On index coding with side information

Dau, Son Hoang

2012

Dau, S. H. (2012). On index coding with side information. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/48642>

<https://doi.org/10.32657/10356/48642>

**NANYANG
TECHNOLOGICAL
UNIVERSITY**

**ON INDEX CODING WITH SIDE
INFORMATION**

DAU SON HOANG

SCHOOL OF PHYSICAL & MATHEMATICAL SCIENCES

2012

ON INDEX CODING WITH SIDE INFORMATION

DAU SON HOANG

School of Physical & Mathematical Sciences

A thesis submitted to the Nanyang Technological University
in partial fulfillment of the requirement for the degree of
Doctor of Philosophy

2012

ACKNOWLEDGEMENTS

This thesis would not have been possible without the help and support of the very kind people.

First and foremost, I would like to express my sincerest gratitude to my supervisors, Professor Chee Yeow Meng and Professor Ling San, for their priceless guidance and support, in both personal aspect and academic aspect. I am extremely grateful for their patience, kindness, and attentiveness towards me during my time at NTU.

I would also like to give a special thank to Vitaly Skachek, who worked very closely to me during my time as a PhD student. It was him who first introduced the Index Coding with Side Information problem to me. Many of the questions, ideas, and suggestions he put to me have actually been turned into important theorems in this thesis. It was the fruitful discussions with Vitaly and Yeow Meng that shaped the main content of the thesis. Vitaly's generous supports and encouragements have been invaluable to me.

I am indebted to many of my friends for their helps with computer related issues, among whom are Birk Bernd Eisermann, Nguyen Luong Dong, Nguyen Quoc Chinh, Punarbasu Purkayastha, Radoslav Kirov, Tran Duc Hieu, and Wei Lei. I also thank Melvin for his help with the use of SPMS servers.

Finally, I thank my Mom and my Dad for their endless encouragements and supports. Their love makes me feel more confident and stronger.

ABSTRACT

Building communication schemes which allow participants to communicate efficiently has always been a challenging yet intriguing problem for information theorists. Index Coding with Side Information (ICSI), first introduced by Birk and Kol (1998) as Coding on Demand by an Informed Source (ISCOD), is a communication scheme dealing with a broadcast channel in which receivers have prior side information about the messages to be broadcast. Exploiting the knowledge about the side information, the sender may significantly reduce the number of required transmissions compared with the naive approach. As a consequence, the efficiency of the communication over this type of broadcast channels could be dramatically improved. Apart from being a special case of the well-known (non-multicast) Network Coding problem (Alhswede *et al.* (2000), Koetter and Médard (2003)), the ICSI problem has also found various potential applications on its own, such as audio- and video-on-demand, daily newspaper delivery, data pushing, and opportunistic wireless networks.

While most of the known works on ICSI focus on the performances of index codes of various kinds, the security and error correction aspects of those codes have never been examined. Although the issues of security and error correction have been thoroughly investigated in multicast Network Coding, on the side of non-multicast Network Coding, where ICSI is one example, those issues are much less understood. In this work, we make an attempt to fill in this gap.

We first analyze the *block security* level of the linear index code based on a matrix $\mathbf{L}$, by examining the minimum distance and the dual distance of the column space $\mathcal{C}(\mathbf{L})$ of $\mathbf{L}$. Here the block security is a generalization of the notion of *weak security*, introduced by Bhattad and Narayanan (2005) in the setups of Network Coding. It turns out that these two distances of $\mathcal{C}(\mathbf{L})$ specify two closely related thresholds, one is for the code to be block secure, the other is for it to be completely insecure. When $\mathcal{C}(\mathbf{L})$ is an MDS code, the two thresholds are actually tight. We then proceed to examine the *strong security* of linear index codes, which is the information-theoretic security in the context of ICSI. We show that the coset coding technique, an indispensable tool in securing a network code in the setting of Network Coding, also provides us with an optimal solution for securing an index code.

Subsequently, we examine the index coding schemes in which errors are involved. Several bounds and constructions for linear error-correcting index codes are established. Syndrome decoding for error-correcting index codes, which is somewhat different in nature from that for classical error-correcting codes, is also investigated. Furthermore, we study the so-called *static* error-correcting index codes. Each of such codes works for a family of instances of the ICSI problem. Analogous bounds and constructions for nonlinear error-correcting index codes, as for linear case, are also developed.

The second line of work in this thesis is on the computational aspects of the ICSI problem. As shown in the work of Bar-Yossef *et al.* (2006), the minimum number of transmissions required by a linear index code is equal to the so-called *minrank* of the corresponding *side information digraph*. However, as shown by Peeters (1996), finding the minrank of a *side information graph* is an NP-hard problem. There are very few known families of (di)graphs for which minranks can be found in polynomial time. They are perfect graphs, odd holes (cycles) and odd anti-holes, outerplanar

graphs, and acyclic digraphs. In this work, we discover several more such families, including connectively reducible digraphs, line digraphs of partially planar digraphs, and graphs possessing a special tree structure. We also characterize (di)graphs of extreme minranks, that is, (di)graphs with minranks close to one or to their orders (i.e., number of vertices). Based on one of these characterizations, we are able to show that deciding whether a *digraph* has minrank two is an NP-complete problem. For comparison, the same decision problem for *graphs* can be solved in polynomial time. Finally, based on the work of Chaudhry and Sprintson (2008), we write a computer program to find the minranks of *all* non-isomorphic graphs of orders up to 10.

CONTENTS

1. <i>Introduction</i>	1
1.1 Background	1
1.2 Contributions of this Thesis	3
1.3 Organization of this Thesis	12
1.4 Notation and Definitions	12
1.4.1 Coding Theory Terminology	12
1.4.2 Linear Algebra Terminology	14
1.4.3 Graph Theory Terminology	15
1.4.4 Information Theory Terminology	19
2. <i>Secure Index Code with Side Information</i>	21
2.1 Index Coding and Some Basic Results	22
2.2 Block Secure Linear Index Codes	26
2.2.1 Block Security and Weak Security	26
2.2.2 Necessary and Sufficient Conditions for Block Security	28
2.2.3 Complete Insecurity	38
2.3 Index Coding with Side and Restricted Information	42
2.4 Strongly Secure Index Code with Side Information	43
2.4.1 A Lower Bound on the Length of a Strongly Secure Index Code	44
2.4.2 A Construction of Optimal Strongly Secure Index Codes	50

3. Index Code and Error Correction	54
3.1 Error-Correcting Index Code with Side Information	55
3.2 The α -Bound and the κ -Bound	61
3.3 The Singleton Bound	69
3.4 Random codes	71
3.5 Strongly Secure Linear Error-Correcting Index Codes	73
3.5.1 A Lower Bound on the Length of a Strongly Secure Linear Error-Correcting Index Code	74
3.5.2 A Construction of an Optimal Strongly Secure Linear Error- Correcting Index Code	75
3.6 Syndrome Decoding	77
3.7 Static Error-Correcting Index Codes	82
3.7.1 Static Error-Correcting Index Codes	82
3.7.2 An Application: Weakly Resilient Functions	84
3.7.3 Bounds and Constructions	87
3.8 Nonlinear Error-Correcting Index Codes	92
4. (Di)Graphs of Extreme MinRanks	97
4.1 (Strongly) Connected Components and MinRanks	97
4.2 (Di)Graphs of Extreme MinRanks	100
4.2.1 (Di)Graphs of MinRank One	100
4.2.2 (Di)Graphs of MinRank Two	100
4.2.3 (Di)Graphs of MinRanks Equal to Their Orders	103
4.2.4 Graphs of MinRanks One Less Than Their Orders	105
5. Computation of MinRanks of Side Information (Di)Graphs	109
5.1 The Hardness of the MinRank Problem for Digraphs	109

5.2	Circuit-Packing Bound	114
5.2.1	The Bound	114
5.2.2	Digraphs Attaining Circuit-Packing Bound	118
5.3	On MinRanks of Graphs Having Tree Structures of Type I	127
5.3.1	Tree Structure of Type I	127
5.3.2	An Algorithm for MinRanks of Graphs in $\mathcal{F}_1(c)$	129
5.3.3	An Algorithm to Recognize a Graph in $\mathcal{F}_2(c)$	141
5.4	MinRanks of (Di)Graphs of Small Orders	145
6.	Conclusion	147
	Bibliography	151
	Appendix	159
A.	An Exact Algorithm for Finding MinRanks of (Di)Graphs over $\mathbb{F}_2$	160
A.1	The Algorithm	160
A.2	Reduction to SAT Problem	163
A.3	Generating Lists of Non-Isomorphic Graphs	169
B.	MinRanks of Graphs of Small Orders	171
B.1	MinRanks of Graphs of Orders One to Four	171
B.2	MinRanks of Graphs of Order Five	172
B.3	MinRanks of Graphs of Order Six	173
	Index	182

1. INTRODUCTION

1.1 Background

The problem of Index Coding with Side Information (ICSI) was introduced by Birk and Kol [10, 11]. The coding scheme studied in the ICSI problem was originally called *Coding on Demand by an Informed Source* (ISCOD). It was motivated by applications such as audio and video-on-demand, and daily newspaper delivery, in which a server (sender) broadcasts a set of messages to a set of clients (receivers). During the transmission, each client might miss a certain part of the data, due to intermittent reception, limited storage capacity or any other reasons. The server has to find a way to deliver to each client all the missing messages, yet spending a minimum number of transmissions. Via a slow backward channel, the clients let the server know which messages they already have in their possession, and which messages they request. By exploiting this information, the amount of the overall transmissions can be significantly reduced [11].

Another possible application of the ICSI problem is in opportunistic wireless networks. These are the networks in which a wireless node can opportunistically listen to the wireless channel. As a result, a node may obtain packets that were not designated to it [28, 41, 42]. In this way, a node obtains some side information about the transmitted data. Exploiting this additional knowledge may help to increase the throughput of the system.

Consider the toy example in Fig. 1.1. It presents a scenario with one sender and four receivers. Each receiver requires a different information packet (or message). The naive approach requires four separate transmissions, one transmission per an information packet. However, by utilizing the knowledge about the subsets of messages that receivers already have, and by coding the transmitted data, the sender can satisfy all the demands by broadcasting just one coded packet.

The ICSI problem has been the subject of several recent studies [2,3,8,15,16,21,22,25,27,28,36,46,69]. On the one hand, this problem is a special case of the well-known (non-multicast) network coding (NC) problem [1,43]. On the other hand, it was shown that every instance of the NC problem can be reduced to an instance of the ICSI problem in the following sense. For each NC instance, we can construct an ICSI instance such that there exists a scalar linear network code for the NC instance if and only if there exists a perfect scalar linear index code for the corresponding ICSI instance (see [27,28] for more details).

Fig. 1.1: An example of the ICSI problem

Several previous works focused on the design of efficient index codes. Given an

instance of the ICSI problem, Bar-Yossef *et al.* [3, 4] proved that finding the best *binary scalar linear index code* is equivalent to finding the so-called *minrank* of a (di)graph (together with a matrix whose rank is equal to the minrank), which is an NP-hard problem [3, 53]. Here scalar linear index codes refer to linear index codes in which each message is a symbol in the underlying field $\mathbb{F}_q$. By contrast, in *vector linear index codes* each message is a vector over $\mathbb{F}_q$. Lubetzky and Stav [46] showed that there exist instances for which scalar linear index codes over nonbinary fields and combinations of linear index codes over two different fields outperform binary scalar linear index codes. El Rouayheb *et al.* [27, 28] showed that for certain instances of the ICSI problem, vector linear index codes achieve strictly higher transmission rate than scalar linear index codes. They also pointed out that there exist instances in which vector nonlinear index codes outperform vector linear index codes. Vector nonlinear index codes were also shown to outperform scalar nonlinear index codes for certain instances by Alon *et al.* [2]. Several heuristic solutions for the ICSI problem were proposed by El Rouayheb *et al.* [25] and by Chaudhry and Sprintson [16]. Optimal scalar linear index codes for an ICSI instance described by a random side information (di)graph was investigated by Haviv and Langberg [36]. A dynamic programming approach (Bellman [5]) was established by Berliner and Langberg [8] to find in polynomial time optimal scalar linear index codes for ICSI instances described by outerplanar graphs.

1.2 Contributions of this Thesis

Security issues and computational issues of the index coding schemes are the main topics to be discussed in this thesis.

In Chapter 2, we study the security aspects of linear index codes. This is the

first known work in this direction. We restrict ourselves to *scalar linear* index codes due to the following reasons. First, although nonlinear codes are theoretically interesting, they are usually not practically attractive, due to the high complexity in encoding and decoding. It is known that vector linear index codes can achieve better transmission rates than their scalar counterparts for certain instances of the ICSI problem [27, 28]. However, if the block length of a vector index code for an instance is fixed, we can model this vector index code as a scalar index code for another instance of the ICSI problem (see the proof of Proposition 5.2.16).

A linear index code maps $\mathbf{x} \in \mathbb{F}_q^n$ onto $\mathbf{x}\mathbf{L}$, where $\mathbf{L}$ is an $n \times N$ matrix over $\mathbb{F}_q$, and $n, N \in \mathbb{N}$. Here $\mathbb{F}_q$ denotes the finite field with q elements. The number of transmissions required when such an index code is used is N . We call N the *length* of this index code.

We show that each linear index code provides a certain level of information security. We describe this result in more details as follows. Let the code $\mathcal{C}(\mathbf{L})$ be spanned by the columns of $\mathbf{L}$, and let d and $d^\perp$ be its minimum distance and dual distance, respectively. Assume that there exists an adversary who can listen to all transmissions from the sender. We say that a particular adversary is of strength t if it has t messages in its possession. An index code is said to be *b-block secure* against all adversaries of strength t if every adversary of strength t has no information (in Shannon's sense) about any block of b messages that are not in his possession. In contrast, an index code is said to be *completely insecure* against any adversary of strength t if an adversary who possesses an arbitrary set of t messages is always able to reconstruct all of the other $n - t$ messages. Then, we show that the index code based on $\mathbf{L}$ is $(d - 1 - t)$ -block secure against all adversaries of strength $t \leq d - 2$ and is completely insecure against any adversary of strength at least $n - d^\perp + 1$ (Theorems 2.2.9 and 2.2.12). If $\mathcal{C}(\mathbf{L})$ is an MDS code, then the two bounds

coincide. The technique used in the proof for this result is reminiscent of that used in the constructions of (multiple) secret sharing schemes from linear error-correcting codes [23, 50]. The results on the security of linear index codes can be further employed to analyze the existence of solutions for a natural generalization of the ICSI problem, so-called the *Index Coding with Side and Restricted Information (ICSRI)* problem. In that problem, it is required that some receivers have no information about some messages.

In the sequel, we also consider linear randomized index codes, which are based on the use of random symbols. Using such index codes, the sender first generates random symbols, mixes them with the original packets, and then broadcasts the mixed packets. We show that the coset coding technique (which has been successfully employed in Secure Network Coding literature [12, 26, 29, 59, 71]) yields an optimal *strongly secure* linear randomized index code of length $\kappa_q + \mu$ (Theorems 2.4.10). This randomized index code is strongly secure against an adversary that

- (i) has t arbitrary messages in advance;
- (ii) eavesdrops at most μ transmissions.

In other words, such an adversary gains no information (in Shannon's sense) about the messages that he does not possess. Observe that the length of the optimal code does not depend on t . In fact, the construction of such a code is independent of t . Here κ_q denotes the minrank over $\mathbb{F}_q$ of the side information hypergraph that corresponds to the instance of the ICSI problem. As shown later in Corollary 2.2.8, κ_q is equal to the length of an optimal scalar linear index code over $\mathbb{F}_q$.

It is worth mentioning that the security models considered in this work (block security and strong security) are special cases of a more general model referred to as *security against guessing* in the context of Network Coding (see Bhattad and

Narayanan [9]). In their model, an adversary can obtain some set M of linear combinations of the (source) messages by eavesdropping on a certain set of links in the communication network that connects the sender and the receivers. Let $\{G_p, U_p\}$ ($p \in P$) be pairs of sets of messages. The adversary can *guess* perfectly all messages in G_p ($p \in P$). In other words, the adversary has some side information about the messages. The security requirement is that the adversary has no information (in Shannon's sense) about U_p given his knowledge of M (eavesdropped information) and G_p (side information) for every $p \in P$. When G_p is any subset of t messages and U_p is any subset of b messages ($G_p \cap U_p = \emptyset$), this model reduces to the b -block security model. When G_p is any subset of t messages and U_p is the set of the remaining $n - t$ messages, the model reduces to the strong security model.

Most of previous works on the security aspects (and on the error-correction aspect, as a special case) of network coding dealt with the multicast scenario. One of the main reasons for this limitation is that the optimal simultaneous transmission rates for non-multicast networks have not been well characterized yet. It is known that the ICSI problem can be modeled as a special case of the non-multicast Network Coding problem [2, 28]. This fact restricts the ability to derive the results on the security of the index coding schemes from the existing results on the security of network coding schemes.

The preceding works on the ICSI problem consider a scenario where the transmissions are error-free. In practice, of course, this might not be the case. In Chapter 3, we assume that the transmitted symbols are subject to errors and extend some known results on index coding to a case where each receiver can correct up to a certain number of errors. It turns out that the problem of designing such error-correcting index codes (ECIC's) naturally generalizes the problem of constructing classical error-correcting codes.

More specifically, assume that the number of messages that the sender possesses is n , and that the maximum number of errors is δ . We show that the problem of constructing an ECIC of minimum length is equivalent to the problem of constructing a matrix $\mathbf{L}$ that has n rows and minimum number of columns, such that

$$\text{wt}(\mathbf{zL}) \geq 2\delta + 1 \text{ for all } \mathbf{z} \in \mathcal{I},$$

where $\mathcal{I}$ is a certain subset of $\mathbb{F}_q^n \setminus \{\mathbf{0}\}$ (Lemma 3.1.3). Here $\text{wt}(\mathbf{x})$ denotes the Hamming weight of the vector $\mathbf{x}$ and $\mathbf{0}$ is the all-zeros vector. If $\mathcal{I} = \mathbb{F}_q^n \setminus \{\mathbf{0}\}$, this problem becomes equivalent to the well-known problem of designing a shortest-length linear error-correcting code of given dimension and minimum distance.

We establish a lower bound (the α -bound) and an upper bound (the κ -bound) on the shortest length of a linear ECIC that is able to correct every error pattern of size up to δ (Theorem 3.2.5 and Proposition 3.2.8). These bounds are described in more details as follows. Let $\mathcal{H}$ be the side information hypergraph that describes an instance of the ICSI problem. Let $\mathcal{N}_q[\delta, \mathcal{H}]$ denote the minimum length of a linear ECIC over $\mathbb{F}_q$ which ensures that every receiver R_i can recover the desired message if the number of errors is at most δ . We use notation $N_q[k, d]$ for the minimum length of a linear error-correcting code of dimension k and minimum distance d over $\mathbb{F}_q$. The α -bound and the κ -bound are the following

$$N_q[\alpha(\mathcal{H}), 2\delta + 1] \leq \mathcal{N}_q[\delta, \mathcal{H}] \leq N_q[\kappa_q(\mathcal{H}), 2\delta + 1], \quad (1.1)$$

where $\alpha(\mathcal{H})$ is the generalized independence number and $\kappa_q(\mathcal{H})$ is the minrank (over $\mathbb{F}_q$) of $\mathcal{H}$. The κ -bound (the second inequality) is obtained simply by concatenating an optimal linear traditional error-correcting code and an optimal linear index code.

For linear index codes, we also derive an analog of the Singleton bound (Theorem 3.3.1). This result implies that (over sufficiently large alphabets) the concatenation of a standard MDS error-correcting code with an optimal linear index code yields an optimal linear error-correcting index code. Finally, we consider a random ECIC. By analyzing its parameters, we obtain an upper bound on its length (Theorem 3.4.1).

When the side information hypergraph is a pentagon, and $\delta = 2$, the inequalities in (1.1) are shown to be strict. This implies that a concatenated scheme based on a classical error-correcting code and on a linear non-error-correcting index code does not necessarily yield an optimal linear error-correcting index code. Since the ICSI problem can also be viewed as a source coding problem [2, 3], this example demonstrates that sometimes designing a single code for both source and channel coding can result in a smaller number of transmissions.

Based on the results obtained in Chapter 2 and the results on ECIC's in this chapter (Chapter 3), we give a construction of an optimal strongly secure linear randomized error-correcting index code of length $\kappa_q + \mu + 2\delta$ (Theorem 3.5.5). This randomized index code is strongly secure against an adversary that

- (i) has t arbitrary messages in advance;
- (ii) eavesdrops at most μ transmissions;
- (iii) corrupts at most δ transmissions.

The decoding of a linear ECIC is somewhat different from that of a classical error-correcting code. There is no longer a need for a complete recovery of the whole information vector. As a consequence, each receiver no longer has to determine the exact error vector in order to decode the desired message. We analyze the

decoding criteria for the ECIC's and show that the syndrome decoding, which might be different for each receiver, outputs the correct result, provided that the number of errors does not exceed the error-correcting capability of the code.

An ECIC is called static under a family of instances of the ICSI problem if it works for all of these instances. Such an ECIC is desirable since it remains relevant as long as the parameters of the problem vary within a particular range. Bounds and constructions for static ECIC's are studied in Section 3.7. Connections between static ECIC's and weakly resilient vectorial Boolean functions are also investigated.

At the end of Chapter 3, in Section 3.8, we briefly discuss the nonlinear ECIC's and establish several analogous results for bounds on minimum lengths of nonlinear ECIC's.

As shown by Bar-Yossef *et al.* [3, 4], the minrank of a digraph is precisely the number of transmissions required in an optimal scalar linear index code for an instance of the ICSI problem described by that digraph. The concept of *minrank* of a graph was first introduced by Haemers [34], which serves as an upper bound for the celebrated Shannon capacity of a graph [57]. This upper bound, as pointed out by Haemers himself, although is usually not as good as the Lovász bound [45], is sometimes tighter and easier to compute. Unfortunately, as shown by Peeters [53], computing the minrank of a general graph (that is, the MinRank problem) is a hard task. More specifically, Peeters showed that deciding whether a graph has minrank three is an NP-complete problem. The interest on the MinRank problem was resumed after the work of Bar-Yossef *et al.* [3], which proved that the binary minrank of a digraph $\mathcal{D}$ characterizes the optimal *binary scalar linear* index code for an ICSI instance with side information digraph $\mathcal{D}$. Exact and heuristic algorithms to find the minrank over $\mathbb{F}_2$ of a hypergraph (and a (di)graph as a special case) were developed in the work of Chaudhry and Sprintson [16]. The minranks of random

(di)graphs are investigated by Haviv and Langberg [36]. A dynamic programming approach was proposed by Berliner and Langberg [8] to compute in polynomial time the minranks of outerplanar graphs. It is also worth noting that approximating minranks of graphs within any constant ratio is known to be NP-hard (see Langberg and Sprintson [44]).

In Chapter 4, we make an attempt to characterize the (di)graphs that have extreme minranks. In particular, it is shown that a digraph has minrank two over $\mathbb{F}_2$ if and only if $\mathcal{D}$ is not a complete digraph and its complement, $\overline{\mathcal{D}}$, is *fairly 3-colorable* (Corollary 4.2.5). Here a digraph is fairly 3-colorable if one can color its vertices by three colors in such a way that not only the endpoints of the same arc have different colors but also all out-neighbors of the same vertex must share the same color. Based on this characterization, we show later in Chapter 5 that the problem of deciding whether a digraph has minrank two is NP-complete. In contrast, for the case of graphs, it is known that the minrank of a graph $\mathcal{G}$ is two if and only if $\mathcal{G}$ is not a complete graph and its complement, $\overline{\mathcal{G}}$, is bipartite, a condition which can be verified in polynomial time (see, for instance West [67, p. 495]). Note that a graph is bipartite if and only if it is 2-colorable.

The other characterizations are the following. A graph has minrank equal to its order if and only if it contains no edges. By contrast, a digraph has minrank equal to its order if and only if it contains no circuits. A graph of order n has minrank $n - 1$ if and only if its largest connected component is a star graph and the other components are one-vertex graphs.

In Chapter 5, we focus on the computational aspects of the ICSI problem. We first establish that deciding whether a digraph has minrank two is already an NP-complete problem (Corollary 5.1.2). Afterwards, a new upper bound on the minrank of a digraph, so-called the circuit-packing bound, is introduced (Proposition 5.2.3).

In certain cases, the circuit-packing bound is shown to be far better than the clique-cover bound. So far, families of (di)graphs whose minranks are either known or computable in polynomial time are the following. For *graphs*, they are odd holes and odd anti-holes, perfect graphs, and outerplanar graphs. For *digraphs*, they are acyclic digraphs. The circuit-packing bound, together with the lower bound based on the order of a maximum acyclic subgraph, enables us to derive several new families of digraphs whose minranks are computable in polynomial time. They are the families of connectively reducible digraphs and line digraphs of partially planar digraphs. For those two families of digraphs, the circuit-packing bound is tight. This bound is also tight for other families of *digraphs*, including digraphs that *pack* (a digraph packs if all of its subgraphs satisfy the *min-max vertex equality*, see Section 5.2.2 for more details), line digraphs of fully reducible flow digraphs, and line digraphs of special Eulerian digraphs. Moreover, for ICSI instances described by the aforementioned families of digraphs, scalar linear index codes are shown to be optimal (Proposition 5.2.16).

Based on the fact that outerplanar graphs can be decomposed into tree structures that support a dynamic programming approach, Berliner and Langberg [8] established an efficient algorithm to find minranks of outerplanar graphs. Inspired by their work, at the end of Chapter 5, we develop a dynamic programming algorithm to compute in polynomial time the minranks of graphs having a special structure. Such a graph can be described as a compound rooted tree, the nodes of which are graphs whose minranks can be computed in polynomial time. The algorithm computes the minranks of the subtrees from the leaves to the root, in a bottom-up manner. The task of computing the minrank of the graph is accomplished when the computation reaches the root of the compound tree. We end Chapter 5 by mentioning that the minranks of *all* non-isomorphic graphs of order up to 10

can be found using a computer program that combines a SAT-based approach and a Branch-and-Bound approach.

1.3 Organization of this Thesis

The main content of the thesis is organized into six chapters.

Chapter 1 provides a background on the ICSI problem, which is the main subject of this work. It also summarizes the key results obtained in the thesis and contains most of the basic definitions and notation used throughout the thesis.

Chapters 2 and 3 discuss the security aspects of scalar linear index codes. Bounds and constructions for error-correcting index codes and strongly secure index codes are established. A decoding algorithm for error-correcting index codes is proposed. Most of the results in these two chapters contain the minrank of a side information hypergraph (and (di)graph as a special case) in their formulations.

A deeper investigation of the minranks of side information (di)graphs is carried out in Chapters 4 and 5. The structures of (di)graphs having extreme minranks are examined. New families of (di)graphs whose minranks can be found in polynomial time are discovered. Details on the computation of minranks of graphs of small orders can be found in the Appendix.

The thesis is concluded in Chapter 6. Several intriguing open problems for future research are proposed.

1.4 Notation and Definitions

1.4.1 Coding Theory Terminology

We use $\mathbb{F}_q$ to denote the finite field of q elements, where q is a power of prime. We denote by $\mathbb{F}_q^*$ the set of all nonzero elements of $\mathbb{F}_q$.

Let $[n] = \{1, 2, \dots, n\}$. For the vectors $\mathbf{u} = (u_1, u_2, \dots, u_n) \in \mathbb{F}_q^n$ and $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \mathbb{F}_q^n$, the (Hamming) distance between $\mathbf{u}$ and $\mathbf{v}$ is defined to be the number of coordinates where $\mathbf{u}$ and $\mathbf{v}$ differ, namely,

$$d(\mathbf{u}, \mathbf{v}) = |\{i \in [n] : u_i \neq v_i\}|.$$

If $\mathbf{u} \in \mathbb{F}_q^n$ and $M \subseteq \mathbb{F}_q^n$ is a set of vectors (for instance, a vector subspace), then the last definition can be extended to

$$d(\mathbf{u}, M) = \min_{\mathbf{v} \in M} d(\mathbf{u}, \mathbf{v}).$$

The *support* of a vector $\mathbf{u} = (u_1, u_2, \dots, u_n) \in \mathbb{F}_q^n$ is defined to be the set

$$\text{supp}(\mathbf{u}) = \{i \in [n] : u_i \neq 0\}.$$

The (Hamming) weight of a vector $\mathbf{u}$, denoted $\text{wt}(\mathbf{u})$, is defined to be $|\text{supp}(\mathbf{u})|$, the number of nonzero coordinates of $\mathbf{u}$. Suppose $E \subseteq [n]$. We write $\mathbf{u} \triangleleft E$ whenever $\text{supp}(\mathbf{u}) \subseteq E$.

A k -dimensional subspace $\mathcal{C}$ of $\mathbb{F}_q^n$ is called a linear $[n, k, d]_q$ code if the minimum distance of $\mathcal{C}$,

$$d(\mathcal{C}) \triangleq \min_{\mathbf{u} \in \mathcal{C}, \mathbf{v} \in \mathcal{C}, \mathbf{u} \neq \mathbf{v}} d(\mathbf{u}, \mathbf{v}),$$

is equal to d . Sometimes we may use the notation $[n, k]_q$ for the sake of simplicity. We call n the *length* and k the *dimension* of the code. The vectors in $\mathcal{C}$ are called codewords. It is easy to see that the minimum weight of a nonzero codeword in a linear code $\mathcal{C}$ is equal to its minimum distance $d(\mathcal{C})$. A *generator matrix* $\mathbf{G}$ of an $[n, k]_q$ code $\mathcal{C}$ is a $k \times n$ matrix whose rows are linearly independent codewords of

$\mathcal{C}$. Then $\mathcal{C} = \{\mathbf{y}\mathbf{G} : \mathbf{y} \in \mathbb{F}_q^k\}$. The *parity-check matrix* of $\mathcal{C}$ is an $(n-k) \times n$ matrix $\mathbf{H}$ over $\mathbb{F}_q$ such that $\mathbf{c} \in \mathcal{C} \Leftrightarrow \mathbf{H}\mathbf{c}^T = \mathbf{0}^T$. Given q , k , and d , let $N_q[k, d]$ denote the length of the shortest linear code over $\mathbb{F}_q$ that has dimension k and minimum distance d .

The *dual code* or *dual space* of $\mathcal{C}$ is defined as $\mathcal{C}^\perp = \{\mathbf{u} \in \mathbb{F}_q^n : \mathbf{u}\mathbf{c}^T = 0 \text{ for all } \mathbf{c} \in \mathcal{C}\}$. The minimum distance of $\mathcal{C}^\perp$, namely $d(\mathcal{C}^\perp)$, is called the *dual distance* of $\mathcal{C}$.

An $(n, M, d)_q$ code is an M -subset $\mathcal{C}$ of $\mathbb{F}_q^n$ that satisfies the property that the (Hamming) distance between every two distinct vectors of $\mathcal{C}$ is at least d . We call M the *size* of the code. We, again, call n and d the length and the distance of the code, respectively. Given q , M , and d , let $N_q(M, d)$ denote the length of a shortest code of size M and distance d .

The following upper bound on the minimum distance of a q -ary linear code is well known [49, Chapter 1].

Theorem 1.4.1 (Singleton bound). *For an $[n, k, d]_q$ code, we have $d \leq n - k + 1$.*

Codes attaining this bound are called *maximum distance separable* (MDS) codes.

1.4.2 Linear Algebra Terminology

We use

$$\mathbf{e}_i = (\underbrace{0, \dots, 0}_{i-1}, 1, \underbrace{0, \dots, 0}_{n-i}) \in \mathbb{F}_q^n$$

to denote the unit vector, which has a one at the i th position, and zeros elsewhere. For a vector $\mathbf{y} = (y_1, y_2, \dots, y_n)$ and a subset $B = \{i_1, i_2, \dots, i_b\}$ of $[n]$, where $i_1 < i_2 < \dots < i_b$, let $\mathbf{y}_B$ denote the vector $(y_{i_1}, y_{i_2}, \dots, y_{i_b})$.

For an $n \times k$ matrix $\mathbf{M}$, let $\mathbf{M}_i$ denote the i th row of $\mathbf{M}$, and $\mathbf{M}[j]$ its j th

column. For a set $E \subseteq [n]$, let $\mathbf{M}_E$ denote the $|E| \times k$ submatrix of $\mathbf{M}$ formed by rows of $\mathbf{M}$ which are indexed by the elements of E . For a set $F \subseteq [k]$, let $\mathbf{M}[F]$ denote the $n \times |F|$ submatrix of $\mathbf{M}$ formed by columns of $\mathbf{M}$ which are indexed by the elements of F . For a matrix $\mathbf{M}$ over $\mathbb{F}_q$, we use $\text{rank}_q(\mathbf{M})$ to denote the rank of $\mathbf{M}$ over $\mathbb{F}_q$.

For a set of vectors $\mathcal{S}$ over $\mathbb{F}_q$, we use notation $\text{span}_q(\mathcal{S})$ to denote the linear space spanned over $\mathbb{F}_q$ by the vectors in $\mathcal{S}$.

1.4.3 Graph Theory Terminology

A simple *graph* is a pair $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ where $\mathcal{V}(\mathcal{G})$ is the set of vertices of $\mathcal{G}$ and $\mathcal{E}(\mathcal{G})$ is a set of *unordered* pairs of distinct vertices of $\mathcal{G}$. We refer to $\mathcal{E}(\mathcal{G})$ as the set of *edges* of $\mathcal{G}$. A typical edge of $\mathcal{G}$ is of the form $\{u, v\}$ where $u \in \mathcal{V}(\mathcal{G})$, $v \in \mathcal{V}(\mathcal{G})$, and $u \neq v$. If $e = \{u, v\} \in \mathcal{E}(\mathcal{G})$ we say that u and v are adjacent. We also refer to u and v as the *endpoints* of e .

A simple *digraph* is a pair $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ where $\mathcal{V}(\mathcal{D})$ is the set of vertices of $\mathcal{D}$, and $\mathcal{E}(\mathcal{D})$ is a set of *ordered* pairs of distinct vertices of $\mathcal{D}$. We refer to $\mathcal{E}(\mathcal{D})$ as the set of arcs (or directed edges) of $\mathcal{D}$. A typical *arc* of $\mathcal{D}$ is of the form $e = (u, v)$ where $u \in \mathcal{V}(\mathcal{D})$, $v \in \mathcal{V}(\mathcal{D})$, and $u \neq v$. The vertices u and v are called the *endpoints* of the arc e . The arc e is called an *out-going* arc of u and an *in-coming* arc of v .

Simple (di)graphs have no loops and no parallel (arcs) edges. In this thesis, only simple (di)graphs are considered. Therefore, we use (di)graphs to refer to simple (di)graphs for succinctness.

For a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, unless specified otherwise, we label the vertices of $\mathcal{D}$ by the natural numbers $1, 2, \dots, |\mathcal{V}(\mathcal{D})|$. In other words, if $|\mathcal{V}(\mathcal{D})| = n$ then we take $\mathcal{V}(\mathcal{D}) = [n]$. The number of vertices $|\mathcal{V}(\mathcal{D})|$ is called the *order* of $\mathcal{D}$, whereas the number of arcs $|\mathcal{E}(\mathcal{D})|$ is called the *size* of $\mathcal{D}$. The *complement* of a digraph

$\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, denoted by $\overline{\mathcal{D}} = (\mathcal{V}(\overline{\mathcal{D}}), \mathcal{E}(\overline{\mathcal{D}}))$, is defined as follows. The vertex set is $\mathcal{V}(\overline{\mathcal{D}}) = \mathcal{V}(\mathcal{D})$. The arc set is

$$\mathcal{E}(\overline{\mathcal{D}}) = \{(u, v) : u \in \mathcal{V}(\mathcal{D}), v \in \mathcal{V}(\mathcal{D}), u \neq v, (u, v) \notin \mathcal{E}(\mathcal{D})\}.$$

Analogous conventions and concepts are also defined for graphs.

Given a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, the graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ where $\mathcal{V}(\mathcal{G}) = \mathcal{V}(\mathcal{D})$ and $\mathcal{E}(\mathcal{G}) = \{\{u, v\} : (u, v) \in \mathcal{E}(\mathcal{D})\}$ is called the *underlying graph* of $\mathcal{D}$. Given a graph $\mathcal{G}$, the digraph $\mathcal{D}^{\mathcal{G}}$ obtained from $\mathcal{G}$ by replacing each edge $\{u, v\}$ of $\mathcal{G}$ by two arcs (u, v) and (v, u) is called the digraph *corresponding* to $\mathcal{G}$.

A digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ is called *symmetric* if it satisfies the property that $(u, v) \in \mathcal{E}(\mathcal{D})$ if and only if $(v, u) \in \mathcal{E}(\mathcal{D})$. A *complete graph* is a graph that contains all possible edges. A *complete digraph* is a digraph that contains all possible arcs. In other words, it is a symmetric digraph whose underlying graph is a complete graph.

A graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ is called *bipartite* if $\mathcal{V}(\mathcal{G})$ can be partitioned into two subsets A and B such that for every edge $\{u, v\} \in \mathcal{E}(\mathcal{G})$, it holds that $u \in A$ and $v \in B$, or vice versa.

A *subgraph* of a graph $\mathcal{G}$ (digraph $\mathcal{D}$, respectively) is a graph (digraph, respectively) whose vertex set V is a subset of that of $\mathcal{G}$ ($\mathcal{D}$, respectively) and whose edge set (arc set, respectively) is a subset of that of $\mathcal{G}$ ($\mathcal{D}$, respectively) restricted on the vertices in V .

Let V be a subset of vertices of a graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ (digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, respectively). The subgraph of $\mathcal{G}$ ($\mathcal{D}$, respectively) *induced* by V is a graph (digraph, respectively) whose vertex set is V , and edge set (arc set, respectively) is $\{\{u, v\} : u \in V, v \in V, \{u, v\} \in \mathcal{E}(\mathcal{G})\}$ ($\{(u, v) : u \in V, v \in V, (u, v) \in \mathcal{E}(\mathcal{D})\}$). We refer to such a subgraph as an *induced subgraph* of $\mathcal{G}$ ($\mathcal{D}$,

respectively).

A *path* in a graph $\mathcal{G}$ (digraph $\mathcal{D}$, respectively) is a sequence of distinct vertices $(u_1, u_2, \dots, u_k)$, such that $\{u_i, u_{i+1}\} \in \mathcal{E}(\mathcal{G})$ ($(u_i, u_{i+1}) \in \mathcal{E}(\mathcal{D})$, respectively) for all $i \in [k - 1]$.

A circuit (cycle, respectively) in a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ (graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$, respectively) is a sequence of pairwise distinct vertices

$$\mathcal{C} = (i_1, i_2, \dots, i_\ell),$$

where $(i_s, i_{s+1}) \in \mathcal{E}(\mathcal{D})$ ($\{i_s, i_{s+1}\} \in \mathcal{E}(\mathcal{G})$, respectively) for all $s \in [\ell - 1]$ and $(i_\ell, i_1) \in \mathcal{E}(\mathcal{D})$ ($\{i_\ell, i_1\} \in \mathcal{E}(\mathcal{G})$, respectively) as well. Let $\mathcal{V}(\mathcal{C}) = \{i_1, i_2, \dots, i_\ell\}$ be the set of vertices of $\mathcal{C}$. A digraph (graph, respectively) is called *acyclic* if it contains no circuits (cycles, respectively).

A (di)graph is called *(strongly) connected* if there is a path from each vertex in the (di)graph to every other vertex. The *(strongly) connected components* of a (di)graph are its maximal (strongly) connected subgraphs.

An edge e of a graph $\mathcal{G}$ is a *bridge* of $\mathcal{G}$ if $\mathcal{G} - e$, the graph obtained from $\mathcal{G}$ by removing e , has more connected components than $\mathcal{G}$. In particular, an edge e in a connected graph $\mathcal{G}$ is a bridge if and only if its removal renders the graph disconnected. A connected graph without any bridge (bridgeless) is called *2-edge connected*.

A *cut* $C = (S, T)$ of a graph $\mathcal{G}$ is a partition of the vertex set $\mathcal{V}(\mathcal{G})$ of $\mathcal{G}$ into two parts S and T . The cut-set of a cut $C = (S, T)$ is the set

$$\left\{ \{u, v\} \in \mathcal{E}(\mathcal{G}) : u \in S, v \in T \right\}.$$

The *size* of a cut is the number of edges in its cut-set. A min-cut is a cut whose size is not larger than the size of any other cut. If a connected graph is bridgeless then the size of each min-cut is at least two.

If (u, v) is an arc in a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, then v is called an *out-neighbor* of u . The set of out-neighbors of a vertex u in the digraph $\mathcal{D}$ is denoted by $N_O^{\mathcal{D}}(u)$. We simply use $N_O(u)$ whenever there is no potential confusion. We refer to $|N_O^{\mathcal{D}}(u)|$ as the *out-degree* of u in $\mathcal{D}$, denoted by $\deg_O^{\mathcal{D}}(u)$. For a graph $\mathcal{G}$, we denote by $N^{\mathcal{G}}(u)$ the set of neighbors of u , namely, the set of vertices adjacent to u .

An *independent set* in a graph $\mathcal{G}$ is a set of vertices of $\mathcal{G}$ with no edges connecting any two of them. An independent set in $\mathcal{G}$ of largest cardinality is called a *maximum independent set* in $\mathcal{G}$. The cardinality of such a maximum independent set is referred to as the *independence number* of $\mathcal{G}$, denoted by $\alpha(\mathcal{G})$.

A clique of a (di)graph is a set of vertices that induces a complete subgraph of that (di)graph. A *clique cover* of a (di)graph is a set of cliques that partition its vertex set. A *minimum clique cover* of a (di)graph is a clique cover of minimum number of cliques. The number of cliques in such a minimum clique cover of a (di)graph is called the clique cover number of that (di)graph. We denote by $\text{cc}(\mathcal{G})$ the clique cover number of a graph $\mathcal{G}$ and $\text{cc}(\mathcal{D})$ the clique cover number of a digraph $\mathcal{D}$.

A *tree* is a connected acyclic graph. A *rooted tree* is a tree with one special vertex designated to be the *root*. In a rooted tree, the *parent* of a vertex v is the vertex connected to it on the path from v to the root. Every vertex except the root has a unique parent. If v is the parent of a vertex u then u is the *child* of v .

Definition 1.4.2 ([34]). Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph of order n .

1. An $n \times n$ matrix $\mathbf{M} = (m_{i,j})$ over an arbitrary field $\mathbb{F}$ is said to *fit* $\mathcal{D}$ if

$$\begin{cases} m_{i,j} \neq 0, & i = j, \\ m_{i,j} = 0, & i \neq j, (i,j) \notin \mathcal{E}(\mathcal{D}). \end{cases}$$

2. The *minrank* of $\mathcal{D}$ over $\mathbb{F}$ is defined to be

$$\text{minrk}_{\mathbb{F}}(\mathcal{D}) = \min\{\text{rank}_{\mathbb{F}}(\mathbf{M}) : \mathbf{M} \text{ is over } \mathbb{F} \text{ and } \mathbf{M} \text{ fits } \mathcal{D}\}$$

When $\mathbb{F} = \mathbb{F}_q$, we simply write $\text{minrk}_q(\mathcal{D})$ instead of $\text{minrk}_{\mathbb{F}_q}(\mathcal{D})$. We also have analogous definitions for a graph.

A (directed) *hypergraph* is a pair $(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a set of vertices and $\mathcal{E}$ is a set of *hyperarcs*. A hyperarc e itself is an ordered pair (T, H) , where T and H are both subsets of $\mathcal{V}$. They respectively represent the *tail* and the *head* of the hyperarc e . If $T = \{v\}$, i.e., T consists of only one vertex v , then we simply write (v, H) instead of $(\{v\}, H)$.

1.4.4 Information Theory Terminology

Let X and Y be discrete random variables taking values in the sets Σ_X and Σ_Y , respectively. Let $\Pr(X = x)$ denote the probability that X takes a particular value $x \in \Sigma_X$. The following notation is standard in Information Theory (see [20] for the background).

The (binary) *entropy* of X is defined as

$$H(X) = - \sum_{x \in \Sigma_X} \Pr(X = x) \times \log_2 \Pr(X = x).$$

The *conditional entropy* of X given Y is defined as

$$H(X | Y) = - \sum_{x \in \Sigma_X, y \in \Sigma_Y} \Pr(X = x, Y = y) \times \log_2 \Pr(X = x | Y = y).$$

This definition can be naturally extended to $H(X | Y_1, Y_2, \dots, Y_n)$ for n discrete random variables $Y_i, i \in [n]$.

If the probability distribution of X is unchanged given the knowledge of Y , that is, $\Pr(X = x | Y = y) = \Pr(X = x)$ for all $x \in \Sigma_X$ and $y \in \Sigma_Y$, then $H(X | Y) = H(X)$. Indeed,

$$\begin{aligned} H(X | Y) &= - \sum_{x \in \Sigma_X} \left(\sum_{y \in \Sigma_Y} \Pr(X = x, Y = y) \right) \times \log_2 \Pr(X = x) \\ &= - \sum_{x \in \Sigma_X} \Pr(X = x) \times \log_2 \Pr(X = x) \\ &= H(X). \end{aligned}$$

2. SECURE INDEX CODE WITH SIDE INFORMATION

Part of the work in this chapter was presented in the 2011 IEEE Symposium on Information Theory [22]. We investigate the security aspects of the Index Coding with Side Information problem. Building on the results of Bar-Yossef *et al.* [3,4], the properties of linear index codes are further explored. The notion of *weak security*, considered by Bhattad and Narayanan [9] in the context of network coding, is generalized to *block security*. It is shown that the linear index code based on a matrix $\mathbf{L}$, whose column space $\mathcal{C}(\mathbf{L})$ has length n , minimum distance d and dual distance $d^\perp$, is $(d - 1 - t)$ -block secure (and hence also weakly secure) if the adversary knows in advance $t \leq d - 2$ messages, and is completely insecure if the adversary knows in advance more than $n - d^\perp$ messages. *Strong security* is examined under the circumstance that the adversary possesses some t messages in advance and eavesdrops at most μ transmissions. We prove that for sufficiently large q , an optimal linear index code that is strongly secure against such an adversary has length $\kappa_q + \mu$. Here κ_q is a generalization (to the case of hypergraph) of the minrank over $\mathbb{F}_q$ of the side information digraph for the ICSI problem in its original formulation [3]. As mentioned earlier in Section 1.2, both block security and strong security are, in fact, special cases of a more general security model called security against guessing, which was also introduced in the work of Bhattad and Narayanan [9].

2.1 Index Coding and Some Basic Results

In Chapter 2 and Chapter 3, we follow the most general model studied by Alon *et al.* [2]. The Index Coding with Side Information (ICSI) problem considers the following communications scenario. There is a unique sender S , who has a vector of messages $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{F}_q^n$ in his possession, which is a realized value of a random vector $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The random variables $X_1, X_2, \dots, X_n$ hereafter are assumed to be independent uniformly distributed random variables over $\mathbb{F}_q$. There are also m receivers $R_1, R_2, \dots, R_m$. For each $i \in [m]$, the receiver R_i has some side information, that is, R_i has a subset of messages $\{x_j\}_{j \in \mathcal{X}_i}$, where $\mathcal{X}_i \subsetneq [n]$. In addition, each R_i , for $i \in [m]$, is interested in receiving the message $x_{f(i)}$, for some *demand function* $f : [m] \rightarrow [n]$. Here we assume that $f(i) \notin \mathcal{X}_i$ for all $i \in [m]$. The assumption that each receiver is interested in exactly one message is not a limitation of the model. In fact, we can consider an equivalent problem by splitting each receiver who requests multiple messages into multiple receivers, each of whom requests exactly one message and have the same set of side information [3, 10]. Let $\mathcal{X} = (\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m)$. An instance of the ICSI problem (or an ICSI instance, for short) is given by a quadruple $(m, n, \mathcal{X}, f)$. It can also be conveniently described by a (directed) hypergraph [2].

Definition 2.1.1. Let $(m, n, \mathcal{X}, f)$ be an ICSI instance. The corresponding *side information (directed) hypergraph* $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$ is defined by the vertex set $\mathcal{V}(\mathcal{H}) = [n]$ and the hyperarc set $\mathcal{E}(\mathcal{H})$, where

$$\mathcal{E}(\mathcal{H}) = \left\{ (f(i), \mathcal{X}_i) : i \in [m] \right\}.$$

We often refer to $(m, n, \mathcal{X}, f)$ as an ICSI instance described by the hypergraph $\mathcal{H}$.

Fig. 2.1: The hypergraphs $\mathcal{H}_1$, $\mathcal{H}_2$ and the digraph $\mathcal{D}_2$

For instance, consider an ICSI instance where $n = 3$ (three messages), $m = 4$ (four receivers), $f(1) = 1$, $f(2) = 2$, $f(3) = 3$, $f(4) = 2$, $\mathcal{X}_1 = \{2, 3\}$, $\mathcal{X}_2 = \{1\}$, $\mathcal{X}_3 = \{1, 2\}$, and $\mathcal{X}_4 = \{3\}$. The hypergraph $\mathcal{H}_1$ that describes this instance has three vertices 1, 2, 3, and has four hyperarcs. These are $e_1 = (1, \{2, 3\})$, $e_2 = (2, 1)$, $e_3 = (3, \{1, 2\})$, and $e_4 = (2, 3)$. This hypergraph is depicted in Fig. 2.1a.

Remark 2.1.2. In the original setting of the ICSI problem [3], we have $m = n$ and $f(i) = i$ for all $i \in [n]$. In that case, the corresponding side information hypergraph has precisely n hyperarcs where each of them has a different origin vertex. Then it is simpler to describe such an ICSI instance by a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}) = [n], \mathcal{E}(\mathcal{D}))$, so-called *side information digraph* [3, 46]. For each hyperarc $(i, \mathcal{X}_i)$ of $\mathcal{H}$, there will be $|\mathcal{X}_i|$ arcs (i, j) of $\mathcal{D}$, for $j \in \mathcal{X}_i$. Equivalently, $\mathcal{E}(\mathcal{D}) = \{(i, j) : i, j \in [n], j \in \mathcal{X}_i\}$. In this case, we refer to $\mathcal{D}$ as the *underlying digraph* of the hypergraph $\mathcal{H}$.

As an example, let $\mathcal{H}_2$ (Fig. 2.1b) be the hypergraph obtained from $\mathcal{H}_1$ by deleting the hyperarc e_4 . The hypergraph $\mathcal{H}_2$ describes the ICSI instance obtained from the aforementioned instance by removing the last receiver. This new ICSI instance can also be described by the side information digraph $\mathcal{D}_2$ (Fig. 2.1c).

Furthermore, when it satisfies that $x_j \in \mathcal{X}_i$ if and only if $x_i \in \mathcal{X}_j$ for all $i \neq j$, the side information digraph $\mathcal{D}$ is symmetric, that is, $(i, j) \in \mathcal{E}(\mathcal{D})$ if and only if $(j, i) \in \mathcal{E}(\mathcal{D})$. In that case, we may consider the underlying graph $\mathcal{G}$ of $\mathcal{D}$ instead,

which is referred to as the *underlying graph* of the hypergraph $\mathcal{H}$. We also call $\mathcal{G}$ the *side information graph* of the corresponding ICSI instance.

Definition 2.1.3. An *index code* over $\mathbb{F}_q$ for an ICSI instance $(m, n, \mathcal{X}, f)$ (or just an $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$), is an encoding function

$$\mathfrak{E} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^N,$$

such that for each receiver R_i , $i \in [m]$, there exists a decoding function

$$\mathfrak{D}_i : \mathbb{F}_q^N \times \mathbb{F}_q^{|\mathcal{X}_i|} \rightarrow \mathbb{F}_q,$$

satisfying

$$\forall \mathbf{x} \in \mathbb{F}_q^n : \mathfrak{D}_i(\mathfrak{E}(\mathbf{x}), \mathbf{x}_{\mathcal{X}_i}) = x_{f(i)}.$$

We also refer to such an $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ as an $\mathcal{H}$ -IC over $\mathbb{F}_q$, where $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$ is the corresponding side information hypergraph. The parameter N is called the *length* of the index code. In the scheme corresponding to this code, S broadcasts a vector $\mathfrak{E}(\mathbf{x})$ of length N over $\mathbb{F}_q$.

Definition 2.1.4. An index code of minimum length is called *optimal*.

Definition 2.1.5. A *linear index code* is an index code for which the encoding function $\mathfrak{E}$ is a linear transformation over $\mathbb{F}_q$. Such a code can be described as

$$\forall \mathbf{x} \in \mathbb{F}_q^n : \mathfrak{E}(\mathbf{x}) = \mathbf{x}\mathbf{L},$$

where $\mathbf{L}$ is an $n \times N$ matrix over $\mathbb{F}_q$. The matrix $\mathbf{L}$ is called the *matrix corresponding*

to the index code $\mathfrak{E}$. We also refer to $\mathfrak{E}$ as the *index code based on $\mathbf{L}$* . Notice that the length of $\mathfrak{E}$ is the number of columns of $\mathbf{L}$.

Let $E \subseteq [n]$ and $\mathbf{u} \in \mathbb{F}_q^n$. Recall that $\mathbf{u} \triangleleft E$ means $\text{supp}(\mathbf{u}) \subseteq E$. If some receiver knows x_j for all $j \in E$ and $\mathbf{u} \triangleleft E$, then this receiver can compute the value of $\mathbf{x}\mathbf{u}^T$.

Hereafter, we assume that the sets $\mathcal{X}_i$'s for $i \in [m]$ are known to S . Moreover, we also assume that the index code $\mathfrak{E}$ is known to each receiver R_i for $i \in [m]$. In practice this can be achieved by a preliminary communication session, when the knowledge of the sets $\mathcal{X}_i$'s for $i \in [m]$ and of the code $\mathfrak{E}$ are disseminated between the participants of the scheme.

Let

$$\mathcal{C}(\mathbf{L}) = \text{span}_q \left(\{ \mathbf{L}[j]^T \}_{j \in [N]} \right),$$

the subspace spanned by the transposed columns of $\mathbf{L}$. The following lemma was implicitly formulated by Bar-Yossef *et al.* [3, 4] for the case where $m = n$, $f(i) = i$ for all $i \in [m]$, and $q = 2$. This lemma specifies a sufficient condition on $\mathcal{C}(\mathbf{L})$ so that $\mathbf{L}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. We reproduce this lemma with its proof in its general form for the sake of completeness of the presentation.

Lemma 2.1.6. *Let $\mathbf{L}$ be an $n \times N$ matrix over $\mathbb{F}_q$ and suppose S broadcast $\mathbf{x}\mathbf{L}$. Then, for each $i \in [m]$, the receiver R_i can reconstruct $x_{f(i)}$ if there exists a vector $\mathbf{u}^{(i)} \in \mathbb{F}_q^n$ satisfying*

1. $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$;
2. $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L})$.

Proof. Assume that $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$ and $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L})$. Since $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L})$, there

exists $\boldsymbol{\beta} \in \mathbb{F}_q^N$ such that

$$\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} = \boldsymbol{\beta} \mathbf{L}^\top.$$

By taking the transpose and pre-multiplying by $\mathbf{x}$, we obtain

$$\mathbf{x}(\mathbf{u}^{(i)} + \mathbf{e}_{f(i)})^\top = (\mathbf{x} \mathbf{L}) \boldsymbol{\beta}^\top.$$

Therefore,

$$x_{f(i)} = \mathbf{x} \mathbf{e}_{f(i)}^\top = (\mathbf{x} \mathbf{L}) \boldsymbol{\beta}^\top - \mathbf{x} \mathbf{u}^{(i)\top}.$$

Observe that R_i is able to find $\mathbf{u}^{(i)}$ and $\boldsymbol{\beta}$ from the knowledge of $\mathbf{L}$. Moreover, R_i is also able to compute $\mathbf{x} \mathbf{u}^{(i)\top}$ since $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$, and receives $\mathbf{x} \mathbf{L}$ from S . Therefore, R_i is able to determine $x_{f(i)}$. ■

Remark 2.1.7. Lemma 2.1.6 implies that $\mathbf{L}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ if $\mathcal{C}(\mathbf{L}) \supseteq \text{span}_q(\{\mathbf{u}^{(i)} + \mathbf{e}_{f(i)}\}_{i \in [m]})$, for some $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$, $i \in [m]$. We show later in Corollary 2.2.6 that this condition is also necessary. Finding such an $\mathbf{L}$ with minimum number of columns by a careful selection of $\mathbf{u}^{(i)}$'s is a difficult task (in fact, the corresponding decision problem is NP-complete [3, 53]), which, however, yields a linear coding scheme with minimum number of transmissions.

2.2 Block Secure Linear Index Codes

2.2.1 Block Security and Weak Security

In this section, we assume the presence of an adversary A who can listen to *all* transmissions. Assume that S employs a linear index code based on $\mathbf{L}$. The adversary is assumed to possess side information $\{x_j\}_{j \in \mathcal{X}_A}$, where $\mathcal{X}_A \subsetneq [n]$. For short, we say that A *knows* or *possesses* $\mathbf{x}_{\mathcal{X}_A}$. The strength of the adversary A is defined to be

$|\mathcal{X}_A|$. Denote $\widehat{\mathcal{X}}_A \triangleq [n] \setminus \mathcal{X}_A$. Note that by listening to S , the adversary also knows $\mathbf{s} = \mathbf{x}\mathbf{L}$. We define below several levels of security for linear index codes.

Definition 2.2.1. Suppose that the sender S possesses a vector of messages $\mathbf{x} \in \mathbb{F}_q^n$, which is a realized value of a random vector $\mathbf{X}$. An adversary A possesses $\mathbf{x}_{\mathcal{X}_A}$. Consider a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$.

1. For $B \subseteq \widehat{\mathcal{X}}_A$, the adversary is said to *have no information about $\mathbf{x}_B$* if

$$\mathrm{H}(\mathbf{X}_B \mid \mathbf{X}\mathbf{L}, \mathbf{X}_{\mathcal{X}_A}) = \mathrm{H}(\mathbf{X}_B). \quad (2.1)$$

In other words, despite the partial knowledge on $\mathbf{x}$ that the adversary has (his side information and the transmissions he eavesdrops), the symbols $\mathbf{x}_B$ still look completely random to him.

2. The index code is said to be *b-block secure against $\mathcal{X}_A$* if for every b -subset $B \subseteq \widehat{\mathcal{X}}_A$ (a b -subset is a subset of size b), the adversary has no information about $\mathbf{x}_B$.
3. The index code is said to be *b-block secure against all adversaries of strength t* ($0 \leq t \leq n - 1$) if it is b -block secure against $\mathcal{X}_A$ for every $\mathcal{X}_A \subset [n]$, where $|\mathcal{X}_A| = t$.
4. The index code is said to be *weakly secure against $\mathcal{X}_A$* if it is 1-block secure against $\mathcal{X}_A$. In other words, after listening to all transmissions, the adversary has no information about each particular message that he does not possess in the first place.
5. The index code is said to be *weakly secure against all adversaries of strength t* ($0 \leq t \leq n - 1$) if it is weakly secure against $\mathcal{X}_A$ for every t -subset $\mathcal{X}_A$ of $[n]$.

6. The index code is said to be *completely insecure against $\mathcal{X}_A$* if an adversary, who possesses $\{x_i\}_{i \in \mathcal{X}_A}$, by listening to all transmissions, is able to determine x_i for all $i \in \widehat{\mathcal{X}}_A$.
7. The index code is said to be *completely insecure against any adversary of strength t* ($0 \leq t \leq n - 1$) if an adversary, who possesses an arbitrary set of t messages, is always able to reconstruct all of the other $n - t$ messages after listening to all transmissions.

Remark 2.2.2. Even when the index code is b -block secure ($b \geq 1$) as defined above, the adversary is still able to obtain information about dependencies between various x_i 's in $\widehat{\mathcal{X}}_A$ (but he gains no information about any group of b particular messages). This definition of b -block security is a generalization of that of weak security [9, 58]. Obviously, if an index code is b -block secure against $\mathcal{X}_A$ ($b \geq 1$) then it is also weakly secure against $\mathcal{X}_A$, but the converse is not always true.

2.2.2 Necessary and Sufficient Conditions for Block Security

In the sequel, we consider the sets $B \subseteq [n]$, where $B \neq \emptyset$, and $E \subseteq [n]$, where $E \neq \emptyset$. Moreover, we assume that the sets $\mathcal{X}_A$, B , and E are disjoint, and that they form a partition of $[n]$, namely $\mathcal{X}_A \cup B \cup E = [n]$. Hence, $\widehat{\mathcal{X}}_A = [n] \setminus \mathcal{X}_A = B \cup E$. Here, $\mathcal{X}_A$ corresponds to the set of messages that the adversary A possesses, B corresponds to the set of messages that A is trying to gain information about, and E corresponds to the set of remaining messages.

Lemma 2.2.3. Assume that for all $\mathbf{u} \triangleleft \mathcal{X}_A$ and for all $\alpha_i \in \mathbb{F}_q$, $i \in B$ (not all α_i 's are zeros),

$$\mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i \notin \mathcal{C}(\mathbf{L}). \quad (2.2)$$

Then,

1. for all $i \in B$:

$$\mathbf{L}_i \in \text{span}_q(\{\mathbf{L}_j\}_{j \in E}); \quad (2.3)$$

2. the system

$$\mathbf{y}\mathbf{L}_E = \mathbf{w}\mathbf{L}_B \quad (2.4)$$

has at least one solution $\mathbf{y} \in \mathbb{F}_q^{|E|}$ for every choice of $\mathbf{w} \in \mathbb{F}_q^{|B|}$.

Proof.

1. If $\text{rank}_q(\mathbf{L}_E) = N$ then the first claim follows immediately. Otherwise, assume that $\text{rank}_q(\mathbf{L}_E) < N$. As the N columns of $\mathbf{L}_E$ are linearly dependent, there exists $\mathbf{y} \in \mathbb{F}_q^N \setminus \{\mathbf{0}\}$ such that $\mathbf{y}\mathbf{L}_E^T = \mathbf{0}$.

- (a) If for all such $\mathbf{y}$ and for all $i \in B$ we have $\mathbf{y}\mathbf{L}_i^T = 0$, then

$$\mathbf{L}_i \in \left((\text{span}_q(\{\mathbf{L}_j\}_{j \in E}))^\perp \right)^\perp = \text{span}_q(\{\mathbf{L}_j\}_{j \in E})$$

for all $i \in B$.

- (b) Otherwise, there exist $\mathbf{y} \in \mathbb{F}_q^N$ and $i \in B$ such that $\mathbf{y}\mathbf{L}_E^T = \mathbf{0}$ and $\mathbf{y}\mathbf{L}_i^T \neq 0$. Without loss of generality, assume that

$$\mathbf{L} = \begin{bmatrix} \mathbf{L}_{\mathcal{X}_A} \\ \mathbf{L}_B \\ \mathbf{L}_E \end{bmatrix}.$$

Let $\mathbf{c} = \mathbf{y}\mathbf{L}^T \in \mathcal{C}(\mathbf{L})$. Then

$$\mathbf{c} = (\mathbf{c}_{\mathcal{X}_A} | \mathbf{c}_B | \mathbf{c}_E) = (\mathbf{y}\mathbf{L}_{\mathcal{X}_A}^T | \mathbf{y}\mathbf{L}_B^T | \mathbf{y}\mathbf{L}_E^T).$$

Hence $\mathbf{c}_B = \mathbf{y}\mathbf{L}_B^T \neq \mathbf{0}$ and $\mathbf{c}_E = \mathbf{y}\mathbf{L}_E^T = \mathbf{0}$. Let $\mathbf{u} = (\mathbf{c}_{\mathcal{X}_A} | \mathbf{0} | \mathbf{0}) \triangleleft \mathcal{X}_A$ and $\alpha_i = c_i$ for all $i \in B$. Then α_i 's are not all zero and $\mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i = \mathbf{c} \in \mathcal{C}(\mathbf{L})$, which contradicts (2.2).

2. By (2.3), each row of $\mathbf{L}_B$ is a linear combination of rows of $\mathbf{L}_E$. Hence $\mathbf{w}\mathbf{L}_B$ is also a linear combination of rows of $\mathbf{L}_E$. Therefore, (2.4) has at least one solution. ■

The following lemma provides us with a criterion to decide whether the index code based on a particular matrix $\mathbf{L}$ is block secure against an adversary A .

Lemma 2.2.4. *Suppose that S employs a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$. For a subset $B \subseteq \widehat{\mathcal{X}}_A = [n] \setminus \mathcal{X}_A$, the adversary, after listening to all transmissions, has no information about $\mathbf{x}_B$ if and only if*

$$\begin{aligned} \forall \mathbf{u} \triangleleft \mathcal{X}_A, \forall \alpha_i \in \mathbb{F}_q \text{ with } \alpha_i, i \in B, \text{ not all zero:} \\ \mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i \notin \mathcal{C}(\mathbf{L}). \end{aligned} \tag{2.5}$$

In particular, for each $i \notin \mathcal{X}_A$, the adversary A has no information about x_i if and only if

$$\forall \mathbf{u} \triangleleft \mathcal{X}_A : \mathbf{u} + \mathbf{e}_i \notin \mathcal{C}(\mathbf{L}).$$

Proof. Assume that (2.5) holds. We need to show that the entropy of $\mathbf{X}_B$ is not changed given the knowledge of $\mathbf{X}\mathbf{L}$ and $\mathbf{X}_{\mathcal{X}_A}$. It suffices to show that for all $\mathbf{g} \in \mathbb{F}_q^{|B|}$:

$$\Pr(\mathbf{X}_B = \mathbf{g} \mid \mathbf{X}\mathbf{L} = \mathbf{s}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) = \frac{1}{q^{|B|}}, \tag{2.6}$$

where $\mathbf{x} \in \mathbb{F}_q^n$ and $\mathbf{s} = \mathbf{x}\mathbf{L}$.

Consider the following linear system with the unknown $\mathbf{z} \in \mathbb{F}_q^n$

$$\begin{cases} \mathbf{z}_B = \mathbf{g} \\ \mathbf{z}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A} \\ \mathbf{z}\mathbf{L} = \mathbf{s} \end{cases},$$

which is equivalent to

$$\begin{cases} \mathbf{z}_B = \mathbf{g} \\ \mathbf{z}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A} \\ \mathbf{z}_E \mathbf{L}_E = \mathbf{s} - \mathbf{g} \mathbf{L}_B - \mathbf{x}_{\mathcal{X}_A} \mathbf{L}_{\mathcal{X}_A} \end{cases}. \quad (2.7)$$

In order to prove that (2.6) holds, it suffices to show that for all choices of $\mathbf{g} \in \mathbb{F}_q^{|B|}$, (2.7) always has the same number of solutions $\mathbf{z}$. Notice that the number of solutions $\mathbf{z}$ of (2.7) is equal to the number of solutions $\mathbf{z}_E$ of

$$\mathbf{z}_E \mathbf{L}_E = \mathbf{s} - \mathbf{g} \mathbf{L}_B - \mathbf{x}_{\mathcal{X}_A} \mathbf{L}_{\mathcal{X}_A}, \quad (2.8)$$

where $\mathbf{s}$, $\mathbf{g}$, and $\mathbf{x}_{\mathcal{X}_A}$ are known. For any $\mathbf{g} \in \mathbb{F}_q^{|B|}$, if (2.8) has a solution, then it has exactly $q^{|E| - \text{rank}_q(\mathbf{L}_E)}$ different solutions. Therefore, it suffices to prove that (2.8) has at least one solution for every $\mathbf{g} \in \mathbb{F}_q^{|B|}$.

Since $\mathbf{s} = \mathbf{x}\mathbf{L}$, we have

$$\mathbf{x}_E \mathbf{L}_E = \mathbf{s} - \mathbf{x}_B \mathbf{L}_B - \mathbf{x}_{\mathcal{X}_A} \mathbf{L}_{\mathcal{X}_A}. \quad (2.9)$$

Subtracting (2.9) from (2.8) we obtain

$$(\mathbf{z}_E - \mathbf{x}_E)\mathbf{L}_E = (\mathbf{x}_B - \mathbf{g})\mathbf{L}_B,$$

which can be rewritten as

$$\mathbf{y}\mathbf{L}_E = \mathbf{w}\mathbf{L}_B, \quad (2.10)$$

where $\mathbf{y} \triangleq \mathbf{z}_E - \mathbf{x}_E$ and $\mathbf{w} \triangleq \mathbf{x}_B - \mathbf{g}$. Due to Lemma 2.2.3, (2.10) always has a solution $\mathbf{y}$, for every choice of $\mathbf{w}$. Therefore (2.8) has at least one solution for every $\mathbf{g} \in \mathbb{F}_q^{|B|}$.

Now we prove the converse. Assume that (2.5) does not hold. Then there exist $\mathbf{u} \triangleleft \mathcal{X}_A$ and $\alpha_i \in \mathbb{F}_q$ for $i \in B$, where α_i 's for $i \in B$ are not all zero, such that

$$\sum_{i \in B} \alpha_i \mathbf{e}_i = \mathbf{c} - \mathbf{u},$$

for some $\mathbf{c} \in \mathcal{C}(\mathbf{L})$. Hence, similarly to the proof of Lemma 2.1.6, the adversary obtains

$$\sum_{i \in B} \alpha_i x_i = \mathbf{x} \left(\sum_{i \in B} \alpha_i \mathbf{e}_i \right)^T = \mathbf{x}(\mathbf{c} - \mathbf{u})^T = \mathbf{x}\mathbf{c}^T - \mathbf{x}\mathbf{u}^T.$$

Note that the adversary can calculate $\mathbf{x}\mathbf{c}^T$ from $\mathbf{s}$, and can also find $\mathbf{x}\mathbf{u}^T$ based on his own side information. Therefore, A is able to compute a nontrivial linear combination of x_i 's, $i \in B$. Hence the entropy $\mathbf{H}(\mathbf{X}_B \mid \mathbf{X}\mathbf{L}, \mathbf{X}_{\mathcal{X}_A}) < \mathbf{H}(\mathbf{X}_B)$. Thus, the adversary gains some information about the $\mathbf{x}_B$. ■

The following corollary generalizes Lemma 2.1.6 by providing a necessary and sufficient condition for a receiver's ability to recover the desired message. Equivalently, it provides a necessary and sufficient condition for a receiver (or an adversary) to have no information about a particular message. Observe that, for one particular

message, a receiver (or an adversary) can recover that message or otherwise has no information about it.

Corollary 2.2.5. *Let $\mathbf{L}$ be an $n \times N$ matrix over $\mathbb{F}_q$ and let S broadcast $\mathbf{xL}$. Then for each $i \in [m]$, the receiver R_i can reconstruct $x_{f(i)}$ if and only if there exists a vector $\mathbf{u}^{(i)} \in \mathbb{F}_q^n$ such that*

1. $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$;
2. $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L})$.

Equivalently, the receiver R_i (or any adversary that owns $\mathbf{x}_{\mathcal{X}_i}$) has no information about $x_{f(i)}$ if and only if the above condition is not satisfied.

Proof. The second statement is a reformulation of the last assertion in Lemma 2.2.4. The first statement is proved as follows. Suppose that there exists a vector $\mathbf{u}^{(i)} \in \mathbb{F}_q^n$ satisfying the two stated conditions. By Lemma 2.1.6, the receiver R_i can determine $x_{f(i)}$. Conversely, suppose that R_i can reconstruct $x_{f(i)}$. Then the aforementioned condition must be satisfied, for otherwise, R_i would have no information about $x_{f(i)}$. ■

Corollary 2.2.6. *The matrix $\mathbf{L}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ if and only if for all $i \in [m]$, there exists a vector $\mathbf{u}^{(i)} \in \mathbb{F}_q^n$ satisfying*

1. $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$;
2. $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L})$.

Definition 2.2.7. Let $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$ be the hypergraph that describes an ICSI instance $(m, n, \mathcal{X}, f)$. The *minrank* over $\mathbb{F}_q$ of $\mathcal{H}$ is defined to be

$$\kappa_q(\mathcal{H}) = \kappa_q(m, n, \mathcal{X}, f) \triangleq \min \left\{ \text{rank}_q(\{\mathbf{u}^{(i)} + \mathbf{e}_{f(i)}\}_{i \in [m]}) : \mathbf{u}^{(i)} \in \mathbb{F}_q^n, \mathbf{u}^{(i)} \triangleleft \mathcal{X}_i \right\}, \quad (2.11)$$

We simply write κ_q when there is no potential confusion.

Fig. 2.2: The hypergraph $\mathcal{H}_1$

For instance, let us consider the hypergraph $\mathcal{H}_1$ (Fig. 2.2) that describes the following ICSI instance: $n = 3$ (three messages), $m = 4$ (four receivers), $f(1) = 1$, $f(2) = 2$, $f(3) = 3$, $f(4) = 2$, $\mathcal{X}_1 = \{2, 3\}$, $\mathcal{X}_2 = \{1\}$, $\mathcal{X}_3 = \{1, 2\}$, and $\mathcal{X}_4 = \{3\}$. It is straightforward to verify that $\kappa_2(\mathcal{H}_1) = 2$. We may select $\mathbf{u}^{(1)} = (0, 1, 1)$, $\mathbf{u}^{(2)} = (0, 0, 0)$, $\mathbf{u}^{(3)} = (1, 1, 0)$, and $\mathbf{u}^{(4)} = (0, 0, 0)$. Then

$$\begin{aligned} & \text{rank}_2(\{\mathbf{u}^{(1)} + \mathbf{e}_1, \mathbf{u}^{(2)} + \mathbf{e}_2, \mathbf{u}^{(3)} + \mathbf{e}_3, \mathbf{u}^{(4)} + \mathbf{e}_2\}) \\ &= \text{rank}_2\{(1, 1, 1), (0, 1, 0), (1, 1, 1), (0, 1, 0)\} \\ &= 2. \end{aligned}$$

It follows from Corollary 2.2.6 that $\mathbf{L}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ if and only if $\mathcal{C}(\mathbf{L}) \supseteq \text{span}_q(\{\mathbf{u}^{(i)} + \mathbf{e}_{f(i)}\}_{i \in [m]})$, for some $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$, $i \in [m]$. Therefore, $\kappa_q(m, n, \mathcal{X}, f)$ is the shortest possible length of a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.

Corollary 2.2.8. *The length of an optimal linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is $\kappa_q = \kappa_q(m, n, \mathcal{X}, f)$.*

When $m = n$ and $f(i) = i$ for all $i \in [n]$, the quantity κ_q is precisely the minrank over $\mathbb{F}_q$ of the side information digraph $\mathcal{D}$, which was introduced by Haemers [34] (see also [3]). Indeed, suppose that $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$ for all $i \in [n]$ and let $\mathbf{A} = (a_{i,j})$ be the $n \times n$ matrix whose i th row is precisely $\mathbf{u}^{(i)} + \mathbf{e}_i$, $i \in [n]$. Recall from Remark 2.1.2 that $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$, where $\mathcal{V}(\mathcal{D}) = [n]$ and

$$\mathcal{E}(\mathcal{D}) = \{e = (i, j) : i, j \in [n], j \in \mathcal{X}_i\}.$$

Then $\mathbf{A}$ fits $\mathcal{D}$, that is,

$$\begin{cases} a_{i,j} \neq 0, & i = j, \\ a_{i,j} = 0, & i \neq j, (i, j) \notin \mathcal{E}(\mathcal{D}). \end{cases}$$

Conversely, if $\mathbf{A}'$ fits $\mathcal{D}$ then by multiplying each row of $\mathbf{A}'$ with a suitable nonzero constant (which does not change the rank of $\mathbf{A}'$), one obtains a matrix $\mathbf{A}$ which is of the aforementioned form. Therefore, κ_q defined as in (2.11) is indeed the minimum rank over $\mathbb{F}_q$ of a matrix that fits the side information graph $\mathcal{D}$. Thus, by Definition 1.4.2, κ_q is precisely the minrank over $\mathbb{F}_q$ of $\mathcal{D}$.

Theorem 2.2.9. *Consider a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$. Let d be the minimum distance of $\mathcal{C}(\mathbf{L})$.*

1. *This index code is $(d - 1 - t)$ -block secure against all adversaries of strength $t \leq d - 2$. In particular, it is weakly secure against all adversaries of strength $t = d - 2$.*

2. This index code is not weakly secure against at least one adversary of strength $t = d - 1$. More generally, if there exists a codeword of $\mathcal{C}(\mathbf{L})$ of weight w , then this index code is not weakly secure against at least one adversary of strength $t = w - 1$.
3. Every adversary of strength $t \leq d - 1$ is able to find a list of q^{n-t-N} vectors in $\mathbb{F}_q^n$ which includes the vector of messages $\mathbf{x}$.

Proof.

1. Assume that $t \leq d - 2$. By Lemma 2.2.4, it suffices to show that for every t -subset $\mathcal{X}_A$ of $[n]$ and for every $(d - 1 - t)$ -subset B of $\widehat{\mathcal{X}}_A = [n] \setminus \mathcal{X}_A$,

$$\begin{aligned} \forall \mathbf{u} \triangleleft \mathcal{X}_A, \forall \alpha_i \in \mathbb{F}_q \text{ with } \alpha_i, i \in B, \text{ not all zero :} \\ \mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i \notin \mathcal{C}(\mathbf{L}). \end{aligned} \tag{2.12}$$

For such $\mathbf{u}$ and α_i 's, we have $\text{wt}(\mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i) \leq \text{wt}(\mathbf{u}) + \text{wt}(\sum_{i \in B} \alpha_i \mathbf{e}_i) \leq t + (d - 1 - t) = d - 1 < d$. Moreover, as $\text{supp}(\mathbf{u}) \cap B = \emptyset$ and α_i 's, $i \in B$, are not all zero, we deduce that $\mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i \neq \mathbf{0}$. We conclude that $\mathbf{u} + \sum_{i \in B} \alpha_i \mathbf{e}_i \notin \mathcal{C}(\mathbf{L})$.

2. We now show that the index code is not weakly secure against at least one adversary of strength $t = d - 1$. The more general statement can be proved in an analogous way.

Pick a codeword $\mathbf{c} = (c_1, c_2, \dots, c_n) \in \mathcal{C}(\mathbf{L})$ such that $\text{wt}(\mathbf{c}) = d$ and let $\text{supp}(\mathbf{c}) = \{i_1, i_2, \dots, i_d\}$. Take $\mathcal{X}_A = \{i_1, i_2, \dots, i_{d-1}\}$, where $|\mathcal{X}_A| = d - 1$. Let

$$\mathbf{u} = (\mathbf{c}/c_{i_d} - \mathbf{e}_{i_d}).$$

Then, $\mathbf{u} \triangleleft \mathcal{X}_A$ and $\mathbf{u} + \mathbf{e}_{i_d} = \mathbf{c}/c_{i_d} \in \mathcal{C}(\mathbf{L})$. By Lemma 2.1.6, A is able to determine x_{i_d} . Hence the index code is not weakly secure against the adversary A , who knows $d - 1$ messages x_i 's in advance.

3. Consider the following linear system of equations with unknown $\mathbf{z} \in \mathbb{F}_q^n$

$$\begin{cases} \mathbf{z}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A} \\ \mathbf{z}\mathbf{L} = \mathbf{s} \end{cases},$$

which is equivalent to

$$\begin{cases} \mathbf{z}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A} \\ \mathbf{z}_{\hat{\mathcal{X}}_A} \mathbf{L}_{\hat{\mathcal{X}}_A} = \mathbf{s} - \mathbf{x}_{\mathcal{X}_A} \mathbf{L}_{\mathcal{X}_A} \end{cases}. \quad (2.13)$$

The adversary A attempts to solve this system. Given that $\mathbf{s}$ and $\mathbf{x}_{\mathcal{X}_A}$ are known, the system (2.13) has $n - t$ unknowns and N equations. Note that $t \leq d - 1$, and thus by applying Theorem 1.4.1 to $\mathcal{C}(\mathbf{L})$ we have $n - t \geq n - d + 1 \geq N$. If $\text{rank}_q(\mathbf{L}_{\hat{\mathcal{X}}_A}) = N$ then (2.13) has exactly q^{n-t-N} solutions, as required.

Next, we show that $\text{rank}_q(\mathbf{L}_{\hat{\mathcal{X}}_A}) = N$. Assume, by contrary, that the N (transposed) columns of $\mathbf{L}_{\hat{\mathcal{X}}_A}$, denoted by $\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \dots, \mathbf{c}^{(N)}$, are linearly dependent. Then there exist $\beta_i \in \mathbb{F}_q$ for $i \in [N]$, not all zero, such that $\sum_{i=1}^N \beta_i \mathbf{c}^{(i)} = \mathbf{0}$. Let

$$\mathbf{c} = \sum_{i=1}^N \beta_i \mathbf{L}[i]^T \in \mathcal{C}(\mathbf{L}) \setminus \{\mathbf{0}\}.$$

(Recall that $\mathbf{L}[i]$ denotes the i th column of $\mathbf{L}$). Then $\mathbf{c}_{\hat{\mathcal{X}}_A} = \sum_{i=1}^N \beta_i \mathbf{c}^{(i)} = \mathbf{0}$ and hence $\text{wt}(\mathbf{c}) = \text{wt}(\mathbf{c}_{\mathcal{X}_A}) \leq t \leq d - 1$. This is a contradiction, which follows

from the assumption that the N rows of $\mathbf{L}_{\hat{\mathcal{X}}_A}$ are linearly dependent. $\blacksquare$

Example 2.2.10. Let $q = 2$. Assume that $\mathcal{X}_A = \emptyset$ and that $\mathcal{X}_i \neq \emptyset$ for all $i \in [m]$. For each $i \in [m]$ choose some $j_i \in \mathcal{X}_i$. Let $\mathbf{L}$ be the binary matrix whose (transposed) columns form a basis of the space $\mathcal{C}(\mathbf{L}) = \text{span}_q(\{\mathbf{e}_{j_i} + \mathbf{e}_{f(i)}\}_{i \in [m]})$. Then $d(\mathcal{C}(\mathbf{L})) = 2$. Since $t = |\mathcal{X}_A| = 0$, we have $d - 1 - t = 1$. Therefore by Theorem 2.2.9 the index code based on $\mathbf{L}$ is weakly secure against A . Moreover, by the Singleton bound, $\mathbf{L}$ has at most $N \leq n - d + 1 = n - 1$ columns. In other words, the index code based on $\mathbf{L}$ requires at most $n - 1$ transmissions.

2.2.3 Complete Insecurity

Theorem 2.2.9 provides a threshold for the security level of a linear index code based on $\mathbf{L}$. If A has a prior knowledge of any $t \leq d - 2$ messages, where $d = d(\mathcal{C}(\mathbf{L}))$, then the index code is still secure, that is, the adversary has no information about any $d - 1 - t$ particular messages from $\{x_j\}_{j \in \hat{\mathcal{X}}_A}$. On the other hand, the index code may no longer be secure against an adversary of strength $t = d - 1$. The last assertion of Theorem 2.2.9 shows us the difference between being block secure and being strongly secure (we define the strong security rigorously later in Definition 2.4.2). More specifically, if the index code is strongly secure, the messages $\mathbf{x}_{\hat{\mathcal{X}}_A}$, which are not leaked to the adversary in advance, look completely random to the adversary, that is, the probability to guess them correctly is $1/q^{n-t}$. However, if the index code is $(d - 1 - t)$ -block secure (for $t \leq d - 2$), then the adversary is able to guess these messages correctly with probability $1/q^{n-t-N}$.

For an adversary of strength $t \geq d$, the security of the index code depends on the properties of $\mathbf{L}$, in particular, it depends on the weight distribution of $\mathcal{C}(\mathbf{L})$. From Theorem 2.2.9, if there exists $\mathbf{c} \in \mathcal{C}(\mathbf{L})$ with $\text{wt}(\mathbf{c}) = w$, then the index code

is not weakly secure against some adversary of strength $t = w - 1$. In general, the index code might still be (b -block or weakly) secure against some adversaries of strength t for $t \geq d$. While we cannot make a general conclusion on the security of the index code when the adversary's strength is larger than $d - 1$, Lemma 2.2.4 is still a useful tool to evaluate the security in that situation. However, as shown in Theorem 2.2.12, if the size of $\mathcal{X}_A$ is sufficiently large, then A is able to determine all the messages in $\{x_j\}_{j \in \hat{\mathcal{X}}_A}$. We first recall the following well-known result in coding theory.

Theorem 2.2.11 ([37], p. 66). *Let $\mathcal{C}$ be an $[n, k, d]_q$ code with dual distance $d^\perp$ and $\mathbf{M}$ denote the $q^k \times n$ matrix whose q^k rows are codewords of $\mathcal{C}$. If $r \leq d^\perp - 1$ then each r -tuple from $\mathbb{F}_q$ appears in an arbitrary set of r columns of $\mathbf{M}$ exactly q^{k-r} times.*

Theorem 2.2.12. *The linear index code based on $\mathbf{L}$ is completely insecure against any adversary of strength $t \geq n - d^\perp + 1$, where $d^\perp$ denotes the dual distance of $\mathcal{C}(\mathbf{L})$.*

Proof. Suppose the adversary knows a subset $\{x_j\}_{j \in \mathcal{X}_A}$, where $\mathcal{X}_A \subsetneq [n]$ and $|\mathcal{X}_A| = t \geq n - d^\perp + 1$. By Corollary 2.2.5, it suffices to show that for all $j \in \hat{\mathcal{X}}_A$, there exists $\mathbf{u} \in \mathbb{F}_q^n$ satisfying simultaneously $\mathbf{u} \triangleleft \mathcal{X}_A$ and $\mathbf{u} + \mathbf{e}_j \in \mathcal{C}(\mathbf{L})$.

Indeed, take any $j \in \hat{\mathcal{X}}_A$, and let $\rho = n - t \leq d^\perp - 1$. Consider the ρ indices that are not in $\mathcal{X}_A$. By Theorem 2.2.11, there exists a codeword $\mathbf{c} \in \mathcal{C}(\mathbf{L})$ with

$$c_\ell = \begin{cases} 1 & \text{if } \ell = j, \\ 0 & \text{if } \ell \notin \mathcal{X}_A \cup \{j\} \end{cases}.$$

Then $\text{supp}(\mathbf{c}) \subseteq \mathcal{X}_A \cup \{j\}$. We define $\mathbf{u} \in \mathbb{F}_q^n$ such that $\mathbf{u} \triangleleft \mathcal{X}_A$ as follows. For

$\ell \in \mathcal{X}_A$, we set $u_\ell = c_\ell$, and for $\ell \notin \mathcal{X}_A$, we set $u_\ell = 0$. It is immediately clear that $\mathbf{c} = \mathbf{u} + \mathbf{e}_j$. Therefore, by Corollary 2.2.5, the adversary can reconstruct x_j . We have shown that the index code is completely insecure against an arbitrary set $\mathcal{X}_A$ satisfying $|\mathcal{X}_A| \geq n - d^\perp + 1$, hence completing the proof. ■

When $\mathcal{C}(\mathbf{L})$ is an MDS code, we have $n - d^\perp + 1 = d - 1$, and hence the two bounds established in Theorems 2.2.9 and 2.2.12 are actually tight. In that case, the third statement in Theorem 2.2.9 implies Theorem 2.2.12 as follows. This statement asserts that an adversary of strength $t = d - 1$ can find a list of $q^{n-N-d+1}$ vectors that includes the vector of messages $\mathbf{x}$. Since $\mathcal{C}(\mathbf{L})$ is an MDS code, we have $n - N - d + 1 = 0$. Therefore, the list contains only one element, namely $\mathbf{x}$ itself. Thus the index code is completely insecure against any adversary of strength $d - 1$. The following example further illustrates the results stated in these theorems.

Example 2.2.13. Let $n = 7$, $m = 7$, $q = 2$, and $f(i) = i$ for all $i \in [m]$. Suppose that the receivers have in their possession sets of messages as appear in the third column of the table below. Suppose also, that the demands of all receivers are as in the second column of the table.

Receiver	Demand	$\{x_j\}_{j \in \mathcal{X}_i}$
R_1	x_1	$\{x_6, x_7\}$
R_2	x_2	$\{x_5, x_7\}$
R_3	x_3	$\{x_5, x_6\}$
R_4	x_4	$\{x_5, x_6, x_7\}$
R_5	x_5	$\{x_1, x_2, x_6\}$
R_6	x_6	$\{x_1, x_3, x_4\}$
R_7	x_7	$\{x_2, x_3, x_6\}$

For $i \in [7]$, let $\mathbf{u}^{(i)} \in \mathbb{F}_2^7$ such that $\text{supp}(\mathbf{u}^{(i)}) = \mathcal{X}_i$. Assume that an index code based on $\mathbf{L}$ with $\mathcal{C}(\mathbf{L}) = \text{span}_q(\{\mathbf{u}^{(i)} + \mathbf{e}_i\}_{i \in [7]})$ is used. For instance, we can take $\mathbf{L}$ to be the matrix whose set of columns is $\{\mathbf{L}[i] \triangleq (\mathbf{u}^{(i)} + \mathbf{e}_i)^T\}_{i \in [4]}$. It is easy to see that $\mathcal{C}(\mathbf{L})$ is a $[7, 4, 3]_2$ Hamming code with $d = 3$ and $d^\perp = 4$.

Following the coding scheme, S broadcasts the following four bits:

$$\begin{aligned} s_1 &= \mathbf{x}(\mathbf{u}^{(1)} + \mathbf{e}_1)^T, \\ s_2 &= \mathbf{x}(\mathbf{u}^{(2)} + \mathbf{e}_2)^T, \\ s_3 &= \mathbf{x}(\mathbf{u}^{(3)} + \mathbf{e}_3)^T, \\ s_4 &= \mathbf{x}(\mathbf{u}^{(4)} + \mathbf{e}_4)^T. \end{aligned}$$

Each R_i for $i \in [7]$, can compute $\mathbf{x}(\mathbf{u}^{(i)} + \mathbf{e}_i)^T$ by using a linear combination of s_1, s_2, s_3, s_4 . Then, each R_i can subtract $\mathbf{x}(\mathbf{u}^{(i)})^T$ (his side information) from $\mathbf{x}(\mathbf{u}^{(i)} + \mathbf{e}_i)^T$ to retrieve $x_i = \mathbf{x}\mathbf{e}_i^T$.

For example, consider R_5 . Since

$$\mathbf{x}(\mathbf{u}^{(5)} + \mathbf{e}_5)^T = \mathbf{x}((\mathbf{u}^{(1)} + \mathbf{e}_1) + (\mathbf{u}^{(2)} + \mathbf{e}_2))^T = s_1 + s_2,$$

R_5 subtracts $x_1 + x_2 + x_6$ from $s_1 + s_2$ to obtain

$$(s_1 + s_2) - (x_1 + x_2 + x_6) = (x_1 + x_2 + x_5 + x_6) - (x_1 + x_2 + x_6) = x_5.$$

If an adversary A has knowledge of a single message x_i , then by Theorem 2.2.9, A is not able to determine any other message x_ℓ , for $\ell \neq i$. Indeed, $\mathbf{d}(\mathcal{C}(\mathbf{L})) = 3$, while $t = 1$, the code is weakly secure against all adversaries of strength $t = 1$. If none of the messages are leaked, then the adversary has no information about any group of 2 messages. On the other hand, the code is completely insecure against any adversary of strength $t \geq 4$; in that case A is able to determine the remaining

$7 - t$ messages.

2.3 Index Coding with Side and Restricted Information

Results on the block security of a linear index code can be employed to study the linear coding schemes for an extension of the ICSI problem, so-called the Index Coding with Side and Restricted Information (ICSRI) problem. An instance $(m, n, \mathcal{X}, \mathcal{Z}, f)$ of the ICSRI problem consists of almost the same parameters as that of the ICSI problem. The only new parameter, which is $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2, \dots, \mathcal{Z}_m)$, represents the sets of messages that the receivers are *not* allowed to obtain. The goal is that at the end of the communication round, for each $i \in [m]$, the receiver R_i obtains the message $x_{f(i)}$ and has no information about x_j for all $j \in \mathcal{Z}_i$. The notion of a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is naturally extended to that of a linear $(m, n, \mathcal{X}, \mathcal{Z}, f)$ -IC over $\mathbb{F}_q$.

Let

$$\mathcal{F}(m, n, \mathcal{X}, \mathcal{Z}, f) \triangleq \bigcup_{i=1}^m \{\mathbf{u} + \mathbf{e}_j : \mathbf{u} \triangleleft \mathcal{X}_i, j \in \mathcal{Z}_i\}.$$

The following proposition provides a necessary and sufficient condition for a linear index code to be also a solution to an instance of the ICSRI problem.

Proposition 2.3.1. *The linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$ is also a linear $(m, n, \mathcal{X}, \mathcal{Z}, f)$ -IC if and only if $\mathcal{C}(\mathbf{L}) \cap \mathcal{F}(m, n, \mathcal{X}, \mathcal{Z}, f) = \emptyset$.*

Proof. Let S employ the $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$. Then clearly R_i can recover $x_{f(i)}$ for all $i \in [m]$. Due to Lemma 2.2.4, for each $i \in [m]$ and $j \in \mathcal{Z}_i$, the receiver R_i has no information about x_j if and only if

$$\forall \mathbf{u} \triangleleft \mathcal{X}_i : \mathbf{u} + \mathbf{e}_j \notin \mathcal{C}(\mathbf{L}).$$

Hence we complete the proof. ■

Example 2.3.2. Consider an instance $(m, n, \mathcal{X}, \mathcal{Z}, f)$ of the ICSRI problem where $m, n, \mathcal{X}$, and f are defined as in Example 2.2.13. Moreover, $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2, \dots, \mathcal{Z}_7)$, where $\mathcal{Z}_1 = \{2, 3, 4, 5\}$, $\mathcal{Z}_2 = \{1, 3, 4, 6\}$, $\mathcal{Z}_3 = \{1, 2, 4, 7\}$, and $\mathcal{Z}_4 = \mathcal{Z}_5 = \mathcal{Z}_6 = \mathcal{Z}_7 = \emptyset$. Consider the index code based on $\mathbf{L}$ constructed in Example 2.2.13. It is straightforward to verify that $\mathcal{C}(\mathbf{L}) \cap \mathcal{F}(m, n, \mathcal{X}, \mathcal{Z}, f) = \emptyset$. Therefore, by Proposition 2.3.1, this index code also provides a solution to this instance of the ICSRI problem.

Let

$$\bar{\kappa}_q = \bar{\kappa}_q(m, n, \mathcal{X}, \mathcal{Z}, f) \triangleq \min \left\{ \text{rank}_q \left(\{ \mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \}_{i \in [m]} \right) \right\},$$

where the minimum is taken over all choices of $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$ that satisfy

$$\text{span}_q \left(\{ \mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \} \right) \cap \mathcal{F}(m, n, \mathcal{X}, \mathcal{Z}, f) = \emptyset. \quad (2.14)$$

Let $\bar{\kappa}_q = +\infty$ if there are no choices of $\mathbf{u}^{(i)}$'s that satisfy (2.14). Then the following proposition is immediate.

Proposition 2.3.3. *If $\bar{\kappa}_q < +\infty$ then the length of an optimal linear $(m, n, \mathcal{X}, \mathcal{Z}, f)$ -IC over $\mathbb{F}_q$ is $\bar{\kappa}_q$. If $\bar{\kappa}_q = +\infty$ then there does not exist a linear $(m, n, \mathcal{X}, \mathcal{Z}, f)$ -IC over $\mathbb{F}_q$.*

2.4 Strongly Secure Index Code with Side Information

In this section, we consider a less powerful adversary, who owns some prior side information and eavesdrops *at most* μ transmissions.

2.4.1 A Lower Bound on the Length of a Strongly Secure Index Code

We start this section with a generalization of the definition of index codes to that of *randomized index codes*. Consider $\eta \in \mathbb{N}$ random variables $G_1, G_2, \dots, G_\eta$, which are distributed independently and uniformly over $\mathbb{F}_q$. Let $\mathbf{G} = (G_1, G_2, \dots, G_\eta)$ and let $\mathbf{g} = (g_1, g_2, \dots, g_\eta)$ be a realization of $\mathbf{G}$.

Definition 2.4.1. An η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ for an ICSI instance $(m, n, \mathcal{X}, f)$ is an encoding function

$$\mathfrak{E} : \mathbb{F}_q^n \times \mathbb{F}_q^\eta \rightarrow \mathbb{F}_q^N,$$

such that for each receiver R_i , $i \in [m]$, there exists a decoding function

$$\mathfrak{D}_i : \mathbb{F}_q^N \times \mathbb{F}_q^{|\mathcal{X}_i|} \rightarrow \mathbb{F}_q,$$

satisfying

$$\forall \mathbf{x} \in \mathbb{F}_q^n : \mathfrak{D}_i(\mathfrak{E}(\mathbf{x}, \mathbf{g}), \mathbf{x}_{\mathcal{X}_i}) = x_{f(i)},$$

for every $\mathbf{g} \in \mathbb{F}_q^\eta$, which is a realization of the random vector $\mathbf{G}$.

An η -randomized index code is linear over $\mathbb{F}_q$ if it has a linear encoding function $\mathfrak{E}$,

$$\mathfrak{E}(\mathbf{x}, \mathbf{g}) = (\mathbf{x}|\mathbf{g})\mathbf{L},$$

where $\mathbf{L}$ is an $(n + \eta) \times N$ matrix over $\mathbb{F}_q$. Observe that by simply treating $x_1, x_2, \dots, x_n, g_1, g_2, \dots, g_\eta$ as messages, the results from previous sections still apply to linear randomized index codes.

Definition 2.4.2. The linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ based on $\mathbf{L}$ is

said to be (μ, t) -strongly secure if the following holds. An adversary A who possesses t arbitrary messages $\mathbf{x}_{\mathcal{X}_A}$, for some t -subset $\mathcal{X}_A$ of $[n]$, and is able to listen to at most μ among N transmissions, gains no information about other messages. Equivalently,

$$\mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A} \mid \mathbf{X}_{\mathcal{X}_A}, (\mathbf{X}|\mathbf{G})\mathbf{L}[W]) = \mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A}),$$

for any $W \subseteq [N]$ and $|W| \leq \mu$.

Remark 2.4.3. If $\mu = t = \eta = 0$, a (μ, t) -strongly secure η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ becomes a normal $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.

We henceforth assume that each message is requested by at least one receiver, for otherwise, that “useless” message can be discarded without doing any harm to the coding scheme.

Lemma 2.4.4. *If $\mathbf{L}$ corresponds to a (μ, t) -strongly secure linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$, then $\eta \geq \mu$.*

Proof. We prove this lemma by contradiction. Suppose that $\mathbf{L}$ corresponds to a (μ, t) -strongly secure η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ and that $\eta < \mu$.

For a subset W of $[N]$ let $\mathcal{C}(\mathbf{L}[W])$ be the space spanned by (transposed) columns of $\mathbf{L}$ indexed by elements of W . Let $E = \{n+1, n+2, \dots, n+\eta\}$. For all subsets W of $[N]$ with $|W| \leq \mu$, it holds that

$$\mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A} \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) = \mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A}),$$

that is, an adversary who owns $\mathbf{x}_{\mathcal{X}_A}$ gains no information about $\mathbf{x}_{\hat{\mathcal{X}}_A}$ after eavesdropping $|W|$ transmissions. Therefore, applying Lemma 2.2.4 with $\mathcal{C}(\mathbf{L})$ being replaced

by $\mathcal{C}(\mathbf{L}[W])$ (equivalently, $\mathbf{L}$ being replaced by $\mathbf{L}[W]$), we conclude that $\mathcal{C}(\mathbf{L}[W])$ does not contain a vector $\mathbf{c}$ that satisfies $\mathbf{c}_{\hat{\mathcal{X}}_A} \neq \mathbf{0}$ and $\mathbf{c}_E = \mathbf{0}$. We refer to this property of $\mathcal{C}(\mathbf{L}[W])$ as Property A.

Let $\mathbf{L}' = (\mathbf{L}_{\hat{\mathcal{X}}_A \cup E})^T$ be the matrix obtained from $\mathbf{L}$ by first deleting rows of $\mathbf{L}$ indexed by $\mathcal{X}_A$, and then taking its transpose. We claim that $\text{rank}_q(\mathbf{L}') \leq \mu - 1$. Indeed, take any μ rows of $\mathbf{L}'$, say $\mathbf{L}'_{j_1}, \dots, \mathbf{L}'_{j_\mu}$, we show that these μ rows are linearly dependent. Let $\mathbf{L}''$ be the submatrix of $\mathbf{L}'$ formed by the last η columns. Since $\eta < \mu$, the μ rows $\mathbf{L}''_{j_1}, \dots, \mathbf{L}''_{j_\mu}$ are linearly dependent. Hence there exist $\alpha_1, \alpha_2, \dots, \alpha_\mu$, not all zeros, such that

$$\sum_{\ell=1}^{\mu} \alpha_\ell \mathbf{L}''_{j_\ell} = \mathbf{0},$$

which implies

$$\sum_{\ell=1}^{\mu} \alpha_\ell \mathbf{L}'_{j_\ell} = \mathbf{0},$$

due to Property A. The claim follows.

Now let $r = \text{rank}_q(\mathbf{L}') < \mu$, and let

$$\{\mathbf{L}'_{j_1}, \mathbf{L}'_{j_2}, \dots, \mathbf{L}'_{j_r}\}$$

be a basis of the space spanned by rows of $\mathbf{L}'$. Suppose that the receiver R_i requests $x_{f(i)}$ where $f(i) \in \hat{\mathcal{X}}_A$. By Corollary 2.2.5, $\mathcal{C}(\mathbf{L})$ contains a vector $\mathbf{c} = \mathbf{u} + \mathbf{e}_{f(i)}$ where $\mathbf{u} \triangleleft \mathcal{X}_i$. Therefore, $\mathbf{c}_E = \mathbf{0}$ and $\mathbf{c}_{\hat{\mathcal{X}}_A} \neq \mathbf{0}$. On the other hand, there exists $\beta_1, \beta_2, \dots, \beta_r$ such that

$$(\mathbf{c}_{\hat{\mathcal{X}}_A} | \mathbf{c}_E) = \sum_{\ell=1}^r \beta_\ell \mathbf{L}'_{j_\ell}.$$

Since $r < \mu$ and $\mathbf{c}_E = \mathbf{0}$, by Property A we have $\mathbf{c}_{\hat{\mathcal{X}}_A} = \mathbf{0}$, which is a contradiction. ■

Remark 2.4.5. By Lemma 2.4.4, a randomized index code requires at least μ random symbols in order to be (μ, t) -strongly secure. We show in the next subsection that using μ random symbols is also sufficient.

Lemma 2.4.6. *Suppose that $\mathbf{L}$ corresponds to a linear μ -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. If this randomized index code is (μ, t) -strongly secure, then for all $i \in [\mu]$, there exists a vector $\mathbf{v}^{(i)} \in \mathbb{F}_q^{n+\mu}$ satisfying*

1. $\mathbf{v}^{(i)} \triangleleft [n]$;
2. $\mathbf{v}^{(i)} + \mathbf{e}_{n+i} \in \mathcal{C}(\mathbf{L})$.

Proof. Assume, by contradiction, that for some $i \in [\mu]$, we have $\mathbf{v}^{(i)} + \mathbf{e}_{n+i} \notin \mathcal{C}(\mathbf{L})$ for all $\mathbf{v}^{(i)} \triangleleft [n]$. By Corollary 2.2.5, a virtual receiver with side information $\{x_j\}_{j \in [n]}$ would have no information about g_i after receiving all transmissions. In other words, we have

$$\mathsf{H}(G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}, \mathbf{X}) = \mathsf{H}(G_i), \quad (2.15)$$

and, in particular, for a smaller set of side information,

$$\mathsf{H}(G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}, \mathbf{X}_{\mathcal{X}_A}) = \mathsf{H}(G_i). \quad (2.16)$$

Since the randomized index code is (μ, t) -strongly secure, for every μ -subset W of $[N]$ and every t -subset $\mathcal{X}_A$ of $[n]$, we have

$$\mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A} \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) = \mathsf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A}). \quad (2.17)$$

We claim that if the realized value of G_i is known to the adversary, this randomized

index code is still (μ, t) -strongly secure. In other words, we aim to show that

$$\mathbf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A} \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}, G_i) = \mathbf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A}), \quad (2.18)$$

for every μ -subset W of $[N]$ and every t -subset $\mathcal{X}_A$ of $[n]$. Indeed, the left-hand side of (2.18) is equal to

$$\mathbf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A} \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) - \mathbf{I}(\mathbf{X}_{\hat{\mathcal{X}}_A}; G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}),$$

which is

$$\mathbf{H}(\mathbf{X}_{\hat{\mathcal{X}}_A}) - \mathbf{I}(\mathbf{X}_{\hat{\mathcal{X}}_A}; G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A})$$

due to (2.17). Hence, it suffices to show that

$$\mathbf{I}(\mathbf{X}_{\hat{\mathcal{X}}_A}; G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) = 0.$$

Using (2.15) we have

$$\begin{aligned} \mathbf{I}(\mathbf{X}_{\hat{\mathcal{X}}_A}; G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) &= \mathbf{H}(G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) \\ &\quad - \mathbf{H}(G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}, \mathbf{X}_{\hat{\mathcal{X}}_A}) \\ &= \mathbf{H}(G_i \mid (\mathbf{X} \mid \mathbf{G})\mathbf{L}[W], \mathbf{X}_{\mathcal{X}_A}) \\ &\quad - \mathbf{H}(G_i \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W], \mathbf{X}) \\ &= \mathbf{H}(G_i) - \mathbf{H}(G_i) \\ &= 0, \end{aligned}$$

where the third transition is due to (2.15) and (2.16). Thus, we have shown that the randomized index code is still (μ, t) -strongly secure if the adversary knows the

realized value of G_i . Equivalently, discarding the random variable G_i from the scheme does not affect its strong security. However, this contradicts Lemma 2.4.4, as the resulting code has less than μ random symbols. ■

The following theorem proves a lower bound on the length of a (μ, t) -strongly secure linear randomized index code.

Theorem 2.4.7. *The length of a (μ, t) -strongly secure linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is at least $\kappa_q + \mu$. In other words, a linear randomized index code requires at least $\kappa_q + \mu$ transmissions in order to be (μ, t) -strongly secure.*

Proof. Suppose the linear randomized index code is based on $\mathbf{L}$. We divide the proof into several cases.

Case 1: $\eta = \mu$. Then, by Corollary 2.2.6 and Lemma 2.4.6, the subspace $\mathcal{C}(\mathbf{L})$ must contain:

1. the vectors $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)}$ for some $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$, for all $i \in [m]$;
2. the vectors $\mathbf{v}^{(i)} + \mathbf{e}_{n+i}$, for some $\mathbf{v}^{(i)} \triangleleft [n]$, for all $i \in [\mu]$.

Due to linear independence of these vectors and to the definition of κ_q , the length of the code is at least

$$\begin{aligned} \dim(\mathcal{C}(\mathbf{L})) &\geq \text{rank}_q \left(\{ \mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \}_{i \in [m]} \right) + \text{rank}_q \left(\{ \mathbf{v}^{(i)} + \mathbf{e}_{n+i} \}_{i \in [\mu]} \right) \\ &\geq \kappa_q + \mu. \end{aligned}$$

Case 2: $\eta > \mu$, and for every $i \in [\eta]$ there exists some vector $\mathbf{v}^{(i)} \triangleleft [n]$ such that $\mathbf{v}^{(i)} + \mathbf{e}_{n+i} \in \mathcal{C}(\mathbf{L})$. In this case, similarly to Case 1, we have

$$\dim(\mathcal{C}(\mathbf{L})) \geq \kappa_q + \eta > \kappa_q + \mu.$$

Therefore, $\mathbf{L}$ has at least $\kappa_q + \mu$ columns.

Case 3: $\eta > \mu$, and for some $i \in [\eta]$, it holds that $\mathbf{v}^{(i)} + \mathbf{e}_{n+i} \notin \mathcal{C}(\mathbf{L})$ for all $\mathbf{v}^{(i)} \triangleleft [n]$.

By following exactly the same argument as in the proof of Lemma 2.4.6, we deduce that discarding G_i does not affect the strong security of the randomized index code. By doing so, we obtain a new (μ, t) -strongly secure randomized index code, which has $\eta - 1$ random variables. This code is based on $\mathbf{L}'$, which is obtained from $\mathbf{L}$ by deleting its $(n + i)$ th row.

The above argument can be applied until either the number of random variables decreases to μ , or the code in consideration satisfies the condition of Case 2. In both cases, the resulting randomized index code has length at least $\kappa_q + \mu$. As the length of the code does not change during the process, we conclude that the length of the original code is at least $\kappa_q + \mu$. $\blacksquare$

2.4.2 A Construction of Optimal Strongly Secure Index Codes

In this subsection, we provide a construction of an optimal (μ, t) -strongly secure μ -randomized linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$, which has length attaining the lower bound established in Theorem 2.4.7. This construction requires q to be as large as $\kappa_q + \mu - 1$. The proposed construction is based on the coset coding technique, originally introduced by Ozarow and Wyner [52]. This technique has been adopted in a variety of network coding applications [12, 26, 29, 59, 71].

Construction A: Let $\mathbf{L}^{(0)}$ correspond to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ of optimal length κ_q . Let $\mathbf{M}$ be an invertible matrix of order $N = \kappa_q + \mu$, so that the last μ rows of $\mathbf{M}$ form a generator matrix of an MDS code. Let $\mathbf{P}$ be the submatrix of $\mathbf{M}$ formed by the first κ_q rows, and $\mathbf{Q}$ the submatrix formed by the last μ rows

of $\mathbf{M}$. Let

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}^{(0)} \mathbf{P} \\ \mathbf{Q} \end{pmatrix}.$$

Lemma 2.4.8. *The matrix $\mathbf{L}$ in Construction A corresponds to a μ -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.*

Proof. Let $\mathbf{g} \in \mathbb{F}_q^\mu$ be a random vector, uniformly distributed over $\mathbb{F}_q^\mu$. The encoding process is as follows

$$\mathbf{x} \mapsto (\mathbf{x}|\mathbf{g})\mathbf{L} = \mathbf{x}\mathbf{L}^{(0)}\mathbf{P} + \mathbf{g}\mathbf{Q} = (\mathbf{x}\mathbf{L}^{(0)}|\mathbf{g})\mathbf{M}.$$

Since $\mathbf{M}$ is invertible, each receiver is able to recover $(\mathbf{x}\mathbf{L}^{(0)}|\mathbf{g})$. Therefore, for each $i \in [m]$, the receiver R_i can recover $\mathbf{x}\mathbf{L}^{(0)}$, and hence, can also recover $x_{f(i)}$, as $\mathbf{L}^{(0)}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. $\blacksquare$

Lemma 2.4.9. *The μ -randomized $(m, n, \mathcal{X}, f)$ -IC based on the matrix $\mathbf{L}$ in Construction A is (μ, t) -strongly secure.*

Proof. Suppose that the adversary A possess a message vector $\mathbf{x}_{\mathcal{X}_A}$, where $|\mathcal{X}_A| = t$. Additionally, A can eavesdrop μ transmissions, that is, it has knowledge of $\mathbf{b} \triangleq (\mathbf{x}|\mathbf{g})\mathbf{L}[W]$, for some $W \subseteq [N]$ and $|W| = \mu$. Below, we show that the entropy of $\mathbf{X}_{\hat{\mathcal{X}}_A}$ is not changed given the knowledge of $(\mathbf{X}|\mathbf{G})\mathbf{L}[W]$ and of $\mathbf{x}_{\mathcal{X}_A}$. It suffices to show that for all $\mathbf{a} \in \mathbb{F}_q^{n-t}$:

$$\Pr(\mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a} \mid (\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) = \frac{1}{q^{n-t}}. \quad (2.19)$$

The probability on the left-hand side of (2.19) can be rewritten as

$$\frac{\Pr(\mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a}, (\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A})}{\Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A})}. \quad (2.20)$$

The numerator in (2.20) is given by

$$\begin{aligned} & \Pr(\mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a}, (\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ &= \Pr(\mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ & \quad \times \Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{x}_{\hat{\mathcal{X}}_A} = \mathbf{a}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ &= \frac{1}{q^{n-t}} \Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ &= \frac{1}{q^{n-t}} \frac{1}{q^\mu} \\ &= \frac{1}{q^{n-t+\mu}}. \end{aligned} \quad (2.21)$$

The penultimate transition can be explained as follows. We have

$$\mathbf{b} = (\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{X}\mathbf{L}^{(0)}\mathbf{P}[W] + \mathbf{G}\mathbf{Q}[W]. \quad (2.22)$$

The matrix $\mathbf{Q}[W]$ is invertible due to the fact that $\mathbf{Q}$ is a generator matrix of an $[N, \mu]$ MDS code. Since $\mathbf{X}$ is known, the system (2.22) has a unique solution given by

$$\mathbf{G} = (\mathbf{b} - \mathbf{X}\mathbf{L}^{(0)}\mathbf{P}[W])(\mathbf{Q}[W])^{-1}.$$

Since $\mathbf{G}$ is uniformly distributed over $\mathbb{F}_q^\mu$,

$$\begin{aligned} & \Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{a}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ &= \Pr\left(\mathbf{G} = (\mathbf{b} - \mathbf{X}\mathbf{L}^{(0)}\mathbf{P}[W])(\mathbf{Q}[W])^{-1}\right) \\ &= \frac{1}{q^\mu}. \end{aligned}$$

Similarly to (2.21), the denominator in (2.20) is

$$\begin{aligned} & \Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ &= \sum_{\mathbf{c} \in \mathbb{F}_q^{n-t}} \Pr(\mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{c} \mid \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \\ & \quad \times \Pr((\mathbf{X}|\mathbf{G})\mathbf{L}[W] = \mathbf{b} \mid \mathbf{X}_{\hat{\mathcal{X}}_A} = \mathbf{c}, \mathbf{X}_{\mathcal{X}_A} = \mathbf{x}_{\mathcal{X}_A}) \quad (2.23) \\ &= q^{n-t} \frac{1}{q^{n-t}} \frac{1}{q^\mu} \\ &= \frac{1}{q^\mu}. \end{aligned}$$

From (2.20), (2.21), and (2.23), we obtain (2.19), as claimed. ■

From Theorem 2.4.7, Lemma 2.4.8, and Lemma 2.4.9, we have the following theorem.

Theorem 2.4.10. *The length of an optimal (μ, t) -strongly secure linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is at least $\kappa_q + \mu$. Moreover, the code based on the matrix $\mathbf{L}$ established in Construction A achieves this bound ($q \geq \kappa_q + \mu - 1$).*

3. INDEX CODE AND ERROR CORRECTION

Part of the work in this chapter was presented in the 2011 IEEE Symposium on Information Theory [21]. In this chapter, a generalization of index coding scheme, where transmitted symbols are subject to errors, is studied. Error-correcting methods for such a scheme, and their parameters, are investigated. In particular, the following question is discussed: given the side information hypergraph of index coding scheme and the maximal number of erroneous symbols δ , what is the shortest (optimal) length of a linear index code which guarantees that every receiver is able to recover the required information? This question turns out to be a generalization of the problem of finding the shortest length of an error-correcting code with a prescribed error-correcting capability in the classical coding theory.

The Singleton bound and two other bounds, referred to as the α -bound and the κ -bound, for the optimal length of a linear error-correcting index code (ECIC) are established. For large alphabets, a construction based on the concatenation of an optimal index code with an MDS classical code, is shown to attain the Singleton bound. For smaller alphabets, however, this construction may not be optimal. A random construction is also analyzed. It yields another implicit bound on the length of an optimal linear ECIC.

The problem of decoding a linear ECIC is studied. It is shown that in order to decode correctly the desired symbol, the decoder is required to find a vector that belongs to an affine space containing the actual error vector. The syndrome

decoding is shown to produce the correct output if the weight of the error pattern is less or equal to the error-correcting capability of the corresponding ECIC.

Furthermore, we introduce the notion of a static ECIC, which is suitable for use with a family of ICSI instances. Several bounds on the length of static ECIC's are derived, and constructions for static ECIC's are discussed. Connections of these codes to weakly resilient Boolean functions are also established.

Finally, we discuss the nonlinear ECIC's. Analogous bounds on the length of an optimal ECIC are established.

3.1 Error-Correcting Index Code with Side Information

Due to noise, the symbols received by $R_i, i \in [m]$, may be subject to errors. Consider an ICSI instance $(m, n, \mathcal{X}, f)$, and assume that S broadcasts a vector $\mathfrak{E}(\mathbf{x}) \in \mathbb{F}_q^N$. Let $\boldsymbol{\epsilon}^{(i)} \in \mathbb{F}_q^N$ be the error affecting the information received by $R_i, i \in [m]$. Then R_i actually receives the vector

$$\mathbf{y}_i = \mathfrak{E}(\mathbf{x}) + \boldsymbol{\epsilon}^{(i)} \in \mathbb{F}_q^N,$$

instead of $\mathfrak{E}(\mathbf{x})$. The following definition is a generalization of Definition 2.1.3.

Definition 3.1.1. Consider an ICSI instance described by $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$. A δ -error-correcting index code (δ -error-correcting $\mathcal{H}$ -IC, or simpler, $(\delta, \mathcal{H})$ -ECIC) over $\mathbb{F}_q$ for this instance is an encoding function

$$\mathfrak{E} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^N,$$

such that for each receiver R_i , $i \in [m]$, there exists a decoding function

$$\mathfrak{D}_i : \mathbb{F}_q^N \times \mathbb{F}_q^{|\mathcal{X}_i|} \rightarrow \mathbb{F}_q,$$

satisfying

$$\forall \mathbf{x}, \boldsymbol{\epsilon}^{(i)} \in \mathbb{F}_q^n, \text{wt}(\boldsymbol{\epsilon}^{(i)}) \leq \delta : \mathfrak{D}_i(\mathfrak{E}(\mathbf{x}) + \boldsymbol{\epsilon}^{(i)}, \mathbf{x}_{\mathcal{X}_i}) = x_{f(i)}.$$

The definitions of the length, of a linear index code, and of the matrix corresponding to a linear index code are naturally extended to an error-correcting index code. Note that if $\mathfrak{E}$ is an $\mathcal{H}$ -IC over $\mathbb{F}_q$, then it is a $(0, \mathcal{H})$ -ECIC, and vice versa.

Definition 3.1.2. An *optimal* linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ is a linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ of the smallest possible length $\mathcal{N}_q[\delta, \mathcal{H}]$.

Consider an ICSI instance described by $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$. We define the set of vectors

$$\mathcal{I}(q, \mathcal{H}) \triangleq \{ \mathbf{z} \in \mathbb{F}_q^n : \exists i \in [m] \text{ such that } \mathbf{z}_{\mathcal{X}_i} = \mathbf{0} \text{ and } z_{f(i)} \neq 0 \}.$$

For all $i \in [m]$, we also define

$$\mathcal{Y}_i \triangleq [n] \setminus (\{f(i)\} \cup \mathcal{X}_i).$$

Then the collection of supports of all vectors in $\mathcal{I}(q, \mathcal{H})$ is given by

$$\mathcal{J}(\mathcal{H}) \triangleq \bigcup_{i \in [m]} \{ \{f(i)\} \cup Y_i : Y_i \subseteq \mathcal{Y}_i \}. \quad (3.1)$$

The necessary and sufficient condition for a matrix $\mathbf{L}$ to be the matrix corresponding to some $(\delta, \mathcal{H})$ -ECIC is given in the following lemma.

Lemma 3.1.3. *The matrix $\mathbf{L}$ corresponds to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ if and only if*

$$\text{wt}(\mathbf{zL}) \geq 2\delta + 1 \text{ for all } \mathbf{z} \in \mathcal{I}(q, \mathcal{H}). \quad (3.2)$$

Equivalently, $\mathbf{L}$ corresponds to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ if and only if

$$\text{wt}\left(\sum_{i \in K} z_i \mathbf{L}_i\right) \geq 2\delta + 1,$$

for all $K \in \mathcal{J}(\mathcal{H})$ and for all choices of $z_i \in \mathbb{F}_q^$, $i \in K$.*

Proof. For each $\mathbf{x} \in \mathbb{F}_q^n$, we define

$$B(\mathbf{x}, \delta) = \{\mathbf{y} \in \mathbb{F}_q^N : \mathbf{y} = \mathbf{xL} + \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \in \mathbb{F}_q^N, \text{wt}(\boldsymbol{\epsilon}) \leq \delta\},$$

the set of all vectors resulting from at most δ errors in the transmitted vector associated with the information vector $\mathbf{x}$. Then the receiver R_i can recover $x_{f(i)}$ correctly if and only if

$$B(\mathbf{x}, \delta) \cap B(\mathbf{x}', \delta) = \emptyset,$$

for every pair $\mathbf{x}, \mathbf{x}' \in \mathbb{F}_q^n$ satisfying:

$$\mathbf{x}_{\mathcal{X}_i} = \mathbf{x}'_{\mathcal{X}_i} \text{ and } x_{f(i)} \neq x'_{f(i)}.$$

(Recall that R_i is interested only in the bit $x_{f(i)}$, not in the whole vector $\mathbf{x}$.)

Therefore, $\mathbf{L}$ corresponds to a $(\delta, \mathcal{H})$ -ECIC if and only if the following condition is satisfied: for all $i \in [m]$ and for all $\mathbf{x}, \mathbf{x}' \in \mathbb{F}_q^n$ such that $\mathbf{x}_{\mathcal{X}_i} = \mathbf{x}'_{\mathcal{X}_i}$ and $x_{f(i)} \neq x'_{f(i)}$,

it holds

$$\forall \epsilon, \epsilon' \in \mathbb{F}_q^N, \text{wt}(\epsilon) \leq \delta, \text{wt}(\epsilon') \leq \delta : \mathbf{x}\mathbf{L} + \epsilon \neq \mathbf{x}'\mathbf{L} + \epsilon'. \quad (3.3)$$

Denote $\mathbf{z} = \mathbf{x}' - \mathbf{x}$. Then, the condition in (3.3) can be reformulated as follows: for all $i \in [n]$ and for all $\mathbf{z} \in \mathbb{F}_q^n$ such that $\mathbf{z}_{\mathcal{X}_i} = \mathbf{0}$ and $z_{f(i)} \neq 0$, it holds

$$\forall \epsilon, \epsilon' \in \mathbb{F}_q^N, \text{wt}(\epsilon) \leq \delta, \text{wt}(\epsilon') \leq \delta : \mathbf{z}\mathbf{L} \neq \epsilon - \epsilon'. \quad (3.4)$$

An equivalent condition is that for all $\mathbf{z} \in \mathcal{I}(q, \mathcal{H})$,

$$\text{wt}(\mathbf{z}\mathbf{L}) \geq 2\delta + 1.$$

Since for $\mathbf{z} \in \mathcal{I}(q, \mathcal{H})$ we have

$$\mathbf{z}\mathbf{L} = \sum_{i \in \text{supp}(\mathbf{z})} z_i \mathbf{L}_i,$$

the condition (3.2) can be restated as

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 2\delta + 1,$$

for all $K \in \mathcal{J}(\mathcal{H})$ and for all choices of nonzero $z_i \in \mathbb{F}_q$, $i \in K$. ■

The next corollary follows from Lemma 3.1.3 in a straightforward manner. It is not hard to see that the conditions stated in Lemma 3.1.3 and in the corollary below are, in fact, equivalent.

Corollary 3.1.4. *For all $i \in [m]$, let*

$$\mathbf{M}_i \triangleq \text{span}_q(\{\mathbf{L}_j : j \in \mathcal{Y}_i\}).$$

Then, the matrix $\mathbf{L}$ corresponds to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ if and only if

$$\forall i \in [m] : d(\mathbf{L}_{f(i)}, \mathbf{M}_i) \geq 2\delta + 1. \quad (3.5)$$

The next corollary also follows directly from Lemma 3.1.3 by considering an error-free setup, that is, $\delta = 0$.

Corollary 3.1.5. *The matrix $\mathbf{L}$ corresponds to an $\mathcal{H}$ -IC over $\mathbb{F}_q$ if and only if*

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 1,$$

for all $K \in \mathcal{J}(\mathcal{H})$ and for all choices of $z_i \in \mathbb{F}_q^*$, $i \in K$, or, equivalently,

$$\forall i \in [m] : \mathbf{L}_{f(i)} \notin \text{span}_q(\{\mathbf{L}_j\}_{j \in \mathcal{Y}_i}).$$

Remark 3.1.6. The conditions stated in Corollary 3.1.5 and Corollary 2.2.6 are, as expected, equivalent. Indeed, for each $i \in [m]$ we have

$$\begin{aligned} \mathbf{L}_{f(i)} \notin \text{span}_q(\{\mathbf{L}_j\}_{j \in \mathcal{Y}_i}) &\iff \exists \mathbf{v}^{(i)} : \mathbf{v}^{(i)} \mathbf{L}_{f(i)}^T = 1 \text{ and } \mathbf{v}^{(i)} \mathbf{L}_{\mathcal{Y}_i}^T = \mathbf{0} \\ &\iff \exists \mathbf{v}^{(i)} : \mathbf{c}^{(i)} = \mathbf{v}^{(i)} \mathbf{L}^T \in \mathcal{C}(\mathbf{L}) \text{ satisfies } \mathbf{c}_{f(i)}^{(i)} = 1 \text{ and} \\ &\quad \mathbf{c}_{\mathcal{Y}_i}^{(i)} = \mathbf{0} \\ &\iff \exists \mathbf{u}^{(i)} \in \mathbb{F}_q^n : \mathbf{u}^{(i)} \triangleleft \mathcal{X}_i \text{ and } \mathbf{c}^{(i)} = \mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in \mathcal{C}(\mathbf{L}). \end{aligned}$$

Example 3.1.7. Let $q = 2$, $m = n = 3$, and $f(i) = i$ for $i \in [3]$. Suppose

$\mathcal{X}_1 = \{2, 3\}$, $\mathcal{X}_2 = \{1, 3\}$, and $\mathcal{X}_3 = \{1, 2\}$. Let

$$\mathbf{L} = \begin{pmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix}.$$

Note that the rows of $\mathbf{L}$ generate a $[4, 3, 1]_2$ code, which has minimum distance one. Nevertheless, the index code based on $\mathbf{L}$ can still correct one error. Indeed, let $\mathcal{H} = \mathcal{H}(3, 3, \mathcal{X}, f)$, we have

$$\mathcal{I}(2, \mathcal{H}) = \{100, 010, 001\}.$$

Since each row of $\mathbf{L}$ has weight at least three, it follows that $\mathbf{wt}(\mathbf{zL}) \geq 3$ for all $\mathbf{z} \in \mathcal{I}(2, \mathcal{H})$. By Lemma 3.1.3, $\mathbf{L}$ corresponds to a $(1, \mathcal{H})$ -ECIC over $\mathbb{F}_2$.

In fact, for this instance, even a simpler index code of length three, based on

$$\mathbf{L}' = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix},$$

is a $(1, \mathcal{H})$ -ECIC over $\mathbb{F}_2$.

Example 3.1.8. Assume that $m = n$ and $f(i) = i$ for all $i \in [m]$. Furthermore, suppose that $\mathcal{X}_i = \emptyset$ for all $i \in [m]$ (that is, there is no side information available to the receivers). Let $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$. Then, $\mathcal{I}(q, \mathcal{H}) = \mathbb{F}_q^n \setminus \{\mathbf{0}\}$. Hence, by Lemma 3.1.3, the $n \times N$ matrix $\mathbf{L}$ corresponding to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ (for some integer $\delta \geq 0$) if and only if it is a generating matrix of an $[N, n, \geq 2\delta + 1]_q$ linear code. Thus, under these settings, the problem of designing an optimal linear ECIC

is reduced to the problem of constructing an optimal classical linear error-correcting code.

Observe however, that for general $\mathcal{X}$, changing the order of rows in $\mathbf{L}$ can lead to ECIC's with different error-correcting capabilities. Therefore, the problem of designing an optimal linear ECIC is essentially the problem of finding the matrix $\mathbf{L}$ corresponding to that code. However, the minimum distance of the code generated by the rows of $\mathbf{L}$ is not necessary a valid indicator for goodness of an ECIC. Sometimes, as Example 3.1.7 shows, matrix $\mathbf{L}$ with redundant rows yields a good ECIC.

3.2 The α -Bound and the κ -Bound

Let $(m, n, \mathcal{X}, f)$ be an ICSI instance, and let $\mathcal{H}$ be the corresponding side information hypergraph. Next, we introduce the following definitions for the hypergraph $\mathcal{H}$.

Definition 3.2.1. A subset H of $[n]$ is called a *generalized independent set* in $\mathcal{H}$ if every nonempty subset K of H belongs to $\mathcal{J}(\mathcal{H})$.

Definition 3.2.2. A generalized independent set of maximum size in $\mathcal{H}$ is called a *maximum generalized independent set*. The size of a maximum generalized independent set in $\mathcal{H}$ is called the *generalized independence number*, and denoted by $\alpha(\mathcal{H})$.

In the following two lemmas, we consider an ICSI instance $(m, n, \mathcal{X}, f)$ in which $m = n$ and $f(i) = i$ for all $i \in [n]$. Let $\mathcal{D}$ be the underlying digraph of the side information hypergraph $\mathcal{H}$, that is, $\mathcal{D}$ is the side information digraph of this ICSI instance (see Remark 2.1.2). A *maximum acyclic induced subgraph* of $\mathcal{D}$ is an acyclic

subgraph which is induced by a set of maximum number of vertices. We use $\alpha(\mathcal{D})$ to denote the order of a maximum acyclic induced subgraph of $\mathcal{D}$. The two lemmas below show that $\alpha(\mathcal{H})$ is indeed a generalization of $\alpha(\mathcal{D})$ and $\alpha(\mathcal{G})$.

Lemma 3.2.3. $\alpha(\mathcal{H}) = \alpha(\mathcal{D})$.

Proof. It suffices to show that a generalized independent set in $\mathcal{H}$ induces an acyclic subgraph of $\mathcal{D}$ and vice versa.

Let H be a generalized independent set in $\mathcal{H}$. If $|H| = 1$, then obviously H induces an acyclic subgraph of $\mathcal{D}$. Assume that $|H| \geq 2$. Take any subset K of H where $|K| \geq 2$. Since $K \in \mathcal{J}(\mathcal{H})$, there exists some $i \in [n]$ such that $i = f(i) \in K$ and $j \in \mathcal{Y}_i$ for all $j \in K \setminus \{i\}$. Recall that $\mathcal{Y}_i = [n] \setminus (\{i\} \cup \mathcal{X}_i)$. As (i, j) 's for $j \in \mathcal{X}_i$ are the only arcs of $\mathcal{D}$ that originate from i , we deduce that $\mathcal{D}$ has no arc of the form (i, j) for every $j \in K \setminus \{i\}$. Therefore, K cannot form a circuit in $\mathcal{D}$. Since this conclusion holds for any such K , we conclude that H induces an acyclic subgraph of $\mathcal{D}$.

Conversely, suppose that H induces an acyclic subgraph of $\mathcal{D}$. We aim to show that H is a generalized independent set in $\mathcal{H}$. Take an arbitrary nonempty subset K of H and let $\mathcal{D}_K$ be the K -induced subgraph of $\mathcal{D}$. If $K = \{i\}$ for some i , then obviously $K \in \mathcal{J}(\mathcal{H})$. Suppose $|K| \geq 2$. Since $\mathcal{D}_K$ is also acyclic, there exists some vertex $i \in K$ with no out-going arcs in $\mathcal{D}_K$. In other words, $(i, j) \notin \mathcal{E}(\mathcal{D}_K)$ for every $j \in K \setminus \{i\}$. Therefore, $j \in \mathcal{Y}_i$ for every $j \in K \setminus \{i\}$, which implies that $K \setminus \{i\} \subseteq \mathcal{Y}_i$. Hence $K = \{f(i) = i\} \cup (K \setminus \{i\}) \in \mathcal{J}(\mathcal{H})$. Since K is an arbitrary nonempty subset of H , we conclude that H is a generalized independent set of $\mathcal{H}$. $\blacksquare$

Lemma 3.2.4. If $\mathcal{D}$ is symmetric, then $\alpha(\mathcal{H}) = \alpha(\mathcal{G})$.

Proof. It suffices to show that if $\mathcal{D}$ is symmetric, then the set of generalized inde-

pendent sets of $\mathcal{H}$ and the set of independent sets of $\mathcal{G}$ coincide.

Let H be a generalized independent set in $\mathcal{H}$. If $|\mathcal{H}| = 1$, then obviously H is an independent set in $\mathcal{G}$. Assume that $|H| \geq 2$. For any pair of vertices i, j in H , the set $\{i, j\}$ belongs to $\mathcal{J}(\mathcal{H})$. By definition of $\mathcal{J}(\mathcal{H})$, either $(i, j) \notin \mathcal{E}(\mathcal{D})$ or $(j, i) \notin \mathcal{E}(\mathcal{D})$. Since $\mathcal{D}$ is symmetric, there are no arcs between i and j , in neither directions. Therefore, H is an independent set in $\mathcal{G}$.

Conversely, let H be an independent set in $\mathcal{G}$. For each $i \in H$, since there are no arcs in $\mathcal{D}$ from i to any of the other vertices in H , we deduce that $H \setminus \{i\} \subseteq \mathcal{Y}_i$. Due to (3.1), every subset of H that contains i belongs to $\mathcal{J}(\mathcal{H})$. This holds for an arbitrary $i \in H$. Therefore, every nonempty subset of H belongs to $\mathcal{J}(\mathcal{H})$. Hence H is a generalized independent set of $\mathcal{H}$. ■

Next, we present a lower bound on the length of a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$. We call this bound the α -bound.

Theorem 3.2.5 (α -bound). *The length of an optimal linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$\mathcal{N}_q[\delta, \mathcal{H}] \geq N_q[\alpha(\mathcal{H}), 2\delta + 1].$$

Proof. Consider an $n \times N$ matrix $\mathbf{L}$, which corresponds to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$. Let $H = \{i_1, i_2, \dots, i_{\alpha(\mathcal{H})}\}$ be a maximum generalized independent set in $\mathcal{H}$. Then, every subset $K \subseteq H$ satisfies $K \in \mathcal{J}(\mathcal{H})$. Therefore,

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 2\delta + 1$$

for all $K \subseteq H$, where $K \neq \emptyset$, and for all choices of $z_i \in \mathbb{F}_q^*$, $i \in K$. Hence, the $\alpha(\mathcal{H})$ rows of $\mathbf{L}$, namely $\mathbf{L}_{i_1}, \mathbf{L}_{i_2}, \dots, \mathbf{L}_{i_{\alpha(\mathcal{H})}}$, form a generator matrix of an

$[N, \alpha(\mathcal{H}), 2\delta + 1]_q$ code. Therefore,

$$N \geq N_q[\alpha(\mathcal{H}), 2\delta + 1]. \quad \blacksquare$$

Example 3.2.6. Let $q = 2$, $m = n = 5$, $f(i) = i$ for all $i \in [m]$, and $\delta = 2$. Assume

$$\mathcal{X}_1 = \{2, 3, 4\}, \mathcal{X}_2 = \{3, 4, 5\}, \mathcal{X}_3 = \{4, 5, 1\}, \mathcal{X}_4 = \{5, 1, 2\}, \mathcal{X}_5 = \{1, 2, 3\}.$$

Let $\mathcal{H} = \mathcal{H}(5, 5, \mathcal{X}, f)$. Then

$$\mathcal{J}(\mathcal{H}) = \left\{ \{1\}, \{1, 5\}, \{2\}, \{2, 1\}, \{3\}, \{3, 2\}, \{4\}, \{4, 3\}, \{5\}, \{5, 4\} \right\}.$$

It is easy to check that $\alpha(\mathcal{H}) = 2$. Therefore, Theorem 3.2.5 implies that

$$\mathcal{N}_2[2, \mathcal{H}] \geq N_2[2, 5] = 8.$$

The last equality can be verified by [31].

Remark 3.2.7. In [3], when $m = n$ and $f(i) = i$ for all $i \in [n]$, the quantity $\alpha(\mathcal{D}) = \alpha(\mathcal{H})$ was shown to be a lower bound for the length of a linear (non-error-correcting) index code. However, the α -bound in Theorem 3.2.5 does not follow from the results in [3]. The reason is that the concatenation of an optimal linear error-correcting code with an optimal linear (non-error-correcting) index code might fail to produce an optimal linear ECIC. This fact is illustrated later in Example 3.2.10.

The following proposition is based on the fact that concatenation of a δ -error-correcting code with an optimal linear (non-error-correcting) $\mathcal{H}$ -IC over $\mathbb{F}_q$ yields a linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$.

Proposition 3.2.8 (κ -bound). *The length of an optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$\mathcal{N}_q[\delta, \mathcal{H}] \leq N_q[\kappa_q(\mathcal{H}), 2\delta + 1].$$

Proof. Let $\mathbf{G}$, which is an $n \times \kappa_q(\mathcal{H})$ matrix, correspond to an optimal $\mathcal{H}$ -IC over $\mathbb{F}_q$. Denote

$$\mathbf{y} = \mathbf{x}\mathbf{G} \in \mathbb{F}_q^{\kappa_q(\mathcal{H})}.$$

Let $\mathbf{M}$ be a generator matrix of an optimal $[N, \kappa_q(\mathcal{H}), 2\delta + 1]_q$ code $\mathcal{C}$, where

$$N = N_q[\kappa_q(\mathcal{H}), 2\delta + 1].$$

Consider a scheme where S broadcasts the vector $\mathbf{y}\mathbf{M} \in \mathbb{F}_q^N$. If less than δ errors occur, then each receiver R_i is able to recover $\mathbf{y}$ by using $\mathcal{C}$. Hence each R_i is able to recover $x_{f(i)}$. Therefore, for the index code based on $\mathbf{L}$,

$$\mathbf{L} = \mathbf{G}\mathbf{M},$$

each receiver R_i is capable to recover $x_{f(i)}$ if the number of errors is at most δ . The length of the corresponding ECIC is $N = N_q[\kappa_q(\mathcal{H}), 2\delta + 1]$. Therefore,

$$\mathcal{N}_q[\delta, \mathcal{H}] \leq N_q[\kappa_q(\mathcal{H}), 2\delta + 1]. \quad \blacksquare$$

By combining the results in Theorem 3.2.5 and in Proposition 3.2.8, we obtain the following corollary.

Corollary 3.2.9. *The length of an optimal linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$N_q[\alpha(\mathcal{H}), 2\delta + 1] \leq \mathcal{N}_q[\delta, \mathcal{H}] \leq N_q[\kappa_q(\mathcal{H}), 2\delta + 1].$$

It is shown in the example below that the inequalities in Corollary 3.2.9 can be strict. In particular, it follows that mere application of an error-correcting code on top of an index code may fail to provide us with an optimal linear ECIC. This fact motivates the study of ECIC's.

Example 3.2.10. Let $q = 2$, $m = n = 5$, $\delta = 2$, and $f(i) = i$ for all $i \in [m]$. Assume

$$\mathcal{X}_1 = \{2, 5\}, \mathcal{X}_2 = \{1, 3\}, \mathcal{X}_3 = \{2, 4\}, \mathcal{X}_4 = \{3, 5\}, \mathcal{X}_5 = \{1, 4\}.$$

Let $\mathcal{H} = \mathcal{H}(5, 5, \mathcal{X}, f)$. Then we have

$$\begin{aligned} \mathcal{J}(\mathcal{H}) = \big\{ & \{1\}, \{1, 3\}, \{1, 4\}, \{1, 3, 4\}, \\ & \{2\}, \{2, 4\}, \{2, 5\}, \{2, 4, 5\}, \\ & \{3\}, \{1, 3\}, \{3, 5\}, \{1, 3, 5\}, \\ & \{4\}, \{1, 4\}, \{2, 4\}, \{1, 2, 4\}, \\ & \{5\}, \{2, 5\}, \{3, 5\}, \{2, 3, 5\} \big\}. \end{aligned}$$

The side information graph $\mathcal{G}$ of this instance is a pentagon. It is easy to verify that $\alpha(\mathcal{H}) = \alpha(\mathcal{G}) = 2$. It follows from [4, Theorem 9] that $\kappa_2(\mathcal{H}) = \text{minrk}_2(\mathcal{G}) = 3$. From [31] we have

$$N_2[2, 5] = 8 \quad \text{and} \quad N_2[3, 5] = 10.$$

Due to Corollary 3.2.9, we have

$$8 \leq \mathcal{N}_2[2, \mathcal{H}] \leq 10.$$

Using a computer search, we obtain that $\mathcal{N}_2[2, \mathcal{H}] = 9$, and the corresponding optimal scheme is based on

$$\mathbf{L} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

It is technical to verify that for all $K \in \mathcal{J}(\mathcal{H})$,

$$\text{wt} \left(\sum_{i \in K} \mathbf{L}_i \right) \geq 5.$$

Therefore by Lemma 3.1.3, for the index code based on $\mathbf{L}$, each receiver R_i is able to recover x_i , if the number of errors is less than or equal to 2. Observe that the length of the ECIC corresponding to $\mathbf{L}$ lies strictly between the α -bound and the κ -bound.

The following bounds on the minranks of graphs are known to hold [34]. Note that the second inequality in the theorem below is also called the *clique-cover bound* for minranks, as $\chi(\overline{\mathcal{G}})$ is nothing other than the clique cover number of $\mathcal{G}$.

Theorem 3.2.11 (Sandwich Theorem). *Let $\chi(\overline{\mathcal{G}})$ denote the chromatic number of*

the complement of the graph $\mathcal{G}$. Then,

$$\alpha(\mathcal{G}) \leq \text{minrk}_q(\mathcal{G}) \leq \chi(\overline{\mathcal{G}}).$$

Consider the basic case when $m = n$ and $f(i) = i$ for all $i \in [m]$, and the side information digraph $\mathcal{D}$ is symmetric. We have $\alpha(\mathcal{H}) = \alpha(\mathcal{G})$ (Lemma 3.2.4) and $\kappa_q(\mathcal{H}) = \text{minrk}_q(\mathcal{G})$ (see the discussion that follows Corollary 2.2.8), where $\mathcal{G}$ is the side information graph (the underlying graph of $\mathcal{H}$). If $\mathcal{G}$ satisfies $\alpha(\mathcal{G}) = \chi(\overline{\mathcal{G}})$, then

$$\alpha(\mathcal{H}) = \alpha(\mathcal{G}) = \text{minrk}_q(\mathcal{G}) = \kappa_q(\mathcal{H}).$$

Hence from Corollary 3.2.9 we have

$$\mathcal{N}_q[\delta, \mathcal{H}] = N_q[\alpha(\mathcal{H}), 2\delta + 1] = N_q[\kappa_q(\mathcal{H}), 2\delta + 1],$$

for all q , and the corresponding bounds in Corollary 3.2.9 are tight.

Definition 3.2.12. A graph $\mathcal{G}$ is called perfect if for every induced subgraph $\mathcal{G}'$ of $\mathcal{G}$, it holds that $\alpha(\mathcal{G}') = \chi(\overline{\mathcal{G}'})$.

Perfect graphs include families of graphs such as trees, bipartite graphs, interval graphs, and chordal graphs. If the side information graph $\mathcal{G}$ is perfect, then the bounds in Corollary 3.2.9 are tight. For the full characterization of perfect graphs, the reader can refer to [18].

3.3 The Singleton Bound

The following bound generalizes the Singleton bound for classical linear error-correcting codes to linear ECIC's.

Theorem 3.3.1 (Singleton bound). *The length of an optimal linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$\mathcal{N}_q[\delta, \mathcal{H}] \geq \kappa_q(\mathcal{H}) + 2\delta.$$

Proof. Let $\mathbf{L}$ be the $n \times \mathcal{N}_q[\delta, \mathcal{H}]$ matrix corresponding to some optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$. Let $\mathbf{L}'$ be the matrix obtained by deleting any 2δ columns from $\mathbf{L}$.

By Lemma 3.1.3, $\mathbf{L}$ satisfies

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 2\delta + 1,$$

for all $K \in \mathcal{J}(\mathcal{H})$ and all choices of $z_i \in \mathbb{F}_q^*$, $i \in K$. We deduce that the rows of $\mathbf{L}'$ also satisfy

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}'_i \right) \geq 1,$$

for all such K and z_i 's. By Corollary 3.1.5, $\mathbf{L}'$ corresponds to a linear $\mathcal{H}$ -IC over $\mathbb{F}_q$. Therefore, by Corollary 2.2.8, $\mathbf{L}'$ has at least $\kappa_q(\mathcal{H})$ columns. We deduce that

$$\mathcal{N}_q[\delta, \mathcal{H}] - 2\delta \geq \kappa_q(\mathcal{H}),$$

which concludes the proof. ■

The following corollary from Proposition 3.2.8 and Theorem 3.3.1 demonstrates that, for sufficiently large alphabets, a concatenation of a classical MDS error-correcting code with an optimal (non-error-correcting) index code yields an optimal

linear ECIC. On the contrary, as it is illustrated in Example 3.2.10, this does not hold for the index coding schemes over small alphabets.

Corollary 3.3.2 (MDS error-correcting index code). *For $q \geq \kappa_q(\mathcal{H}) + 2\delta - 1$,*

$$\mathcal{N}_q[\delta, \mathcal{H}] = \kappa_q(\mathcal{H}) + 2\delta. \quad (3.6)$$

Proof. From Theorem 3.3.1, we have

$$\mathcal{N}_q[\delta, \mathcal{H}] \geq \kappa_q(\mathcal{H}) + 2\delta.$$

On the other hand, from Proposition 3.2.8,

$$\mathcal{N}_q[\delta, \mathcal{H}] \leq N_q[\kappa_q(\mathcal{H}), 2\delta + 1] = \kappa_q(\mathcal{H}) + 2\delta,$$

for $q \geq \kappa_q(\mathcal{H}) + 2\delta - 1$ (by taking doubly-extended Reed-Solomon codes). Therefore, for these q , (3.6) holds. ■

Remark 3.3.3. Let $q = 2$, $m = n = 2\ell + 1$ ($\ell \geq 2$), and $f(i) = i$ for all $i \in [n]$. Let $\mathcal{X}_1 = \{2, n\}$ and $\mathcal{X}_n = \{1, n - 1\}$. For $2 \leq i \leq n$, let $\mathcal{X}_i = \{i - 1, i + 1\}$. Let $\mathcal{H} = \mathcal{H}(n, n, \mathcal{X}, f)$. Notice that the side information graph $\mathcal{G}$ is the odd cycle of length n . Therefore, $\alpha(\mathcal{H}) = \alpha(\mathcal{G}) = \ell$. From [4], $\kappa_2(\mathcal{H}) = \text{minrk}_2(\mathcal{G}) = \ell + 1$. From the α -bound,

$$\mathcal{N}_2[\delta, \mathcal{H}] \geq N_2[\ell, 2\delta + 1].$$

From the Singleton bound,

$$\mathcal{N}_2[\delta, \mathcal{H}] \geq (\ell + 1) + 2\delta.$$

As there are no nontrivial binary MDS codes, we have

$$N_2[\ell, 2\delta + 1] \geq \ell + 2\delta + 1,$$

for all $\delta > 0$. Therefore, in this case the α -bound is at least as good as the Singleton bound.

3.4 Random codes

In this section we prove an implicit upper bound on the optimal length of the ECIC's. The proof is based on constructing a random ECIC and analyzing its parameters.

Theorem 3.4.1. *Let $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$ describe an ICSI instance. Then there exists a linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ of length N if*

$$\sum_{i \in [m]} q^{n-|\mathcal{X}_i|-1} < \frac{q^N}{V_q(N, 2\delta)},$$

where

$$V_q(N, 2\delta) = \sum_{\ell=0}^{2\delta} \binom{N}{\ell} (q-1)^\ell$$

is the volume of the q -ary sphere in $\mathbb{F}_q^N$.

Proof. We construct a random $n \times N$ matrix $\mathbf{L}$ over $\mathbb{F}_q$, row by row. Each row is selected independently of other rows, uniformly over $\mathbb{F}_q^N$. Define vector spaces

$$\mathbf{M}_i \triangleq \text{span}_q(\{\mathbf{L}_j : j \in \mathcal{Y}_i\})$$

for all $i \in [m]$. We also define the following events:

$$\forall i \in [m] : \text{Event } E_i \triangleq \{d(\mathbf{L}_{f(i)}, \mathbf{M}_i) < 2\delta + 1\},$$

and

$$\text{Event } E_{Fail} \triangleq \{\mathbf{L} \text{ does not correspond to a } (\delta, \mathcal{H})\text{-ECIC over } \mathbb{F}_q\}.$$

The event E_i represents the situation when the receiver R_i cannot recover $x_{f(i)}$. Then, by Corollary 3.1.4, the event E_{Fail} is equivalent to $\cup_{i \in [m]} E_i$. Therefore,

$$\Pr(E_{Fail}) = \Pr\left(\bigcup_{i \in [m]} E_i\right) \leq \sum_{i \in [m]} \Pr(E_i). \quad (3.7)$$

For a particular event E_i , $i \in [m]$,

$$\Pr(E_i) \leq \frac{q^{|\mathcal{Y}_i|} V_q(N, 2\delta)}{q^N}. \quad (3.8)$$

There exists a matrix $\mathbf{L}$ corresponding to a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ if $\Pr(E_{Fail}) < 1$. It is enough to require that the right-hand side of (3.7) is smaller than 1. By plugging in the expression in (3.8), we obtain a sufficient condition on the existence of a linear $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$:

$$\frac{V_q(N, 2\delta)}{q^N} \sum_{i \in [m]} q^{|\mathcal{Y}_i|} < 1. \quad \blacksquare$$

Remark 3.4.2. The bound in Theorem 3.4.1 does not take into account the structure of the sets $\mathcal{X}_i$'s, other than their cardinalities. Therefore, this bound generally is weaker than the κ -bound. On the other hand, for a particular ICSI instance, it is easier to calculate this bound, while computing the κ -bound in general is a hard problem.

Remark 3.4.3. The bound in Theorem 3.4.1 implies a bound on $\kappa_q(\mathcal{H})$, which is tight for some $\mathcal{H}$. Indeed, fix $\delta = 0$. The bound implies that there exists a linear index code of length N whenever

$$\sum_{i \in [m]} q^{n-|\mathcal{X}_i|-1} < q^N. \quad (3.9)$$

Let $m = n = 2\ell + 1$ ($\ell \geq 2$), and $f(i) = i$ for all $i \in [n]$. Let $\mathcal{X}_1 = [n] \setminus \{1, 2, n\}$ and $\mathcal{X}_n = [n] \setminus \{1, n-1, n\}$. For $2 \leq i \leq n-1$, let $\mathcal{X}_i = [n] \setminus \{i-1, i, i+1\}$. Let $\mathcal{H} = \mathcal{H}(n, n, \mathcal{X}, f)$ be the corresponding side information hypergraph. The side information graph $\mathcal{G}$ is the complement of the odd cycle of length n . We have $|\mathcal{X}_i| = 2\ell - 2$ for all $i \in [n]$. Then (3.9) becomes

$$N > 2 + \log_q(2\ell + 1).$$

If $q > 2\ell + 1$ then we obtain $N \geq 3$. Observe that in this case $\kappa_q(\mathcal{H}) = \text{minrk}_q(\mathcal{G}) = 3$ (see [2, Claim A.1]), and thus the bound is tight.

3.5 Strongly Secure Linear Error-Correcting Index Codes

In this section we consider an active adversary who not only owns some side information and eavesdrops some transmissions, but also corrupts some transmissions from S . Note that a simple concatenation of an error-correcting scheme and a security scheme may not necessarily work. Indeed, applying an error-correcting code on top of a secure index code may sabotage the security of the scheme. Conversely, if error-correcting index coding is followed by security coding, the scheme is no longer resistant to errors.

3.5.1 A Lower Bound on the Length of a Strongly Secure Linear Error-Correcting Index Code

The notion of randomized index codes introduced in Subsection 2.4.1 can be naturally extended to that of δ -error-correcting randomized index codes. Then we have the following definition.

Definition 3.5.1. A randomized index code is referred to as a (μ, t) -strongly secure δ -error-correcting randomized index code if it is not only (μ, t) -strongly secure (see Definition 2.4.2) but also capable of correcting δ errors (see Definition 3.1.1).

The next theorem establishes a lower bound on the length of a (μ, t) -strongly secure δ -error-correcting linear randomized index code.

Theorem 3.5.2. *The length of a (μ, t) -strongly secure δ -error-correcting linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is at least $\kappa_q + \mu + 2\delta$.*

Proof. Let $\mathbf{L}$ correspond to a (μ, t) -strongly secure δ -error-correcting η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. Let $\mathbf{L}'$ be the matrix obtained from $\mathbf{L}$ by deleting any 2δ columns of $\mathbf{L}$. Since $\mathbf{L}$ corresponds to a δ -error-correcting index code, by Lemma 3.1.3 it satisfies

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 2\delta + 1,$$

for all $K \in \mathcal{J}(\mathcal{H})$ and all choices of nonzero $z_i \in \mathbb{F}_q$, $i \in K$. We deduce that the rows $\mathbf{L}'_i$'s of $\mathbf{L}'$ also satisfy

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}'_i \right) \geq 1.$$

By Corollary 3.1.5, $\mathbf{L}'$ still corresponds to an η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. Since less information about $\mathbf{x}$ is broadcast if $\mathbf{L}'$ is used, we deduce that $\mathbf{L}'$ also corresponds to a (μ, t) -strongly secure η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.

Therefore, by Theorem 3.2.5, $\mathbf{L}'$ has at least $\kappa_q + \mu$ columns. We then deduce that $\mathbf{L}$ has at least $\kappa_q + \mu + 2\delta$ columns. $\blacksquare$

3.5.2 A Construction of an Optimal Strongly Secure Linear Error-Correcting Index Code

In this subsection we provide a construction that produces randomized index codes attaining the bound established in Theorem 3.5.2. This construction requires q to be as large as $\kappa_q + \mu + 2\delta$.

Construction B: Let $\mathbf{L}^{(0)}$ correspond to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ of optimal length κ_q . Let $\mathbf{M}$ be a generator matrix of an $[N = \kappa_q + \mu + 2\delta, \kappa_q + \mu, 2\delta + 1]_q$ -MDS code, so that the last μ rows of $\mathbf{M}$ form a generator matrix of another MDS code. For instance, we can take a generator matrix of a Generalized Reed-Solomon code (see, for instance [39])

$$\mathbf{M} = \left(\begin{array}{cccc} 1 & 1 & \cdots & 1 \\ \alpha_1 & \alpha_2 & \cdots & \alpha_N \\ \vdots & \vdots & \cdots & \vdots \\ \alpha_1^{\kappa_q-1} & \alpha_2^{\kappa_q-1} & \cdots & \alpha_N^{\kappa_q-1} \\ \hline \alpha_1^{\kappa_q} & \alpha_2^{\kappa_q} & \cdots & \alpha_N^{\kappa_q} \\ \alpha_1^{\kappa_q+1} & \alpha_2^{\kappa_q+1} & \cdots & \alpha_N^{\kappa_q+1} \\ \vdots & \vdots & \cdots & \vdots \\ \alpha_1^{\kappa_q+\mu-1} & \alpha_2^{\kappa_q+\mu-1} & \cdots & \alpha_N^{\kappa_q+\mu-1} \end{array} \right),$$

where $\alpha_1, \alpha_2, \dots, \alpha_N$ are pairwise distinct elements in $\mathbb{F}_q$. Let $\mathbf{P}$ be the submatrix of $\mathbf{M}$ formed by the first κ_q rows, and $\mathbf{Q}$ the submatrix formed by the last μ rows

of $\mathbf{M}$. Let

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}^{(0)}\mathbf{P} \\ \mathbf{Q} \end{pmatrix}.$$

Lemma 3.5.3. *The matrix $\mathbf{L}$ in Construction B corresponds to a δ -error-correcting μ -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.*

Proof. Let $\mathbf{g} \in \mathbb{F}_q^\mu$ be a randomized vector, uniformly distributed over $\mathbb{F}_q^\mu$. The encoding process is as follows

$$\mathbf{x} \mapsto (\mathbf{x}|\mathbf{g})\mathbf{L} = \mathbf{x}\mathbf{L}^{(0)}\mathbf{P} + \mathbf{g}\mathbf{Q} = \left(\mathbf{x}\mathbf{L}^{(0)}|\mathbf{g}\right)\mathbf{M}.$$

Since $\mathbf{M}$ is a generator matrix of a δ -error-correcting code, if at most δ errors occur, then each receiver is able to recover $(\mathbf{x}\mathbf{L}^{(0)}|\mathbf{g})$. Therefore, for each $i \in [m]$, the receiver R_i can recover $\mathbf{x}\mathbf{L}^{(0)}$, and hence, can also recover $x_{f(i)}$, as $\mathbf{L}^{(0)}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$. ■

Lemma 3.5.4. *The matrix $\mathbf{L}$ in Construction A corresponds to a (μ, t) -strongly secure μ -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$.*

Proof. The proof of this lemma is the same as that of Lemma 2.4.9. ■

From Theorem 3.5.2, Lemma 3.5.3, and Lemma 3.5.4, we have the following theorem.

Theorem 3.5.5. *The length of an optimal (μ, t) -strongly secure δ -error-correcting linear η -randomized $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ is at least $\kappa_q + \mu + 2\delta$. Moreover, the code based on the matrix $\mathbf{L}$ established in Construction B achieves this bound ($q \geq \kappa_q + \mu + 2\delta$).*

3.6 Syndrome Decoding

Consider the $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ based on a matrix $\mathbf{L}$. Suppose that the receiver R_i , $i \in [m]$, receives the vector

$$\mathbf{y}^{(i)} = \mathbf{x}\mathbf{L} + \boldsymbol{\epsilon}^{(i)}, \quad (3.10)$$

where $\mathbf{x}\mathbf{L}$ is the codeword transmitted by S , and $\boldsymbol{\epsilon}^{(i)}$ is the error pattern affecting this codeword.

In the classical coding theory, the transmitted vector $\mathbf{c}$, the received vector $\mathbf{y}$, and the error pattern $\mathbf{e}$ are related by $\mathbf{y} = \mathbf{c} + \mathbf{e}$. Therefore, if $\mathbf{y}$ is known to the receiver, then there is a one-to-one correspondence between the values of unknown vectors $\mathbf{c}$ and $\mathbf{e}$. For index coding, however, this is no longer the case. The following lemma shows that, in order to recover the message $x_{f(i)}$ from $\mathbf{y}^{(i)}$ using (3.10), it is sufficient to find just one vector from a set of possible error patterns. This set is defined as follows:

$$\mathcal{L}_i(\boldsymbol{\epsilon}^{(i)}) = \left\{ \boldsymbol{\epsilon}^{(i)} + \mathbf{z} : \mathbf{z} \in \text{span}_q(\{\mathbf{L}_j\}_{j \in \mathcal{Y}_i}) \right\}.$$

We henceforth refer to the set $\mathcal{L}_i(\boldsymbol{\epsilon}^{(i)})$ as the *set of relevant error patterns*.

Lemma 3.6.1. *Assume that the receiver R_i receives $\mathbf{y}^{(i)}$.*

1. *If R_i knows the message $x_{f(i)}$ then it is able to determine the set $\mathcal{L}_i(\boldsymbol{\epsilon}^{(i)})$.*
2. *If R_i knows some vector $\hat{\boldsymbol{\epsilon}} \in \mathcal{L}_i(\boldsymbol{\epsilon}^{(i)})$ then it is able to determine $x_{f(i)}$.*

Proof.

1. From (3.10), we have

$$\mathbf{y}^{(i)} = x_{f(i)}\mathbf{L}_{f(i)} + \mathbf{x}_{\mathcal{X}_i}\mathbf{L}_{\mathcal{X}_i} + \mathbf{x}_{\mathcal{Y}_i}\mathbf{L}_{\mathcal{Y}_i} + \boldsymbol{\epsilon}^{(i)}. \quad (3.11)$$

If R_i knows $x_{f(i)}$, then it is also able to determine

$$\boldsymbol{\epsilon}^{(i)} + \mathbf{x}_{\mathcal{Y}_i}\mathbf{L}_{\mathcal{Y}_i} = \mathbf{y}^{(i)} - x_{f(i)}\mathbf{L}_{f(i)} - \mathbf{x}_{\mathcal{X}_i}\mathbf{L}_{\mathcal{X}_i} \in \mathcal{L}_i(\boldsymbol{\epsilon}^{(i)}).$$

Since R_i has knowledge of $\mathbf{L}$, it is also able to determine the whole $\mathcal{L}_i(\boldsymbol{\epsilon}^{(i)})$.

2. Suppose that R_i knows a vector

$$\hat{\boldsymbol{\epsilon}} = \boldsymbol{\epsilon}^{(i)} + \sum_{j \in \mathcal{Y}_i} z_j \mathbf{L}_j \in \mathcal{L}_i(\boldsymbol{\epsilon}^{(i)}),$$

for some $\mathbf{z} \in \mathbb{F}_q^{|\mathcal{Y}_i|}$. We show that R_i is able then to determine $x_{f(i)}$. Indeed, we rewrite (3.11) as

$$\mathbf{y}^{(i)} = x_{f(i)}\mathbf{L}_{f(i)} + \mathbf{x}_{\mathcal{X}_i}\mathbf{L}_{\mathcal{X}_i} + (\mathbf{x}_{\mathcal{Y}_i} - \mathbf{z})\mathbf{L}_{\mathcal{Y}_i} + \hat{\boldsymbol{\epsilon}}. \quad (3.12)$$

The receiver R_i can find some solution of the equation

$$\mathbf{y}^{(i)} = \hat{x}_{f(i)}\mathbf{L}_{f(i)} + \mathbf{x}_{\mathcal{X}_i}\mathbf{L}_{\mathcal{X}_i} + \hat{\mathbf{x}}_{\mathcal{Y}_i}\mathbf{L}_{\mathcal{Y}_i} + \hat{\boldsymbol{\epsilon}}, \quad (3.13)$$

with respect to the unknowns $\hat{x}_{f(i)}$ and $\hat{\mathbf{x}}_{\mathcal{Y}_i}$. Observe that (3.13) has at least one solution due to (3.12).

From (3.12) and (3.13), we deduce that

$$\mathbf{0} = (\hat{x}_{f(i)} - x_{f(i)})\mathbf{L}_{f(i)} + (\hat{\mathbf{x}}_{\mathcal{Y}_i} - \mathbf{x}_{\mathcal{Y}_i} + \mathbf{z})\mathbf{L}_{\mathcal{Y}_i}.$$

This equality implies that $\hat{x}_{f(i)} = x_{f(i)}$ (otherwise, by Corollary 3.1.4, the sum in the right-hand side will have nonzero weight). Hence, R_i is able to determine $x_{f(i)}$, as claimed. $\blacksquare$

We now describe a syndrome decoding algorithm for linear error-correcting index codes. From (3.11), we have

$$\mathbf{y}^{(i)} - \mathbf{x}_{\mathcal{X}_i} \mathbf{L}_{\mathcal{X}_i} - \boldsymbol{\epsilon}^{(i)} \in \text{span}_q\left(\{\mathbf{L}_{f(i)}\} \cup \{\mathbf{L}_j\}_{j \in \mathcal{Y}_i}\right).$$

Let $\mathcal{C}_i = \text{span}_q(\{\mathbf{L}_{f(i)}\} \cup \{\mathbf{L}_j\}_{j \in \mathcal{Y}_i})$, and let $\mathbf{H}^{(i)}$ be a parity check matrix of $\mathcal{C}_i$. We obtain that

$$\mathbf{H}^{(i)} \boldsymbol{\epsilon}^{(i)\top} = \mathbf{H}^{(i)} (\mathbf{y}^{(i)} - \mathbf{x}_{\mathcal{X}_i} \mathbf{L}_{\mathcal{X}_i})^\top. \quad (3.14)$$

Let $\boldsymbol{\beta}^{(i)}$ be a column vector defined by

$$\boldsymbol{\beta}^{(i)} = \mathbf{H}^{(i)} (\mathbf{y}^{(i)} - \mathbf{x}_{\mathcal{X}_i} \mathbf{L}_{\mathcal{X}_i})^\top. \quad (3.15)$$

Observe that each R_i is capable of determining $\boldsymbol{\beta}^{(i)}$. Then we can rewrite (3.14) as

$$\mathbf{H}^{(i)} \boldsymbol{\epsilon}^{(i)\top} = \boldsymbol{\beta}^{(i)}.$$

This leads us to the formulation of the following decoding procedure for R_i (see Fig. 3.1).

Remark 3.6.2. Gaussian elimination can be used to solve (3.17) for $\hat{x}_{f(i)}$. However, since $\mathbf{L}$ also corresponds to an $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$, there is more efficient way to do so. From Corollary 2.2.6, there exists a vector $\mathbf{u}^{(i)} \triangleleft \mathcal{X}_i$ satisfying $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} \in$

- *Input:* $\mathbf{y}^{(i)}$, $\mathbf{x}_{\mathcal{X}_i}$, $\mathbf{L}$.

- *Step 1:* Compute the syndrome

$$\boldsymbol{\beta}^{(i)} = \mathbf{H}^{(i)} (\mathbf{y}^{(i)} - \mathbf{x}_{\mathcal{X}_i} \mathbf{L}_{\mathcal{X}_i})^T.$$

- *Step 2:* Find the lowest Hamming weight solution $\hat{\boldsymbol{\epsilon}}$ of the system

$$\mathbf{H}^{(i)} \hat{\boldsymbol{\epsilon}}^T = \boldsymbol{\beta}^{(i)}. \quad (3.16)$$

- *Step 3:* Given that $\hat{\mathbf{x}}_{\mathcal{X}_i} = \mathbf{x}_{\mathcal{X}_i}$, solve the system for $\hat{x}_{f(i)}$:

$$\mathbf{y}^{(i)} = \hat{\mathbf{x}} \mathbf{L} + \hat{\boldsymbol{\epsilon}}. \quad (3.17)$$

- *Output:* $\hat{x}_{f(i)}$.

Fig. 3.1: Syndrome decoding procedure.

$\mathcal{C}(\mathbf{L})$. Hence $\mathbf{u}^{(i)} + \mathbf{e}_{f(i)} = \boldsymbol{\alpha} \mathbf{L}^T$ for some $\boldsymbol{\alpha} \in \mathbb{F}_q^N$. Therefore

$$\hat{x}_{f(i)} = \hat{\mathbf{x}} (\mathbf{u}^{(i)} + \mathbf{e}_{f(i)})^T - \hat{\mathbf{x}} \mathbf{u}^{(i)T} = \hat{\mathbf{x}} \mathbf{L} \boldsymbol{\alpha}^T - \hat{\mathbf{x}} \mathbf{u}^{(i)T} = (\mathbf{y}^{(i)} - \hat{\boldsymbol{\epsilon}}) \boldsymbol{\alpha}^T - \hat{\mathbf{x}} \mathbf{u}^{(i)T}.$$

With the knowledge of $\mathbf{L}$ and $\mathbf{x}_{\mathcal{X}_i}$, the receiver R_i can determine $\boldsymbol{\alpha}$ and $\hat{\mathbf{x}} \mathbf{u}^{(i)T}$. Therefore, it can also determine $\hat{x}_{f(i)}$. Note that (3.17) may have more than one solution $\hat{\mathbf{x}}$ with $\hat{\mathbf{x}}_{\mathcal{X}_i} = \mathbf{x}_{\mathcal{X}_i}$. However, as shown in the next theorem, if at most δ errors occur in $\mathbf{y}^{(i)}$, then it always holds that $\hat{x}_{f(i)} = x_{f(i)}$.

Theorem 3.6.3. *Let $\mathbf{y}^{(i)} = \mathbf{x} \mathbf{L} + \boldsymbol{\epsilon}^{(i)}$ be the vector received by R_i , and let $\text{wt}(\boldsymbol{\epsilon}^{(i)}) \leq \delta$. Assume that the procedure in Fig. 3.1 is applied to $(\mathbf{y}^{(i)}, \mathbf{x}_{\mathcal{X}_i}, \mathbf{L})$. Then, its output satisfies $\hat{x}_{f(i)} = x_{f(i)}$.*

Proof. By Lemma 3.6.1, it is sufficient to prove that $\hat{\epsilon} \in \mathcal{L}_i(\epsilon^{(i)})$. Indeed, since

$$\mathbf{H}^{(i)} \epsilon^{(i)\top} = \mathbf{H}^{(i)} \hat{\epsilon}^\top = \beta^{(i)},$$

we have

$$\mathbf{H}^{(i)} (\hat{\epsilon} - \epsilon^{(i)})^\top = \mathbf{0}.$$

Hence, $\hat{\epsilon} - \epsilon^{(i)} \in \mathcal{C}_i$, and therefore,

$$\hat{\epsilon} - \epsilon^{(i)} = z_{f(i)} \mathbf{L}_{f(i)} + \sum_{j \in \mathcal{Y}_i} z_j \mathbf{L}_j, \quad (3.18)$$

for some $z_{f(i)} \in \mathbb{F}_q$ and $z_j \in \mathbb{F}_q, j \in \mathcal{Y}_i$.

Since $\epsilon^{(i)}$ is a solution of (3.16), and $\text{wt}(\epsilon^{(i)}) \leq \delta$, we deduce that $\text{wt}(\hat{\epsilon}) \leq \delta$ as well. Hence,

$$\text{wt} \left(z_{f(i)} \mathbf{L}_{f(i)} + \sum_{j \in \mathcal{Y}_i} z_j \mathbf{L}_j \right) = \text{wt}(\hat{\epsilon} - \epsilon^{(i)}) \leq 2\delta.$$

Therefore, by Corollary 3.1.4, $z_{f(i)} = 0$. Hence, $\hat{\epsilon} \in \mathcal{L}_i(\epsilon^{(i)})$, as desired, and therefore $\hat{x}_{f(i)} = x_{f(i)}$. ■

Remark 3.6.4. Step 2 in Fig. 3.1 is computationally hard. Indeed, the problem of finding $\hat{\epsilon}$ over $\mathbb{F}_2$ of the lowest weight satisfying

$$\mathbf{H}^{(i)} \hat{\epsilon}^\top = \beta^{(i)}, \quad (3.19)$$

for a given binary vector $\beta^{(i)}$ corresponds to the decision problem COSET WEIGHTS, which was shown by Berlekamp *et al.* [7] to be NP-complete.

3.7 Static Error-Correcting Index Codes

3.7.1 Static Error-Correcting Index Codes

In previous sections of this chapter, we focus on designing an optimal linear δ -error-correcting index code for a *particular* ICSI instance. When the parameters m , n , $\mathcal{X}$, and f are changed, it is very likely that an error-correcting index code for the current instance can no longer be used for the new instance. Therefore, an interesting task is to design an error-correcting index code that works for a *family* of ICSI instances.

Definition 3.7.1. Let Γ be a set of instances $(m, n, \mathcal{X}, f)$ for the ICSI problem. A δ -error-correcting index code over $\mathbb{F}_q$ is said to be *static* under the set Γ if it is a δ -error-correcting $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ for all instances $(m, n, \mathcal{X}, f) \in \Gamma$.

Recall that an instance $(m, n, \mathcal{X}, f)$ can be described by the side information hypergraph $\mathcal{H}(m, n, \mathcal{X}, f)$. For a set Γ of instances $(m, n, \mathcal{X}, f)$, let

$$\mathfrak{J}(\Gamma) \triangleq \bigcup_{(m, n, \mathcal{X}, f) \in \Gamma} \mathcal{J}(\mathcal{H}(m, n, \mathcal{X}, f)), \quad (3.20)$$

where $\mathcal{J}(\mathcal{H}(m, n, \mathcal{X}, f))$ is defined as in (3.1). We also define

$$n(\Gamma) \triangleq \max \{n : (m, n, \mathcal{X}, f) \in \Gamma\}.$$

Lemma 3.7.2. *The $n(\Gamma) \times N$ matrix $\mathbf{L}$ corresponds to a δ -error-correcting index code that is static under Γ if and only if*

$$\text{wt} \left(\sum_{i \in K} z_i \mathbf{L}_i \right) \geq 2\delta + 1,$$

for all $K \in \mathfrak{J}(\Gamma)$ and for all choices of $z_i \in \mathbb{F}_q^*$, $i \in K$.

Proof. The proof is immediate from Definition 3.7.1 and Lemma 3.1.3. Note that when $\mathbf{L}$ is used for an instance $(m, n, \mathcal{X}, f) \in \Gamma$ with $n < n(\Gamma)$, then the last $n(\Gamma) - n$ rows of $\mathbf{L}$ are simply discarded. ■

One particular family of interest is $\Gamma(n, \rho)$, the family that contains all instances where each receiver owns at least $n - \rho$ messages as its side information. More formally,

$$\Gamma(n, \rho) = \{(m, n', \mathcal{X}, f) : n' \leq n \text{ and } \forall i \in [m], |\mathcal{X}_i| \geq n - \rho\}.$$

A δ -error-correcting index code that is static under $\Gamma(n, \rho)$ will provide successful communication between the sender and the receivers under the presence of at most δ errors, despite a possible change of the sets of side information, a change of the set of receivers, and a change of the demand function, as long as each receiver still possesses at least $n - \rho$ messages.

In the rest of this section, we assume that $N \geq 1$, $n \geq \rho \geq 1$ and $\delta \geq 0$.

Definition 3.7.3. An $n \times N$ matrix $\mathbf{L}$ is said to satisfy the (ρ, δ) -Property if any nontrivial linear combination of at most ρ rows of $\mathbf{L}$ has weight at least $2\delta + 1$.

Proposition 3.7.4. The $n \times N$ matrix $\mathbf{L}$ corresponds to a δ -error-correcting linear index code that is static under $\Gamma(n, \rho)$, if and only if $\mathbf{L}$ satisfies the (ρ, δ) -Property.

Proof. By Lemma 3.7.2, it suffices to show that $\mathfrak{J}(\Gamma(n, \rho))$ is the collection of all nonempty subsets of $[n]$, whose cardinalities are not greater than ρ .

Consider an instance $(m, n', \mathcal{X}, f) \in \Gamma(n, \rho)$. For all $i \in [m]$, we have $|\mathcal{X}_i| \geq n - \rho$

and $\mathcal{Y}_i = [n'] \setminus (f(i) \cup \mathcal{X}_i)$, and thus we deduce that

$$|\mathcal{Y}_i| \leq n' - 1 - (n - \rho) \leq n' - 1 - (n' - \rho) = \rho - 1.$$

Hence by (3.1), the cardinality of each set in $\mathcal{J}(\mathcal{H}(m, n', \mathcal{X}, f))$ is at most

$$1 + (\rho - 1) = \rho.$$

Therefore, due to (3.20), every set in $\mathfrak{J}(\Gamma(n, \rho))$ has at most ρ elements.

It remains to show that every nonempty subset of $[n]$ whose cardinality is at most ρ belongs to $\mathfrak{J}(\Gamma(n, \rho))$. Consider an arbitrary ρ' -subset $K = \{i_1, i_2, \dots, i_{\rho'}\}$ of $[n]$, with $1 \leq \rho' \leq \rho$. Consider an instance $(m = 1, n, \mathcal{X}, f) \in \Gamma(n, \rho)$ with $\mathcal{X}_1 = [n] \setminus K$ and $f(1) = i_1$. Since

$$\mathcal{Y}_1 = K \setminus \{i_1\},$$

we have

$$K = \{i_1\} \cup \mathcal{Y}_1 \in \mathcal{J}(\mathcal{H}(m, n, \mathcal{X}, f)) \subseteq \mathfrak{J}(\Gamma(n, \rho)).$$

The proof follows. ■

3.7.2 An Application: Weakly Resilient Functions

In this section we introduce the notion of weakly resilient functions. Hereafter, we restrict the discussion to the binary alphabet.

The concept of *binary resilient functions* was first introduced by Chor *et al.* [17] and independently by Bennet *et al.* [6].

Definition 3.7.5. A function $\mathbf{f} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^n$ is said to be *t-resilient* if $\mathbf{f}$ satisfies the property that it runs through every possible output n -tuple an equal number of

times when t arbitrary inputs are fixed and the remaining $N - t$ inputs run through all the 2^{N-t} input tuples once. Moreover, if $\mathbf{f}$ is a linear transformation then it is called a *linear t -resilient function*. We also refer to t as the *resiliency* of $\mathbf{f}$.

The applications of resilient functions can be found in fault-tolerant distributed computing, quantum cryptographic key distribution [17], privacy amplification [6] and random sequence generation for stream ciphers [13]. The equivalence between a linear error-correcting code and a resilient function was established by Chor [17].

Theorem 3.7.6 ([17]). *Let $\mathbf{L}$ be an $n \times N$ binary matrix. Then $\mathbf{L}$ is a generator matrix of a linear error-correcting code with minimum distance $d = t + 1$ if and only if $\mathbf{f}(\mathbf{z}) = \mathbf{z}\mathbf{L}^T$ is t -resilient.*

Remark 3.7.7. Vectorial boolean functions with certain properties are useful for design of stream ciphers. These properties include high resiliency and high nonlinearity (see, for instance, [13]). Nevertheless, linear resilient functions are still particularly interesting, since they can be transformed into highly nonlinear resilient functions with the same parameters. This can be achieved by a composition of the linear function with a highly nonlinear permutation (see [33, 70] for more details).

Below we introduce a definition of a ρ -weakly t -resilient function, which is a weaker version of a t -resilient function.

Definition 3.7.8. A function $\mathbf{f} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^n$ is said to be *ρ -weakly t -resilient* if $\mathbf{f}$ satisfies the property that every set of ρ coordinates of $\mathbf{f}$ runs through every possible output ρ -tuple an equal number of times when t arbitrary inputs are fixed and the remaining $N - t$ inputs run through all the 2^{N-t} input tuples once.

Remark 3.7.9. A ρ -weakly t -resilient function $\mathbf{f} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^n$ can be viewed as a collection of $\binom{n}{\rho}$ different t -resilient functions $\mathbb{F}_2^N \rightarrow \mathbb{F}_2^\rho$, each such function is obtained by taking some ρ coordinates in the image of $\mathbf{f}$. Similarly to [17], consider a scenario, in which two parties are sharing a secret key, which consists of N randomly selected bits. Suppose that at some moment t out of the N bits of the key are leaked to an adversary. By applying a t -resilient function to the current N -bit key, two parties are able to obtain a completely new and secret key of n bits, without requiring any communication or randomness generation. However, if the parties use various parts of the key for various purposes, they may require several secret ρ -bit keys rather than one secret n -bit key. In that case a ρ -weakly t -resilient function can be used. By applying a ρ -weakly t -resilient function to the current N -bit key, the parties obtain a set of $\binom{n}{\rho}$ different ρ -bit keys, each key is new and secret (however these keys might not be independent of each other). With the same number of inputs N and outputs n , a weakly resilient function may have strictly higher resiliency than a resilient function (see Remark 3.7.17).

Theorem 3.7.10. *Let $\mathbf{L}$ be an $n \times N$ binary matrix. Then $\mathbf{L}$ satisfies the (ρ, δ) -Property if and only if the function $\mathbf{f} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^n$ defined by $\mathbf{f}(\mathbf{z}) = \mathbf{z}\mathbf{L}^T$ is ρ -weakly 2δ -resilient.*

Proof.

1. Suppose that $\mathbf{L}$ satisfies the (ρ, δ) -Property. Take any ρ -subset $K \subseteq [n]$. By Definition 3.7.3, the $\rho \times N$ submatrix $\mathbf{L}_K$ of $\mathbf{L}$ is a generating matrix of the error-correcting code with the minimum distance $\geq 2\delta + 1$. By Theorem 3.7.6, the function $\mathbf{f}_K : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^\rho$ defined by $\mathbf{f}_K(\mathbf{z}) = \mathbf{z}\mathbf{L}_K^T$ is 2δ -resilient. Since K is an arbitrary ρ -subset of $[n]$, the function $\mathbf{f}$ is ρ -weakly 2δ -resilient.

2. Conversely, assume that the function $\mathbf{f}$ is ρ -weakly 2δ -resilient. Take any subset $K \subseteq [n]$ where $|K| = \rho$. Then the function $\mathbf{f}_K : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^\rho$ defined by $\mathbf{f}_K(\mathbf{z}) = \mathbf{z}\mathbf{L}_K^\top$ is 2δ -resilient. Therefore, by Theorem 3.7.6, $\mathbf{L}_K$ is a generating matrix of a linear code with minimum distance $2\delta + 1$. Since K is an arbitrary ρ -subset of $[n]$, by Proposition 3.7.4 $\mathbf{L}$ satisfies the (ρ, δ) -Property. ■

3.7.3 Bounds and Constructions

In this section we study the problem of constructing a matrix $\mathbf{L}$ satisfying the (ρ, δ) -Property. Such an $\mathbf{L}$ with the minimum number of columns is called *optimal*. It turns out that this is a special case of the problem of finding an optimal linear error-correcting index code. However, in this special case, as more structure is given to the sets of side information, the bounds and constructions based on $\alpha(\mathcal{H})$ and $\kappa_q(\mathcal{H})$ (see Section 3.2) are much simpler.

Recall that in the proof of Proposition 3.7.4 it is shown that

$$\mathfrak{J}(\Gamma(n, \rho)) = \bigcup_{i=1}^{\rho} \binom{[n]}{i},$$

the set of all nonempty subsets of $[n]$ of cardinality at most ρ . Consider an instance $(m^*, n, \mathcal{X}^*, f^*)$ satisfying

$$\mathcal{J}(\mathcal{H}^*) = \mathfrak{J}(\Gamma(n, \rho)) = \bigcup_{i=1}^{\rho} \binom{[n]}{i}, \quad (3.21)$$

where $\mathcal{H}^* = \mathcal{H}(m^*, n, \mathcal{X}^*, f^*)$ is the side information hypergraph corresponding to this instance. Such an instance can be constructed as follows. For each subset $K = \{i_1, i_2, \dots, i_{\rho'}\} \subseteq [n]$ ($1 \leq \rho' \leq \rho$), we introduce a receiver which requests the message x_{i_1} , and has side information $\{x_j : j \in [n] \setminus K\}$. It is straightforward to verify that we

indeed obtain an instance $(m^*, n, \mathcal{X}^*, f^*)$ satisfying (3.21). The problem of designing an optimal matrix $\mathbf{L}$ satisfying the (ρ, δ) -Property then becomes equivalent to the problem of finding an optimal $(\delta, \mathcal{H}^*)$ -ECIC. Thus, $\mathcal{N}_q[\mathcal{H}^*, \delta]$ is equal to the number of columns in an optimal matrix that satisfies the (ρ, δ) -Property. In other words, $\mathcal{N}_q[\mathcal{H}^*, \delta]$ is the length of an optimal δ -error-correcting linear index code that is static under $\Gamma(n, \rho)$.

The corresponding α -bound and κ -bound for $\mathcal{N}_q[\mathcal{H}^*, \delta]$ can be stated as follows.

Theorem 3.7.11. *Let ρ^* be the smallest number such that an $[n, n - \rho^*, \geq \rho + 1]_q$ code exists. Then we have*

$$N_q[\rho, 2\delta + 1] \leq \mathcal{N}_q[\delta, \mathcal{H}^*] \leq N_q[\rho^*, 2\delta + 1].$$

Proof. The first inequality follows from the α -bound and from the fact that $\alpha(\mathcal{H}^*) = \rho$, which is due to (3.21).

For the second inequality, it suffices to show that $\kappa_q(\mathcal{H}^*) = \rho^*$. By Corollary 3.1.5, a matrix $\mathbf{L}$ corresponds to an $\mathcal{H}^*$ -IC if and only if $\{\mathbf{L}_i : i \in K\}$ is linear independent for every $K \in \mathcal{J}(\mathcal{H}^*)$. Since $\mathcal{J}(\mathcal{H}^*)$ is the set of all nonempty subsets of cardinality at most ρ , this is equivalent to saying that every set of at most ρ rows of $\mathbf{L}$ is linear independent. This condition is equivalent to the condition that $\mathbf{L}^T$ is a parity check matrix of an $[n, n - N, \geq \rho + 1]_q$ code [49, Chapter 1]. Therefore, a linear $\mathcal{H}^*$ -IC of length N exists if and only if an $[n, n - N, \geq \rho + 1]_q$ code exists. Since ρ^* is the smallest number such that an $[n, n - \rho^*, \geq \rho + 1]_q$ code exists, we conclude that $\kappa_q(\mathcal{H}^*) = \rho^*$. ■

Corollary 3.7.12. *The length of an optimal δ -error-correcting linear index code*

over $\mathbb{F}_q$ that is static under $\Gamma(n, \rho)$ satisfies

$$\mathcal{N}_q[\delta, \mathcal{H}^*] \geq \rho^* + 2\delta,$$

where ρ^* is the smallest number such that an $[n, n - \rho^*, \geq \rho + 1]_q$ code exists.

Proof. This is a straightforward corollary of Theorem 3.3.1 (the Singleton bound) and Theorem 3.7.11. ■

Remark 3.7.13. We can see that the three bounds on the length of an optimal linear error-correcting index code that is static under $\Gamma(n, \rho)$ established above can be computed explicitly by solely using the results from classical error-correcting codes. By contrast, the original bounds for general ECIC involve the computation of $\alpha(\mathcal{H})$ and $\kappa_q(\mathcal{H})$, which are hard problems.

Corollary 3.7.14. For $q \geq \max\{n - 1, \rho + 2\delta - 1\}$, the length of an optimal δ -error-correcting linear index code over $\mathbb{F}_q$ that is static under $\Gamma(n, \rho)$ is $\rho + 2\delta$.

Proof. When $q \geq n - 1$, an (optimal) $[n, n - \rho, \rho + 1]_q$ code exists (one may choose an extended Reed-Solomon code [49, Chapter 11]). Therefore $\rho^* = \rho$. Following the lines of the proof of Theorem 3.7.11, there exists a δ -error-correcting index code of length $N_q[\rho, 2\delta + 1]$, which is static under $\Gamma(n, \rho)$. As $q \geq \rho + 2\delta - 1$, we have

$$N_q[\rho, 2\delta + 1] = \rho + 2\delta$$

(for example, by taking an extended RS code). Due to Corollary 3.7.12, this static error-correcting index code is optimal. ■

Remark 3.7.15. We observe from the proof of Theorem 3.7.11 that the problem of constructing an optimal linear (non-error-correcting) index code which is static under $\Gamma(n, \rho)$, is, in fact, equivalent to the problem of constructing a parity check matrix of a classical linear error-correcting code.

Example 3.7.16. Let $n = 20$, $\rho = 10$, $\delta = 1$ and $q = 2$. From [31], an optimal binary linear code of length 20 and minimum distance 11 has dimension 3. We deduce that $\rho^* = 17$. We also have $N_2[17, 3] = 22$. Theorem 3.7.11 implies the existence of a one-error-correcting index code of length 22 which can be used for any instance in which each receiver owns at least 10 out of (at most) 20 messages as side information. It also implies that the length of any such static error-correcting index code is at least $N_2[10, 3] = 14$. Corollary 3.7.12 provides a better lower bound, which is $19 = 17 + 2$.

Remark 3.7.17. Below we show that with the same number of inputs N and outputs n , a weakly resilient function may have strictly higher resiliency t . From Example 3.7.16, there exists a linear vectorial Boolean function $\mathbf{f} : (\mathbb{F}_2)^{22} \rightarrow (\mathbb{F}_2)^{20}$ which is 10-weakly 2-resilient. According to [31], an optimal linear $[22, 20]_2$ code has minimum distance $d = 2$. Hence, due to Theorem 3.7.6, the resiliency of any linear vectorial Boolean function $\mathbf{g} : (\mathbb{F}_2)^{22} \rightarrow (\mathbb{F}_2)^{20}$ cannot exceed one.

The problem of constructing an $n \times N$ matrix $\mathbf{L}$ that satisfies the (ρ, δ) -Property is a natural generalization of the problem of constructing the parity check matrix $\mathbf{H}$ of a linear $[n, k, d \geq \rho + 1]_q$ code. Indeed, $\mathbf{H}$ is a parity check matrix of an $[n, k, d \geq \rho + 1]_q$ code if and only if every set of ρ columns of $\mathbf{H}$ is linearly independent. Equivalently, any nontrivial linear combination of at most ρ columns of $\mathbf{H}$ has weight at least one. For comparison, $\mathbf{L}$ satisfies the (ρ, δ) -Property if and only if

any nontrivial linear combination of at most ρ columns of $\mathbf{L}^T$ has weight at least $2\delta + 1$.

Some classical methods for deriving bounds on the parameters of error-correcting codes can be generalized to the case of linear static error-correcting index codes. Below we present a Gilbert-Varshamov-like bound.

Theorem 3.7.18. *If*

$$\sum_{i=1}^{\rho-1} \binom{n-1}{i} (q-1)^i \sum_{j=0}^{2\delta} \binom{N}{j} (q-1)^j < q^N,$$

then there exists an $n \times N$ matrix $\mathbf{L}$ that satisfies the (ρ, δ) -Property.

Proof. We build up the set $\mathcal{R}$ of rows of $\mathbf{L}$ one by one. The first row can be any vector in $\mathbb{F}_q^N$ of weight at least $2\delta + 1$. Now suppose we have chosen r rows so that no nontrivial linear combination of at most ρ among these r rows have weight less than $2\delta + 1$. Recall that

$$V_q(N, 2\delta) = \sum_{\ell=0}^{2\delta} \binom{N}{\ell} (q-1)^\ell$$

denotes the volume of the q -ary sphere in $\mathbb{F}_q^N$. There are at most

$$V_q(N, 2\delta) \sum_{i=0}^{\rho-1} \binom{r}{i} (q-1)^i$$

vectors that are at distance less than $2\delta + 1$ from any linear combination of at most $\rho - 1$ among r chosen rows (this includes vectors at distance less than $2\delta + 1$ from $\mathbf{0}$). If this quantity is smaller than q^N , then we can add another row to the set $\mathcal{R}$ so that no nontrivial linear combination of at most ρ rows in $\mathcal{R}$ has weight less than $2\delta + 1$. We repeat this process until $r = n$. ■

Remark 3.7.19. If we apply Theorem 3.4.1 to the instance $(m^*, n, \mathcal{X}^*, f^*)$ defined in the beginning of this section, then we obtain a bound, which is somewhat weaker than its counterpart in Theorem 3.7.18, namely the $n \times N$ matrix $\mathbf{L}$ as above exists if

$$\sum_{i=1}^{\rho} \binom{n}{i} q^{i-1} < \frac{q^N}{V_q(N, 2\delta)}.$$

3.8 Nonlinear Error-Correcting Index Codes

We now turn our attention to nonlinear ECIC's. Analogous results for bounds on the length of optimal nonlinear ECIC's can be established.

Definition 3.8.1. Let $\mathfrak{E}$ be a nonlinear ECIC over $\mathbb{F}_q$. The *size* of $\mathfrak{E}$ is defined to be

$$\text{size}(\mathfrak{E}) = |\{\mathfrak{E}(\mathbf{x}) : \mathbf{x} \in \mathbb{F}_q^n\}|.$$

Definition 3.8.2. Let $(m, n, \mathcal{X}, f)$ be an ICSI instance and $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$.

1. The two vectors $\mathbf{x}$ and $\mathbf{x}'$ in $\mathbb{F}_q^n$ are called *confusable* if there exists an $i \in [n]$ such that $x_{f(i)} \neq x'_{f(i)}$ and yet $x_j = x'_j$ for all $j \in \mathcal{X}_i$.
2. The *confusion graph* of $\mathcal{H} = \mathcal{H}(m, n, \mathcal{X}, f)$, denoted $\mathfrak{C}(\mathcal{H})$, is the graph on the vertex set $\mathbb{F}_q^n$, where two vertices $\mathbf{x}$ and $\mathbf{x}'$ are adjacent if and only if they are confusable.

Lemma 3.8.3. Let $(m, n, \mathcal{X}, f)$ be an ICSI instance described by the side information hypergraph $\mathcal{H}$. Then $\mathfrak{E} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^N$ is a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ if and only if $d(\mathfrak{E}(\mathbf{x}), \mathfrak{E}(\mathbf{x}')) \geq 2\delta + 1$ whenever $\mathbf{x}$ and $\mathbf{x}'$ are adjacent in $\mathfrak{C}(\mathcal{H})$.

Proof. Suppose that $d(\mathfrak{E}(\mathbf{x}), \mathfrak{E}(\mathbf{x}')) \geq 2\delta + 1$ for every pair of adjacent vertices $\mathbf{x}$ and $\mathbf{x}'$ of $\mathfrak{C}(\mathcal{H})$. In other words, suppose that confusable vectors are encoded to

codewords of distance at least $2\delta + 1$ from each other. Then a nearest neighbor decoding will always be able to recover the transmitted vector of messages if the number of erroneous symbols in the received vector is at most δ . Hence $\mathfrak{C}$ is a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$.

Conversely, suppose that $\mathfrak{C}$ is a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$. Consider a pair of confusable vectors $\mathbf{x}$ and $\mathbf{x}'$. If $d(\mathfrak{C}(\mathbf{x}), \mathfrak{C}(\mathbf{x}')) \leq 2\delta$, there exists a vector $\mathbf{y} \in \mathbb{F}_q^N$ such that

$$d(\mathfrak{C}(\mathbf{x}), \mathbf{y}) \leq \delta,$$

and

$$d(\mathfrak{C}(\mathbf{x}'), \mathbf{y}) \leq \delta,$$

If $\mathbf{y}$ is received, since $\mathbf{x}_{\mathcal{X}_i} = \mathbf{x}'_{\mathcal{X}_i}$ and $x_{f(i)} \neq x'_{f(i)}$, it is impossible for the receiver to determine its desired message. This contradicts our assumption that $\mathfrak{C}$ is a $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$. ■

Lemma 3.8.4. *A subset $H \subseteq [n]$ is a generalized independent set of $\mathcal{H}$ if and only if for all nonempty subsets $K \subseteq H$, there exists some $i \in [m]$ such that $f(i) \in K$ and $\mathcal{X}_i \cap K = \emptyset$.*

Proof. By definition, H is a generalized independent set of $\mathcal{H}$ if and only if every nonempty subset K of H belongs to

$$\mathcal{J}(\mathcal{H}) = \bigcup_{i \in [m]} \{\{f(i)\} \cup Y_i : Y_i \subset \mathcal{Y}_i\}.$$

Equivalently, for every nonempty subset K of H , there exists $i \in [m]$ such that $f(i) \in K$ and $K \setminus \{i\} \subseteq \mathcal{Y}_i$. In other words, for such a K , there exists $i \in [m]$ such that $f(i) \in K$ and $\mathcal{X}_i \cap K = \emptyset$. ■

The following lemma generalizes the second statement in [3, Lemma 19].

Lemma 3.8.5. *If H is a generalized independent set of $\mathcal{H}$, then the set of vertices*

$$V(H) = \{\mathbf{x} \in \mathbb{F}_q^n : \mathbf{x} \triangleleft H\}$$

forms a clique in the confusion graph $\mathfrak{C}(\mathcal{H})$.

Proof. It suffices to show that for any $\mathbf{x} \triangleleft H$ and $\mathbf{y} \triangleleft H$, where $\mathbf{x} \neq \mathbf{y}$, it holds that $\mathbf{x}$ and $\mathbf{y}$ are adjacent in $\mathfrak{C}(\mathcal{H})$.

Let $T = \{j : x_j = y_j\}$ and $S = \text{supp}(\mathbf{x}) \cup \text{supp}(\mathbf{y})$. Then $K = S \setminus T \neq \emptyset$ since $\mathbf{x} \neq \mathbf{y}$. As $K \subseteq H$, by Lemma 3.8.4, there exists some $i \in [m]$ such that $f(i) \in K$ and $\mathcal{X}_i \cap K = \emptyset$. Since $f(i) \in K$, we have $x_{f(i)} \neq y_{f(i)}$. On the other hand, since $\mathcal{X}_i \cap K = \emptyset$, by definition of K , we have $x_j = y_j$ for all $j \in \mathcal{X}_i$. Hence, $\mathbf{x}$ and $\mathbf{y}$ are adjacent in $\mathfrak{C}(\mathcal{H})$. ■

The α -bound for nonlinear ECIC's is stated as follows.

Proposition 3.8.6. *The length of an optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$\mathcal{N}_q(\delta, \mathcal{H}) \geq N_q(q^{\alpha(\mathcal{H})}, 2\delta + 1),$$

where $N_q(M, d)$ denotes the length of a shortest code of size M and distance d .

Proof. Let $\mathfrak{C}$ be an optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ of length $N = \mathcal{N}_q(\delta, \mathcal{H})$. Let H be a maximum generalized independent set of $\mathcal{H}$. Then $|H| = \alpha(\mathcal{H})$. By Lemma 3.8.5, the set $V(H) = \{\mathbf{x} \in \mathbb{F}_q^n : \mathbf{x} \triangleleft H\}$ forms a clique in $\mathcal{H}$. Moreover $|V(H)| = q^{\alpha(\mathcal{H})}$. By Lemma 3.8.3, the set $\mathcal{C}(H) \triangleq \{\mathfrak{C}(\mathbf{x}) : \mathbf{x} \in V(H)\}$ forms an $(N, q^{\alpha(\mathcal{H})}, 2\delta + 1)_q$ code. Therefore,

$$N \geq N_q(q^{\alpha(\mathcal{H})}, 2\delta + 1).$$

The proof follows. ■

Remark 3.8.7. This lower bound for *general* ECIC's is certainly less than or equal to the respective bound for *linear* ECIC's, which is $N_q[\alpha(\mathcal{H}), 2\delta + 1]$.

The Singleton bound for nonlinear ECIC's is stated as follows.

Proposition 3.8.8. *The length of an optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ satisfies*

$$\mathcal{N}_q(\delta, \mathcal{H}) \geq \kappa_q^* + 2\delta,$$

where κ_q^* is the length of an optimal $\mathcal{H}$ -IC over $\mathbb{F}_q$.

Proof. Let $\mathfrak{E}$ be an optimal $(\delta, \mathcal{H})$ -ECIC over $\mathbb{F}_q$ of length $N = \mathcal{N}_q(\delta, \mathcal{H})$. By Lemma 3.8.3,

$$\mathbf{d}(\mathfrak{E}(\mathbf{x}), \mathfrak{E}(\mathbf{x}')) \geq 2\delta + 1, \tag{3.22}$$

for every pair of confusable vectors $\mathbf{x}$ and $\mathbf{x}'$ in $\mathbb{F}_q^n$. Let $\mathfrak{E}^d : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{N-2\delta}$ be obtained from $\mathfrak{E}$ as follows

$$(\mathfrak{E}^d(\mathbf{x}))_i = (\mathfrak{E}(\mathbf{x}))_i, \quad i \in [N - 2\delta].$$

In other words, the values of $\mathfrak{E}^d$ is obtained from the corresponding values of $\mathfrak{E}$ by discarding the last 2δ coordinates. Then from (3.22) we have

$$\mathbf{d}(\mathfrak{E}^d(\mathbf{x}), \mathfrak{E}^d(\mathbf{x}')) \geq 1,$$

for every pair of confusable vectors $\mathbf{x}$ and $\mathbf{x}'$ in $\mathbb{F}_q^n$. Again by Lemma 3.8.3, $\mathfrak{E}^d$ is an $\mathcal{H}$ -IC over $\mathbb{F}_q$. Therefore, its length is bounded from below by κ_q^* . Thus, $N - 2\delta \geq \kappa_q^*$,

which concludes the proof. ■

An immediate corollary of this bound is that for sufficiently large q , an optimal ECIC can be obtained by applying an MDS code on top of an optimal nonlinear IC.

Corollary 3.8.9. *For $q \geq \kappa_q^* + 2\delta - 1$,*

$$\mathcal{N}_q(\delta, \mathcal{H}) = \kappa_q^* + 2\delta,$$

where κ_q^* is the length of an optimal $\mathcal{H}$ -IC over $\mathbb{F}_q$.

There has been very limited knowledge on the length κ_q^* of an optimal $\mathcal{H}$ -IC over $\mathbb{F}_q$. It was shown by Bar-Yossef *et al.* [3, 4] (in the original setting of the ICSI problem) that κ_q^* is bounded from below by $\alpha(\mathcal{D})$ where $\mathcal{D}$ is the corresponding side information (di)graph. Therefore, for ICSI instances whose side information (di)graphs $\mathcal{D}$ satisfy $\alpha(\mathcal{D}) = \text{minrk}_q(\mathcal{D})$, we have $\kappa_q^* = \alpha(\mathcal{D}) = \text{minrk}_q(\mathcal{D})$. For such instances, there exist optimal ECIC's which are linear. The instances described by perfect graphs and acyclic digraphs are the only known families of ICSI instances which possess this special feature [3]. In Chapter 5, several new families of such instances are discovered.

4. (DI)GRAPHS OF EXTREME MINRANKS

The *minrank* of a *graph* was introduced by Haemers [34] to bound the Shannon capacity [57] of a graph. This parameter of a *digraph* has recently drawn much more attention from the research community after the work of Bar-Yossef *et al.* [3]. In their paper, it was shown that the binary minrank of a digraph $\mathcal{D}$ characterizes the length of an optimal binary scalar linear index code for an ICSI instance described by the digraph $\mathcal{D}$. A generalization of this result for an arbitrary finite field $\mathbb{F}_q$ is straightforward. Since a graph can be regarded as a symmetric digraph, the same result also holds for graphs. In this chapter, we characterize the (di)graphs that have extreme minranks. Based on these characterizations, it is shown in the next chapter that the problem of deciding whether the minrank of a *digraph* is equal to two is NP-complete. In contrast, the same question for *graphs* can be answered in polynomial time. In the context of index coding, we only study minranks of (di)graphs over a *finite* field $\mathbb{F}_q$. However, all of our results presented in Chapter 4 and Chapter 5 (except Theorem 4.2.4, Corollary 4.2.5, and Corollary 5.1.2) still hold for an *arbitrary* field $\mathbb{F}$. This is because the characteristic of the field does not play any role in their proofs.

4.1 (Strongly) Connected Components and MinRanks

Lemma 4.1.1 (Folklore). *Let $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ be a graph. Suppose that $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_k$ are subgraphs of $\mathcal{G}$ that satisfy the following conditions*

1. The sets $\mathcal{V}(\mathcal{G}_i)$, $i \in [k]$, partition $\mathcal{V}(\mathcal{G})$;
2. There is no edge of the form $\{u, v\}$ where $u \in \mathcal{V}(\mathcal{G}_i)$ and $v \in \mathcal{V}(\mathcal{G}_j)$ for $i \neq j$.

Then

$$\text{minrk}_q(\mathcal{G}) = \sum_{i=1}^k \text{minrk}_q(\mathcal{G}_i).$$

In particular, the above equality holds if $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_k$ are all connected components of $\mathcal{G}$.

Proof. The proof follows directly from the fact that a matrix fits $\mathcal{G}$ if and only if it is a block diagonal matrix (relabeling the vertices if necessary), the blocks of which fit the corresponding subgraphs $\mathcal{G}_i$'s, $i \in [k]$. ■

Lemma 4.1.2 (Folklore). *Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph. If $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_k$ are all strongly connected components of $\mathcal{D}$, then*

$$\text{minrk}_q(\mathcal{D}) = \sum_{i=1}^k \text{minrk}_q(\mathcal{D}_i).$$

Proof. Suppose that $\mathcal{V}_i$ is the set of vertices that induces $\mathcal{D}_i$, $i \in [k]$. Then $\{\mathcal{V}_i\}_{i \in [k]}$ forms a partition of $\mathcal{V}(\mathcal{D})$. Relabeling the vertices of $\mathcal{D}$ if necessary, we may assume without loss of generality that for every $i < j$

1. $v_i < v_j$ whenever $v_i \in \mathcal{V}_i$ and $v_j \in \mathcal{V}_j$;
2. There are no arcs of the form (v_j, v_i) where $v_i \in \mathcal{D}_i$ and $v_j \in \mathcal{D}_j$.

If $\mathbf{M}_i$ is a matrix that fits $\mathcal{D}_i$, $i \in [k]$, then the diagonal block matrix $\mathbf{M}$ whose diagonal blocks are $\mathbf{M}_i$ clearly fits $\mathcal{D}$. Moreover, $\text{rank}_q(\mathbf{M}) = \sum_{i=1}^k \text{rank}_q(\mathbf{M}_i)$. Hence $\text{minrk}_q(\mathcal{D}) \leq \sum_{i=1}^k \text{minrk}_q(\mathcal{D}_i)$.

It remains to show that $\text{minrk}_q(\mathcal{D}) \geq \sum_{i=1}^k \text{minrk}_q(\mathcal{D}_i)$. Suppose that the matrix $\mathbf{M}$ fits $\mathcal{D}$. Then $\mathbf{M}$ is an upper-triangular block matrix, as shown in Fig. 4.1. If we let $\mathbf{M}_{i,i}$ be the sub-matrix of $\mathbf{M}$ formed by the rows and columns indexed by the elements of $\mathcal{V}_i$, then $\mathbf{M}_{i,i}$ fits $\mathcal{D}_i$ and hence,

$$\text{rank}_q(\mathbf{M}) \geq \sum_{i=1}^k \text{rank}_q(\mathbf{M}_{i,i}) \geq \sum_{i=1}^k \text{minrk}_q(\mathcal{D}_i).$$

Thus, $\text{minrk}_q(\mathcal{D}) \geq \sum_{i=1}^k \text{minrk}_q(\mathcal{D}_i)$.

	$\mathcal{V}_1$	$\mathcal{V}_2$	$\mathcal{V}_3$		$\mathcal{V}_{k-1}$	$\mathcal{V}_k$
$\mathcal{V}_1$				• • •		
$\mathcal{V}_2$	0			• • •		
$\mathcal{V}_3$	0	0				
	• • •	• • •	• • •	• • •	• • •	• • •
$\mathcal{V}_{k-1}$	0	0	0	• • •		
$\mathcal{V}_k$	0	0	0	• • •	0	

Fig. 4.1: Matrix $\mathbf{M}$ that fits $\mathcal{D}$

■

These two lemmas suggest that it is sufficient to study the minranks of connected graphs and strongly connected digraphs.

4.2 (Di)Graphs of Extreme MinRanks

In this section, we investigate the structural properties of (di)graphs that have extreme minranks.

4.2.1 (Di)Graphs of MinRank One

Lemma 4.2.1 (Folklore). *Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph. Then $\text{minrk}_q(\mathcal{D}) = 1$ if and only if $\mathcal{D}$ is a complete digraph. The same statement holds for a graph.*

Proof. Suppose $\mathcal{D}$ is a digraph. If $\text{minrk}_q(\mathcal{D}) = 1$, by definition there exists a matrix $\mathbf{M} = (m_{i,j})$ of rank one over $\mathbb{F}_q$ that fits $\mathcal{D}$. Then the rows of $\mathbf{M}$ must be scalar multiples of each other. Moreover, by definition, $a_{i,i} \neq 0$ for all $i \in \mathcal{V}(\mathcal{D})$. Hence $a_{i,j} \neq 0$ for all $i \in \mathcal{V}(\mathcal{D})$ and all $j \in \mathcal{V}(\mathcal{D})$. Therefore, $(i, j) \in \mathcal{E}(\mathcal{D})$ for all $i \neq j$, $i \in \mathcal{V}(\mathcal{D})$ and $j \in \mathcal{V}(\mathcal{D})$. In other words, $\mathcal{D}$ is a complete digraph.

Conversely, suppose that $\mathcal{D}$ is a complete digraph. Then $\mathbf{J}$, the all-one matrix, fits $\mathcal{D}$ and $\text{minrk}_q(\mathbf{J}) = 1$, which implies that $\text{minrk}_q(\mathcal{D}) = 1$. The same arguments hold for graphs. ■

4.2.2 (Di)Graphs of MinRank Two

We first study (di)graphs of minranks at most two. In this section, only *binary* alphabet is considered. We first introduce the following concept of a *fair coloring* of a digraph. Recall that a k -coloring of a graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ is a mapping $\phi : \mathcal{V}(\mathcal{G}) \rightarrow [k]$ which satisfies the condition that $\phi(u) \neq \phi(v)$ whenever $\{u, v\} \in \mathcal{E}(\mathcal{G})$. We often refer to $\phi(u)$ as the *color* of u . If there exists a k -coloring of $\mathcal{G}$, then we say that $\mathcal{G}$ is k -colorable.

Definition 4.2.2. Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph. A *fair k -coloring* of $\mathcal{D}$ is a mapping $\phi : \mathcal{V}(\mathcal{D}) \rightarrow [k]$ that satisfies the following conditions:

- (C1) If $(u, v) \in \mathcal{E}(\mathcal{D})$ then $\phi(u) \neq \phi(v)$;
- (C2) For each vertex u of $\mathcal{D}$, it holds that $\phi(v) = \phi(\omega)$ for all out-neighbors v and ω of u .

If there exists a fair k -coloring of $\mathcal{D}$, we say that we can *color $\mathcal{D}$ fairly by k colors*, or, $\mathcal{D}$ is *fairly k -colorable*.

We refer to the condition (C2) as the *fairness* of the coloring, since this condition guarantees that all out-neighbors of each vertex share the same color.

Lemma 4.2.3. *A digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ is fairly 3-colorable if and only if there exists a partition of $\mathcal{V}(\mathcal{D})$ into three subsets A , B , and C that satisfy the following conditions:*

1. *For every $i \in A$: either $N_O(i) \subseteq B$ or $N_O(i) \subseteq C$;*
2. *For every $i \in B$: either $N_O(i) \subseteq A$ or $N_O(i) \subseteq C$;*
3. *For every $i \in C$: either $N_O(i) \subseteq A$ or $N_O(i) \subseteq B$.*

Proof. If $\mathcal{D}$ is fairly 3-colorable, let A , B , and C respectively be the sets of vertices of $\mathcal{D}$ that share the same color. Then clearly A , B , and C partition $\mathcal{V}(\mathcal{D})$. Moreover, since all out-neighbors of each vertex must have the same color, the three conditions above are obviously satisfied. Conversely, if those conditions are satisfied, then $\phi : \mathcal{V}(\mathcal{D}) \rightarrow [3]$ defined by

$$\phi(u) = \begin{cases} 1, & u \in A, \\ 2, & u \in B, \\ 3, & u \in C, \end{cases}$$

is a fair 3-coloring of $\mathcal{D}$. ■

Theorem 4.2.4. *Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph. Then $\text{minrk}_2(\mathcal{D}) \leq 2$ if and only if $\overline{\mathcal{D}}$, the complement of $\mathcal{D}$, is fairly 3-colorable.*

Proof. Suppose $\mathcal{V}(\mathcal{D}) = [n]$. By definition, $\text{minrk}_2(\mathcal{D}) \leq 2$ if and only if there exists an $n \times n$ binary matrix $\mathbf{M}$ of rank at most two which fits $\mathcal{D}$. The matrix $\mathbf{M}$ has rank less than or equal to two if and only if there are two rows of $\mathbf{M}$ that span its row space. Without loss of generality, suppose that they are the first two rows of $\mathbf{M}$, namely, $\mathbf{M}_1$ and $\mathbf{M}_2$. Let

$$B = \text{supp}(\mathbf{M}_1) \cap \text{supp}(\mathbf{M}_2)$$

and

$$A = \text{supp}(\mathbf{M}_1) \setminus B, \quad C = \text{supp}(\mathbf{M}_2) \setminus B.$$

Since the binary alphabet is considered, for every $i \in [n]$, one of the following must hold: (1) $\mathbf{M}_i = \mathbf{M}_1$; (2) $\mathbf{M}_i = \mathbf{M}_2$; (3) $\mathbf{M}_i = \mathbf{M}_1 + \mathbf{M}_2$. Hence for every $i \in [n]$

$$i \in \text{supp}(\mathbf{M}_i) \subseteq A \cup B \cup C.$$

This implies that $A \cup B \cup C = [n]$.

Suppose that $i \in A$. Then either $\mathbf{M}_i = \mathbf{M}_1$ or $\mathbf{M}_i = \mathbf{M}_1 + \mathbf{M}_2$. The former holds if and only if $\text{supp}(\mathbf{M}_i) = A \cup B$. Equivalently, we obtain that $(i, j) \in \mathcal{E}(\mathcal{D})$ for all $j \in A \cup B \setminus \{i\}$. In other words, $(i, j) \notin \mathcal{E}(\overline{\mathcal{D}})$ for all $j \in A \cup B$. The latter holds if and only if $\text{supp}(\mathbf{M}_i) = A \cup C$. Similarly, this equality in turn is equivalent to the property that $(i, j) \notin \mathcal{E}(\overline{\mathcal{D}})$ for all $j \in A \cup C$. In summary, for every $i \in A$

1. $(i, j) \notin \mathcal{E}(\overline{\mathcal{D}})$, for all $j \in A$;

2. Either $(i, j) \notin \mathcal{E}(\overline{\mathcal{D}})$, for all $j \in B$, or $(i, j) \notin \mathcal{E}(\overline{\mathcal{D}})$, for all $j \in C$; In other words, either $\mathcal{N}_O^{\overline{\mathcal{D}}}(i) \subseteq B$ or $\mathcal{N}_O^{\overline{\mathcal{D}}}(i) \subseteq C$.

Analogous conditions hold for every $i \in B$ and for every $i \in C$ as well. Therefore, by Lemma 4.2.3, $\text{minrk}_2(\mathcal{D}) \leq 2$ if and only if $\overline{\mathcal{D}}$ is fairly 3-colorable. ■

The following corollary characterizes the digraphs of minrank two.

Corollary 4.2.5. *A digraph $\mathcal{D}$ has minrank two over $\mathbb{F}_2$ if and only if $\overline{\mathcal{D}}$ is fairly 3-colorable and $\mathcal{D}$ is not a complete digraph.*

4.2.3 (Di)Graphs of MinRanks Equal to Their Orders

For graphs, it is simple to characterize graphs of minranks equal to their orders.

Proposition 4.2.6. *Let $\mathcal{G}$ be a graph of order n . Then $\text{minrk}_q(\mathcal{G}) = n$ if and only if $\mathcal{G}$ has no edges.*

Proof. If $\mathcal{G}$ has no edges, a matrix fits $\mathcal{G}$ if and only if it is a diagonal matrix whose entries on the main diagonal are all nonzero. The rank of such a matrix is n . Therefore, $\text{minrk}_q(\mathcal{G}) = n$. Suppose that $\text{minrk}_q(\mathcal{G}) = n$ and (i, j) is an edge of $\mathcal{G}$. Then the matrix $\mathbf{M}$, where $\mathbf{M}_k = \mathbf{e}_k$ for all $k \notin \{i, j\}$ and $\mathbf{M}_i = \mathbf{M}_j = \mathbf{e}_i + \mathbf{e}_j$, fits $\mathcal{G}$. Moreover, $\text{rank}_q(\mathbf{M}) = n - 1$. Hence, $\text{minrk}_q(\mathcal{G}) \leq n - 1$. We obtain a contradiction. ■

For digraphs, the characterization is not that obvious.

Theorem 4.2.7. *Let $\mathcal{D}$ be a digraph of order n . Then $\text{minrk}_q(\mathcal{D}) = n$ if and only if $\mathcal{D}$ is acyclic.*

Proof. Equivalently, we show that $\text{minrk}_q(\mathcal{D}) \leq n - 1$ if and only if $\mathcal{D}$ has a circuit.

Suppose without loss of generality that $\mathcal{C} = (1, 2, \dots, \ell)$ is a circuit in $\mathcal{D}$. We construct a matrix $\mathbf{M}$ fitting $\mathcal{D}$ as follows. For $j \notin \mathcal{V}(\mathcal{C})$, let $\mathbf{M}_j = \mathbf{e}_j$. For $i \in [\ell - 1]$, let $\mathbf{M}_i = \mathbf{e}_i - \mathbf{e}_{i+1}$. Finally, let $\mathbf{M}_\ell = \mathbf{e}_1 - \mathbf{e}_\ell$. Clearly, $\mathbf{M}$ fits $\mathcal{D}$. Moreover, as $\mathbf{M}_\ell = \sum_{i=1}^{\ell-1} \mathbf{M}_i$, we have $\text{rank}_q(\mathbf{M}_{\mathcal{V}(\mathcal{C})}) = \ell - 1$. Hence

$$\begin{aligned} \text{rank}_q(\mathbf{M}) &= \text{rank}_q(\mathbf{M}_{\mathcal{V}(\mathcal{C})}) + \text{rank}_q(\mathbf{M}_{[n] \setminus \mathcal{V}(\mathcal{C})}) \\ &= (\ell - 1) + (n - \ell) \\ &= n - 1. \end{aligned}$$

Therefore, $\text{minrk}_q(\mathcal{D}) \leq n - 1$.

Conversely, suppose that $\text{minrk}_q(\mathcal{D}) \leq n - 1$. Then there exists a matrix $\mathbf{M}$ fitting $\mathcal{D}$ whose rows are linearly dependent. In other words, $\sum_{i \in I} \alpha_i \mathbf{M}_i = \mathbf{0}$ for some nonempty subset $I \subseteq [n]$ and for some $\alpha_i \in \mathbb{F}_q^*$, $i \in I$. Let $\mathcal{D}'$ be the subgraph of $\mathcal{D}$ induced by the vertices in I and $\mathbf{M}'$ the sub-matrix of $\mathbf{M}$ restricted to the rows and columns indexed by the elements of I . Obviously $\mathbf{M}'$ fits $\mathcal{D}'$. We show that there exists a circuit in $\mathcal{D}'$. Since $\sum_{i \in I} \alpha_i \mathbf{M}'_i = \mathbf{0}$, each column of $\mathbf{M}'$ has at least two nonzero entries. Therefore, for each vertex v of $\mathcal{D}'$, there exists another vertex u of $\mathcal{D}'$ such that (u, v) is an arc in $\mathcal{D}'$. Starting from an arbitrary vertex v_1 of $\mathcal{D}'$ and applying this property recursively, we obtain a sequence of vertices in $\mathcal{D}'$

$$v_1, v_2, \dots, v_s, v_{s+1}, \dots,$$

where (v_{s+1}, v_s) is an arc in $\mathcal{D}'$ for every $s \geq 1$. Since $\mathcal{D}'$ has a finite number of vertices, there must be a time (we care about the first time only) when a vertex appears twice in the above sequence. This vertex, together with the other vertices lying between its two occurrences, form a circuit inside $\mathcal{D}'$, which is also a circuit

inside $\mathcal{D}$. ■

The existence of a circuit in a digraph can be detected by using a depth-first search, the time complexity of which is polynomial with respect to the order of the digraph. Hence, as a consequence of Theorem 4.2.7, the problem of deciding whether a digraph has minrank equal to its order can be solved in polynomial time.

4.2.4 Graphs of MinRanks One Less Than Their Orders

We only consider graphs.

Definition 4.2.8. A graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ is called a *star graph* if $|\mathcal{V}(\mathcal{G})| \geq 2$ and there exists a vertex $v \in \mathcal{V}(\mathcal{G})$ such that $\mathcal{E}(\mathcal{G}) = \{\{u, v\} : u \in \mathcal{V}(\mathcal{G}) \setminus \{v\}\}$.

Fig. 4.2: A star graph

Theorem 4.2.9. Let $\mathcal{G}$ be a connected graph of order $n \geq 2$. Then $\text{minrk}_q(\mathcal{G}) = n - 1$ if and only if $\mathcal{G}$ is a star graph.

Proof. We first suppose that $\text{minrk}_q(\mathcal{G}) = n - 1$. If $n = 2$ then $\mathcal{G}$ must be a complete graph, which is also a star graph. We assume that $n \geq 3$. As $\mathcal{G}$ is connected, there exists a vertex v of degree at least two. Let v_1 and v_2 be any two distinct vertices adjacent to v . Our goal is to show that for every vertex $u \neq v$, we have $\{u, v\} \in \mathcal{E}(\mathcal{G})$, and those are all possible edges.

1. Suppose for a contradiction that $\{u, v\} \notin \mathcal{E}(\mathcal{G})$ for some $u \neq v$. Since $\mathcal{G}$ is connected, there exists ω such that $\{u, \omega\} \in \mathcal{E}(\mathcal{G})$. Then either $\omega \neq v_1$ or $\omega \neq v_2$. Suppose that $\omega \neq v_2$. We create a matrix $\mathbf{M}$ as follows. Let $\mathbf{M}_u = \mathbf{M}_\omega = \mathbf{e}_u + \mathbf{e}_\omega$, $\mathbf{M}_v = \mathbf{M}_{v_2} = \mathbf{e}_v + \mathbf{e}_{v_2}$, and $\mathbf{M}_i = \mathbf{e}_i$ for $i \notin \{u, v, v_2, \omega\}$. Then $\mathbf{M}$ fits $\mathcal{G}$ and $\text{rank}_q(\mathbf{M}) = n - 2$. Therefore, $\text{minrk}_q(\mathcal{G}) < n - 1$. We obtain a contradiction. Thus, all other vertices are adjacent to v .
2. Suppose for a contradiction that there exist two adjacent vertices, namely u and ω , both are different from v . As we just proved, both u and ω must be adjacent to v . We create a matrix $\mathbf{M}$ as follows. We take $\mathbf{M}_u = \mathbf{M}_v = \mathbf{M}_\omega = \mathbf{e}_u + \mathbf{e}_v + \mathbf{e}_\omega$, and $\mathbf{M}_i = \mathbf{e}_i$ for $i \notin \{u, v, \omega\}$. Clearly $\mathbf{M}$ fits $\mathcal{G}$ and moreover, $\text{rank}_q(\mathbf{M}) = n - 2$, which implies that $\text{minrk}_q(\mathcal{G}) < n - 1$. We obtain a contradiction.

Conversely, assume that $\mathcal{G}$ is a star graph, where $\mathcal{E}(\mathcal{G}) = \{\{v, v_i\} : i \in [n - 1]\}$, $v_i \in \mathcal{V}(\mathcal{G}) \setminus \{v\}$ for all $i \in [n - 1]$. We create a matrix $\mathbf{M}$ fitting $\mathcal{G}$ by taking $\mathbf{M}_{v_i} = \mathbf{e}_v + \mathbf{e}_{v_i}$ for $i \in [n - 1]$, and $\mathbf{M}_v = \mathbf{e}_v + \mathbf{e}_{v_1}$. Since $\mathbf{M}_v \equiv \mathbf{M}_{v_1}$, we deduce that $\text{rank}_q(\mathbf{M}) = n - 1$. Hence $\text{minrk}_q(\mathcal{G}) \leq n - 1$. On the other hand, since $\{v_i : i \in [n - 1]\}$ is a maximum independent set in $\mathcal{G}$, we obtain that $\alpha(\mathcal{G}) = n - 1$. By Theorem 3.2.11, $\text{minrk}_q(\mathcal{G}) \geq \alpha(\mathcal{G}) = n - 1$. Thus, $\text{minrk}_q(\mathcal{G}) = n - 1$. $\blacksquare$

Corollary 4.2.10. *Let $\mathcal{G}$ be a graph of order $n \geq 2$ and let $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_k$ be all of its connected components, with*

$$|\mathcal{V}(\mathcal{G}_1)| \geq |\mathcal{V}(\mathcal{G}_2)| \geq \dots \geq |\mathcal{V}(\mathcal{G}_k)|.$$

Then $\text{minrk}_q(\mathcal{G}) = n - 1$ if and only if the following conditions are satisfied

1. $\mathcal{G}_1$ is a star graph having at least two vertices,

2. $\mathcal{G}_i$ is a one-vertex graph for every $i \geq 2$.

Proof. Suppose that $\text{minrk}_q(\mathcal{G}) = n - 1$. Then $\mathcal{G}$ must have some connected component that contains at least two vertices. Otherwise, by Proposition 4.2.6, $\text{minrk}_q(\mathcal{G})$ would be n . Hence, as the largest component, $\mathcal{G}_1$ must contain at least two vertices. Suppose for a contradiction that there exists another connected component, say $\mathcal{G}_2$, that contains more than one vertex. As each of $\mathcal{G}_1$ and $\mathcal{G}_2$ contains at least one edge, by Lemma 4.1.1 and Proposition 4.2.6, we obtain

$$\begin{aligned} \text{minrk}_q(\mathcal{G}) &= \text{minrk}_q(\mathcal{G}_1) + \text{minrk}_q(\mathcal{G}_2) + \sum_{i=3}^k \text{minrk}_q(\mathcal{G}_i) \\ &\leq (|\mathcal{V}(\mathcal{G}_1)| - 1) + (|\mathcal{V}(\mathcal{G}_2)| - 1) + \sum_{i=3}^k |\mathcal{V}(\mathcal{G}_i)| \\ &< n - 1. \end{aligned}$$

We have a contradiction. Therefore, except from $\mathcal{G}_1$, all other connected components each contains precisely one vertex. Moreover, $\text{minrk}_q(\mathcal{G}_1) = |\mathcal{V}(\mathcal{G}_1)| - 1$, for otherwise,

$$\begin{aligned} \text{minrk}_q(\mathcal{G}) &= \text{minrk}_q(\mathcal{G}_1) + \sum_{i=2}^k \text{minrk}_q(\mathcal{G}_i) \\ &\leq (|\mathcal{V}(\mathcal{G}_1)| - 2) + \sum_{i=2}^k |\mathcal{V}(\mathcal{G}_i)| \\ &< n - 1. \end{aligned}$$

Therefore, by Theorem 4.2.9, $\mathcal{G}_1$ must be a star graph.

Conversely, suppose that $\mathcal{G}_1$ is a star graph with at least two vertices, and that

$\mathcal{G}_i$ for $i > 1$ contains precisely one vertex. Then $|\mathcal{V}(\mathcal{G}_1)| = n - k + 1$ and

$$\begin{aligned} \text{minrk}_q(\mathcal{G}) &= \text{minrk}_q(\mathcal{G}_1) + \sum_{i=2}^k \text{minrk}_q(\mathcal{G}_i) \\ &= (|\mathcal{V}(\mathcal{G}_1)| - 1) + (k - 1) \\ &= n - 1. \end{aligned}$$

■

5. COMPUTATION OF MINRANKS OF SIDE INFORMATION (DI)GRAPHS

In this chapter, the computational aspects of the ICSI problem are investigated. Firstly, we show that deciding whether a *digraph* has minrank two is an NP-complete problem. For comparison, there is a polynomial time to decide if a *graph* has minrank two. Secondly, we establish a new upper bound for the minrank of a digraph, namely the *circuit-packing bound*, which is arguably more suitable for digraphs than the best currently known upper bound. Employing this new bound, we point out several families of digraphs whose minranks can be found in polynomial time. Moreover, for ICSI instances described by such digraphs, scalar linear index codes are shown to be optimal. Thirdly, a polynomial time dynamic programming algorithm is developed to compute the minranks of a family of graphs possessing a special tree structure. Intuitively, such graphs are obtained by gluing together, in a tree-like structure, several graphs for which the minranks can be determined in polynomial time. Finally, using a computer program, we compute the minranks of all non-isomorphic graphs of order up to 10.

5.1 *The Hardness of the MinRank Problem for Digraphs*

In this section, we first prove that it is an NP-complete problem to decide whether a given digraph is fairly k -colorable (see Definition 4.2.2), for $k \geq 3$. The hardness

of this problem, by Lemma 4.2.1 and Corollary 4.2.5, leads to the hardness of the problem of deciding whether a given digraph has minrank two over $\mathbb{F}_2$. The fair k -coloring problem is defined formally as follows.

Problem: **FAIR k -COLORING**

Instance: A digraph $\mathcal{D}$

Question: Is $\mathcal{D}$ fairly k -colorable?

Theorem 5.1.1. *The fair k -coloring problem is NP-complete for $k \geq 3$.*

Proof. This problem is obviously in NP. For NP-hardness, we reduce the k -coloring problem to the fair k -coloring problem. Suppose that $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ is an arbitrary graph. We aim to build a digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ so that $\mathcal{G}$ is k -colorable if and only if $\mathcal{D}$ is fairly k -colorable. Suppose that $\mathcal{V}(\mathcal{G}) = [n]$. For each $i \in [n]$, we build the following gadget, which is a digraph $\mathcal{D}_i = (\mathcal{V}_i, \mathcal{E}_i)$. The vertex set of $\mathcal{D}_i$ is

$$\mathcal{V}_i = \{i\} \cup \{\omega_{i,j} : j \in N^{\mathcal{G}}(i)\},$$

where $\omega_{i,j}$ are newly introduced vertices. We refer to $\omega_{i,j}$ as a *clone* (in $\mathcal{D}_i$) of the vertex $j \in [n]$. The arc set of $\mathcal{D}_i$ is

$$\mathcal{E}_i = \{(\omega_{i,j}, i) : j \in N^{\mathcal{G}}(i)\}.$$

Let $N^{\mathcal{G}}(i) = \{i_1, i_2, \dots, i_{n_i}\}$. Then $\mathcal{D}_i$ can be drawn as in Fig. 5.1.

Moreover, we also introduce n new vertices, which are $p_1, p_2, \dots, p_n$. The digraph $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ is built as follows. The vertex set of $\mathcal{D}$ is

$$\mathcal{V}(\mathcal{D}) = \left(\bigcup_{i=1}^n \mathcal{V}_i \right) \cup \{p_1, p_2, \dots, p_n\}.$$

Fig. 5.1: Gadget $\mathcal{D}_i$ for each vertex i of $\mathcal{G}$

The arc set of $\mathcal{D}$ is

$$\mathcal{E}(\mathcal{D}) = \left(\bigcup_{i=1}^n \mathcal{E}_i \right) \cup \left(\bigcup_{i=1}^n \mathcal{Q}_i \right),$$

where $\mathcal{Q}_i$ consists of (p_i, i) and the arcs that connect p_i and all the clones of i . More formally,

$$\mathcal{Q}_i = \{(p_i, i)\} \cup \{(p_i, \omega_{i',i}) : i' \in [n], i \in N^{\mathcal{G}}(i')\}.$$

For example, if $\mathcal{G}$ is the graph in Fig. 5.2, then $\mathcal{D}$ is the digraph in Fig. 5.3.

Fig. 5.2: The graph $\mathcal{G}$

Our goal now is to show that $\mathcal{G}$ is k -colorable if and only if $\mathcal{D}$ is fairly k -colorable.

Suppose that $\mathcal{G}$ is k -colorable and $\phi_{\mathcal{G}} : [n] \rightarrow [k]$ is a k -coloring of $\mathcal{G}$. We consider the mapping $\phi_{\mathcal{D}} : \mathcal{V}(\mathcal{D}) \rightarrow [k]$ defined as follows

1. For every $i \in [n]$, $\phi_{\mathcal{D}}(i) \triangleq \phi_{\mathcal{G}}(i)$;
2. If $i \in N^{\mathcal{G}}(i')$ then $\phi_{\mathcal{D}}(\omega_{i',i}) \triangleq \phi_{\mathcal{D}}(i) = \phi_{\mathcal{G}}(i)$, in other words, clones of i have

Fig. 5.3: The digraph $\mathcal{D}$ built from the graph $\mathcal{G}$ in Fig. 5.2

the same color as i ;

3. For every $i \in [n]$, $\phi_{\mathcal{D}}(p_i)$ can be chosen arbitrarily, as long as it is different from $\phi_{\mathcal{D}}(i)$.

We claim that $\phi_{\mathcal{D}}$ is a fair k -coloring for $\mathcal{D}$. We first verify the condition (C1) (see Definition 4.2.2). It is straightforward from the definition of $\phi_{\mathcal{D}}$ that the endpoints of each of the arcs of the forms (p_i, i) for $i \in [n]$, and $(p_i, \omega_{i',i})$ for $i \in N^{\mathcal{G}}(i')$, have different colors. It remains to check if i and $\omega_{i,j}$ for $j \in N^{\mathcal{G}}(i)$ have different colors. On the one hand, $\omega_{i,j}$ is a clone of j , and hence has the same color as j . In other words,

$$\phi_{\mathcal{D}}(\omega_{i,j}) = \phi_{\mathcal{D}}(j) = \phi_{\mathcal{G}}(j).$$

On the other hand, since $j \in N^{\mathcal{G}}(i)$, we obtain that

$$\phi_{\mathcal{G}}(j) \neq \phi_{\mathcal{G}}(i) = \phi_{\mathcal{D}}(i).$$

Therefore, $\phi_{\mathcal{D}}(\omega_{i,j}) \neq \phi_{\mathcal{D}}(i)$ for all $i \in [n]$ and $j \in N^{\mathcal{G}}(i)$. Thus, (C1) is satisfied.

We now check if (C2) (see Definition 4.2.2) is also satisfied. The out-neighbors

of p_i are i and its clones $\omega_{i',i}$ ($i \in N^{\mathcal{G}}(i')$). These vertices have the same color in $\mathcal{D}$, namely $\phi_{\mathcal{G}}(i)$, by the definition of $\phi_{\mathcal{D}}$. Thus (C2) is also satisfied. Therefore $\phi_{\mathcal{D}}$ is a fair k -coloring of $\mathcal{D}$.

Conversely, suppose that $\phi_{\mathcal{D}} : \mathcal{V}(\mathcal{D}) \rightarrow [k]$ is a fair k -coloring of $\mathcal{D}$. Condition (C2) guarantees that all clones of i have the same color as i , namely, $\phi_{\mathcal{D}}(\omega_{i',i}) = \phi_{\mathcal{D}}(i)$ if $i \in N^{\mathcal{G}}(i')$. Therefore, by (C1), if $\{i, j\} \in \mathcal{E}(\mathcal{G})$, that is, $j \in N^{\mathcal{G}}(i)$, then

$$\phi_{\mathcal{D}}(i) \neq \phi_{\mathcal{D}}(\omega_{i,j}) = \phi_{\mathcal{D}}(j).$$

Hence, if we define $\phi_{\mathcal{G}} : [n] \rightarrow [k]$ by $\phi_{\mathcal{G}}(i) = \phi_{\mathcal{D}}(i)$ for all $i \in [n]$, then it is a k -coloring of $\mathcal{G}$. Thus $\mathcal{G}$ is k -colorable.

Finally, notice that the order of $\mathcal{D}$ is a polynomial with respect to the order of $\mathcal{G}$. More specifically, $\mathcal{V}(\mathcal{D}) = 2|\mathcal{V}(\mathcal{G})| + 2|\mathcal{E}(\mathcal{G})|$ and $\mathcal{E}(\mathcal{D}) = |\mathcal{V}(\mathcal{G})| + 4|\mathcal{E}(\mathcal{G})|$. Moreover, building $\mathcal{D}$ from $\mathcal{G}$, and also obtaining a coloring of $\mathcal{G}$ from a coloring of $\mathcal{D}$, can be done in polynomial time with respect to the order of $\mathcal{G}$. Since the k -coloring problem ($k \geq 3$) is NP-hard [40], we conclude that the fair k -coloring problem is also NP-hard. ■

Corollary 5.1.2. *Given an arbitrary digraph $\mathcal{D}$, deciding whether $\text{minrk}_2(\mathcal{D}) = 2$ is NP-complete.*

In contrast, for a graph $\mathcal{G}$, it was observed by Peeters [53] that $\mathcal{G}$ has minrank two if and only if $\overline{\mathcal{G}}$ is a bipartite graph and $\mathcal{G}$ is not a complete graph, which can be verified in polynomial time. This fact can also be derived by applying Theorem 4.2.4 to $\mathcal{D}^{\mathcal{G}}$, the digraph corresponding to $\mathcal{G}$ ($\mathcal{D}^{\mathcal{G}}$ is obtained from $\mathcal{G}$ by replacing each edge of $\mathcal{G}$ by two opposite arcs). Indeed, we first observe that a symmetric digraph is fairly 3-colorable if and only if it is fairly 2-colorable. Therefore, a symmetric

digraph is fairly 3-colorable if and only if its underlying graph is 2-colorable, i.e., is a bipartite graph. Hence, by Theorem 4.2.4, $\text{minrk}_2(\mathcal{D}^{\mathcal{G}}) \leq 2$ if and only if $\overline{\mathcal{G}}$ is bipartite. As $\mathcal{G}$ and $\mathcal{D}^{\mathcal{G}}$ have the same minrank, we conclude that $\text{minrk}_2(\mathcal{G}) \leq 2$ if and only if $\overline{\mathcal{G}}$ is bipartite.

5.2 Circuit-Packing Bound

In this section we introduce a new upper bound for the minrank of a digraph. This bound reveals some new families of digraphs whose minranks are computable in polynomial time.

5.2.1 The Bound

We first introduce an easy lemma.

Lemma 5.2.1. *If $\mathcal{D}'$ is a subgraph of a digraph $\mathcal{D}$ then*

$$\text{minrk}_q(\mathcal{D}) \geq \text{minrk}_q(\mathcal{D}').$$

The same conclusion also holds for graphs.

Proof. Let $\mathbf{M}$ be a matrix that fits $\mathcal{D}$ and has rank equal to the minrank of $\mathcal{D}$. Then the sub-matrix $\mathbf{M}'$ of $\mathbf{M}$ restricted on the rows and columns indexed by the vertices in $\mathcal{V}(\mathcal{D}')$ is a matrix that fits $\mathcal{D}'$. Then

$$\text{minrk}_q(\mathcal{D}') \leq \text{rank}_q(\mathbf{M}') \leq \text{rank}_q(\mathbf{M}) = \text{minrk}_q(\mathcal{D}). \quad \blacksquare$$

The following lower bound on the minrank of a *digraph* is an immediate corollary of Theorem 4.2.7, which was also established by Bar-Yossef *et al.* [3, 4].

Corollary 5.2.2. *For a digraph $\mathcal{D}$ we have*

$$\text{minrk}_q(\mathcal{D}) \geq \alpha(\mathcal{D}),$$

where $\alpha(\mathcal{D})$ denotes the order of a maximum acyclic induced subgraph of $\mathcal{D}$.

Proof. Let $\mathcal{D}'$ be a maximum acyclic induced subgraph of $\mathcal{D}$ of order $\alpha(\mathcal{D})$. Since $\mathcal{D}'$ is acyclic, by Lemma 5.2.1 and Theorem 4.2.7 we have

$$\text{minrk}_q(\mathcal{D}) \geq \text{minrk}_q(\mathcal{D}') = \alpha(\mathcal{D}). \quad \blacksquare$$

Let $\nu_0(\mathcal{D})$ be the *circuit packing number* of $\mathcal{D}$, namely, the maximum number of vertex-disjoint circuits in $\mathcal{D}$. We establish below an upper bound on minranks of digraphs. This bound was also obtained independently by Chaudhry *et al.* [14, 15].

Proposition 5.2.3 (Circuit-packing bound). *The following holds for every digraph $\mathcal{D}$ of order n :*

$$\text{minrk}_q(\mathcal{D}) \leq n - \nu_0(\mathcal{D}).$$

Proof. Suppose $\mathcal{D}$ contains $\nu_0(\mathcal{D})$ vertex-disjoint circuits $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{\nu_0(\mathcal{D})}$, where

$$\mathcal{C}_i = (u_{i,1}, u_{i,2}, \dots, u_{i,n_i}), \quad i \in [\nu_0(\mathcal{D})].$$

We construct a matrix $\mathbf{M}$ fitting $\mathcal{D}$ as follows. Let

$$\mathcal{A} \triangleq [n] \setminus \bigcup_{i \in [\nu_0(\mathcal{D})]} \mathcal{V}(\mathcal{C}_i).$$

For $v \in \mathcal{A}$ let $\mathbf{M}_v = \mathbf{e}_v$. For $i \in [\nu_0(\mathcal{D})]$ and $s \in [n_i - 1]$, let

$$\mathbf{M}_{u_{i,s}} = \mathbf{e}_{u_{i,s}} - \mathbf{e}_{u_{i,s+1}},$$

and let

$$\mathbf{M}_{u_{i,n_i}} = \mathbf{e}_{u_{i,1}} - \mathbf{e}_{u_{i,n_i}}.$$

Clearly, $\mathbf{M}$ fits $\mathcal{D}$. Moreover, as

$$\mathbf{M}_{u_{i,n_i}} = \sum_{s=1}^{n_i-1} \mathbf{M}_{u_{i,s}},$$

we have

$$\text{rank}_q(\mathbf{M}_{\mathcal{V}(\mathcal{C}_i)}) = n_i - 1$$

for all $i \in [\nu_0(\mathcal{D})]$. Since $\mathcal{V}(\mathcal{C}_i)$'s, $i \in [\nu_0(\mathcal{D})]$, are pairwise disjoint, we have

$$\begin{aligned} \text{rank}_q(\mathbf{M}) &= \sum_{i=1}^{\nu_0(\mathcal{D})} \text{rank}_q(\mathbf{M}_{\mathcal{V}(\mathcal{C}_i)}) + \text{rank}_q(\mathbf{M}_{\mathcal{A}}) \\ &= \sum_{i=1}^{\nu_0(\mathcal{D})} (n_i - 1) + \left(n - \sum_{i=1}^{\nu_0(\mathcal{D})} n_i \right) \\ &= n - \nu_0(\mathcal{D}). \end{aligned}$$

Thus, $\text{minrk}_q(\mathcal{D}) \leq n - \nu_0(\mathcal{D})$. ■

The clique-cover bound for graphs (see Theorem 3.2.11) can be easily extended to a bound for digraphs as follows.

Proposition 5.2.4. *It holds for any digraph $\mathcal{D}$ that $\text{minrk}_q(\mathcal{D}) \leq \text{cc}(\mathcal{D})$.*

Proof. Let $\mathcal{V}(\mathcal{D}) = [n]$ and $\mathcal{W} = \{\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_{\text{cc}(\mathcal{D})}\}$ a minimum clique cover of

Fig. 5.4: Example where the circuit-packing bound outperforms the clique-cover bound

$\mathcal{D}$. Then $[n] = \cup_{i=1}^{\text{cc}(\mathcal{D})} \mathcal{W}_i$. We aim to construct a matrix $\mathbf{M}$ that fits $\mathcal{D}$ and has rank at most $\text{cc}(\mathcal{D})$. Notice that by definition, $\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_{\text{cc}(\mathcal{D})}$ partition $\mathcal{V}(\mathcal{D})$. For each $i \in [\text{cc}(\mathcal{D})]$ and $j \in \mathcal{W}_i$, let $\mathbf{M}_j = \sum_{s \in \mathcal{W}_i} \mathbf{e}_s$. Then $\mathbf{M}$ fits $\mathcal{D}$ and

$$\text{rank}_q(\mathbf{M}) \leq \sum_{i=1}^{\text{cc}(\mathcal{D})} \text{rank}_q(\mathbf{M}_{\mathcal{W}_i}) \leq \sum_{i=1}^{\text{cc}(\mathcal{D})} 1 = \text{cc}(\mathcal{D}). \quad \blacksquare$$

Whereas for graphs the clique-cover bound is the best known bound, for digraphs that are not symmetric, this is not the case. The worst scenario for the clique-cover bound is when the digraph has no two arcs of opposite directions. For such a digraph, this bound becomes trivial, as the size of the smallest clique cover is equal to the order of the digraph. The following example emphasizes the fact that for certain digraphs, the circuit-packing bound can be *significantly better* than the clique-cover bound.

Example 5.2.5. Let $\mathcal{D}_k$ be the digraph of order $n = 3k$ depicted in Fig. 5.4. As there are no arcs of opposite directions, all cliques in $\mathcal{D}_k$ are of cardinality one. Therefore, the clique-cover bound gives $\text{minrk}_q(\mathcal{D}_k) \leq 3k$. On the other hand, as $\mathcal{D}_k$ contains k vertex-disjoint circuits, namely $\mathcal{C}_i = (3i + 1, 3i + 2, 3i + 3)$ for $i = 0, 1, \dots, k - 1$, the circuit-packing bound yields $\text{minrk}_q(\mathcal{D}_k) \leq 2k = 3k - k$. The gap between the two bounds is a third of the order of the digraph.

5.2.2 Digraphs Attaining Circuit-Packing Bound

A *feedback vertex (arc, respectively) set* of $\mathcal{D}$ is a set of vertices (arcs, respectively) whose removal destroys all circuits in $\mathcal{D}$. Let $\tau_0(\mathcal{D})$ ($\tau_1(\mathcal{D})$, respectively) denote the *minimum size* of a feedback vertex (arc, respectively) set of $\mathcal{D}$. Then it is clear that $\alpha(\mathcal{D}) = n - \tau_0(\mathcal{D})$.

Corollary 5.2.6. *If $\nu_0(\mathcal{D}) = \tau_0(\mathcal{D})$ then*

$$\text{minrk}_q(\mathcal{D}) = n - \nu_0(\mathcal{D}) = n - \tau_0(\mathcal{D}).$$

Proof. By Corollary 5.2.2 and Proposition 5.2.3 we have

$$n - \tau_0(\mathcal{D}) \leq \text{minrk}_q(\mathcal{D}) \leq n - \nu_0(\mathcal{D}).$$

Hence, the proof follows. ■

When $\mathcal{D}$ satisfies $\nu_0(\mathcal{D}) = \tau_0(\mathcal{D})$, we say that $\mathcal{D}$ satisfies the *min-max vertex equality*. In that case, the circuit-packing bound is actually tight. Similarly, let $\nu_1(\mathcal{D})$ denotes the maximum number of arc-disjoint circuits in $\mathcal{D}$. We say that $\mathcal{D}$ satisfies the *min-max arc equality* if $\nu_1(\mathcal{D}) = \tau_1(\mathcal{D})$.

The first family of digraphs for which the circuit-packing bound is tight is the family of *fully reducible flow digraphs* [30]. A flow digraph is a digraph that contains a special vertex called root, from which any vertex is reachable by a path. A fully reducible flow digraph is a flow digraph that satisfies the property that every circuit $\mathcal{C}$ in the digraph has a unique vertex $v_{\mathcal{C}}$ such that every path from the root to a vertex of $\mathcal{C}$ must contain $v_{\mathcal{C}}$. Interestingly, it was proved by Shamir [56] that there is a *linear time* algorithm to find $\nu_0(\mathcal{D}) = \tau_0(\mathcal{D})$ for a fully reducible flow digraph

$\mathcal{D}$. As a consequence, the minrank of a fully reducible flow digraph (recognizable in polynomial time with respect to its size [63]) can be calculated in linear time with respect to its size.

The second family of digraphs that satisfy the min-max vertex equality is the family of *connectively reducible digraphs* [61]. This family actually generalizes both the family of fully reducible flow digraphs and the family of *cyclically reducible digraphs* [66]. A polynomial time algorithm was provided by Szwarcfiter [61] to recognize a member of this family and subsequently find a maximum set of vertex-disjoint circuits as well as a minimum feedback vertex set. Therefore, by Corollary 5.2.6,

$$\text{minrk}_q(\mathcal{D}) = n - \nu_0(\mathcal{D}) = n - \tau_0(\mathcal{D})$$

for a connectively reducible digraph D and moreover, it can be found in polynomial time.

The third family of digraphs for which the circuit-packing bound is tight is the family of digraphs that *pack* [32]. A digraph packs if the min-max vertex equality holds for all of its subgraphs. The digraphs in this family are exactly ones that have no minor isomorphic to an odd double circuit or F_7 , a special digraph of order 7 (interested readers may refer to [32] for more details, also for a structural characterization of this family of digraphs). For instance, *strongly planar* digraphs [32] belong to this family. As far as we know, there are no known polynomial time algorithms to find a minimum feedback vertex set of a digraph that packs.

The last three families of digraphs for which the circuit-packing bound is tight are families of *line digraphs* of planar digraphs, of fully reducible flow digraphs, and of (special) Eulerian digraphs.

Definition 5.2.7. Let $\mathcal{D} = (\mathcal{V}(\mathcal{D}), \mathcal{E}(\mathcal{D}))$ be a digraph. Then the digraph $\mathcal{L} =$

$(\mathcal{V}(\mathcal{L}), \mathcal{E}(\mathcal{L}))$ with $\mathcal{V}(\mathcal{L}) = \mathcal{E}(\mathcal{D})$ and

$$\mathcal{E}(\mathcal{L}) = \{(e, e') : e = (u, v) \in \mathcal{E}(\mathcal{D}), e' = (v, w) \in \mathcal{E}(\mathcal{D})\},$$

is called the *line digraph* of $\mathcal{D}$. We denote the line digraph of $\mathcal{D}$ by $\mathcal{L}(\mathcal{D})$. The digraph $\mathcal{D}$ is called a *root digraph* of $\mathcal{L}(\mathcal{D})$.

Lemma 5.2.8. $\nu_0(\mathcal{L}(\mathcal{D})) = \nu_1(\mathcal{D})$.

Proof.

1. $\nu_0(\mathcal{L}(\mathcal{D})) \geq \nu_1(\mathcal{D})$. It suffices to show that the existence of a set of arc-disjoint circuits in $\mathcal{D}$ implies the existence of a set of vertex-disjoint circuits of the same size in $\mathcal{L}(\mathcal{D})$. Let $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k\}$ be a set of arc-disjoint circuits in $\mathcal{D}$, where $\mathcal{C}_i = (v_{i,1}, v_{i,2}, \dots, v_{i,r_i})$, $r_i \geq 2$, $i \in [k]$. Let $e_{i,j} = (v_{i,j}, v_{i,j+1})$, for $i \in [k]$ and $j \in [r_i - 1]$. Moreover, let $e_{i,r_i} = (v_{i,r_i}, v_{i,1})$ for $i \in [k]$. Let $\mathcal{C}'_i = (e_{i,1}, e_{i,2}, \dots, e_{i,r_i})$ for $i \in [k]$. Then $\mathcal{C}'_i$ is also a circuit in $\mathcal{L}(\mathcal{D})$ for every $i \in [k]$. Moreover, as the circuits $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$ share no common edges in $\mathcal{D}$, we deduce that $\mathcal{C}'_1, \mathcal{C}'_2, \dots, \mathcal{C}'_k$ share no common vertices in $\mathcal{L}(\mathcal{D})$. Therefore, they form a set of k vertex-disjoint circuits in $\mathcal{L}(\mathcal{D})$.
2. $\nu_0(\mathcal{L}(\mathcal{D})) \leq \nu_1(\mathcal{D})$. It suffices to show that the existence of a set of vertex-disjoint circuits in $\mathcal{L}(\mathcal{D})$ implies the existence of a set of arc-disjoint circuits of the same size in $\mathcal{D}$. Let $\{\mathcal{C}'_1, \mathcal{C}'_2, \dots, \mathcal{C}'_k\}$ be a set of vertex-disjoint circuits in $\mathcal{L}(\mathcal{D})$, where $\mathcal{C}'_i = \{e_{i,1}, e_{i,2}, \dots, e_{i,r_i}\}$ for $i \in [k]$. Suppose that $e_{i,j} = (v_{i,j}, v_{i,j+1}) \in \mathcal{E}(\mathcal{D})$ for $i \in [k]$ and $j \in [r_i]$, where $v_{i,j}$ and $v_{i,j+1}$ are vertices of $\mathcal{D}$. Then $v_{i,r_i+1} \equiv v_{i,1}$ for $i \in [k]$. For each $i \in [k]$, consider the sequence of

(possibly repeated) vertices

$$v_{i,1}, v_{i,2}, \dots, v_{i,r_i+1}.$$

Since $v_{i,1} \equiv v_{i,r_i+1}$ and $(v_{i,j}, v_{i,j+1}) \in \mathcal{E}(\mathcal{D})$ for all $j \in [r_i]$, there exist j_0 and j_1 such that

- (a) $1 \leq j_0 < j_1 \leq r_i$;
- (b) $v_{i,j_0} \equiv v_{i,j_1+1}$;
- (c) $v_{i,j_0}, v_{i,j_0+1}, \dots, v_{i,j_1}$ are distinct.

Then $\mathcal{C}_i = (v_{i,j_0}, v_{i,j_0+1}, \dots, v_{i,j_1})$ is a circuit in $\mathcal{D}$. Since the circuits $\mathcal{C}'_1, \mathcal{C}'_2, \dots, \mathcal{C}'_k$ share no common vertices in $\mathcal{L}(\mathcal{D})$, we obtain that the circuits $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$ share no common edges in $\mathcal{D}$. ■

Lemma 5.2.9. $\tau_0(\mathcal{L}(\mathcal{D})) = \tau_1(\mathcal{D})$.

Proof. Let $F = \{e_1, e_2, \dots, e_k\}$, where $e_i \in \mathcal{E}(\mathcal{D})$ for $i \in [k]$, be an arbitrary set of arcs of $\mathcal{D}$. We can also consider F as a set of vertices of $\mathcal{L}(\mathcal{D})$. It suffices to show that F is a feedback arc set of $\mathcal{D}$ if and only if F is a feedback vertex set of $\mathcal{L}(\mathcal{D})$, for every such set F .

Let $\mathcal{D} - F$ be the digraph obtained from $\mathcal{D}$ by removing all arcs in F . Let $\mathcal{L}(\mathcal{D}) - F$ be the digraph obtained from $\mathcal{L}(\mathcal{D})$ by removing all vertices in F . Then $\mathcal{L}(\mathcal{D}) - F = \mathcal{L}(\mathcal{D} - F)$. As shown in the proof of Lemma 5.2.8, the existence of a circuit in $\mathcal{D} - F$ would result in the existence of a circuit in $\mathcal{L}(\mathcal{D} - F)$ and vice versa. Therefore, $\mathcal{D} - F$ is acyclic if and only if $\mathcal{L}(\mathcal{D}) - F$ is acyclic. Thus, F is a feedback arc set of $\mathcal{D}$ if and only if F is a feedback vertex set of $\mathcal{L}(\mathcal{D})$. ■

Proposition 5.2.10. *Let $\mathcal{D}$ be a digraph. If $\nu_1(\mathcal{D}) = \tau_1(\mathcal{D})$ then $\nu_0(\mathcal{L}(\mathcal{D})) = \tau_0(\mathcal{L}(\mathcal{D}))$ and*

$$\text{minrk}_q(\mathcal{L}(\mathcal{D})) = |\mathcal{V}(\mathcal{L}(\mathcal{D}))| - \nu_0(\mathcal{L}(\mathcal{D})) = |\mathcal{E}(\mathcal{D})| - \nu_1(\mathcal{D}).$$

Proof. Suppose that $\nu_1(\mathcal{D}) = \tau_1(\mathcal{D})$. By Lemma 5.2.8 and Lemma 5.2.9, $\nu_0(\mathcal{L}(\mathcal{D})) = \tau_0(\mathcal{L}(\mathcal{D}))$. Therefore, applying Corollary 5.2.6 to $\mathcal{L}(\mathcal{D})$ we obtain

$$\text{minrk}_q(\mathcal{L}(\mathcal{D})) = |\mathcal{V}(\mathcal{L}(\mathcal{D}))| - \nu_0(\mathcal{L}(\mathcal{D})) = |\mathcal{E}(\mathcal{D})| - \nu_1(\mathcal{D}). \quad \blacksquare$$

Definition 5.2.11. A (di)graph that can be drawn on a plane in such a way that its (arcs) edges intersect only at their endpoints is called *planar*.

It is known that the min-max arc equality is satisfied for planar digraphs [48], for fully reducible flow digraphs [54], and for a special family of Eulerian digraphs [55]. Therefore, by Proposition 5.2.10, the min-max vertex equality is satisfied for the line digraphs of the members of these families. In summary, we have the following.

Corollary 5.2.12. *The circuit-packing bound is tight for the following families of digraphs: connectively reducible digraphs, digraphs that pack, line digraphs of planar digraphs, line digraphs of fully reducible flow digraphs, and line digraphs of special Eulerian digraphs.*

Definition 5.2.13. A digraph is called *partially planar* if all of its strongly connected components are planar.

Since the strongly connected components of a planar digraph are also planar, a planar digraph is partially planar. However, the converse is not always true, as

shown in Fig. 5.5.

Fig. 5.5: A partially planar digraph that is not planar

Proposition 5.2.14. *There is a polynomial time algorithm to recognize the line digraph of a partially planar digraph and subsequently determine its minrank.*

Proof.

1. Recognition Phase:

There is a one-to-one correspondence between the set of strongly connected components of order at least two of $\mathcal{D}$ and the set of strongly connected components of $\mathcal{L}(\mathcal{D})$ in the following sense. If $\mathcal{D}_i$'s, $i \in [k]$, are all strongly connected components of $\mathcal{D}$ each of which contains at least two vertices, then $\mathcal{L}(\mathcal{D}_i)$'s, $i \in [k]$, are all strongly connected components of $\mathcal{L}(\mathcal{D})$. Therefore, to determine whether a given digraph $\mathcal{L}$ is the line digraph of a partially planar digraph, it suffices to determine whether each of its strongly connected components $\mathcal{L}_i$ ($i \in [k]$) is the line digraph of a planar digraph. Note also that we can find all strongly connected components of a digraph in linear time [62].

For each $i \in [k]$, employing a polynomial time algorithm, we can determine

whether $\mathcal{L}_i$ is a line digraph of a digraph [60]. If the answer is YES, then the algorithm also outputs a digraph $\mathcal{D}'_i$, which is a root digraph of $\mathcal{L}_i$ and is strongly connected.

Suppose $\mathcal{L} = \mathcal{L}(\mathcal{D})$, where $\mathcal{D}$ is a digraph. Moreover, let $\mathcal{L}_i = \mathcal{L}(\mathcal{D}_i)$, where $\mathcal{D}_i$'s, $i \in [k]$, are all strongly connected components of $\mathcal{D}$ of order at least two. By [35, Theorem 3], $\mathcal{D}'_i$ and $\mathcal{D}_i$ are actually isomorphic, $i \in [k]$. Hence, to complete the Recognition Phase, one needs to test the planarity of $\mathcal{D}'_i$ for every $i \in [k]$. It is well known that this task can be done in linear time [38]. Thus, the Recognition Phase can be done in polynomial time.

2. MinRank Computation Phase:

Upon the completion of the Recognition Phase, if it is confirmed that $\mathcal{L}$ is indeed the line digraph of a partially planar digraph, then the second phase is carried out to compute $\text{minrk}_q(\mathcal{L})$. We show that this phase can also be done in polynomial time. Indeed, by Lemma 4.1.2, it suffices to show that $\text{minrk}_q(\mathcal{L}_i)$ for $i \in [k]$ can be found in polynomial time.

On the one hand, since $\mathcal{D}'_i$ (which is isomorphic to $\mathcal{D}_i$) is planar, as shown by Lucchesi and Younger [48], $\nu_1(\mathcal{D}'_i) = \tau_1(\mathcal{D}'_i)$. Therefore, by Proposition 5.2.10,

$$\text{minrk}_q(\mathcal{L}_i) = |\mathcal{V}(\mathcal{L}_i)| - \nu_0(\mathcal{L}_i) = |\mathcal{E}(\mathcal{D}'_i)| - \nu_1(\mathcal{D}'_i).$$

On the other hand, $\nu_1(\mathcal{D}'_i)$ can be computed in polynomial time [47]. Therefore $\text{minrk}_q(\mathcal{L}_i)$ for each $i \in [k]$ can be computed in polynomial time. Thus, $\text{minrk}_q(\mathcal{L})$ can be found in polynomial time. ■

In summary, we have the following.

Corollary 5.2.15. *There are polynomial time algorithms to recognize a member and subsequently determine the minrank of that member of the following families of digraphs: connectively reducible digraphs (which includes fully reducible flow digraphs and cyclically reducible digraphs), and line digraphs of partially planar digraphs.*

Note that since a strongly connected component of a fully reducible flow digraph may no longer be a fully reducible flow digraph, the arguments in the proof of Proposition 5.2.14 do not work for line digraphs of fully reducible flow digraphs.

We now discuss vector index codes and their transmission rates. Consider an ICSI instance $(n, n, \mathcal{X}, f)$. We treat each message x_i ($i \in [n]$) as a vector in $\mathbb{F}_q^t$ for some $t \geq 1$. An encoding function $\mathfrak{E} : \mathbb{F}_q^{tn} \rightarrow \mathbb{F}_q^k$ that enables each receiver R_i to recover x_i based on $\mathfrak{E}(\mathbf{x})$ and its side information is called a vector index code of block length t . We call the ratio k/t the *transmission rate* of the vector index code (see [2, 28]). An interesting task is to find an *optimal* vector index code, that is, a vector index code that achieves minimum transmission rate. When the block length t equals one, the vector index codes become the scalar index codes, and the aforementioned problem becomes the familiar problem of finding a scalar index code of minimum length.

Consider an ICSI instance described by a digraph $\mathcal{D}$. In the remainder of this section, we show that if $\text{minrk}_q(\mathcal{D}) = \alpha(\mathcal{D})$ then an optimal scalar linear index code (block length $t = 1$) always achieves the minimum transmission rate. As a consequence, we obtain several families of side information digraphs, listed in Corollary 5.2.12, for ICSI instances described by which, scalar linear index codes are already optimal. Before this work, only perfect graphs and acyclic digraphs are known to have this property.

Proposition 5.2.16. *If $\text{minrk}_q(\mathcal{D}) = \alpha(\mathcal{D})$ then the optimal transmission rate of an*

ICSI instance described by $\mathcal{D}$ is $\alpha(\mathcal{D})$. As a result, scalar linear index codes achieve the optimal rate.

Proof. First, observe that a vector index code for an ICSI instance described by $\mathcal{D}$ with block length $t \geq 1$ is a scalar index code for another ICSI instance described by $\mathcal{D}_t$ and vice versa, where $\mathcal{D}_t = (\mathcal{V}(\mathcal{D}_t), \mathcal{E}(\mathcal{D}_t))$ is defined as follows. The vertex set is $\mathcal{V}(\mathcal{D}_t) = \{(i, r) : i \in [n], r \in [t]\}$, and the arc set is

$$\mathcal{E}(\mathcal{D}_t) = \left\{ ((i, r), (j, s)) : (i, j) \in \mathcal{E}(\mathcal{D}), r, s \in [t] \right\}.$$

Note that in the ICSI instance described by $\mathcal{D}_t$, there are nt receivers, each requests precisely one bit of information. There are now n groups of receivers, in each of which all t receivers own the same side information. Each of these groups of t receivers corresponds to one receiver in the original instance.

Suppose that $\mathcal{I} = \{i_1, i_2, \dots, i_{\alpha(\mathcal{D})}\}$ induces a maximum acyclic subgraph of $\mathcal{D}$. By the definition of $\mathcal{D}_t$, the set

$$\mathcal{I}_t = \left\{ (i_j, r) : j \in [\alpha(\mathcal{D})], r \in [t] \right\}$$

induces an acyclic subgraph of $\mathcal{D}_t$. Therefore

$$\alpha(\mathcal{D}_t) \geq |\mathcal{I}_t| = |t\alpha(\mathcal{D})|.$$

Hence, by [3, Theorem 6], any scalar index code for the ICSI instance described by $\mathcal{D}_t$ has length at least $t\alpha(\mathcal{D})$. Therefore, any vector index code for the ICSI instance described by $\mathcal{D}$ has transmission rate at least $\alpha(\mathcal{D}) = (t\alpha(\mathcal{D}))/t$, which is equal to the rate of an optimal scalar linear index code. We recall that a straightforward generalization of [4, Theorem 5] shows that $\text{minrk}_q(\mathcal{D})$ ($= \alpha(\mathcal{D})$ in this case) equals

the rate of an optimal scalar linear index code over $\mathbb{F}_q$. The proof is complete. ■

5.3 On MinRanks of Graphs Having Tree Structures of Type I

We present in this section a new family of graphs whose minranks can be found in polynomial time.

5.3.1 Tree Structure of Type I

Let $\mathcal{P}$ be a collection of finitely many families of graphs that satisfy the following requirements:

1. Each family is closed under the operation of taking induced subgraphs, that is, every induced subgraph of a member of a family of $\mathcal{P}$ also belongs to that family;
2. There is a polynomial time algorithm to recognize a member of each family;
3. There is a polynomial time algorithm to find the minrank of every member of each family.

For instance, we may choose $\mathcal{P}$ to be the collection of the following three families: perfect graphs [3, 19], outerplanar graphs [8, 68], and graphs of orders bounded by a constant. Instead of saying that a graph $\mathcal{G}$ belongs to a family in $\mathcal{P}$, with a slight abuse of notation, we often simply say that $\mathcal{G} \in \mathcal{P}$.

Let U and V be two disjoint nonempty sets of vertices of $\mathcal{G}$. Let

$$s_{\mathcal{G}}(U, V) = \left| \left\{ \{u, v\} : u \in U, v \in V, \{u, v\} \in \mathcal{E}(\mathcal{G}) \right\} \right|,$$

denotes the number of edges each of which has one endpoint in U and the other endpoint in V .

Definition 5.3.1. A connected graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ is said to have a *tree structure of Type I* if there exists a partition $\Gamma = [V_1, V_2, \dots, V_k]$ of the vertex set $\mathcal{V}(\mathcal{G})$ that satisfies the following three requirements:

(R1) The V_i -induced subgraph $\mathcal{G}_i$ of $\mathcal{G}$ belongs to a family in $\mathcal{P}$, for every $i \in [k]$;

(R2) $s_{\mathcal{G}}(V_i, V_j) \in \{0, 1\}$ for every $i \neq j$;

(R3) The graph $T = (\mathcal{V}(T), \mathcal{E}(T))$, where $\mathcal{V}(T) = [k]$ and

$$\mathcal{E}(T) = \left\{ \{i, j\} : s_{\mathcal{G}}(V_i, V_j) = 1 \right\},$$

is a rooted tree; The tree T can also be thought of as a graph obtained from $\mathcal{G}$ by contracting each $\mathcal{V}_i$ to a single vertex.

The 2-tuple $\mathcal{T} = (\Gamma, T)$ is called a *tree structure of Type I* of $\mathcal{G}$.

Fig. 5.6: A tree structure of Type I of a graph $\mathcal{G}$

If a tree structure of Type I $\mathcal{T} = (\Gamma, T)$ of $\mathcal{G}$ is given, where $\Gamma = [V_1, V_2, \dots, V_k]$, then we can define the following terms:

1. Each V_i -induced subgraph $\mathcal{G}_i$ of $\mathcal{G}$ is called a node of $\mathcal{T}$;

2. If i is the parent of j in T , then G_i is called the parent (node) of G_j in $\mathcal{T}$;
We also refer to G_j as a child (node) of G_i ; Nodes with no children are called leaves; The node with no parent is called the root;
3. For each $i \in [k]$ let $\mathcal{S}_i$ be the subgraph of $\mathcal{G}$ induced by $V_i \cup (\cup_{j \in \text{des}_T(i)} V_j)$,
where $\text{des}_T(i)$ denotes the set of descendants of i in T ;
4. If G_j is a child of G_i , and $\{u, v\} \in \mathcal{E}(\mathcal{G})$, where $u \in V_i$ and $v \in V_j$, then u is called a *downward connector* (DC) of G_i and v is called the *upward connector* (UC) of G_j ; Each node may have several DCs but at most one UC; DCs and UC of a node are called *connectors* of that node.
5. Let $\text{mdc}(\mathcal{T})$ denote the maximum number of DCs of a node of $\mathcal{T}$.

For any $c > 0$ we define the following two families of connected graphs

$$\mathcal{F}_1(c) \triangleq \left\{ \mathcal{G} : \mathcal{G} \text{ is connected and has a tree structure } \mathcal{T} = (\Gamma, T) \text{ of Type I} \right. \\ \left. \text{satisfying } \text{mdc}(\mathcal{T}) \leq c \right\},$$

and

$$\mathcal{F}_2(c) \triangleq \left\{ \mathcal{G} : \mathcal{G} \text{ is connected and has a tree structure } \mathcal{T} = (\Gamma, T) \text{ of Type I} \right. \\ \left. \text{satisfying } \text{mdc}(\mathcal{T}) \leq c \text{ and each node of } \mathcal{T} \text{ is 2-edge connected} \right\},$$

A tree structure of Type I of a graph $\mathcal{G}$ that proves the membership of $\mathcal{G}$ in $\mathcal{F}_1(c)$ or $\mathcal{F}_2(c)$ is called a *relevant tree structure* of $\mathcal{G}$.

5.3.2 An Algorithm for MinRanks of Graphs in $\mathcal{F}_1(c)$

In this section we establish that the minrank of a member of $\mathcal{F}_1(c)$ can be found in polynomial time.

Theorem 5.3.2. *Let $c > 0$ be a constant and $\mathcal{G} \in \mathcal{F}_1(c)$. Suppose further that a tree structure of Type I $\mathcal{T} = (\Gamma, T)$ of $\mathcal{G}$ with $\text{mdc}(\mathcal{T}) \leq c$ is known. Then there is an algorithm that computes the minrank of $\mathcal{G}$ in polynomial time.*

To prove Theorem 5.3.2, we describe below an algorithm that computes the minrank of $\mathcal{G}$ when $\mathcal{G} \in \mathcal{F}_1(c)$ and investigate its complexity.

First, we introduce some notation which is used throughout this section. If v is any vertex of a graph $\mathcal{G}$, then $\mathcal{G} - v$ denotes the graph obtained from $\mathcal{G}$ by removing v and all edges incident to v . In general, if V is any set of vertices, then $\mathcal{G} - V$ denotes the graph obtained from $\mathcal{G}$ by removing all vertices in V and all edges incident to any vertex in V . In other words, $\mathcal{G} - V$ is the subgraph of $\mathcal{G}$ induced by $\mathcal{V}(\mathcal{G}) \setminus V$. Note that if $\mathcal{G} \in \mathcal{P}$ then the minrank of $\mathcal{G} - V$ can be computed in polynomial time for every subset $V \subset \mathcal{V}(\mathcal{G})$. If $e = \{u, v\} \in \mathcal{E}(\mathcal{G})$ then $\mathcal{G} - e$ denotes the graph obtained from $\mathcal{G}$ by removing the edge e without removing its endpoints u and v . In general, if E is any set of edges, then $\mathcal{G} - E$ denotes the graph obtained from $\mathcal{G}$ by removing all edges in E without removing their endpoints. The union of two or more graphs is a graph whose vertex set and edge set are the unions of the vertex sets and of the edge sets of the original graphs, respectively.

The following results from [8] are particularly useful in our discussion. Their proofs can be found in the full version of [8] at <http://www.openu.ac.il/home/mikel/papers/outer.pdf>.

Lemma 5.3.3 ([8]). *Let v be a vertex of a graph $\mathcal{G}$. Then*

$$\text{minrk}_q(\mathcal{G} - v) \leq \text{minrk}_q(\mathcal{G}) \leq \text{minrk}_q(\mathcal{G} - v) + 1.$$

Lemma 5.3.4 ([8]). *Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be two graphs with one common vertex v . Then*

$$\begin{aligned} \text{minrk}_q(\mathcal{G}_1 \cup \mathcal{G}_2) &= \text{minrk}_q(\mathcal{G}_1 - v) + \text{minrk}_q(\mathcal{G}_2 - v) \\ &\quad + (\text{minrk}_q(\mathcal{G}_1) - \text{minrk}_q(\mathcal{G}_1 - v))(\text{minrk}_q(\mathcal{G}_2) - \text{minrk}_q(\mathcal{G}_2 - v)). \end{aligned}$$

In other words, the minrank of $\mathcal{G}_1 \cup \mathcal{G}_2$ can be computed explicitly based on the minranks of $\mathcal{G}_1$, $\mathcal{G}_1 - v$, $\mathcal{G}_2$, and $\mathcal{G}_2 - v$.

ALGO-1:

Suppose $\mathcal{G} \in \mathcal{F}_1(c)$ and a relevant tree structure of $\mathcal{G}$ is given. The algorithm computes the minrank by dynamic programming in a bottom-up manner, from the leaves of $\mathcal{T}$ to its root. Suppose that $\Gamma = [V_1, V_2, \dots, V_k]$ and $\mathcal{G}_i$ is induced by V_i for $i \in [k]$. Let v_i be the UC (if any) of $\mathcal{G}_i$ for $i \in [k]$. For each i , ALGO-1 maintains a table which contains the two values, namely, minranks of $\mathcal{S}_i$ and $\mathcal{S}_i - v_i$. The minrank of the latter is omitted if $\mathcal{G}_i$ is the root node of $\mathcal{T}$. A key point is that the minranks of $\mathcal{S}_i$ and $\mathcal{S}_i - v_i$ can be computed in polynomial time from the minranks of $\mathcal{S}_j$'s and of $(\mathcal{S}_j - v_j)$'s where $\mathcal{G}_j$'s are children of $\mathcal{G}_i$, and from the minranks of at most 2^c subgraphs of $\mathcal{G}_i$. Each of these subgraphs is obtained from $\mathcal{G}_i$ by removing a subset of a set that consists of at most c vertices of $\mathcal{G}$. When the minrank of $\mathcal{S}_{i_0}$ is determined, where $\mathcal{G}_{i_0}$ is the root of $\mathcal{T}$, the minrank of $\mathcal{G}$ is found.

At the leaf-nodes:

Suppose $\mathcal{G}_i$ is a leaf and v_i is its UC. Since $\mathcal{G}_i$ has no children, $\mathcal{S}_i \equiv \mathcal{G}_i$. Hence,

$$\text{minrk}_q(\mathcal{S}_i) = \text{minrk}_q(\mathcal{G}_i),$$

and

$$\text{minrk}_q(\mathcal{S}_i - v_i) = \text{minrk}_q(\mathcal{G}_i - v_i).$$

Both of these values can be computed in polynomial time, as $\mathcal{G}_i \in \mathcal{P}$.

At the intermediate nodes:

Suppose the minranks of $\mathcal{S}_j$ and $\mathcal{S}_j - v_j$ are known for all j such that $\mathcal{G}_j$ is a child of $\mathcal{G}_i$. The goal of the algorithm at this step is to compute the minranks of $\mathcal{S}_i$ and $\mathcal{S}_i - v_i$ in polynomial time. It is complicated to analyze directly the general case where $\mathcal{G}_i$ has an arbitrary number (at most c) of downward connectors. Therefore, we first consider a special case where $\mathcal{G}_i$ has only one downward connector (Case 1). The results established in this case are then used to investigate the general case (Case 2).

Case 1: $\mathcal{G}_i$ has only one DC u and has r children, namely $\mathcal{G}_{j_1}, \mathcal{G}_{j_2}, \dots, \mathcal{G}_{j_r}$, all of which are connected to $\mathcal{G}_i$ via u .

Let $\mathcal{K}$ be the subgraph of $\mathcal{G}$ induced by the following set of vertices

$$\mathcal{V}(\mathcal{K}) = \mathcal{V}(\mathcal{S}_{j_1}) \cup \mathcal{V}(\mathcal{S}_{j_2}) \cup \dots \cup \mathcal{V}(\mathcal{S}_{j_r}) \cup \{u\}.$$

Notice that the graphs $\mathcal{G}_i$ and $\mathcal{K}$ share exactly one vertex, namely, u . Hence by Lemma 5.3.4, once the minranks of $\mathcal{G}_i$, $\mathcal{G}_i - u$, $\mathcal{K}$, and $\mathcal{K} - u$ are known, the minrank of $\mathcal{S}_i = \mathcal{G}_i \cup \mathcal{K}$ can be explicitly computed. Similarly, if $v_i \neq u$ and the minranks of $\mathcal{G}_i - v_i$, $\mathcal{G}_i - v_i - u$, $\mathcal{K}$, and $\mathcal{K} - u$ are known, the minrank of $\mathcal{S}_i - v_i = (\mathcal{G}_i - v_i) \cup \mathcal{K}$ can be explicitly computed. Observe also that if $v_i \equiv u$

Fig. 5.7: $\mathcal{G}_i$ has only one downward connector

then by Lemma 4.1.1,

$$\text{minrk}_q(\mathcal{S}_i - v_i) = \text{minrk}_q(\mathcal{G}_i - u) + \text{minrk}_q(\mathcal{K} - u).$$

Again by Lemma 4.1.1,

$$\text{minrk}_q(\mathcal{K} - u) = \sum_{\ell=1}^r \text{minrk}_q(\mathcal{S}_{j_\ell}),$$

which is known. Moreover, as $\mathcal{G}_i \in \mathcal{P}$, the minranks of $\mathcal{G}_i$, $\mathcal{G}_i - v_i$, $\mathcal{G}_i - u$, and $\mathcal{G}_i - v_i - u$ can be determined in polynomial time. Therefore it remains to compute the minrank of $\mathcal{K}$ efficiently. According to the following claim, the minrank of $\mathcal{K}$ can be explicitly computed based on the knowledge of the minranks of $\mathcal{S}_{j_\ell}$ and $\mathcal{S}_{j_\ell} - v_{j_\ell}$ for $\ell \in [r]$.

Claim 5.3.5. *The minrank of $\mathcal{K}$ is equal to*

$$\begin{cases} \text{minrk}_q(\mathcal{K} - u), & \text{if } \exists h \in [r] \text{ s.t. } \text{minrk}_q(\mathcal{S}_{j_h} - v_{j_h}) = \text{minrk}_q(\mathcal{S}_{j_h}) - 1, \\ \text{minrk}_q(\mathcal{K} - u) + 1, & \text{otherwise.} \end{cases}$$

Proof. Suppose there exists $h \in [r]$ such that

$$\text{minrk}_q(\mathcal{S}_{j_h} - v_{j_h}) = \text{minrk}_q(\mathcal{S}_{j_h}) - 1.$$

By Lemma 5.3.3,

$$\text{minrk}_q(\mathcal{K}) \geq \text{minrk}_q(\mathcal{K} - u).$$

Therefore, in this case it suffices to show that a matrix that fits $\mathcal{K}$ and has rank equal to $\text{minrk}_q(\mathcal{K} - u)$ exists. Indeed, such a matrix $\mathbf{M}$ can be constructed as follows:

1. Its sub-matrix restricted on $\mathcal{V}(\mathcal{S}_{j_\ell})$ for $\ell \neq h$ has rank equal to $\text{minrk}_q(\mathcal{S}_{j_\ell})$;
2. Its sub-matrix restricted on $\mathcal{V}(\mathcal{S}_{j_h}) \setminus \{v_{j_h}\}$ has rank equal to $\text{minrk}_q(\mathcal{S}_{j_h} - v_{j_h})$;
3. $\mathbf{M}_u = \mathbf{M}_{v_{j_h}} = \mathbf{e}_u + \mathbf{e}_{v_{j_h}}$;
4. All other entries are zero.

Since the sets $\mathcal{V}(\mathcal{S}_{j_\ell})$ ($\ell \neq h$), $\mathcal{V}(\mathcal{S}_{j_h}) \setminus \{v_{j_h}\}$, and $\{u, v_{j_h}\}$ are pairwise disjoint, the above requirements can be met without any contradiction arising.

Moreover,

$$\begin{aligned}
 \text{rank}_q(\mathbf{M}) &= \sum_{\ell \neq h} \text{minrk}_q(\mathcal{S}_{j_\ell}) + \text{minrk}_q(\mathcal{S}_{j_h} - v_{j_h}) + 1 \\
 &= \sum_{\ell \neq h} \text{minrk}_q(\mathcal{S}_{j_\ell}) + \text{minrk}_q(\mathcal{S}_{j_h}) \\
 &= \text{minrk}_q(\mathcal{K} - u).
 \end{aligned}$$

We now suppose that $\text{minrk}_q(\mathcal{S}_{j_\ell} - v_{j_\ell}) = \text{minrk}_q(\mathcal{S}_{j_\ell})$ for all $\ell \in [r]$. We prove that

$$\text{minrk}_q(\mathcal{K}) = \text{minrk}_q(\mathcal{K} - u) + 1$$

by induction on r .

1. The base case: $r = 1$.

Fig. 5.8: The base case when $r = 1$

Let $\mathcal{J} = (\mathcal{V}(\mathcal{J}), \mathcal{E}(\mathcal{J}))$ where $\mathcal{V}(\mathcal{J}) = \{u, v_{j_1}\}$ and $\mathcal{E}(\mathcal{J}) = \{\{u, v_{j_1}\}\}$.

Then $\mathcal{K} = \mathcal{J} \cup \mathcal{S}_{j_1}$ and $\mathcal{V}(\mathcal{J}) \cap \mathcal{V}(\mathcal{S}_{j_1}) = \{v_{j_1}\}$. Moreover,

$$\text{minrk}_q(\mathcal{J}) = \text{minrk}_q(\mathcal{J} - v_{j_1}) = 1.$$

Therefore by Lemma 5.3.4,

$$\begin{aligned}
 \text{minrk}_q(\mathcal{K}) &= \text{minrk}_q(\mathcal{S}_{j_1} - v_{j_1}) + \text{minrk}_q(\mathcal{I} - v_{j_1}) \\
 &= \text{minrk}_q(\mathcal{S}_{j_1}) + 1 \\
 &= \text{minrk}_q(\mathcal{K} - u) + 1.
 \end{aligned}$$

2. The inductive step: suppose that the assertion holds for $r \geq 1$. We aim to show that it also holds for $r + 1$.

Fig. 5.9: The inductive step

Let $\mathcal{J}$ be the subgraph of $\mathcal{G}$ induced by

$$\{u\} \cup \left(\bigcup_{\ell=1}^r \mathcal{V}(\mathcal{S}_{j_\ell}) \right).$$

Since $\text{minrk}_q(\mathcal{S}_{j_\ell} - v_{j_\ell}) = \text{minrk}_q(\mathcal{S}_{j_\ell})$ for all $\ell \in [r]$, by the induction

hypothesis, we have

$$\text{minrk}_q(\mathcal{J}) = \text{minrk}_q(\mathcal{J} - u) + 1.$$

Let $\mathcal{I}$ be the subgraph of $\mathcal{G}$ induced by $\{u\} \cup \mathcal{V}(\mathcal{S}_{j_{r+1}})$. As

$$\text{minrk}_q(\mathcal{S}_{j_{r+1}} - v_{j_{r+1}}) = \text{minrk}_q(\mathcal{S}_{j_{r+1}}),$$

similar arguments as in the base case yield

$$\text{minrk}_q(\mathcal{I}) = \text{minrk}_q(\mathcal{I} - u) + 1.$$

Applying Lemma 5.3.4 to the graphs $\mathcal{I}$ and $\mathcal{J}$ we obtain

$$\begin{aligned} \text{minrk}_q(\mathcal{K}) &= \text{minrk}_q(\mathcal{I} \cup \mathcal{J}) \\ &= \text{minrk}_q(\mathcal{I} - u) + \text{minrk}_q(\mathcal{J} - u) + 1 \\ &= \sum_{\ell=1}^{r+1} \text{minrk}_q(\mathcal{S}_{j_\ell}) + 1, \end{aligned}$$

which is equal to $\text{minrk}_q(\mathcal{K} - u) + 1$. ■

Case 2: $\mathcal{G}_i$ has d DCs, namely, $u_1, u_2, \dots, u_d$. Let $\{\mathcal{G}_j : j \in I_t\}$ for $1 \leq t \leq d$, be the set of children of $\mathcal{G}_i$ connected to $\mathcal{G}_i$ via u_t .

Recall that the goal of the algorithm is to compute the minranks of $\mathcal{S}_i$ and $\mathcal{S}_i - v_i$ in polynomial time, given that the minranks of $\mathcal{S}_j$ and $\mathcal{S}_j - v_j$ are known for all children $\mathcal{G}_j$'s of $\mathcal{G}_i$.

For each $t \in [d]$ let $\mathcal{K}_t$ be the subgraph of $\mathcal{G}$ induced by the following set of

Fig. 5.10: $\mathcal{G}_i$ has several downward connectors

vertices

$$\{u_t\} \cup \left(\bigcup_{j \in I_t} \mathcal{V}(\mathcal{S}_j) \right).$$

As in Case 1, based on the minranks of $\mathcal{S}_j$'s and $\mathcal{S}_j - v_j$ for $j \in I_t$, it is possible to compute the minranks of $\mathcal{K}_t$ and $\mathcal{K}_t - u_t$ explicitly for all $t \in [d]$.

Let

$$\mathcal{N}_1 = \mathcal{K}_1 \cup \mathcal{G}_i,$$

and

$$\mathcal{N}_t = \mathcal{N}_{t-1} \cup \mathcal{K}_t,$$

for every $t \in [d]$ and $t \geq 2$. Observe that $\mathcal{N}_d \equiv \mathcal{S}_i$. Below we show how the algorithm computes the minranks of $\mathcal{N}_d$ and $\mathcal{N}_d - v_i$ inductively in polynomial time.

1. At the base case, the minranks of $\mathcal{N}_1 - U$, for every subset $U \subseteq \{v_i, u_2, u_3,$

$\dots, u_d\}$, are computed as follows.

If $v_i \equiv u_1 \in U$, then

$$\mathcal{N}_1 - U = (\mathcal{K}_1 - u_1) \cup (\mathcal{G}_i - U).$$

Since

$$\mathcal{V}(\mathcal{K}_1 - u_1) \cup \mathcal{V}(\mathcal{G}_i - U) = \emptyset,$$

by Lemma 4.1.1,

$$\text{minrk}_q(\mathcal{N}_1 - U) = \text{minrk}_q(\mathcal{K}_1 - u_1) + \text{minrk}_q(\mathcal{G}_i - U),$$

which is computable in polynomial time.

Suppose that either $v_i \not\equiv u_1$ or $v_i \notin U$. By Lemma 5.3.4, since

$$\mathcal{N}_1 - U = \mathcal{K}_1 \cup (\mathcal{G}_i - U),$$

and

$$\mathcal{V}(\mathcal{K}_1) \cap \mathcal{V}(\mathcal{G}_i - U) = \{u_1\},$$

the minrank of $\mathcal{N}_1 - U$ can be determined based on the minranks of $\mathcal{K}_1$, $\mathcal{K}_1 - u_1$, $\mathcal{G}_i - U$, and $\mathcal{G}_i - U - u_1$. The minranks of these graphs are either known or computable in polynomial time. As $\text{mdc}(\mathcal{T}) \leq c$, there are at most $2^d \leq 2^c$ (a constant) such subsets U . Hence, the total computation in the base case can be done in polynomial time.

2. At the inductive step, suppose that the minrank of $\mathcal{N}_{t-1} - U$, $t \geq 2$, for every subset $U \subseteq \{v_i, u_t, u_{t+1}, \dots, u_d\}$ are known. Our goal is to show that the minrank of $\mathcal{N}_t - V$ for every subset $V \subseteq \{v_i, u_{t+1}, u_{t+2}, \dots, u_d\}$

can be determined in polynomial time. Observe again that there are at most 2^c such subsets V .

If $v_i \equiv u_t$ and $v_i \in V$, then

$$\mathcal{N}_t - V = (\mathcal{N}_{t-1} - V) \cup (\mathcal{K}_t - u_t).$$

Moreover, as we have

$$\mathcal{V}(\mathcal{N}_{t-1} - V) \cap \mathcal{V}(\mathcal{K}_t - u_t) = \emptyset,$$

by Lemma 4.1.1,

$$\text{minrk}_q(\mathcal{N}_t - V) = \text{minrk}_q(\mathcal{N}_{t-1} - V) + \text{minrk}_q(\mathcal{K}_t - u_t),$$

which is known.

Suppose that either $v_i \not\equiv u_t$ or $v_i \notin V$. Since

$$\mathcal{N}_t - V = (\mathcal{N}_{t-1} - V) \cup \mathcal{K}_t,$$

and

$$\mathcal{V}(\mathcal{N}_{t-1} - V) \cap \mathcal{V}(\mathcal{K}_t) = \{u_t\},$$

the minrank of $\mathcal{N}_t - V$ can be computed based on the minranks of $\mathcal{N}_{t-1} - V$, $\mathcal{N}_{t-1} - V - u_t$, $\mathcal{K}_t$, and $\mathcal{K}_t - u_t$, which are all available from the previous inductive step.

When the inductive process reaches $t = d$, the minranks of $\mathcal{N}_d$ and $\mathcal{N}_d - v_i$ are found, as desired. The analysis of Case 2 is completed.

When the algorithm reaches the root node of $\mathcal{T}$, the minrank of $\mathcal{G}$ is found.

5.3.3 An Algorithm to Recognize a Graph in $\mathcal{F}_2(c)$

In order for ALGO-1 to work, it is assumed that a relevant tree structure of the input graph $\mathcal{G} \in \mathcal{F}_1(c)$ is given. Therefore, the next question is how to design an algorithm that recognizes a graph in that family and subsequently finds a relevant tree structure for that graph in polynomial time. Unfortunately, we are unable to provide such an algorithm at the moment. However, the same task is possible for a sub-family of $\mathcal{F}_1(c)$, namely, $\mathcal{F}_2(c)$.

Theorem 5.3.6. *Let $c > 0$ be any constant. Then there is a polynomial time algorithm that recognizes a member of $\mathcal{F}_2(c)$. Moreover, this algorithm also outputs a relevant tree structure of that member.*

In order to prove Theorem 5.3.6, we introduce ALGO-2 (Fig. 5.11). We show below that ALGO-2 indeed does what it is claimed to do.

Claim 5.3.7. *If ALGO-2 terminates successfully, that is, it terminates without any error message printed out, then $\mathcal{G} \in \mathcal{F}_2(c)$ and the output is a relevant tree structure of $\mathcal{G}$.*

Proof. Suppose that ALGO-2 terminates successfully. In the Splitting Phase, the algorithm first splits $\mathcal{G}$ into two vertex-disjoint components (subgraphs) that are connected to each other by exactly one edge in $\mathcal{G}$. It then keeps splitting the existing components, whenever possible, each into two new smaller (vertex-disjoint) components that are connected to each other by exactly one edge in the original

ALGO-2:

Input: A connected graph $\mathcal{G} = (\mathcal{V}(\mathcal{G}), \mathcal{E}(\mathcal{G}))$ and a constant $c > 0$.

Output: If $\mathcal{G} \in \mathcal{F}_2(c)$, the algorithm prints out a confirmation message, namely “ $\mathcal{G} \in \mathcal{F}_2(c)$ ”, and then returns a relevant tree structure of $\mathcal{G}$. Otherwise, it prints out an error message “ $\mathcal{G} \notin \mathcal{F}_2(c)$ ”.

Initialization: Create two empty queues, $\mathcal{Q}_1$ and $\mathcal{Q}_2$, which contains graphs as their elements. Push $\mathcal{G}$ into $\mathcal{Q}_1$.

Splitting Phase:

```

while  $\mathcal{Q}_1 \neq \emptyset$  do
  for  $\mathcal{A} = (\mathcal{V}(\mathcal{A}), \mathcal{E}(\mathcal{A})) \in \mathcal{Q}_1$  do
    if there exist  $U$  and  $V$  that partition  $\mathcal{V}(\mathcal{A})$  and  $s_{\mathcal{A}}(U, V) = 1$  then
      Let  $\mathcal{B}$  and  $\mathcal{C}$  be subgraphs of  $\mathcal{A}$  induced by  $U$  and  $V$ , respectively
      Push  $\mathcal{B}$  and  $\mathcal{C}$  into  $\mathcal{Q}_1$ 
    else if  $\mathcal{A} \in \mathcal{P}$  then
      Push  $\mathcal{A}$  into  $\mathcal{Q}_2$ 
    else
      Print the error message “ $\mathcal{G} \notin \mathcal{F}_2(c)$ ” and exit
    end if
  end for
end while

```

Verifying Phase:

Suppose $\mathcal{Q}_2$ contains k' graphs $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{k'}$. The number of connectors of each graphs $\mathcal{A}_i$ for $i \in [k']$ is computed. If $\mathcal{A}_i$ has at most $c + 1$ connectors for every i then print out the message “ $\mathcal{G} \in \mathcal{F}_2(c)$ ” and output the vertex partition $[\mathcal{V}(\mathcal{A}_1), \dots, \mathcal{V}(\mathcal{A}_{k'})]$ together with the rooted tree $T' = (\mathcal{V}(T'), \mathcal{E}(T'))$ constructed as follows. Let $\mathcal{V}(T') = [k']$ and $\mathcal{E}(T') = \{\{i, j\} : s_{\mathcal{G}}(\mathcal{V}(\mathcal{A}_i), \mathcal{V}(\mathcal{A}_j)) = 1\}$. The root of T' is i_0 , where $\mathcal{A}_{i_0}$ is any node that has at most c connectors. Otherwise, print out the error message “ $\mathcal{G} \notin \mathcal{F}_2(c)$ ” and exit.

Fig. 5.11: ALGO-2

Fig. 5.12: Splitting Phase of ALGO-2

component (see Fig. 5.12). A straightforward inductive argument shows the following

1. Throughout the Splitting Phase, the vertex sets that induce the components of $\mathcal{G}$ partition $\mathcal{V}(\mathcal{G})$; Hence $\mathcal{V}(\mathcal{A}_i)$'s, $i \in [k']$, partition $\mathcal{V}(\mathcal{G})$;
2. Throughout the Splitting Phase, any two different components of $\mathcal{G}$ are connected to each other by at most one edge in $\mathcal{G}$; Therefore, $s_{\mathcal{G}}(\mathcal{V}(\mathcal{A}_i), \mathcal{V}(\mathcal{A}_j)) \in \{0, 1\}$ for every $i \neq j$;
3. At any time during the Splitting Phase, the graph that is obtained from $\mathcal{G}$ by contracting the vertex set of each component of $\mathcal{G}$ to a single vertex is a tree; Therefore, T' is a tree;
4. In the Splitting Phase, every component of $\mathcal{G}$ is connected; Hence, as there are no bridges in each of the subgraph $\mathcal{A}_i$ for $i \in [k']$, they are all 2-edge connected.

It is also clear that each $\mathcal{A}_i$ belongs to a family in $\mathcal{P}$. Since $\mathcal{G}$ passes the Verifying Phase successfully, $\mathcal{T}' = (\Gamma' = [\mathcal{V}(\mathcal{A}_1), \dots, \mathcal{V}(\mathcal{A}_{k'})], T')$ is already qualified to be a tree structure of Type I of $\mathcal{G}$. It remains to show that $\text{mdc}(\mathcal{T}') \leq c$. Indeed, firstly, the root node $\mathcal{A}_{i_0}$ has at most c connectors, hence at most c downward connectors. Secondly, the other nodes $\mathcal{A}_i$ for $i \neq i_0$, each has a parent node and has at most $c + 1$ connectors. Hence, each of them has at most c DCs. Hence $\text{mdc}(\mathcal{T}') \leq c$, as desired. ■

Claim 5.3.8. *If $\mathcal{G} \in \mathcal{F}_2(c)$ then ALGO-2 terminates successfully.*

Proof. Suppose $\mathcal{G} \in \mathcal{F}_2(c)$ and $\mathcal{T} = (\Gamma, T)$, where $\Gamma = [V_1, V_2, \dots, V_k]$, is a relevant tree structure of $\mathcal{G}$. As each $\mathcal{G}_i$ is 2-edge connected, throughout the Splitting Phase,

the splits always occur at and only at the edges of $\mathcal{G}$ each of which connects two different nodes of $\mathcal{T}$. The Splitting Phase stops when and only when there are no more components that contain such an edge within each of them. Therefore, at the end of the Splitting Phase, the set of components of $\mathcal{G}$ generated by algorithm, namely $\{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{k'}\}$, coincide with the set of nodes $\{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_k\}$ of $\mathcal{T}$. The Verifying Phase also terminates successfully. Indeed, firstly, since each node $\mathcal{G}_j$ ($j \in [k]$) has at most c downward connectors in $\mathcal{T}$, the node $\mathcal{A}_i$ ($i \in [k]$) has at most $c + 1$ connectors in $\mathcal{T}'$. Secondly, a node $\mathcal{A}_{i_0}$ that has at most c connectors always exists since we can choose, for instance, the root node of $\mathcal{T}$. ■

Claim 5.3.9. *The running time of ALGO-2 is polynomial with respect to the order of $\mathcal{G}$.*

Proof. Notice that every single task in ALGO-2 can be accomplished in polynomial time. Those tasks include: finding a min-cut of size one in a graph, deciding whether a graph belongs to $\mathcal{P}$, counting the number of connectors of each component of $\mathcal{G}$, and building a rooted tree based on the components of $\mathcal{G}$.

Let consider the “while loop” and the “for loop”. After each intermediate iteration in the while loop, as at least one component gets split into two smaller components, the number of components of $\mathcal{G}$ is increased by at least one. Since the vertex sets of the components are pairwise disjoint, there are no more than $n = |\mathcal{V}(\mathcal{G})|$ components at any time. Hence, there are no more than n iterations in the while loop. Since the number of graphs in $\mathcal{Q}_1$ cannot exceed n , the number of iterations in the for loop is also at most n . Thus, in total, the running time of ALGO-2 is polynomial with respect to n . ■

5.4 MinRanks of (Di)Graphs of Small Orders

To aid further research on the behavior of minranks of graphs, we have carried out a computation of binary minranks of *all* non-isomorphic graphs of orders up to 10.

Order	Number of non-isomorphic graphs	Total running time
1	1	< 1 seconds
2	2	< 1 seconds
3	4	< 1 seconds
4	11	< 1 seconds
5	34	< 1 seconds
6	156	< 1 seconds
7	1,044	< 1 seconds
8	12,346	25 seconds
9	274,668	56 minutes
10	12,005,168	4.3 days

Fig. 5.13: Running time for finding minranks of graphs of small orders

A reduction to the Satisfiability (SAT) problem [16] provides us with an elegant method to compute the binary minrank of a (hyper,di)graph. We observed that while the SAT-based approach is very efficient for (di)graphs having many edges, it does not perform well for simple instances, such as a graph on 10 vertices with no edges (minrank 10), or a digraph on seven vertices $0, 1, \dots, 6$, with six arcs $(0, 1), (1, 2), (2, 3), (3, 4), (4, 5), (5, 6)$ (minrank seven). For such naive instances, the SAT solver that we used, Minisat [24], was not able to terminate after hours. This weirdness might be attributable to the fact that the SAT instances corresponding to a (di)graph with fewer (arcs) edges contain more variables than those corresponding to a (di)graphs with more (arcs) edges on the same set of vertices (see the Appendix A.2 for more details).

To achieve our goal, we wrote a sub-program which used a Branch-and-Bound algorithm to find minranks in an exhaustive manner. When the input graph was

of large size, that is, its size surpassed a given threshold, a sub-program using a SAT solver was invoked; Otherwise, the Branch-and-Bound sub-program was used. We noticed that there were graphs of order 10 that have around 21–22 edges, for which the Branch-and-Bound sub-program could find the minranks in less than one second, while the SAT-based sub-program could not after 3 or 4 hours. For graphs of order 10, we observed that the threshold 24, which we actually used, did work well. The most time-consuming task is to compute the minranks of all 12,005,618 non-isomorphic graphs of order 10. This task actually took more than four days to finish.

We put all details of the SAT-based approach [16] and the Branch-and-Bound approach in Appendix A. The minranks of all non-isomorphic graphs of orders at most six are listed in Appendix B for illustration. The minranks and the corresponding matrices that achieve the minranks of all non-isomorphic graphs of orders up to 10 are available at web.spms.ntu.edu.sg/~daus0001/mr-small-graphs.html. Computation of the minrank of a single (di)graph is available at web.spms.ntu.edu.sg/~daus0001/mr.html.

6. CONCLUSION

We have studied the security aspects and the computational aspects of the Index Coding with Side Information (ICSI) problem.

For *block security*, given a linear index code based on a matrix $\mathbf{L}$, we have established two bounds on the security level of the code. These bounds employ the minimum distance and the dual distance of $\mathcal{C}(\mathbf{L})$, the subspace spanned by the (transposed) columns of $\mathbf{L}$. While the dimension of this subspace, which is equal to the number of columns of $\mathbf{L}$ in this setting, corresponds to the number of transmissions in the scheme (that is, the efficiency of the code), the minimum distance determines the security of the scheme.

Open Problem I: Find a trade-off between the efficiency and the block security of a linear index code.

We now elaborate more on this problem. Consider an ICSI instance $(m, n, \mathcal{X}, f)$. Let $\mathcal{W}_i \triangleq \{c(\mathbf{u}^{(i)} + \mathbf{e}_{f(i)}) : c \in \mathbb{F}_q^*, \mathbf{u}^{(i)} \triangleleft \mathcal{X}_i\}$ for $i \in [m]$. We have established (in Corollary 2.2.6) that $\mathbf{L}$ corresponds to a linear $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$ if and only if $\mathcal{C}(\mathbf{L}) \cap \mathcal{W}_i \neq \emptyset$ for all $i \in [m]$. Here $\mathcal{C}(\mathbf{L})$ denotes the (transposed) column space of $\mathbf{L}$. Then the Open Problem I can be restated as follows: given an instance $(m, n, \mathcal{X}, f)$, what is the trade-off between the dimension and the minimum distance of a vector space $\mathcal{C} \subseteq \mathbb{F}_q^n$ that satisfies $\mathcal{C} \cap \mathcal{W}_i \neq \emptyset$ for all $i \in [m]$?

Open Problem II: Find constructions of families of linear index codes that possess good efficiency and good block security.

At one extreme, when only the efficiency is considered, it is known that designing an optimal linear index code is an NP-hard problem. At the other extreme, if we ignore the efficiency, we may ask the following question.

Open Problem III: How hard is it to find a linear index code that provides the highest level of block security?

We discuss below a different view of the Open Problem III. Let us consider the original setting of the ICSI problem [3], in which $q = 2$, $m = n$, and $f(i) = i$ for all $i \in [n]$. Let $\mathbf{A} = (a_{i,j})_{n \times n}$ be the binary *side information matrix*, defined by $a_{i,j} = 1$ if and only if either $j = i \in [n]$ or $j \in \mathcal{X}_i$. Equivalently, $\mathbf{A}$ is obtained by taking the sum of the adjacency matrix of the side information digraph and the identity matrix. It is not hard to see that the task specified in the Open Problem III is equivalent to the task of finding a way to turn certain off-diagonal 1's into 0's in $\mathbf{A}$, so that the rows of the resulting matrix generate an error-correcting code of minimum distance as large as possible. Note that the row space of such a matrix corresponds to the transposed column space of a matrix $\mathbf{L}$ associated with a linear index code. It is very likely that this task is a hard problem. For comparison, even finding the minimum distance of an error-correcting code given its generating matrix corresponds to an NP-complete decision problem [65].

For *strong security*, we have introduced a lower bound on the length of a linear index code that is resistant to eavesdropping, information leaking, and errors. Index codes that achieve this bound have been constructed for sufficiently large alphabets.

Open Problem IV: Investigate the strong security of linear index codes over small alphabets, such as the binary alphabet.

It seems that one cannot obtain the perfect strong security for linear index codes over small alphabets. However, we may consider the asymptotic strong security instead. In that case, vector linear index codes with block lengths approaching

infinity are more suitable. A valid question is: does the concatenation of binary wiretap codes (studied by Ozarow and Wyner [52]) and optimal binary vector linear index codes yield (asymptotically) optimal security levels and optimal transmission rates? Note that this question has been answered in the affirmative for the case of wiretap codes and scalar linear index codes that are over sufficiently large alphabets (Chapter 2 and 3).

We have also constructed a number of bounds on the length of an optimal error-correcting index code. As it is shown in Example 3.2.10, a separation of error-correcting code and index code sometimes leads to a non-optimal scheme. This raises a question of designing coding schemes in which the two layers are treated as a whole. We have discussed a general decoding procedure for linear error-correcting index codes. The difference between decoding of a classical error-correcting code and decoding of an error-correcting index code is that in the latter case, each receiver does not require a complete knowledge of the error vector. This difference may help to ease the decoding process. We state our sixth open problem below.

Open Problem VI: Find constructions of error-correcting index codes with good parameters and efficient decoding methods.

It has been shown that the minrank of the side information hypergraph is actually the length of an optimal scalar linear index code. However, it is, in general, extremely hard to compute the minrank of a hypergraph. Therefore, it would be desirable to find families of graphs, digraphs, and hypergraphs whose minranks can be determined in polynomial time. The second main contribution of this thesis is the revelation of some new families of (di)graphs for which the minrank computation can be done in polynomial time. We state below several open problems related to minranks of hypergraphs ((di)graphs).

Open Problem VII: Examine the hardness of the following decision problem:

decide whether a digraph has minrank two over a *nonbinary* field $\mathbb{F}_q$.

Open Problem VIII: Find a polynomial time algorithm to recognize a member of $\mathcal{F}_1(c)$ (see Section 5.3).

Open Problem IX: Find efficient exact algorithms for computing minrank of a hypergraph over a *nonbinary* field $\mathbb{F}_q$.

Open Problem X: Find new families of hypergraphs ((di)graphs) whose minranks can be found in polynomial time.

BIBLIOGRAPHY

- [1] R. Ahlswede, N. Cai, S. Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, 46:1204–1216, 2000.
- [2] N. Alon, A. Hassidim, E. Lubetzky, U. Stav, and A. Weinstein. Broadcasting with side information. In *Proc. 49th Annu. IEEE Symp. on Found. of Comput. Sci. (FOCS)*, pages 823–832, 2008.
- [3] Z. Bar-Yossef, Z. Birk, T. S. Jayram, and T. Kol. Index coding with side information. In *Proc. 47th Annu. IEEE Symp. on Found. of Comput. Sci. (FOCS)*, pages 197–206, 2006.
- [4] Z. Bar-Yossef, Z. Birk, T. S. Jayram, and T. Kol. Index coding with side information. *IEEE Trans. Inform. Theory*, 57(3):1479–1494, 2011.
- [5] R. Bellman. The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60:503–516, 1954.
- [6] C. H. Bennet, G. Brassard, and J. M. Robert. Privacy amplification by public discussion. *SIAM J. Computing*, 17:210–229, 1988.
- [7] E. R. Berlekamp, R. J. McEliece, and H. C. A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Trans. Inform. Theory*, IT-24(3):384–386, 1978.

-
- [8] Y. Berliner and M. Langberg. Index coding with outerplanar side information. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 869–873, Saint Petersburg, Russia, 2011.
 - [9] K. Bhattad and K. R. Narayanan. Weakly secure network coding. In *Proc. 1st Workshop on Network Coding, Theory, and Applicat. (NetCod)*, 2005.
 - [10] Y. Birk and T. Kol. Informed-source coding-on-demand (ISCOD) over broadcast channels. In *Proc. IEEE Conf. on Comput. Commun. (INFOCOM)*, pages 1257–1264, San Francisco, CA, 1998.
 - [11] Y. Birk and T. Kol. Coding-on-demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients. *IEEE Trans. Inform. Theory*, 52(6):2825–2830, 2006.
 - [12] N. Cai and R. W. Yeung. Secure network coding. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Lausanne, Switzerland, 2002.
 - [13] C. Carlet. *Vectorial Boolean Functions for Cryptography*. In the book *Boolean Models and Methods in Mathematics, Computer Science and Engineering*, edited by Y. Crama and P. L. Hammer. Cambridge University Press, 2010.
 - [14] M. A. R. Chaudhry, Z. Asad, A. Sprintson, and M. Langberg. Finding sparse solutions for the index coding problem. In *Proc. IEEE Globecom*, pages 1–5, Houston, Texas, USA, 2011.
 - [15] M. A. R. Chaudhry, Z. Asad, A. Sprintson, and M. Langberg. On the complementary index coding problem. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 306–310, Saint Petersburg, Russia, 2011.

-
- [16] M. A. R. Chaudhry and A. Sprintson. Efficient algorithms for index coding. In *Proc. IEEE Conf. on Comput. Commun. (INFOCOM)*, pages 1–4, 2008.
 - [17] B. Chor, O. Goldreich, J. Hastad, J. Freidmann, S. Rudich, and R. Smolensky. The bit extraction problem or t -resilient functions. In *Proc. 26th Annu. IEEE Symp. on Found. of Comput. Sci. (FOCS)*, pages 396–407, 1985.
 - [18] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas. The strong perfect graph theorem. *Annals of Mathematics*, 164:51–229, 2006.
 - [19] G. Cornuejols, X. Liu, and K. Vuskovic. A polynomial time algorithm for recognizing perfect graphs. *Proc. 48th Annu. IEEE Symp. on Found. of Comput. Sci. (FOCS)*, pages 20–27, 2003.
 - [20] T. M. Cover and J. A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, 1991.
 - [21] S. H. Dau, V. Skachek, and Y. M. Chee. Index coding and error correction. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 1753–1757, Saint Petersburg, Russia, 2011.
 - [22] S. H. Dau, V. Skachek, and Y. M. Chee. On secure index coding with side information. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 1046–1050, Saint Petersburg, Russia, 2011.
 - [23] C. Ding, R. Laihonon, and A. Renvall. Linear multisecret-sharing schemes and error-correcting codes. *J. UCS*, 3(9):1023–1036, 1997.
 - [24] N. Eén and N. Sörensson. An extensible SAT-solver. In *Proc. 6th Int. Conf. Theory Appl. Satisf. Testing*, 2003.

-
- [25] S. El Rouayheb, M. A. R. Chaudhry, and A. Sprintson. On the minimum number of transmissions in single-hop wireless coding networks. In *Proc. IEEE Inform. Theory Workshop (ITW)*, pages 120–125, 2007.
 - [26] S. El Rouayheb and E. Soljanin. On wiretap networks II. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 551–555, Nice, France, 2007.
 - [27] S. El Rouayheb, A. Sprintson, and C. Georghiades. On the relation between the index coding and the network coding problems. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 1823–1827, Toronto, Canada, 2008.
 - [28] S. El Rouayheb, A. Sprintson, and C. Georghiades. On the index coding problem and its relation to network coding and matroid theory. *IEEE Trans. Inform. Theory*, 56(7):3187–3195, 2010.
 - [29] J. Feldman, T. Malkin, C. Stein, and R. A. Servedio. On the capacity of secure network coding. In *Proc. Annu. Allerton Conf. Commun., Control, and Comput.*, 2004.
 - [30] A. Frank and A. Gyárfas. Directed graphs and computer programs. *Problèmes Combinatoires et Théorie des Graphes*, (260):157–158, 1976.
 - [31] M. Grassl. Bounds on the minimum distance of linear codes and quantum codes. available online at <http://www.codetables.de>.
 - [32] B. Guenin and R. Thomas. Packing directed circuits exactly. Manuscript, 2001.
 - [33] K.C. Gupta and P. Sarkar. Improved construction of nonlinear resilient s-boxes. *IEEE Trans. Inform. Theory*, 51(1):339 – 348, 2005.
 - [34] W. Haemers. An upper bound for the Shannon capacity of a graph. *Algebr. Methods Graph Theory*, 25:267–272, 1978.

-
- [35] F. Harary and R. Norman. Some properties of line digraphs. *Rendiconti del Circolo Matematico Palermo*, 9:149–163, 1960.
 - [36] I. Haviv and M. Langberg. On linear index coding for random graphs. <http://arxiv.org/abs/1107.0390>, 2011.
 - [37] A. S. Hedayat, N. J. A. Sloane, and J. Stufken. *Orthogonal Arrays: Theory and Applications*. Springer-Verlag, New York, 1999.
 - [38] J. Hopcroft and R. Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, 1974.
 - [39] W. Huffman and V. Pless. *Fundamentals of Error-Correcting Codes*. Cambridge University Press, 2003.
 - [40] R. M. Karp. Reducibility among combinatorial problems. *Complexity Comput. Comput.*, 40(4):85–103, 1972.
 - [41] S. Katti, D. Katabi, H. Balakrishnan, and M. Médard. Symbol-level network coding for wireless mesh networks. *ACM SIGCOMM Comput. Commun. Review*, 38(4):401–412, 2008.
 - [42] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft. Xors in the air: Practical wireless network coding. In *Proc. ACM SIGCOMM*, pages 243–254, 2006.
 - [43] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Trans. Netw.*, 11:782–795, 2003.
 - [44] M. Langberg and A. Sprintson. On the hardness of approximating the network coding capacity. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 315–319, Toronto, Canada, 2008.

-
- [45] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25:1–7, 1979.
 - [46] E. Lubetzky and U. Stav. Non-linear index coding outperforming the linear optimum. *Proc. 48th Annu. IEEE Symp. on Found. of Comput. Sci. (FOCS)*, pages 161–168, 2007.
 - [47] C. Lucchesi. *A minimax equality for directed graphs*. PhD thesis, University of Waterloo, Ontario, 1976.
 - [48] C. Lucchesi and D. Younger. A minimax theorem for directed graphs. *J. London Math. Soc.*, 17(2):369–374, 1978.
 - [49] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. Amsterdam: North-Holland, 1977.
 - [50] J. L. Massey. Minimal codewords and secret sharing. In *Proc. Joint Swedish-Russian Int. Workshop on Inform. Theory*, pages 276–279, 1993.
 - [51] B. D. McKay. Practical graph isomorphism. *Congressus Numerantium*, 30:45–87, 1981.
 - [52] L. H. Ozarow and A. D. Wyner. The wire-tap channel II. *Bell Syst. Tech. J.*, 63:2135–2157, 1984.
 - [53] R. Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica*, 16(3):417–431, 1996.
 - [54] V. Ramachandran. A minimax arc theorem for reducible flow graphs. *SIAM J. Discrete Math.*, pages 554–560, 1990.

-
- [55] P. D. Seymour. Packing circuits in Eulerian digraphs. *Combinatorica*, 16(2):223–231, 1996.
- [56] A. Shamir. A linear time algorithm for finding minimum cutsets in reducible graphs. *J. Comput.*, 8(4):645–655, 1979.
- [57] C. E. Shannon. The zero-error capacity of a noisy channel. *IRE Trans. Inform. Theory*, 3:3–15, 1956.
- [58] D. Silva and F. R. Kschischang. Universal weakly secure network coding. In *Proc. Inform. Theory Workshop on Networking and Inform. Theory (ITW)*, pages 281–285, 2009.
- [59] D. Silva and F. R. Kschischang. Universal secure error-correcting schemes for network coding. In *Proc. IEEE Symp. on Inform. Theory (ISIT)*, pages 2428–2432, Austin, Texas, USA, 2010.
- [60] M. Syslo. A labeling algorithm to recognize a line digraph and output its root graph. *Inform. Process. Lett.*, 15(1):28–30, 1982.
- [61] J. L. Szwarcfiter. On minimum cuts of cycles and maximum disjoint cycles. In *Graphs and algorithms*, vol. 89 of *Contemp. Math.*, 153–166. Amer. Math. Soc., Providence, RI, 1989.
- [62] R. E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
- [63] R. E. Tarjan. Testing flow graph reducibility. *J. Comput. System Sci.*, (9):355–365, 1974.

-
- [64] G. S. Tseitin. *On the complexity of derivation in propositional calculus*. In the book *Studies in Constructive Mathematics and Mathematical Logic, Part 2*, edited by A. O. Slisenko. Consultants Bureau, New York, 1970.
- [65] A. Vardy. The intractability of computing the minimum distance of a code. *IEEE Trans. Inform. Theory*, 43(6):1757–1766, 1997.
- [66] C.-C. Wang, E. L. Lloyd, and M. L. Soffa. Feedback vertex sets and cyclically reducible graphs. *J. ACM*, 32(2):296–313, 1985.
- [67] D. B. West. *Introduction to graph theory*. Prentice Hall, 2 edition, 1996.
- [68] M. Wiegiers. Recognizing outerplanar graphs in linear time. In *Int. Workshop WG '86 on Graph-Theoretic Concepts in Comput. Sci.*, pages 165–176, 1987.
- [69] Y. Wu, J. Padhye, R. Chandra, V. Padmanabhan, and P. A. Chou. The local mixing problem. In *Proc. Inform. Theory and Applicat. Workshop*, San Diego, CA, 2006.
- [70] X-M. Zhang and Y. Zheng. On nonlinear resilient functions. In *Proc. 14th Annu. Int. Conf. on Theory and Appl. of Cryptographic Tech. (EUROCRYPT)*, pages 274–288, 1995.
- [71] Z. Zhuang, Y. Luo, and A. J. Han Vinck. Secure error-correcting network codes with side information from source. In *Proc. Int. Conf. Commun. Intell. Inform. Secur.*, pages 55–59, 2010.

APPENDIX

A. AN EXACT ALGORITHM FOR FINDING MINRANKS OF (DI)GRAPHS OVER $\mathbb{F}_2$

A.1 *The Algorithm*

We described below the algorithm for digraphs. For graphs, it is completely the same.

```

Main()
{
  Input a digraph  $\mathcal{D}$  and a threshold  $\theta$ ;
  if  $|\mathcal{E}(\mathcal{D})| \leq \theta$  then
    Find-MinRank-Using-BranchAndBound( $\mathcal{D}$ );
  else
    Find-MinRank-Using-SAT-Solver( $\mathcal{D}$ );
  end if
}

```

Fig. A.1: Main module

The Branch-and-Bound (BB) module runs through the set of all matrices that fit $\mathcal{D}$ (possibly ignores some) and finds out a matrix that has smallest minrank. Suppose $\mathcal{D}$ is of order n .

The BB module builds up a search tree as follows. The nodes of the tree are grouped into $n + 1$ different layers. The root node belongs to Layer 0. The children of the root node belong to Layer 1. The children of each of the children of the root node belong to Layer 2, and so forth. Each node in Layer i ($i \in [n]$) is labeled by a row vector $\mathbf{u} = (u_1, u_2, \dots, u_n) \in \mathbb{F}_2^n$ satisfying $u_i = 1$ and $u_j = 0$ if $(i, j) \notin \mathcal{E}(\mathcal{D})$.

In other words, the label of each node in Layer i is a candidate for the i th row of a matrix that fits $\mathcal{D}$. The sets of labels of the children of each node in Layer $i - 1$ ($2 \leq i \leq n$) is precisely the set of all such row vectors $\mathbf{u}$. The n labels of the nodes on an arbitrary path from the root to a leaf (excluding the root) form the rows of a matrix that fits $\mathcal{D}$. Traversing all such paths of the search tree, we can find all matrices that fit $\mathcal{D}$. The following example illustrates the search tree described above.

Fig. A.2: An example of the search tree built by the BB module

The BB module during its running time maintains a global bound, which is a “current best” candidate for the minrank of $\mathcal{D}$. This bound is the rank of a matrix that fits $\mathcal{D}$ and has lowest rank among all previously examined matrices. The global bound is updated every time a matrix with lower rank is examined. Based on this bound, the algorithm may prune the branches (i.e., it bypasses those branches and goes upwards to examine other branches) that surely do not contain matrices with lower ranks than the current best rank. Hence, the search space can be reduced significantly. At the beginning, we can initiate this bound to be any upper bound

for $\text{minrk}_2(\mathcal{D})$, found by some fast heuristic algorithm. Some candidates are those established by Chaudhry and Sprintson [16]. In the case of graphs, we actually used a simple upper bound for $\text{cc}(\mathcal{G})$ (hence it is also an upper bound for $\text{minrk}_2(\mathcal{G})$), which was obtained by coloring $\overline{\mathcal{G}}$ greedily. Note that we should choose different heuristic algorithms in different scenarios. Indeed, if we want to compute the min-rank of only one (di)graph, an algorithm that outputs a good bound in half a minute seems to be fine. However, if we are computing the minranks of millions of (di)graphs, such an algorithm would be a very poor choice. In such a case, an algorithm that quickly produces a not-so-good bound might be preferred.

```

Find-MinRank-Using-SAT-Solver( $\mathcal{D}$ )
{
  lower-bound := 1; upper-bound :=  $|\mathcal{V}(\mathcal{D})|$ ;
  while upper-bound > lower-bound + 1 do
     $k := \lfloor (\text{upper-bound} + \text{lower-bound})/2 \rfloor$ ;
    Create the corresponding SAT instance  $I_k(\mathcal{D})$ ;
    sat := SAT-Solver( $I_k(\mathcal{D})$ );
    if sat = SATISFIABLE then
      upper-bound :=  $k$ ;
    else
      lower-bound :=  $k$ ;
    end if
  end while
  Print out upper-bound as the minrank of  $\mathcal{D}$ ;
}

```

Fig. A.3: SAT-based module

The issue of creating the SAT instance $I_k(\mathcal{D})$ is discussed in details in the next section.

A.2 Reduction to SAT Problem

In this section, we discuss a reduction of the ICSI problem to the SAT problem (Fig. A.3). The reduction is attributed to Chaudhry and Sprintson [16]. Although the reduction works for ICSI instances described by hypergraphs, we here restrict ourselves to the case of (di)graphs, and provide more details on the implementation of the reduction.

Suppose the digraph $\mathcal{D}$ describes the following ICSI instance. There are n receivers $R_1, R_2, \dots, R_n$, and n messages $x_1, x_2, \dots, x_n$ in $\mathbb{F}_2$. The receiver R_i ($i \in [n]$) possesses $\{x_j\}_{j \in \mathcal{X}_i}$ as side information, and requests x_i . Given a value $k > 0$, based on a SAT solver, we want to decide whether the sender can satisfy the demands from all receivers by broadcasting at most k linear combinations of the x_i 's.

Suppose the k encoding vectors are

$$\mathbf{g}_j = (g_{1,j}, g_{2,j}, \dots, g_{n,j}) \in \mathbb{F}_2^n, \quad j \in [k].$$

The sender broadcasts k bits

$$\mathbf{x}\mathbf{g}_j^T = \sum_{i=1}^n x_i g_{i,j}, \quad j \in [k].$$

Suppose R_i ($i \in [n]$) uses the decoding vector

$$\mathbf{q}_i = (q_{i,1}, q_{i,2}, \dots, q_{i,k}) \in \mathbb{F}_2^k.$$

He applies the decoding vector to the received bits

$$\mathbf{q}_i \left(\mathbf{x}\mathbf{g}_1^T, \mathbf{x}\mathbf{g}_2^T, \dots, \mathbf{x}\mathbf{g}_k^T \right)^T = \sum_{r=1}^n \left(\sum_{j=1}^k g_{r,j} q_{i,j} \right) x_r.$$

The receiver R_i can retrieve x_i successfully if and only if

$$\begin{cases} \sum_{j=1}^k g_{i,j} q_{i,j} = 1, \\ \sum_{j=1}^k g_{r,j} q_{i,j} = 0, \quad \forall r \notin \{i\} \cup \mathcal{X}_i. \end{cases}$$

If we regard 0 as FALSE and 1 as TRUE, then the summation and the product in $\mathbb{F}_2$ are equivalent to the XOR $\oplus$ and the AND $\wedge$ operators in Boolean logic, respectively. Therefore, it is possible for the sender to satisfy all demands with k linear combinations of the messages if and only if the following Boolean formula is satisfiable

$$B_k \triangleq \bigwedge_{i=1}^n \left\{ \left[\bigoplus_{j=1}^k (g_{i,j} \wedge q_{i,j}) \right] \wedge \left\{ \bigwedge_{r \notin \{i\} \cup \mathcal{X}_i} \left[\bigoplus_{j=1}^k (g_{r,j} \wedge q_{i,j}) \right] \right\} \right\}. \quad (\text{A.1})$$

Most of the current SAT solvers require the input to be in CNF (Conjunctive Normal Form), that is, the input formula should be a conjunction of clauses

$$C_1 \wedge C_2 \wedge \cdots \wedge C_s,$$

where each clause C_i is a disjunction of literals

$$C_i = L_{i,1} \vee L_{i,2} \vee \cdots \vee L_{i,r},$$

and each literal is either a plain variable or a negated variable.

Any Boolean formula can be transformed into an equivalent formula, that is, a formula with the same set of variables and the same satisfiability property, based on rules about logical equivalences. However, a straightforward transformation can lead to an exponential explosion of the formula, that is, the number of clauses is

increased exponentially. To avoid the exponential blowup in the size (number of variables plus number of clauses) of the formula, the Tseitin transformation is often used. The Tseitin transformation [64] transforms an arbitrary Boolean formula into a new one such that not only the satisfiability is preserved but also the size of the new formula is linear with respect to the size of the original one. Given a formula ϕ , the idea is to introduce a fresh variable τ_θ for every sub-formula θ of ϕ . Then one adds constraints in CNF which ensure that $\tau_\theta \longleftrightarrow \theta$ holds. Finally, one adds $\wedge \tau_\phi$ to the formula to make sure that the satisfiability is preserved.

The Tseitin transformations of the basic formulas that we need are

$$\begin{aligned}\tau &\longleftrightarrow (a \wedge b) \rightsquigarrow (a \vee \bar{\tau}) \wedge (b \vee \bar{\tau}) \wedge (\bar{a} \vee \bar{b} \vee \tau), \\ \tau &\longleftrightarrow (\bar{a}) \rightsquigarrow (a \vee \tau) \wedge (\bar{a} \vee \bar{\tau}), \\ \tau &\longleftrightarrow (a \oplus b) \rightsquigarrow (a \vee b \vee \bar{\tau}) \wedge (\bar{a} \vee \bar{b} \vee \bar{\tau}) \wedge (a \vee \bar{b} \vee \tau) \wedge (\bar{a} \vee b \vee \tau).\end{aligned}\tag{A.2}$$

For each $i \in [n]$ we consider the following formula

$$\bigoplus_{j=1}^k (g_{i,j} \wedge q_{i,j}) = \underbrace{\underbrace{(g_{i,1} \wedge q_{i,1})}_{h_{i,1}} \oplus \underbrace{(g_{i,2} \wedge q_{i,2})}_{h_{i,2}}}_{s_{i,1}} \oplus \underbrace{(g_{i,3} \wedge q_{i,3})}_{h_{i,3}} \oplus \cdots \oplus \underbrace{(g_{i,k} \wedge q_{i,k})}_{h_{i,k}}.$$

$$\underbrace{\hspace{15em}}_{s_{i,2}}$$

$$\underbrace{\hspace{15em}}_{\vdots}$$

$$\underbrace{\hspace{15em}}_{s_{i,k-1}}$$

By introducing fresh variables $h_{i,j}$'s and $s_{i,j}$'s for the sub-formulas of the formula

above, we obtain a new formula

$$\begin{aligned} & (h_{i,1} \longleftrightarrow (g_{i,1} \wedge q_{i,1})) \wedge \cdots \wedge (h_{i,k} \longleftrightarrow (g_{i,k} \wedge q_{i,k})) \\ & \wedge (s_{i,1} \longleftrightarrow (h_{i,1} \oplus h_{i,2})) \wedge (s_{i,2} \longleftrightarrow (s_{i,1} \oplus h_{i,3})) \wedge \cdots \wedge (s_{i,k-1} \longleftrightarrow (s_{i,k-2} \oplus h_{i,k})) \\ & \wedge s_{i,k-1}. \end{aligned}$$

We now apply (A.2) to the above formula and obtain a formula in CNF

$$\begin{aligned} C_i \triangleq & (g_{i,1} \vee \overline{h_{i,1}}) \wedge (q_{i,1} \vee \overline{h_{i,1}}) \wedge (\overline{g_{i,1}} \vee \overline{q_{i,1}} \vee h_{i,1}) \\ & \wedge (g_{i,2} \vee \overline{h_{i,2}}) \wedge (q_{i,2} \vee \overline{h_{i,2}}) \wedge (\overline{g_{i,2}} \vee \overline{q_{i,2}} \vee h_{i,2}) \\ & \vdots \\ & \wedge (g_{i,k} \vee \overline{h_{i,k}}) \wedge (q_{i,k} \vee \overline{h_{i,k}}) \wedge (\overline{g_{i,k}} \vee \overline{q_{i,k}} \vee h_{i,k}) \\ & \wedge (h_{i,1} \vee h_{i,2} \vee \overline{s_{i,1}}) \wedge (\overline{h_{i,1}} \vee \overline{h_{i,2}} \vee \overline{s_{i,1}}) \wedge (h_{i,1} \vee \overline{h_{i,2}} \vee s_{i,1}) \wedge (\overline{h_{i,1}} \vee h_{i,2} \vee s_{i,1}) \\ & \wedge (h_{i,3} \vee s_{i,1} \vee \overline{s_{i,2}}) \wedge (\overline{h_{i,3}} \vee \overline{s_{i,1}} \vee \overline{s_{i,2}}) \wedge (h_{i,3} \vee \overline{s_{i,1}} \vee s_{i,2}) \wedge (\overline{h_{i,3}} \vee s_{i,1} \vee s_{i,2}) \\ & \vdots \\ & \wedge (h_{i,k} \vee s_{i,k-2} \vee \overline{s_{i,k-1}}) \wedge (\overline{h_{i,k}} \vee \overline{s_{i,k-2}} \vee \overline{s_{i,k-1}}) \wedge (h_{i,k} \vee \overline{s_{i,k-2}} \vee s_{i,k-1}) \\ & \wedge (\overline{h_{i,k}} \vee s_{i,k-2} \vee s_{i,k-1}) \\ & \wedge s_{i,k-1}. \end{aligned}$$

Similarly, for each $i \in [n]$ and $r \notin \{i\} \cup \mathcal{X}_i$ we consider the following formula

$$\bigoplus_{j=1}^k (g_{r,j} \wedge q_{i,j}) = \overbrace{\underbrace{(g_{r,1} \wedge q_{i,1})}_{t_{i,r,1}} \oplus \underbrace{(g_{r,2} \wedge q_{i,2})}_{t_{i,r,2}} \oplus \underbrace{(g_{r,3} \wedge q_{i,3})}_{t_{i,r,3}} \oplus \cdots \oplus \underbrace{(g_{r,k} \wedge q_{i,k})}_{t_{i,r,k}}}_{\underbrace{\quad}_{u_{i,r,1}}} \oplus \underbrace{\quad}_{u_{i,r,2}} \oplus \cdots \oplus \underbrace{\quad}_{u_{i,r,k-1}} \oplus \underbrace{\quad}_{u_{i,r,k}}.$$

By introducing fresh variables $t_{i,r,j}$'s and $u_{i,r,j}$'s to the sub-formulas of this formula, we obtain a new formula

$$\begin{aligned} D_{i,r} &\triangleq (t_{i,r,1} \longleftrightarrow (g_{r,1} \wedge q_{i,1})) \wedge \cdots \wedge (t_{i,r,k} \longleftrightarrow (g_{r,k} \wedge q_{i,k})) \\ &\quad \wedge (u_{i,r,1} \longleftrightarrow (t_{i,r,1} \oplus t_{i,r,2})) \wedge (u_{i,r,2} \longleftrightarrow (u_{i,r,1} \oplus t_{i,r,3})) \wedge \cdots \\ &\quad \quad \quad \wedge (u_{i,r,k-1} \longleftrightarrow (u_{i,r,k-2} \oplus t_{i,r,k})) \\ &\quad \wedge (u_{i,r,k} \longleftrightarrow \overline{u_{i,r,k-1}}) \\ &\quad \wedge u_{i,r,k}. \end{aligned}$$

We can rewrite $D_{i,r}$ in CNF as follows

$$\begin{aligned}
D_{i,r} = & (g_{r,1} \vee \overline{t_{i,r,1}}) \wedge (q_{i,1} \vee \overline{t_{i,r,1}}) \wedge (\overline{g_{r,1}} \vee \overline{q_{i,1}} \vee t_{i,r,1}) \\
& \wedge (g_{r,2} \vee \overline{t_{i,r,2}}) \wedge (q_{i,2} \vee \overline{t_{i,r,2}}) \wedge (\overline{g_{r,2}} \vee \overline{q_{i,2}} \vee t_{i,r,2}) \\
& \vdots \\
& \wedge (g_{r,k} \vee \overline{t_{i,r,k}}) \wedge (q_{i,k} \vee \overline{t_{i,r,k}}) \wedge (\overline{g_{r,k}} \vee \overline{q_{i,k}} \vee t_{i,r,k}) \\
& \wedge (t_{i,r,1} \vee t_{i,r,2} \vee \overline{u_{i,r,1}}) \wedge (\overline{t_{i,r,1}} \vee \overline{t_{i,r,2}} \vee \overline{u_{i,r,1}}) \wedge (t_{i,r,1} \vee \overline{t_{i,r,2}} \vee u_{i,r,1}) \\
& \qquad \qquad \qquad \wedge (\overline{t_{i,r,1}} \vee t_{i,r,2} \vee u_{i,r,1}) \\
& \wedge (t_{i,r,3} \vee u_{i,r,1} \vee \overline{u_{i,r,2}}) \wedge (\overline{t_{i,r,3}} \vee \overline{u_{i,r,1}} \vee \overline{u_{i,r,2}}) \wedge (t_{i,r,3} \vee \overline{u_{i,r,1}} \vee u_{i,r,2}) \\
& \qquad \qquad \qquad \wedge (\overline{t_{i,r,3}} \vee u_{i,r,1} \vee u_{i,r,2}) \\
& \vdots \\
& \wedge (t_{i,r,k} \vee u_{i,r,k-2} \vee \overline{u_{i,r,k-1}}) \wedge (\overline{t_{i,r,k}} \vee \overline{u_{i,r,k-2}} \vee \overline{u_{i,r,k-1}}) \wedge (t_{i,r,k} \vee \overline{u_{i,r,k-2}} \vee u_{i,r,k-1}) \\
& \qquad \qquad \qquad \wedge (\overline{t_{i,r,k}} \vee u_{i,r,k-2} \vee u_{i,r,k-1}) \\
& \wedge (u_{i,r,k-1} \vee u_{i,r,k}) \wedge (\overline{u_{i,r,k-1}} \vee \overline{u_{i,r,k}}) \\
& \wedge u_{i,r,k}.
\end{aligned}$$

We obtain the following formula (in CNF)

$$B'_k = \bigwedge_{i=1}^n \left\{ C_i \wedge \left(\bigwedge_{r \notin \{i\} \cup \mathcal{X}_i} D_{i,r} \right) \right\},$$

which is a Tseitin transformation of the formula B_k given in (A.1). Therefore, B'_k is satisfiable if and only if B_k is satisfiable. Thus, the sender can satisfy all requests from the receivers by sending out at most k linear combinations of the messages if and only if B'_k is satisfiable. Since B'_k is in CNF, its satisfiability can be tested using any SAT solver.

The formula B' has $n(2kn + 2k - 1) - 2k \sum_{i=1}^n |\mathcal{X}_i|$ variables and $n(7kn - n - 2) - (7k - 1) \sum_{i=1}^n |\mathcal{X}_i|$. Observe that $\sum_{i=1}^n |\mathcal{X}_i|$ is the number of arcs in the case of digraphs and is two times the number of edges in the case of graphs. Therefore, for two (di)graphs of the same order, the (di)graph with fewer (arcs) edges produces a CNF formula consisting of more variables. That might explain why the SAT-based algorithm seems to be not very good at dealing with (di)graphs having very few (arcs) edges.

A.3 Generating Lists of Non-Isomorphic Graphs

In order to compute the minranks of all graphs of certain order, we first need to produce a list of all non-isomorphic graphs of that order. The Nauty package, written by McKay [51], is a good tool for that purpose. Once the package (<http://cs.anu.edu.au/~bdm/nauty/>) is installed, we can first use the command (in Linux)

```
./geng n graphs-n-vertices.txt
```

to generate a list of all non-isomorphic graphs over n vertices and write the list to the file `graphs-n-vertices.txt`. If we also use the option `-l` then the graphs are canonically labeled. The graphs in this list are written in the compact graph6 format. If we want to write them in a readable format, we can call

```
./listg -q -e graphs-n-vertices.txt graphs-n-vertices-readable.txt
```

Then in the file `graphs-n-vertices-readable.txt`, each graph will be written in the following format

$$n \ m \ a_1 \ b_1 \ a_2 \ b_2 \ \cdots \ a_m \ b_m,$$

where n and m are the number of vertices and edges, respectively, and $\{a_i, b_i\}$ is the i th edge of the graph. For more information on how to use other functionalities of the Nauty package, please refer to its user guide.

B. MINRANKS OF GRAPHS OF SMALL ORDERS

B.1 MinRanks of Graphs of Orders One to Four

Graphs of minrank one:

Graphs of minrank two:

Graphs of minrank three:

Graphs of minrank four:

B.2 MinRanks of Graphs of Order Five

Graphs of minrank one:

Graphs of minrank two:

Graphs of minrank three:

Graphs of minrank four:

Graphs of minrank five:

B.3 MinRanks of Graphs of Order Six

Graphs of minrank one:

Graphs of minrank two:

Graphs of minrank three:

Graphs of minrank four:

Graphs of minrank five:

Graphs of minrank six:

INDEX

(ρ, δ) -Property, 83
2-edge connected, 17
$N_O^{\mathcal{D}}(u)$, 18
$N^{\mathcal{G}}(u)$, 18
$N_O(u)$, <i>see</i> $N_O^{\mathcal{D}}(u)$
$N_q(M, d)$, 14
$N_q[k, d]$, 14
$[n]$, 13
$\mathcal{F}_1(c)$, 129
$\mathcal{F}_2(c)$, 129
$\mathcal{F}(m, n, \mathcal{X}, \mathcal{Z}, f)$, 42
$\mathcal{G} - E$, 130
$\mathcal{G} - V$, 130
$\mathcal{G} - e$, 130
$\mathcal{G} - v$, 130
$\Gamma(n, \rho)$, 83
$\mathcal{I}(q, \mathcal{H})$, 56
$\mathcal{J}(\mathcal{H})$, 56
$\mathcal{L}_i$, 77
$\mathcal{N}_q[\delta, \mathcal{H}]$, 56
$\mathcal{P}$, 127
$\mathcal{T}$, 128 | $\mathcal{X}$, 22
$\mathcal{X}_A$, 26
$\mathcal{X}_i$, 22
$\mathcal{Y}_i$, 56
$\mathcal{Z}$, 42
$\mathcal{Z}_i$, 42
$\alpha(\mathcal{D})$, 62
$\alpha(\mathcal{G})$, 18
α -bound, 63, 94
$\alpha(\mathcal{H})$, 61
$\mathbf{e}_i$, 14
$\mathcal{C}(\mathbf{L})$, 25
$\mathbf{H}(\cdot \mid \cdot)$, <i>see</i> conditional entropy
$\mathbf{H}(\cdot)$, <i>see</i> entropy
$\mathfrak{J}(\Gamma)$, 82
$\mathbb{F}_q$, 12
$\mathbb{F}_q^*$, 12
κ -bound, 65
κ_q , 34
$\triangleleft$, 13
$\mathbf{mdc}(\mathcal{T})$, 129
$\mathbf{minrk}$, 19 |
|---|--|

-
- $\nu_0(\mathcal{D})$, 115
 - $\nu_1(\mathcal{D})$, 118
 - $\bar{\kappa}_q$, 43
 - $\mathcal{S}_i$, 129
 - $\mathbf{s}_{\mathcal{G}}(U, V)$, 127
 - span_q , 15
 - $\tau_0(\mathcal{D})$, 118
 - $\tau_1(\mathcal{D})$, 118
 - $\hat{\mathcal{X}}_A$, 28
 - $n(\Gamma)$, 82
 - $\mathcal{H}$ -IC over $\mathbb{F}_q$, *see* index code
 - $(m, n, \mathcal{X}, f)$ -IC over $\mathbb{F}_q$, *see* index code
 - acyclic, 17
 - adjacent, 15
 - adversary, 27
 - strength, 27
 - arc, 15
 - in-coming, 15
 - out-going, 15
 - bridge, 17
 - circuit, 17
 - circuit packing number, 115
 - circuit-packing bound, 115
 - clique, 18
 - clique cover, 18
 - minimum, 18
 - number, 18
 - clique-cover bound
 - digraph, 116
 - graph, 67
 - code
 - $(n, M, d)_q$, 14
 - $[n, k, d]_q$, 13
 - $[n, k]_q$, 13
 - dimension, 13
 - generator matrix, 13
 - length, 13
 - linear, 13
 - MDS, 14
 - parity-check matrix, 14
 - size, 14
 - codeword, 13
 - coloring, 100
 - fair, 100
 - conditional entropy, 20
 - confusable, 92
 - connected component, 17
 - connector
 - downward, 129
 - upward, 129
 - Construction

-
- A, 50
 - B, 75
 - cut, 17
 - cycle, 17
 - DC, *see* connector, downward
 - digraph, 15
 - complement, 16
 - complete, 16
 - connectively reducible, 119
 - fully reducible flow, 119
 - order, 15
 - partially planar, 122
 - planar, 122
 - size, 15
 - strongly connected, 17
 - symmetric, 16
 - that pack, 119
 - underlying graph, 16
 - distance
 - dual, 14
 - Hamming, 13
 - ECIC, 55
 - $(\delta, \mathcal{H})$ -, 55
 - optimal, 56
 - edge, 15
 - endpoint, 15
 - entropy, 19
 - fairly k -colorable, 101
 - feedback set
 - arc, 118
 - vertex, 118
 - fit, 18
 - generalized independent number, 61
 - graph, 15
 - bipartite, 16
 - bridgeless, 17
 - complement, 16
 - complete, 16
 - confusion, 92
 - connected, 17
 - corresponding digraph, 16
 - order, 16
 - perfect, 68
 - size, 16
 - star, 105
 - hyperarc, 19
 - hypergraph, 22
 - underlying digraph, 23
 - underlying graph, 24
 - IC, *see* index code

-
- ICSI instance, 22
 - ICSRI, 42
 - independence number, 18
 - independent set, 18
 - generalized, 61
 - maximum, 18
 - maximum generalized, 61
 - index code, 24
 - based on $\mathbf{L}$, 25
 - block length, 125
 - error-correcting, 55
 - length, 24
 - linear, 24
 - matrix corresponding to, 25
 - optimal, 24
 - randomized, 44
 - static, 82
 - transmission rate, 125
 - line digraph, 119
 - minrank
 - digraph, 19
 - graph, 19
 - hypergraph, 34
 - no information, 27
 - node, 128
 - child, 129
 - leaf, 129
 - parent, 129
 - root, 129
 - out-degree, 18
 - out-neighbor, 18
 - path, 17
 - Pr, 19
 - relevant error pattern, 77
 - relevant tree structure, 129
 - resilient function, 84
 - weakly, 85
 - security
 - block, 27
 - completely insecure, 28
 - strong, 44
 - weak, 27
 - side information
 - digraph, 23
 - graph, 24
 - hypergraph, 22
 - Singleton bound, 14, 69, 95
 - strongly connected component, 17
 - subgraph, 16

-
- induced, 17
 - maximum acyclic induced, 62
 - support, 13
 - tree, 18
 - rooted, 18
 - tree structure of Type I, 128
 - UC, *see* connector, upward
 - vertex, 15
 - child, 18
 - parent, 18
 - weight, 13