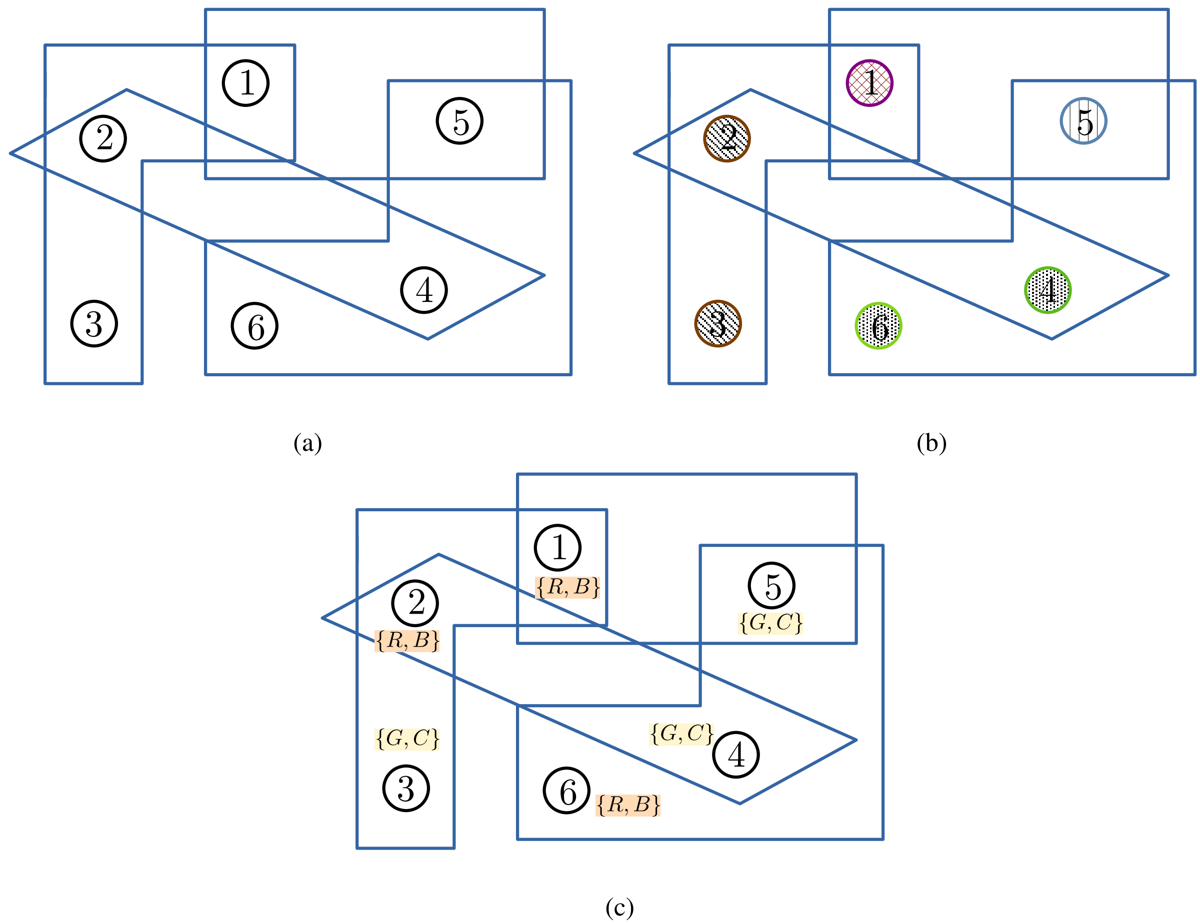

Figure 6: Figure (a) shows a hypergraph $\mathcal{H}$ with 6 vertices and edge set $\mathcal{E} = \{\{1,2,3\},\{1,5\},\{2,4\},\{4,5,6\}\}$. Figure (b) represents a 1-fold conflict-free coloring with 4 colors, with the color classes $\{1\},\{2,3\},\{5\},\{4,6\}$. Figure (c) shows a 2-fold conflict-free coloring using the colors $\{R,G,B,C\}$.

---

from a 1-fold conflict-free coloring by expanding each color into k unique colors. However, we show an example here for which this inequality is strict.

Example 8. *Consider the hypergraph $\mathcal{H}$ given by vertex set $V = \{a,\ldots,e\}$ and $\mathcal{E} = \{\{a,b\},\{b,c\},\{c,d\},\{d,e\},\{e,a\}\}$. Consider any 1-fold coloring of this graph. It is easy to see that two colors are not sufficient to give a 1-fold conflict-free coloring. It is also easy to find a conflict-free coloring with 3 colors, for instance, give color 1 to vertices $\{a,c\}$, color 2 to $\{b,d\}$ and color 3 to vertex e. Thus $\chi_{1,CF}(\mathcal{H}) = 3$.*

Similarly, we can show that there cannot be a 2-fold conflict-free coloring with 4 colors. Now consider the following 2-fold coloring with 5 colors denoted by $\{1,\ldots,5\}$. Let set $\{1,2\}$ be *assigned to vertex* a, $\{3,4\}$ *to* b, $\{5,1\}$ *to* c, $\{2,3\}$ *to* d *and* $\{4,5\}$ *to* e. *It is easy to check that this is a* 2*-fold conflict-free coloring. Thus,* $\chi_{2,CF}(\mathcal{H}) = 5 < 6 = 2(\chi_{1,CF}(\mathcal{H}))$. ■

We now define the indicator matrix of k-fold coloring of $\mathcal{H}$, which leads to a k-vector PIC achievability scheme for $\mathcal{H}$.

Definition 8. *Let* $C : V \to \binom{[L]}{k}$ *denote a* k*-fold coloring of* $\mathcal{H}(V, \mathfrak{I})$. *Let* $C(i) = \{C_{i,1}, \dots, C_{i,k}\}$ *denote the subset assigned to the vertex* $i \in V = [m]$. *Consider a standard basis of the* L*-dimensional vector space over* $\mathbb{F}$, *denoted by* $\{e_1, \dots, e_L\}$. *Now consider the* $L \times mk$ *matrix* G *(with columns indexed as* $\{G_{i,j} : i \in [m], j \in [k]\}$*) constructed as follows.*

- *For each* $i \in [m], j \in [k]$, *column* $G_{i,j}$ *of* G *is fixed to be* $e_{C_{i,j}}$.

We call G *as the indicator matrix associated with the coloring* C.

Using Definition 8, we have the following achievability result.

Theorem 10. $\ell^*_k(\mathcal{H}) \leq \chi_{k,CF}(\mathcal{H})$.

Proof: Let G denote the indicator matrix associated with a k-fold conflict-free coloring C as defined in Definition 8. Let $C(i) = \{C_{i,1}, \dots, C_{i,k}\}$ be the set assigned to vertex i.

Since C is conflict-free, any edge I_r of $\mathcal{H}$ has a vertex d such that $C(d) \cap C(i) = \emptyset, \forall i \in I_r \setminus d$. Then, we have $\{e_{C_{d,j}} : j \in [k]\} \cap \{e_{C_{i,j}} : j \in [k]\} = \emptyset$, for any $i \in I_r \setminus d$. This also means $\mathsf{span}(\{e_{C_{d,j}} : j \in [k]\}) \cap \mathsf{span}(\{e_{C_{i,j}} : i \in I_r \setminus d, j \in [k]\}) = \{\mathbf{0}\}$, since the vectors involved are basis vectors. Further, as $|\{C_{d,j} : j \in [k]\}| = k$, hence $\mathsf{dim}(\mathsf{span}(\{e_{C_{d,j}} : j \in [k]\})) = k$. Thus, G satisfies receiver r by Lemma 1. As r is arbitrary, G is a valid encoder for $\mathcal{H}$. By definition of $\chi_{k,CF}(\mathcal{H})$, the proof is complete. ■

The conflict-free collection can similarly be generalized from Definition 3.

Definition 9 (k*-fold conflict-free collection,* k*-fold conflict-free covering number*)**.** *Let* $\mathcal{H} = (V, \mathcal{E})$ *be a hypergraph. Let* $\mathfrak{C} = \{C^1, \dots, C^P\}$ *where each* $C^p : V \to \binom{[L_p]}{k}$ *is a* k*-fold colorings of the hypergraph* $\mathcal{H}$. *We say* $\mathfrak{C}$ *is a* conflict-free collection of k-fold colorings of $\mathcal{H}$, *if there exists a collection of* P *subgraphs* $\mathcal{H}_p : p \in [P]$ *such that* $\mathcal{H} = \cup_{p \in [P]} \mathcal{H}_p$ *and the coloring* C^p *(when restricted to vertices in* $\mathcal{H}_p$*) is a* k*-fold conflict-free coloring for* $\mathcal{H}_p$.

The minimum value of the sum $\sum_{p=1}^{P} L_p$ *over all possible collections* $\mathfrak{C}$ *(over all* P*) as defined above, is called the* k-fold conflict-free covering number of $\mathcal{H}$ *denoted by* $\alpha_{k,CF}(\mathcal{H})$.

Note that $\alpha_{1,CF}(\mathcal{H}) = \alpha_{CF}(\mathcal{H})$. The notion of local conflict-free chromatic number can also be naturally extended to its k-fold version.

Definition 10 (Local k-fold conflict-free chromatic number)**.** *Given a hypergraph $\mathcal{H}(V, \mathcal{E})$, the local k-fold conflict-free chromatic number of $\mathcal{H}$ is given by*

$$\Delta_k(\mathcal{H}) = \min_{\substack{C:C \text{ is a } k\text{-fold CF} \\ \text{coloring of } \mathcal{H}}} \underbrace{\max_{E\in\mathcal{E}} |\bigcup_{v\in E} C(v)|}_{\Delta_{k,C}(\mathcal{H})}.$$

For convenience, we define $\Delta_{k,C}(\mathcal{H}) = \max_{E\in\mathcal{E}} |\bigcup_{v\in E} C(v)|$, where C is a k-fold conflict-free coloring of $\mathcal{H}$. Therefore, $\Delta_k(\mathcal{H}) = \min_C \Delta_{k,C}(\mathcal{H})$, where the minimum is over all such k-fold colorings C of $\mathcal{H}$.

Similarly, we can define the local k-fold conflict-free covering number of $\mathcal{H}$, denoted by $\lambda_k(\mathcal{H})$, extending Definition 6. Using similar arguments as in Section III, we get the bounds in Theorem 11 below. We leave the details to the reader.

Theorem 11. $\ell_k^*(\mathcal{H}) \overset{(a)}{\leq} \lambda_k(\mathcal{H}) \overset{(b)}{\leq} \min(\Delta_k(\mathcal{H}), \alpha_{k,CF}(\mathcal{H})) \overset{(c)}{\leq} \chi_{k,CF}(\mathcal{H})$. *Further, $\beta_k(\mathcal{H}) \leq k\beta_1(\mathcal{H})$ where β_k refers to any of the parameters $\lambda_k, \Delta_k, \alpha_{k,CF}, \chi_{k,CF}$.*

VII. Discussion

We have presented a hypergraph coloring framework for the pliable index coding problem. We have presented new upper bounds on the optimal PICOD length for the PICOD problem and the t-request PICOD problem via probabilistic methods. These results are accompanied by easy-to-implement randomized algorithms. Our results provide new perspectives on the design of low-length PICOD schemes, which in some cases are near-optimal. Our randomized algorithms can be derandomized using existing techniques. However, such deterministic algorithms may be cumbersome to implement. It would be interesting to give simpler deterministic polynomial-time algorithms for the same. Explicit algorithms for k-vector pliable index coding which give non-trivial improvements over simple extensions of scalar index codes would certainly be interesting. Finally, more investigation is needed into the gaps between the parameters presented in this work and the optimal PICOD length.

References

[1] P. Krishnan, R. Mathew, and S. Kalyanasundaram, “Pliable index coding via conflict-free colorings of hypergraphs,” in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 214–219.

[2] Y. Birk and T. Kol, “Coding on demand by an informed source (iscod) for efficient broadcast of different supplemental data to caching clients,” *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2825–2830, 2006.

[3] Z. Bar-Yossef, Y. Birk, T. S. Jayram, and T. Kol, “Index coding with side information,” *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1479–1494, 2011.

[4] K. Shanmugam, A. G. Dimakis, and M. Langberg, “Local graph coloring and index coding,” in *2013 IEEE International Symposium on Information Theory*, 2013, pp. 1152–1156.

[5] A. Blasiak, R. Kleinberg, and E. Lubetzky, “Broadcasting with side information: Bounding and approximating the broadcast rate,” *IEEE Transactions on Information Theory*, vol. 59, no. 9, pp. 5811–5823, 2013.

[6] H. Maleki, V. R. Cadambe, and S. A. Jafar, “Index coding—an interference alignment perspective,” *IEEE Transactions on Information Theory*, vol. 60, no. 9, pp. 5402–5432, 2014.

[7] S. Brahma and C. Fragouli, “Pliable index coding,” *IEEE Transactions on Information Theory*, vol. 61, no. 11, pp. 6192–6203, 2015.

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

[9] T. Liu and D. Tuninetti, “Tight information theoretic converse results for some pliable index coding problems,” *IEEE Transactions on Information Theory*, vol. 66, no. 5, pp. 2642–2657, 2020.

[10] S. Sasi and B. S. Rajan, “Code construction for pliable index coding,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 527–531.

[11] L. Ong, B. N. Vellambi, and J. Kliewer, “Optimal-rate characterisation for pliable index coding using absent receivers,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 522–526.

[12] L. Ong, B. N. Vellambi, J. Kliewer, and P. Sadeghi, “Improved lower bounds for pliable index coding using absent receivers,” 2019.

[13] T. Liu and D. Tuninetti, “Secure decentralized pliable index coding,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1729–1734.

[14] ——, “Decentralized pliable index coding,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 532–536.

[15] L. Song, C. Fragouli, and T. Zhao, “A pliable index coding approach to data shuffling,” *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1333–1353, 2020.

[16] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky, “Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks,” *SIAM Journal on Computing*, vol. 33, no. 1, pp. 94–136, 2003.

[17] S. Smorodinsky, “Conflict-free coloring and its applications,” in *Geometry—Intuitive, Discrete, and Convex*. Springer, 2013, pp. 331–389.

[18] T. S. B., V. Subramanian, and P. Krishnan, “Bounding the optimal length of pliable index coding via a hypergraph-based approach,” 2022. [Online]. Available: https://arxiv.org/abs/2208.10389

[19] J. Pach and G. Tardos, “Conflict-free colourings of graphs and hypergraphs,” *Combinatorics, Probability and Computing*, vol. 18, no. 5, pp. 819–834, 2009.

[20] T. S. B., V. Subramanian, and P. Krishnan, “Bounding the optimal length of pliable index coding via a hypergraph-based approach,” in *2022 IEEE Information Theory Workshop (ITW)*, 2022, pp. 696–701.

[21] P. Erdős and L. Lovász, “Problems and results on 3-chromatic hypergraphs and some related questions,” *Infinite and finite sets*, vol. 10, pp. 609–627, 1975.

[22] R. A. Moser and G. Tardos, “A constructive proof of the general lovász local lemma,” *J. ACM*, vol. 57, no. 2, Feb. 2010. [Online]. Available: https://doi.org/10.1145/1667053.1667060

[23] K. Chandrasekaran, N. Goyal, and B. Haeupler, "Deterministic algorithms for the Lovász local lemma," *SIAM Journal on Computing*, vol. 42, no. 6, pp. 2132–2155, 2013.

[24] D. G. Harris, "Deterministic algorithms for the Lovász local lemma: simpler, more general, and more parallel," *arXiv preprint arXiv:1909.08065*, 2019.

[25] E. Horev, R. Krakovski, and S. Smorodinsky, "Conflict-free coloring made stronger," in *Algorithm Theory - SWAT 2010*, H. Kaplan, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 105–117.

[26] M. Mitzenmacher and E. Upfal, *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge Univ Pr, 2005.

APPENDIX A
TOOLS FROM PROBABILITY

Below we state the Local Lemma, due to Erdős and Lovász, which is required in some of our proofs.

Lemma 14 (*The Local Lemma*, [21])**.** *Let $A_1, \dots, A_n$ be events in an arbitrary probability space. Suppose that each event A_i is mutually independent of a set of all the other events A_j but at most d, and that $\mathsf{Pr}[A_i] \le p$ for all $i \in [n]$. If $\mathsf{e}p(d+1) \le 1$, then $\mathsf{Pr}[\cap_{i=1}^{n} \overline{A_i}] > 0$.*

Moser and Tardos [22] demonstrated an algorithmic version of the Local Lemma. They showed the following.

Theorem 12 (Algorithmic Local Lemma [22])**.** *Let $\mathcal{P}$ be a finite set of mutually independent random variables in a probability space. Let $A_1, \dots, A_n$ be events that are determined by these variables. Suppose that each event A_i is mutually independent of a set of all the other events A_j but at most d, and that $\mathsf{Pr}[A_i] \le p$ for all $i \in [n]$. If*

$$\mathsf{e}p(d+1) \le 1,$$

then there exists an assignment of the random variables in $\mathcal{P}$ such that none of the events A_i occur. Moreover, there is a randomized algorithm (described below) that finds such an assignment, that uses at most n/d resampling steps in expectation.

The paper [22] states the result and algorithm in a general setting. To avoid clutter, we state a specific symmetric case which suffices our requirements.

Algorithmic Local Lemma

First, sample the random variables in $\mathcal{P}$ as per the distribution. If at least one of the events A_i occur, choose an arbitrary A_j that occurs. Then resample the random variables in $\mathcal{P}$ that determine A_j. Repeat this until none of the events A_i occur.

Finally, we need the following Chernoff-type bound for some of our proofs.

Theorem 13 (Chernoff Bound, Corollary 4.6 in [26]). *Let* $X_1, \ldots, X_n$ *be independent Poisson trials such that* $\mathsf{Pr}[X_i] = p_i$. *Let* $X = \sum_{i=1}^{n} X_i$ *and* $\mu = E[X]$. *For* $0 < \delta < 1$,

$$\mathsf{Pr}[|X - \mu| \geq \delta\mu] \leq 2\mathsf{e}^{-\mu\delta^2/3}.$$