

Index Coding With Side Information

Ziv Bar-Yossef, Yitzhak Birk, *Senior Member, IEEE*, T. S. Jayram, and Tomer Kol, *Member, IEEE*

Abstract—Motivated by a problem of transmitting supplemental data over broadcast channels (Birk and Kol, INFOCOM 1998), we study the following coding problem: a sender communicates with n receivers R_1, \dots, R_n . He holds an input $x \in \{0, 1\}^n$ and wishes to broadcast a single message so that each receiver R_i can recover the bit x_i . Each R_i has prior side information about x , induced by a directed graph G on n nodes; R_i knows the bits of x in the positions $\{j \mid (i, j) \text{ is an edge of } G\}$. G is known to the sender and to the receivers. We call encoding schemes that achieve this goal INDEX codes for $\{0, 1\}^n$ with side information graph G . In this paper we identify a measure on graphs, the *minrank*, which exactly characterizes the minimum length of linear and certain types of nonlinear INDEX codes. We show that for natural classes of side information graphs, including directed acyclic graphs, perfect graphs, odd holes, and odd anti-holes, minrank is the optimal length of arbitrary INDEX codes. For arbitrary INDEX codes and arbitrary graphs, we obtain a lower bound in terms of the size of the maximum acyclic induced subgraph. This bound holds even for randomized codes, but has been shown not to be tight.

Index Terms—Broadcast channels, code length, error correction coding, information cost.

I. INTRODUCTION

A. Background and Problem Statement

SOURCE coding is one of the central areas of coding and information theory. Shannon's famous source coding theorem states that the average number of bits necessary and sufficient to encode a source is equal (up to one bit) to the entropy of the source. In many distributed applications, though, the receiver may have some prior *side information* about the source message, before it is sent. Source coding with side information addresses encoding schemes that exploit the side information in order to reduce the length of the code. Classical results in this area [1]–[3] describe how to achieve optimal rates with respect to the joint entropy of the source and the side information.

Zero-error codes, without side information, were discussed by Shannon [4] as early as 1948. Witsenhausen [5] initiated the study of the zero-error side information problem. For every

source input $x \in \mathcal{X}$, the receiver gets an input $y \in \mathcal{Y}$ that gives some information about x . This is captured by restricting the pairs (x, y) to belong to a fixed set $\mathcal{L} \subseteq \mathcal{X} \times \mathcal{Y}$. Both the sender and the receiver know \mathcal{L} , and thus each of them, given his own input, has information about the other's input. Witsenhausen showed that fixed-length side information codes are equivalent to colorings of a related object called the *confusion graph*, and thus the logarithm of the chromatic number of this graph tightly characterizes the minimum number of bits needed to encode the source. Further results by Alon and Orlitsky [6] and Koulgi *et al.* [7] showed that graph-theoretic information measures could be used to characterize both the average length of variable-length codes, as well as asymptotic rates of codes that simultaneously encode multiple inputs drawn from the same source.

In this paper, we study a new variant of source coding with side information, first proposed by Birk and Kol [30], [8] in the context of a server that disseminates a set of data blocks over a broadcast channel to a set of caching clients. Each client possesses in its cache only a subset of the transmitted blocks, due to reception problems, limited storage space, rejection by an interest filter, etc. The client needs a certain subset of the data blocks, yet some of these blocks may be missing from its cache. The client uses a backward channel to request blocks that it needs but has not cached and to advise the server of the blocks it already has in its cache. The challenge is to minimize the amount of supplemental information that must be broadcast by the server in order to enable every client to derive all its requested blocks. See Section II for more details on this problem. In [8], “coding on demand by an informed source” (ISCOD) is proposed: the server is made aware of the requests and the cache contents of every client, constructs an appropriate systematic erasure correcting code, and transmits only the “redundant” data. Each client then uses its cached data along with the supplemental data to derive its requested block. Upper bounds on the amount of supplemental data that must be transmitted are presented, but are shown to be sub-optimal. In the model considered, each client may request multiple blocks (in which case it is equivalently represented as multiple single-request clients, each with the full cache content of the original one). Also, the work focuses on the case wherein any given block is requested by at most one client.

In this paper, we assume block-level erasure; i.e., any given client either possesses an entire error-free block or does not have it at all. This well represents practical situations wherein intra block error detection and correction are used. With this assumption, it suffices to consider blocks of size one bit. (More will be said about this later.) We formalize the above as a source coding with side information problem as follows (cf. [8]). There is a sender (server) who has an input (data) x from a source alphabet $\mathcal{X} = \{0, 1\}^n$. There are n receivers (clients) R_1, \dots, R_n , where

Manuscript received August 25, 2007; revised July 12, 2010; accepted July 22, 2010. Date of current version February 18, 2011. The work of Z. Bar-Yossef was supported by the European Commission Marie Curie International Re-integration Grant. The material in this paper was presented at the 47th IEEE Annual Symposium on Foundations of Computer Science (FOCS), Berkeley, CA, October 2006.

Z. Bar-Yossef and T. Kol are with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel. They are also with Google Haifa Engineering Center, Haifa 20692, Israel (e-mail: zivby@ee.technion.ac.il; tomer.kol@gmail.com).

Y. Birk is with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: birk@ee.technion.ac.il).

T. S. Jayram is with the IBM Almaden Research Center, San Jose, CA 95120 USA (e-mail: jayram@almaden.ibm.com).

Communicated by W. Szpankowski, Associate Editor for Source Coding.

Digital Object Identifier 10.1109/TIT.2010.2103753




Fig. 1. Sample side information graphs influencing encoding and decoding.

for each i , R_i is interested in the bit (block) x_i . The side information is characterized by a simple directed graph G (no self loops or parallel edges) on $\{1, 2, \dots, n\}$. For a subset $S \subseteq [n]$, $x[S]$ denotes the projection of x on the coordinates in S . The side information of R_i (cached blocks) equals $x[N(i)]$, where $N(i) \triangleq \{j \in V \mid (i, j) \text{ is an edge}\}$ denotes the set of out-neighbors of i in the graph G . The sender and the receivers are both assumed to know G .

Remark: the case wherein a given bit may be requested by multiple clients is a topic for further research.

Example 1: For an input $x \in \{0, 1\}^n$, each receiver R_i is interested in the value x_i but knows x_{i-1} as side information. (Abusing notation slightly, receiver R_1 knows x_n .) The side information graph is thus a directed cycle of length n . Since x_{i-1} is “independent” of x_i , it may not be clear at first how the sender can take advantage of the side information of the receivers to shorten the broadcast message. However, there is a strategy in which the sender can save one bit: rather than sending all the bits of x , the sender broadcasts the $n - 1$ parities $x_1 \oplus x_2, x_2 \oplus x_3, \dots, x_{n-1} \oplus x_n$. Now, each receiver R_i for $i > 1$ can recover x_i by taking the parity of $x_{i-1} \oplus x_i$ with x_{i-1} . Finally, receiver R_1 XORs the $n - 1$ parities broadcast by the sender together with x_n to recover x_1 .

Definition 1 (INDEX Codes): A deterministic INDEX code \mathcal{C} for $\{0, 1\}^n$ with side information graph G on n nodes, abbreviated as “INDEX code for G ”, is a set of codewords in $\{0, 1\}^\ell$ together with:

- 1) An encoding function E mapping inputs in $\{0, 1\}^n$ to codewords, and
- 2) A set of decoding functions D_1, D_2, \dots, D_n such that $D_i(E(x), x[N(i)]) = x_i$ for every i .

The length of \mathcal{C} , denoted by $\text{len}(\mathcal{C})$, is defined to be ℓ . Both the encoding and decoding functions depend on G , so the sender typically has to broadcast the graph.

Consider the sample side information graphs depicted in Fig. 1. When encoding the three unrelated bits requested by clients 1, 2, and 3, an encoder lacking knowledge of the side information must send all three bits (the entropy is 3 bits). Knowing G , a possible encoding is $\{x_1 \oplus x_2, x_2, x_3\}$ for the case depicted in Fig. 1(a), and $\{x_1, x_2 \oplus x_3\}$ for case (b). A receiver must know G and the sender’s encoding algorithm in order to utilize the received supplementary transmissions. Receiver 1 in Fig. 1 would not be able to “understand” the two received bits without knowledge of G and the encoding scheme. (It would not be able to tell whether this is case 1(a) or 1(b).) The cost of transmitting the graph is amortized by working with larger blocks. Note that the side-information graph for all the bits in the block is the same (we assume a block is protected by ECC as a single entity and that it is either correctly received or completely lost). In the remainder of this paper we assume

that the encoder and receivers all know G . The question when partial knowledge of G suffices is beyond the scope for the current paper.

This problem can also be cast in an equivalent setting with a single receiver: The receiver is given an index i and the side information $x[N(i)]$ as inputs and wants to recover the value x_i . (The equivalence follows from the fact the sender does not know the index i given to the receiver, and thus has to use an encoding that enables the recovery of x_i , for any i .) Using this equivalent form, we can contrast our side information problem with Witsenhausen’s zero-error side information problem. A first notable difference is that while in Witsenhausen’s setting the entire input x has to be recovered, in our setting only a single bit x_i is needed. This allows significant savings in the encoding length, as the following example demonstrates: suppose the side information graph is a perfect matching on n nodes (here, and below, when for simplicity we give an example of an undirected side information graph G' , it represents a directed graph G with the same set of vertices s.t. for each $(i, j) \in G'$ there are edges $(i, j) \in G$ and $(j, i) \in G$). Since the receiver has only a single bit of side information, then $n - 1$ bits are necessary to recover the entire input. If, however, only a single bit is needed, then the sender can encode his input by the $n/2$ parities of pairs of matched bits. A second difference from Witsenhausen’s setting is that the type of side information addressed in our problem is restricted to side information graphs. This natural restriction emanates from the broadcast application mentioned above and also imposes more structure that enables us to obtain an interesting combinatorial characterization of the minimum length of INDEX codes in terms of the side information graphs.

We also consider in this paper randomized INDEX codes, in which the encoding and decoding functions are allowed to be randomized and are even allowed to use a common public random string. Decoding needs to succeed only with *high probability*, taken over the random choices made by the encoding and decoding functions.

B. Our Contributions

In this paper we identify a graph functional, called *minrank*, which we show to characterize the minimum length of INDEX codes, for natural types of codes and for wide classes of side information graphs. Let G be a directed graph on n vertices without self-loops. We say that a 0–1 matrix $A = (a_{ij})$ fits G if for all i and j : (i) $a_{ii} = 1$, and (ii) $a_{ij} = 0$ whenever (i, j) is not an edge of G . Thus, $A - I$ is the adjacency matrix of an edge subgraph of G , where I denotes the identity matrix. Let $\text{rk}_2(\cdot)$ denote the 2-rank of a 0–1 matrix, namely, its rank (the maximum number of linearly independent rows) over the field $GF(2)$.

Definition 2: $\text{minrk}_2(G) \triangleq \min \{\text{rk}_2(A) : A \text{ fits } G\}$.

The above measure for undirected graphs was considered by Haemers [9] in the context of proving bounds for the Shannon capacity Θ of undirected graphs. (The Shannon capacity $\Theta(G)$ of a graph G is defined as $\lim_{k \rightarrow \infty} \alpha(G^k)^{1/k}$, where G^k denotes the (strong) k th power of G and $\alpha(G)$ the independence number of G – the size of largest independent set.) For an undirected graph G whose adjacency matrix is M , the 2-rank of




Fig. 2. Sample side information graph G and a matrix A that fits G .

$M + I$ (which fits G) has also been studied in the algebraic graph theory community. For example, Brouwer and van Eijl [10] and Peeters [11] study this quantity for strongly regular and distance-regular graphs, respectively. It has been shown by Peeters [12] that computing $\text{minrk}_2(G)$ is NP-hard. Finally, it is known that minrk_2 has the “sandwich property,” similar to other natural quantities such as the Lovász Theta function.

Proposition 1 ([9], [13]): For any undirected graph G , $\omega(\overline{G}) \leq \Theta(G) \leq \text{minrk}_2(G) \leq \chi(\overline{G})$, where \overline{G} is the complement of G and $\omega(\cdot)$, $\Theta(\cdot)$, and $\chi(\cdot)$ are, respectively, the clique number (size of the largest clique in the graph), the Shannon capacity, and the chromatic number (the smallest number of colors needed for coloring the graph with no two neighboring vertices having the same color). Moreover, each of these inequalities is strict.

Our first result (see Section III) shows that $\text{minrk}_2(G)$ completely characterizes the minimum length of *linear* INDEX codes (i.e., ones whose encoding function is linear), for *arbitrary* directed side information graphs G .

Theorem 1: For any side information graph G , there exists a linear INDEX code for G whose length equals $\text{minrk}_2(G)$. This bound is optimal for all linear INDEX codes for G .

Example 2: Consider the side information graph G and the matrix A that fits G , both depicted in Fig. 2. $\text{minrank}_2(A) = 2$ and we can derive from A a 2-row matrix $A' = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 \end{pmatrix}$. Using A' we can construct an optimal linear code

$$\text{codeword to be sent} = A'x = \begin{pmatrix} x_0 \oplus x_1 \oplus x_4 \\ x_1 \oplus x_2 \oplus x_3 \end{pmatrix}$$

wherein x is the original input vector.

This bound strictly improves a previous upper bound of Birk and Kol [8]. Birk and Kol showed a construction of a linear INDEX code, whose length is the “cover cost” of the side information graph (and showed that the construction is suboptimal). For undirected graphs, the cover cost is the same as the chromatic number of the complement graph. Since the minrank can be strictly smaller than this chromatic number, it immediately follows that the minrank bound beats the Birk and Kol bound. The lower bound for linear codes is of interest, since linear codes are possibly the most natural type of codes. In fact, all the existing INDEX codes (with or without side information) we are aware of are linear.

In Section IV we prove that $\text{minrk}_2(G)$ characterizes not only the optimal length of linear codes, but also the optimal length of a wide class of *nonlinear* codes. An INDEX code is called *linearly decodable*, if all its n decoding functions are linear. A linearly decodable code need not be linearly encodable. A

simple argument shows that the length of a linearly decodable INDEX code for any graph G is at least $\text{minrk}_2(G)$. We relax the notion of linearly decodable codes to “semilinearly decodable” codes. An INDEX code is *k-linearly decodable* if at least k of its decoding functions are linear. Note that n -linearly decodable codes are simply linearly decodable, while 0-linearly decodable codes are unrestricted. We are able to prove that $\text{minrk}_2(G)$ is the optimal length of k -linearly decodable codes when $k \geq n - 2$:

Theorem 2: For any graph G , and for any $k \geq n - 2$, the length of any k -linearly decodable INDEX code for G is at least $\text{minrk}_2(G)$.

As our new linear INDEX code (proof of Theorem 1 in Section III) is also linearly decodable (and thus k -linearly decodable, for any k), the bound in Theorem 2 is tight. Note, however, that this result is mostly of theoretical value in view of the restrictions on k .

Our third contribution is a lower bound that holds for general INDEX codes including deterministic and randomized INDEX codes. This result is presented in Section V.

Theorem 3: The length of any δ -error randomized INDEX code for G is at least $\text{MAIS}(G) \cdot (1 - H_2(\delta))$, where $\text{MAIS}(G)$ is the size of the maximum acyclic induced subgraph of G and $H_2(\cdot)$ is the binary entropy function.

This lower bound immediately gives a tight bound for directed acyclic graphs and undirected graphs G that satisfy $\omega(\overline{G}) = \text{minrk}_2(G) = \chi(\overline{G})$. In particular, it holds for perfect graphs.¹ In Section V-D.I, we are able to prove that minrank characterizes the minimum length of INDEX codes, even for nonperfect graphs, namely *odd holes* (undirected odd-length cycles of length at least 5) and *odd anti-holes* (complements of odd holes).

Theorem 4: Let G be any graph, which is either a DAG, a perfect graph, an odd hole, or an odd anti-hole. Then, the length of any INDEX code for G is at least $\text{minrk}_2(G)$.

This theorem implies that our lower bound for general codes (Theorem 3) is not tight. For odd holes, $\text{MAIS}(G)$ is the size of the largest independent set, i.e., $\omega(\overline{G})$, which we show to be strictly smaller than $\text{minrk}_2(G)$.

The Strong Perfect Graph Theorem (conjectured by Berge [14] and proved by Chudnovsky *et al.* [15]) states that a graph is perfect if and only if it contains no (induced) odd hole or odd anti-hole. It follows that every undirected graph can be partitioned into induced subgraphs, each of which is either perfect, an odd hole, or an odd anti-hole. This motivated us to study the following direct sum-type problem: if a graph G can be partitioned into k induced sub-graphs G_1, \dots, G_k , then is the length of the best INDEX code for G equal to the sum of the lengths of the best codes for G_1, \dots, G_k ? While we believe the answer to this general question to be negative, we were able to prove it for the case wherein G_1, \dots, G_k are disconnected components (i.e., there is no edge connecting G_i and G_j , for any $i \neq j$). A direct proof of this result seems to be elusive. In fact, an argument based on the techniques of Feder *et al.* [16] incurs a loss of an

¹Recall that an undirected graph G is called *perfect*, if for every induced sub-graph G' of G , $\omega(\overline{G'}) = \chi(G')$. Perfect graphs include a wide class of graphs such as trees, bipartite graphs, interval graphs, chordal graphs, etc.

additive term that depends linearly on k . After lower bounding the length of a code by its *information cost* [17], [18], we are able to prove a tight direct sum theorem w.r.t. the information cost measure. We note that almost all our lower bounds hold not only for the length of INDEX codes but also for their information cost. This result is presented in Section V.

C. Techniques

We resort to a multitude of techniques from linear algebra, information theory, Fourier analysis, and combinatorics to prove the results presented in this paper.

The lower bounds for linearly encodable and linearly decodable codes are based on dimension arguments from linear algebra. To extend the lower bound for linearly decodable codes to semilinearly decodable codes, we used an intriguing “balance property” of Boolean functions: if all linear Boolean functions are “balanced” on some set U (i.e., get the same number of 0’s and 1’s on the set), then all Boolean functions (whether linear or not) are balanced on U . To prove this property, we use Fourier analysis to represent arbitrary Boolean functions as linear combinations of linear functions. We then introduce the notion of “*minimum dimension*”, which is dual to *minrank*, and explore its properties using the balance property. This in turn allows us to extend the lower bound for linearly decodable codes to $(n - 2)$ -linearly decodable codes.

The lower bound for general (randomized) codes and the direct sum theorem are proved via information theory arguments. We extend previous arguments from [18], [19] to obtain a direct sum theorem for the *information cost* of codes.

Finally, our lower bounds for odd holes and odd anti-holes are purely combinatorial. We employ a connection between vertex covers of a graph G and the structure of the confusion graph corresponding to INDEX codes for G . We note that dealing with odd holes, and with the pentagon in particular, turned out to be very challenging, because the standard technique of lower bounding the chromatic number of the corresponding confusion graph via its independence number does not work.

D. Related Work

Prior Art: There are settings other than source coding in which INDEX codes have been addressed. Ambainis *et al.* [20] considered what they called “random access codes”, which are identical to randomized INDEX codes without side information. Their main thrust was proving tight bounds on the length of the codes in the quantum setting, where inputs can be encoded by qubits rather than classical bits; their result applied to the classical setting is a special case of our Theorem 3 for the case when G is the empty graph.

The problem of INDEX coding with side information can also be cast as a one-way communication complexity problem of the *indexing* function [21] (from which the term INDEX codes was coined) with the additional twist of side information. Alice (the sender) is given an input x and sends a single message to Bob. Bob is given an index i and the side information $x[N(i)]$, and uses Alice’s message to learn x_i . Another formulation of INDEX coding is in terms of *network coding* [22], [23]. As such, it represents a restricted case of a single source, a single encoder and

a single channel, but with the important addition of a special flavor of side information. Parts of this information are known to different decoders, and the encoder is fully aware of this knowledge.

Subsequent Work: Following the publication of the extended abstract of this work [24], Lubetzky and Stav [25] were able to make remarkable progress and prove that there could be an unbounded gap between $\text{minrk}_2(G)$ and the length of an optimal INDEX code for G . They achieved this by constructing a new family of Ramsey graphs. It thus remains an open problem to find an exact characterization of the optimal length of INDEX codes for general codes and arbitrary graphs. It should nonetheless be noted that the commonly used codes are linear, for which the bounds presented in the current paper are tight. The current paper formally addresses single-bit transmissions. When multi-bit packets are transmitted, this corresponds to encoding the different bits of the packets independently (scalar codes), albeit sharing the same side-information graph. El Rouayheb, Sprintson and Georgiades [26] explored jointly encoding multi-bit packets and showed that vector linear codes can outperform scalar linear codes.

Notation: Throughout the paper, we use the following notation. Let $[n]$ denote the set $\{1, 2, \dots, n\}$. Let e_i denote the i th standard basis vector. The dimension of this vector is understood from the context. For a subset $S \subseteq [n]$, we denote by $x[S]$ the projection of a vector $x \in \{0, 1\}^n$ on the coordinates in S .

II. MOTIVATING APPLICATION

Many important data dissemination systems employ a broadcast channel at the physical level. Prominent examples include satellite, terrestrial wireless systems, and systems employing coaxial cable. Broadcast channels are frequently used to disseminate high volumes of media-rich content, such as movies, episodes of TV series, and “multimedia” newspapers containing images and video clips.

The broadcast channel is extremely efficient for sending the same data to a large number of receivers, but its use for sending different data to different users is wasteful. For the case wherein users (may) need the same data but at different times (e.g., on-demand viewing of a “hot” movie), service providers can speculatively “push” data to clients; upon demand by a client, the already present data is presented immediately, as if it were sent on demand. XTV by News Data Systems² is an early example.

The above example demonstrates how abundant client storage capacity can be used in lieu of true communication capability in order to increase perceived communication-related quality of service. With the rapid decline in the cost of nonvolatile storage (disk drives and Flash memory), it is interesting to look for additional ways of exploiting its abundance in order to reduce demand on less abundant or more expensive resources. We next describe such an application, which has motivated the current work.

Despite the use of a broadcast channel, not all information is received by all clients. This may result from intermittent connectivity due to bad weather, a power outage at some receiver

²<http://www.nds.com>

locations, intermittent reception by mobile receivers, or due to equipment being temporarily switched off. Finally, even data that is received by a given client may be discarded, be it for lack of space or by an “interest filter”.

Following the broadcast of various content, arranged in blocks, each client thus typically has in its local storage some subset of the transmitted blocks. (An erasure channel is assumed at block granularity, so a client either possesses all bits of a block or none of them.) As for the remainder of this data, a given client may request some of it while not being interested in the rest. The question is then how to use a broadcast channel, which is very effective at sending the same data to all recipients, in order to efficiently send different *supplemental* data to the different clients.

Whenever each client stores a substantial fraction of the transmitted data while only requesting a single block, the probability that a pair of clients each have the block requested by the other is much higher than the probability that they both request the same block. This gave rise to the idea of using source coding for this purpose [8]: each client uses a (slow) “upstream” channel to request blocks it has not cached and to advise the server of the blocks it already has in its cache (for sufficiently large blocks, the amount metadata is negligible); the server then uses this information to derive the required supplemental information and broadcasts it. Finally, note that unlike in conventional multicast settings, it is not required that every client be able to derive all the originally transmitted blocks.

III. LINEAR CODES

In this section we obtain a tight characterization of the length of linear INDEX codes for all side information graphs G .

Theorem 1 (Restated): For any side information graph G , there exists a linear INDEX code for G whose length equals $\text{minrk}_2(G)$. This bound is optimal for all linear INDEX codes for G .

Proof: Let A be the matrix that fits G whose 2-rank equals $\text{minrk}_2(G) \triangleq k$. Assume without loss of generality that the span of the first k rows A_1, \dots, A_k equals the span of all the rows of A . The encoding function is simply the k bits $b_j \triangleq A_j \cdot x$ for $1 \leq j \leq k$.

Decoding proceeds as follows. Fix a receiver R_i for some $i \in [n]$ and let $A_i = \sum_{j=1}^k \lambda_j A_j$ for some choice of λ_j 's. The receiver first computes $A_i \cdot x = \sum_{j=1}^k \lambda_j b_j$ using the k -bit encoding of x . Now, consider the vector $c_i = A_i - e_i$, where e_i is the i th standard basis vector. Observe that the only nonzero entries in c_i correspond to coordinates that are among the neighbors of i in G . This means that the receiver can compute $c_i \cdot x$ using the side information. Receiver R_i can now recover x_i via $(A_i \cdot x) - (c_i \cdot x) = e_i \cdot x = x_i$.

For the lower bound, suppose \mathcal{C} is an arbitrary linear INDEX code for G defined by the set $S = \{u_1, u_2, \dots, u_k\}$, i.e., x is encoded by the taking its inner product with each vector in S .

Claim 1: For every i , e_i belongs to the span of $S \cup \{e_j : j \in N(i)\}$.

Before we prove the claim, we show how to finish the proof of the lower bound. For each $i \in [n]$, the claim shows that

$e_i = \sum_{j=1}^k \lambda_j u_j + \sum_{j \in N(i)} \mu_j e_j$, for some choice of λ and μ . Rearranging, we have $\sum_{j=1}^k \lambda_j u_j = e_i - \sum_{j \in N(i)} \mu_j e_j \triangleq A_i$. It follows that A_i has value 0 in coordinates outside $N(i) \cup \{i\}$, A_i has value 1 in its i th coordinate, and A_i belongs to the span of S . Therefore, the matrix A whose rows are given by A_1, A_2, \dots, A_n fits G and has rank at most k . We conclude that $k \geq \text{rk}_2(A) \geq \text{minrk}_2(G)$.

It remains to prove the claim. Fix an i and suppose to the contrary that e_i is *not* in the subspace W spanned by the vectors in $S \cup \{e_j : j \in N(i)\}$. Recall that the *dual* of W , denoted by W^\perp , consists of the set of vectors orthogonal to every vector in W , i.e., $W^\perp = \{v : v \cdot w = 0 \text{ for all } w \in W\}$. It is well-known that $W^{\perp\perp} = W$. Therefore, the assumption $e_i \notin W$ implies

that there is a vector $x \in W^\perp$ such that $x \cdot e_i \stackrel{(*)}{\neq} 0$. On the other hand, since $x \in W^\perp$, we have that x is orthogonal to every vector in $S \cup \{e_j : j \in N(i)\}$. It follows that (i) the encoding for x equals 0^k , and (ii) the side information x_j available to receiver R_i equals 0 for all $j \in N(i)$. This violates the correctness of the encoding because the input 0^n also satisfies (i) and (ii), yet Equation (*) shows that it differs from x in coordinate i . ■

IV. SEMILINEARLY DECODABLE CODES

In this section, we show that $\text{minrk}_2(G)$ is a lower bound on the minimum length of *semilinearly decodable* INDEX codes for arbitrary graphs G .

Let \mathcal{C} be an INDEX code for G . Let D_1, \dots, D_n be the n decoding functions of \mathcal{C} . Fix a codeword $c \in \mathcal{C}$, and for each index $i \in [n]$, we denote by D_i^c the function induced by fixing c as input to D_i : $D_i^c(x[N(i)]) = D_i(c, x[N(i)])$. Although D_i^c is applied only to the side information bits $x[N(i)]$, it will be convenient for us to view it as acting on the whole input x with the restriction that it depends only³ on the set of coordinates $N(i)$. Thus, from now on, $D_i^c : \{0, 1\}^n \rightarrow \{0, 1\}$.

An INDEX code \mathcal{C} is said to be *k-linearly decodable*, if for every codeword $c \in \mathcal{C}$, at least k of the decoding functions D_1^c, \dots, D_n^c are linear. Note that the smaller k is, the less restricted is the class of k -linearly decodable codes. When $k = n$, these codes are simply called *linearly decodable*, while 0-linearly decodable are unrestricted codes. Our upper bound (Theorem 1) is a linearly decodable INDEX code (and thus also k -linearly decodable, for any k).

Our goal is to obtain lower bounds on the length of k -linearly decodable codes for a value of k as small as possible.

Theorem 2 (Restated): For any graph G , and for any $k \geq n - 2$, the length of any k -linearly decodable INDEX code for G is at least $\text{minrk}_2(G)$.

A. Kernel Size

To prove the lower bound, we introduce the notion of *kernel*. The kernel of a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is the set of inputs it maps to 0: $\ker(f) = \{x \mid f(x) = 0\}$. By extension, the kernel of a family of Boolean functions $\mathcal{F} = \langle f_i \mid i \in T \rangle$ (T is some index set) is the set of inputs that are mapped to 0 by all of the functions in the family: $\ker(\mathcal{F}) = \{x \mid \forall i f_i(x) = 0\}$. We

³A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to *depend only* on a set of coordinates $S \subseteq [n]$, if $f(x) = f(y)$ for every two inputs x, y with $x[S] = y[S]$.

next show a connection between the length of INDEX codes and the size of the kernel of a suitably chosen family of functions.

Note that $D_i^c(x) = x_i$ for every x whose encoding $E(x)$ equals c . This can be also written as $D_i^c(x) + e_i \cdot x = 0$. If we view the vector e_i as a linear function operating over $\{0, 1\}^n$, then we can say that x belongs to the kernel of the function $D_i^c + e_i$, i.e., $(D_i^c + e_i)(x) = 0$. As this holds for every i , we conclude the following.

Proposition 2: For every codeword c

$$\begin{aligned} \{x \mid E(x) = c\} \\ \subseteq \ker(D_1^c + e_1, \dots, D_n^c + e_n) = \{x \mid \forall i D_i^c(x) = x_i\}. \end{aligned}$$

We obtain as an immediate corollary the following lower bound on the length of INDEX codes in terms of kernel size.

Proposition 3: If $|\ker(D_1^c + e_1, \dots, D_n^c + e_n)| \leq M$ for every codeword $c \in \mathcal{C}$, then $\text{len}(\mathcal{C}) \geq \lceil n - \log M \rceil$.

Proof: Consider any codeword $c \in \mathcal{C}$. Let $E^{-1}(c) = \{x \mid E(x) = c\}$ be the set of inputs whose corresponding codeword is c . By Proposition 2, $|E^{-1}(c)| \leq M$. Hence, the number of distinct codewords in \mathcal{C} is at least $2^n/M$, and thus its length must be at least $\lceil n - \log M \rceil$. ■

Thus, to prove Theorem 2, it suffices to prove the following.

Theorem 5: Let c be a codeword in a k -linearly decodable code \mathcal{C} with side information graph G , where $k \geq n - 2$. Then, $|\ker(D_1^c + e_1, \dots, D_n^c + e_n)| \leq 2^{n - \text{minrk}_2(G)}$.

We will in fact prove a more general version of Theorem 5. To state this more general form, we first need to extend the notion of *fitting*.

Fix a graph G on n nodes. We say that a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ fits an index $i \in [n]$, if $f = g + e_i$ for some function g that depends only on $N(i)$ (g is not necessarily linear). Note that $f(x) = g(x) + x_i$. Extending the definition, we say that a family of (not necessarily distinct) functions $\langle f_j : \{0, 1\}^n \rightarrow \{0, 1\} \mid j \in T \rangle$ fits a subset T of the indices in $[n]$, if f_j fits j for every $j \in T$.

Every linear function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ corresponds to a vector v so that $f(x) = v \cdot x$. Therefore, f fits index i if and only if v can be written as $v = d + e_i$, where d is a vector whose value in every coordinate $j \notin N(i)$ equals 0. A matrix A fits $[n]$ (or, G), if the i th row of A , for every i , fits index i . As the value of this row must be 1 in the i th coordinate and 0 in every coordinate $j \notin N(i) \cup \{i\}$, this definition is consistent with our earlier definition for a matrix fitting a graph.

Fix an INDEX code \mathcal{C} for G and a codeword $c \in \mathcal{C}$. Let D_1^c, \dots, D_n^c be the n decoding functions associated with c . Note that each function $D_i^c + e_i$ fits index i , for all i , and thus the family $\langle D_1^c + e_1, \dots, D_n^c + e_n \rangle$ fits $[n]$.

We say that a family of functions $\langle f_j \mid j \in T \rangle$ is k -linear, if at least k of the functions in the family are linear. Note that if \mathcal{C} is k -linearly decodable, then the family $\langle D_1^c + e_1, \dots, D_n^c + e_n \rangle$ is k -linear.

The stronger version of Theorem 5 we will prove is as follows.

Theorem 6: Let G be a graph on n nodes and let $k \geq n - 2$. Then, for any k -linear family $\mathcal{F} = \langle f_j \mid j \in [n] \rangle$ of Boolean functions that fits $[n]$, $|\ker(\mathcal{F})| \leq 2^{n - \text{minrk}_2(G)}$.

Theorem 5 follows by setting $f_j = D_j^c + e_j$ for every j . The rest of this section is devoted to the proof of Theorem 6.

B. Maximum Dimension

In this section we explore a new notion—the *maximum dimension*—which is dual to the minrank and plays a key role in the proof of Theorem 6.

To motivate the proof, consider the following simple argument for the case $k = n$ (i.e., all the functions f_j are linear). Since f_j is linear and fits index j , it is associated with a vector v_j so that $f_j(x) = v_j \cdot x$. Let A be the $n \times n$ Boolean matrix whose rows are v_1, \dots, v_n . Since f_j fits index j , it follows that A fits G , so $\text{rk}_2(A) \geq \text{minrk}_2(G)$. Next, observe that $\ker(\mathcal{F})$ is exactly the kernel of the matrix A . By standard linear algebra, the dimension of this kernel is $n - \text{rk}_2(A) \leq n - \text{minrk}_2(G)$, and therefore the size of the kernel is at most $2^{n - \text{minrk}_2(G)}$.

To deal with the case $k < n$, we would like to generalize the above argument. When some of the functions in \mathcal{F} are not linear, $\ker(\mathcal{F})$ is no longer a linear space and thus does not have a properly defined dimension. In order to address this difficulty, we introduce the new notion of *maximum dimension*.

Let S be any subset of $[n]$ and let $\mathcal{H}_S = \langle h_j \mid j \in S \rangle$ be any family of *linear* functions that fits S . Let $T \subseteq [n] \setminus S$. For any family $\mathcal{H}_T = \langle h_j \mid j \in T \rangle$ of (not necessarily linear) functions that fits T , we denote by $\mathcal{H}_{S \cup T}$ the union of the two families: $\langle h_j \mid j \in S \cup T \rangle$. When \mathcal{H}_T is also a family of linear functions, $\ker(\mathcal{H}_{S \cup T})$ is a linear space and thus has a dimension. We define the maximum dimension of T relative to \mathcal{H}_S , denoted $\text{maxdim}_2(T \mid \mathcal{H}_S)$, to be the maximum value of $\dim(\ker(\mathcal{H}_{S \cup T}))$, where the maximum is taken over all families \mathcal{H}_T of linear functions that fit T . Note that when $S = \emptyset$, $T = [n]$, $\text{maxdim}_2(T \mid \emptyset) = n - \text{minrk}_2(G)$, and thus the maximum dimension can be viewed as dual to the minrank. The following are basic facts about the maximum dimension that will be used later in our analysis.

Proposition 4: Fix any set $S \subseteq [n]$, any family \mathcal{H}_S of linear functions that fits S , and any set $T \subseteq [n] \setminus S$. For simplicity, we shorthand $\text{maxdim}_2(T)$ for $\text{maxdim}_2(T \mid \mathcal{H}_S)$. The following are properties of $\text{maxdim}_2(T)$:

- 1) $\text{maxdim}_2(\emptyset) = \dim(\ker(\mathcal{H}_S))$.
- 2) For any $i \in T$, $\text{maxdim}_2(T) \leq \text{maxdim}_2(T \setminus \{i\}) \leq \text{maxdim}_2(T) + 1$.
- 3) More generally, $\text{maxdim}_2(T) \leq \text{maxdim}_2(T') \leq \text{maxdim}_2(T) + |T| - |T'|$ for any $T' \subseteq T$.
- 4) $\dim(\ker(\mathcal{H}_S)) - |T| \leq \text{maxdim}_2(T) \leq \dim(\ker(\mathcal{H}_S))$.
- 5) If $\text{maxdim}_2(T) = \dim(\ker(\mathcal{H}_S)) - |T|$, then $\text{maxdim}_2(T') = \dim(\ker(\mathcal{H}_S)) - |T'|$ for every $T' \subseteq T$.
- 6) Suppose $\text{maxdim}_2(\{j\}) = \dim(\ker(\mathcal{H}_S))$ for every $j \in T$. Then $\text{maxdim}_2(T) = \dim(\ker(\mathcal{H}_S))$ as well.
- 7) Let $T = [n] \setminus S$. Then, $\text{maxdim}_2(T) \leq n - \text{minrk}_2(G)$.

Proof: Part 1 follows simply by definition. Part 2 follows from the standard linear algebra fact that adding a single constraint to any subspace can only decrease its dimension, but by

at most 1; an inductive argument yields Part 3. Setting $T' = \emptyset$ in Part 3 and then using Part 1 yields Part 4.

For Part 5, note that Part 4 implies that $\text{maxdim}_2(T') \geq \dim(\ker(\mathcal{H}_S)) - |T'|$. By Part 3

$$\begin{aligned} \text{maxdim}_2(T') &\leq \text{maxdim}_2(T) + |T| - |T'| \\ &= \dim(\ker(\mathcal{H}_S)) - |T'| \end{aligned}$$

using the premise of Part 5. Therefore, $\text{maxdim}_2(T') = \dim(\ker(\mathcal{H}_S)) - |T'|$ as well.

For Part 6, the premise says that there exist linear functions h_j for all $j \in T$ such that $h_j(x) = 0$ for all $x \in \ker(\mathcal{H}_S)$. Define the family $\mathcal{H}_T = \langle h_j : j \in T \rangle$. It can be seen that $\ker(\mathcal{H}_{S \cup T}) = \ker(\mathcal{H}_S)$ and thus $\dim(\ker(\mathcal{H}_{S \cup T})) = \dim(\ker(\mathcal{H}_S))$, which is the maximum value that $\text{maxdim}_2(T)$ can attain by Part 4.

Finally, for Part 7, let \mathcal{H}_T be the family of linear functions such that $\dim(\ker(\mathcal{H}_{S \cup T})) = \text{maxdim}_2(T)$. Recall that $\mathcal{H}_{S \cup T}$ fits $S \cup T = [n]$, so let A be the matrix whose rows consist of the vectors that correspond to the functions in $\mathcal{H}_{S \cup T}$. It follows that A fits G . Since its kernel equals $\ker(\mathcal{H}_{S \cup T})$, we conclude

$$\dim(\ker(\mathcal{H}_{S \cup T})) = n - \text{rk}_2(A) \leq n - \text{minrk}_2(G).$$

■

The following lemma is the main technical result that will be used to prove Theorem 6.

Lemma 1: Let G be a graph on n nodes. Then, for any $S \subseteq [n]$, any family \mathcal{H}_S of linear functions that fits S , any $T \subseteq [n] \setminus S$ with $|T| \leq 2$, and any family \mathcal{H}_T of (not necessarily linear) functions that fits T , $|\ker(\mathcal{H}_{S \cup T})| \leq 2^{\text{maxdim}_2(T|\mathcal{H}_S)}$.

To derive Theorem 6, we choose S to be the set of indices of the k linear functions in \mathcal{F} , \mathcal{H}_S to be these linear functions, $T = [n] \setminus S$, and \mathcal{H}_T to be the rest of the functions in \mathcal{F} . Note that $\mathcal{F} = \mathcal{H}_{S \cup T}$. By Proposition 4, Part 7, we have $\text{maxdim}_2(T|\mathcal{H}_S) \leq n - \text{minrk}_2(G)$, which immediately yields Theorem 6.

Note that the restriction we have on k ($k \geq n-2$) in theorems 2, 5, and 6 derives from the restriction we have in Lemma 1 on $|T|$ ($|T| \leq 2$). It remains an open problem to find the largest value of $|T|$ (and thus the smallest value of k) for which the bound holds.

We first prove a stronger version of Lemma 1 for the special case when $\text{maxdim}_2(T|\mathcal{H}_S)$ has the smallest possible value $\dim(\ker(\mathcal{H}_S)) - |T|$ (Proposition 4, Part 4), in which case the bound given by Lemma 1 is achieved for every T (even $|T| > 2$).

Lemma 2: Let G, S, \mathcal{H}_S, T , and \mathcal{H}_T be as defined above (except that $|T|$ need not be at most 2). If $\text{maxdim}_2(T|\mathcal{H}_S) = \dim(\ker(\mathcal{H}_S)) - |T|$, then $|\ker(\mathcal{H}_{S \cup T})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T|}$.

Proof: As we will see later, proving the lemma for the case \mathcal{H}_T is a family of linear functions is easy (follows from standard dimension arguments). To extend the proof to hold for unrestricted functions, we will use a ‘‘Balance Lemma,’’ which is proved in the next section via Fourier analysis.

The lemma will be proved by gradually moving from a family \mathcal{H}_T of linear functions to a family \mathcal{H}_T of unrestricted functions. Formally, we will show the following.

Claim 2: Let G, S, \mathcal{H}_S, T , and \mathcal{H}_T be as defined in Lemma 2. Let $\ell \leq |T|$. If \mathcal{H}_T is ℓ -linear and $\text{maxdim}_2(T|\mathcal{H}_S) = \dim(\ker(\mathcal{H}_S)) - |T|$, then $|\ker(\mathcal{H}_{S \cup T})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T|}$.

Applying Claim 2 with $\ell = 0$ will yield Lemma 2. We prove the claim by a double induction: an outer induction on $|T|$ and an inner induction on $|T| - \ell$.

The base case of the outer induction $|T| = 0$ follows from standard linear algebra, because $|\ker(\mathcal{H}_S)| = 2^{\dim(\ker(\mathcal{H}_S))}$. For the base case of the inner induction $\ell = |T|$ note that \mathcal{H}_T is a linear family of functions. Therefore, $\ker(\mathcal{H}_{S \cup T})$ is a linear space and $\dim(\ker(\mathcal{H}_{S \cup T})) \leq \text{maxdim}_2(T|\mathcal{H}_S) = \dim(\ker(\mathcal{H}_S)) - |T|$. On the other hand, $\dim(\ker(\mathcal{H}_{S \cup T})) \geq \dim(\mathcal{H}_S) - |T|$, because each constraint added to a linear sub-space can reduce its dimension by at most 1. Hence, $\dim(\ker(\mathcal{H}_{S \cup T})) = \dim(\mathcal{H}_S) - |T|$. Now, define $S' = S \cup T$ and $T' = \emptyset$. As all the functions in $\mathcal{H}_{S'}$ are linear and as $\text{maxdim}_2(T'|\mathcal{H}_{S'}) = \dim(\ker(\mathcal{H}_{S'}))$ (Proposition 4, Part 1), we can apply the base case of the outer induction to conclude that

$$\begin{aligned} |\ker(\mathcal{H}_{S \cup T})| &= |\ker(\mathcal{H}_{S'})| = 2^{\dim(\ker(\mathcal{H}_{S'}))} \\ &= 2^{\dim(\ker(\mathcal{H}_S)) - |T|}. \end{aligned}$$

Let $1 \leq t \leq n$ and let $0 \leq \ell \leq t - 1$. For the induction step, assume that the claim holds for the following cases: (1) every T with $|T| < t$ and every family \mathcal{H}_T that fits T (no linearity restrictions on \mathcal{H}_T); (2) every T with $|T| = t$ and every ℓ' -linear family \mathcal{H}_T that fits T , where $\ell' > \ell$. We will show that the claim holds also for the case $|T| = t$ and \mathcal{H}_T is ℓ -linear.

Let \mathcal{H}_T be any ℓ -linear family of functions that fits T . At least ℓ of the functions in \mathcal{H}_T are linear. If \mathcal{H}_T has $\ell + 1$ linear functions or more, then it is in fact $(\ell + 1)$ -linear, and therefore the statement of the claim follows in this case from the induction hypothesis. So suppose exactly ℓ of the functions in \mathcal{H}_T are linear. As $\ell < |T|$, \mathcal{H}_T has at least one nonlinear function. Let h_i , where $i \in T$, be one such function.

Let $T_{-i} = T \setminus \{i\}$ and let $\mathcal{H}_{T_{-i}}$ be the family of functions obtained by removing h_i from \mathcal{H}_T . We will prove that $|\ker(\mathcal{H}_{S \cup T})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T|}$ in two steps. First, we will show that $|\ker(\mathcal{H}_{S \cup T_{-i}})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T| + 1}$. Then, we will prove that h_i is *balanced* on the set $\ker(\mathcal{H}_{S \cup T_{-i}})$.

Definition 3 (Balanced Function): A Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to be *balanced* on a subset U of its domain, if it is 0 on half of the inputs in U and it is 1 on the other half. That is, $|\ker(f) \cap U| = |U|/2$.

Having proved that $|\ker(\mathcal{H}_{S \cup T_{-i}})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T| + 1}$ and that h_i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$, we will obtain the desired equality

$$\begin{aligned} |\ker(\mathcal{H}_{S \cup T})| &= |\ker(h_i) \cap \ker(\mathcal{H}_{S \cup T_{-i}})| = \\ &= \frac{1}{2} |\ker(\mathcal{H}_{S \cup T_{-i}})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T|}. \end{aligned}$$

We start by showing that $|\ker(\mathcal{H}_{S \cup T_{-i}})| = 2^{\dim(\ker(\mathcal{H}_S)) - |T| + 1}$. Using Proposition 4, Part 5, since $\text{maxdim}_2(T|\mathcal{H}_S) = \dim(\ker(\mathcal{H}_S)) - |T|$, then $\text{maxdim}_2(T_{-i}|\mathcal{H}_S) = \dim(\ker(\mathcal{H}_S)) - |T| + 1$. As

$|T_{-i}| < t$, we can apply the induction hypothesis and obtain what we wanted

$$|\ker(\mathcal{H}_{S \cup T_{-i}})| = 2^{\ker(\mathcal{H}_S) - |T| + 1}. \quad (1)$$

Showing that h_i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$ is harder. To this end, we first prove that every linear function that fits index i must be balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$. We then prove a Balance Lemma, which shows that every function that fits index i , h_i included, must be balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$.

Let us start by proving that every linear function that fits index i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$. Let g_i be any such linear function and let $\mathcal{H}'_T = \mathcal{H}_{T_{-i}} \cup \{g_i\}$. Note that \mathcal{H}'_T fits T and that it is $(\ell+1)$ -linear. Let $\mathcal{H}'_S = \mathcal{H}_S$. Applying the induction hypothesis we obtain

$$|\ker(\mathcal{H}'_{S \cup T})| = 2^{\ker(\mathcal{H}'_S) - |T|} = 2^{\ker(\mathcal{H}_S) - |T|}. \quad (2)$$

We can rewrite $\ker(\mathcal{H}'_{S \cup T})$ as follows:

$$\begin{aligned} \ker(\mathcal{H}'_{S \cup T}) &= \ker(g_i) \cap \ker(\mathcal{H}'_{S \cup T_{-i}}) \\ &= \ker(g_i) \cap \ker(\mathcal{H}_{S \cup T_{-i}}). \end{aligned} \quad (3)$$

Combining (1), (2), and (3), we have

$$|\ker(g_i) \cap \ker(\mathcal{H}_{S \cup T_{-i}})| = \frac{1}{2} |\ker(\mathcal{H}_{S \cup T_{-i}})|. \quad (4)$$

Equation (4) implies that the function g_i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$. As all we used is the linearity of g_i and the fact it fits index i , we conclude that every linear function that fits index i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$. The following Balance Lemma, which is proved in the next section, shows that every function that fits index i , whether linear or not, must be balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$.

Lemma 3 (Balance Lemma): Let G be a graph on n nodes, let $i \in [n]$, and let $U \subseteq \{0, 1\}^n$. If every linear function that fits index i is balanced on U , then every function that fits index i (whether linear or not) is balanced on U .

We conclude that in particular h_i is balanced on $\ker(\mathcal{H}_{S \cup T_{-i}})$, which is what we wanted. Claim 2 and Lemma 2 follow. ■

We can now prove Lemma 1:

Proof of Lemma 1: For brevity of notation, throughout this proof we shorthand $\maxdim_2(T)$ for $\maxdim_2(T|\mathcal{H}_S)$.

We prove the lemma by induction on the size of T . The case $|T| = 0$, meaning $T = \emptyset$, follows simply from the fact that $\maxdim_2(\emptyset) = \dim(\mathcal{H}_S)$ (Proposition 4, Part 1):

$$|\ker(\mathcal{H}_{S \cup T})| = |\ker(\mathcal{H}_S)| = 2^{\dim(\mathcal{H}_S)} = 2^{\maxdim_2(T)}.$$

Let $t \in \{1, 2\}$. Assume that the statement of the lemma holds for all T such that $|T| < t$. We will prove it for $|T| = t$.

For $i \in T$, let $T_{-i} = T \setminus \{i\}$. By Proposition 4, Part 2, for every $i \in T$, $\maxdim_2(T_{-i}) \in \{\maxdim_2(T), \maxdim_2(T) + 1\}$. We split our analysis into two cases.

Case 1: For some $i \in T$, $\maxdim_2(T_{-i}) = \maxdim_2(T)$. In this case

$$\begin{aligned} |\ker(\mathcal{H}_{S \cup T})| &\leq |\ker(\mathcal{H}_{S \cup T_{-i}})| \\ &\leq 2^{\maxdim_2(T_{-i})} = 2^{\maxdim_2(T)} \end{aligned}$$

where the second inequality follows from the induction hypothesis and the last equality follows from our assumption in Case 1.

Case 2: For all $i \in T$, $\maxdim_2(T_{-i}) = \maxdim_2(T) + 1$. This is the case we know how to handle only for $|T| = 1, 2$. Suppose, first, that $|T| = 1$. Then, by the assumption of this case, $\maxdim_2(\emptyset) = \maxdim_2(T) + 1$. Since $\maxdim_2(\emptyset) = \dim(\ker(\mathcal{H}_S))$ (Proposition 4, Part 1), we rearrange and obtain $\maxdim_2(T) = \dim(\ker(\mathcal{H}_S)) - 1$. Hence, the statement follows in this case from Lemma 2.

Consider now the case $|T| = 2$ and let $T = \{i, j\}$. By the premise of Case 2, $\maxdim_2(\{i\}) = \maxdim_2(\{j\}) = \maxdim_2(\{i, j\}) + 1$. By Proposition 4, Part 2, either both $\maxdim_2(\{i\})$ and $\maxdim_2(\{j\})$ equal $\maxdim_2(\emptyset) = \dim(\ker(\mathcal{H}_S))$ or both are 1 less than $\dim(\ker(\mathcal{H}_S))$. The first case is impossible because by Proposition 4, Part 6, $\maxdim_2(\{i, j\}) = \dim(\mathcal{H}_S)$ as well, violating the premise of Case 2. Therefore, $\maxdim_2(\{i\}) = \maxdim_2(\{j\}) = \dim(\mathcal{H}_S) - 1$ implying that $\maxdim_2(\{i, j\}) = \dim(\mathcal{H}_S) - 2$. Hence, the statement follows in this case once again from Lemma 2. ■

C. Proof of the Balance Lemma

We next prove the Balance Lemma used in the proof of the lower bound for semilinearly decodable codes.

Lemma 3 (Restated): Let G be a graph on n nodes, let $i \in [n]$, and let $U \subseteq \{0, 1\}^n$. If every linear function that fits index i is balanced on U , then every function that fits index i (whether linear or not) is balanced on U .

The proof of the lemma relies on a simple principle: under the mapping $0 \mapsto 1$ and $1 \mapsto -1$, a Boolean function f is balanced on the set U if and only if $\sum_{x \in U} f(x) = 0$. The linear Boolean functions in the ± 1 world are exactly the characters of the group \mathbb{Z}_2^n and thus the lemma tells us that each of these characters sums to 0 on U . Fourier transform allows us to write any Boolean function f as a linear combination of characters. Therefore, if all characters sum to 0 on U , then also f must sum to 0 on U , and thus f is balanced.

To prove the lemma, we need to prepare some machinery from Fourier analysis of Boolean functions. Consider the group \mathbb{Z}_2^n , whose elements are the vectors $\{0, 1\}^n$. By mapping the standard 0 to 1, the standard 1 to -1 , and the XOR operation to multiplication, we view the elements of the group as vectors in $\{-1, 1\}^n$, where coordinate-wise multiplication is the group operation. A complex function $f : \mathbb{Z}_2^n \rightarrow \mathbb{C}$ over this group can be viewed as a vector in \mathbb{C}^{2^n} . The inner product between two functions $f, g \in \mathbb{C}^{2^n}$ is defined as $\langle f, g \rangle = \frac{1}{2^n} \sum_{x \in \mathbb{Z}_2^n} f(x)g(x)$.

\mathbb{Z}_2^n has 2^n characters. Each subset $S \subseteq [n]$ is associated with the character χ_S defined as: $\chi_S(x) = \prod_{i \in S} x_i$. The characters form an orthonormal basis of \mathbb{C}^{2^n} . The expansion of a function $f \in \mathbb{C}^{2^n}$ in this basis is its *Fourier Transform*. The coefficient of χ_S in this expansion is $\hat{f}(S) = \langle f, \chi_S \rangle$. Thus

$$f = \sum_{S \subseteq [n]} \hat{f}(S) \chi_S.$$

A *Boolean function* is a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ (recall the mapping $0 \mapsto 1$ and $1 \mapsto -1$). The *kernel* of a

Boolean function f is the set of inputs that is mapped to 1: $\ker(f) = \{x | f(x) = 1\}$. It is easy to verify that the characters of \mathbb{Z}_2^n are exactly the set of all Boolean linear functions on \mathbb{Z}_2^n .

To prove Lemma 3, we show two simple properties of Boolean functions.

Proposition 5: Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a Boolean function that depends only on a set $S \subseteq [n]$. Then, the Fourier transform of f has nonzero coefficients only for characters χ_T with $T \subseteq S$.

Proof: Let T be any subset of the coordinates that is not contained in S . We show that $\hat{f}(T) = 0$.

Since $T \not\subseteq S$, there exists a coordinate $i \in T \setminus S$. For each vector $x \in \mathbb{Z}_2^n$, let $x^{(i)}$ denote the vector obtained from x by flipping its i th bit (from 1 to -1 or vice versa). Let Z^1 be the set of vectors in \mathbb{Z}_2^n that have 1 at the i th coordinate, and let Z^{-1} be the set of vectors in \mathbb{Z}_2^n that have -1 at the i th coordinate. The mapping $x \mapsto x^{(i)}$ induces a perfect matching of vectors in Z^1 with vectors in Z^{-1} .

Note that for a pair $(x, x^{(i)})$, $f(x) = f(x^{(i)})$, because the two inputs differ only outside the set S . However, $\chi_T(x) \neq \chi_T(x^{(i)})$ because x and $x^{(i)}$ differ only at the i th coordinate and $i \in T$.

Consider now the coefficient $\hat{f}(T)$: $\hat{f}(T) = \frac{1}{2^n} \sum_{x \in \mathbb{Z}_2^n} f(x) \chi_T(x)$. We reorder the terms in the sum according to the above matching: $\hat{f}(T) = \frac{1}{2^n} \sum_{x \in Z^1} (f(x) \chi_T(x) + f(x^{(i)}) \chi_T(x^{(i)}))$. Since $f(x) = f(x^{(i)})$ and since $\chi_T(x) \neq \chi_T(x^{(i)})$, each of the terms in the above sum is 0. Therefore, $\hat{f}(T) = 0$, as desired. ■

Next, we characterize the set of Boolean linear functions that depend only on a set S .

Proposition 6: The set of Boolean linear functions that depend only on S is exactly the set of characters $\{\chi_T\}_{T \subseteq S}$.

Proof: Suppose $T \subseteq S$. We show that χ_T depends only on S . Let $x, x' \in \{-1, 1\}^n$ be two inputs s.t. $x[S] = x'[S]$. Since $T \subseteq S$, it follows that also $x[T] = x'[T]$. Therefore, $\prod_{i \in T} x_i = \prod_{i \in T} x'_i$, implying $\chi_T(x) = \chi_T(x')$.

For the other direction, suppose $T \not\subseteq S$. Let $i \in T \setminus S$. Let $\mathbf{1}$ be the all-one input (corresponding to the all-zero input in the 0-1 world) and let e_i be the standard unit vector (e_i is 1 in every coordinate, except for the i th coordinate in which it is -1). Since $i \notin S$, $\mathbf{1}[S] = e_i[S]$. Clearly, $\chi_T(\mathbf{1}) = 1$. Since there is a single coordinate in T in which e_i is -1 , then $\chi_T(e_i) = -1$. Thus, $\chi_T(\mathbf{1}) \neq \chi_T(e_i)$, implying χ_T does not depend only on S . ■

We can now prove Lemma 3.

Proof of Lemma 3: A Boolean function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ is balanced on U if and only if the number of inputs that it maps to 1 equals the number of inputs that it maps to -1 . This in turn happens if and only if $\sum_{x \in U} f(x) = 0$.

By Proposition 6, the set of linear Boolean functions that depend only on $N(i)$ is the family of characters $\{\chi_T\}_{T \subseteq N(i)}$. Therefore, the set of linear functions that fit index i are of the form $\{\chi_T \cdot e_i\}_{T \subseteq N(i)}$. (Since we moved to the ± 1 world, summation is mapped to multiplication, and the standard unit vector e_i is the all-one vector, except for the i th coordinate which is -1 .) The premise given in Lemma 3 implies that all these

functions are balanced on U . That is, for every $T \subseteq N(i)$, $\sum_{x \in U} \chi_T(x) \cdot x_i = 0$.

Let f be any function (not necessarily linear) that fits index i . We can write $f = g \cdot e_i$, where g is a function that depends only on $N(i)$. By Proposition 5, g is a linear combination of the characters $\{\chi_T\}_{T \subseteq N(i)}$. Therefore

$$\begin{aligned} \sum_{x \in U} f(x) &= \sum_{x \in U} g(x) \cdot x_i = \sum_{x \in U} \left(\sum_{T \subseteq S} \hat{g}(T) \chi_T(x) \right) \cdot x_i \\ &= \sum_{x \in U} \sum_{T \subseteq S} \hat{g}(T) (\chi_T(x) \cdot x_i) \\ &= \sum_{T \subseteq S} \hat{g}(T) \sum_{x \in U} \chi_T(x) \cdot x_i \\ &= \sum_{T \subseteq S} \hat{g}(T) \cdot 0 = 0. \end{aligned}$$

Therefore, also f is balanced on U . ■

V. GENERAL CODES

In this section, we prove lower bounds for the class of general randomized INDEX codes. Let us first formally define these codes.

Definition 4 (Randomized INDEX Codes): Let $0 \leq \delta < 1$. A δ -error randomized INDEX code \mathcal{C} for $\{0, 1\}^n$ with side information graph G on n nodes is a set of codewords in $\{0, 1\}^\ell$ together with:

- 1) A public random string R for both encoding and decoding.
- 2) A private random string R_E for encoding.
- 3) n private random strings R_{D_1}, \dots, R_{D_n} for decoding.
- 4) An encoding function E that given a source input x in $\{0, 1\}^n$ maps the triple (x, R, R_E) into a codeword.
- 5) A set of decoding functions D_1, D_2, \dots, D_n . For each i , D_i maps the quadruple $(E(x, R, R_E), x[N(i)], R, R_{D_i})$ into a bit, satisfying the following:

$$\Pr(D_i(E(x, R, R_E), x[N(i)], R, R_{D_i}) \neq x_i) \leq \delta.$$

The probability is over the three random strings R , R_E , and R_{D_i} .

The $n + 2$ random strings $R, R_E, R_{D_1}, \dots, R_{D_n}$ have finite domains and are mutually independent of each other. Usually these are uniformly distributed strings of some fixed length.

The distributions of the private random strings are known in advance to all parties (the sender and the receivers), yet the specific instances chosen are known only to the respective parties. Therefore, the encoding function E may depend on the distributions of R_{D_1}, \dots, R_{D_n} , but not on the specific instances chosen. Similarly, D_i may depend on the distribution of R_E , but not on the specific instance. As usual, the graph G is known in advance to the sender and the receivers and thus the encoding and decoding functions can depend on G . The length of \mathcal{C} , denoted by $\text{len}(\mathcal{C})$, is defined to be ℓ .

The main technical statement of this section is a direct-sum result for the information cost of a randomized INDEX code. A corollary of this result will be the lower bound on the length of randomized INDEX codes. We start with a brief overview of the information theory notions and facts used in this section (see [27] for a more extensive background).

A. Information Theory Background

In the following $X \sim \mu_X$, $Y \sim \mu_Y$, $Z \sim \mu_Z$ are random variables on domains $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$, respectively.

Entropy and Mutual Information: The entropy of X (or, equivalently, of μ_X) is $H(X) = \sum_{x \in \mathcal{X}} \mu_X(x) \log \frac{1}{\mu_X(x)}$. The binary entropy function $H_2(p)$ is the entropy of a Bernoulli random variable with probability of success p .

The *joint entropy* of X and Y is the entropy of the joint distribution (μ_X, μ_Y) . The *conditional entropy* of X given an event A , denoted $H(X|A)$, is the entropy of the conditional distribution of μ_X given A . The *conditional entropy* of X given Y is $H(X|Y) = \sum_{y \in \mathcal{Y}} \mu_Y(y) H(X|Y=y)$.

The *mutual information* between X and Y is $I(X; Y) = H(X) - H(X|Y)$. The conditional mutual information between X and Y given Z is $I(X; Y|Z) = H(X|Z) - H(X|Y, Z)$.

The following are basic facts about entropy and mutual information.

Proposition 7 (Entropy Bound): Let X be any random variable and let $\text{supp}(X)$ be the support of X . Then, $H(X) \leq \log |\text{supp}(X)|$. Equality iff X is uniform on $\text{supp}(X)$.

Proposition 8 (Chain Rule Conditional Mutual Information): For any sequence of random variables X_1, \dots, X_n

$$I(X_1, \dots, X_n; Y) = \sum_{i=1}^n I(X_i; Y|X_1, \dots, X_{i-1}).$$

Fano's Inequality: Fano's inequality [28] gives a lower bound on the error probability of predicting the value of a random variable X from the observation of another random variable Y . We consider a special case where X is uniformly distributed over a binary domain.

Theorem 7 (Fano's Inequality): Let Y be a random variable and let X be uniformly distributed over $\{0, 1\}$. Let $g(Y) \in \{0, 1\}$ be a function whose prediction error probability $\Pr(g(Y) \neq X) \leq \delta \leq 1/2$. Then, $I(X; Y) \geq 1 - H_2(\delta)$.

B. Direct Sum for Information Cost

Suppose G' is a vertex-induced subgraph of G . An INDEX code for G easily yields an INDEX code for G' of the same length by arbitrarily fixing the bits of $V(G) \setminus V(G')$. Thus,

Proposition 9: If G' is a vertex-induced subgraph of G , then the optimal length of an INDEX code for G' is no more than that of G .

What about the other direction? Suppose we can split G into k mutually disjoint vertex-induced subgraphs G_1, \dots, G_k and suppose we have k INDEX codes $\mathcal{C}_1, \dots, \mathcal{C}_k$ for these subgraphs. Clearly, by concatenating $\mathcal{C}_1, \dots, \mathcal{C}_k$ we can obtain an INDEX code for G whose length is $\sum_{i=1}^k \text{len}(\mathcal{C}_i)$. But is this always the optimal length code for G ?

In general, it looks like one could obtain shorter INDEX codes for G , by exploiting the edges connecting the different subgraphs G_1, \dots, G_k . But what if these graphs are disconnected from each other? In this case, it seems that the optimal length of the INDEX code for G must equal the sum of the optimal lengths of the INDEX codes for G_1, \dots, G_k . In other words, the optimal length of INDEX codes should admit a direct sum property. Nevertheless, proving this property for the measure of code length is

elusive. The techniques of Feder *et al.* [16] yield a weaker result, which incurs a loss of an additive term that depends linearly on k . We are able to prove the direct sum property not directly for code length, but rather for the "information cost" of codes.

Definition 5 (Information Cost): Let \mathcal{C} be a randomized index code for G . Let R denote the public random string of \mathcal{C} , and let $E(x, R)$ denote the encoding of x in \mathcal{C} .⁴ Let X be uniformly distributed in $\{0, 1\}^n$. The *information cost* of \mathcal{C} , denoted by $\text{icost}(\mathcal{C})$, equals $I(X; E(X, R) | R)$.

As the information cost of a code is always at most the entropy of the codewords, the entropy bound (Proposition 7) implies that information cost is a lower bound on the code length.

We prove that the information cost of an INDEX code admits a direct sum property. The property holds not only when G_1, \dots, G_k are totally disconnected from each other; it suffices that there are no edges directed from G_i to G_j for all $i < j$:

Theorem 8: Let G_1, G_2, \dots, G_k be vertex-induced subgraphs of a directed graph G such that:

- 1) The vertices of G_1, G_2, \dots, G_k partition the vertices of G .
- 2) For any $i < j$ and vertices $v_i \in V(G_i)$ and $v_j \in V(G_j)$, there is no directed edge in G from v_i to v_j .

Let \mathcal{C} be a δ -error randomized INDEX code for G . Then, there exist δ -error randomized INDEX codes $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$ for G_1, G_2, \dots, G_k such that $\text{icost}(\mathcal{C}) \geq \sum_j \text{icost}(\mathcal{C}_j)$.

Proof: For $j = 1 \dots k$, define $V_j = V(G_j)$ and $U_j = \bigcup_{i=1}^j V_i$. Let $E(x, R)$ be the encoding function of the INDEX code \mathcal{C} and let X be uniformly distributed on $\{0, 1\}^n$. By definition, $\text{icost}(\mathcal{C}) = I(X; E(X, R) | R)$. Using the chain rule for conditional mutual information

$$I(X; E(X, R) | R) = \sum_{j=1}^k I(X[V_j]; E(X, R) | X[U_{j-1}], R).$$

(Slightly abusing notation, $U_0 = \emptyset$ and $x[U_0]$ is an empty string.) To complete the proof of the theorem, it suffices to show the following claim.

Claim 3: For every j , there is a δ -error randomized INDEX code \mathcal{C}_j for G_j such that

$$\text{icost}(\mathcal{C}_j) = I(X[V_j]; E(X, R) | X[U_{j-1}], R).$$

Proof: The proof is based on a reduction lemma proven in [18]. Fix a value of j and we will construct \mathcal{C}_j using \mathcal{C} as follows. Let $a \in \{0, 1\}^{|V_j|}$ denote the source input. As we want to use \mathcal{C} , we need to transform a into some input x for \mathcal{C} . The transformation will be randomized. That is, x will be a random string, created from a , from the public random string, and from the private random string of the encoder.

x will be equal to a in the coordinates corresponding to vertices in G_j . The other coordinates of x will be filled randomly as follows. Let B have the same distribution as $X[V \setminus U_j]$ and let C be independent of B and have the same distribution as $X[U_{j-1}]$. The public random string for \mathcal{C}_j consists of (R, C) while B will be part of the private randomness of the sender. The random input x will be defined to be the tuple $\langle C, a, B \rangle$.

⁴The dependence of E on the sender's private randomness is suppressed for ease of presentation.

The encoding of a is then $E(\langle C, a, B \rangle, R)$. Note that this encoding is a function of a , of the sender's private random string, and of the public random string.

Let $i \in V_j$ be any coordinate. When applying the decoding function D_i of \mathcal{C} in order to recover a_i , the receiver needs to know the bits of x corresponding to neighbors of i in the graph G . By the property of G_j , it can be seen that the neighbors of i in G are either among the neighbors of i in V_j or belong to U_{j-1} . Now, the values for the former are part of the side information for coordinate i while the values for the latter can be found in the public random string C .

For any instantiation of B and C , the decoding error is simply the error of \mathcal{C} on the input x obtained from a and from the instantiations of B and C . As this error is at most δ , then also averaging over all choices of B and C , the error of \mathcal{C}_j on a is at most δ .

Next, we calculate the information cost of \mathcal{C}_j as follows. Let A be uniformly distributed over $\{0, 1\}^{|V_j|}$ and be independent of B and C . Then

$$\begin{aligned} \text{icost}(\mathcal{C}_j) &= I(A; E(\langle C, A, B \rangle, R) \mid C, R) \\ &= I(X[V_j]; E(X, R) \mid X[U_{j-1}], R) \end{aligned}$$

completing the proof of the claim. \blacksquare

Applying the above claim for all j completes the proof of the theorem. \blacksquare

C. Lower Bound for Randomized Codes

Theorem 3 can now be shown as a simple application of Theorem 8.

Theorem 3 (Restated): The length of any δ -error randomized INDEX code for G is at least $\text{MAIS}(G) \cdot (1 - H_2(\delta))$, where $\text{MAIS}(G)$ is the size of the maximum acyclic induced subgraph of G and $H_2(\cdot)$ is the binary entropy function.

Proof: Let G' be a maximal acyclic induced subgraph of G . By Proposition 9, it suffices to consider any INDEX code \mathcal{C} for G' . Let u_1, u_2, \dots, u_k denote the vertices of G' such that there is no edge from u_i to u_j whenever $i < j$. Apply Theorem 8 with $G = G'$ and where G_j is a graph with a single vertex u_j . We have $\text{icost}(\mathcal{C}) \geq \sum_j \text{icost}(\mathcal{C}_j)$. Now, since \mathcal{C}_j is a INDEX code for a single vertex graph, therefore, it encodes just a single bit that can be decoded with probability of error at most δ . By Fano's inequality, it must have at least $1 - H_2(\delta)$ bits of information. \blacksquare

VI. LOWER BOUNDS FOR RESTRICTED GRAPHS

In this section we show that for certain natural classes of graphs, the minrank bound is tight w.r.t. *arbitrary* INDEX codes.

Theorem 4 (Restated): Let G be any graph, which is either a DAG, a perfect graph, an odd hole, or an odd anti-hole. Then, the length of any INDEX code for G is at least $\text{minrk}_2(G)$.

A. Directed Acyclic Graphs

A directed acyclic graph (DAG) is one without directed cycles.

Proposition 10: Let G be any DAG on n nodes. Then, the length of any INDEX code for G is at least $\text{minrk}_2(G)$.

Proof: Let \mathcal{C} be any INDEX code for G . Since G is a DAG, then $\text{MAIS}(G) = n$. Hence, by Theorem 3, $\text{len}(\mathcal{C}) \geq n$. Clearly, $\text{minrk}_2(G) \leq n$, and thus $\text{len}(\mathcal{C}) \geq \text{minrk}_2(G)$. \blacksquare

B. Perfect Graphs

An undirected graph G is called *perfect*, if for any induced subgraph G' of G , $\omega(G') = \chi(G')$.

Proposition 11: Let G be any perfect graph on n nodes. Then, the length of any INDEX code for G is at least $\text{minrk}_2(G)$.

Proof: Let \mathcal{C} be any INDEX code for G . By Theorem 3, $\text{len}(\mathcal{C}) \geq \text{MAIS}(G)$. Since G is undirected, then $\text{MAIS}(G) = \alpha(G)$, i.e., the independence number of G . Clearly, $\alpha(G) = \omega(\overline{G})$, implying that $\text{len}(\mathcal{C}) \geq \omega(\overline{G})$.

Lovász [29] proved in 1972 the ‘‘Perfect Graph Theorem,’’ stating that a graph G is perfect if and only if its complement is perfect. Now, since G is perfect, then by this theorem also \overline{G} is perfect, implying that in particular $\omega(\overline{G}) = \chi(\overline{G})$. Hence, $\text{len}(\mathcal{C}) \geq \chi(\overline{G})$. However, by the sandwich property of minrank (Proposition 1), $\text{minrk}_2(G) \leq \chi(\overline{G})$ and thus $\text{len}(\mathcal{C}) \geq \text{minrk}_2(G)$. \blacksquare

C. Odd Holes

Before we prove the lower bound for odd holes, we first characterize their minrank.

Theorem 9: Let G be an odd hole of length $2n + 1$ ($n \geq 2$). Then, $\text{minrk}_2(G) = n + 1$.

Note that since for an odd hole G , $\omega(\overline{G}) = n$, odd holes are examples of graphs for which $\omega(\overline{G}) < \text{minrk}_2(G)$.

Proof: As $\chi(\overline{G}) = n + 1$ for an odd hole of length $2n + 1$ and as $\text{minrk}_2(G) \leq \chi(\overline{G})$ (Proposition 1), it suffices to prove that $\text{minrk}(G) \geq n + 1$.

Fix any matrix A that fits G . For convenience, we number the rows and columns of A as $0, 1, \dots, 2n$ and make all the index arithmetic below modulo $2n$. Let A_0, \dots, A_{2n} be the $2n + 1$ rows of A . A has the following three properties, for every i :

- 1) $A_i[i] = 1$.
- 2) $A_i[i - 1], A_i[i + 1] \in \{0, 1\}$.
- 3) $A_i[j] = 0$, for $j \notin \{i - 1, i, i + 1\}$.

For a row A_i , we call the rows A_0, \dots, A_{i-1} the ‘‘predecessors of A_i .’’ Note that A_0 has no predecessors. We next prove the following two claims.

Claim 4: For $i = 1, \dots, 2n - 2$, either A_i is linearly independent of its predecessors or A_{i+1} is linearly independent of its predecessors.

Proof: Suppose, to reach a contradiction, that the claim is false. Hence, there exists some $i \in \{1, \dots, 2n - 2\}$ s.t. both A_i and A_{i+1} linearly depend on their predecessors. It follows that both A_i and A_{i+1} linearly depend on A_0, \dots, A_{i-1} . Since $A_0[i + 1] = \dots = A_{i-1}[i + 1] = 0$ (using Property 3 of A and the fact $i + 1 < 2n$), then also $A_i[i + 1] = 0$ and $A_{i+1}[i + 1] = 0$. This contradicts the fact $A_{i+1}[i + 1] = 1$ (Property 1 of A). \blacksquare

Claim 5: At least one among A_1, A_{2n-1}, A_{2n} is linearly independent of its predecessors.

Proof: If at least one of A_{2n-1}, A_{2n} is independent of its predecessors, then we are done. So suppose both depend on their

predecessors. As argued earlier, this means that A_{2n-1}, A_{2n} both depend on A_0, \dots, A_{2n-2} .

By Property 1 of A , $A_{2n}[2n] = 1$. The only vector among A_0, \dots, A_{2n-2} that can have a 1 at the $2n$ th coordinate is A_0 . Thus, we must have: $A_0[2n] = 1$. By Property 3 of A , $A_1[2n] = 0$. Hence, A_1 cannot depend on its sole predecessor, A_0 . We thus obtained that A_1 is linearly independent of its predecessors. ■

Note that in Claim 5 we implicitly use the fact $n \geq 2$, because we assume the indices $1, 2n-1, 2n$ are distinct.

We next use the above two claims to count the number of rows of A that must be linearly independent of their predecessors. For each i , let $Z_i = 1$ if the i th row is independent of its predecessors and $Z_i = 0$ otherwise. The number of rows that are linearly independent of their predecessors is $\sum_{i=0}^{2n} Z_i$. Note that this number is exactly the 2-rank of the matrix A .

We know the following three facts about the sequence Z_0, \dots, Z_{2n} :

- 1) $Z_0 = 1$, because A_0 simply does not have any predecessors.
- 2) For each $i = 1, \dots, 2n-3$, $Z_i + Z_{i+1} \geq 1$ (Claim 4).
- 3) $Z_1 + Z_{2n-1} + Z_{2n} \geq 1$ (Claim 5).

We now write the sum $2 \sum_{i=0}^{2n} Z_i$ as follows:

$$2 \sum_{i=0}^{2n} Z_i = 2Z_0 + Z_1 + \sum_{i=1}^{2n-2} (Z_i + Z_{i+1}) + Z_{2n-1} + 2Z_{2n}.$$

Using the above three facts, we have

$$2 \sum_{i=0}^{2n} Z_i \geq 2 + 2n - 2 + 1 + Z_{2n} \geq 2n + 1.$$

Therefore, $\sum_{i=0}^{2n} Z_i \geq (2n+1)/2$. However, since $\sum_{i=0}^{2n} Z_i$ is an integer we have the stronger bound:

$$\sum_{i=0}^{2n} Z_i \geq \lceil (2n+1)/2 \rceil = n+1.$$

Hence, $\text{rk}_2(A) = \sum_{i=0}^{2n} Z_i \geq n+1$. As this holds for any A that fits G , also $\text{minrk}_2(G) \geq n+1$. ■

The lower bound for odd holes is then the following:

Theorem 10: Let G be an odd hole on $2n+1$ nodes ($n \geq 2$). Then, the length of any INDEX code for G is at least $\text{minrk}_2(G) = n+1$.

As for an odd hole G , $\text{MAIS}(G) = \omega(\overline{G}) = n < \text{minrk}_2(G)$, this theorem implies that our lower bound for general INDEX codes (Theorem 3) is not tight.

The proof of this lower bound is considerably harder than the proofs for DAGs and perfect graphs. To this end, we need to study some combinatorial properties of the *confusion graph* associated with INDEX coding.

Definition 6 (Confusion Graph): The *confusion graph* $C(G)$ associated with INDEX coding for a directed graph G (abbreviated “confusion graph for G ”) is an *undirected* graph on $\{0, 1\}^n$ such that x and x' are connected by an edge if for some i , we have $x[N(i)] = x'[N(i)]$ but $x_i \neq x'_i$.

If x and x' are connected by an edge in $C(G)$, then no INDEX code \mathcal{C} for G can map x and x' to the same codeword, implying $\log \chi(C(G))$ is a lower bound on $\text{len}(\mathcal{C})$.

Notation: Let $\mathbf{0}$ and $\mathbf{1}$ denote, respectively, the all-zero and the all-one vectors. Let $\mathbf{1}_S$ denotes the characteristic vector of a set $S \subseteq [n]$.

Lemma 4: Let G be an undirected graph on n nodes and let $C(G)$ be the confusion graph corresponding to INDEX coding for G . Then,

- 1) If S is a vertex cover of G , then any two inputs $x, x' \in \{0, 1\}^n$ that agree on S (i.e., $x[S] = x'[S]$) are connected by an edge in $C(G)$.
- 2) If S is an independent set in G , then the set $X_S = \{\mathbf{1}_T \mid T \subseteq S\}$ forms a clique in $C(G)$.
- 3) If S, T are two disjoint and independent sets in G , and there exists some $i \in S$ that has no neighbors in T or some $j \in T$ that has no neighbors in S , then the inputs $\mathbf{1}_S$ and $\mathbf{1}_T$ are connected by an edge in $C(G)$.

Proof of Part 1: Since $x \neq x'$, there exists some index $i \in [n]$ s.t. $x_i \neq x'_i$. This means that $i \notin S$. If a node does not belong to a vertex cover, then all its neighbors must belong to the vertex cover. We conclude that $N(i) \subseteq S$ and thus $x[N(i)] = x'[N(i)]$. This implies that x and x' are connected by an edge in the confusion graph. ■

Proof of Part 2: Define $U = [n] \setminus S$. Since S is an independent set in G , then U is a vertex cover. Note that any two input $x, x' \in X_S$ agree on U , and thus by Part 1 must be connected by an edge in the confusion graph. ■

Proof of Part 3: Suppose, for example, there is $i \in S$ that has no neighbors in T . Since S, T are disjoint, $\mathbf{1}_S$ and $\mathbf{1}_T$ disagree on the i th coordinate. Since S is independent, $N(i) \subseteq [n] \setminus S$, and thus $\mathbf{1}_S[N(i)] = \mathbf{0}$. Since $N(i) \cap T = \emptyset$, then also $\mathbf{1}_T[N(i)] = \mathbf{0}$. This implies that $\mathbf{1}_S[N(i)] = \mathbf{1}_T[N(i)]$ and therefore $\mathbf{1}_S, \mathbf{1}_T$ must be connected by an edge in the confusion graph. ■

We can now prove Theorem 10.

Proof of Theorem 10: Let G be an odd hole on $2n+1$ nodes ($n \geq 2$). Let \mathcal{C} be any INDEX code for G . We will prove that $|\mathcal{C}|$, the number of codewords in \mathcal{C} , is greater than 2^n , implying that $\text{len}(\mathcal{C}) \geq n+1 = \text{minrk}_2(G)$.

Consider the following coloring of G : $S_1 = \{1, 3, \dots, 2n-1\}$, $S_2 = \{2, 4, \dots, 2n\}$ and $S_3 = \{2n+1\}$. For each $i \in \{1, 2, 3\}$, since S_i is an independent set, then by Part 2 of Lemma 4, \mathcal{C} must use $2^{|S_i|}$ different codewords to encode inputs in X_{S_i} . Since $|S_1| = |S_2| = n$, this already implies $|\mathcal{C}| \geq 2^n$. Assume, to the contradiction, that $|\mathcal{C}| = 2^n$.

Since S_1, S_2, S_3 are pairwise disjoint, then the sets $X_{S_1}, X_{S_2}, X_{S_3}$ have only $\mathbf{0}$ as a common input and are otherwise pairwise disjoint. Since $|\mathcal{C}| = 2^n$, and no codeword can encode two different inputs in X_{S_i} ($i = 1, 2, 3$), then there must be at least one codeword encoding a nonzero input from X_{S_1} , a nonzero input from X_{S_2} , and a nonzero input from X_{S_3} . We call these inputs x_1, x_2, x_3 .

We view x_1, x_2, x_3 as characteristic vectors of sets $T_1, T_2, T_3 \subseteq [n]$. Since $x_1, x_2, x_3 \neq \mathbf{0}$, then $T_1, T_2, T_3 \neq \emptyset$.

Furthermore, they are all independent and pairwise disjoint. Since the only nonzero vector in X_{S_3} is e_{2n+1} , $T_3 = \{2n+1\}$.

Since x_1, x_2, x_3 are encoded by the same codeword, no two of them can be connected by an edge in the confusion graph. Consider any $i \in T_1$. By Part 3 of Lemma 4, i must have a neighbor $j \in T_2$. Similarly, both i and j must have neighbors in T_3 . Since $T_3 = \{2n+1\}$, both are neighbors of $2n+1$. We conclude that $(i, j, 2n+1)$ forms a triangle in G . However, all odd holes are triangle-free. This is a contradiction, and thus $|\mathcal{C}| > 2^n$. ■

The above theorem provides a tight lower bound on the *length* of INDEX codes for odd holes, but not on their *size*. Our upper bound (Theorem 1) gives a code whose size is 2^{n+1} , while the above proof only shows a lower bound of $|\mathcal{C}| > 2^n$. Optimal code size lower bounds are important for deriving lower bounds on the average encoding length and on the information cost. Resorting to a more involved combinatorial argument, we are able to prove tight bounds (i.e., 2^{n+1}) on the size of INDEX codes for odd holes of length at least 7.

Theorem 11: Let G be an odd hole on $2n+1$ nodes ($n \geq 3$). Then, the size of any INDEX code for G is at least $2^{\text{minrk}(G)} = 2^{n+1}$.

The proof of this theorem appears in Appendix A.

Dealing with the pentagon (a hole of length 5) turns out to be very tricky. The difficulty of handling the pentagon stems from the fact that the corresponding confusion graph has a rather peculiar property. In most cases, one can obtain tight lower bounds on the chromatic number of the confusion graph by obtaining tight upper bounds on the graph's independence number. It turns out that this approach fails for the pentagon. The size of the pentagon's confusion graph is 32 and its chromatic number is 8. Yet, the code we show below for the pentagon demonstrates that the independence number of its confusion graph is 5, implying that $32/5 < 8$ is not a tight lower bound on the chromatic number.

Codeword	Inputs
C_1	00000, 00110, 10001, 11101, 11110
C_2	11111, 11001, 01110, 00010, 00001
C_3	01010, 01100, 11011, 10111, 10100
C_4	00100, 01011, 10010, 10101
C_5	00111, 01001, 10110, 11010
C_6	01000, 01111, 10000, 10011
C_7	00011, 00101, 11000
C_8	01101, 11100

By applying arguments from the proof of Theorem 11 we can obtain a lower bound of 7 on the size of codes for the pentagon, one short of the upper bound of 8. By the same arguments, any INDEX code of size 7 for the pentagon must adhere to certain structural constraints. By a brute force exhaustive search over such codes, we verified that 8 is the tight lower bound.

D. Odd Anti-Holes

Recall that an odd anti-hole is the complement graph of an odd hole. We prove a tight lower bound on the minimum length

of codes for odd anti-holes. This bound does not give a tight lower bound on the size of codes for odd anti-holes. Unfortunately, we could not prove tight bounds on the size.

Theorem 12: Let G be an odd anti-hole on $2n+1$ nodes ($n \geq 2$). Then, the length of any INDEX code for G is at least $\text{minrk}(G) = 3$.

Proof: We use the same notation and propositions as in the proof for odd holes. Let \mathcal{C} be any INDEX code for G . We would like to show that $|\mathcal{C}| \geq 5$. That would imply that $\text{len}(\mathcal{C}) \geq 3$.

An odd anti-hole of length $2n+1$ is $(n+1)$ -colorable. Consider the following coloring of G : $S_1 = \{1\}, S_2 = \{2, 3\}, \dots, S_{n+1} = \{2n, 2n+1\}$. Let $X_{S_1}, \dots, X_{S_{n+1}}$ be the input sets corresponding to S_1, \dots, S_{n+1} . Note that these sets share a single input (the all-zero input) and are otherwise pairwise disjoint.

For $i = 2, \dots, n+1$, $|X_{S_i}| = 4$, and thus by Part 2 of Lemma 4, \mathcal{C} must use 4 different codewords for each of these sets. Assume, to reach a contradiction, that $|\mathcal{C}| = 4$. Therefore, there must be a single codeword that encodes a nonzero input from X_{S_i} , for each $i = 1, \dots, n+1$. Let us denote these inputs by x_1, \dots, x_{n+1} . We view x_1, \dots, x_{n+1} as characteristic vectors of sets T_1, \dots, T_{n+1} . These sets are all independent and pairwise disjoint. Furthermore, $T_1 = \{1\}$, because the only nonzero input in X_{S_1} is e_1 .

Since x_1, \dots, x_{n+1} are all encoded by a single input, they must form an independent set in the confusion graph $C(G)$. We next prove by induction that for every $i = 1, \dots, n+1$, T_i must be the set $\{2i-1\}$.

For $i = 1$, we already know that $T_1 = \{1\}$. Assume correctness for i . We will show correctness for $i+1$. Since x_i and x_{i+1} are not connected by an edge in the confusion graph, then by Part 3 of Lemma 4, every node in T_i must have a neighbor in T_{i+1} and vice versa. Since $T_i = \{2i-1\}$ and since the only neighbor of $2i-1$ in the set $S_{i+1} = \{2i, 2i+1\}$ is $2i+1$, then T_{i+1} must be $\{2i+1\}$.

It follows that $T_{n+1} = \{2n+1\}$. However, since nodes 1 and $2n+1$ are not neighbors in G , it follows that no node in T_1 has neighbors in T_{n+1} . Thus, by Part 3 of Lemma 4, x_1 and x_{n+1} must be connected by an edge in the confusion graph, in contradiction to the fact x_1, \dots, x_{n+1} is an independent set in the confusion graph. Therefore, $|\mathcal{C}| \geq 5$. ■

VII. CONCLUSION

In this paper, we explored upper and lower bounds on the length of INDEX codes for $\{0, 1\}^n$ with side information graph G . We identified a measure on graphs, the *minrank*, which we showed to characterize the length of INDEX codes for natural classes of graphs (DAGs, perfect graphs, odd holes, and odd anti-holes). We also proved that minrank characterizes the minimum length of natural types of INDEX codes (linear, linearly decodable, and semilinearly decodable) for *arbitrary* graphs. For general codes and general graphs, we were able to obtain a weaker bound in terms of the maximum acyclic induced subgraph. Finally, we proved a direct sum theorem for the information cost of INDEX codes with side information.

As Lubetzky and Stav [25] have recently shown, the minrank is not a tight lower bound on the length of a general INDEX code

for arbitrary graphs. Characterizing the optimal length of INDEX codes for arbitrary graphs therefore remains an open problem. It is nonetheless important to note that virtually all codes presently in use are linear, and for those our bounds are tight.

The minrank by itself is an interesting subject of study. We know that for undirected graphs, it is bounded from below by the Shannon capacity and from above by the chromatic number of the complement graph. It would be interesting to explore further properties of minrank with respect to other graph measures such as the Lovász Theta function.

The paper only considered binary codes; it might be interesting to extend the ideas of the paper to larger field sizes.

Heuristics for constructing the code for a given side information graph is an interesting research topic given the resulting benefits. To this end, knowing the length of the code may be helpful, so efficient heuristics for estimating the *minrank* of a given matrix (an NP-hard problem) may be beneficial.

Finally, a practical conclusion is that keeping “junk” (unneeded information) may be beneficial, as it can serve as side information and save communication. This is particularly true in view of the growing abundance and declining cost of storage space.

APPENDIX

SIZE LOWER BOUND FOR ODD HOLES

Theorem 11 (Restated): Let G be an odd hole on $2n + 1$ nodes ($n \geq 3$). Then, the size of any INDEX code for G is at least $2^{\minrk(G)} = 2^{n+1}$.

Proof: Our strategy for proving the theorem will be by showing that the independence number of the confusion graph $C(G)$ is at most 2^n . Since $|C(G)| = 2^{2n+1}$, it will immediately follow that $\chi(C(G)) \geq 2^{n+1}$, giving us the desired lower bound on the size of the INDEX code.

In the derivations below, we assume the nodes of G are numbered $0, 1, 2, \dots, 2n$. All the index arithmetics are done modulo $2n$.

Let X be any independent set of inputs in the confusion graph. We would like to prove that $|X| \leq 2^n$. For any set $S \subseteq [n]$ of coordinates, define $X[S]$ to be the set obtained by projecting all the inputs in X on the coordinates in S : $X[S] = \{x[S] \mid x \in X\}$. (When S is a sequence of indices $\{i_1, \dots, i_k\}$, we write $X[i_1, \dots, i_k]$ as a shorthand for $X[\{i_1, \dots, i_k\}]$.)

As an immediate corollary of Part 1 of Lemma 4, we obtain the following.

Proposition 12: Let G be an undirected graph on n nodes, and let $X \subseteq \{0, 1\}^n$ be an independent set in the confusion graph $C(G)$ for G . Then, for any vertex cover $S \subseteq [n]$ of G , $|X| \leq |X[S]|$.

It is possible to strengthen Proposition 12. Given an undirected graph G on n nodes and a subset T of its vertices, we say that a subset S of T is a vertex cover of T , if for every node $v \in T$, either $v \in S$ or $N(v) \subseteq S$ (where $N(v)$ is the set of all the neighbors of v in G). For example, if G is a cycle, and $T = \{i, i+1, i+2\}$, then $S = \{i, i+2\}$ is a vertex cover of T . The following is a straightforward generalization of Proposition 12.

Proposition 13: Let G be an undirected graph on n nodes, and let $X \subseteq \{0, 1\}^n$ be an independent set in the confusion graph $C(G)$ for G . Then, for any subset $T \subseteq [n]$, and for any vertex cover S of T , $|X[T]| \leq |X[S]|$.

Proof: Suppose, to reach a contradiction, $|X[T]| > |X[S]|$. Then, by the Pigeonhole Principle, there exist two inputs $x, x' \in X$ s.t. $x[S] = x'[S]$ but $x[T] \neq x'[T]$. That is, there is some $j \in T \setminus S$, s.t. $x[j] \neq x'[j]$. Yet, since $j \notin S$, $N(j) \subseteq S$, and thus x and x' agree on all the neighbors of j but disagree on j . This means that x and x' are connected by an edge in the confusion graph, in contradiction to the assumption they are both members of an independent set. ■

Now, in order to bound $|X|$, we consider three cases.

Case 1: $\exists i \in [0, 2n]$ s.t. $|X[i, i+1]| \leq 2$:

Without loss of generality, assume $i = 0$. We construct a vertex cover S of G as follows: $S = \{0, 1, 3, 5, 7, \dots, 2n-1\}$. Note that S is indeed a vertex cover of G and that $|S| = n+1$. We split it into two parts: $S_1 = \{0, 1\}$ and $S_2 = \{3, 5, \dots, 2n-1\}$. Clearly, $|X[S]| \leq |X[S_1]| \cdot |X[S_2]|$. Since $|S_2| = n-1$, then $|X[S_2]| \leq 2^{n-1}$. Therefore, $|X[S]| \leq |X[S_1]| \cdot |X[S_2]| \leq 2 \cdot 2^{n-1} \leq 2^n$. Using Proposition 12 we have in this case: $|X| \leq 2^n$.

Case 2: $\exists i \in [0, 2n]$ s.t. $|X[i, i+1]| = 4$: WLOG, assume $i = 1$. Hence, $|X[1, 2]| = 4$. Since $\{1, 3\}$ is a vertex cover of the set $\{1, 2, 3\}$, then by Proposition 13, $|X[1, 2, 3]| \leq |X[1, 3]| \leq 4$. On the other hand, since $\{1, 2\}$ is a subset of $\{1, 2, 3\}$, $4 = |X[1, 2]| \leq |X[1, 2, 3]| \leq 4$. Hence, $|X[1, 2, 3]| = |X[1, 2]| = 4$, and thus if any two inputs $x, x' \in X$ agree on positions 1 and 2, they also must agree on position 3. Analogously, we have $|X[0, 1, 2]| = |X[1, 2]| = 4$, and thus any two inputs that agree on positions 1 and 2 also agree on position 0. Now, since the bits at positions 1 and 2 completely determine the bits at positions 0 and 3, we have: $|X[0, 1, 2, 3]| = 4$.

Consider now the following vertex cover of G : $S = \{0, 1, 3, 5, 7, \dots, 2n-1\}$. We split it into two parts: $S_1 = \{0, 1, 3\}$ and $S_2 = \{5, 7, \dots, 2n-1\}$. We obtain: $|X[S]| \leq |X[S_1]| \cdot |X[S_2]|$. Since $S_1 \subseteq \{0, 1, 2, 3\}$, $|X[S_1]| \leq |X[0, 1, 2, 3]| = 4$. Since $|S_2| = n-2$, $|X[S_2]| \leq 2^{n-2}$. Therefore, $|X[S]| \leq 4 \cdot 2^{n-2} = 2^n$. Applying now Proposition 12, we have: $|X| \leq |X[S]| \leq 2^n$.

Case 3: $\forall i \in [0, 2n]$, $|X[i, i+1]| = 3$:

We split the analysis of this case into two subcases.

Subcase 3.1: $n \geq 4$: That is, G is an odd hole of length at least 9. Our goal in this case is to show that $|X[0, 1, \dots, 7]| \leq 16$. If we do that, then we can construct a vertex cover of G as follows: $S = S_1 \cup S_2$, where $S_1 = \{0, 1, 3, 5, 7\}$ and $S_2 = \{9, 11, \dots, 2n-1\}$. We obtain

$$\begin{aligned} |X| &\leq |X[S]| \leq |X[S_1]| \cdot |X[S_2]| \\ &\leq |X[0, 1, \dots, 7]| \cdot 2^{|S_2|} \leq 16 \cdot 2^{n-4} = 2^n. \end{aligned}$$

In order to bound $|X[0, 1, \dots, 7]|$, we develop a recursive expression for $|X[0, 1, \dots, i+1]|$, for any $i = 0, \dots, 2n-1$.

Fix some $i \in [0, 2n-1]$ and denote the three bitstrings in $X[i, i+1]$ by $a_i b_i$, $a_i \bar{b}_i$, and $\bar{a}_i b_i$, where $a_i, b_i \in \{0, 1\}$. Note

that exactly two of these bitstrings agree on position i and exactly two agree on position $i + 1$. a_i is the majority bit at position i and b_i is the majority bit at position $i + 1$. We split the set $X[0, 1, \dots, i + 1]$ into three parts accordingly

$$\begin{aligned} \mathcal{A}_i &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[i, i + 1] = a_i b_i\}, \\ \mathcal{B}_i &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[i, i + 1] = a_i \bar{b}_i\}, \text{ and} \\ \mathcal{C}_i &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[i, i + 1] = \bar{a}_i b_i\}. \end{aligned}$$

Clearly, $|X[0, 1, \dots, i + 1]| = |\mathcal{A}_i| + |\mathcal{B}_i| + |\mathcal{C}_i|$. We will develop recursive expressions for $|\mathcal{A}_i|, |\mathcal{B}_i|, |\mathcal{C}_i|$ and use them to bound $|X[0, 1, \dots, i + 1]|$.

Recall that $a_i b_i, a_i \bar{b}_i,$ and $\bar{a}_i b_i$ are the three bitstrings constituting $X[i, i + 1]$. We would like to explore next how these bitstrings can be extended into bitstrings in $X[i, i + 1, i + 2]$.

We next argue that $a_i b_i$ and $a_i \bar{b}_i$ cannot be extended using the same bit into bitstrings in $X[i, i + 1, i + 2]$. If there exists a bit c_i s.t. both $a_i b_i c_i$ and $a_i \bar{b}_i c_i$ belong to $X[i, i + 1, i + 2]$, then there exist two inputs $x, x' \in X$ s.t. $x[i, i + 1, i + 2] = a_i b_i c_i$ and $x'[i, i + 1, i + 2] = a_i \bar{b}_i c_i$. That is, $x[i + 1] \neq x'[i + 1]$, but x, x' agree on the two neighbors of $i + 1$. Therefore, x, x' are connected by an edge in the confusion graph $C(G)$, in contradiction to the assumption both belong to an independent set. We conclude that $a_i b_i$ and $a_i \bar{b}_i$ must be extended by complementary bits into bitstrings in $X[i, i + 1, i + 2]$.

The above implies that there exists some bit c_i s.t. $X[i, 1 + 1, i + 2]$ consists of the following bitstrings:

$$\{a_i b_i c_i, a_i \bar{b}_i \bar{c}_i, \bar{a}_i b_i \bar{c}_i\}.$$

It may optionally contain also the following bitstring:

$$\bar{a}_i b_i c_i.$$

(Note that it cannot consist only of the bitstrings $\{a_i b_i c_i, a_i \bar{b}_i \bar{c}_i, \bar{a}_i b_i \bar{c}_i\}$, because then $|X[i + 1, i + 2]| = 2$.)

Since there is a single bitstrings in $X[i, i + 1, i + 2]$ in which bit $i + 1$ is \bar{b}_i , b_i must be the majority bit at position $i + 1$ w.r.t. the bitstrings in $X[i + 1, i + 2]$. In other words, $a_{i+1} = b_i$. Similarly, it can be seen that \bar{c}_i must be the majority bit at position $i + 2$ w.r.t. the bitstrings in $X[i + 1, i + 2]$. Hence, $b_{i+1} = \bar{c}_i$.

We conclude the following:

- 1) If $x \in X$ and $x[i + 1, i + 2] = a_{i+1} b_{i+1} = b_i \bar{c}_i$, then necessarily $x[i, i + 1] = \bar{a}_i b_i$. Therefore, $|\mathcal{A}_{i+1}| \leq |\mathcal{C}_i|$.
- 2) If $x \in X$ and $x[i + 1, i + 2] = a_{i+1} \bar{b}_{i+1} = b_i c_i$, then $x[i, i + 1] = a_i b_i$ or $x[i, i + 1] = \bar{a}_i b_i$. Therefore, $|\mathcal{B}_{i+1}| \leq |\mathcal{A}_i| + |\mathcal{C}_i|$.
- 3) If $x \in X$ and $x[i + 1, i + 2] = \bar{a}_{i+1} b_{i+1} = \bar{b}_i \bar{c}_i$, then necessarily $x[i, i + 1] = a_i \bar{b}_i$. Therefore, $|\mathcal{C}_{i+1}| \leq |\mathcal{B}_i|$.

For $i = 0$, we have $|\mathcal{A}_0| = |\mathcal{B}_0| = |\mathcal{C}_0| = 1$. Hence, $\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i$ form Fibonacci-like series. Since we need the value of the series only at $i = 6$, we expand their prefixes explicitly in Table I.

The last row of the table gives our desired upper bound: $|X[0, 1, \dots, 7]| \leq 16$.

Subcase 3.2: $n = 3$: That is, G is a hole of length 7. We show that in this case $|X[0, 1, \dots, 5]| \leq 8$. This would imply that $|X| \leq |X[0, 1, 3, 5]| \leq |X[0, 1, \dots, 5]| \leq 8 = 2^n$.

By what we proved in the previous sub-case, we know $|X[0, 1, \dots, 5]| \leq 9$. We assume then, to reach a contradiction,

TABLE I
UPPER BOUNDS OBTAINED USING THE RECURSION

i	$ \mathcal{A}_i $	$ \mathcal{B}_i $	$ \mathcal{C}_i $	$ X[0, 1, \dots, i + 1] $
0	1	1	1	3
1	1	2	1	4
2	1	2	2	5
3	2	3	2	7
4	2	4	3	9
5	3	5	4	12
6	4	7	5	16

that $|X[0, 1, \dots, 5]| = 9$. We partition X into three sets as follows:

$$\begin{aligned} X_{\mathcal{A}} &= \{x \in X \mid x[0, 1] = a_0 b_0\}, \\ X_{\mathcal{B}} &= \{x \in X \mid x[0, 1] = a_0 \bar{b}_0\}, \text{ and} \\ X_{\mathcal{C}} &= \{x \in X \mid x[0, 1] = \bar{a}_0 b_0\}. \end{aligned}$$

We will prove that if $|X[0, 1, \dots, 5]| = 9$, then $|X_{\mathcal{A}}[5]| = |X_{\mathcal{B}}[5]| = |X_{\mathcal{C}}[5]| = 2$. That would imply that

$$\begin{aligned} |X[5, 6, 0, 1]| &= |X_{\mathcal{A}}[5, 6, 0, 1]| + |X_{\mathcal{B}}[5, 6, 0, 1]| + |X_{\mathcal{C}}[5, 6, 0, 1]| \\ &\geq 2 + 2 + 2 = 6. \end{aligned}$$

By symmetry, the upper bound we proved in the previous subcase on $|X[0, 1, 2, 3]|$ holds for any sequence of four consecutive positions. Therefore, $|X[5, 6, 0, 1]| \leq 5$, contradicting $|X[5, 6, 0, 1]| \geq 6$.

So how do we prove $|X_{\mathcal{A}}[5]| = |X_{\mathcal{B}}[5]| = |X_{\mathcal{C}}[5]| = 2$? Define, for any $i \in [0, 2n - 1]$, $\mathcal{A}_i^R = X_{\mathcal{A}}[0, 1, \dots, i + 1]$, $\mathcal{B}_i^R = X_{\mathcal{B}}[0, 1, \dots, i + 1]$, and $\mathcal{C}_i^R = X_{\mathcal{C}}[0, 1, \dots, i + 1]$. Note that

$$\begin{aligned} \mathcal{A}_i^R &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[0, 1] = a_0 b_0\}, \\ \mathcal{B}_i^R &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[0, 1] = a_0 \bar{b}_0\}, \text{ and} \\ \mathcal{C}_i^R &= \{x[0, 1, \dots, i + 1] \mid x \in X, x[0, 1] = \bar{a}_0 b_0\}. \end{aligned}$$

By reversing the order of indices and then applying the same argument as in the proof of the previous subcase, we can obtain that:

- 1) The upper bound on $|\mathcal{A}_i|$ applies to $|\mathcal{A}_i^R|$.
- 2) The upper bound on $|\mathcal{B}_i|$ applies to $|\mathcal{C}_i^R|$.
- 3) The upper bound on $|\mathcal{C}_i|$ applies to $|\mathcal{B}_i^R|$.

Since $|X[0, 1, \dots, 5]| = 9$ meets its upper bound, then also $|\mathcal{A}_i^R|, |\mathcal{B}_i^R|, |\mathcal{C}_i^R|$ must meet their upper bounds for $i = 0, 1, \dots, 4$. (By definition $|X[0, 1, \dots, i + 1]| = |\mathcal{A}_i| + |\mathcal{B}_i| + |\mathcal{C}_i|$, thus it follows from the relations between $|\mathcal{A}_{i+1}|, |\mathcal{B}_{i+1}|, |\mathcal{C}_{i+1}|$ and $|\mathcal{A}_i|, |\mathcal{B}_i|, |\mathcal{C}_i|$, and between $|\mathcal{A}_i|, |\mathcal{B}_i|, |\mathcal{C}_i|$ and $|\mathcal{A}_i^R|, |\mathcal{B}_i^R|, |\mathcal{C}_i^R|$.) We now show separately that $|X_{\mathcal{A}}[5]| = |X_{\mathcal{B}}[5]| = |X_{\mathcal{C}}[5]| = 2$:

- 1) Using Table I, we have $|\mathcal{A}_2^R| = 1$, while $|\mathcal{A}_4^R| = 2$. Therefore, $|X_{\mathcal{A}}[3]| = 1$ while $|X_{\mathcal{A}}[3, 4, 5]| = 2$. Since $\{3, 5\}$ is a vertex cover of $\{3, 4, 5\}$, then $|X_{\mathcal{A}}[3, 5]| \geq |X_{\mathcal{A}}[3, 4, 5]| = 2$. However, since $|X_{\mathcal{A}}[3]| = 1$, then $|X_{\mathcal{A}}[5]| = 2$.
- 2) Using Table I, we have $|\mathcal{B}_3^R| = 3$ while $|\mathcal{B}_4^R| = 4$. That is, $|X_{\mathcal{B}}[0, 1, 2, 3, 4]| = 3$ while $|X_{\mathcal{B}}[0, 1, 2, 3, 4, 5]| = 4$. This can happen only if $|X_{\mathcal{B}}[5]| = 2$.
- 3) Using Table I, we have $|\mathcal{C}_3^R| = 2$ while $|\mathcal{C}_4^R| = 3$. That is, $|X_{\mathcal{C}}[0, 1, 2, 3, 4]| = 2$ while $|X_{\mathcal{C}}[0, 1, 2, 3, 4, 5]| = 3$. This can happen only if $|X_{\mathcal{C}}[5]| = 2$.

■

REFERENCES

- [1] D. Slepian and J. K. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inf. Theory*, vol. IT-19, pp. 471–480, 1973.
 - [2] A. Wyner and J. Ziv, "A theorem on the entropy of certain binary sequences and applications I," *IEEE Trans. Inf. Theory*, vol. IT-19, pp. 769–771, 1973.
 - [3] A. Wyner, "A theorem on the entropy of certain binary sequences and applications II," *IEEE Trans. Inf. Theory*, vol. IT-19, pp. 772–777, 1973.
 - [4] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 1948.
 - [5] H. S. Witsenhausen, "The zero-error side information problem and chromatic numbers," *IEEE Trans. Inf. Theory*, vol. 22, pp. 592–593, 1976.
 - [6] N. Alon and A. Orlitsky, "Source coding and graph entropies," *IEEE Trans. Inf. Theory*, vol. 42, pp. 1329–1339, 1996.
 - [7] P. Koulgi, E. Tuncel, S. L. Regunathan, and K. Rose, "On zero-error source coding with decoder side information," *IEEE Trans. Inf. Theory*, vol. 49, pp. 99–111, 2003.
 - [8] Y. Birk and T. Kol, "Coding-on-demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients," *IEEE Trans. Inf. Theory*, vol. 52, pp. 2825–2830, 2006.
 - [9] W. H. Haemers, "On some problems of Lovász concerning the Shannon capacity of a graph," *IEEE Trans. Inf. Theory*, vol. 25, pp. 231–232, 1979.
 - [10] A. E. Brouwer and C. A. van Eijl, "On the p-ranks of the adjacency matrices of strongly regular graphs," *J. Algebr. Combinator.*, vol. 1, pp. 329–346, 1992.
 - [11] R. Peeters, "On the p-ranks of the adjacency matrices of distance-regular graphs," *J. Algebr. Combinator.*, vol. 15, no. 2, pp. 127–149, 2002.
 - [12] R. Peeters, "Orthogonal representations over finite fields and the chromatic number of graphs," *Combinatorica*, vol. 16, no. 3, pp. 417–431, 1996.
 - [13] W. H. Haemers, "An upper bound for the Shannon capacity of a graph," *Algebr. Methods in Graph Theory*, vol. 25, pp. 267–272, 1978.
 - [14] C. Berge, "Färbung von graphen, deren sämtliche bzw. deren ungerade kreise starr sind," *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, vol. 10, 1961.
 - [15] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas, "The strong perfect graph theorem," *Ann. Math.*, vol. 164, pp. 51–229, 2006.
 - [16] T. Feder, E. Kushilevitz, M. Naor, and N. Nisan, "Amortized communication complexity," *SIAM J. Computing*, vol. 24, no. 4, pp. 736–750, 1995.
 - [17] A. Chakrabarti, Y. Shi, A. Wirth, and A. C.-C. Yao, "Informational complexity and the direct sum problem for simultaneous message complexity," in *Proc. 42nd IEEE Ann. Symp. Found. Comput. Sci. (FOCS)*, 2001, pp. 270–278.
 - [18] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar, "An information statistics approach to data stream and communication complexity," *J. Comput. Syst. Sci.*, vol. 68, no. 4, pp. 702–732, 2004.
 - [19] Z. Bar-Yossef, T. S. Jayram, R. Krauthgamer, and R. Kumar, "The sketching complexity of pattern matching," in *Proc. 8th Int. Workshop on Randomiz. Computat. (RANDOM)*, 2004, pp. 261–272.
 - [20] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani, "Dense quantum coding and quantum finite automata," *J. ACM*, vol. 49, no. 4, pp. 496–511, 2002.
 - [21] I. Kremer, N. Nisan, and D. Ron, "On randomized one-round communication complexity," *Computat. Complex.*, vol. 8, no. 1, pp. 21–49, 1999.
 - [22] R. W. Yeung and Z. Zhang, "Distributed source coding for satellite communications," *IEEE Trans. Inf. Theory*, vol. 45, pp. 1111–1120, 1999.
 - [23] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inf. Theory*, vol. 46, pp. 1204–1216, 2000.
 - [24] Z. Bar-Yossef, Y. Birk, T. S. Jayram, and T. Kol, "Index coding with side information," in *Proc. 47th IEEE Ann. Symp. Found. Comput. Sci. (FOCS)*, 2006, pp. 197–206.
 - [25] E. Lubetzky and U. Stav, "Non-linear index coding outperforming the linear optimum," *Preprint*, 2007.
 - [26] S. Y. E. Rouayheb, A. Sprintson, and C. N. Georghiades, "On the index coding problem and its relation to network coding and matroid theory," *CoRR*, vol. abs/0810.0068, 2008.
 - [27] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
 - [28] R. M. Fano, *Class Notes for Transmission of Information* MIT, Cambridge, MA, 1952, course 6.574.
 - [29] L. Lovász, "A characterization of perfect graphs," *J. Combinator. Theory, Ser. B*, vol. 13, pp. 95–98, 1972.
 - [30] Y. Birk and T. Kol, "Informed-source coding-on-demand (ISCOD) over broadcast channels," in *Proc. IEEE INFOCOM '98*, 1998, vol. 13, pp. 1257–1264.
- Ziv Bar-Yossef** received the B.Sc. degree in computer science and mathematics and the M.Sc. degree in computer science from the Hebrew University, Jerusalem, Israel. He received the Ph.D. degree in computer science from the University of California at Berkeley.
- His main research interests are web search, information retrieval, data mining, computational complexity, and algorithms. He currently works as a Research Scientist and Technical Lead Manager at Google's Engineering Center, Israel. He holds an Adjunct Senior Lecturer position with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa, Israel. Prior to that, he was a full-time Senior Lecturer with the Technion and a Research Staff Member with the IBM Almaden Research Center, San Jose, CA. He is an author of more than 30 publications and five patents.
- Dr. Bar-Yossef won the Best Paper Award at the WWW 2006 conference and was a runner-up for the Best Application Paper Award at ICDM 2006.
- Yitzhak Birk** (M'82–SM'02) received the B.Sc. (*cum laude*) and M.Sc. degrees from the Technion—Israel Institute of Technology, Haifa, Israel, in 1975 and 1982, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, in 1987, all in electrical engineering.
- He has been on the faculty of the Electrical Engineering Department, Technion, since October 1991, and heads its Parallel Systems Laboratory. He is also Director of the Center For Information and Communication Technologies. From 1976 to 1981, he was Project Engineer with the Israel Defense Forces. From 1986 to 1991, he was a Research Staff Member with IBM's Almaden Research Center, where he worked on parallel architectures, computer subsystems, and passive fiber-optic interconnection networks. From 1993 to 1997, he also served as a consultant to HP Labs in the areas of storage systems and video servers, and was later involved with several companies. His research interests include computer systems, with much attention to the storage subsystem and to the interplay between storage and communication, e.g., in information dissemination systems. The true application requirements are considered in each case. The judicious exploitation of redundancy for performance enhancement in these contexts has been the subject of much of his recent work. He is also engaged in research into various facets of processor architecture, including many-core systems, attempting "cross fertilization" between his various areas of research. Recently, he has taken interest in algorithmic issues of electric-car charging infrastructure.
- T. S. Jayram** received the B.Tech. degree in computer science from the Indian Institute of Technology Madras in 1990 and the M.S. and Ph.D. degrees in computer science and engineering in 1996 and 1998, respectively, from the University of Washington, Seattle.
- His research interests are in massive data sets, probabilistic databases, and computational complexity. He joined IBM Research in 1998 and is currently a manager of the Algorithms and Computation group at IBM Almaden Research Center, San Jose, CA.
- Dr. Jayram received the IBM Outstanding Innovation Award for his contributions to the foundations of massive data sets. Previously, he received the Machtey award for Best Student Paper in FOCS 1995.
- Tomer Kol** (M'98) received the B.Sc. (*summa cum laude*), M.Sc., and Ph.D. degrees in electrical engineering, all from the Technion—Israel Institute of Technology, Haifa, in 1994, 1997, and 2007, respectively.
- His research interests include parallel and distributed architectures (HW & SW) for information systems, and algorithms. He has been working since 2007 as a software engineer and technical lead with Google's Engineering Center, Israel. From 2002 to 2007, he was a Research Staff Member with IBM's Haifa Research Lab, where he worked on medical IT. From 2000 to 2002, he was with a start-up company developing routers. Prior to that, he worked as an intern with the Intel Development Center, Haifa, and with IBM's Haifa Research Lab.