

# Redundancy Is All You Need (for CSP Sparsification)\*

Joshua Brakensiek<sup>†</sup>Venkatesan Guruswami<sup>‡</sup>

## Abstract

The seminal work of Benczúr and Karger demonstrated cut sparsifiers of near-linear size, with several applications throughout theoretical computer science. Subsequent extensions have yielded sparsifiers for hypergraph cuts and more recently linear codes over Abelian groups. A decade ago, Kogan and Krauthgamer asked about the sparsifiability of arbitrary constraint satisfaction problems (CSPs). For this question, a trivial lower bound is the size of a *non-redundant* CSP instance, which admits, for each constraint, an assignment satisfying only that constraint (so that no constraint can be dropped by the sparsifier). For instance, for graph cuts, spanning trees are non-redundant instances.

Our main result is that redundant clauses are sufficient for sparsification: for any CSP predicate  $R$ , every unweighted instance of  $\text{CSP}(R)$  has a sparsifier of size at most its non-redundancy (up to polylog and  $1/\epsilon$  factors). For weighted instances, we similarly pin down the sparsifiability to the so-called chain length of the predicate. These results precisely determine the extent to which any CSP can be sparsified.

Our result is established in the general setting of non-linear codes, or equivalently set families, yielding a VC-type theorem for multiplicative error approximation. This unifies and extends known sparsification results for graph cuts, linear codes, and broad CSP classes. A key technical ingredient in our work is a novel application of the entropy method from Gilmer’s recent breakthrough on the union-closed sets conjecture.

As an immediate consequence of our main theorem, a number of results in the non-redundancy literature immediately extend to CSP sparsification. We also contribute new techniques for understanding the non-redundancy of CSP predicates. By adapting methods from the matching vector codes literature in coding theory, we are able to construct an explicit predicate whose non-redundancy lies between  $\Omega(n^{1.5})$  and  $\tilde{O}(n^{1.6})$ , the first example with a provably non-integral exponent. The tools we introduce in this paper, including a conditional version of non-redundancy, have spurred several follow-ups on non-redundancy as well as new connections to extremal combinatorics, sparsification, and streaming.

---

\*An extended abstract of this work appeared in STOC 2025 [BG25].

<sup>†</sup>Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. Contact: josh.brakensiek@berkeley.edu

<sup>‡</sup>Simons Institute for the Theory of Computing, and the Departments of EECS and Mathematics, University of California, Berkeley. Contact: venkatg@berkeley.edu

# Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Introduction</b>  | <b>1</b>  |
| 1.1      | Non-linear Code Sparsification . . . . .   | 1         |
| 1.2      | CSP Sparsification . . . . .   | 3         |
| 1.3      | Weighted CSP Sparsification . . . . .  | 3         |
| 1.4      | Non-redundancy and Chain Length of Specific Relations . . . . .                    | 4         |
| 1.5      | Technical Overview . . . . .   | 5         |
| 1.6      | Related Work . . . . .   | 8         |
| 1.7      | Subsequent Work . . . . .  | 11        |
| 1.8      | Open Questions . . . . .   | 13        |
| <b>2</b> | <b>Preliminaries</b>   | <b>14</b> |
| 2.1      | CSPs . . . . .   | 14        |
| 2.2      | Non-linear Codes . . . . .   | 16        |
| 2.3      | Theorem 1.1 Implies Theorem 1.2 . . . . .  | 17        |
| 2.4      | Chernoff Bounds . . . . .  | 18        |
| <b>3</b> | <b>Warmup: A Simple Sparsifier</b>   | <b>19</b> |
| <b>4</b> | <b>The Entropy Sparsifier</b>  | <b>21</b> |
| 4.1      | OR-closure, VC Dimension, and Entropy . . . . .                                    | 21        |
| 4.2      | Gilmer’s Entropy Method . . . . .  | 23        |
| 4.3      | Improved Entropy Bound for $\theta$ -sparse Distributions . . . . .                | 25        |
| 4.4      | Minimax . . . . .  | 27        |
| 4.5      | Decomposition Theorem . . . . .  | 29        |
| 4.6      | Proof of Theorem 1.1 . . . . .   | 31        |
| <b>5</b> | <b>Extension to Weighted Sparsification</b>  | <b>33</b> |
| 5.1      | Chain Length of Codes . . . . .  | 35        |
| 5.2      | Chain Length of CSPs . . . . .   | 36        |
| 5.3      | Proof of Theorem 5.4 . . . . .   | 38        |
| <b>6</b> | <b>Immediate Applications of Non-redundancy and Chain Length to Sparsification</b> | <b>42</b> |
| 6.1      | Code Sparsification for Non-Abelian Groups . . . . .                               | 42        |
| 6.2      | Maximum-size Dichotomy . . . . .   | 42        |
| 6.3      | Generalized Group Equations: Mal’tsev Sparsification . . . . .                     | 43        |
| 6.4      | Gadget Reductions . . . . .  | 44        |
| 6.5      | Applications of Chain Length to Weighted Code Sparsifiers . . . . .                | 44        |
| <b>7</b> | <b>New results on Non-redundancy via Matching Vector Families</b>                  | <b>46</b> |

|          |  |           |
|----------|--|-----------|
| 7.1      | Conditional Non-redundancy . . . . .             | 47        |
| 7.2      | Connection to Matching Vector Families . . . . . | 48        |
| 7.3      | Constructing a Non-redundant Instance . . . . .  | 49        |
| 7.4      | Upper Bounding Non-redundancy . . . . .          | 50        |
| <b>8</b> | <b>Predicates with Unresolved Non-Redundancy</b> | <b>53</b> |
| 8.1      | The Simplest Unresolved Predicate . . . . .      | 53        |
| 8.2      | Modulo Predicates. . . . .                       | 56        |
| 8.3      | Near-polynomial Predicates . . . . .             | 58        |
| <b>9</b> | <b>Conclusion and Open Questions</b>             | <b>58</b> |

# 1 Introduction

The broad goal in sparsification is to replace an object by a more compact surrogate, typically a carefully chosen subsample, that preserves the behavior of the object under some metric of interest. For instance, for preserving cuts in undirected graphs, the influential works of Karger [Kar93] and Benczúr and Karger [BK96] showed that every graph has an edge-weighted subgraph with near-linear number of edges that preserves the value of all (edge) cuts up to a  $(1 \pm \epsilon)$  multiplicative factor. These papers have had a substantial impact in shaping the last thirty years of work in areas such as spectral sparsifiers [ST11, BSS12, LS18], clustering [KVV04, SPR11], hypergraph sparsifiers [KK15, CKN20, KKTY21, KK24, KPS24b], linear solvers [ST04, Vis13, KMP14], convex optimization [LS14, AK16, Tod16], sketching/streaming algorithms [AG09, AGM12b, AGM12a, ACK<sup>+</sup>16, McG14, KLM<sup>+</sup>17, BHM<sup>+</sup>21], max-flow/min-cut algorithms [LR99, CKM<sup>+</sup>11, KLOS14, CKL<sup>+</sup>22], machine learning [LCY<sup>+</sup>21, CSZ22, ZSW<sup>+</sup>23, GBY<sup>+</sup>24], submodular functions [KK24, Sch24, Raf24, Qua24], differential privacy [BBDS12, AU19], PageRank [Chu14], and even theoretical physics [HKTH16, Van18, TN22], among many other works.

Among the multiple exciting dimensions in which cut sparsification has been generalized, we now highlight two which form the backdrop for our work. Note that the graph cut problem can be modeled by the arity-two Boolean constraint  $x + y = 1 \pmod{2}$ . One can thus generalize cut sparsification by allowing for arbitrary constraints (of any arity over some finite domain) as considered in the field of constraint satisfaction problems (CSPs), leading to *CSP sparsification*. This direction was proposed by Kogan and Krauthgamer [KK15] in their work on hypergraph cut sparsifiers, where the not-all-equal constraint captures hypergraph cut. As a special case, arbitrary binary CSPs (where each constraint has two variables) were studied in [FK17] for the Boolean domain and in [BŽ20] for general domains, leading to a dichotomy: either near-linear sized sparsifiers exist, or no improvement over quadratic is possible.

In another direction, one can instead look toward more general structures to sparsify. For instance, a recent line of work by Khanna, Putterman, and Sudan turned toward sparsifying linear codes [KPS24a], or more generally subgroups of powers of Abelian groups [KPS25a]. Beyond being algorithmically efficient [KPS25a], these structural results have led to exciting new results in CSP sparsification by constructing optimal sparsifiers when the constraints can be embedded into linear/Abelian equations.

In this work, we obtain sparsifiers encompassing both these generalizations via a unified approach to sparsification of *non-linear* codes. The resulting sparsifiers for CSPs have *optimal asymptotic size* up to polylogarithmic factors, for *every choice of predicate* defining the CSP. In other words, we *pinpoint the optimal extent to which an arbitrary CSP can be sparsified*.<sup>1</sup>

## 1.1 Non-linear Code Sparsification

We first state our result for codes as it is very general and crisply stated, and then turn to the consequences and further new results for CSPs. For a non-linear code<sup>2</sup>  $C \subseteq \{0, 1\}^m$ , an  $\epsilon$ -sparsifier

---

<sup>1</sup>In this work we focus on the *existence* of sparsifiers, which is already highly non-trivial (e.g., [KPS24a, BŽ20] are also non-algorithmic). Future directions (and barriers) for algorithmic aspects are briefly discussed in Sections 1.8 and 9.

<sup>2</sup>Such a problem can be equivalently phrased as the sparsification of *set families*. However, we find the code framework notationally more convenient for our applications.

(for a parameter  $\varepsilon \in (0, 1)$ ) is a weight function  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$  such that for every codeword  $c$ , adding up the weights of its nonzero positions, i.e.,  $\sum_i w(i)c_i$ , is an accurate estimate of the Hamming weight of  $c$  (i.e.,  $\text{Ham}(c) := \sum_i c_i$ ) to within a  $(1 \pm \varepsilon)$  multiplicative factor (Definition 2.8). The goal is to minimize the support of  $w$  (i.e., the number of nonzero entries  $w(i)$ ), and the minimum value is called  $\varepsilon$ -sparsifiability of  $C$  and is denoted  $\text{SPR}(C, \varepsilon)$ . One of our main results is an upper bound on the sparsifiability in terms of a natural combinatorial parameter of the code called its *non-redundancy*  $\text{NRD}(C)$ , defined as follows:

$\text{NRD}(C)$  is the size of the largest subset of indices  $I \subseteq [m]$  such that for each  $i \in I$ , there is a codeword  $c \in C$  with  $c_i = 1$  and  $c_{i'} = 0$  for  $i' \in I \setminus \{i\}$ .

In other words, if we imagine the code as a matrix whose rows are codewords, its non-redundancy is largest square submatrix which is a permutation matrix.

Our main result can then be stated compactly as follows.

**Theorem 1.1 (Main).** *For all  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ ,*

$$\text{SPR}(C, \varepsilon) = O(\text{NRD}(C)(\log m)^6 / \varepsilon^2).$$

This theorem is highly versatile and can cover the many motivating examples for studying such sparsification. For the sparsification of constraint satisfaction problems, we can let the codewords in our code  $C$  correspond to indicators vectors of assignments to clauses; we discuss this in more detail in Section 1.2. For sparsification of linear codes and more broadly group codes, Theorem 1.1 can be used to give essentially tight bounds, matching the results of [KPS24a, KPS25a] (modulo the efficiency of [KPS25a]); see Section 6 for further discussion and connections.

**Interpreting Theorem 1.1 as a multiplicative-error VC theorem.** If the goal is to estimate the Hamming weight of every  $c \in C$  within an *additive* error of  $\pm \varepsilon m$ , then it well-known that the sample complexity is characterized by the *VC-dimension* of  $C$ . Recall that the VC-dimension of  $C \subseteq \{0, 1\}^m$  is largest subset of indices  $I \subseteq [m]$  such that every pattern of 0's and 1's occurs amongst the codewords of  $C$  in the positions  $I$ , i.e.,  $C|_I = \{0, 1\}^{|I|}$ .<sup>3</sup> A random subset  $I \subset [m]$  of size  $O(\text{VC-dim}(C)/\varepsilon^2)$  [VC15, AB99] is known to satisfy  $\frac{1}{|I|}(\sum_{i \in I} c_i) \in \frac{\text{Ham}(c)}{m} \pm \varepsilon$  for every  $c \in C$ . Conversely, one needs to sample at least  $\Omega_\varepsilon(\text{VC-dim}(C))$  coordinates to obtain  $\pm \varepsilon$  additive error.

With this view, Theorem 1.1 in effect says that the non-redundancy  $\text{NRD}(C)$  plays the role of the VC dimension of  $C$  when the goal is to estimate the Hamming weight of the set  $c \in C$  within a *multiplicative*  $(1 \pm \varepsilon)$  error. This is a harder goal, since one has to be pay special attention to low-weight codewords (uniform sampling will clearly fail in general). Rather surprisingly, when one performs the sampling carefully and re-weights the sampled coordinates appropriately, Theorem 1.1 says that multiplicative error sparsification *is* in fact possible for any code, albeit with sparsity roughly the non-redundancy of the code.

The NRD of a code can in general be much larger than its VC-dimension. In fact, it is easily seen that  $\text{NRD}(C)$  is precisely the VC-dimension of the *union-closure* of  $C$  (see Proposition 4.3). This connection to union-closed families plays a crucial role in the proof of Theorem 1.1. See the technical overview (Section 1.5) for more details, including discussion of a significantly simpler  $\tilde{O}_\varepsilon(\text{NRD}(C) \log |C|)$ -sized sparsifier.

<sup>3</sup>In the set family view, for  $\mathcal{F} \subseteq 2^U$ , this is the largest subset of the universe  $U$  that is *shattered* by the family  $\mathcal{F}$ .

## 1.2 CSP Sparsification

We now turn to (unweighted<sup>4</sup>) CSP sparsification. For a relation  $R \subseteq D^r$  of arity  $r$  over a finite domain  $D$ , an instance  $\Psi$  of the CSP( $R$ ) problem consists a variable set  $X$  and a constraint set  $Y \subseteq X^r$ . An assignment  $\sigma : X \rightarrow D$  satisfies a constraint  $y = (x_1, x_2, \dots, x_r) \in Y$  if  $(\sigma(x_1), \sigma(x_2), \dots, \sigma(x_r)) \in R$ . The value  $\text{val}(\Psi, \sigma)$  of an assignment  $\sigma$  is the number of constraints  $y \in Y$  that it satisfies. Similarly, for a weight function  $w : Y \rightarrow \mathbb{R}_{\geq 0}$ , the weighted value  $\text{val}(\Psi, w, \sigma)$  is the sum of weights  $w(y)$  of all constraints  $y \in Y$  that  $\sigma$  satisfies. The goal in CSP sparsification is to output a weight function  $w : Y \rightarrow \mathbb{R}_{\geq 0}$  of small support, such that for *every* assignment  $\sigma : X \rightarrow D$ ,

$$(1 - \varepsilon) \text{val}(\Psi, \sigma) \leq \text{val}(\Psi, w, \sigma) \leq (1 + \varepsilon) \text{val}(\Psi, \sigma) ,$$

and minimum such support size is denoted  $\text{SPR}(\Psi, \varepsilon)$ .

The  $\varepsilon$ -sparsifiability of the relation  $R \subseteq D^r$ , as a function of number of variables, is defined to the maximum (i.e., worst-case) value of  $\text{SPR}(\Psi, \varepsilon)$  over all  $n$ -variables instances  $\Psi$  of CSP( $R$ ). We denote it by  $\text{SPR}(R, n, \varepsilon)$  and it is the chief object of our study. Note that this is for the unweighted case, see Section 1.3 how this result can be (tightly) applied to the weighted case.

Let us note an obvious obstruction to sparsification. Suppose we have an instance  $\Psi = (X, Y)$  of CSP( $R$ ) such that for each of its constraints  $y \in Y$ , there is an assignment  $\sigma_y : X \rightarrow D$  that satisfies only  $y$  and no other constraint. Then clearly  $\Psi$  cannot be sparsified at all—dropping any constraint  $y$  would make the value of  $\sigma_y$  drop from 1 to 0. We call such an instance a *non-redundant* instance of CSP( $\bar{R}$ ), where  $\bar{R} = D^r \setminus R$  (cf., [BCH<sup>+</sup>13, BCK20]).<sup>5</sup> As introduced by Bessiere, Carbonnel, and Katsirelos [BCK20], we denote the size of the largest such non-redundant instance of CSP( $\bar{R}$ ) on  $n$ -variables by  $\text{NRD}(\bar{R}, n)$  and call it the non-redundancy of  $\bar{R}$ . Thus a trivial lower bound on sparsifiability of CSP( $R$ ), regardless of the choice of  $\varepsilon \in (0, 1)$ , is given by

$$\text{SPR}(R, n, \varepsilon) \geq \text{NRD}(\bar{R}, n) , \tag{1}$$

and this holds even if the goal is merely to preserve which assignments have nonzero value.

Rather remarkably, this simplistic lower bound can be met and one can sparsify all the way down to  $\text{NRD}(\bar{R}, n)$  times polylogarithmic factors! In fact, this turns out to be an easy corollary of Theorem 1.1. One can associate a canonical code  $C_\Psi \subseteq \{0, 1\}^Y$  with any CSP( $R$ ) instance  $\Psi = (X, Y)$  whose codewords  $c_\sigma$  correspond to the assignments  $\sigma : X \rightarrow D$ , and  $c_{\sigma, y}$  is 1 precisely when  $\sigma$  satisfies  $y$ . It is easy to check that CSP sparsification of  $\Psi$  reduces to code sparsification of  $C_\Psi$ , and the non-redundancy of  $C$  equals the size of the largest non-redundant sub-instance of  $\Psi$  (viewed as an instance of CSP( $\bar{R}$ )). Combining Theorem 1.1 and (1), we therefore have our main result pinning down the sparsifiability of every CSP up to polylogarithmic factors.

**Theorem 1.2.** *For every nonempty  $R \subsetneq D^r$  and  $\varepsilon \in (0, 1)$ , we have that*

$$\text{NRD}(\bar{R}, n) \leq \text{SPR}(R, n, \varepsilon) \leq O(\text{NRD}(\bar{R}, n)(r \log n)^6 / \varepsilon^2).$$

## 1.3 Weighted CSP Sparsification

The discussion so far has focused on unweighted CSP instances, and we now shift our focus to the weighted case, where each constraint of  $Y$  comes with a weight. We also get a tight characterization

<sup>4</sup>The weighted case is discussed in Section 1.3.

<sup>5</sup>We use  $\bar{R}$  rather than  $R$  due to the conventions of each community. See Remark 2.6 for deeper technical reasons.

of weighted CSP sparsifiability, in terms of a parameter called the *chain length*, which was defined by Lagerkvist and Wahlström [LW17, LW20] in the context of *CSP kernelization* and later utilized by Bessiere, Carbonnel, and Katsirelos [BCK20] in the context of learning CSPs in a certain query model (see Section 1.6 for more details on these connections). As before, the result is obtained in the setting of weighted non-linear codes, with the consequence for weighted CSPs being an easy corollary. We just state the result for codes here (see Section 5 for the full treatment of weighted CSPs).

For weighted sparsification of a code  $C \subseteq \{0, 1\}^m$ , we might have an arbitrary input weighting  $\zeta : [m] \rightarrow \mathbb{R}_{\geq 0}$  of its coordinates, and we must find a sparsifier  $\tilde{w} : [n] \rightarrow \mathbb{R}_{\geq 0}$  of low support that sparsifies  $C$  with respect to the weighting  $\zeta$ , i.e.,  $\langle \tilde{w}, c \rangle \in (1 \pm \varepsilon)\langle \zeta, c \rangle$ . The minimum possible support of sparsifiers over all weightings  $\zeta$  is called the weighted  $\varepsilon$ -sparsity  $\text{wSPR}(C, \varepsilon)$ .

Now we define chain length. If we line up the codewords of  $C$  as rows of an  $|C| \times m$  matrix and allow arbitrary column permutations, the chain length of  $C$ , denoted  $\text{CL}(C)$ , is the dimension of the largest upper triangular square submatrix with 1's on the diagonal.<sup>6</sup>

In our main result for the weighted setting, we pin the sparsifiability of a weighted code to its chain length. Note that in the weighted case  $\text{CL}(C)$  is also a lower bound.

**Theorem 1.3.** *For all  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ , we have*

$$\text{CL}(C) \leq \text{wSPR}(C, \varepsilon) = O(\text{CL}(C)(\log m)^6/\varepsilon^2).$$

The upper bound proceeds by using Theorem 1.1 as a black-box together with a geometric weight bucketing technique from [KPS25a]. The lower bound proceeds by applying an exponential sequence of weights to the indices  $i_1, \dots, i_{\text{CL}(C)} \in [m]$  forming a maximal chain. Of note, if for a particular set of weights, the ratio between maximum and minimal weights is  $\lambda \ll \exp(\text{CL}(C)/\text{NRD}(C))$ , we get a sharper upper bound of  $\tilde{O}_\varepsilon(\text{NRD}(C) \log \lambda)$  (see Corollary 5.19).

## 1.4 Non-redundancy and Chain Length of Specific Relations

As shown in our main results, to understand the sparsifiability of various CSPs, we must understand their non-redundancy and chain length. These quantities are readily computed in some simple cases. For example, for the relation  $\text{OR}_r := D^r \setminus \{0^r\}$ , we have that  $\text{NRD}(\text{OR}_r, n) = \text{CL}(\text{OR}_r, n) = \Theta(n^r)$ . Indeed  $Y = \binom{X}{r}$  is a non-redundant instance because setting all but  $r$  variables to 1 fails to satisfy exactly that  $r$ -tuple (see [FK17, Car22, KPS25a]). When  $R$  is affine,  $\text{NRD}(R, n) = \text{CL}(R, n) = \Theta(n)$ , and when  $R$  is defined as the zero set of a degree  $k$  polynomial,  $\text{NRD}(R, n) = \text{CL}(R, n) = O(n^k)$ ; these follow from simple rank arguments (e.g., [LW20]). Via Theorem 1.2, these special cases (plus simple gadget reductions) already capture all the previously known upper and lower bounds for CSP sparsification (see Section 1.6 for more details on the CSP sparsification literature).

Furthermore, there are also some non-trivial upper bounds known on  $\text{NRD}$  and  $\text{CL}$  in the literature, which we can now import to sparsifiability for free courtesy Theorem 1.2. For instance, the so-called Mal'tsev relations, which generalize affine predicates (i.e., cosets) over Abelian groups, have been shown to have  $O_D(n)$  non-redundancy and chain length [LW20, BCK20], and therefore by Theorem 1.2 their complements have near-linear sparsifiability. Carbonnel [Car22] showed that

---

<sup>6</sup>In this view  $\text{NRD}(C)$  is the dimension of the largest identity submatrix, so clearly  $\text{NRD}(C) \leq \text{CL}(C)$ . The quantity  $\text{CL}(C)$  was called *visible rank* in [AG21] and served as a field independent lower bound on the rank of  $C$ .

if  $R$  is an arity  $r$  relation that does not contain<sup>7</sup> any copy of  $\text{OR}_r$ , then  $\text{NRD}(R, n) \leq O(n^{r-\delta_r})$  for  $\delta_r = 2^{1-r}$  (the specific bound arises from a classic hypergraph Turán result [Erd64]). By Theorem 1.2 this immediately implies  $\text{SPR}(\overline{R}, n, \varepsilon) \leq \tilde{O}_\varepsilon(n^{r-\delta_r})$ , where  $\tilde{O}(\cdot)$  hides polylogarithmic factors in  $n$ , yielding an  $\Omega(n^r)$  vs  $\tilde{O}(n^{r-\delta_r})$  dichotomy for sparsification of arity  $r$  CSPs. (This was known for the Boolean case [KPS25a]; see the related work subsection.)

The non-redundancy (and chain length) of a relation can in general be difficult to estimate. Thus while in principle Theorem 1.2 pins down the sparsifiability of every CSP, for specific relations, it can still be non-trivial to actually determine the asymptotic behavior of its sparsifiability. Our next set of results makes progress in this direction via novel methods to bound non-redundancy.

Given that the non-redundancy of linear predicates is easy to pin down, we consider a natural family of relations which are very close to being linear. Specifically, let  $3\text{LIN}_G = \{(x, y, z) \mid x + y + z = 0\}$  over an Abelian group  $G$ , and consider  $3\text{LIN}_G^* = 3\text{LIN}_G \setminus \{(0, 0, 0)\}$ . (We pick arity 3 since the arity 2 case is already fully resolved [FK17, BŽ20].) Being defined by a linear equation over an Abelian group, we already know that  $\text{NRD}(3\text{LIN}_G, n) = \Theta_G(n)$ . However the non-redundancy of  $3\text{LIN}_G^*$  seems challenging to understand. Existing methods in the literature only yield  $\text{NRD}(3\text{LIN}_G^*, n) \in [\Omega_G(n), O_G(n^2)]$ .

We introduce a new method for bounding the non-redundancy of predicates like  $3\text{LIN}_G^*$  by connecting them to the theory of *matching vector (MV) families* [Yek08, DGY11] that have been used in the construction of locally decodable codes. Exploiting this connection, we construct a non-redundant instance to establish that  $\text{NRD}(3\text{LIN}_G, n) \geq \Omega(n^{1.5})$  for all Abelian groups of order  $\geq 3$ . Adapting ideas from the analysis of MV families together with some combinatorial ideas, we also prove an upper bound  $\text{NRD}(3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}, n) = \tilde{O}_p(n^{2-\varepsilon_p})$  for  $\varepsilon_p = \frac{2}{2p-1}$  and  $p$  prime. Specializing for  $p = 3$ , we have the following result, which also gives the first examples of relations whose non-redundancy and sparsifiability have a non-integral exponent.

**Theorem 1.4.** *We have*

$$\text{NRD}(3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*, n) \in [\Omega(n^{1.5}), \tilde{O}(n^{1.6})], \quad \text{and} \quad \text{SPR}(\overline{3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*}, n, \varepsilon) \in [\Omega(n^{1.5}), \tilde{O}(n^{1.6}/\varepsilon^2)].$$

We now transition to discussing the broader context of our work in the literature.

## 1.5 Technical Overview

We next describe the primary techniques we use to prove Theorem 1.1 and Theorem 1.4.

**A Simple Sparsifier.** To begin, we discuss a warm-up version of Theorem 1.1 which proves a weaker upper bound of  $\text{SPR}(C, \varepsilon) \leq \tilde{O}_\varepsilon(\text{NRD}(C) \cdot \log |C|)$  (see Theorem 3.1), which for CSPs corresponds to an extra factor of the number of variables  $n$ . The key technical insight (Lemma 3.3) is that for all  $d \in [m]$ , the set of codewords of  $C$  with Hamming weight at most  $d$  (denoted by  $C_{\leq d}$ ) has total support size at most  $d \cdot \text{NRD}(C)$ . This can be proved inductively by noticing that dropping a suitable non-redundant set of coordinates decreases the Hamming weight of every codeword of  $C$  by at least one.

With this lemma, we can recursively construct a sparsifier as follows, similar to the divide-and-conquer framework in [KPS24a, KPS25a] for linear codes. Pick  $d \approx \tilde{\Theta}_\varepsilon(\log |C|)$  and let  $I \subseteq [m]$  be the support of  $C_{\leq d}$ . Every  $i \in I$  is given weight 1 in our sparsifier. For the rest of  $[m]$ , let

---

<sup>7</sup>See Theorem 6.2 for a precise definition.

$J \subseteq [m] \setminus I$  be a subsample where each  $i \in [m] \setminus I$  is kept independently with probability  $1/3$ . Using a standard Chernoff bound, we can show that with positive<sup>8</sup> probability the following holds for all  $c \in C$ :

$$3\text{Ham}(c|_J) + \text{Ham}(c|_I) \in \left[1 - \frac{\varepsilon}{2 \log_2 m}, 1 + \frac{\varepsilon}{2 \log_2 m}\right] \cdot \text{Ham}(c).$$

By induction, we can find a  $\tilde{O}_{\varepsilon'}(\text{NRD}(C') \cdot \log |C'|)$   $\varepsilon'$ -sparsifier for  $C' := C|_J$  with  $\varepsilon' := (1 - 1/\log_2 m)\varepsilon$ . Scaling this sparsifier by 3 and adding weights for  $I$  gives us an  $\varepsilon$ -sparsifier of  $C$ .

**Entropy-based Sparsification.** The key inefficiency of the  $\tilde{O}_\varepsilon(\text{NRD}(C) \cdot \log |C|)$  bound is that the use of Lemma 3.3 is too conservative. For the purposes of this overview, assume that all codewords of  $C$  have the same Hamming weight  $d \approx \text{NRD}(C)$  as that is the most representative case. Naively, Lemma 3.3 says we should set aside  $d^2$  coordinates of  $[m]$  to “sparsify” all codewords of weight  $d$ . However, we can give a heuristic argument that far fewer than  $d^2$  of these potential coordinates contain useful information for our sparsifier.

Assume without loss of generality that the support of  $C$  lies in  $[d^2]$ . For each  $i \in [d^2]$ , let  $p_i$  be the probability that a codeword  $c \in C$  selected uniformly at random has  $c_i = 1$ . Since each codeword of  $C$  has Hamming weight  $d$ , we have that  $p_1 + \dots + p_{d^2} = d$ . Thus, the average value of  $p_i$  is  $1/d$ . Consider the case in which each  $p_i = O(1/d)$ . In particular, no coordinate is distinguishing itself as a “must” to add to the sparsifier. A priori, the size of  $C$  may be  $\exp(\tilde{\Omega}(d))$ , so we cannot immediately use Chernoff bounds to analyze a random subsampling of the coordinates.

To get around this issue, we need to prove a much stronger upper bound on the size of  $C$ , similar to Benczúr and Karger’s cut-counting bound [BK96] and its adaptation to linear codes [KPS24a, KPS25a]. However, we use an entirely new method for proving such bounds based on the entropy method Gilmer [Gil22] developed to prove the union-closed sets conjecture up to a constant factor. In our context, pick  $t = \tilde{\Theta}(d)$  and sample uniformly and independently  $t$  codewords  $c_1, \dots, c_t \in C$ . Let  $c$  be the bitwise OR of these  $t$  codewords, and let  $\mathcal{D}$  be the distribution of  $c$  over  $\{0, 1\}^{d^2}$  (recall that the weight  $d$  codewords are supported on  $d^2$  coordinates). Since each  $p_i = O(1/d)$ , by adapting Gilmer’s method (or more precisely, a refinement due to Sawin [Saw23]), we can show the entropy of  $\mathcal{D}$  is at least  $\tilde{\Theta}(t) = \tilde{\Theta}(d)$  times the entropy of the uniform distribution over  $C$  (i.e.,  $\log_2 |C|$ )—a similar inequality appears in [Wak24].

To apply this fact, observe that each sample of  $\mathcal{D}$  lies in the “OR-closure” of  $C$  (denoted by  $\text{span}_{\text{OR}}(C)$ ). As such, the entropy of  $\mathcal{D}$  is at most  $\log |\text{span}_{\text{OR}}(C)|$ , which by the Sauer-Shelah-Peres lemma is at most (up to log factors) the VC dimension of  $\text{span}_{\text{OR}}(C)$ . It is easily seen that the VC dimension of  $\text{span}_{\text{OR}}(C)$  equals the non-redundancy of  $C$  [BCK20]. Therefore, we have proved that  $\tilde{\Theta}(t) \cdot \log_2(C) \leq \tilde{O}(\text{NRD}(C))$ . Since  $t \approx d \approx \text{NRD}(C)$ ,  $C$  is actually at most quasipolynomial in size! Thus we can now use a Chernoff bound to prove that  $C$  can be subsampled to  $\tilde{O}_\varepsilon(d)$  coordinates while approximately preserving all Hamming weights.

Recall this discussion was purely about the “uniform” case  $p_i = O(1/d)$ . In general, we apply minimax theorem to prove the following “skewed” versus “sparse” dichotomy (see Proposition 4.14): for every code  $C$  and parameter choice  $\theta \geq 1$  there is either a probability distribution  $\mathcal{P}$  over  $C$  for which each coordinate equals 1 with probability at most  $1/\theta$  (i.e.,  $\mathcal{P}$  is “ $\theta$ -sparse”); or, there

---

<sup>8</sup>We only need positive probability since we are focused on existence. This can easily be amplified to  $1 - 1/m^{\Omega(1)}$  probability by making  $d$  a factor of  $\log m$  bigger. In applications to CSPs, the main algorithmic bottleneck is (approximately) finding  $I$ , which appears to be similar in difficulty to an open problem in CSP kernelization (see Section 1.8).

is a probability distribution  $\mathcal{Q}$  over the coordinates of  $C$  such that for every (nonzero)  $c \in C$ , we have that  $\mathcal{Q}$ 's measure of  $\text{supp}(c)$  is at least  $1/\theta$  (i.e.,  $\mathcal{Q}$  is a “ $\theta$ -cover.”) For a suitable choice of  $\theta$ , we repeatedly apply Proposition 4.14 to recursively build the sparsifier: in the  $\theta$ -sparse case, we use the entropy method to prove that a “small” number of codewords of  $C$  can be removed to put us in the  $\theta$ -cover case (see Lemma 4.15); and in the  $\theta$ -cover case, we sample from the  $\theta$ -cover to get a coordinate to add to our sparsifier. This procedure culminates in showing that we can set aside  $\tilde{O}_\varepsilon(\text{NRD}(C))$  coordinates to have weight 1 in our sparsifier with the remainder of the code being sufficiently sparse that subsampling can be used (Theorem 4.16). Note that the statement of Theorem 4.16 resembles the analogous decompositions for linear codes [KPS24a, KPS25a]. However, their method found all the coordinates to set aside in “one pass,” whereas we iteratively understand the dense and sparse structure of our non-linear code. With Theorem 4.16 in hand, we construct the sparsifier with a recursive argument similar to that of Theorem 3.1.

As mentioned earlier, extended these ideas to weighted sparsification (Theorem 1.3) is relatively straightforward. We adapt a weight-binning argument of [KPS25a] by essentially computing an (unweighted) sparsifier for each group of coordinates that is similar in weight (within  $\text{poly}(m)$ ). We then analyze the aggregated size of these sparsifiers by comparing the sum of the non-redundancies of the groups of coordinates to the chain length of the code.

**Connections to Matching Vector Families.** We now switch gears to briefly discussing the key ideas behind Theorem 1.4. Let  $G := \mathbb{Z}/3\mathbb{Z}$  and recall that  $3\text{LIN}_G = \{(x, y, z) \mid x + y + z = 0\}$  and  $3\text{LIN}_G^* = 3\text{LIN}_G \setminus \{(0, 0, 0)\}$ . It is well-known that since  $3\text{LIN}_G$  is an affine predicate, we have that  $\text{NRD}(3\text{LIN}_G, n) = \Theta(n)$ , which is much smaller than our bound on  $\text{NRD}(3\text{LIN}_G^*, n)$ . As such, we prove that to understand the asymptotics of  $\text{NRD}(3\text{LIN}_G^*, n)$  it suffices to look at specially-structured non-redundant instances.

Recall that an instance  $\Psi := (X, Y)$  of  $\text{CSP}(3\text{LIN}_G^*)$  is non redundant if for every clause  $y \in Y$  there is an assignment  $\sigma_y$  which satisfies every clause of  $\Psi$  except  $y$ . We show that with at most an additive  $\Theta(n)$  change in size, we can assume that  $\sigma_y$  maps  $y$  to  $(0, 0, 0)$ . In other words, each  $\sigma_y$  is a satisfying assignment to  $\Psi$  when viewed as an instance of  $\text{CSP}(3\text{LIN}_G)$  (see Proposition 7.3). This idea of “conditional” non-redundancy abstracts and generalizes an approach from [BCK20].

Since the set of solutions to an instance of  $\text{CSP}(3\text{LIN}_G)$  form a vector space (of some dimension, say  $d$ ) over  $\mathbb{F}_3$ , we can think of each variable  $x \in X$  of  $\Psi$  as a vector  $v_x \in \mathbb{F}_3^d$  and the assignments as linear maps on the vectors. Because we are studying satisfying assignment to  $\text{CSP}(3\text{LIN}_G)$ , these vectors are highly structured: for each  $y := (x_1, x_2, x_3) \in Y$ , we have that  $v_{x_1} + v_{x_2} + v_{x_3} = 0$ . Further,  $\sigma_y$  can be viewed as a linear map taking each of  $v_{x_1}, v_{x_2}, v_{x_3}$  to 0, while mapping at least one vector in every other triple in  $Y$  to a nonzero value. We call this family of vectors together with these assignments a  $G$ -ensemble (Definition 7.5), and note that it bears a strong resemblance to matching vector families.

In particular, we adapt techniques used by Dvir, Gopalan, and Yekhanin [DGY11] for constraining the size of matching vector families to give nontrivial upper and lower bounds on the size of  $G$ -ensembles. For the lower bound (Theorem 7.8), we directly construct a non-redundant instance with  $\Omega(n^{1.5})$  clauses. The proof is self-contained and elementary.

The upper bound (Theorem 7.13) is slightly more technical. We break the proof into cases based on whether the embedding dimension  $d$  of the vectors is small ( $d = \tilde{O}(n^{0.4})$ ) or large ( $d = \tilde{\Omega}(n^{0.4})$ ). For small  $d$ , we adapt the polynomial method used in [DGY11] to prove there can be at most  $O(d^4) = \tilde{O}(n^{1.6})$  non-redundant clauses. On the other hand, when  $d$  is large, we ignore

the assignments  $\sigma_y$  and use a careful induction (Lemma 7.11) to show that the geometry of the vectors imply that some  $x \in X$  is a member of at most  $\tilde{O}(n/d) = \tilde{O}(n^{0.6})$  clauses, thereby leading to a bound of at most  $\tilde{O}(n^{1.6})$  clauses total. Closing the gap between  $\Omega(n^{1.5})$  and  $\tilde{O}(n^{1.6})$  for  $\text{NRD}(3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*, n)$  is a tantalizing open question.

## 1.6 Related Work

Our results and techniques have connections to many areas including computational complexity theory, extremal combinatorics, coding theory, and learning theory. We now give a general overview of these connections.

**CSP Sparsification.** Since we already discussed the history of CSP sparsification, we give a comprehensive list of known results about CSP sparsification (up to polylog factors).

- The case of binary CSPs ( $r = 2$ ) is fully classified. In particular, for every finite domain  $D$  and  $R \subseteq D^2$ , we either have that  $\text{SPR}(R, n, \varepsilon) = O(n/\varepsilon^2)$  or  $\text{SPR}(R, n, \varepsilon) = \Omega(n^2)$  [BŽ20]. However, the sparsification routine is only efficient in the Boolean case [FK17]. Of note,  $\text{SPR}(R, n, \varepsilon) = \Omega(n^2)$  if and only if there exist  $D_1, D_2 \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times D_2)| = 1$  (informally  $R$  has an “induced copy” of  $\text{AND}_2$ ).
- For  $r \geq 3$ , much less is known. Kogan and Krauthgamer [KK15] contributed near-linear hypergraph cut sparsifiers (i.e., the predicate is  $\text{NAE}_r := \{0, 1\}^r \setminus \{0^r, 1^r\}$ ). Since then, there have been multiple improvements in efficiently constructing hypergraph sparsifiers/sketches (e.g., [CKN20, KKTY21, KPS24b]).
- The breakthroughs of Khanna, Putterman, and Sudan [KPS24a, KPS25a] construct near-linear sparsifiers for any predicate which can be defined by a system of linear (in)equations (possibly over a higher domain). For example  $\text{NAE}_r = \{x \in \{0, 1\}^r : x_1 + \dots + x_r \not\equiv 0 \pmod r\}$ . Of note, their first paper [KPS24a] only proved the result over finite fields (and was nonalgorithmic), whereas their second paper [KPS25a] extended the result to all Abelian groups and was computationally efficient.
- The framework of Khanna, Putterman, and Sudan [KPS25a] produced numerous corollaries. In particular, if a predicate can be expressed as the nonzero set of a degree  $k$  polynomial, then it has a sparsifier of size  $\tilde{O}_\varepsilon(n^k)$ . Furthermore, they show if a predicate  $R$  can express<sup>9</sup>  $\text{AND}_k := \{1^k\}$ , then  $\text{SPR}(R, n, \varepsilon) = \Omega(n^k)$ . As a consequence, they also classify all ternary Boolean predicates ( $r = 3$ ) as well as which Boolean predicates of arity  $r$  cannot be sparsified below  $\Omega(n^r)$  (just  $\text{AND}_r$  and its bit flips), while also constructing a sparsifier of size  $\tilde{O}_\varepsilon(n^{r-1})$  in the other cases.
- It appears that lower bounds with a nontrivial dependence on  $\varepsilon$  are only known for cut sparsifiers (and thus hypergraph cut sparsifiers via a simple gadget reduction). See [ACK<sup>+</sup>16, CKST19] as well as Section 9 for further discussion.

---

<sup>9</sup>More specifically, we say that  $R \subseteq \{0, 1\}^r$  can express  $\text{AND}_k$  if there exists a map  $z : [r] \rightarrow \{0, 1, x_1, \dots, x_k, \bar{x}_1, \dots, \bar{x}_k\}$  such that  $R(z(1), \dots, z(r)) = \text{AND}_k(x_1, \dots, x_k)$ . We discuss a more general framework of gadget reductions in Section 6.4.

**CSP Kernelization.** Another question similar in spirit to CSP sparsification is that of *CSP kernelization*.<sup>10</sup> The basic question is to, given an instance  $\Psi$  of  $\text{CSP}(R)$ , *efficiently* find as small of an instance  $\Psi'$  of  $\text{CSP}(R)$  as possible (not necessarily a subinstance) such that  $\Psi$  and  $\Psi'$  are either both satisfiable or both unsatisfiable. This particular question can be attributed to Dell and van Melkebeek [DvM14], who were particularly inspired by Impagliazzo, Paturi, and Zane’s sparsification lemma [IPZ01] and Harnik and Naor’s compression framework [HN10]. See the literature review in [DvM14] for further motivations.

At first, the problem seems rather unrelated to CSP sparsification. For example, if  $\text{CSP}(R)$  is polynomial-time tractable, then there trivially exists a kernel of size  $O(1)$ . When  $\text{CSP}(R)$  is NP-hard, however, the size of the smallest possible kernelization seems to much more closely track with the non-redundancy of  $R$ . In particular, Dell and van Melkebeek [DvM14], proved that assuming  $\text{coNP} \not\subseteq \text{NP/poly}$ , the problem  $k$ -SAT cannot be kernelized below  $\Omega(n^{k-\varepsilon})$  for any constant  $\varepsilon > 0$ , which is close to  $k$ -SAT’s non-redundancy of  $\Theta(n^k)$ .

Furthermore, most upper bounds on the kernelization of NP-hard predicates follow from upper bounds on non-redundancy (see [Car22]). For example the works of Chen, Jansen, and Pieterse [CJP20] as well as Lagerkvist and Wahlström [LW17, LW20] develop various kernelization methods that happen to just be “efficient” non-redundancy upper bounds. For example, these works show that if the predicate  $R$  can be expressed as the zero set of a polynomial of degree  $k$ , then there exist a kernel of size  $O(n^k)$ . This kernel happens to preserve every solution to  $R$ , so it is also a non-redundancy upper bound. Using techniques like these, they are able to prove a number of results similar to the state-of-the-art in CSP sparsification, such as a complete classification of ternary Boolean predicates and a  $O(n^{r-1})$  vs  $\Omega(n^{r-\varepsilon})$  Boolean dichotomy [CJP20]. See [JP19, JW24, Jan20, Tak23, Beu21] and citations therein for related work.

We seek to emphasize that any *efficient* CSP sparsification algorithm for  $\text{CSP}(R)$  is by design a kernelization algorithm for  $\text{CSP}(\overline{R})$  (since all codewords with weight 0 are preserved). As such, making Theorem 1.2 efficient would require explicitly proving that every CSP can be kernelized to (approximately) its non-redundancy, which is a significant open question in the CSP kernelization community (see [Car22]). See Section 1.8 and Section 9 for further discussion.

**The Union-closed Sets Conjecture.** A family  $\mathcal{F}$  of subsets of  $[n]$  is *union-closed* if  $A, B \in \mathcal{F}$  imply that  $A \cup B \in \mathcal{F}$ . In 1979, Frankl [Fra95] conjectured that there always exists  $i \in [n]$  which appears in at least half of the sets of  $\mathcal{F}$ . For decades, progress on the conjecture was minimal, with the best general result being that some  $i \in [n]$  appears in  $\Omega(1/\log_2 |\mathcal{F}|)$  of the sets [Kni94, Wój99, Gil22]. However, in 2022, Gilmer [Gil22] shocked the combinatorics community by using an entropy-based approach to prove that some  $i \in [n]$  appears in  $1/100$  of the sets. This immediately led to a large number of follow-up works refining Gilmer’s entropy method [AHS24, CL22, Peb22, Saw23, Cam22]. In particular, we can now replace ‘ $1/100$ ’ with ‘ $0.382\dots$ ’, leaving Frankl’s conjecture (technically) still open.

For our application to CSP sparsification, the entropy method used by Gilmer (and its subsequent refinements by many other researchers) is the key idea needed to show that non-redundancy is essentially the optimal size for a CSP sparsifier. In particular, the improvement from  $1/\log_2 |\mathcal{F}|$

---

<sup>10</sup>More commonly, CSP kernelization is referred to as CSP sparsification (e.g., [DvM14, LW20]). However, we refer to this line of work by the former name to reduce ambiguity. This similarity in name has been noted before in the literature (e.g., [BŽ20]), but we appear to be the first work to notice both variants of “CSP sparsification” can be analyzed with similar techniques.

to  $\Omega(1)$  is precisely the same “gain” we utilize to go from a very simple  $\tilde{O}_\varepsilon(\text{NRD}(C) \cdot \log_2 |C|)$  sparsifier (see Section 3) to our  $\tilde{O}_\varepsilon(\text{NRD}(C))$  sparsifier. See the technical overview (Section 1.5) for more details. To the best of our knowledge, our work is the first application of Gilmer’s entropy method to sparsification.<sup>11</sup>

**Matching Vector Families and Locally Decodable Codes.** In coding theory, locally decodable codes (LDCs) are a class of codes which allow for the reliable recovery of any message symbol based on a small sample of codeword symbols, even in the presence of a constant fraction of errors. A particularly interesting family of constructions of LDCs has arisen out of a theory of *matching vector codes* [Yek08] and follow-ups [Rag07, Gop09, Efr12, DGY11]. See [DGY11] for a literature survey. Simply stated, a matching vector (MV) family over a (finite) ring  $\mathcal{R}$  is a pair of lists of vectors  $u_1, \dots, u_k, v_1, \dots, v_k \in \mathcal{R}^d$  such that the inner products  $\langle u_i, v_j \rangle$  are nonzero<sup>12</sup> if and only if  $i = j$ . Informally, the  $u_i$ ’s play a role in the encoding of the  $i$ ’th message symbol, with the matching vector  $v_i$  helping with its local decoding. Given a choice of  $\mathcal{R}$  and  $d$ , the primary question of interest is to find the maximal possible value of  $k$ . This “spin off” question about LDCs has become a topic of interest in its own right [DGY11, Yek12, GHSY12, BDL12].

In this work, we demonstrate a novel application of matching vector families to the study of non-redundancy and thus (by Theorem 1.2) sparsification. In particular, we construct an explicit family of predicates such that their non-redundant instances can be viewed as a generalized MV family. We then use techniques developed for MV families to give nontrivial bounds on the non-redundancy of the predicates. See Section 7 and the technical overview (Section 1.5) for more details.

**Extremal Combinatorics.** Computing the non-redundancy of a predicate can be viewed as a problem in extremal combinatorics known as a *hypergraph Turán problem*. In particular, for an instance of a CSP to be non-redundant, every instance induced by a subset of the variables must also be non-redundant. In particular, if  $\mathcal{F}$  is a family of hypergraphs which can never appear in non-redundant instances of  $\text{CSP}(R)$ , then  $\text{NRD}(R, n) \leq \text{ex}_r(n, \mathcal{F})$ , where the hypergraph Turán number  $\text{ex}_r(n, \mathcal{F})$  is the size of the largest  $r$ -uniform hypergraph on  $n$  vertices without any  $F \in \mathcal{F}$  as a subgraph. This observation was first made explicit by Carbonnel [Car22] although the technique was also used in earlier work [BCK20]. As far as we are aware, ours is the first work to observe that these insights can also benefit the study of CSP sparsification.

The literature on hypergraph Turán numbers is quite rich. For instance, Keevash [Kee11] surveys the vast body of work on the “non-degenerate” case in which  $\text{ex}_r(n, \mathcal{F}) = \Omega_r(n^r)$ . However, for our applications, we are mostly interested in the “degenerate” case in which  $\text{ex}_r(n, \mathcal{F}) = O(n^c)$  for some  $c \in [1, r)$ . The works [BCK20, Car22] apply some of the most well-known works in this setting [Erd64, SEB73, RS78] to get some nontrivial results such as classifying precisely which predicates  $R$  have  $\text{NRD}(R, n) = \Theta(n^r)$ , extending Chen, Jansen, and Pieterse’s result for the Boolean case [CJP20]. See Sections 6.2, 8.1, and 6.5.2 for more details on specific applications.

**Query Complexity and Learning Theory.** Rather surprisingly, the definition of non-redundancy appears to have come out of the artificial intelligence community [BCK20]. In particular, a rather broad and well-studied question (e.g., [FW02, PBS08, LLMV10, BK12, BCH<sup>+</sup>13, BCK20]) is that

<sup>11</sup>Gilmer’s breakthrough is cited in the literature review of [CDL<sup>+</sup>24], but the property-testing question they study on union-closed families has no technical connection to Gilmer’s entropy method. See also [Wak24] for applications of the entropy method to learning theory and statistical physics.

<sup>12</sup>Or, more generally the inner products lie in some restricted subset of  $\mathcal{R}$ .

of *constraint acquisition*: how can an agent learn the constraints defining an instance of a constraint satisfaction problem?

A model specifically relevant to our work is the *partial membership queries* model studied by Bessiere, Carbonnel, and Katsirelos [BCK20]. In this model, the domain  $D$ , the constraint type  $R$  (or types), and the set of variables  $X$  are known but the constraints are hidden. For each query, the agent picks some subset of variables  $X' \subseteq X$  as well as a partial assignment  $\sigma : X' \rightarrow D$ . The response to the query is ‘YES’ if  $\sigma$  satisfies every constraint induced by  $X'$ , and ‘NO’ otherwise. The goal is to construct an instance of  $\text{CSP}(R)$  with the same solution set as the hidden CSP. For every CSP predicate  $R$ , they prove that the query complexity of an instance of  $\text{CSP}(R)$  on  $n$  variables is bounded between  $\Omega(\text{NRD}(R, n))$  and  $O(\text{CL}(R, n) \cdot \log n)$ . Notably, the lower bound is proved by showing that the VC dimension of the query complexity problem equals  $\text{NRD}(R, n)$ .<sup>13</sup>

## 1.7 Subsequent Work

Since the initial version of our paper was posted, numerous follow-up works have emerged which expand on the many directions covered in this paper.

### 1.7.1 Sparsifier Improvements

**Improvements to the Unweighted Sparsifier.** Very recently, [LMW26] gave an improved bound for Theorem 1.1. They coin the term “moonflowers” to refer to non-redundant subsets of a code  $C \subseteq \{0, 1\}^m$  and adapt ideas from the study of sunflowers (e.g., [ALWZ21]) to study the maximum density of a code which avoids particular moonflowers. In particular, their main sparsifier has a bound of the form.

$$\text{SPR}(C, \varepsilon) \leq O\left(\frac{\text{NRD}(C) \cdot \log m}{\varepsilon^2}\right) \cdot (\log(\text{NRD}(C)/\varepsilon))^{O(1)} (\log \log m)^{O(1)}$$

The overall method of constructing the sparsifier is rather similar to our work in the sense that it crucially uses the entropy method of Gilmer as a key step in the sparsification process, although also have a number of technical refinements of our techniques. By suitably adapting our Example 5.8 they also show a lower bound of  $\Omega(\text{NRD}(C) \cdot \log m/\varepsilon)$  for an adversarial choice of  $C$ . Despite the improvements, the impact on Theorem 1.2 is negligible as the parameters  $\log m$  and  $\log \text{NRD}(C)$  are related by a factor of  $r$  (the arity of the underlying CSP) which we consider to be a constant. Furthermore, the lower bound example does not arise from a CSP, so no limitation beyond  $\text{NRD}(C)$  itself is known for Theorem 1.2.

**Improvements to the Weighted Sparsifier.** Another recent work [BGP26b] improved Theorem 1.3 on optimal weighted code sparsifiers. By adapting the “contraction”-based techniques for linear code sparsifiers [KPS24a, KPS25a] in combination with some of the techniques developed in Section 5, the authors of [BGP26b] showed that Theorem 1.3 can be improved to

$$\text{wSPR}(C, \varepsilon) \leq O\left(\frac{\text{CL}(C) \cdot \log^2(\text{CL}(S)/\varepsilon) \cdot (\log \log(\text{CL}(S)/\varepsilon))^2}{\varepsilon^2}\right).$$

Note that the dependence on  $m$  is removed entirely. That said, since we do not understand the relationship between the non-redundancy and chain length of CSPs, these techniques are unable to

---

<sup>13</sup>This observation is directly used in proving our main result, see Section 4.1.

give any improvements to Theorem 1.2. Nor can they recover the interpolation between Theorem 1.1 and Theorem 1.3 in Corollary 5.19.

### 1.7.2 Further Study of Non-redundancy

Building off of the contributions this paper makes to the study of CSP non-redundancy, multiple follow-up works [BGJ<sup>+</sup>25, BGP26a, SV26b] have further expanded our knowledge of non-redundancy.

**New Techniques.** In particular, the lower-bound technique of Theorem 1.4 was abstracted in [BGJ<sup>+</sup>25] in a logical framework known as generalized “functionally guarded primitive positive” (fgpp) definitions. They also introduce using the Kruskal-Katona and Shearer’s inequalities to give novel upper bounds on CSP non-redundancy. As a result, we now know for any rational  $\beta \geq 1$ , there exists an explicit CSP predicate  $R$  such that  $\text{NRD}(R, n) = \Theta_R(n^\beta)$ . Our notion of conditional non-redundancy (Definition 7.2) was critical in proving this fact.

In addition, [BGJ<sup>+</sup>25] made a number of other contributions to the study of CSP non-redundancy. They show that many extremal hypergraph problems (such as understanding the maximum density of graphs with a given girth) can be encoded into suitable CSP non-redundancy problems. They also introduce a novel tool known as *Catalan polymorphisms* to study the existence of Mal’tsev embeddings (see [LW20, BCK20]).

**Further Classifications.** We mentioned earlier that the non-redundancy of every Boolean predicate of arity at most 3 is known [CJP20, KPS25a]. The recent work [BGP26a] extends this result to arity 4 by classifying the non-redundancy of every arity-4 Boolean predicate except (up to isomorphism) one. Of interest, they identify a predicate  $R \subseteq \{0, 1\}^4$  for which

$$\text{NRD}(R, n) \in \left[ \frac{n^3}{2^{\mathcal{O}(\sqrt{\log n})}}, \frac{n^3}{2^{\Omega(\log^*(n))}} \right],$$

the first example of non-redundancy growth which provably does not follow a simple power law (e.g.,  $\Theta(n^\beta)$ ).

The work of [SV26b] investigates the non-redundancy of symmetric Boolean predicate of arity 4 and 5. Notably, they identify two interesting symmetric predicates of arity 5 for which state-of-the-art techniques leave a polynomial gap between the upper and lower bounds.

### 1.7.3 Broader Connections

**Spectral Sparsifiers.** Although Theorem 1.2 resolves the optimal size of CSP sparsifiers, there is still much ongoing work to understand *spectral* variants of CSP sparsification, analogous to the theory of spectral graph and hypergraph cut sparsifiers [ST11, BSS12, LS18, CKN20, KKTY21, KK24, KPS24b, Yos26].

In one such direction, [KPS25b] define the notion of a spectral sparsifier for any CSP predicate. Their work is mostly focused on the study of spectral sparsifiers for linear equations (i.e., “spectral code sparsifiers”) where they show such sparsifiers exist of size  $\tilde{O}_\varepsilon(n^2)$ . They leave reducing the sparsifier size to near-linear as an open question.

In another direction, [BKLM26] motivated by the theory of Cayley graph sparsifiers (see also [KPS24a, HLM<sup>+</sup>26]) developed a theory of sparsification of a family  $\mathcal{A}$  of positive semi-definite

(PSD) matrices. In particular, they introduce a quantity  $N^*(\alpha, \mathcal{A})$  which measures how tightly any matrix  $A \in \mathcal{A}$  can be dominated by sums of other matrices (up to a factor of  $\alpha$ ). Their main result is that up to a logarithmic factor in the dimension of the matrices,  $N^*(\alpha, \mathcal{A})$  controls the sparsifiability of the family.

In the setting of our code sparsification problem, their result can recover our warm-up bound Theorem 3.1 by representing a code  $C \subseteq \{0, 1\}^m$  as a family of  $m$  diagonal matrices which are each  $|C|$ -dimensional. However, their union bound is too coarse to recover Theorem 1.2.

**Streaming Algorithms.** Very recently, [SV26a] showed that (up to logarithmic factors), non-redundancy of a CSP also governs the *streaming* decidability of CSPs. That is, fix a CSP predicate  $R$  and consider a streaming of clauses corresponding to an instance of  $\text{CSP}(R)$  on  $n$  variables. The goal is to minimize the amount of space needed to determine if the stream corresponds to a satisfiable instance. The reason that non-redundancy is connected is that if the streaming algorithm just keeps track of a non-redundant set of clauses, any other (dropped) clause can be logically deduced from these clauses, so decidability is preserved. Using a communication complexity lower bound, they show this analysis is essentially tight. Such questions are also closely related to the study of approximating min-CSPs. As such, any future improvements to the study of streaming decidability will also impact the study of CSP non-redundancy and CSP sparsifiability.

**Average-case Sparsification.** From Theorem 1.2, we know that every instance  $\text{CSP}(R)$  has a sparsifier of size approximately its own non-redundancy. A previous version of the paper asked whether the sparsifier size can be substantially improved in the *average case* setting. That is, if the CSP predicate  $R$  is fixed but the instance is sampled randomly, can the  $\text{NRD}(\bar{R}, n)$  barrier be broken? This was generally resolved in the affirmative by [BGP25] where they compute optimal sparsifier sizes for random instances of all CSPs and *valued* CSPs. See Section 9 for further discussion.

## 1.8 Open Questions

We conclude the introduction with a few directions of further study. See Section 8 and Section 9 for a more thorough discussion of directions for future exploration.

- **Making Theorem 1.2 efficient.** Note that the underlying construction for Theorem 1.1, if made algorithmic, runs in polynomial time with respect to the size of the code, yielding an  $\exp(O(n))$ -time algorithm<sup>14</sup> for Theorem 1.2. The primary barrier in constructing our sparsifier in  $\text{poly}(n)$  time is the fact that an efficient sparsifier is also a kernelization algorithm, but kernelizing every CSP instance to its non-redundancy is a significant open question in the kernelization community [Car22].
- **Computing  $\text{NRD}(R, n)$ .** For a general predicate  $R \subseteq D^r$ , there is no simple (even conjectured) expression for  $\text{NRD}(R, n)$ . In fact, even determining when  $\text{NRD}(R, n) = \Theta(n)$  is a major open question (e.g., [BCK20, Car22]). In Section 8, we explore a number of predicates from the various parts of the literature whose status is unresolved, including a predicate we categorize as the “simplest unresolved predicate.”
- **Non-redundancy versus Chain Length.** Recall we show that unweighted sparsification is closely tied to non-redundancy while weighted sparsification is closely tied to chain length. For

---

<sup>14</sup>This is already nontrivial, as a naive guess-and-check algorithm would require  $\exp(\tilde{O}(\text{NRD}(\bar{R}, n)))$  time.

non-linear codes, NRD and CL can be very different (e.g., Example 5.8), but the relationship for CSPs is unknown [BCK20, Car22]. In particular, it seems quite possible that there exists a CSP predicate  $R$  for which  $\text{wSPR}(R, n, \varepsilon) / \text{SPR}(R, n, \varepsilon) = n^{\Omega(1)}$ .

- **Spectral CSP Sparsification.** As mentioned, [KPS25b] recently defined a notion of spectral CSP sparsifiers. Proving an analogue of Theorem 1.2 in this setting would be a rather interesting result.

## Organization

In Section 2, we prove some basic facts about non-redundancy, sparsification and their relationship. In Section 3, we give a straightforward proof that  $\text{SPR}(C, \varepsilon) = \tilde{O}_\varepsilon(\text{NRD}(C) \log_2 |C|)$ . In Section 4, we prove Theorem 1.2 by connecting CSP sparsification to non-redundancy via Gilmer’s entropy method. In Section 5, we extend Theorem 1.2 to weighted instances. In Section 6, we discuss the immediate applications of Theorem 1.2 (and its weighted variant) based on what is known about non-redundancy and chain length in the literature. In Section 7, we bound the non-redundancy of a family of predicates via methods related to matching vector families. In Section 8, we give examples of CSP predicates in the literature whose non-redundancy is unresolved. In Section 9, we wrap up with other directions of exploration. We emphasize that the material in Sections 3, 4, and 5 are largely independent of the material in Sections 6, 7, 8.

## Acknowledgments

We thank Libor Barto, Dmitry Zhuk, Madhu Sudan, and Louie Putterman for valuable conversations. We thank Arpon Basu, Pravesh Kothari, Cosmas Kravaris, and Louie Putterman for helping to correct some errors in the original manuscript. We also thank anonymous STOC reviewers for suggestions which improved the quality of this manuscript. This research was supported in part by a Simons Investigator award, NSF grant CCF-2211972, and NSF grant DMS-2503280.

## 2 Preliminaries

In this section, we give precise definitions of the non-redundancy and sparsification of CSPs and non-linear codes and discuss how they related to each other. In particular, we show that the main sparsification result for codes (Theorem 1.1) implies the main sparsification result for CSPs (Theorem 1.2). We conclude with a few tail inequalities useful for sparsification.

### 2.1 CSPs

Given a finite set  $D$  called the *domain* and integer  $r \in \mathbb{N} := \{1, 2, \dots\}$  called the *arity*, we call any subset  $R \subseteq D^r$  a *predicate* (interchangeably called a *relation*). We say that  $R$  is *non-trivial* if  $R \neq \emptyset$  and  $R \neq D^r$ . We define an *instance*  $\Psi$  of  $\text{CSP}(R)$  to be a pair  $(X, Y)$ , where  $X$  is a (finite) set of variables and  $Y \subseteq X^r$  is the set of clauses. Typically, we let  $n := |X|$  parameterize the number of variables and  $m := |Y|$  parameterize the number of clauses. Note that  $m \leq n^r$ , so  $\log m \leq r \log n$ . We now clarify a couple points about our model of CSPs.

- In general, a CSP may consist of a family of multiple predicates. In the context of finding solutions to a CSPs, this detail may significantly change the complexity, but for sparsification and related concepts (like non-redundancy), each predicate can be sparsified separately (e.g., [KPS25a]) resulting in a multiplicative error of at most the number of predicates (e.g., [BCK20, Car22]). Since we think of predicates as being of constant size in this paper, this has no effect on our results.
- The choice  $Y \subseteq X^r$  explicitly allows for repeated variables in clauses.<sup>15</sup> All our results also apply to the restricted family of instances without repeated variables. To see why, let  $X' := X \times [r]$  and map each  $(y_1, \dots, y_r) \in Y$  to  $((y_1, 1), \dots, (y_r, r)) \in Y' \subseteq X'^r$ . One can verify that any sparsifier of  $(X', Y')$  is also a sparsifier of  $(X, Y)$ , and the number of variables only differs by a constant factor.

We say that an *assignment* is a map  $\sigma : X \rightarrow D$ . A clause  $y := (y_1, \dots, y_r) \in Y$  is *R-satisfied* by  $\sigma$  if and only if  $\sigma(y) := (\sigma(y_1), \dots, \sigma(y_r)) \in R$ . We say that  $\sigma$  is an *R-satisfying assignment* to  $\Psi$  (or just  $Y$ ) if every clause is satisfied. We let  $\text{sat}(R, \Psi)$  (or  $\text{sat}(R, Y)$  if  $X$  is fixed) denote the set of satisfying assignments to  $\Psi$ . We now formally define *non-redundancy*.

**Definition 2.1** (Non-redundancy of an instance, adapted from [BCK20]). We say that an instance  $\Psi := (X, Y)$  of  $\text{CSP}(R)$  is *non-redundant* if for all  $Y' \subseteq Y$  of size  $|Y| - 1$ , we have that  $\text{sat}(R, Y') \neq \text{sat}(R, Y)$ . In other words, for all  $y \in Y$ , there exists an assignment  $\sigma_y$  which *R-satisfies*  $Y \setminus \{y\}$  but not  $y$ . We define<sup>16</sup> the non-redundancy of an arbitrary instance  $(X, Y)$  to be the size of the largest  $Y' \subseteq Y$  such that  $(X, Y')$  is non-redundant.

**Definition 2.2** (Non-redundancy of a predicate [BCK20]). Given a relation  $R \subseteq D^r$  and  $n \in \mathbb{N}$ , we define  $\text{NRD}(R, n)$  to be the maximum number of clauses of any non-redundant instance of  $\text{CSP}(R)$  on  $n$  variables.

We further define the notion of a CSP sparsifier. Given an instance  $\Psi := (X, Y)$  of  $\text{CSP}(R)$ , we define the *R-weight* of an assignment  $\sigma : X \rightarrow D$ , denoted by  $\text{wt}(R, \Psi, \sigma)$ , to be the number of clauses of  $Y$  *R-satisfied* by  $\sigma$ . Given a weight function  $w : Y \rightarrow \mathbb{R}_{\geq 0}$ , we define the  $(w, R)$ -*weight* of  $\sigma$  to be

$$\text{wt}(R, \Psi, w, \sigma) := \sum_{y \in Y} w(y) \cdot \mathbf{1}[\sigma(y) \in R].$$

We can now define a CSP sparsifier.

**Definition 2.3** (CSP sparsifiers (e.g., Definition 3.11 [KPS24a])). Given  $R \subseteq D^r$ , we say that  $w : Y \rightarrow \mathbb{R}_{\geq 0}$  is an  $\varepsilon$ -*sparsifier* of an instance  $\Psi$  of  $\text{CSP}(R)$  if for all assignments  $\sigma : X \rightarrow D$ , we have that

$$(1 - \varepsilon) \text{wt}(R, \Psi, \sigma) \leq \text{wt}(R, \Psi, w, \sigma) \leq (1 + \varepsilon) \text{wt}(R, \Psi, \sigma).$$

<sup>15</sup>Various papers in the literature make different (and sometimes ambiguous) choices regarding allowing repeated variables. As we argue, as long as  $r$  is a constant, the answer can vary by at most a constant factor. Thus, we apply results from the literature (e.g., in Section 6) without concern to their specific convention.

<sup>16</sup>We purposely deviate from the definition of the non-redundancy of an instance in [BCK20]. In their work, they define the non-redundancy of an instance to be the number of  $Y' \subseteq Y$  of size  $|Y| - 1$  for which  $\text{sat}(R, Y') \neq \text{sat}(R, Y)$ . For non-redundant instances our definitions coincide (so, Definition 2.2 always gives the same value), but for “highly redundant” instances our non-redundancy is much greater. We do this to ensure that our definition of non-redundancy is monotone with respect to adding clauses.

We define the  $\varepsilon$ -sparsity of  $\Psi$  to be the minimum support (i.e.,  $|w^{-1}(\mathbb{R}_{>0})|$ ) of any  $\varepsilon$ -sparsifier. We define  $\text{SPR}(R, n, \varepsilon)$  to be the maximum  $\varepsilon$ -sparsity of any instance of  $\text{CSP}(R)$  on  $n$  elements.

**Remark 2.4.** One may also define sparsifiers for *weighted instances* in essentially the same manner. See Section 5 for a discussion of this variant. The main takeaway is that we also prove a tight characterization in terms of another combinatorial quantity called *chain length*.

We now prove a lower bound for sparsity in terms of non-redundancy. Observe that in an instance  $(X, Y)$  of  $\text{CSP}(R)$  if some  $y \in Y$  has a corresponding assignment  $\sigma_y$  which only satisfies  $y$ , then  $y$  must have nonzero weight in any sparsifier (or else the weight of  $\sigma_y$  will fall from 1 to 0). In particular, this implies that non-redundant instances of the *complement* of  $R$  cannot be sparsified at all.

**Proposition 2.5.** *For any  $R \subseteq D^r$ , let  $\bar{R} := D^r \setminus R$ . For all  $n \in \mathbb{N}$  and  $\varepsilon \in (0, 1)$ , we have that*

$$\text{SPR}(R, n, \varepsilon) \geq \text{NRD}(\bar{R}, n). \quad (2)$$

*Proof.* Consider any non-redundant instance  $\Psi := (X, Y)$  of  $\text{CSP}(\bar{R})$  with  $n = |X|$ . For each  $y \in Y$ , let  $\sigma_y$  be a solution in  $\text{sat}(\bar{R}, Y \setminus \{y\}) \setminus \text{sat}(\bar{R}, Y)$ .

View  $\Psi$  also as an instance of  $\text{CSP}(R)$  and observe that for all  $y \in Y$ ,  $\sigma_y$  has a  $R$ -weight of 1. Further, for any  $w : Y \rightarrow \mathbb{R}_{\geq 0}$ ,  $\sigma_y$  has a  $(R, w)$ -weight of  $w(y)$ . Thus, any  $\varepsilon$ -sparsifier  $w : Y \rightarrow \mathbb{R}_{\geq 0}$  of  $\Psi$  must have that  $w(y) \in (1 - \varepsilon, 1 + \varepsilon)$  for all  $y \in Y$ . Therefore,  $\text{supp}(w) = Y$ .

This proves that  $\text{SPR}(R, n, \varepsilon) \geq |Y|$  for every non-redundant instance  $(X, Y)$  of  $\text{CSP}(\bar{R})$  on  $n$  variables. Taking the maximum over all choices of  $Y$  proves (2).  $\square$

Theorem 1.2 establishes that this simple lower bound is essentially optimal.

**Remark 2.6.** Based on the complement in (2), one might think that  $\text{NRD}$  is defined “backwards.” However, as we shall see in Section 6.4, the definition of  $\text{NRD}$  allows for a rich (universal) algebraic framework, which should greatly aid with the classification of  $\text{NRD}$  for all predicates.

## 2.2 Non-linear Codes

An important abstraction for studying  $\text{CSP}$  sparsification is what we call *non-linear code sparsification*. We define a (non-linear) Boolean *code* to be an arbitrary  $C \subseteq \{0, 1\}^m$ . We say that  $C$  is *non-trivial* if  $C \neq \emptyset, \{0^m\}$ . For any  $c \in C$ , we define its *Hamming weight*, to be  $\text{Ham}(c) := c_1 + \dots + c_m$ . We define the support  $c \in \{0, 1\}^m$ , denoted by  $\text{supp}(c) \subseteq [m]$  to be the set of nonzero coordinates. We further define  $\text{supp}(C) = \bigcup_{c \in C} \text{supp}(c)$ .

Given  $S \subseteq [m]$  and  $c \in \{0, 1\}^m$ , we define  $c|_S \in \{0, 1\}^S$  to be the *list*  $(c_i : i \in S)$ . Likewise, we define *punctured* code  $C|_S := \{c|_S : c \in C\} \subseteq \{0, 1\}^S$ . Next, we formally define the non-redundancy of a code.

**Definition 2.7.** A subset  $I \subseteq [m]$  is *non-redundant* for a code  $C \subseteq \{0, 1\}^m$  if for each  $i \in I$ , there exists  $c \in C$  such that for all  $i' \in I$ ,  $c_{i'} = 1$  if and only if  $i = i'$ . We define the *non-redundancy* of  $C$ , denoted by  $\text{NRD}(C)$ , to be the size of the largest non-redundant set that is non-redundant for  $C$ .

In other words, if we line up the codewords of  $C$  as rows of an  $|C| \times m$  matrix,  $\text{NRD}(C)$  is the dimension of the largest identity submatrix within that matrix. When  $C$  is a linear code,  $\text{NRD}(C)$  equals the dimension of  $C$ . For non-trivial  $C$ , we have that  $1 \leq \text{NRD}(C) \leq m$ .

We now formally define a Hamming-weight sparsifier of a code. Given a *weight function*,  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$ , we define the  $w$ -weight of a codeword  $c \in C \subseteq \{0, 1\}^m$  to be

$$\langle w, c \rangle := \sum_{i=1}^m w(i)c_i.$$

**Definition 2.8.** For  $\varepsilon \in (0, 1)$ , we say that  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$  is a  $\varepsilon$ -sparsifier of  $C \subseteq \{0, 1\}^m$  if for all  $c \in C$ , we have that

$$(1 - \varepsilon) \text{Ham}(c) \leq \langle w, c \rangle \leq (1 + \varepsilon) \text{Ham}(c).$$

We define the  $\varepsilon$ -sparsity of  $C$ , denoted by  $\text{SPR}(C, \varepsilon)$ , to be the minimum support size (i.e., number of nonzero coordinates) of any  $\varepsilon$ -sparsifier of  $C$ .

### 2.3 Theorem 1.1 Implies Theorem 1.2

We now connect the non-redundancy notions of CSPs and codes together. Given a relation  $R \subseteq D^r$  and any instance  $\Psi := (X, Y)$  of  $\text{CSP}(R)$ , we can define a *satisfiability code*, denoted by  $C_{R, \Psi} \subseteq \{0, 1\}^Y$  as follows. For each assignment  $\sigma : X \rightarrow D$ , we define a codeword  $c_\sigma$  such that for each  $y \in Y$ , we have

$$c_{\sigma, y} := \mathbf{1}[\sigma(y) \in R]. \quad (3)$$

We then define  $C_{R, \Psi} := \{c_\sigma \mid \sigma : X \rightarrow D\}$ . We now show that our definitions of sparsity and non-redundancy of codes corresponds to that of CSPs.

**Proposition 2.9.** *Given  $R \subseteq D^r$  and  $n \in \mathbb{N}$ , we have that*

$$\text{SPR}(R, n, \varepsilon) = \max_{\substack{\Psi \text{ instance of } \text{CSP}(R) \\ \text{on } n \text{ variables}}} \text{SPR}(C_{R, \Psi}, \varepsilon) \quad (4)$$

$$\text{NRD}(\overline{R}, n) = \max_{\substack{\Psi \text{ instance of } \text{CSP}(R) \\ \text{on } n \text{ variables}}} \text{NRD}(C_{R, \Psi}) \quad (5)$$

*Proof.* We first prove (4). Observe that for any instance  $\Psi := (X, Y)$  of  $\text{CSP}(R)$  and any weight function  $w : Y \rightarrow \mathbb{R}_{\geq 0}$  we have that the  $(R, w)$ -weight of  $\sigma$  is equal to the  $w$ -weight of  $c_\sigma$  due to (3). Thus,  $\text{SPR}(C_{R, \Psi}, \varepsilon)$  is precisely the  $\varepsilon$ -sparsity of  $\Psi$ . Therefore,  $\text{SPR}(R, n, \varepsilon)$  is the maximum such  $\varepsilon$ -sparsity, we have proved (4).

Finally, we prove (5). Fix an instance  $\Psi := (X, Y)$  of  $\text{CSP}(R)$ . Let  $Z \subseteq Y$  be a maximum-sized non-redundant index set of  $C_{R, \Psi}$ . In particular, for each  $z \in Z$ , there is an assignment  $\sigma_z : X \rightarrow D$  such that  $\sigma_z$   $R$ -satisfies  $z$  but no other  $z' \in Z \setminus \{z\}$ . Thus,  $\sigma_z$   $\overline{R}$ -satisfies  $Z \setminus \{z\}$  but not  $z$ . Thus,  $(X, Z)$  is a non-redundant instance of  $\text{CSP}(R)$ , so  $\text{NRD}(\overline{R}, n) \geq \text{NRD}(C_{R, \Psi})$  for all  $\Psi$ .

Likewise, if we let  $(X, Z)$  be a maximum-sized non-redundant instance of  $\text{CSP}(\overline{R})$  with witnessing assignments  $\sigma_z$  for  $z \in Z$ , we can view  $\Psi := (X, Z)$  as an instance of  $\text{CSP}(\overline{R})$ . Thus,  $\{c_{\sigma_z} : z \in Z\}$  witness that  $Z$  is a non-redundant index of  $C_{R, \Psi}$ . Therefore,  $\text{NRD}(C_{R, \Psi}) = |Z|$ . This proves (5).  $\square$

Thus, we can now prove that Theorem 1.1 implies Theorem 1.2.

**Proposition 2.10.** *Theorem 1.1 implies Theorem 1.2.*

*Proof.* The lower bound follows from Proposition 2.5. For the upper bound, consider any instance  $\Psi$  of  $\text{CSP}(R)$  on  $n$  variables and  $m$  clauses. Since  $m \leq n^r$ , by Theorem 1.1, we have that

$$\text{SPR}(C_{R,\Psi}, \varepsilon) \leq O(\text{NRD}(C_{R,\Psi}(r \log n)^6 / \varepsilon^2)).$$

Take the maximum of both sides over all instances  $\Psi$  on  $n$  variables. Then, Proposition 2.9 implies Theorem 1.2.  $\square$

## 2.4 Chernoff Bounds

We now state a few Chernoff-type bounds that will be useful in our sparsification arguments. We let  $\exp$  and  $\log$  denote the natural exponential and logarithm, respectively. For any base  $b > 1$ , we let  $\exp_b(x) := b^x$  and  $\log_b(x) := \log(x) / \log(b)$ .

As our sparsifiers utilize i.i.d. subsampling, we use the following Chernoff bound to analyze the probability that Hamming weight is preserved.

**Theorem 2.11** (e.g., [MR95, MU05, FHHP19]). *Let  $X_1, \dots, X_n$  be i.i.d. samples of the Bernoulli distribution with probability  $p \in [0, 1]$ . Then, for all  $\delta > 0$ ,*

$$\Pr[X_1 + \dots + X_n < (1 - \delta)pn] \leq \exp(-\delta^2 np / 2) \tag{6}$$

$$\Pr[X_1 + \dots + X_n > (1 + \delta)pn] \leq \exp(-\delta^2 np / (2 + \delta)) \tag{7}$$

The key applications of the Chernoff bound in our sparsifier constructions are as follows.

**Lemma 2.12.** *Let  $S \subseteq [m]$  be a random subset of  $[m]$  such that each element is included independently with probability at most  $1/3$ . Then,*

$$\Pr[|S| > m/2] \leq \exp(-m/30).$$

*Proof.* By a standard coupling argument, the probability that  $|S| > m/2$  is maximized when each element is included with probability exactly  $1/3$ . We then apply (7) with  $p = 1/3$ ,  $n = m$ , and  $\delta = 1/2$ . In that case,  $\delta^2 np / (2 + \delta) = m/30$ .  $\square$

**Lemma 2.13.** *Let  $S \cup T$  be a partition of  $[m]$ . For  $p \in (0, 1]$ , let  $S_p$  be a random subset of  $S$  where each element of  $S$  is included independently with probability  $p$ . For any codeword  $c \in \{0, 1\}^m$  and any  $\varepsilon \in (0, 1)$ , we have that*

$$\Pr_{S_p} \left[ \frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_T) \notin [1 - \varepsilon, 1 + \varepsilon] \cdot \text{Ham}(c) \right] < 2 \exp(-\varepsilon^2 \text{Ham}(c)p/3).$$

*Proof.* Let  $w_S = \text{Ham}(c|_S)$  and  $w_T = \text{Ham}(c|_T)$ , so  $\text{Ham}(c) = w_S + w_T$ . If  $w_S = 0$ , then  $\text{Ham}(c) = w_T$ . Thus, the probability of failure is 0. Otherwise, apply Theorem 2.11 with  $n := w_S$  and

$\delta := \varepsilon \cdot \frac{w_S + w_T}{w_S}$  (note that  $\delta$  may be larger than 1) to have that

$$\begin{aligned}
\Pr_{S_p}[\text{Ham}(c|_{S_p}) \notin [1 - \delta, 1 + \delta] \cdot pn] &< 2 \exp\left(-\frac{\delta^2 w_S p}{2 + \delta}\right) \\
&= 2 \exp\left(-\varepsilon^2 \left(\frac{w_S + w_T}{w_S}\right)^2 \frac{w_S}{2 + \varepsilon \cdot \frac{w_S + w_T}{w_S}} p\right) \\
&= 2 \exp\left(-\varepsilon^2 \frac{(w_S + w_T)^2}{2w_S + \varepsilon(w_S + w_T)} p\right) \\
&\leq 2 \exp\left(-\varepsilon^2 \frac{(w_S + w_T)^2}{3(w_S + w_T)} p\right) \\
&= 2 \exp(-\varepsilon^2 \text{Ham}(c)p/3),
\end{aligned}$$

where the second-to-last line uses the fact that  $\varepsilon \in (0, 1)$ . Finally, to complete the proof, note that

$$\begin{aligned}
\text{Ham}(c|_{S_p}) \notin [1 - \delta, 1 + \delta]pn &\iff \frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_T) \notin [1 - \delta, 1 + \delta] \cdot w_S + w_T \\
&\iff \frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_T) \notin [1 - \varepsilon, 1 + \varepsilon] \text{Ham}(c). \quad \square
\end{aligned}$$

### 3 Warmup: A Simple Sparsifier

As a warm-up for our main result, we prove that every code  $C \subseteq \{0, 1\}^m$  has an  $\varepsilon$ -sparsifier of size within  $\tilde{O}_\varepsilon(\log |C| \cdot \text{NRD}(C))$ .

**Theorem 3.1.** *For all nonempty  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ , we have that*

$$\text{SPR}(C, \varepsilon) = O(\text{NRD}(C) \log |C| (\log^3 m) / \varepsilon^2).$$

*More precisely,*

$$\text{SPR}(C, \varepsilon) \leq 36 \text{NRD}(C) \log_2(4|C|) \log_2^3(2m) / \varepsilon^2. \quad (8)$$

**Remark 3.2.** In terms of CSP sparsification, we have that for any  $R \subseteq D^r$  that

$$\text{SPR}(R, n, \varepsilon) = O(n \cdot \text{NRD}(\bar{R}, n) (\log |D|) (r \log n)^3 / \varepsilon^2).$$

since any instance  $\Psi$  of  $\text{CSP}(R)$  on  $n$  variables has at most  $n^r$  clauses and at most  $|D|^n$  assignments (so  $|C_{\Psi, R}| \leq |D|^n$ ). Ignoring logarithmic factors, we are within a factor of  $n$  of Proposition 2.5, which is highly nontrivial given the fact that a priori we only know that  $\text{SPR}(R, n, \varepsilon) \in [\Omega(n), n^r]$ . However, Theorem 3.1 has no nontrivial implications in the  $r = 2$  case.

For any  $d \in [m]$ , let  $C_{\leq d}$  be the set of codewords of  $C$  with Hamming weight at most  $d$ . The key observation of this simple sparsifier is that  $\text{NRD}(C)$  carefully controls the support size of  $C_{\leq d}$ .

**Lemma 3.3.** *For all  $C \subseteq \{0, 1\}^m$  and all  $d \in \{0, 1, \dots, m\}$ , we have that*

$$|\text{supp}(C_{\leq d})| \leq d \cdot \text{NRD}(C).$$

**Remark 3.4.** For linear codes, this lemma is trivial: the non-redundancy of a linear code is precisely its dimension and each time we expand the support of a linear code with a new codeword, the rank increases by 1. A variant of this observation appears in [KPS24a] as part of their divide-and-conquer framework for linear code sparsification.

*Proof.* We prove this result by induction on  $d$ . The base case of  $d = 0$  is trivial. Now assume  $d \geq 1$ . Pick  $I \subseteq [m]$  minimal such that<sup>17</sup> for every non-zero  $c \in C$ ,  $|\text{supp}(c) \cap I| \geq 1$ . We claim that  $I$  is non-redundant. If not, there exists  $i \in I$  such that for every  $c \in C$  with  $c_i = 1$ , we have that  $|\text{supp}(c) \cap I| \geq 2$ . In that case,  $I \setminus \{i\}$  is a smaller set with the prescribed property, contradicting  $I$ 's minimality. Therefore,  $I$  is non-redundant and

$$(C_{\leq d})|_I \subseteq (C|_I)_{\leq d-1}.$$

Therefore, by the induction hypothesis,

$$|\text{supp}(C|_I)_{\leq d-1}| \leq (d-1) \text{NRD}(C|_I) \leq (d-1) \text{NRD}(C).$$

Thus, since  $I$  is non-redundant,

$$|\text{supp}(C_{\leq d})| \leq |I| + |\text{supp}(C|_I)_{\leq d-1}| \leq \text{NRD}(C) + (d-1) \text{NRD}(C) = d \text{NRD}(C). \quad \square$$

We now proceed to construct the sparsifier. We use a fairly standard divide-and-conquer technique (e.g., [KPS24a, KPS25a]).

*Proof of Theorem 3.1.* We seek to prove (8) by strong induction on  $m$ . The base cases of  $m \leq 36$  is trivial. Define  $\varepsilon_0 := \varepsilon / (2 \log_2(2m))$  and

$$d_0 := 9 \log_2(4|C|) / \varepsilon_0^2 = 36 \log_2(4|C|) \log_2^2(2m) / \varepsilon^2.$$

Let  $T = \text{supp}(C_{\leq d_0})$  and  $S = [m] \setminus T$ . By Lemma 3.3, we have that  $|T| \leq d_0 \cdot \text{NRD}(C)$ . Let  $p := 1/3$  and let  $S_p$  be a random subset of  $S$  where each element is included independently with probability  $p$ . We seek to show that with positive probability, we have that  $|S_p| \leq m/2$  and for all  $c \in C$ ,

$$\frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_T) \in [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c). \quad (9)$$

First, by Lemma 2.12 we have that  $\Pr[|S_p| > m/2] < \exp(-m/30) < 1/3$  since  $m > 36$ . To prove (9), first observe that if  $\text{Ham}(c) \leq d_0$ , then  $\text{supp}(c) \subseteq T$ , so (9) trivially follows. Otherwise, by Lemma 2.13, the probability that (9) fails is at most

$$2 \exp\left(-\frac{\varepsilon_0^2 \text{Ham}(c)}{9}\right) < 2 \exp\left(-\frac{d_0 \varepsilon_0^2}{9}\right) = 2 \exp(-\log_2(4|C|)) \leq \frac{1}{2|C|}.$$

In particular, by the union bound the probability that (9) holds for all  $c \in C$  is at least  $1/2$ . Therefore, we have proved the existence of  $S' \subseteq S$  of size at most  $m/2$  such that for all  $c \in C$ , the map

$$\forall i \in [m], w(i) := \begin{cases} 1 & i \in T \\ 3 & i \in S' \\ 0 & i \in S \setminus S', \end{cases}$$

---

<sup>17</sup>This is known as a *hitting set* (e.g., [ERS05]).

is an  $\varepsilon_0$ -sparsifier of  $C$ . We now apply the induction hypothesis to  $C' := C|_{S'}$  to find an  $\varepsilon' := \varepsilon - 2\varepsilon_0$ -sparsifier  $w' : S' \rightarrow \mathbb{R}_{\geq 0}$  of size at most  $36 \text{NRD}(C') \log_2(4|C'|) \log_2^3(2|\text{supp } C'|)/\varepsilon'^2$ . Now define the sparsifier  $\tilde{w} : [m] \rightarrow \mathbb{R}_{\geq 0}$  by

$$\forall i \in [m], \tilde{w}(i) := \begin{cases} 1 & i \in T \\ 3w'(i) & i \in S' \\ 0 & i \in S \setminus S'. \end{cases}$$

Observe that for any  $c \in C$ , we have that

$$\begin{aligned} \langle \tilde{w}, c \rangle &= 3\langle w', c|_{S'} \rangle + \text{Ham}(c|_T) \\ &\in 3[1 - \varepsilon', 1 + \varepsilon'] \cdot \text{Ham}(c|_{S'}) + \text{Ham}(c|_T) \\ &\subseteq [1 - \varepsilon', 1 + \varepsilon'] \cdot \langle w, c \rangle \\ &\subseteq [1 - \varepsilon', 1 + \varepsilon'] \cdot [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c) \\ &\subseteq [1 - \varepsilon, 1 + \varepsilon] \cdot \text{Ham}(c), \end{aligned}$$

where the second-to-last inclusion follows from  $w$  being an  $\varepsilon_0$ -sparsifier. Thus,  $\tilde{w}$  is an  $\varepsilon$ -sparsifier. To finish, we bound the size of  $\tilde{w}$  as follows:

$$\begin{aligned} |\text{supp}(\tilde{w})| &= |\text{supp}(w')| + |T| \\ &\leq \frac{36 \text{NRD}(C') \log_2(4|C'|) \log_2^3(2|\text{supp } C'|)}{\varepsilon'^2} + \frac{36 \text{NRD}(C) \log_2(4|C|) \log_2^2(2m)}{\varepsilon^2} \\ &\leq 36 \text{NRD}(C) \log_2(4|C|) \left[ \frac{\log_2^3 m}{\left(\frac{\log_2 m}{\log_2(2m)} \varepsilon\right)^2} + \frac{\log_2^2(2m)}{\varepsilon^2} \right] \\ &= \frac{36 \text{NRD}(C) \log_2(4|C|) \log_2^2(2m)}{\varepsilon^2} [\log_2 m + 1] \\ &= \frac{36 \text{NRD}(C) \log_2(4|C|) \log_2^3(2m)}{\varepsilon^2}. \quad \square \end{aligned}$$

## 4 The Entropy Sparsifier

In this section, we prove Theorem 1.1. As mentioned in the introduction, this proof crucially needs bounds on the entropy of probability distributions over our non-linear code. We begin by defining the necessary notation for discussing such concepts and then proceed to analyze Gilmer's entropy method (and its subsequent refinements) and how these can assist with sparsification.

### 4.1 OR-closure, VC Dimension, and Entropy

We begin by defining a few basic properties of codes.

#### 4.1.1 OR-closure and VC Dimension

Given two strings  $a, b \in \{0, 1\}^m$ , we define their bitwise-OR  $a \vee b$  to be the string  $c \in \{0, 1\}^m$  such that  $c_i = a_i \vee b_i$  for all  $i \in [m]$ . Likewise, given a list  $a_1, \dots, a_k \in \{0, 1\}^m$ , we let  $\bigvee_{i=1}^k a_i$  denote their collective bitwise-OR.

**Definition 4.1.** Given  $C \subseteq \{0, 1\}^m$  we define its OR-closure (or OR-span) to be

$$\text{span}_{\text{OR}}(C) := \bigcup_{k=0}^{\infty} \left\{ \bigvee_{i=1}^k c_i : c_1, \dots, c_k \in C \right\},$$

where the empty OR is defined to be  $0^m$ . If  $C = \text{span}_{\text{OR}}(C)$ , we say that  $C$  is *OR-closed*.

We now discuss the VC dimension of  $C$ .

**Definition 4.2.** For  $C \subseteq \{0, 1\}^m$ , we define the VC dimension of  $C$ , denoted by  $\text{VC}(C)$ , to be the size of largest  $I \subseteq [m]$  such that  $C|_I = \{0, 1\}^I$ . In particular, any code with at most one codeword has VC dimension 0.

We now restate an observation of [BCK20] that non-redundancy of a code is precisely the VC dimension of its OR-closure.

**Proposition 4.3** (Obs. 4 [BCK20], restated). *For all  $C \subseteq \{0, 1\}^m$ ,  $\text{NRD}(C) = \text{VC}(\text{span}_{\text{OR}}(C))$ .*

*Proof.* Observe that for all  $I \subseteq [m]$ ,  $\text{span}_{\text{OR}}(C|_I) = \{0, 1\}^I$  if and only if for all  $i \in I$ , there exists  $c_i \in C_I$  such that for all  $i' \in I$ ,  $c_{i,i'} = 1$  iff  $i = i'$ . In other words,  $\text{span}_{\text{OR}}(C|_I) = \{0, 1\}^I$  if and only if  $I$  is non-redundant in  $C|_I$  (and thus  $C$ ). Since  $\text{span}_{\text{OR}}(C|_I) = \text{span}_{\text{OR}}(C)|_I$  for all  $I \subseteq [m]$ , we have  $\text{VC}(\text{span}_{\text{OR}}(C))$  is precisely the non-redundancy of  $C$ .  $\square$

Our main use of VC dimension is to bound the size of the corresponding code via the Sauer-Shelah-Peres lemma.

**Lemma 4.4** (Sauer-Shelah-Peres [Sau72, She72]). *For all  $C \subseteq \{0, 1\}^m$ ,*

$$|C| \leq \sum_{i=0}^{\text{VC}(C)} \binom{m}{i} \leq (m+1)^{\text{VC}(C)}.$$

As an immediate corollary of Proposition 4.3 and Lemma 4.4, we can control the size of  $\text{span}_{\text{OR}}(C)$  in terms of  $\text{NRD}(C)$  via the Sauer-Shelah-Peres lemma.

**Corollary 4.5.** *For any  $C \subseteq \{0, 1\}^m$ ,*

$$|\text{span}_{\text{OR}}(C)| \leq (m+1)^{\text{NRD}(C)}.$$

### 4.1.2 Entropy

We define the binary entropy function to be

$$h(x) := x \log_2(1/x) + (1-x) \log_2(1/(1-x))$$

for  $x \in (0, 1)$  with  $h(0) = h(1) = 0$  with maximum value at  $h(1/2) = 1$ . Given a probability distribution  $\mathcal{D}$  over a finite set, we define the entropy of the distribution as

$$H(\mathcal{D}) := \sum_{u \in \text{supp}(\mathcal{D})} \Pr_{\mathcal{D}}[u] \log_2 \left( \frac{1}{\Pr_{\mathcal{D}}[u]} \right).$$

Since  $x \log_2(1/x)$  is a concave function of  $x$ , we have that every distribution  $\mathcal{D}$  over a finite set  $S$  has  $H(\mathcal{D}) \leq \log_2 |S|$ , with equality when  $\mathcal{D}$  is the uniform distribution over  $S$ . As such, the following is implied by Corollary 4.5.

**Proposition 4.6.** *For any  $C \subseteq \{0, 1\}^m$  and any probability distribution  $\mathcal{D}$  over  $\text{span}_{\text{OR}}(C)$ , we have that*

$$H(\mathcal{D}) \leq \text{NRD}(C) \cdot \log_2(m + 1). \quad (10)$$

Given a distribution  $\mathcal{D}$  over  $\{0, 1\}^m$ , we define  $\mathbb{E}[\mathcal{D}] \in [0, 1]^m$  to be the coordinate-wise expected value of the distribution. Given  $I \subseteq [m]$ , we let the punctured distribution  $\mathcal{D}|_I$  to be

$$\Pr_{\mathcal{D}|_I}[c] := \sum_{\substack{c' \in \{0, 1\}^m \\ c'|_I = c}} \Pr_{\mathcal{D}}[c'].$$

We next discuss the significance of entropy to the structure of codes.

## 4.2 Gilmer’s Entropy Method

In 2022, Gilmer [Gil22] made the following breakthrough on the structure of union-closed sets, which we restate in terms of the OR-closed codes.

**Theorem 4.7** ([Gil22], restated). *Let  $C \subseteq \{0, 1\}^m$  be OR-closed, let  $\mathcal{D}$  be the uniform distribution over  $C$ . Then, there exists  $i \in [m]$  for which  $\mathbb{E}[\mathcal{D}]_i \geq 1/100$ .*

As mentioned in the introduction, numerous technical improvements have now risen the constant  $1/100$  to more than  $0.382$  [AHS24, CL22, Peb22, Saw23, Cam22], with  $1/2$  still conjectured as being the best possible improvement. Of note, the previous best version of Theorem 4.7 was a lower bound of  $\Omega(1/\log_2 |C|)$  [Kni94, Wój99, Gil22]. We now give some intuition as to why the  $\Theta(\log_2 |C|)$  “gain” in Gilmer’s theorem is essentially the same  $\Theta(\log_2 |C|)$  we seek to shave in Theorem 3.1.

From the perspective of code sparsification, Gilmer’s theorem appears quite useful, as adding the coordinate  $i$  with  $\mathbb{E}[\mathcal{D}]_i \geq 1/100$  to our sparsifier allowed us to make nontrivial progress on sparsifying a constant fraction of  $\text{span}_{\text{OR}}(C)$ . If we make the (bold) assumption that this constant-fraction of codewords need not be revisited, we can recursively apply Gilmer’s theorem to the remaining  $99/100$  fraction of  $\text{span}_{\text{OR}}(C)$ . Then, since  $\log_2 |\text{span}_{\text{OR}}(C)| \leq \text{NRD}(C) \cdot \log_2(m + 1)$ , we will have “sparsified” all of  $\text{span}_{\text{OR}}(C)$  (and thus  $C$ ) using only  $\tilde{O}(\text{NRD}(C))$  coordinates.

Although this informal sketch has serious technical issues (e.g., why can we “forget” a codeword after saving a single coordinate?), it turns out that we can make a rigorous argument affirming this intuition. For simplicity, assume all codewords in  $C \subseteq \{0, 1\}^m$  have the same Hamming weight  $d$ . What we eventually prove (see Theorem 4.16) is that for any  $\lambda \geq 1$  there exist  $\approx \lambda \text{NRD}(C)$  coordinates  $I \subseteq [m]$  such that the punctured code  $C|_{\bar{I}}$  has very few codewords (i.e., at most  $\approx \exp(d/\lambda)$ ). In other words, the coordinates of  $I$  capture the most important distinguishing features among codewords of  $C$ . We can then subsample the remaining coordinates of  $C$  to make progress toward the sparsifier, like in one step of the recursive argument proving Theorem 3.1.

We note that the linear code sparsifiers of Khanna, Putterman, and Sudan [KPS24a, KPS25a] also prove a similar bound (e.g., see Theorem 2.2 in [KPS24a]). However, the proof methods are very different. In [KPS24a], they recursively remove “dense” subcodes of their linear code and then use a “Benczúr-Karger-style” contraction algorithm to prove the result code is sparse. In our case, we instead simultaneously understand the dense and sparse structures of our code  $C$  using Gilmer’s entropy method. In particular, by applying the minimax theorem in Proposition 4.14, we show

that our code either has a probability distribution over it in which each coordinate is unlikely to equal 1 (“ $\theta$ -sparse”) or there is probability distribution over the coordinates has nontrivial overlap with *every* codeword (“ $\theta$ -cover”). If there is a “ $\theta$ -cover,” we prove we can add a coordinate to our sparsifier which makes substantial progress toward properly sparsifying every codeword. Conversely, if there is a  $\theta$ -sparse distribution, we use Gilmer’s entropy method to prove that the “important part” of the distribution is quite sparse, and thus can be removed. We apply Proposition 4.14 many times to repeatedly refine our understanding of sparse and dense structures in  $C$ . When the process terminates, we obtain the decomposition theorem (Theorem 4.16), from which the remainder of the proof follows using standard techniques.

**The Entropy Bound.** We now return to discussing concrete technical details. The key entropy bound we use is a refined version of Gilmer’s method due to Sawin [Saw23]. Like for Theorem 4.7, we restate the bound in terms of non-linear codes.

**Theorem 4.8** (Theorem 2, [Saw23], adapted). *Let  $\mathcal{D}$  be a distribution on  $\{0,1\}^m$  such that  $\max_{i \in [m]} \mathbb{E}[\mathcal{D}]_i \leq p \leq \frac{3-\sqrt{5}}{2} \approx 0.3819$ . Let  $A, B \sim \mathcal{D}$  be i.i.d. samples from  $\mathcal{D}$ . Then,*

$$H(A \vee B) \geq H(A) \cdot \frac{h(1 - (1 - p)^2)}{h(p)}. \quad (11)$$

Note that (11) is tight when  $\mathcal{D}$  is the product of  $m$  independent Bernoulli distributions with probability  $p$ . The key idea used by Gilmer and others to prove results like (11) is to inductively show for  $i \in [n]$  that

$$H(A_i \vee B_i | A_{[i-1]}, B_{[i-1]}) \geq H(A_i | A_{[i-1]}) \cdot \frac{h(1 - (1 - p)^2)}{h(p)}. \quad (12)$$

The crucial property is that the conditional random variable  $\mathbb{E}[A_i | A_{[i-1]}]$  is independently and identically distributed to  $\mathbb{E}[B_i | B_{[i-1]}]$ . Thus, (12) can be reduced to proving a simple (but ingenious) inequality about two i.i.d.  $[0, 1]$ -valued random variables with mean at most  $p$ . See [Saw23] and related works (such as the exposition by Boppana [Bop23]) for more details.

Although the application of (11) to the union-closed sets conjecture is mostly interested in the regime in which  $p$  is a constant, we shall focus on using this bound in the regime in which  $p$  tends to 0 as  $p \rightarrow \infty$ . We recursively apply Theorem 4.8 to amplify the growth in entropy.

**Corollary 4.9.** *Let  $\mathcal{D}$  be a distribution over  $\{0,1\}^m$  with  $\max_{i \in [m]} \mathbb{E}[\mathcal{D}]_i \leq p$ . Let  $N$  be a power of two such that  $1 - (1 - p)^{N/2} \leq \frac{3-\sqrt{5}}{2}$ . Let  $A_1, \dots, A_N$  be i.i.d. samples from  $\mathcal{D}$ , then*

$$H\left(\bigvee_{i=1}^N A_i\right) \geq H(\mathcal{D}) \cdot \frac{h(1 - (1 - p)^N)}{h(p)}.$$

**Remark 4.10.** We only prove our result when  $N$  power of two as more precision would not significantly improve our main result. See Conjecture 1 and Lemma 26 in [Wak24] for a comparable bound for all  $N$ .

**Remark 4.11.** Note that if  $p = \Theta(1/N)$ , then  $h(1 - (1 - p)^N)/h(p) \approx \frac{N}{\log N}$  (see Lemma 4.12), so we get a nearly factor of  $N$  boost in the entropy. As we shall soon see, this “boost” is analogous to the exponential savings in the Benczúr-Karger cut bound.

*Proof.* We prove this by induction on  $\log_2 N$ . If  $N = 2$ , we apply Theorem 4.8 directly. Otherwise, we consider the random variable  $\mathcal{D}_{N/2}$  corresponding to the distribution  $\bigvee_{i=1}^{N/2} A_i$ , and note that by independence, for all  $i \in [m]$ ,

$$\Pr_{c \sim \mathcal{D}_{N/2}} [c_i = 1] = 1 - \Pr_{c \sim \mathcal{D}} [c_i = 0]^{N/2} \leq 1 - (1-p)^{N/2} \leq \frac{3 - \sqrt{5}}{2}.$$

Therefore,

$$\begin{aligned} H\left(\bigvee_{i=1}^N A_i\right) &\geq H(\mathcal{D}_{N/2}) \cdot \frac{h(1 - (1-p)^N)}{h(1 - (1-p)^{N/2})} \\ &\geq H(\mathcal{D}) \cdot \frac{h(1 - (1-p)^{N/2})}{h(p)} \cdot \frac{h(1 - (1-p)^N)}{h(1 - (1-p)^{N/2})}, \\ &= H(\mathcal{D}) \cdot \frac{h(1 - (1-p)^N)}{h(p)}, \end{aligned}$$

where the second line invokes the induction hypothesis because  $1 - (1-p)^{N/4} \leq 1 - (1-p)^{N/2} \leq \frac{3 - \sqrt{5}}{2}$ .  $\square$

### 4.3 Improved Entropy Bound for $\theta$ -sparse Distributions

We say that a distribution  $\mathcal{D}$  over  $\{0, 1\}^m$  is  $\theta$ -sparse if for all  $\max_{i \in [m]} \mathbb{E}[\mathcal{D}]_i \leq 1/\theta$ . By using Gilmer's entropy method, we can show that the entropy of any  $\theta$ -sparse distributions is approximately  $\frac{\log \theta}{\theta}$  times smaller than the RHS of (10). In some sense, this bound is similar to [KPS24a]'s code-counting lemma for sparse linear codes, although their bound is proved using a Benczúr-Karger-style contraction algorithm whereas ours is proved using Gilmer's entropy method.

We begin by proving the following bound for any distribution over any OR-closed code.

**Lemma 4.12.** *Let  $C \subseteq \{0, 1\}^m$  be an OR-closed code and let  $\theta \geq 1$ . Let  $\mathcal{D}$  be a  $\theta$ -sparse distribution over  $C$ . Then,*

$$H(\mathcal{D}) \leq \frac{3 \log_2(3\theta)}{\theta} \cdot \log_2 |C|. \quad (13)$$

*Proof.* If  $\theta < 4$ , we directly see that

$$H(\mathcal{D}) \leq \log_2 |C| \leq \frac{3 \log_2(3\theta)}{\theta} \cdot \log_2 |C|.$$

Otherwise, let  $p := 1/\theta$  and  $N := 2^{\lceil \log_2 \theta \rceil - 1} \in (\theta/4, \theta/2]$ . Since  $\theta \geq 4$ , we have that  $N \geq 2$ . By Bernoulli's inequality,

$$\begin{aligned} 1 - (1-p)^{N/2} &\leq pN/2 \leq 1/4 < \frac{3 - \sqrt{5}}{2}, \text{ and} \\ 1 - (1-p)^N &\leq pN \leq 1/2. \end{aligned}$$

Furthermore, since  $1 - p \leq e^{-p}$  and  $pN > 1/4$ , we have that

$$1 - (1-p)^N \geq 1 - e^{-pN} > 1 - e^{-1/4} > 0.22.$$

Thus, since the binary entropy function is increasing in the range  $[0, 1/2]$ ,

$$h(1 - (1 - p)^N) > h(0.22) > 1/3.$$

Therefore, by Corollary 4.9 and the fact that  $C$  is OR-closed, we have that if  $A_1, \dots, A_N$  are i.i.d. samples of  $\mathcal{D}$ , then

$$\begin{aligned} \log_2 |C| &\geq H\left(\bigvee_{i=1}^N A_i\right) \\ &\geq H(\mathcal{D}) \cdot \frac{h(1 - (1 - p)^N)}{h(p)} \\ &\geq H(\mathcal{D}) \cdot \frac{1/3}{\log_2(3\theta)/\theta}. \end{aligned}$$

Thus, (13) holds.  $\square$

As an immediate corollary of Lemma 4.12, we can bound the entropy of any  $\theta$ -sparse distribution in terms of the support's non-redundancy.

**Corollary 4.13.** *Let  $C \subseteq \{0, 1\}^m$  and  $\theta \geq 1$ . For any  $\theta$ -sparse distribution  $\mathcal{D}$  over  $C$  we have that*

$$H(\mathcal{D}) \leq H(C, \theta) := \frac{3 \log_2(3\theta)}{\theta} \cdot \text{NRD}(C) \cdot \log_2(m + 1). \quad (14)$$

*Proof.* By Lemma 4.12, Lemma 4.4, and Proposition 4.3, we have that

$$\begin{aligned} H(\mathcal{D}) &\leq \frac{3 \log_2(3\theta)}{\theta} \cdot \log_2 |\text{span}_{\text{OR}} C| \\ &= \frac{3 \log_2(3\theta)}{\theta} \cdot \text{VC}(\text{span}_{\text{OR}} C) \cdot \log_2(m + 1) \\ &= \frac{3 \log_2(3\theta)}{\theta} \cdot \text{NRD}(C) \cdot \log_2(m + 1) \end{aligned}$$

can directly apply Proposition 4.6 to obtain that  $\square$

As an informal application, for  $d \leq \text{NRD}(C)$ , let  $C_d$  be the set of codewords of weight exactly  $d$  in  $C$ . Recall that the warm-up sparsifier (Theorem 3.1) adds the entire support of  $C_d$  to the sparsifier, so we are currently getting no nontrivial savings. However, if the uniform distribution  $\mathcal{D}$  over  $C_d$  is  $\theta$ -sparse, then we know that  $|C_d| = H(\mathcal{D}) \leq \tilde{O}(\text{NRD}(C)/\theta)$ . As such, by subsampling  $\text{supp}(C_d)$  at a rate of  $\tilde{O}_\varepsilon(1/\theta)$ , we can still get an  $\varepsilon$ -sparsifier for  $C_d$ , beating the bound given by Lemma 3.3. In the next section, we discuss the scenario in which  $\mathcal{D}$  (or in fact any other distribution on  $C_d$ ) is not  $\theta$ -sparse. In that case, we show using the minimax theorem that the codewords of  $C_d$  have a “ $\theta$ -cover.”

## 4.4 Minimax

For any  $\theta \geq 1$ , we say that a code  $C \subseteq \{0, 1\}^m$  has a  $\theta$ -cover if there exists a probability distribution  $\mathcal{Q}$  over  $[m]$  such that

$$\forall c \in C, \mathbb{E}_{i \sim \mathcal{Q}} [c_i] \geq 1/\theta. \quad (15)$$

By definition, code  $\{\}$  has a  $\theta$ -cover for all  $\theta \geq 1$ , while the other trivial code  $\{0^m\}$  has no  $\theta$ -cover for all  $\theta \geq 1$ . We now use the minimax theorem [vN28] to show that  $\theta$ -covers and  $\theta$ -sparse distributions are dual to each other.

**Proposition 4.14.** *For every  $C \subseteq \{0, 1\}^m$  and every  $\theta \geq 1$ , at least one of the following is true.*

- (1)  $C$  has a  $\theta$ -cover.
- (2) There exists a probability distribution over  $C$  which is  $\theta$ -sparse.

*Proof.* Since the empty code has a  $\theta$ -cover for all  $\theta \geq 1$ , we may assume that  $C$  is nonempty. Consider the following zero-sum game. Have Alice pick a codeword  $c \in C$  and Bob pick an index  $i \in [m]$ . The payoff of the game is  $c_i$ , which Alice is seeking to minimize and Bob is seeking to maximize. By the minimax theorem [vN28], there exists  $\eta \in [0, 1]$ , a probability distribution  $\mathcal{P}$  over  $C$ , and a probability distribution  $\mathcal{Q}$  over  $[m]$  with the following properties.

$$\forall i \in [m], \mathbb{E}_{c \sim \mathcal{P}} [c_i] \leq \eta \quad (16)$$

$$\forall c \in C, \mathbb{E}_{i \sim \mathcal{Q}} [c_i] \geq \eta \quad (17)$$

In particular, if  $1/\theta \geq \eta$ , then (16) implies that  $\mathcal{P}$  is a  $\theta$ -sparse distribution supported on  $C$ . Otherwise, if  $1/\theta \leq \eta$ , then (17) implies that  $\mathcal{Q}$  is a  $\theta$ -cover of  $C$ .  $\square$

We now use Proposition 4.14 to prove that if we remove a “small” number of codewords from  $C$ , then the remainder of  $C$  has a  $\theta$ -cover.

**Lemma 4.15.** *For any  $C \subseteq \{0, 1\}^m$  and  $\theta \geq 1$ . Recall that  $H(C, \theta)$  is the RHS of (14). There exists  $S \subseteq C$  of size at most  $\exp_2(H(C, \theta))$  such that  $C \setminus S$  has a  $\theta$ -cover.*

*Proof.* Intuitively, we seek to iteratively delete some codewords of  $C$  until Proposition 4.14 guarantees the remainder of  $C$  has a  $\theta$ -cover. Each codeword removed has a corresponding  $\theta$ -sparse distribution associated to it. To prove not too many codewords are deleted, we take a (careful) weighted average of the  $\theta$ -sparse distributions in such a way that every removed codeword has equal probability of being sampled. In that case, the entropy of this  $\theta$ -sparse distribution is at least the logarithm of the number of codewords removed, so we can bound the number of codewords by Corollary 4.13.

Given a probability distribution  $\mathcal{D}$  over  $C$ , we let  $\mu_{\mathcal{D}} : C \rightarrow R_{\geq 0}$  be the probability density function (PDF) of  $\mathcal{D}$ . That is,  $\mu_{\mathcal{D}}(c)$  is the probability that  $c$  is sampled from  $\mathcal{D}$ . We say that  $\mu_{\mathcal{D}}$  is a  $\theta$ -sparse PDF if  $\mathcal{D}$  is  $\theta$  sparse. Further, we define  $\max \mu$  to be the highest probability attained by  $\mu$ , and  $\arg \max \mu := \{c \in C : \mu(c) = \max \mu\}$ .

If  $C$  has a  $\theta$ -cover, take  $S = \emptyset$ . Otherwise, by Proposition 4.14, there is a  $\theta$ -sparse PDF  $\mu_1$  supported on  $C$ . We now inductively define a sequence of probability distributions  $\nu_1, \dots, \nu_t$ , where  $t \leq |C|$  is decided by a stopping criterion. As our base case, we define  $\nu_1 := \mu_1$ .

Now repeat the following process. Consider  $i \in \{1, 2, \dots\}$ . Let  $S_i := \operatorname{argmax} \nu_i$ . If  $C \setminus S_i$  has a  $\theta$ -cover, we stop, and set  $t := i$ . Otherwise, if  $C \setminus S_i$  has no  $\theta$ -cover, by Proposition 4.14, there exist a  $\mu_{i+1}$  be a  $\theta$ -sparse PDF on  $C \setminus S_i$ . By convention, we assume that  $\mu_i(c) = 0$  for all  $c \in S_i$ . Let  $p_i := \max \nu_i$ , and for every  $c \in C \setminus S_i$  define

$$q_c := \frac{p_i - \nu_i(c)}{p_i - \nu_i(c) + \mu_{i+1}(c)}.$$

Note that  $q_c \in (0, 1]$  because  $p_i > \nu_i(c)$  by definition of  $S_i$ . Let  $q := \min_{c \in C \setminus S_i} q_c$ . Further define  $\nu_{i+1} := (1 - q)\nu_i + q\mu_{i+1}$  and  $S_{i+1} := \operatorname{argmax} \nu_{i+1}$ . We claim that

$$S_{i+1} = S_i \cup \operatorname{argmin}\{q_c : c \in C \setminus S_i\} \supseteq S_i. \quad (18)$$

To see why, first note that for all  $c \in S_i$ ,

$$\nu_{i+1}(c) = (1 - q)\nu_i(c) + q\mu_{i+1}(c) = (1 - q)p_i.$$

Likewise, for all  $c \in \operatorname{argmin}\{q_c : c \in C \setminus S_i\}$ , we have that

$$\begin{aligned} \nu_{i+1}(c) &= (1 - q)\nu_i(c) + q\mu_{i+1}(c) \\ &= \nu_i(c) + q(\mu_{i+1}(c) - \nu_i(c)) \\ &= \frac{(p_i - \nu_i(c) + \mu_{i+1}(c))\nu_i(c) + (p_i - \nu_i(c))(\mu_{i+1}(c) - \nu_i(c))}{p_i - \nu_i(c) + \mu_{i+1}(c)} \\ &= \frac{p_i\mu_{i+1}(c)}{p_i - \nu_i(c) + \mu_{i+1}(c)} \\ &= (1 - q)p_i. \end{aligned}$$

However, for all other  $c \in C \setminus S_i$ , we have that

$$\begin{aligned} \nu_{i+1}(c) &= (1 - q)\nu_i(c) + q\mu_{i+1}(c) \\ &= (1 - q_c)\nu_i(c) + q_c\mu_{i+1}(c) + (q_c - q)(\nu_i(c) - \mu_{i+1}(c)) \\ &= \frac{(p_i - \nu_i(c) + \mu_{i+1}(c))\nu_i(c) + (p_i - \nu_i(c))(\mu_{i+1}(c) - \nu_i(c))}{p_i - \nu_i(c) + \mu_{i+1}(c)} + (q_c - q)(\nu_i(c) - \mu_{i+1}(c)) \\ &= (1 - q_c)p_i + (q_c - q)(\nu_i(c) - \mu_{i+1}(c)) \\ &< (1 - q_c)p_i + (q_c - q)(p_i - 0) \\ &= (1 - q)p_i, \end{aligned}$$

where the inequality uses the fact that  $q_c - q > 0$ ,  $p_i > \nu_i(c)$ , and  $\mu_{i+1}(c) \geq 0$ . Therefore, we have proved (18), so  $S_i := \operatorname{argmax} \nu_i$  is strictly increasing as  $i$  increases.

Since  $C$  is finite in size and the empty code has a  $\theta$ -cover, this process must terminate in  $t \leq |C|$  steps. That is,  $C \setminus S_t$  has a  $\theta$ -cover. It suffices to prove that  $|S_t| \leq 2^{H(C, \theta)}$ , where  $H(C, \theta)$  is the RHS of (14). Since  $S_t = \operatorname{argmax} \nu_t$ , we have that  $\max \nu_t \leq \frac{1}{|S_t|}$ , so

$$H(\nu_t) = \sum_{c \in \operatorname{supp}(\nu_t)} \nu_t(c) \log_2 \left( \frac{1}{\nu_t(c)} \right) \geq \sum_{c \in \operatorname{supp}(\nu_t)} \nu_t(c) \log_2 |S_t| = \log_2 |S_t|.$$

Further,  $\nu_t$  is a convex combination of  $\theta$ -sparse distributions  $\mu_1, \dots, \mu_t$ , so  $\nu_t$  is also a  $\theta$ -sparse over  $C$ . Thus, by Corollary 4.13, we have that

$$\log_2 |S_t| \leq H(\nu_t) \leq H(C, \theta),$$

as desired. □

We next utilize Lemma 4.15 to prove our key decomposition theorem.

## 4.5 Decomposition Theorem

In Khanna, Putterman, and Sudan’s linear code sparsification [KPS24a], their key technical result (their Theorem 2.2) is that after puncturing a “dense” collection of coordinates, the remaining code has a sparsity property. Their proof proceeds by recursively removing dense subcodes until the remainder is sparse. However, as we shall see, the situation is a bit more nuanced for non-linear codes. Instead, we alternate between removing sparse and dense portions of our code using Lemma 4.15. By keeping track of a suitable monovariant, we can carefully control the tradeoff between the dense and sparse portions of our decomposition. More precisely, by conditioning on  $\approx \lambda \text{NRD}(C)$  coordinates to puncture, we save a factor of approximately  $\lambda$  in the exponent on the number of codewords remaining.

**Theorem 4.16.** *For any  $C \subseteq \{0, 1\}^m$ , integer  $d \geq 1$ , and real  $\lambda \geq 1$ , there exists  $I \subseteq [m]$  of size at most  $2\lambda \text{NRD}(C) \log_2(4m)$  such that  $C_{\leq d|_{\bar{I}}}$  has at most  $m \cdot \exp_2(3d \log_2^2(2m)/\lambda)$  codewords.*

*Proof.* Since deleting codewords can only decrease the non-redundancy, assume without loss of generality that  $C = C_{\leq d}$ . We first handle a few simple edge cases.

**Case 1,**  $d \geq \lambda \text{NRD}(C)$ . In this case, we let  $I := \emptyset$  and note that by Lemma 4.4,  $C$  has at most

$$\exp_2(\text{NRD}(C) \cdot \log_2(m+1)) \leq \exp_2(d \log_2(m+1)/\lambda) \leq m \cdot \exp_2(3d \log_2^2(2m)/\lambda)$$

codewords, as desired.<sup>18</sup>

**Case 2,**  $m \leq 2\lambda \text{NRD}(C) \log_2(4m)$ . In this case, we set  $I = [m]$  and observe that  $C|_{\bar{I}}$  has at most 1 codeword (empty string).

We now move onto the main case.

**Case 3,**  $d < \lambda \text{NRD}(C)$  and  $m > 2\lambda \text{NRD}(C) \log_2(4m)$ .

Let  $\theta := \lambda \text{NRD}(C)/d > 1$ . We inductively build sequences  $\emptyset = I_0 \subsetneq \dots \subsetneq I_t \subseteq [m]$  and  $\emptyset = T_0 \subseteq \dots \subseteq T_t \subseteq C$  for a currently unspecified  $t \in \mathbb{N}$  such that for all  $j \in \{0, 1, \dots, t\}$ ,

$$|I_j| = j \tag{19}$$

$$|T_j|_{I_j} \leq j \cdot \exp(H(C, \theta)) \tag{20}$$

We build the sequence using the following (randomized) procedure.

- **Base Case:** Set  $I_0 := \emptyset$  and  $T_0 := \emptyset$ .

---

<sup>18</sup>This easy case is similar to code-counting bound for the proof of Theorem 3.1.

- **For**  $j \in \{0, 1, \dots, m-1\}$ ,
  - **If**  $C = T_j$ , **halt** with  $t := j$ .
  - Let  $C_j := (C \setminus T_j)|_{\overline{I_j}}$ .
  - Invoke Lemma 4.15 to find  $S_{j+1} \subseteq C_j$  of size at most  $\exp_2(H(C, \theta))$  such that  $C_j \setminus S_{j+1}$  has a  $\theta$ -cover  $\mathcal{Q}_{j+1}$ .
  - Randomly sample  $i_{j+1}$  from  $\mathcal{Q}_{j+1}$ .
  - Set  $I_{j+1} := I_j \cup \{i_{j+1}\}$ .
  - Set  $T_{j+1} := T_j \cup \{c \in C : c|_{\overline{I_j}} \in S_{j+1}\}$ .

As written, it seems possible that we may exhaust the **For** loop without meeting the halt condition. However, we shall prove in Claim 4.17 that with positive probability, we meet the halt condition for some  $j = t \leq 2\lambda \text{NRD}(C) \log_2(4m) < m$ . Assuming this is the case, it suffices to show that  $C|_{\overline{I_t}}$  has at most  $m \cdot \exp_2(3d \log_2^2(2m)/\lambda)$  codewords. By the halting condition, we have that  $C = T_t$ . Therefore, we have that

$$C|_{\overline{I_t}} = \bigcup_{i=1}^t S_i|_{\overline{I_t}}.$$

Hence, the size of  $C|_{\overline{I_t}}$  is at most

$$\begin{aligned} \sum_{i=1}^t |S_i| &\leq \sum_{i=1}^t H(C, \theta) \\ &\leq t \cdot \exp_2\left(\frac{3 \log_2(3\theta)}{\theta} \cdot \text{NRD}(C) \cdot \log_2(m+1)\right) \\ &\leq m \cdot \exp_2\left(\frac{3d \log_2(3\lambda \text{NRD}(C)/d)}{\lambda} \cdot \log_2(2m)\right), \end{aligned}$$

recalling that  $\theta = \frac{\lambda \cdot \text{NRD}(C)}{d}$ . By assumption on the size of  $m$ , we have that

$$\frac{3\lambda \text{NRD}(C)}{d} \leq \frac{3m}{2 \log_2(4m)d} < 2m.$$

Thus,  $C|_{\overline{I_t}}$  has size at most  $m \cdot \exp_2(3d \log_2^2(2m)/\lambda)$ , as desired. To finish, we prove Claim 4.17.

**Claim 4.17.** *With nonzero probability, the procedure halts after at most  $2\lambda \text{NRD}(C) \log_2(4m) < m$  steps.*

*Proof.* Consider the following monovariant defined for any  $I \subseteq [m]$  and  $T \subseteq [m]$ .

$$f(I, T) := \sum_{c \in C \setminus T} \exp_2(\text{Ham}(c|_{\overline{I}})) \tag{21}$$

In particular, if  $I' \supseteq I$  and  $T' \supseteq T$ , then  $f(I', T') \leq f(I, T)$ . Observe that since every codeword of  $C$  has Hamming weight at most  $d$ , we have that  $|C|$  has size at most  $\binom{m}{0} + \dots + \binom{m}{d} \leq (m+1)^d$ . Thus,

$$f(I_0, T_0) = f(\emptyset, \emptyset) = \sum_{c \in C} \exp_2(\text{Ham}(c)) \leq |C| \cdot 2^d \leq (2m+2)^d \leq \exp(d \log(2m+2)).$$

Assume that for some  $t_0$ , the probability that the procedure terminates within  $t_0$  steps is zero. For any  $j \in \{0, 1, \dots, t_0 - 1\}$ , we have that

$$\begin{aligned}
\mathbb{E}_{i_{j+1} \sim \mathcal{Q}_{j+1}} [f(I_{j+1}, T_{j+1})] &= \mathbb{E}_{i \sim \mathcal{Q}_{j+1}} [f(I_j \cup \{i\}, T_{j+1})] \\
&= \sum_{c \in C_d \setminus T_{j+1}} \mathbb{E}_{i \sim \mathcal{Q}_{j+1}} [\exp_2(\text{Ham}(c|_{\overline{I_j}}) - c_i)] \\
&= \sum_{c \in C \setminus T_{j+1}} \exp_2(\text{Ham}(c|_{\overline{I_j}})) \mathbb{E}_{i \sim \mathcal{Q}} [\exp_2(-c_i)] \\
&\leq \sum_{c \in C \setminus T_{j+1}} \exp(\text{Ham}(c|_{\overline{I_j}})) \left(1 - \frac{1}{\theta} + \frac{1}{\theta} \cdot \frac{1}{2}\right) \\
&\leq \left(1 - \frac{1}{2\theta}\right) f(I_j, T_{j+1}) \leq \left(1 - \frac{1}{2\theta}\right) f(I_j, T_j),
\end{aligned}$$

Thus, with nonzero probability, we have that for all  $j \in [t_0 - 1]$

$$f(I_j, T_j) \leq \left(1 - \frac{1}{2\theta}\right)^j f(I_0, T_0) \leq \exp\left(d \log(2m + 2) - \frac{j}{2\theta}\right).$$

In particular, if  $j \geq (2 \log_2(4m) - 1)\lambda \text{NRD}(C)$ , then

$$\begin{aligned}
f(I_j, T_j) &\leq \exp(d \log(2m + 2) - (2 \log_2(4m) - 1)\lambda \text{NRD}(C)/(2\theta)) \\
&= \exp(d(\log(2m + 2) - \log_2(4m) + 1/2)) \\
&< \exp(-d/9) < 1.
\end{aligned}$$

However, each term in the sum (21) defining  $f(I_j, T_j)$  is at least 1. Therefore,  $f(I_j, T_j) = 0$  and  $C = T_j$ . Thus, the halting condition will be met at step  $j$ . In other words, with nonzero probability the halting condition must be met in

$$[(2 \log_2(4m) - 1)\lambda \text{NRD}(C)] \leq 2\lambda \text{NRD}(C) \log_2(4m) < m$$

steps, as desired. □

This completes the proof of Theorem 4.16. □

We now have all the ingredients we need to prove our main result.

## 4.6 Proof of Theorem 1.1

We now prove Theorem 1.1. We use a recursive argument similar to that of Theorem 3.1, which itself was inspired by [KPS24a].

**Theorem 4.18** (Theorem 1.1, more precise version). *For all  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ ,*

$$\text{SPR}(C, \varepsilon) \leq 800 \text{NRD}(C) \log^6(4m)/\varepsilon^2. \tag{22}$$

*Proof.* Like in the proof of Theorem 3.1, we prove (22) by strong induction on  $m$ . The base cases of  $m \leq 800$  are trivial, and if  $\text{NRD}(C) = 0$ , then  $C \subseteq \{0^m\}$  so the zero sparsifier of size 0 works.

Now consider the following parameters:

$$\begin{aligned}\varepsilon_0 &:= \frac{\varepsilon}{2 \log_2(4m)} \\ \lambda &:= \frac{100 \log_2^3(4m)}{\varepsilon_0^2} = \frac{400 \log_2^5(4m)}{\varepsilon^2} \\ d_0 &:= 2\lambda \log_2(4m) \\ \ell &:= \lceil \log_2(m/d_0) \rceil \\ p &:= \frac{1}{3}.\end{aligned}$$

By Lemma 3.3, we have that there exists  $I_0 \subseteq [m]$  of size at most  $d_0 \text{NRD}(C) = 2\lambda \text{NRD}(C) \log_2(4m)$  which contains the support of all codewords  $c \in C$  of Hamming weight at most  $d_0$ . For all  $j \in [\ell]$ , let

$$C_j := \{c \in C : \text{Ham}(c) \in (2^{j-1}d_0, 2^j d_0]\}.$$

Observe that  $C = C_{\leq d_0} \cup C_1 \cup \dots \cup C_\ell$ . Now, for each  $j \in [\ell]$ , apply Theorem 4.16 to  $C_j$  with  $d := 2^j d_0$  and our choice of  $\lambda$ . As a result, there exists  $I_j \subseteq [m]$  such that  $|I_j| \leq 2\lambda \text{NRD}(C) \log_2(4m)$  and

$$|C_j|_{\overline{I_j}} \leq m \cdot \exp_2 \left( \frac{3 \cdot 2^j d_0 \log_2^2(2m)}{\lambda} \right) \leq m \cdot \exp_2(6 \cdot 2^j \log_2^3(4m)). \quad (23)$$

Let  $I = I_0 \cup I_1 \cup \dots \cup I_\ell$ . Then,  $|I| \leq 2(\ell + 1)\lambda \text{NRD}(C) \log_2(4m) \leq 2\lambda \text{NRD}(C) \log_2^2(4m)$ . Further, for all  $j \in [\ell]$ , we have that  $|C_j|_I \leq |C_j|_{\overline{I_j}}$  and so is bounded from above by the RHS of (23).

Let  $S_p \subseteq [m] \setminus I$  be a random subset where each element is included independently with probability  $p := 1/3$ . We claim that with positive probability, we have that  $|S_p| \leq m/2$  and for all  $c \in C$ ,

$$\frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_I) \in [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c). \quad (24)$$

As in Theorem 3.1, by Lemma 2.12, we have that  $\Pr[|S| > m/2] \leq \exp(-m/30) < 1/100$  since  $m > 800$ . We now turn to proving (24). If  $c \in C_{\leq d_0}$ , then  $\text{supp}(c) \subseteq I_0 \subseteq I$ , so (24) always holds. Now pick  $i \in [\ell]$  and consider  $c \in C_j$ . By Lemma 2.13, we have that the probability that (24) does not hold is at most

$$2 \exp(-\varepsilon_0^2 \text{Ham}(c)p/3) < 2 \exp(-\varepsilon_0^2 2^{j-1} d_0/9) = 2 \exp(-(100/9)2^j \log_2^4(4m)).$$

Naively, we need to now take a union bound of size  $|C_j|$ . However, we only need to take a union bound of size  $|C_j|_{\overline{I_j}}$ . To see why, for a fixed  $\tilde{c} \in C_j|_{\overline{I_j}}$  consider any  $c \in C_j$  with  $c|_{\overline{I_j}} = \tilde{c}$  and  $\text{Ham}(c)$  is minimal. Assume that (24) holds for  $c$ , for any other  $c' \in C_j$  with  $c'|_{\overline{I_j}} = \tilde{c}$  we have that

$$\begin{aligned}\frac{1}{p} \text{Ham}(c'|_{S_p}) + \text{Ham}(c'|_I) &= \frac{1}{p} \text{Ham}(c|_{S_p}) + \text{Ham}(c|_I) + (\text{Ham}(c'|_I) - \text{Ham}(c|_I)) \\ &\in [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c) + (\text{Ham}(c') - \text{Ham}(c)) \\ &\subseteq [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c'),\end{aligned}$$

since  $\text{Ham}(c') \geq \text{Ham}(c)$ . Therefore by (23), the probability that (24) fails for any  $c \in C_j$  is at most

$$\begin{aligned} 2 \exp(-(100/9)2^j \log_2^4(4m)) \cdot m \cdot \exp_2(6 \cdot 2^j \log_2^3(4m)) &\leq 2m \exp_2(2^j \log_2^3(4m) (6 - 11 \log_2(4m))) \\ &\leq 2m \exp_2(-5 \log_2^3(4m)) \\ &\leq \frac{1}{2m}. \end{aligned}$$

In particular, the probability that (23) fails for any  $c \in C$  is at most  $\ell/(2m) \leq 1/2$ . Thus, there exists  $S \subseteq [m] \setminus I$  with  $|S| \leq m/2$  for which  $\frac{1}{p} \text{Ham}(c|_S) + \text{Ham}(c|_I) \in [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c)$  for all  $c \in C$ . To finish, let  $\varepsilon' := \varepsilon - 2\varepsilon_0$  and apply the induction hypothesis to  $\varepsilon'$ -sparsify  $C|_S$ . Let  $w : S \rightarrow \mathbb{R}_{\geq 0}$  be the resulting sparsifier. We then build a sparsifier  $\tilde{w} : [m] \rightarrow \mathbb{R}_{\geq 0}$  such that

$$\forall i \in [m], \tilde{w}(i) := \begin{cases} 1 & i \in I \\ w(i)/p = 3w(i) & i \in S \\ 0 & \text{otherwise.} \end{cases}$$

By the same logic as in the proof of Theorem 3.1, for any  $c \in C$ , we have that

$$\begin{aligned} \langle \tilde{w}, c \rangle &= 3\langle w, c|_S \rangle + \text{Ham}(c|_I) \\ &\in 3[1 - \varepsilon', 1 + \varepsilon'] \cdot \text{Ham}(c|_S) + \text{Ham}(c|_I) \\ &\subseteq [1 - \varepsilon', 1 + \varepsilon'] \cdot (3 \text{Ham}(c|_S) + \text{Ham}(c|_I)) \\ &\subseteq [1 - \varepsilon', 1 + \varepsilon'] \cdot [1 - \varepsilon_0, 1 + \varepsilon_0] \cdot \text{Ham}(c) \\ &\subseteq [1 - \varepsilon, 1 + \varepsilon] \cdot \text{Ham}(c), \end{aligned}$$

as desired. To finish, we bound the size of  $\tilde{w}$ . In particular,

$$\begin{aligned} |\text{supp}(\tilde{w})| &= |\text{supp}(w)| + |I| \\ &\leq 800 \text{NRD}(C|_S) \log_2^6(4|\text{supp}(C|_S)|) / \varepsilon'^2 + 2\lambda \text{NRD}(C) \log_2^2(4m) \\ &\leq \text{NRD}(C) \left[ \frac{800 \log_2^6(2m)}{\varepsilon^2(1 - 1/\log_2(4m))^2} + \frac{800 \log_2^5(4m)}{\varepsilon^2} \right] \\ &= \frac{800 \text{NRD}(C)}{\varepsilon^2} [\log_2^4(2m) \log_2^2(4m) + \log_2^5(4m)] \\ &\leq \frac{800 \text{NRD}(C) \log_2^5(4m)}{\varepsilon^2} [\log_2(2m) + 1] \\ &= \frac{800 \text{NRD}(C) \log_2^6(4m)}{\varepsilon^2}. \quad \square \end{aligned}$$

## 5 Extension to Weighted Sparsification

We now prove our main result for sparsifying *weighted* CSPs. More precisely, given a predicate  $R \subseteq D^r$ , a weighted instance of  $\text{CSP}(R)$  can be viewed as an ordinary instance  $\Psi := (X, Y \subseteq X^r)$  of  $\text{CSP}(R)$  along with a weight function  $w : Y \rightarrow \mathbb{R}_{>0}$  for the clauses. Note that we assume all

weights are nonnegative, else the set of clauses  $Y$  can be reduced in size. Given an assignment  $\sigma : X \rightarrow D$ , recall that its  $(R, w)$ -weight is defined to be

$$\text{wt}(R, \Psi, w, \sigma) := \sum_{y \in Y} w(y) \mathbf{1}[(\sigma(y_1), \dots, \sigma(y_r)) \in R].$$

Then, a  $(w, \varepsilon)$ -sparsifier of  $\Psi$  is a function  $\tilde{w} : Y \rightarrow \mathbb{R}_{\geq 0}$  such that

$$(1 - \varepsilon) \text{wt}(R, \Psi, w, \sigma) \leq \text{wt}(R, \Psi, \tilde{w}, \sigma) \leq (1 + \varepsilon) \text{wt}(R, \Psi, w, \sigma).$$

The  $\varepsilon$ -sparsity of  $(\Psi, w)$  is then the minimal support size of any  $(w, \varepsilon)$ -sparsifier of  $\Psi$ . We can now define the weighted sparsity of a predicate.

**Definition 5.1.** For any  $R \subseteq D^r$ ,  $n \in \mathbb{N}$  and  $\varepsilon \in (0, 1)$ , we define  $\text{wSPR}(R, n, \varepsilon)$  to be the maximal  $\varepsilon$ -sparsity of an weighted instance  $(\Psi, w)$  of  $\text{CSP}(R)$  on  $n$  variables.

Likewise, given a code  $C \subseteq \{0, 1\}^m$  and a weight function  $w : [m] \rightarrow \mathbb{R}_{> 0}$ , we say that  $\tilde{w} : [m] \rightarrow \mathbb{R}_{\geq 0}$  is  $(w, \varepsilon)$ -sparsifier of  $C$  is for all  $c \in C$ ,

$$(1 - \varepsilon) \langle w, c \rangle \leq \langle \tilde{w}, c \rangle \leq (1 + \varepsilon) \langle w, c \rangle, \quad (25)$$

with the  $\varepsilon$ -sparsity of  $(C, w)$  being the minimum support size of any  $(w, \varepsilon)$ -sparsifier. This leads to an analogous definition of weighted code sparsity.

**Definition 5.2.** For any  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ , we define  $\text{wSPR}(C, \varepsilon)$  to be the maximum  $\varepsilon$ -sparsity of  $(C, w)$  among all weight functions  $w : [m] \rightarrow \mathbb{R}_{> 0}$ .

**Proposition 5.3.** For all  $R \subseteq D^r$ ,  $n \in \mathbb{N}$ , and  $\varepsilon \in (0, 1)$ , we have that

$$\text{wSPR}(R, n, \varepsilon) = \max_{\substack{\Psi \text{ instance of } \text{CSP}(R) \\ \text{on } n \text{ variables}}} \text{wSPR}(C_{R, \Psi}, \varepsilon),$$

with  $C_{R, \Psi}$  defined as in Section 2.3.

*Proof.* The proof is essentially identical to the proof of (4) in Proposition 2.9. □

It is easy to see that

$$\text{wSPR}(R, n, \varepsilon) \geq \text{SPR}(R, n, \varepsilon) \geq \text{NRD}(\bar{R}, n).$$

However, we can prove a stronger (tight!) lower bound for both in terms of *chain length*. We denote the chain length of a code  $C$  by  $\text{CL}(C)$  and the chain length of a CSP predicate  $R$  by  $\text{CL}(R, n)$  for instances with  $n$  variables. These terms are defined in Section 5.1 and Section 5.2, respectively. However, we can immediately state the main results of this section.

**Theorem 5.4.** For all  $C \subseteq \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ , we have that

$$\text{CL}(C) \leq \text{wSPR}(C, \varepsilon) = O(\text{CL}(C)(\log m)^6 / \varepsilon^2).$$

**Theorem 5.5.** For all  $R \subseteq D^r$ ,  $n \in \mathbb{N}$  and  $\varepsilon \in (0, 1)$ , we have that

$$\text{CL}(\bar{R}, n) - 1 \leq \text{wSPR}(R, n, \varepsilon) = O(\text{CL}(\bar{R}, n)(r \log n)^6 / \varepsilon^2)$$

## 5.1 Chain Length of Codes

As it is more elementary to state, we begin by defining the chain length of a code.

**Definition 5.6.** Let  $C \subseteq \{0, 1\}^m$  be a code. A *chain* of length  $\ell$  is a pair of injective maps  $a : [\ell] \rightarrow [m]$  and  $c : [\ell] \rightarrow C$  such that the following conditions hold.

$$\begin{aligned} \forall i \in [\ell], c(i)_{a(i)} &= 1. \\ \forall 1 \leq i < j \leq \ell, c(j)_{a(i)} &= 0. \end{aligned}$$

The *chain length* of  $C$ , denoted by  $\text{CL}(C)$  is the maximum-length chain.

Recall that if we line up the codewords of  $C$  as rows of an  $|C| \times m$  matrix and allow arbitrary column permutations,  $\text{NRD}(C)$  is the dimension of the largest identity submatrix within that matrix. In this setup,  $\text{CL}(C)$  is the dimension of the largest upper triangular square submatrix with 1's on the diagonal. When  $C$  is a linear code, both  $\text{NRD}(C)$  and  $\text{CL}(C)$  equal the dimension of  $C$ . Further, for nontrivial  $C$ , we have that  $1 \leq \text{NRD}(C) \leq \text{CL}(C) \leq m$ .

If the set of indices  $a : [\ell] \rightarrow [m]$  is clear from context, we refer to just the list of codewords  $c : [\ell] \rightarrow C$  as the chain. We now prove that chain-length of a code is always at least its non-redundancy.

**Proposition 5.7.** For all  $C \subseteq \{0, 1\}^m$ ,  $\text{CL}(C) \geq \text{NRD}(C)$ .

*Proof.* Let  $I \subseteq [m]$  be a maximum-sized non-redundant index set of  $C$ . Let  $c : I \rightarrow C$  witness that  $I$  is non-redundant. Let  $\ell := |I|$ . For any map bijection  $a : [\ell] \rightarrow I$ , we have that  $(a, c \circ a)$  is a chain of length  $\ell$ .  $\square$

We now consider an example where  $\text{CL}(C)$  and  $\text{NRD}(C)$  can be as different as possible.

**Example 5.8.** Consider the code  $C := \{1^n 0^{m-n} : n \in [m]\}$ . It is straightforward to check that the defining codewords of  $C$  form a chain, so  $\text{CL}(C) = m$ . However,  $\text{NRD}(C) = 1$  because for any  $c, c' \in C$  either  $\text{supp}(c) \subseteq \text{supp}(c')$  or  $\text{supp}(c') \subseteq \text{supp}(c)$ . Therefore, by Theorem 1.1<sup>19</sup>  $\text{SPR}(C, \varepsilon) \leq (\log^{O(1)} m) / \varepsilon^2$ . However, if we consider the weighting  $w(i) = 2^i$ , then any pair of codewords has weight different by a factor of 2. Hence, for  $\varepsilon < 1/2$ , every  $\varepsilon$ -sparsifier must keep every coordinate of the code or else two codewords will be given identical weight by the sparsifier. Therefore,  $\text{wSPR}(C, \varepsilon) = m$ .

We now formalize this example into a lower bound of  $\text{wSPR}(C, \varepsilon)$  for every code  $C$ .

**Lemma 5.9.** For every  $C \in \{0, 1\}^m$  and  $\varepsilon \in (0, 1)$ , we have that  $\text{wSPR}(C, \varepsilon) \geq \text{CL}(C)$ .

*Proof.* Let  $\ell := \text{CL}(C)$ . By permuting  $[m]$ , we may assume without loss of generality that there exist  $c(1), \dots, c(\ell) \in C$  such that  $c(i)_j = 1$  if  $i = j$  and  $c(i)_j = 0$  if  $j < i$ . Pick  $\lambda := \frac{4}{1-\varepsilon}$  and consider the following weight function  $w : [m] \rightarrow \mathbb{R}_{>0}$ .

$$w(i) := \begin{cases} \lambda^{\ell+1-i} & i \in [\ell] \\ \frac{1}{m} & \text{otherwise} \end{cases}$$

<sup>19</sup>This can be done in a more elementary manner by subsampling the  $i$ th coordinate with probability  $p_i := \min(1, O(\frac{\log m}{i\varepsilon^2}))$  and giving it a weight of  $1/p_i$  if kept.

Observe then that for each  $i \in [\ell]$ , that

$$\langle w, c(i) \rangle \geq w(i)c(i)_i = \lambda^{\ell+1-i}, \text{ and}$$

$$\langle w, c(i) \rangle \leq \sum_{j=i}^m w(j) = \sum_{j=i}^{\ell} \lambda^{\ell+1-j} + \frac{m-\ell}{m} < 1 + \lambda + \dots + \lambda^{\ell+1-i} < \frac{1}{1-1/\lambda} \cdot \lambda^{\ell+1-i}.$$

Now consider any  $\varepsilon$ -sparsifier  $\tilde{w} : [m] \rightarrow \mathbb{R}_{\geq 0}$ . Let  $S := \text{supp}(\tilde{w})$  and for each  $i \in [\ell]$ , let  $S_i := S \cap \text{supp } c(i)$ . We seek to prove that  $|S| \geq \ell$ . To prove this, it suffices to show for all  $i \in [\ell]$  that

$$S_i \setminus \bigcup_{j=i+1}^{\ell} S_j \neq \emptyset \quad (26)$$

First, observe that since  $\langle w, c(\ell) \rangle$  has nonzero weight, so  $S_\ell \neq \emptyset$ . Now assume for sake of contradiction that for some  $i \in [\ell-1]$ , has  $S_i \subseteq S_{i+1} \cup \dots \cup S_\ell$ . Then, since  $\tilde{w}$  is an  $\varepsilon$ -sparsifier and  $\lambda := 4/(1-\varepsilon)$ ,

$$\begin{aligned} 4\lambda^{\ell-i} &= (1-\varepsilon)\lambda^{\ell+1-i} \leq (1-\varepsilon)\langle w, c(i) \rangle \leq \langle \tilde{w}, c(i) \rangle \\ &= \sum_{a \in S_i} \tilde{w}(a) && \text{(definition of } S_i) \\ &\leq \sum_{j=i+1}^{\ell} \sum_{a \in S_j} \tilde{w}(a) = \sum_{j=i+1}^{\ell} \langle \tilde{w}, c(j) \rangle \\ &\leq (1+\varepsilon) \sum_{j=i+1}^{\ell} \langle w, c(j) \rangle \\ &\leq (1+\varepsilon) \sum_{j=i+1}^{\ell} \frac{1}{1-1/\lambda} \cdot \lambda^{\ell+1-j} \\ &\leq \frac{1+\varepsilon}{(1-1/\lambda)^2} \lambda^{\ell-i} = \frac{4^2(1+\varepsilon)}{(3+\varepsilon)^2} \lambda^{\ell-i} \leq \frac{32}{9} \lambda^{\ell-i} \\ &< 4\lambda^{\ell-i}, \end{aligned}$$

a contradiction. Therefore (26) holds for all  $i \in [\ell]$ , so  $|S| \geq \ell$ . Therefore,  $\text{wSPR}(C) \geq \text{CL}(C)$ .  $\square$

## 5.2 Chain Length of CSPs

We define the chain length of a CSP predicate in terms of sets of satisfying assignments. Recall for an instance  $\Psi$  of  $\text{CSP}(R)$  on  $n$  variables,  $\text{sat}(R, \Psi)$  is the set of satisfying assignments to the  $\Psi$ . Lagerkvist and Wahlström [LW20] first defined the concept, although our definition is closer to that of Bessiere, Carbonnel, and Katsirelos [BCK20].

**Definition 5.10** ([LW20, BCK20]). Given  $R \subseteq D^r$  and  $n \in \mathbb{N}$ , we define a *chain* to be a sequence of instances  $\Psi_1, \dots, \Psi_\ell$  of  $\text{CSP}(R)$  on  $n$  variables such that

$$\text{sat}(R, \Psi_1) \subsetneq \text{sat}(R, \Psi_2) \subsetneq \dots \subsetneq \text{sat}(R, \Psi_\ell). \quad (27)$$

The *chain length* of  $\text{CSP}(R)$  on  $n$  variables, denoted by  $\text{CL}(R, n)$ , is the maximum length  $\ell$  of such a chain.

It is not immediately obvious how Definition 5.10 relates to Definition 5.6. We prove this as follows

**Proposition 5.11.** *For all  $R \subseteq D^r$  and  $n \in \mathbb{N}$ , we have that*

$$\text{CL}(\overline{R}, n) - 1 = \max_{\substack{\Psi \text{ instance of CSP}(R) \\ \text{on } n \text{ variables}}} \text{CL}(C_{R,\Psi}), \quad (28)$$

with  $C_{R,\Psi}$  defined as in Section 2.3.

**Remark 5.12.** Note that chain lengths for CSPs and codes are off by one for the same reason that a path graph with  $n$  vertices has only  $n - 1$  edges. In particular, the code's chain length is effectively counting the number of  $\zeta$ 's in (27).

*Proof.* We first prove that the LHS of (28) is at least the RHS. Fix an instance  $\Psi := (X, Y \subseteq X^r)$  with  $|X| = n$  and let  $\ell := \text{CL}(C_{R,\Psi})$ . Thus, there exists  $y : [\ell] \rightarrow Y$  and  $c : [\ell] \rightarrow C_{R,\Psi}$  such that for all  $i \in [\ell]$ ,  $c(i)_{y(i)} = 1$  and for all  $1 \leq i < j \leq \ell$ ,  $c(j)_{y(i)} = 0$ .

Let  $\sigma_1, \dots, \sigma_\ell : X \rightarrow D$  correspond to  $c(1), \dots, c(j)$ . In particular, for all  $i \in [\ell]$ ,  $\sigma_i(y(i)) \in R$  but for all  $1 \leq i < j \leq \ell$ ,  $\sigma_j(y(i)) \notin R$ . For all  $i \in \{0, 1, \dots, \ell\}$  define the instance

$$\Psi_i := (X, \{y(j) : j \in [i]\}).$$

as an instance of  $\text{CSP}(\overline{R})$ . Since adding clauses can only decrease the number of satisfying assignments, we have that

$$\text{sat}(\overline{R}, \Psi_\ell) \subseteq \text{sat}(\overline{R}, \Psi_{\ell-1}) \subseteq \dots \subseteq \text{sat}(\overline{R}, \Psi_1) \subseteq \text{sat}(\overline{R}, \Psi_0) \quad (29)$$

To prove these inclusions are strict, for each  $i \in [\ell]$ , Observe that

$$\sigma_i \in \text{sat}(\overline{R}, \Psi_{i-1}) \setminus \text{sat}(\overline{R}, \Psi_i),$$

as  $\sigma_i(y(1)), \dots, \sigma_{i+1}(y(i-1)) \notin R$  but  $\sigma_{i+1}(y(i)) \in R$ . Therefore, all inclusions in (29) are strict, so  $\text{CL}(\overline{R}, n) \geq \text{CL}(C_{R,\Psi}) + 1$ .

We now prove that the LHS of (28) is at most the RHS. Let  $\ell := \text{CL}(\overline{R}, n) - 1$ . Fix a set  $X$  of size  $n$  and consider instances  $\Psi_0 := (X, Y_0), \dots, \Psi_\ell := (X, Y_\ell)$  of  $\text{CSP}(\overline{R})$  on the variable set  $X$  such that

$$\text{sat}(\overline{R}, \Psi_\ell) \subsetneq \text{sat}(\overline{R}, \Psi_{\ell-1}) \subsetneq \dots \subsetneq \text{sat}(\overline{R}, \Psi_1) \subsetneq \text{sat}(\overline{R}, \Psi_0)$$

For each  $i \in [\ell]$ , pick  $\sigma_i \in \text{sat}(\overline{R}, \Psi_{i-1}) \setminus \text{sat}(\overline{R}, \Psi_i)$ . Since  $\sigma_i \notin \text{sat}(\overline{R}, \Psi_i)$ , there exists  $y(i) \in Y_i$  which  $\sigma_i$  does not satisfy.

Now consider  $Y := \{y(1), \dots, y(\ell)\}$  and let  $\Psi := (X, Y)$ . It suffices to prove that  $\text{CL}(C_{R,\Psi}) \geq \ell$ . For each  $i \in [\ell]$ , there exists  $c(i) \in C_{R,\Psi}$  corresponding to the assignment  $\sigma_i$ . Since  $\sigma_i \in \text{sat}(\overline{R}, \Psi_{i-1}) \setminus \text{sat}(\overline{R}, \Psi_i)$ , we have that  $c(i)_{y(1)} = \dots = c(i)_{y(i-1)} = 0$  but  $c(i)_{y(i)} = 1$ . Therefore,  $(y : [\ell] \rightarrow Y, c : [\ell] \rightarrow C_{R,\Psi})$  is a chain of length  $\ell$  in  $C_{R,\Psi}$ .  $\square$

As an immediate corollary of the proof, we have the following more combinatorial interpretation of chain length. This is somewhat closer to Lagerkvist and Wahlström's [LW20] definition of chain length.

**Corollary 5.13.** *Let  $R \subseteq D^r$  be a relation and  $X$  a variable set of size  $n$ . We have that  $\text{CL}(R, n) \geq \ell + 1$  if and only if there exist assignments  $\sigma_1, \dots, \sigma_\ell : X \rightarrow D$  and clauses  $y_1, \dots, y_\ell \in X^r$  such that for all  $i \in [\ell]$ ,*

$$\sigma_i(y_1), \dots, \sigma_i(y_{i-1}) \in R \text{ and } \sigma_i(y_i) \notin R.$$

We can now show that Theorem 5.4 implies Theorem 5.5.

**Proposition 5.14.** *Theorem 5.4 implies Theorem 5.5.*

*Proof.* Fix  $R \subseteq D^r, n \in \mathbb{N}, \varepsilon > 0$ . Apply Theorem 5.4 to  $C_{R, \Psi}$  for all instances  $\Psi$  of  $\text{CSP}(R)$  on  $n$  variables. Then, take the maximum of these resulting inequalities for all such  $\Psi$ . By applying Proposition 5.3 and Proposition 5.11, we obtain Theorem 5.5.  $\square$

### 5.3 Proof of Theorem 5.4

We now prove Theorem 5.4. We use a bucketing approach by [KPS25a] to split the weights into groups of roughly equal weight (within a factor of  $m^{O(1)}$ ). We then use a standard repetition trick and Theorem 1.1 to sparsify these groups. Finally, we argue that the total size of the sparsifier cannot be significantly more than the length of the maximal chain in our code.

As a warmup, we first prove that non-redundancy is an upper bound if the ratio between the maximum and minimum weight is at most  $m^{O(1)}$ . (this observation is also in [KPS25a] for linear codes).

**Lemma 5.15.** *Let  $C \subseteq \{0, 1\}^m$  and  $w : [m] \rightarrow R_{>0}$  be a positive weight function such that  $\max(w)/\min(w) \leq m^3$ . Then, for all  $\varepsilon \in (0, 1)$ ,  $C$  has a  $(w, \varepsilon)$ -sparsifier of size at most*

$$10^7 \text{NRD}(C) \log^6(4m)/\varepsilon^2.$$

*Proof.* If  $\text{NRD}(C) = 0$ , return the empty sparsifier. If  $\varepsilon \leq 1/m$ , return  $w$  itself as  $m \leq 1/\varepsilon^2$ . Thus, assume  $\text{NRD}(C) \geq 1$  and  $\varepsilon > 1/m$ . For each  $i \in [m]$  define

$$b_i := \left\lfloor \frac{2w(i)}{\varepsilon \min w} \right\rfloor \leq 2m^5.$$

Let  $\tilde{m} = \sum_{i=1}^m b_i \leq 2m^6$ . Let  $f : \{0, 1\}^m \rightarrow \{0, 1\}^{\tilde{m}}$  be such that

$$f(x) := \underbrace{x_1, \dots, x_1}_{b_1}, \underbrace{x_2, \dots, x_2}_{b_2}, \dots, \underbrace{x_m, \dots, x_m}_{b_m}.$$

We also let  $S_i \subseteq [\tilde{m}]$  be the set of  $b_i$  coordinates which are equal to  $x_i$ . Observe that for all

$x \in \{0, 1\}^m$ , we have that

$$\begin{aligned}
\text{Ham}(f(x)) &= \sum_{i=1}^n \left\lfloor \frac{2w(i)}{\varepsilon \min w} \right\rfloor x_i \\
&\in \sum_{i=1}^n \left[ \frac{2w(i)}{\varepsilon \min w} - 1, \frac{2w(i)}{\varepsilon \min w} \right] \cdot x_i \\
&\subseteq \sum_{i=1}^n [1 - \varepsilon/2, 1] \cdot \frac{w(i)x_i}{(\varepsilon/2) \min(w)} && (\text{since } w(i) \geq \min(w)) \\
&\subseteq [1 - \varepsilon/2, 1] \cdot \frac{\langle w, x \rangle}{(\varepsilon/2) \min(w)}.
\end{aligned}$$

Now define  $\tilde{C} := \{f(c) : c \in C\}$ . Note that  $\text{NRD}(\tilde{C}) = \text{NRD}(C)$  as the repetition of coordinates cannot increase non-redundancy. Let  $\tilde{w} : [\tilde{m}] \rightarrow \mathbb{R}_{\geq 0}$  be an  $(\varepsilon/2)$ -sparsifier of  $\tilde{C}$ . By Theorem 1.1, we have that

$$|\text{supp}(\tilde{w})| \leq 800 \text{NRD}(\tilde{C}) \log_2^6(4\tilde{m}) / (\varepsilon/2)^2 \leq 10^7 \text{NRD}(C) \log_2^6(4m) / \varepsilon^2.$$

Now consider the map  $w' : [m] \rightarrow \mathbb{R}_{\geq 0}$  defined by

$$\forall i \in [m], w'(i) := \frac{\varepsilon \min(w)}{2} \sum_{j \in S_i} \tilde{w}(j).$$

Clearly,  $|\text{supp}(w')| \leq |\text{supp}(\tilde{w})|$  so it suffices to prove that  $w'$  is a  $(w, \varepsilon)$ -sparsifier of  $C$ . To see why, for all  $c \in C$ ,

$$\begin{aligned}
\langle w', c \rangle &= \frac{\varepsilon \min(w)}{2} \langle \tilde{w}, f(c) \rangle \\
&\in [1 - \varepsilon/2, 1 + \varepsilon/2] \cdot \frac{\varepsilon \min(w)}{2} \text{Ham}(f(c)) \\
&\subseteq [1 - \varepsilon/2, 1 + \varepsilon/2] \cdot [1 - \varepsilon/2, 1] \cdot \langle w, x \rangle \\
&\subseteq [1 - \varepsilon, 1 + \varepsilon] \cdot \langle w, x \rangle,
\end{aligned}$$

as desired. □

*Proof of Theorem 5.4.* Fix our code  $C \subseteq \{0, 1\}^m$ , our positive weight function  $w : [m] \rightarrow \mathbb{R}_{>0}$ , and  $\varepsilon \in (0, 1)$ . Similar to the proof of Lemma 5.15, we may assume that  $\text{NRD}(C) \geq 1$ ,  $\varepsilon > 6/m$ , and  $m \geq 2$ .

Define a function  $t : [m] \rightarrow \mathbb{Z}$  be

$$t(i) := \left\lfloor \frac{\log w(i)}{3 \log m} \right\rfloor.$$

Let  $T := \{t(i) : i \in \mathbb{Z}\}$  be a finite set. For each  $c \in C$ , define its *type* to be

$$\text{type}(c) := \max_{i \in \text{supp}(c)} t(i) \in T.$$

For each  $t \in \mathbb{Z}$ , define

$$I_t := \{i \in [m] : t(i) = t\}, \text{ and}$$

$$C_t := \{c \in C : \text{type}(c) = t\}.$$

Note that these sets are only nonempty for  $t \in T$ . For each  $t \in T$ , observe that

$$\frac{\max w|_{I_t}}{\min w|_{I_t}} \leq \frac{m^{3t+3}}{m^{3t}} = m^3.$$

For each  $t \in T$ , let  $\tilde{w}_t : [m] \rightarrow \mathbb{R}_{\geq 0}$  be an  $(\varepsilon/2)$ -sparsifier of  $(C_t \cup C_{t+1} \cup \{1_{I_t}\})|_{I_t}$ , where  $1_{I_t} \in \{0, 1\}^m$  is the indicator of the set  $I_t$ . If  $t \notin T$ , define  $\tilde{w}_t : [m] \rightarrow \mathbb{R}_{\geq 0}$  to be the zero map.

For all  $t \in T$ , note that  $(C_t \cup C_{t+1})|_{I_t}$  has at least one nonzero codeword. Therefore,

$$\text{NRD}((C_t \cup C_{t+1} \cup \{1_{I_t}\})|_{I_t}) \leq \text{NRD}((C_t \cup C_{t+1})|_{I_t}) + 1 \leq 2 \text{NRD}((C_t \cup C_{t+1})|_{I_t}).$$

Thus, by Lemma 5.15, we can ensure that

$$|\text{supp}(\tilde{w}_t)| \leq 8 \cdot 10^7 \text{NRD}((C_t \cup C_{t+1})|_{I_t}) \log^6(4m)/\varepsilon^2. \quad (30)$$

Now consider  $\tilde{w} : [m] \rightarrow \mathbb{R}_{\geq 0}$  defined by

$$\forall i \in [m], \tilde{w}(i) := \tilde{w}_{t(i)}(i).$$

We claim that  $\tilde{w}$  is an  $(w, \varepsilon)$ -sparsifier of  $C$ . To verify this, for any  $c \in C$  with  $t := \text{type}(c)$ , we have that

$$\begin{aligned} \langle \tilde{w}, c \rangle &= \langle \tilde{w}_{t-1}, c|_{I_{t-1}} \rangle + \langle \tilde{w}_t, c|_{I_t} \rangle + \sum_{\substack{t' \in T \\ t' \leq t-2}} \langle \tilde{w}_{t'}, c|_{I_{t'}} \rangle \\ &\in [1 - \varepsilon/2, 1 + \varepsilon/2] \cdot \sum_{i \in I_{t-1} \cup I_t} w(i)c_i + \left[ 0, \sum_{\substack{t' \in T \\ t' \leq t-2}} \langle \tilde{w}_{t'}, 1_{I_{t'}} \rangle \right] \\ &= [1 - \varepsilon/2, 1 + \varepsilon/2] \cdot \left[ \langle w, c \rangle - \sum_{i \in \text{supp}(c) \setminus (I_{t-1} \cup I_t)} w(i) \right] \\ &\quad + (1 + \varepsilon/2) \left[ 0, \sum_{\substack{t' \in T \\ t' \leq t-2}} \langle w, 1_{I_{t'}} \rangle \right] \end{aligned}$$

Since  $t = \text{type}(c)$ , we have that  $\langle w, c \rangle \geq m^{3t}$ . Further, observe that

$$\sum_{i \in \text{supp}(c) \setminus (I_{t-1} \cup I_t)} w(i) \leq \sum_{\substack{t' \in T \\ t' \leq t-2}} \sum_{i \in I_{t'}} w(i) \leq m \cdot m^{3t-3} \leq m^{3t} \cdot \frac{\varepsilon}{6} \leq \frac{\varepsilon}{6} \langle w, c \rangle,$$

where we used the fact that  $\varepsilon > 6/m$ . Therefore,

$$\langle \tilde{w}, c \rangle \in ([1 - \varepsilon/2, 1 + \varepsilon/2] \cdot [1 - \varepsilon/6, 1] + (1 + \varepsilon/2) \cdot [0, \varepsilon/6]) \cdot \langle w, c \rangle \subseteq [1 - \varepsilon, 1 + \varepsilon] \langle w, c \rangle,$$

as desired. It thus suffices to bound  $|\text{supp}(\tilde{w})|$ . By (30), we have that

$$|\text{supp}(\tilde{w})| \leq \sum_{t \in T} |\text{supp}(\tilde{w}_t)| \leq 8 \cdot 10^7 \left[ \sum_{t \in T} \text{NRD}((C_t \cup C_{t+1})|I_t) \right] \log^6(4m)/\varepsilon^2. \quad (31)$$

To finish, we prove the following claim.

**Claim 5.16.**  $\sum_{t \in T} \text{NRD}((C_t \cup C_{t+1})|I_t) \leq 2 \text{CL}(C)$ .

*Proof.* For each  $t \in T$ , let  $j_t := \text{NRD}((C_t \cup C_{t+1})|I_t)$  and let  $i_1^{(t)}, \dots, i_{j_t}^{(t)} \in I_t$  be distinct indices such that there exist  $c^{(t)}(1), \dots, c^{(t)}(j_t) \in C_t \cup C_{t+1}$  such that

$$\forall a, b \in [j_t], c^{(t)}(a)_{i_a^{(t)}} = \mathbf{1}[a = b].$$

We can view this non-redundant set as a chain in  $C$ . Call this chain  $\mathcal{C}_t$ . The key observation is that if  $t, t' \in T$  are such that  $t \geq t' + 2$ , then the concatenation of  $\mathcal{C}_t$  and  $\mathcal{C}_{t'}$  is a chain. Since  $\mathcal{C}_t$  and  $\mathcal{C}_{t'}$  are each chains in  $C$ , it suffices to prove for any  $a \in [j_t]$  and  $b \in [j_{t'}]$  that

$$c^{(t')}(b)_{i_a^{(t)}} = 0. \quad (32)$$

To see why, for any  $c \in C_{t'} \cup C_{t'+1}$ , we have that  $\text{type}(c) \leq t' + 1 < t$ , so  $\text{supp}(c) \cap I_t = \emptyset$ . Therefore, (32) holds.

Thus, for any decreasing subsequence  $t_1, \dots, t_\ell$  of  $T$  with consecutive terms having difference at least 2, we have that the concatenation of  $\mathcal{C}_{t_1}, \dots, \mathcal{C}_{t_\ell}$  is a chain in  $C$ , so  $j_{t_1} + \dots + j_{t_\ell} \leq \text{CL}(C)$ . Since  $T$  can be partitioned into two such subsequences (e.g., even/odd), we have that  $j_1 + \dots + j_\ell \leq 2 \text{CL}(C)$ , as desired.  $\square$

Applying Claim 5.16 to (31) proves the theorem.  $\square$

**Remark 5.17.** A mistake in the proof of Theorem 5.4 was caught while preparing the follow-up work [BGP26b]. In particular, the initial (incorrect) version of the proof claimed that  $\langle \tilde{w}, c \rangle = \langle \tilde{w}_{t-1}, c|_{I_{t-1}} \rangle + \langle \tilde{w}_t, c|_{I_t} \rangle$ , which is only approximately true. The precise error terms are now tracked and bounded in the corrected proof.

**Remark 5.18.** Alternatively, we could bound the sum in (31) by  $|T| \text{NRD}(C)$ . Since  $|T| \leq 2 + (\log \frac{\max w}{\min w}) / (3 \log m)$ , we get the following immediate corollary which gives a savings if  $\max w / \min w$  is small. It can also be seen as an interpolation between Theorem 1.1 and Theorem 5.4.

**Corollary 5.19.** *Let  $C \subseteq \{0, 1\}^m$  and let  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$  be a weight function. For all  $\varepsilon \in (0, 1)$ , there exist a  $(w, \varepsilon)$ -sparsifier of  $C$  of size at most*

$$O \left( \log \frac{m \max w}{\min w} \cdot \text{NRD}(C) \log^5 m / \varepsilon^2 \right).$$

**Remark 5.20.** As observed in [BGP26b], one can recursively apply Theorem 5.4 to itself to get a better sparsifier whose size is independent of the length  $m$  of the codewords. We defer to [BGP26b] for details on how this can be done.

## 6 Immediate Applications of Non-redundancy and Chain Length to Sparsification

In this section, we prove many corollaries of Theorem 1.1, Theorem 1.2, and Theorem 1.3 based on established facts in the non-redundancy literature.

### 6.1 Code Sparsification for Non-Abelian Groups

Given a finite group  $G$  with identity element  $e$ , we define a  $G$ -code to be a subgroup  $H \leq G^m$  for some  $m \in \mathbb{N}$ . Define the Hamming weight of  $h \in H$  to be the number of coordinates  $i \in [m]$  not equal to  $e$ . A  $G$ -code  $\varepsilon$ -sparsifier is a map  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$  such that for every  $h \in H$ , the  $w$ -weight of  $h$  (i.e., the sum of the weights of the non-identity coordinates) is within a factor of  $1 \pm \varepsilon$  of its Hamming weight.

The works of Khanna, Putterman, and Sudan [KPS24a, KPS25a] proved the existence (and in fact efficient construction) of  $G$ -code  $\varepsilon$ -sparsifiers of size  $\log |H| \log^{O(1)} m / \varepsilon^2$  for any Abelian group  $G$ . In particular  $G = \mathbb{Z}/2\mathbb{Z}$  captures cut sparsification. We show as an immediate corollary of Theorem 1.1 that comparably-sized sparsifiers exist for any *non-Abelian* group  $G$ . The proof adapts techniques from [LW20].

**Theorem 6.1.** *For any finite group  $G$ , any  $H \leq G^m$ , and any  $\varepsilon \in (0, 1)$ , there is an  $\varepsilon$ -sparsifier of  $H$  of size  $O(\log |H|(\log^6 m)/\varepsilon^2)$ .*

*Proof.* If  $|G| = 1$ , there is nothing to prove, so assume  $|G| \geq 2$ . Let  $C \subseteq \{0, 1\}^m$  be the code such that for every  $h \in H$ , the binary string  $(\mathbf{1}[h_i \neq e] : i \in [m])$  is an element of  $C$ . Observe that any  $\varepsilon$ -sparsifier of  $C$  is an  $\varepsilon$ -sparsifier of  $H$  (and vice-versa). Therefore, by Theorem 1.1, it suffices to prove that  $\text{NRD}(C) \leq \log_2 |H|$ .

Let  $I := \{i_1, \dots, i_\ell\} \subseteq [m]$  be non-redundant. Thus, there exists  $h_1, \dots, h_\ell$  such that  $h_{j, i_{j'}} \neq e$  if and only if  $j = j'$ . Now consider the set

$$S := \left\{ h_1^{b_1} \cdots h_\ell^{b_\ell} \in H : b \in \{0, 1\}^\ell \right\}.$$

Note that for each  $b \in \{0, 1\}^\ell$ , the coordinates within  $I$  of  $h_1^{b_1} \cdots h_\ell^{b_\ell}$  which do not equal  $e$  are precisely the support of  $b$ . Therefore,  $|S| = 2^\ell$ . Thus,  $\ell \leq \log_2 |H|$ , so  $\text{NRD}(C) \leq \log_2 |H|$ , as desired.  $\square$

For non-Abelian groups  $G$ , Section 7 of [KPS25a] observed that for any two non-commuting elements  $a, b \in G$  that the predicate  $P := \{x \in \{0, 1\}^4 : a^{x_1} b^{x_2} a^{-x_3} b^{-x_4} = e\}$  has  $\text{NRD}(P, n) \geq \Omega(n^2)$ , suggesting that sparsification for (Abelian) affine CSPs does not extend to non-Abelian groups. However, this is not exactly the case—there are predicates corresponding to non-Abelian code sparsification, falling under the more general framework of “Mal’tsev predicates.” See Section 6.3 for more details.

### 6.2 Maximum-size Dichotomy

Observe that for any  $R \subseteq D^r$ , we have that  $\text{NRD}(R, n), \text{SPR}(R, n, \varepsilon) \leq n^r$  as there are at most  $n^r$  possible clauses on  $n$  variables. A natural question is then which predicates have non-redundancy/sparsification of size  $\Omega(n^r)$ . When  $D = \{0, 1\}$ , it is known that  $\text{NRD}(R, n) = \Omega(n^r)$

if and only if  $R = \{0, 1\}^r \setminus \{b\}$  for some  $b \in \{0, 1\}^r$ , otherwise  $\text{NRD}(R, n) = O(n^{r-1})$  [CJP20]. Using linear code sparsification, the analogous result is also known for sparsification [KPS25a]. In particular,  $\text{SPR}(R, n, \varepsilon) = \Omega(n^r)$  if and only if  $R = \{b\}$  for some  $b \in \{0, 1\}^r$ , otherwise  $\text{SPR}(R, n, \varepsilon) = \tilde{O}_\varepsilon(n^{r-1})$ .

For non-Boolean  $D$ , an analogous result is known for non-redundancy due to Carbonnel [Car22]. The main proof technique is applying a hypergraph Turán result due to Erdős [Erd64].

**Theorem 6.2** (Carbonnel [Car22]). *For  $r \in \mathbb{N}$ , let  $\delta_r := 2^{1-r}$ . For all  $R \subseteq D^r$ , we have that  $\text{NRD}(R, n) = \Omega(n^r)$  if and only if there exists subsets  $D_1, \dots, D_r \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times \dots \times D_r)| = 2^r - 1$ . Otherwise,  $\text{NRD}(R, n) = O(n^{r-\delta_r})$ .*

As an immediate corollary of our main result Theorem 1.1 tying sparsification to non-redundancy, we can now extend this dichotomy to sparsification.

**Theorem 6.3.** *For  $r \in \mathbb{N}$ , let  $\delta_r := 2^{1-r}$ . For all  $R \subseteq D^r$  and  $\varepsilon \in (0, 1)$ , we have that  $\text{SPR}(R, n, \varepsilon) = \Omega(n^r)$  if and only if there exists subsets  $D_1, \dots, D_r \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times \dots \times D_r)| = 1$ . Otherwise,  $\text{SPR}(R, n, \varepsilon) = O(n^{r-\delta_r} \log^6 n / \varepsilon^2)$ .*

*Proof.* Apply Theorem 1.2 to Theorem 6.2. □

### 6.3 Generalized Group Equations: Mal'tsev Sparsification

For over a half-century (e.g., [Ber11]), the field of universal algebra has sought to generalize various algebraic structures, such as Abelian and non-Abelian groups. One particular notion that has been vital to the study of CSPs is that of *Mal'tsev operators* [Mal54]. Succinctly, for a domain  $D$ , a Mal'tsev operator is a ternary function  $\varphi : D^3 \rightarrow D$  such that for any  $x, y \in D$ , we have that  $\varphi(x, y, y) = \varphi(y, y, x) = x$ . As a concrete example, for any group  $G$ , consider  $\varphi_G(x, y, z) := x \cdot y^{-1} \cdot z$ . It is clear that  $\varphi_G(x, y, y) = \varphi_G(y, y, x) = x$ , so  $\varphi_G$  is a Mal'tsev operator.

Every finite domain  $D$  and Mal'tsev operator  $\varphi$  has a corresponding family of CSP predicates (e.g., [Ber11, BKW17]). In particular, we say that  $R \subseteq D^r$  is a *Mal'tsev predicate* if for any three tuples  $a, b, c \in R$ , we have that<sup>21</sup>

$$(\varphi(a_1, b_1, c_1), \dots, \varphi(a_r, b_r, c_r)) \in R.$$

In particular, for an Abelian group  $G$ , the Mal'tsev predicates corresponding to  $(G, \varphi_G)$  are precisely the affine predicates, i.e., cosets of  $G^r$ . For non-Abelian groups and more general Mal'tsev operators, the corresponding predicates are much more subtle (see e.g., [LW20]). However, enough is known about Mal'tsev predicates to prove that their non-redundancy is quite small. In fact, [LW20, BCK20] prove that the non-redundancy of every Mal'tsev predicate is linear by utilizing a characterization of Mal'tsev predicates by [BD06].

**Theorem 6.4** ([LW20, BCK20], simplified). *For every Mal'tsev predicate  $R \subseteq D^r$ , we have that  $\text{NRD}(R, n) = O_D(n)$ .*

It turns out that Theorem 6.4 can be strengthened by considering gadget reductions which preserve non-redundancy up to a constant factor. See Section 6.4 and Section 9 for more details. As an immediate corollary of Theorem 1.2, we get a nearly identical result for sparsification.

<sup>20</sup>Spelling variants include Maltsev, Malcev, and Mal'cev.

<sup>21</sup>In CSP parlance, we say that  $\varphi$  is a *Mal'tsev polymorphism* of  $R$ .

**Theorem 6.5** (“Mal’tsev sparsification”). *For every Mal’tsev predicate  $R \subseteq D^r$ , we have that  $\text{SPR}(\overline{R}, n, \varepsilon) = O_D(n \log^6 n / \varepsilon^2)$ .*

## 6.4 Gadget Reductions

An excellent property of non-redundancy is that there is a rich gadget reduction framework for describing the relationships between the non-redundancy of predicates.<sup>22</sup> In particular, the framework allows for three types of operations.

**Theorem 6.6** (NRD gadget reductions, adapted from Proposition 2 [Car22]). *Let  $P, Q \subseteq D^r$  be predicates. The following gadget reductions are (approximately) monotonic with respect to non-redundancy.*

1. **Projection (or minor).** *For any  $s \in \mathbb{N}$  and  $\pi : [s] \rightarrow [r]$ , let  $R \subseteq D^s$  be defined by  $R := \{(a_{\pi(1)}, \dots, a_{\pi(s)}) \in D^s : a \in P\}$ . Then,  $\text{NRD}(R, n) = O(\text{NRD}(P, n))$ .*
2. **Domain restriction.** *For any maps  $f_1, \dots, f_r : E \rightarrow D$ , consider the predicate  $R := \{a \in E^r : (f_1(a_1), \dots, f_r(a_r)) \in P\}$ . Then,  $\text{NRD}(R, n) = O(\text{NRD}(P, n))$ .*
3. **Conjunction.** *Let  $R := P \wedge Q$ . Then,  $\text{NRD}(R, n) = O(\text{NRD}(P, n) + \text{NRD}(Q, n))$ .*

As an immediate consequence of Theorem 1.2, we obtain a comparable gadget framework for sparsification, although somewhat less transparent due to the appearance of complemented relations.

**Theorem 6.7.** *Let  $P, Q \subseteq D^r$  be predicates and  $\varepsilon > 0$ . The following gadget reductions are (approximately) monotonic with respect to sparsification.*

1. **Projection (or minor).** *For any  $s \in \mathbb{N}$  and  $\pi : [s] \rightarrow [r]$ , let  $R \subseteq D^s$  be defined by  $R := \{(a_{\pi(1)}, \dots, a_{\pi(s)}) \in D^s : a \in P\}$ . Then,  $\text{SPR}(\overline{R}, n, \varepsilon) = \tilde{O}_\varepsilon(\text{SPR}(\overline{P}, n, \varepsilon))$ .*
2. **Domain restriction.** *For any maps  $f_1, \dots, f_r : E \rightarrow D$ , consider the predicate  $R := \{a \in E^r : (f_1(a_1), \dots, f_r(a_r)) \in P\}$ . Then,  $\text{SPR}(\overline{R}, n) = \tilde{O}_\varepsilon(\text{SPR}(\overline{P}, n, \varepsilon))$ .*
3. **Disjunction.** *Let  $R := P \vee Q$ . Then,  $\text{SPR}(R, n, \varepsilon) = \tilde{O}_\varepsilon(\text{SPR}(P, n, \varepsilon) + \text{SPR}(Q, n, \varepsilon))$ .*

Of note, the gadget operations are sufficiently rich that the relative non-redundancy/sparsification of a predicate can be described in terms of a set of (universal) algebraic operations it is closed under called *pattern partial polymorphisms*. See [JLR17, LW18, LW20, Car22] for more details.

## 6.5 Applications of Chain Length to Weighted Code Sparsifiers

We now discuss how the previously-mentioned results for unweighted sparsification can be extended to weighted sparsification.

---

<sup>22</sup>Formally these reductions are called *functionally guarded primitive positive (fgpp) definitions* [LW17, Car22].

### 6.5.1 Weighted Code Sparsification for Non-Abelian Groups

Recall for any finite group  $G$  and any subgroup  $H \leq G^m$  that an  $\varepsilon$ -sparsifier of  $H$  is a (sparse) reweighting of  $[m]$  which preserves the Hamming weight over every  $h \in H$  up to a factor of  $1 \pm \varepsilon$ . Likewise, for any weight function  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$ , we can define an  $(w, \varepsilon)$ -sparsifier analogous to (25). Using Theorem 5.4, we can prove an analogue of Theorem 6.1 for weighted sparsification.

**Theorem 6.8.** *For any finite group  $G$ , any  $H \leq G^m$ , any weight function  $w : [m] \rightarrow \mathbb{R}_{\geq 0}$ , and any  $\varepsilon \in (0, 1)$ , there is an  $(w, \varepsilon)$ -sparsifier of  $H$  of size  $O(\log |H|(\log^6 m)/\varepsilon^2)$ .*

*Proof.* If  $|G| = 1$ , there is nothing to prove, so assume  $|G| \geq 2$ . Let  $C \subseteq \{0, 1\}^m$  be the code such that for every  $h \in H$ , the binary string  $(\mathbf{1}[h_i \neq e] : i \in [m])$  is an element of  $C$ . Observe that any  $(w, \varepsilon)$ -sparsifier of  $C$  is a  $(w, \varepsilon)$ -sparsifier of  $H$  (and vice-versa). Therefore, by Theorem 5.4, it suffices to prove that  $\text{CL}(C) \leq \log_2 |H|$ .

Consider a sequence  $i_1, \dots, i_\ell \subseteq [m]$  along with  $h_1, \dots, h_\ell \in H$  such that for all  $1 \leq j \leq j' \leq \ell$ , we have that  $h_{j, i_{j'}} \neq e$  if and only if  $j = j'$ . Now consider the set

$$S := \left\{ h_1^{b_1} \cdots h_\ell^{b_\ell} \in H : b \in \{0, 1\}^\ell \right\}.$$

We claim that all the elements defining  $S$  are distinct. To see why, assume there exist distinct  $b, b' \in \{0, 1\}^\ell$  with  $h_1^{b_1} \cdots h_\ell^{b_\ell} = h_1^{b'_1} \cdots h_\ell^{b'_\ell}$ . Let  $j \in [\ell]$  be the first coordinate for which  $b_j \neq b'_j$ . Then, we have that  $h_j^{b_j} \cdots h_\ell^{b_\ell} = h_j^{b'_j} \cdots h_\ell^{b'_\ell}$ . Note that the  $i_j$ th coordinate of the LHS is equal to  $h_{j, i_j}^{b_j}$ . Likewise, the  $i_j$ th coordinate of the RHS is equal to  $h_{j, i_j}^{b'_j}$ . Since  $h_{j, i_j} \neq e$  and  $b_j \neq b'_j$ , we have a contradiction. Thus, the elements defining  $S$  are distinct, so  $|S| = 2^\ell$ . Thus,  $\ell \leq \log_2 |H|$ , so  $\text{CL}(C) \leq \log_2 |H|$ , as desired.  $\square$

### 6.5.2 Maximum-size Dichotomy

Although Theorem 6.2 due to Carbonnel [Car22] is only stated in terms of non-redundancy, it turns out that it also applies to chain length. We give a proof for completeness.

**Theorem 6.9.** *For  $r \in \mathbb{N}$ , let  $\delta_r := 2^{1-r}$ . For all  $R \subseteq D^r$ , we have that  $\text{CL}(R, n) = \Omega(n^r)$  if and only if there exists subsets  $D_1, \dots, D_r \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times \cdots \times D_r)| = 2^r - 1$ . Otherwise,  $\text{CL}(R, n) = O(n^{r-\delta_r})$ .*

*Proof.* If there exists subsets  $D_1, \dots, D_r \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times \cdots \times D_r)| = 2^r - 1$ , then by Proposition 5.7 and Theorem 6.2, we have that  $\text{CL}(R, n) \geq \text{NRD}(R, n) = \Omega(n^r)$ , as desired.

Now assume such subsets do not exist. Let  $X$  be a set of variables of size  $n$ , and let  $\ell := \text{CL}(R, n) - 1$ . By Corollary 5.13, there exist assignments  $\sigma_1, \dots, \sigma_\ell : X \rightarrow D$  and clauses  $y_1, \dots, y_\ell \in X^r$  such that for all  $i \in [\ell]$ ,  $\sigma_i(y_1), \dots, \sigma_i(y_{i-1}) \in R$  and  $\sigma_i(y_i) \notin R$ .

Let  $Y := \{y_1, \dots, y_\ell\}$ . We claim there do not exist  $X_1, \dots, X_r \subset X$  of size 2 such that  $X_1 \times \cdots \times X_r \subseteq Y$ . Otherwise, there exists  $y_0 \in X_1 \times \cdots \times X_r$  such that there is an assignment  $\sigma : X \rightarrow D$  for which  $\sigma(y_0) \notin R$  but  $\sigma(y) \in R$  for all  $y \in X_1 \times \cdots \times X_r$  not equal to  $y_0$ . For that to be the case, the sets  $D_1 := \sigma(X_1), \dots, D_r := \sigma(X_r)$  must each have size exactly 2 (since something must differ from  $\sigma(y_{0,i})$  for each  $i \in [r]$ ), and

$$D_1 \times \cdots \times D_r = \{\sigma(y) : y \in X_1 \times \cdots \times X_r\}.$$

Therefore,  $(D_1 \times \cdots \times D_r) \cap R$  has size exactly  $2^r - 1$  (since  $\sigma(y_0)$  is excluded), a contradiction. Hence, the  $r$ -uniform hypergraph on vertex set  $X$  and edge set  $Y$  lacks any complete  $r$ -partite subgraph with each part having size exactly 2. Thus, by a theorem of Erdős [Erd64], we have that  $\ell = O(n^{r-\delta_r})$ , so  $\text{CL}(R, n) = O(n^{r-\delta_r})$ .  $\square$

By Theorem 5.4, this immediately extends to weighted sparsification.

**Theorem 6.10.** *For  $r \in \mathbb{N}$ , let  $\delta_r := 2^{1-r}$ . For all  $R \subseteq D^r$  and  $\varepsilon \in (0, 1)$ , we have that  $\text{wSPR}(R, n, \varepsilon) = \Omega(n^r)$  if and only if there exists subsets  $D_1, \dots, D_r \subseteq D$  of size exactly 2 such that  $|R \cap (D_1 \times \cdots \times D_r)| = 1$ . Otherwise,  $\text{wSPR}(R, n, \varepsilon) = O(n^{r-\delta_r}(r \log n)^6/\varepsilon^2)$ .*

### 6.5.3 Generalized Group Equations: Mal'tsev Sparsification

Both Lagerkvist and Wahlström [LW20] and Bessiere, Carbonnel, and Katsirelos [BCK20] proved a linear upper bound on the chain length of any Mal'tsev predicate.

**Theorem 6.11** ([LW20, BCK20]). *For every Mal'tsev predicate  $R \subseteq D^r$ , we have that  $\text{CL}(R, n) = O_D(n)$ .*

Thus, as an immediate corollary of Theorem 5.4, we have that

**Theorem 6.12.** *For every Mal'tsev predicate  $R \subseteq D^r$ , we have that  $\text{wSPR}(\overline{R}, n, \varepsilon) = O_D(n \log^6 n/\varepsilon^2)$ .*

### 6.5.4 Gadget Reductions

Unlike for non-redundancy, a formal gadget framework is not explicitly stated in the literature for chain length. Bessiere, Carbonnel, and Katsirelos [BCK20] did observe (Observation 2) that a restricted family of gadget reductions called “ $c$ -definitions” is monotone with respect to chain length. In particular, using the notation of Theorem 6.6, they allow arbitrary conjunctions as well as projection reductions from maps  $\pi : [s] \rightarrow [r]$  which are injective.

However, in practice (i.e., in [LW20]), chain length is closely used in association with pattern partial polymorphisms, suggesting that the full suite of gadget reductions in Theorem 6.6 also apply to chain length. We leave filling this gap in the literature as future work.

## 7 New results on Non-redundancy via Matching Vector Families

In this section, we consider the following families of arity-3 predicates. Let  $G$  be a finite Abelian group of order at least 3 with operator  $+$  and identity element 0.

$$\begin{aligned} 3\text{LIN}_G &:= \{(x, y, z) \in G^3 : x + y + z = 0\} \\ 3\text{LIN}_G^* &:= 3\text{LIN}_G \setminus \{(0, 0, 0)\}. \end{aligned}$$

**Remark 7.1.** The inspiration for studying  $3\text{LIN}_G^*$  came from investigating a related predicate  $\text{BCK} := \{111, 222, 012, 120, 201\}$  from [BCK20]. See Section 8 for more details.

Since  $3\text{LIN}_G$  is defined by a linear equation over an Abelian group, it is already known (e.g., [LW20, BCK20]) that  $\text{NRD}(3\text{LIN}_G, n) = \Theta_G(n)$  and  $\text{SPR}(\overline{3\text{LIN}_G}, n) = \tilde{\Theta}_G(n/\varepsilon^2)$  [KPS25a].

However, much less is known about  $3\text{LIN}_G^*$ . In particular, existing methods in the literature can only deduce  $\text{NRD}(3\text{LIN}_G^*, n) \in [\Omega_G(n), O_G(n^2)]$ . In this section, we introduce a new method for bounding the non-redundancy of predicates like  $3\text{LIN}_G^*$  by connecting them to the theory of *matching vector families* [Yek08, Rag07, Gop09, Efr12, DGY11] used primarily in the construction of locally decodable codes (LDCs). In particular, this is the first example in the non-redundancy/sparsification literature of a predicate whose optimal exponent as a function of  $n$  is non-integral.

## 7.1 Conditional Non-redundancy

Recall (Definition 2.2) that for a predicate  $R \subseteq D^r$ , an instance  $\Psi := (X, Y \subseteq X^r)$  is non-redundant if for all  $y \in Y$ , there exists an assignment  $\sigma_y : X \rightarrow D$  such that  $\sigma_y(y) \notin R$  but  $\sigma_y(y') \in R$  for all  $y' \in Y \setminus \{y\}$ . We now modify the definition of non-redundancy slightly to capture conditional information about these assignments  $\sigma_y$ .

**Definition 7.2** (Conditional Non-redundancy). Given  $P \subseteq Q \subseteq D^r$ , we say that an instance  $\Psi := (X, Y \subseteq X^r)$  is a non-redundant instance of  $\text{CSP}(P \mid Q)$  if for all  $y \in Y$ , there exists an assignment  $\sigma_y : X \rightarrow D$  such that  $\sigma_y(y) \in Q \setminus P$  and  $\sigma_y(y') \in P$  for all  $y' \in Y \setminus \{y\}$ . We let  $\text{NRD}(P \mid Q, n)$  be the maximum number of clauses in a non-redundant instance of  $\text{CSP}(P \mid Q)$  with  $n$  variables.

Observe that each  $\sigma_y$  is a satisfying assignment to  $\text{CSP}(Q)$ . Also note that if  $Q = D^r$ , then  $\text{NRD}(P \mid Q, n) = \text{NRD}(P, n)$ . The key motivation for conditional non-redundancy is that it satisfies a form of the triangle inequality.

**Proposition 7.3.** *Let  $P \subseteq Q \subseteq R \subseteq D^r$  be predicates, then we have that*

$$\text{NRD}(P \mid Q, n) \leq \text{NRD}(P \mid R, n) \leq \text{NRD}(P \mid Q, n) + \text{NRD}(Q \mid R, n). \quad (33)$$

*In particular, if  $R = D^r$ , we have that*

$$\text{NRD}(P \mid Q, n) \leq \text{NRD}(P, n) \leq \text{NRD}(P \mid Q, n) + \text{NRD}(Q, n). \quad (34)$$

In particular, if we seek to understand  $\text{NRD}(P, n)$ , it suffices to identify a predicate  $Q \supseteq P$  with  $\text{NRD}(Q, n) \ll \text{NRD}(P, n)$  and to bound  $\text{NRD}(P \mid Q, n)$ . Then, we can use the structure of satisfying assignments to  $\text{CSP}(Q)$  to better constrain the non-redundancy of  $\text{CSP}(P)$ . As we shall see, this is particularly useful when  $Q$  is an affine predicate.

*Proof.* Note that  $Q \setminus P \subseteq R \setminus P$ . Thus, any non-redundant instance of  $\text{CSP}(P \mid Q)$  is a non-redundant instance of  $\text{CSP}(P \mid R)$ . This proves the left inequality of (33).

Let  $(X, Y \subseteq X^r)$  be a non-redundant instance of  $\text{CSP}(P \mid R)$ . Let  $\{\sigma_y : y \in Y\}$  be a set of assignments satisfying the conditions of Definition 7.2. Let  $Y_R \subseteq Y$  be the set of  $y \in Y$  such that there exists an assignment  $\sigma'_y : X \rightarrow D$  for which  $\sigma'_y(y) \in R \setminus Q$  and  $\sigma'_y(y') \in P$  for all  $y' \in Y \setminus \{y\}$ . Since  $P \subseteq Q$ , we have that  $(X, Y_R)$  is a non-redundant instance of  $\text{CSP}(Q \mid R)$ , so  $|Y_R| \leq \text{NRD}(Q \mid R, n)$ .

By definition of  $Y_R$ , for all  $y \in Y \setminus Y_R$ , we have that  $\sigma'_y(y) \in Q \setminus P$ . Therefore,  $(X, Y \setminus Y_R)$  is a non-redundant instance of  $\text{CSP}(P \mid Q)$ . Thus,  $|Y \setminus Y_R| \leq \text{NRD}(P \mid Q, n)$ . Since  $|Y| \leq |Y \setminus Y_R| + |Y_R|$ , we have proved (33).  $\square$

As a consequence, we can prove a simple  $O_G(n^2)$  upper bound on  $\text{NRD}(3\text{LIN}_G^*, n)$ . A similar upper bound for a similar predicate was given by [BCK20] via a somewhat different technique, see Section 8 for more details.

**Proposition 7.4.**  $\text{NRD}(3\text{LIN}_G^*, n) = O_G(n^2)$ .

*Proof.* As previously discussed,  $\text{NRD}(3\text{LIN}_G, n) = O_G(n)$ . Hence, by (34), it suffices to prove that  $\text{NRD}(3\text{LIN}_G^* \mid 3\text{LIN}_G, n) \leq n^2$ . Let  $(X, Y)$  be a non-redundant instance of  $\text{CSP}(3\text{LIN}_G^* \mid 3\text{LIN}_G)$ .

The key observation is that any two non-redundant clauses must share at most one variable in common. In particular, if both  $(x_1, x_2, x_3) \in Y$  and  $(x_1, x_2, x'_3) \in Y$  with  $x_3 \neq x'_3$ , then any satisfying assignment  $\sigma$  to  $(X, Y)$  as in instance of  $\text{CSP}(3\text{LIN}_G)$  must impose that  $\sigma(x_3) = \sigma(x'_3)$ . Therefore, either both clauses satisfy  $3\text{LIN}_G^*$  or neither do. This contradicts that  $(X, Y)$  is a non-redundant instance. Hence for any  $(x_1, x_2, x_3) \in Y$ ,  $(x_1, x_2)$  uniquely determines the clause. Thus,  $|Y| \leq n^2$ , as desired.  $\square$

In Section 7.4, we shall substantially improve on this upper bound for many finite Abelian groups  $G$ .

## 7.2 Connection to Matching Vector Families

Recall from the introduction that a matching vector (MV) family is two lists of vectors of the same length such that the dot product of a pair of vectors for the two lists equals 0 if and only if they have the same index in their respective lists. We now define a generalization of MV families where the homomorphisms are applied to triples of vectors at a time instead a single vector.

**Definition 7.5.** For an Abelian group  $G$ , a  $G$ -ensemble consists of a dimension  $d \in \mathbb{N}$ , a collection of vectors  $V \subseteq G^d$ , a collection of edges  $E \subseteq V^3$ , and a collection of homomorphisms  $\phi_e \in \text{Hom}(G^d, G)$  for all  $e \in E$ . We require that each  $e := (v_1, v_2, v_3) \in E$ , we have that  $v_1 + v_2 + v_3 = 0$ . Finally, for any  $e, e' \in E$ , we have that  $\phi_e$  maps all the vectors of  $e'$  to 0 if and only if  $e = e'$ . We define the *size* of a  $G$ -ensemble to be  $|E|$ , the number of edges.

We now connect the existence of these  $G$ -ensembles to the non-redundancy problem.

**Proposition 7.6.** *The maximum size of a  $G$ -ensemble with  $n$  vectors (in any dimension) equals  $\text{NRD}(3\text{LIN}_G^* \mid 3\text{LIN}_G, n)$ .*

*Proof.* First, we show for any  $G$ -ensemble  $(d, V, E, \{\phi_e : e \in E\})$  that  $(V, E)$  is a non-redundant instance of  $\text{CSP}(3\text{LIN}_G^* \mid 3\text{LIN}_G)$ . To see why, note that for any  $e, e' = (v'_1, v'_2, v'_3) \in E$ , we have that  $\phi_e(e') = (\phi_e(v'_1), \phi_e(v'_2), \phi_e(v'_3))$ . These values must sum to 0 because  $v'_1 + v'_2 + v'_3 = 0$  and  $\phi_e$  is a homomorphism. Further, observe that  $\phi_e(e') = (0, 0, 0) \in 3\text{LIN}_G \setminus 3\text{LIN}_G^*$  if and only if  $e = e'$ . Therefore,  $(V, E)$  is indeed non-redundant.

Second, we show that any non-redundant instance  $(X, Y)$  of  $\text{CSP}(3\text{LIN}_G^* \mid 3\text{LIN}_G)$  can be turned into a  $G$ -ensemble  $(d, V, E, \{\phi_e : e \in E\})$  with  $|V| = |X|$  and  $|E| = |Y|$ . Let  $\{\sigma_y : X \rightarrow G \mid y \in Y\}$  be the assignments guaranteed by Definition 7.2. Let  $H := \text{sat}(3\text{LIN}_G, Y)$ , the set of satisfying assignments to  $(X, Y)$  when viewed as an instance of  $\text{CSP}(3\text{LIN}_G)$ .

We now highlight the key step in the proof. Observe that  $H$  can be viewed as a subgroup of  $G^X$ . Let  $h_1, \dots, h_d \in H$  be any generating set of  $H$ . For each  $x \in X$ , we can construct  $v_x \in G^d$

by letting  $v_{x,i} = h_{i,x}$  for all  $i \in [d]$  (i.e., the transpose). For each  $y := (x_1, x_2, x_3) \in Y$ , we let  $e_y := (v_{x_1}, v_{x_2}, v_{x_3})$  with  $E := \{e_y : y \in Y\}$ . Since for every  $h \in H$ , we have  $h_{x_1} + h_{x_2} + h_{x_3} = 0$ , and in particular this holds for each  $h_i$ , we can conclude that  $v_{x_1} + v_{x_2} + v_{x_3} = 0$ .

Observe that for any  $y \in Y$  each assignment  $\sigma_y$  can be viewed as an element of  $H$ . Thus, there exists  $z_{y,1}, \dots, z_{y,d} \in \mathbb{Z}$  (not necessarily unique) such that  $\sigma_y = z_{y,1}h_1 + \dots + z_{y,d}h_d$ . We can thus define the map  $\phi_{e_y} \in \text{Hom}(G^d, G)$  by

$$\phi_{e_y}(g_1, \dots, g_d) := \sum_{i=1}^d z_{y,i}g_i.$$

Thus, by definition of  $\sigma_y$ , each corresponding  $\phi_{e_y}$  must map  $e_{y'}$  to  $(0, 0, 0)$  if and only if  $y = y'$ , as desired. Note this further implies that each  $e_y$  is distinct, so  $|E| = |Y|$ , as desired.

As a last detail, note that  $|V| \leq |X|$  by definition of  $X$ , but  $|V| < |X|$  if two  $v_x$ 's are identical. We can get around this issue by padding. Increase  $d$  to  $d + |X| - |V|$  and pad each  $v \in V$  by appending  $|X| - |V|$  copies of 0. We extend each homomorphism  $\phi_e$  by ignoring these last  $|X| - |V|$  coordinates. Pad  $V$  with any  $|X| - |V|$  additional vectors. It is straightforward to check we still have a  $G$ -ensemble of the same size, but now with  $|X|$  vectors.  $\square$

We now heavily use Proposition 7.6 to construct novel upper and lower bounds on  $\text{NRD}(3\text{LIN}_G^*, n)$ .

### 7.3 Constructing a Non-redundant Instance

We now show a lower bound of size  $\Omega(n^{1.5})$  for  $\text{NRD}(3\text{LIN}_G^*, n)$  for all Abelian groups  $G$  of order at least 3. The cardinality requirement is necessary as  $3\text{LIN}_{\mathbb{Z}/1\mathbb{Z}}^*$  is the empty relation and  $3\text{LIN}_{\mathbb{Z}/2\mathbb{Z}}^*$  is equivalent to 2-in-3 SAT, which has linear non-redundancy (e.g., [LW20, BCK20, KPS25a]). The key property of  $G$  we need is quite simple:

**Lemma 7.7.** *If  $G$  is an Abelian group with  $|G| \geq 3$ , then there exists  $g_1, g_2, g_3 \in G \setminus \{0\}$  such that  $g_1 + g_2 + g_3 = 0$ .*

*Proof.* If there exists some element  $g \in G \setminus \{0\}$  of order  $t \geq 3$ , then we can take  $g_1 = g_2 = g$  and  $g_3 = (t - 2)g$ . Otherwise, every element of  $g \in G \setminus \{0\}$  has order 2. Since  $|G| \geq 3$ , there exists distinct  $g_1, g_2 \in G \setminus \{0\}$ . Set  $g_3 = g_1 + g_2 \neq 0$ .  $\square$

We can now prove our lower bound on non-redundancy. The construction presented is inspired by a method in the matching vector code literature (cf. Lemma 16 of [DGY11]).

**Theorem 7.8.** *For every Abelian group  $G$  with  $|G| \geq 3$ , we have that*

$$\text{NRD}(3\text{LIN}_G^*, n) \geq \text{NRD}(3\text{LIN}_G^* | 3\text{LIN}_G, n) = \Omega(n^{1.5}).$$

*Proof.* By Proposition 7.3, we have that  $\text{NRD}(3\text{LIN}_G^*, n) \geq \text{NRD}(3\text{LIN}_G^* | 3\text{LIN}_G, n)$ . By Proposition 7.6, to lower bound  $\text{NRD}(3\text{LIN}_G^* | 3\text{LIN}_G, n)$ , it suffices to construct a  $G$ -ensemble on  $n$ -vectors of size  $\Omega(n^{1.5})$ . We describe the construction as follows.

Let  $t = \lfloor \sqrt{n/3} \rfloor$  and let  $d = 3t + 2$ . Let  $g_1, g_2, g_3 \in G \setminus \{0\}$  be chosen as in Lemma 7.7. Let  $I_1 := \{1, \dots, t\}$ ,  $I_2 := \{t + 1, \dots, 2t\}$ , and  $I_3 := \{2t + 1, \dots, 3t\}$ . For all  $i \in [d]$  and  $g \in G$ , let

$b_i(g) \in G^d$  be the vector which has all coordinates equal to zero except for the  $i$ th coordinate which has value  $g$ . Now consider the following sets of vectors

$$\begin{aligned} V_{12} &:= \{u_{i_1, i_2} := b_{i_1}(g_1) + b_{i_2}(g_2) + b_{3t+1}(g_3) : i_1 \in I_1, i_2 \in I_2\}, \\ V_{23} &:= \{v_{i_2, i_3} := -b_{i_2}(g_2) - b_{i_3}(g_3) - b_{3t+2}(g_1) : i_2 \in I_2, i_3 \in I_3\}, \\ V_{13} &:= \{w_{i_1, i_3} := -b_{i_1}(g_1) + b_{i_3}(g_3) - b_{3t+1}(g_3) + b_{3t+2}(g_1) : i_1 \in I_1, i_3 \in I_3\}. \end{aligned}$$

Let  $V = V_{12} \cup V_{23} \cup V_{13}$  which has size  $3t^2 \leq n$ , which can be made of size exactly  $n$  via padding. For each triple  $(i_1, i_2, i_3) \in I_1 \times I_2 \times I_3$  observe that  $u_{i_1, i_2} + v_{i_2, i_3} + w_{i_1, i_3} = 0$ . Therefore, we can define  $E := \{e_{i_1, i_2, i_3} := (u_{i_1, i_2}, v_{i_2, i_3}, w_{i_1, i_3}) \mid (i_1, i_2, i_3) \in I_1 \times I_2 \times I_3\} \subset V^3$ . The size of  $|E|$  is  $t^3 = \Omega(n^{3/2})$ .

To finish, it suffices to construct maps  $\phi_{i_1, i_2, i_3} \in \text{Hom}(G^d, G)$  which map only  $e_{i_1, i_2, i_3}$  to  $(0, 0, 0)$ . To do this, we define

$$\phi_{i_1, i_2, i_3}(z) := z_{i_1} + z_{i_2} + z_{i_3} + z_{3t+1} + z_{3t+2}.$$

Now observe that for any  $(i'_1, i'_2, i'_3) \in I_1 \times I_2 \times I_3$ , we have that

$$\phi_{i_1, i_2, i_3}(u_{i_1, i_2}) = g_1 \mathbf{1}[i'_1 = i_1] + g_2 \mathbf{1}[i'_2 = i_2] + g_3, \quad (35)$$

$$\phi_{i_1, i_2, i_3}(v_{i_2, i_3}) = -g_2 \mathbf{1}[i'_2 = i_2] - g_3 \mathbf{1}[i'_3 = i_3] - g_1. \quad (36)$$

Since each  $g_i$  is nonzero and their sum equals zero, the only way (35) equals 0 is for  $i'_1 = i_1$  and  $i'_2 = i_2$ . Likewise, the only way (36) equals 0 is for  $i'_2 = i_2$  and  $i'_3 = i_3$ . Thus, the only way that (35) and (36) are both zero is for  $(i'_1, i'_2, i'_3) = (i_1, i_2, i_3)$ . Since  $w_{i_1, i_3} = -u_{i_1, i_2} - v_{i_2, i_3}$  and  $\phi_{i_1, i_2, i_3}$  is a homomorphism, we have that  $\phi_{i_1, i_2, i_3}(e_{i'_1, i'_2, i'_3}) = (0, 0, 0)$  if and only if  $(i_1, i_2, i_3) = (i'_1, i'_2, i'_3)$ . Therefore, we have indeed constructed a  $G$ -ensemble of size  $\Omega(n^{1.5})$ .  $\square$

**Remark 7.9.** Note that the only assignments any  $\phi_{i_1, i_2, i_3}$  assigns to a clause are from the list

$$\begin{aligned} R := &\{(0, 0, 0), (0, g_3, -g_3), (-g_1, 0, g_1), (-g_1, g_3, g_1 - g_3), \\ &(-g_2, -g_1, -g_3), (-g_2, g_2, 0), (g_3, -g_1, g_1 - g_3), (g_3, g_2, g_1)\} \end{aligned}$$

In particular, this implies our non-redundancy lower bound also applies to  $R \setminus \{(0, 0, 0)\}$ . We explore this in more detail in Section 8. This observation also suggests that for large groups  $G$ , a construction which uses “more” of  $3\text{LIN}_G^*$  could yield a stronger lower bound.

## 7.4 Upper Bounding Non-redundancy

In this section, we show how the theory of matching vector codes can also lead to upper bounds on  $\text{NRD}(3\text{LIN}_G^*, n)$ . Due to the existence of very large MV families for Abelian groups  $G$  divisible by at least two primes [DGY11], we assume that  $G = \mathbb{Z}/p\mathbb{Z}$ , where  $p$  is a prime.

In particular, we can identify vectors in  $G^d$  with vectors in the vector space  $\mathbb{F}_p^d$ . As a consequence, for our collection of vectors  $V \subseteq G^d$ , we can quantify the “spread” of  $V$  by  $\dim(\text{span}(V))$ .<sup>23</sup>

As a first step, we give a bound on the size of any  $\mathbb{F}_p$ -ensemble as a function of  $\dim(\text{span}(V))$ .

<sup>23</sup>In [KPS25a], the authors use  $\log_2|\text{span}(V)|$  as a surrogate for dimension for arbitrary finite Abelian groups. However, to extend Theorem 7.13 to all finite Abelian groups we would need a suitable analogue of Lemma 7.10 for arbitrary groups, which in many cases is false (see [DGY11]).

**Lemma 7.10.** *Let  $p$  be a prime. There exists a constant  $c_p > 0$  with the following property. Let  $(d, V, E, \{\phi_e : e \in E\})$  be an  $\mathbb{F}_p$ -ensemble. For any subspace  $U \leq \mathbb{F}^d$ , we have that*

$$|E \cap U^3| \leq \binom{\dim(U) + 2p - 2}{2p - 2}.$$

In particular, for  $U = \text{span}(V)$ ,

$$|E| \leq \binom{\dim(\text{span}(V)) + 2p - 2}{2p - 2}. \quad (37)$$

*Proof.* We mimic the proof of Theorem 23 of [DGY11] (themselves inspired by [BF24]).

Let  $m = |E \cap U^3|$ . Let  $e_i = (u_i, v_i, w_i) \in E \cap U^3$  be an enumeration of these edges. For each  $i \in [m]$ , define the degree  $2p - 2$  polynomial  $P_i : U \rightarrow \mathbb{F}$  as follows:

$$P_i(x) := (1 - \langle x, u_i \rangle^{p-1})(1 - \langle x, v_i \rangle^{p-1}).$$

For each  $e \in E \cap U^3$ , there exists a unique  $z_e \in U$  such that  $\phi_e(x) = \langle z_e, x \rangle$  for all  $x \in U$ . In particular, for each  $i, j \in [m]$ , we have that

$$\begin{aligned} P_i(z_{e_j}) &= (1 - \langle z_{e_j}, u_i \rangle^{p-1})(1 - \langle z_{e_j}, v_i \rangle^{p-1}) \\ &= (1 - \phi_{e_j}(u_i)^{p-1})(1 - \phi_{e_j}(v_i)^{p-1}) \\ &= \mathbf{1}[(\phi_{e_j}(u_i), \phi_{e_j}(v_i)) = (0, 0)] \\ &= \mathbf{1}[i = j]. \end{aligned}$$

Thus,  $P_1, \dots, P_m$  are linearly independent in the vector space of degree  $2p - 2$  polynomials on  $\dim(U)$  variables. Therefore, by a standard counting argument,  $m \leq \binom{\dim(U) + 2p - 2}{2p - 2}$ , as desired.  $\square$

The key takeaway of Lemma 7.10, is that if  $\dim(\text{span}(V)) \leq n^{1/(p-1)-\varepsilon}$  for some  $\varepsilon > 0$ , then we have proved a subquadratic upper bound on  $|E|$ . To handle the case in which  $V$  is more “spread out,” we diverge from the standard techniques in the matching vector codes literature and prove a more combinatorial upper bound. In particular, we utilize the fact that in the spread-out case,  $V$  has large subsets of linearly independent vectors.

**Lemma 7.11.** *Let  $\mathbb{F}$  be a field and let  $V \subseteq \mathbb{F}^d$  be a collection of vectors. Let  $I \subseteq V$  be a linearly independent subset of vectors. Then, there are at most  $6|V| \log_2 |V|$  triples of vectors  $v_1, v_2, v_3 \in V$  summing to 0 with at least one vector in  $I$ .*

**Remark 7.12.** This bound is tight up to constant factors for the Hadamard code. In particular, let  $V = \mathbb{F}^s$  and let  $I$  be the standard basis. There are  $\Theta(|I||V|) = \Theta(|V| \log |V|)$  triples of vectors summing to 0 with at least one vector in  $I$ . However, due to the limitations imposed by Lemma 7.10, this construction does not correspond to an  $\mathbb{F}$ -ensemble.

*Proof.* Let  $e_1, \dots, e_t$  be an enumeration of the elements of  $I$ . For  $i \in [t]$ , let  $E_i \subseteq V^2$  be the pairs of vectors which sum to  $-e_i$ . Observe that each  $E_i$  is a partial matching and  $E_i \cap E_j = \emptyset$  for all  $i \neq j$ . Let  $E_{\leq t} = E_1 \cup \dots \cup E_t$ . Note that  $3|E_{\leq t}| = 3 \sum_{i=1}^t |E_i|$  is an upper bound on the number of edges from  $E$  incident with exactly one edge of  $I$ . Thus, it suffices to prove that  $|E_{\leq t}| \leq 2|V| \log_2 |V|$ . The key observation is as follows.

**Observation.** View  $E_{\leq i}$  as the edges of a bipartite graph on  $V \sqcup V$ . For every  $j > i$ , and for every connected component  $C$  of  $E_{\leq i}$ , no edge of  $E_j$  has both ends in  $C$ .

To see why this is true, consider any potential edge  $(v_1, v_2) \notin E_{\leq i}$  lying in  $C$ . Since  $E_{\leq t}$  is bipartite, there exists a path  $v_1 = u_1, u_2, \dots, u_{2k} = v_2$  with  $(u_a, u_{a+1}) \in E_{\leq i}$  for all  $a \in [2k-1]$ , where we use the fact that  $(p, q) \in E_{\leq i}$  if and only if  $(q, p) \in E_{\leq i}$ . Now observe that

$$v_2 + v_1 = (u_1 + u_2) - (u_2 + u_3) + (u_3 + u_4) - \dots + (u_{2k-1} + u_{2k}) \in \text{span}\{e_1, \dots, e_i\},$$

where the inclusion is by the definition of  $E_{\leq i}$ . Since  $e_1, \dots, e_t$  are linearly independent, we must have that  $(v_1, v_2) \notin E_j$  for all  $j > i$ , as desired.

**Structural Induction.** We now prove by induction the following claim: for every  $i \in [t]$ , and every connected component  $C$  of  $E_{\leq i}$  (viewed as a bipartite graph on  $V \sqcup V$ ), we have that the subgraph of  $E_{\leq i}$  induced by  $C$ ,  $E_{\leq i} \upharpoonright_C$ , has at most  $|C| \log_2 |C|$  edges.

For the base case of  $i = 1$ , note that every connected component has either exactly one vertex and zero edges or two vertices and one edge. In both cases, the bound of  $|C| \log_2 |C|$  holds.

Now assume the induction hypothesis holds for  $E_{\leq i}$ . Let  $C$  be a new (compared to  $E_{\leq i}$ ) connected component of  $E_{\leq i+1}$ . Note that  $C$  must be the union of components  $C'_1, \dots, C'_a$  of  $E_{\leq i}$ . Further, due to the observation, the edges of  $E_{i+1}$  must be between different components. Assume WLOG that  $|C'_1| \geq \dots \geq |C'_a|$ . By the observation, at least one vertex of every edge of  $E_{i+1} \upharpoonright_C$  must lie in  $C'_2 \cup \dots \cup C'_a$ . Thus, we have a bound

$$\begin{aligned} |E_{\leq i+1} \upharpoonright_C| &\leq \sum_{j=2}^a |C'_j| + \sum_{j=1}^a |E_{\leq i} \upharpoonright_{C'_j}| \\ &\leq \sum_{j=2}^a |C'_j| + \sum_{j=1}^a |C'_j| \log_2 |C'_j| && \text{(induction hypothesis)} \\ &= |C'_1| \log_2 |C'_1| + \sum_{j=2}^a |C'_j| \log_2 (2|C'_j|) \\ &\leq |C'_1| \log_2 |C| + \sum_{j=2}^a |C'_j| \log_2 |C| && \text{(since } |C| \geq |C_1| + |C_j| \geq 2|C_j|) \\ &= |C| \log_2 |C|, \end{aligned}$$

as desired.

To finish, each component  $C$  of  $E_{\leq t}$  has at most  $|C| \log_2 |C| \leq |C| \log_2 |V|$  edges. Since the sum of the sizes of the components is  $2|V|$ , we have that  $|E_{\leq t}| \leq 2|V| \log_2 |V|$ , as desired.  $\square$

We now conclude with the main result of this section.

**Theorem 7.13.** *For any prime  $p$ , we have  $\text{NRD}(3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}^*, n) = O_p(n^{2-\varepsilon_p} \log n)$ , for  $\varepsilon_p = \frac{2}{2p-1}$ .*

*Proof.* By Proposition 7.3 and the fact that  $\text{NRD}(3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}, n) = O(n)$ , it suffices to prove that  $\text{NRD}(3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}^* | 3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}, n) = O_p(n^{2-\varepsilon_p} \log^{1-\varepsilon_p} n)$ . By Proposition 7.6, it suffices to prove this upper bound on the size of any  $\mathbb{F}_p$ -ensemble with  $n$  vectors. To do this, we shall show that every  $\mathbb{F}_p$ -ensemble on  $n'$  vectors for  $n' \leq n$  has a vector involved in at most  $\Delta := O_p((n \log n)^{1-\varepsilon_p})$  edges;

i.e., the hypergraph is  $\Delta$ -degenerate. In particular, this implies that any  $\mathbb{F}_p$ -ensemble on  $n$  vectors has at most  $n\Delta = O_p(n(n \log n)^{1-\varepsilon_p})$  edges, as desired.

To do this, we perform casework on the size of  $\dim(\text{span}(V))$ .

**Case 1**,  $\dim(\text{span}(V)) \leq n^{\varepsilon_p}$ . By Lemma 7.10, we have our graph has at most

$$\binom{\dim(\text{span}(V)) + 2p - 2}{2p - 2} \leq \binom{n^{\varepsilon_p} + 2p - 2}{2p - 2} = O_p(n^{2-\varepsilon_p})$$

edges. Thus, some vertex has degree at most  $O_p(n^{1-\varepsilon_p}) \leq \Delta$ .

**Case 2**,  $\dim(\text{span}(V)) > n^{\varepsilon_p}$ . Let  $I \subseteq V$  be a basis of  $\text{span}(V)$ . By Lemma 7.11 there are at most  $O(n \log n)$  edges incident with  $I$ . Thus, some vertex of  $I$  has degree at most  $O((n/|I|) \log n) = O_p(n^{1-\varepsilon_p} \log n) \leq \Delta$ , as desired.  $\square$

**Remark 7.14** (Efficient Kernelization). We observe that this upper bound proof can be turned into an efficient kernelization algorithm (see Section 1.6). More precisely, given any instance  $(X, Y)$  of  $\text{CSP}(3\text{LIN}_{\mathbb{Z}/p\mathbb{Z}}^*)$  on  $n$  variables, we can efficiently identify  $Y' \subseteq Y$  of size at most  $\tilde{O}(n^{2-2/(2p-1)})$  such that  $Y$  and  $Y'$  have the same solution set.

**Remark 7.15**. As an immediate corollary of Theorem 7.8 and Theorem 7.13, we have identified the first predicate  $R$  for which there exist  $a \in \mathbb{N}$  for which

$$a < \liminf_{n \rightarrow \infty} \frac{\log \text{NRD}(R, n)}{\log n} \leq \limsup_{n \rightarrow \infty} \frac{\log \text{NRD}(R, n)}{\log n} < a + 1.$$

In particular, we now know that

$$\begin{aligned} \text{NRD}(3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*, n) &\in [\Omega(n^{1.5}), \tilde{O}(n^{1.6})], \text{ and} \\ \text{SPR}(\overline{3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*}, n, \varepsilon) &\in [\Omega(n^{1.5}), \tilde{O}(n^{1.6}/\varepsilon^2)]. \end{aligned}$$

## 8 Predicates with Unresolved Non-Redundancy

In this section, we give an overview of various predicates discussed in the literature whose non-redundancy is unknown. Essentially all of these predicates are some modification of a system of linear equations, perhaps indicating some modification of the techniques in Section 7 could be useful in their resolution.

### 8.1 The Simplest Unresolved Predicate

A central open question is many works on non-redundancy and related questions (e.g., [LW20, BCK20, Car22]) is characterizing which predicates  $R \subseteq D^r$  have linear/near-linear non-redundancy (or related quantity). The best characterization to date ([BCK20, Car22]) are the Mal'tsev predicates (Theorem 6.4) and derived predicates via gadget reductions (see Section 6.4). This characterization is tight for binary relations ( $r = 2$ ) [BCK20] as well as for Boolean predicates with  $r \leq 3$  [CJP20, KPS25a].

However, multiple works [LW20, BCK20, KPS25a] have identified predicates which fall outside of this characterization for which no nontrivial lower bound is known. Surprisingly, assuming Mal'tsev

predicates do characterize near-linear redundancy (and thus up to complement also sparsification), it suffices to resolve all of these examples in the literature by looking at a single ternary predicate (i.e.,  $r = 3$ ) over a domain  $D = \{0, 1, 2\}$  of size 3:

$$\text{BCK} := \{111, 222, 012, 120, 201\},$$

where tuples are written as concatenated integers for succinctness. In particular, we know<sup>24</sup> from Example 3 and Observation 12 of [BCK20] that one cannot establish that BCK has near-linear non-redundancy via Theorem 6.4. In particular, Bessiere, Carbonnel, and Katsirelos [BCK20], adapting techniques of [LW20], show that every predicate  $P \supseteq \text{BCK}$  definable from Mal'tsev predicates via gadget reductions must also contain the tuple 000. That is, the relation

$$\text{BCK}^+ := \{000, 111, 222, 012, 120, 201\}$$

is in a strong sense the ‘‘Mal'tsev closure’’ of BCK. It is easy to verify that  $\text{BCK}^+$  is Mal'tsev as it turns out that

$$\widetilde{\text{BCK}^+} := \{0, 1, 3\}^3 \cap \{(x, y, z) \in \mathbb{F}_7^3 : x + 2y + 4z = 0\} = \{000, 111, 333, 013, 130, 301\}, \quad (38)$$

which is identical to  $\text{BCK}^+$  up to swapping 2 and 3.

**Booleanization.** Although BCK lives over domain size 3, its non-redundancy also has nontrivial implications in the Boolean setting. In particular, consider the following ‘‘Booleanization’’ of BCK:

$$\begin{aligned} \text{BCK}_{\mathbb{B}} &:= \{(\mathbf{1}[x = 0], \mathbf{1}[x = 1], \mathbf{1}[x = 2], \\ &\quad \mathbf{1}[y = 2], \mathbf{1}[y = 0], \mathbf{1}[y = 1], \\ &\quad \mathbf{1}[z = 1], \mathbf{1}[z = 2], \mathbf{1}[z = 0]) : (x, y, z) \in R\} \subset \{0, 1\}^9, \\ &= \{010001100, 001100010, 100001010, 010100001, 001010100\} \end{aligned}$$

and one can define  $\text{BCK}_{\mathbb{B}}^+$  analogously by adding 100010001. If we view  $\{0, 1\}^9$  as  $3 \times 3$  Boolean matrices, then  $\text{BCK}_{\mathbb{B}}^+$  is precisely the six permutation matrices, and  $\text{BCK}_{\mathbb{B}}$  is  $\text{BCK}_{\mathbb{B}}^+$  with the identity matrix removed.

It is straightforward to observe that  $\text{NRD}(\text{BCK}_{\mathbb{B}}, 3n) \geq \text{NRD}(\text{BCK}, n)$ . For any non-redundant instance  $(X, Y \subset X^3)$  of  $\text{CSP}(\text{BCK})$ , let  $X' = X \times \{0, 1, 2\}$  and map any  $(x_1, x_2, x_3) \in Y$  to

$$((x_1, 0), (x_1, 1), (x_1, 2), (x_2, 2), (x_2, 0), (x_2, 1), (x_3, 1), (x_3, 2), (x_3, 0)) \in Y' \subset (X')^3.$$

Then,  $(X', Y')$  must be a non-redundant instance of  $\text{BCK}_{\mathbb{B}}$ . Thus, proving that  $\text{NRD}(\text{BCK}, n) = \omega(n)$ , would also show that  $\text{NRD}(\text{BCK}_{\mathbb{B}}, n) = \omega(n)$ .

The reason this observation is significant is that  $\text{BCK}_{\mathbb{B}}$  has appeared three times in the literature. In [KPS25a], the predicate

$$P_{\text{KPS}} := \{000000000, 111111000, 111000111, 110001001, 101010010\}$$

<sup>24</sup>This predicate appears in Example 3 of [BCK20] as  $R := \{001, 020, 122, 202, 210\}$ , but one can show  $\text{NRD}(\text{BCK}, n) = \Theta(\text{NRD}(R, n))$  by applying to each coordinate a suitable bijection from  $D$  to itself (see Section 6.4). Explicitly, let  $g_1, g_3$  be permutations which swap 0 and 1 and let  $g_2$  be the cycle  $0 \rightarrow 1 \rightarrow 2 \rightarrow 0$ . Then,  $\text{BCK} = \{(x, y, z) : (g_1(x), g_2(y), g_3(z)) \in R\}$ .

is constructed as an example for which  $\text{SPR}(\overline{P}_{\text{KPS}}, n, \varepsilon)$  cannot be near-linear due to their linear code sparsification framework. In particular, any affine predicate containing  $P_{\text{KPS}}$  must also contain 011100100. Up to a permutation of the coordinates and flipping bits,<sup>25</sup> it turns out that  $P_{\text{KPS}}$  and  $\text{BCK}_{\mathbb{B}}$  are the same, with the same automorphism sending 011100100 to the identity matrix in  $\text{BCK}_{\mathbb{B}}^+$ .

Even earlier, [CJP20] considered the predicate

$$P_{\text{CJP}} := \{100111001, 000010011, 010110110, 000101100, 001001111\}.$$

Here, any affine predicate containing  $P_{\text{CJP}}$  also has 11111111. If we negate the last 6 bits and permute the coordinates suitably, we again arrive at  $\text{BCK}_{\mathbb{B}}$ .

As one last example, [LW20] considers the following predicate of arity  $r = 10$ ,

$$P_{\text{LW}} = \{0000000001, 1000100010, 0100011000, 0011000100, 1000010100, 0010101000\}$$

If one applies the gadget reduction where the last coordinate of  $P_{\text{LW}}$  is set to 0, then this reduced version of  $P_{\text{LW}}$  is equal to  $\text{BCK}_{\mathbb{B}}$  up to a permutation of the coordinates.<sup>26</sup> Thus,  $\text{NRD}(P_{\text{LW}}, n) \geq \text{NRD}(\text{BCK}_{\mathbb{B}}, n) \geq \text{NRD}(\text{BCK}, n)$ . Thus, again, it suffices<sup>27</sup> to prove that  $\text{BCK}$  has super-linear non-redundancy to prove that  $P_{\text{LW}}$  has super-linear non-redundancy.

**Remark 8.1.** The reason why  $P_{\text{LW}}$  has an extra tuple compared to  $\text{BCK}_{\mathbb{B}}$  is due to how [LW20] constructed the predicate. Like [BCK20], they also showed that  $P_{\text{LW}}$  cannot be a gadget construction from Mal'tsev predicates (in fact, [LW17] appears to be the first work to develop this technique). However, the style of argument they used (in some sense a “depth 2” circuit refutation) was only strong enough to work with  $P_{\text{LW}}$  but not  $\text{BCK}_{\mathbb{B}}$ . In contrast, the argument used by [BCK20] for  $\text{BCK}$  has depth 3. Such circuits are studied in much more detail in [BGJ<sup>+</sup>25].

**Known Bounds for  $\text{NRD}(\text{BCK}, n)$ .** We now shift our focus to discussing what is currently known about  $\text{NRD}(\text{BCK}, n)$ . Since  $\text{BCK}$  is nontrivial, we know that  $\text{NRD}(\text{BCK}, n) = \Omega(n)$ , which is still the best known lower bound to date.

For upper bounds, a baseline upper bound is  $\text{NRD}(\text{BCK}, n) = O(n^2)$ . This can be proved in a few ways (e.g., write out as a quadratic polynomial, or use the techniques of [BCK20]), but perhaps the conceptually simplest way uses the conditional non-redundancy framework established in Section 7. In particular, since  $\text{NRD}(\text{BCK}^+, n) = O(n)$  by (38) and Theorem 6.4, we have that  $\text{NRD}(\text{BCK}, n) = \Theta(\text{NRD}(\text{BCK} \mid \text{BCK}^+, n))$ .

It is straightforward to see that  $\text{NRD}(\text{BCK} \mid \text{BCK}^+, n) = O(n^2)$  via the techniques of Proposition 7.4, as any such non-redundant instance cannot have two clauses sharing two variables (or else one is assigned 000 exactly when the other is). Thus, every pair of variables appears in at most one clause, proving the quadratic upper bound. More strongly, [BCK20] essentially showed that any tripartite non-redundant instance of  $\text{CSP}(\text{BCK} \mid \text{BCK}^+)$  cannot have a 3-cycle. As such, by invoking a hypergraph Turán bound of Sós, Erdős, and Brown [SEB73], they could then prove that

<sup>25</sup>Explicitly, flip the first three bits of  $P_{\text{KPS}}$ , then permute the coordinates of  $P$  according to the cycles  $2 \rightarrow 6 \rightarrow 3 \rightarrow 8 \rightarrow 2$  and  $4 \rightarrow 5 \rightarrow 7 \rightarrow 9 \rightarrow 4$ .

<sup>26</sup>In fact, [LW20]’s combinatorial interpretation of  $Q$  inspired our definition of  $\text{BCK}_{\mathbb{B}}$ .

<sup>27</sup>However, to answer the open question in [LW20], we would need to show that  $\text{BCK}$  has no near-linear kernelization algorithm, however every instance of  $\text{BCK}$  is satisfiable by assigning all 1’s. Such trivialities can be circumvented by considering a variant such as  $\widehat{\text{BCK}} := \{120, 201, 021, 210, 102\}$ .

$\text{NRD}(\text{BCK} \mid \text{BCK}^+, n) = o(n^2)$ , however the improvement over quadratic is known to be at best subpolynomial [RS78].

In Section 7, we indirectly gave a new upper bound for  $\text{NRD}(\text{BCK}, n)$ . In particular, observe that  $\text{BCK} \subsetneq 3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*$ , but  $\text{BCK}^+ \setminus \text{BCK} = 3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}} \setminus 3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^*$ . Thus, any non-redundant instance of  $\text{CSP}(\text{BCK} \mid \text{BCK}^+)$  is also a non-redundant instance of  $\text{CSP}(3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}}^* \mid 3\text{LIN}_{\mathbb{Z}/3\mathbb{Z}})$ . Therefore, as an immediate corollary of Theorem 7.13, we have that  $\text{NRD}(\text{BCK}, n) = O(n^{1.6} \log n)$ . In fact, one can do slightly better by observing in Lemma 7.10 that the task of checking whether  $(\langle x, u_i \rangle, \langle x, v_i \rangle, \langle x, w_i \rangle) \in \text{BCK}$  for  $u_i + v_i + w_i = 0$  can be achieved by the *cubic* polynomial

$$P_i(x) := (1 - \langle x, u_i \rangle^2)(1 - \langle x, v_i \rangle).$$

Leaving the other steps in the proof of Theorem 7.13 unchanged, we can deduce that  $\text{NRD}(\text{BCK}, n) = O(n^{1.5} \log n)$ . The authors doubt the exponent of 1.5 is tight.

**Remark 8.2.** Subsequent to the initial posting of this work, [BGJ<sup>+</sup>25] proved a slightly stronger bound  $O(n^{1.5})$  by connecting non-redundant instances of  $\text{CSP}(\text{BCK})$  to the structure of graphs with girth at least 6. We leave combining these methods to get an even sharper upper bound as a question for future work.

One last observation is that Theorem 7.8 cannot be applied to  $\text{NRD}(\text{BCK}, n)$ . Although the instance constructed lacks 3-cycles, one can check it contains other configurations which make it fail to be non-redundant. However, by Remark 7.9 and using  $g_1 = g_2 = g_3 = 1$ , we deduce that Theorem 7.8 gives an  $\Omega(n^{1.5})$  lower bound on the non-redundancy of

$$R := \{111, 222, 012, 120, 201, 210\} = \text{BCK} \cup \{210\},$$

so in some sense we are “only” one tuple away from a superlinear lower bound for  $\text{NRD}(\text{BCK}, n)$ .

**Question 8.3.** *Does there exist  $\delta > 0$  such that  $\text{NRD}(\text{BCK}, n) = \Omega(n^{1+\delta})$ ?*

## 8.2 Modulo Predicates.

A class of predicates frequently examined in the literature [JP19, LW20, KPS25a] are of the following form

$$\text{LIN}_{k,m,S} := \{x \in \{0, 1\}^k : \text{Ham}(x) \in S \pmod{m}\},$$

where  $k$  and  $m$  are positive integers, and nonempty  $S \subsetneq \{0, 1, \dots, m-1\}$ . Some special cases are understood. For example, if  $|S| = 1$ , then  $\text{NRD}(\text{LIN}_{k,m,S}, n) = \Theta(n)$  because the predicate is affine. For general  $S$ , when  $m$  is a prime, we know that  $\text{NRD}(\text{LIN}_{k,m,S}, n) = O(n^{|S|})$  because  $\text{LIN}_{k,m,S}$  has a degree- $k$  representation of the form

$$\text{LIN}_{k,m,S} = \{0, 1\}^k \cap \left\{ x \in \mathbb{F}_m^k : \prod_{s \in S} \left( \sum_{i=1}^k x_i - s \right) = 0 \right\}.$$

However, even in the prime case, this upper bound does not appear tight in general, as finding a copy of  $\text{OR}_{|S|}$  in  $\text{LIN}_{k,m,S}$  seems to depend on the additive combinatorics properties of  $S \pmod{m}$  (e.g., whether  $S$  is an arithmetic progression).

Even more confounding, the upper bound of  $O(n^{|S|})$  does not apply in general when  $m$  is non-prime. For instance, if  $m = 6$  and  $S = \{0, 1\}$  [KPS25a, LW20, JP19], the polynomial  $\text{Ham}(x)(\text{Ham}(x) - 1)$  has zeros whenever  $\text{Ham}(x) \in \{0, 1, 3, 4\}$  because  $2 \cdot 3 \equiv 0 \pmod{6}$ . However, by setting  $k - 2$  bits of  $\text{LIN}_{k,6,\{0,1\}}$  to zero and flipping the other two bits, we get a copy of  $\text{OR}_2$ , so  $\text{NRD}(\text{LIN}_{k,6,\{0,1\}}) = \Omega(n^2)$ . Furthermore,  $\text{LIN}_{k,6,\{0,1\}}$  has a cubic representation

$$\sum_{1 \leq i_1 < i_2 < i_3 \leq k} x_{i_1} x_{i_2} x_{i_3} - 2 \sum_{1 \leq i_1 < i_2 \leq k} x_{i_1} x_{i_2} \equiv 0 \pmod{6},$$

so  $\text{NRD}(\text{LIN}_{k,6,\{0,1\}}) = O(n^3)$ . That said, [KPS25a] proved  $\text{LIN}_{k,6,\{0,1\}}$  cannot similarly define  $\text{OR}_3$  and conjectured that  $\text{LIN}_{k,6,\{0,1\}}$  fails to have a quadratic representation for all sufficiently large  $k$ . However, they only ruled out a *symmetric* polynomial representation (see Claim D.2 [KPS25a]). We complete the proof here.

**Proposition 8.4.** *For  $k \geq 16$ ,  $\text{LIN}_{k,6,\{0,1\}}$  has no quadratic representation over any Abelian group  $G$ .*

*Proof.* Assume for sake of contradiction that there exists a map  $g : \{0, 1\}^k \rightarrow G$  such that

$$g(x) := g_0 + \sum_{i=1}^k g_i \cdot x_i + \sum_{1 \leq i < j \leq k} g_{ij} \cdot x_i x_j = \mathbf{1}[x \notin \text{LIN}_{k,6,\{0,1\}}],$$

where for any  $h \in G$ ,  $h \cdot 0 = 0$  and  $h \cdot 1 = h$ . In particular, since  $g(x) = 0$  whenever  $\text{Ham}(x) \in \{0, 1\}$ , we know that the constant ( $g_0$ ) and linear ( $g_i, i \in [k]$ ) coefficients of  $g(x)$  all equal 0.

For  $S \subseteq [k]$ , Let  $1_S \in \{0, 1\}^k$  be the indicator vector of  $S$ . Pick  $|S| = 16$  and consider all subsets  $T \subset S$  of size exactly 6. Note that  $g(1_S) \neq 0$  but  $g(1_T) = 0$ . In particular,

$$\begin{aligned} 0 &= \sum_{T \in \binom{S}{6}} g(1_T) = \sum_{T \in \binom{S}{6}} \sum_{\{i,j\} \subset T} g_{ij} \\ &= \sum_{\{i,j\} \subset S} \sum_{T: \{i,j\} \subset T \in \binom{S}{6}} g_{ij} \\ &= \sum_{\{i,j\} \subset S} \binom{14}{4} g_{ij} \\ &= 1001 \cdot g(1_S). \end{aligned}$$

Now, from the theory of combinatorial designs (see Table 1.32 in [CD06]), for  $|S| = 16$ , there exists  $T_1, \dots, T_{16} \subseteq S$  of size exactly 6 such that each pair  $\{i, j\} \subseteq S$  is contained in exactly two of the  $T_i$ 's. By similar logic as before, we have that

$$0 = \sum_{i=1}^{16} g(1_{T_i}) = \sum_{\{i,j\} \subset S} 2g_{ij} = 2 \cdot g(1_S).$$

Thus,

$$g(1_S) = 1001 \cdot g(1_S) - 500 \cdot 2 \cdot g(1_S) = 0,$$

a contradiction. □

### 8.3 Near-polynomial Predicates

We also briefly mention that predicates like  $3\text{LIN}_G^*$  in Section 7 can also be vastly generalized. For simplicity, we consider the case where  $G$  is a cyclic group of order a prime  $p$ , so  $G$  can be identified with the field  $\mathbb{F}_p$ .

Pick integers  $r \geq k \geq 1$  and let  $f : \mathbb{F}_p^r \rightarrow \mathbb{F}_p$  be a polynomial in which every monomial has at most  $k$  distinct variables. Assume also that  $f(0^r) = 0$ . Now consider the following pair of predicates

$$\begin{aligned} \text{POLY}_f &:= \{x \in \mathbb{F}_p^r : f(x) = 0\}, \text{ and} \\ \text{POLY}_f^* &:= \text{POLY}_f \setminus \{0^r\}. \end{aligned}$$

Since computing  $f$  only depends on the behavior of tuples of at most  $k$  variables, we have that  $\text{NRD}(\text{POLY}_f, n) = O(n^k)$  (e.g., [LW20, KPS25a]). However, the behavior of  $\text{NRD}(\text{POLY}_f^*, n)$  appears to be much more difficult to control.

As a concrete example, consider  $p = 3$  and  $f(x) = (x_1 + x_2 + x_3)(x_1 + x_2 + x_3 - 1)$ . Observe that for  $z_1, z_2 \in \{0, 1\}$ , we have that  $R(z_1, z_2, 2)$  is equivalent to  $\text{OR}_2(z_1, z_2)$ . Thus,  $\text{NRD}(\text{POLY}_f^*, n) = \Omega(n^2)$ . However, one can show that  $\text{POLY}_f^*$  cannot gadget-construct  $\text{OR}_3$ , so by Theorem 6.2, we have that  $\text{NRD}(\text{POLY}_f^*, n) = O(n^{2.75})$ . We conjecture that  $\text{NRD}(\text{POLY}_f^*, n) = \Omega(n^{2+\delta})$  for some  $\delta > 0$ .

## 9 Conclusion and Open Questions

In this paper, we closely linked the sparsification of CSPs to combinatorial properties of the corresponding predicate: its non-redundancy (for unweighted sparsification) and its chain length (for weighted sparsification). In addition to the new sparsification bounds we obtained by directly applying known bounds for non-redundancy and chain length, we also developed brand new tools for studying non-redundancy by linking a family of predicates to the structure of matching vector families.

Beyond classifying the non-redundancy and chain length for specific predicates, such as those suggested in Section 8, we now discuss a few directions in which this work could be extended.

**Efficient Sparsification for CSPs.** From an algorithmic perspective, the primary open question is making the sparsification procedures in this paper efficient in terms of the number of variables in the CSP instance. Any such procedure cannot be purely from the perspective of non-linear code sparsification, as our proofs (if turned into algorithms) already run in polynomial time with respect to the size of the code.

In Khanna, Putterman, and Sudan’s recent line of work on linear code sparsification [KPS24a, KPS25a], their first sparsification method was inefficient (as it required solving the minimum-weight codeword problem), but was later made efficient. However, their algorithms heavily used techniques specialized to the structure of linear codes, such as efficiently finding a spanning set for a vector space and conditioning on certain linear conditions being satisfied (i.e., “contraction”). Since Mal’tsev predicates behave similarly to systems of linear equations [BD06], making a result like Theorem 6.5 efficient seems well within reach, but an analogous result for general NP-hard

CSPs<sup>28</sup> seems much more daunting.

**Question 9.1.** *Can the upper bounds of Theorem 1.2 and Theorem 5.5 be made efficient?*

In fact, an efficient algorithm for Theorem 1.2 would imply that every CSP can be kernelized to (approximately) its non-redundancy which would resolve an important open question in the CSP kernelization community (e.g., [Car22]).

**Non-redundancy Versus Chain Length for CSPs.** Another interesting direction question is to understand the relationship between non-redundancy and chain length for CSPs (see the open problems of [Car22]). Recall by Proposition 5.7 that chain length is always at least non-redundancy. However, as discussed in Example 5.8, there are non-linear codes for which non-redundancy and chain length are as distant from each other as possible. Conversely, for linear codes, chain length and non-redundancy are very closely<sup>29</sup> linked. Since the particular codes which are to be sparsified for a CSP predicate  $R$  are highly structured, it remains to be seen what the general relationship between  $\text{CL}(R, n)$  and  $\text{NRD}(R, n)$  is.

**Question 9.2.** *For every nontrivial  $R \subseteq D^r$ , is it the case that  $\text{CL}(R, n) \leq \text{NRD}(R, n) \cdot \log^{O(1)}(n)$ ? In fact, can we even show that there exists a universal constant  $\alpha \geq 1$  such that  $\text{CL}(R, n) \leq O(\text{NRD}(R, n)^\alpha)$ ?*

As discussed in the introduction, the concepts of non-redundancy and chain length appear as the currently best upper and lower bounds on a query complexity problem for CSPs [BCK20]. As such, a positive resolution of Question 9.2 could lead to a resolution of that question as well.

**Extending to Valued CSPs.** A natural variant of CSP sparsification that has been studied in the literature (e.g., [KK24]) is to assign different weights to each assignment to a CSP constraint. In the CSP literature, this is referred to as a *valued CSP* (VCSP) (e.g., [TŽ15, TŽ17]). One method of abstracting a question like this would be to consider *non-Boolean* code sparsification: we have a code  $C \subseteq \Sigma^m$  along with a weight function  $w : \Sigma \rightarrow \mathbb{R}_{\geq 0}$  so that the weight of a codeword  $c \in C$  is  $\sum_{i=1}^m w(c_i)$  (or more generally, a separate weight function from some family for each coordinate). To represent an “unsatisfied” assignment, we also assume there is  $0 \in \Sigma$  with  $w(0) = 0$ .

As a baseline upper bound, split  $C$  into  $|\Sigma| - 1$  codes  $\{C_\alpha : \alpha \in \Sigma \setminus \{0\}\}$  such that each  $c \in C$  corresponds to  $(\mathbf{1}[c_i = \alpha] : i \in [m])$  in  $C_\alpha$ . Then, we can combine the sparsifiers for each  $C_\alpha$  to get a sparsifier for  $C$ . In general, this is a poor bound, as correctly sparsifying the count of each symbol is a much stronger property than estimating their weighted sum.

Likewise, replacing each nonzero symbol with 1 to reduce to the Boolean case does not in general yield a tight lower bound. For example, consider  $\Sigma = \{0, 1, 2\}$  and  $w(\alpha) = \alpha$ . The Booleanization of the code  $C := \{1, 2\}^m$  gives every codeword a weight of  $m$ , so it has a trivial  $O(1)$  sparsifier. However, one can work out that any  $\varepsilon$ -sparsifier of  $C$  requires  $\Omega(m)$  support size.

**Question 9.3.** *What is the optimal sparsifiability of valued CSPs / non-Boolean codes?*

<sup>28</sup>Or more precisely, CSPs which do not have a tight gadget reduction from any efficiently-solvable CSP.

<sup>29</sup>For codes over fields, chain length and non-redundancy are equal, but for codes over an Abelian group  $G$  they can differ by a factor of  $\log_2 |G|$ . For instance, consider for  $G := \mathbb{Z}/2^k\mathbb{Z}$  and let  $C \subset G^k$  be the span of  $\{2^{k-1}e_1, 2^{k-2}e_1 + 2^{k-1}e_2, \dots, e_1 + 2e_2 + \dots + 2^{k-1}e_k\}$ . The non-redundancy of  $C$  is 1 but the chain length is  $k = \log_2 |G|$ .

**Improved Lower Bounds.** For both unweighted and weighted sparsification, our current lower bounds on  $\varepsilon$ -sparsification are purely a function of  $n$ . Can we improve the upper bound as a function of  $\varepsilon$ ? For instance, a lower bound of  $n/\varepsilon^2$  is known for cut sparsifiers, which is tight [ACK<sup>+</sup>16, CKST19]. However, the lower bound techniques seem to crucially use that the non-redundancy/chain length of CUT is linear, so many nearly-isolated clauses can be used in a lower bound. For general CSPs, the optimal relationship between  $n$  and  $\varepsilon$  is unclear. For instance, if  $\text{NRD}(R, n) = \Omega(n^r)$  (i.e.,  $R$  is  $\text{OR}_r$  or  $\overline{R}$  is  $\text{AND}_r$ ), then no improvement as a function of  $\varepsilon$  is possible as every sparsifier has size at most  $n^r$ .

**Question 9.4.** For every  $R \subseteq D^r$  and  $\varepsilon \in (0, 1)$ , is it the case that

$$\Omega(\min(\text{NRD}(R, n)/\varepsilon^2, n^r)) \leq \text{SPR}(R, n, \varepsilon)?, \text{ and}$$

$$\Omega(\min(\text{CL}(R, n)/\varepsilon^2, n^r)) \leq \text{wSPR}(R, n, \varepsilon)?$$

**Average-case Behavior.** As our final area of investigation, it would be interesting to understand the behavior of the non-redundancy (and thus sparsity<sup>30</sup>) of random instances of  $\text{CSP}(R)$ . In particular, the size of the sparsifier may be significantly better than the guarantees of Theorem 1.2. More formally for a variable set  $X$  of size  $n$ , pick  $m \in [n, n^r]$  and sample uniformly at random  $m$  clauses  $Y$  from  $X^r$ . What is the non-redundancy? If  $R$  is an affine CSP like CUT or  $3\text{LIN}_G$ , then it is not hard to see that if  $m = \tilde{\Omega}(n)$ , then the space of assignments to the clauses will have rank  $\Omega(n)$ , so average-case behavior is essentially the same as worst-case behavior. One can make a similar argument for certain “higher degree” clauses like  $\text{AND}_r$ .

However, for a predicate like  $3\text{LIN}_G^*$ , the non-redundant instance constructed in Theorem 7.8 seems rather “fragile.” In particular, a random sample of say  $m := n^{1.5}$  clauses seems to not include  $\tilde{\Omega}(n)$  clauses from any copy of the construction in Theorem 7.8. As such, it could be possible that the non-redundancy of  $3\text{LIN}_G^*$  is near-linear in this regime.

**Question 9.5.** For random instances of  $\text{CSP}(R)$ , what is the relationship between clause count and non-redundancy?

Subsequent to the initial posting of our work, the intent behind Question 9.5 was largely solved by [BGP25] who computed an optimal sparsifier size for random instances of any *valued* CSP. Two important features of this result are (1) the answer is always an integral polynomial (i.e.,  $\tilde{O}_\varepsilon(n^k)$  for some integer  $k$ ) and (2) depending on the model the optimal sparsifier size can be *non-monotone* in the number of edges! However, since the sparsification criterion use by [BGP25] was not exactly non-redundancy, Question 9.5 is technically still open.

## References

- [AB99] Martin Anthony and Peter L. Bartlett. *Neural network learning: theoretical foundations*. Cambridge University Press, Cambridge, 1999. doi:10.1017/CB09780511624216.

---

<sup>30</sup>Note that via the reduction from CSP sparsification to code sparsification, there exists a sparsifier of an instance  $(X, Y)$  of  $\text{CSP}(\overline{R})$  of size approximately the non-redundancy of  $(X, Y)$  when viewed as an instance of  $\text{CSP}(R)$ .

- [ACK<sup>+</sup>16] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On Sketching Quadratic Forms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS '16*, pages 311–319, New York, NY, USA, January 2016. Association for Computing Machinery. doi:10.1145/2840728.2840753.
- [AG09] Kook Jin Ahn and Sudipto Guha. Graph Sparsification in the Semi-streaming Model. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris Nikolettseas, and Wolfgang Thomas, editors, *Automata, Languages and Programming*, volume 5556, pages 328–338. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. doi:10.1007/978-3-642-02930-1\_27.
- [AG21] Omar Alrabiah and Venkatesan Guruswami. Visible rank and codes with locality. In Mary Wootters and Laura Sanità, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, volume 207 of *LIPICs*, pages 57:1–57:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.APPROX/RANDOM.2021.57.
- [AGM12a] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 459–467. SIAM, 2012. doi:10.1137/1.9781611973099.40.
- [AGM12b] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: sparsification, spanners, and subgraphs. In Michael Benedikt, Markus Krötzsch, and Maurizio Lenzerini, editors, *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2012, Scottsdale, AZ, USA, May 20-24, 2012*, pages 5–14. ACM, 2012. doi:10.1145/2213556.2213560.
- [AHS24] Ryan Alweiss, Brice Huang, and Mark Sellke. Improved lower bound for Frankl’s union-closed sets conjecture. *Electron. J. Combin.*, 31(3):Paper No. 3.35, 11, 2024. doi:10.37236/12232.
- [AK16] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. *J. ACM*, 63(2):12:1–12:35, 2016. doi:10.1145/2837020.
- [ALWZ21] Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. Improved bounds for the sunflower lemma. *Ann. of Math. (2)*, 194(3):795–815, 2021. doi:10.4007/annals.2021.194.3.5.
- [AU19] Raman Arora and Jalaj Upadhyay. On differentially private graph sparsification and applications. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13378–13389, 2019. URL: <https://proceedings.neurips.cc/paper/2019/hash/e44e875c12109e4fa3716c05008048b2-Abstract.html>.

- [BBDS12] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The johnson-lindenstrauss transform itself preserves differential privacy. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 410–419. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.67.
- [BCH<sup>+</sup>13] Christian Bessiere, Remi Coletta, Emmanuel Hebrard, George Katsirelos, Nadjib Lazaar, Nina Narodytska, Claude-Guy Quimper, and Toby Walsh. Constraint acquisition via partial queries. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 475–481. IJCAI/AAAI, 2013. URL: <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6659>.
- [BCK20] Christian Bessiere, Clément Carbonnel, and George Katsirelos. Chain length and cps learnable with few queries. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 1420–1427. AAAI Press, 2020. doi:10.1609/aaai.v34i02.5499.
- [BD06] Andrei Bulatov and Víctor Dalmau. A Simple Algorithm for Mal’tsev Constraints. *SIAM Journal on Computing*, 36(1):16–27, January 2006. doi:10.1137/050628957.
- [BDL12] Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New lower bounds for matching vector codes. *CoRR*, abs/1204.1367, 2012. URL: <http://arxiv.org/abs/1204.1367>, arXiv:1204.1367.
- [Ber11] Clifford Bergman. *Universal Algebra: Fundamentals and Selected Topics*. Chapman and Hall/CRC, New York, November 2011. doi:10.1201/9781439851302.
- [Beu21] Stijn Beukers. Extending the Scope of Algebraic Kernelization for Constraint Satisfaction Problems. Master’s thesis, Eindhoven University of Technology, 2021.
- [BF24] Laszlo Babai and Peter Frankl. *Linear Algebra Methods in Combinatorics*. 2024.
- [BG25] Joshua Brakensiek and Venkatesan Guruswami. Redundancy is all you need. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 1614–1625. ACM, 2025. doi:10.1145/3717823.3718212.
- [BGJ<sup>+</sup>25] Joshua Brakensiek, Venkatesan Guruswami, Bart M. P. Jansen, Victor Lagerkvist, and Magnus Wahlström. The richness of CSP non-redundancy. *CoRR*, abs/2507.07942, 2025. arXiv:2507.07942, doi:10.48550/arXiv.2507.07942.
- [BGP25] Joshua Brakensiek, Venkatesan Guruswami, and Aaron Putterman. Tight bounds for sparsifying random cps. *CoRR*, abs/2508.13345, 2025. arXiv:2508.13345, doi:10.48550/arXiv.2508.13345.

- [BGP26a] Joshua Brakensiek, Venkatesan Guruswami, and Aaron Putterman. Classification of non-redundancy of boolean predicates of arity 4. *CoRR*, abs/2603.21353, 2026. arXiv: 2603.21353, doi:10.48550/arXiv.2603.21353.
- [BGP26b] Joshua Brakensiek, Venkatesan Guruswami, and Aaron Putterman. Multiplicative error set system sparsification: A simpler proof via chain length contraction, 2026. URL: <https://arxiv.org/abs/2605.01508>, arXiv:2605.01508.
- [BHM<sup>+</sup>21] Vladimir Braverman, Avinatan Hassidim, Yossi Matias, Mariano Schain, Sandeep Silwal, and Samson Zhou. Adversarial robustness of streaming algorithms through importance sampling. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 3544–3557, 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/1d01bd2e16f57892f0954902899f0692-Abstract.html>.
- [BK96] András A. Benczúr and David R. Karger. Approximating  $s$ - $t$  minimum cuts in  $\tilde{O}(n^2)$  time. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 47–55. ACM, 1996. doi:10.1145/237814.237827.
- [BK12] Christian Bessiere and Frédéric Koriche. Non learnability of constraint networks with membership queries. Technical report, Technical report, Coconut, Montpellier, France, 2012.
- [BKLM26] Arpon Basu, Pravesh K. Kothari, Yang P. Liu, and Raghu Meka. Sparsifying sums of positive semidefinite matrices. In Kasper Green Larsen and Barna Saha, editors, *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2026, Vancouver, BC, Canada, January 11-14, 2026*, pages 6042–6064. SIAM, 2026. doi:10.1137/1.9781611978971.216.
- [BKW17] Libor Barto, Andrei A. Krokhin, and Ross Willard. Polymorphisms, and how to use them. In Andrei A. Krokhin and Stanislav Živný, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, Dagstuhl Follow-Ups, pages 1–44. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. URL: <https://doi.org/10.4230/DFU.Vol17.15301.1>, doi:10.4230/DFU.VOL7.15301.1.
- [Bop23] Ravi B. Boppana. A useful inequality for the binary entropy function, 2023. URL: <https://arxiv.org/abs/2301.09664>, arXiv:2301.09664.
- [BSS12] Joshua D. Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. *Siam Journal On Computing*, 41(6):1704–1721, 2012. doi:10.1137/090772873.
- [BŽ20] Silvia Butti and Stanislav Živný. Sparsification of Binary CSPs. *SIAM Journal on Discrete Mathematics*, 34(1):825–842, January 2020. doi:10.1137/19M1242446.
- [Cam22] Stijn Cambie. Better bounds for the union-closed sets conjecture using the entropy approach. *arXiv preprint arXiv:2212.12500*, 2022.

- [Car22] Clément Carbonnel. On Redundancy in Constraint Satisfaction Problems. In *28th International Conference on Principles and Practice of Constraint Programming (CP 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.CP.2022.11.
- [CD06] Charles J. Colbourn and Jeffrey H. Dinitz, editors. *Handbook of Combinatorial Designs*. Chapman and Hall/CRC, New York, 2 edition, November 2006. doi:10.1201/9781420010541.
- [CDL<sup>+</sup>24] Xi Chen, Anindya De, Yuhao Li, Shivam Nadimpalli, and Rocco A. Servedio. Testing intersecting and union-closed families. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, January 30 to February 2, 2024, Berkeley, CA, USA*, volume 287 of *LIPIcs*, pages 33:1–33:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ITCS.2024.33.
- [Chu14] Fan Chung. A brief survey of pagerank algorithms. *IEEE Trans. Netw. Sci. Eng.*, 1(1):38–42, 2014. doi:10.1109/TNSE.2014.2380315.
- [CJP20] Hubie Chen, Bart M. P. Jansen, and Astrid Pieterse. Best-Case and Worst-Case Sparsifiability of Boolean CSPs. *Algorithmica*, 82(8):2200–2242, August 2020. doi:10.1007/s00453-019-00660-y.
- [CKL<sup>+</sup>22] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.
- [CKM<sup>+</sup>11] Paul Christiano, Jonathan A. Kelner, Aleksander Madry, Daniel A. Spielman, and Shang-Hua Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11*, pages 273–282, New York, NY, USA, June 2011. Association for Computing Machinery. doi:10.1145/1993636.1993674.
- [CKN20] Yu Chen, Sanjeev Khanna, and Ansh Nagda. Near-linear Size Hypergraph Cut Sparsifiers. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 61–72, November 2020. doi:10.1109/FOCS46700.2020.00015.
- [CKST19] Charles Carlson, Alexandra Kolla, Nikhil Srivastava, and Luca Trevisan. Optimal lower bounds for sketching graph cuts. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2565–2569. SIAM, Philadelphia, PA, 2019. doi:10.1137/1.9781611975482.158.
- [CL22] Zachary Chase and Shachar Lovett. Approximate union closed conjecture, 2022. URL: <https://arxiv.org/abs/2211.11689>, arXiv:2211.11689.
- [CSZ22] Jie Chen, Yousef Saad, and Zechen Zhang. Graph coarsening: From scientific computing to machine learning. *SeMA Journal*, 79(1):187–223, March 2022. doi:10.1007/s40324-021-00282-x.

- [DGY11] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching Vector Codes. *SIAM Journal on Computing*, 40(4):1154–1178, January 2011. doi:10.1137/100804322.
- [DvM14] Holger Dell and Dieter van Melkebeek. Satisfiability Allows No Nontrivial Sparsification unless the Polynomial-Time Hierarchy Collapses. *J. ACM*, 61(4):23:1–23:27, July 2014. doi:10.1145/2629620.
- [Efr12] Klim Efremenko. 3-query locally decodable codes of subexponential length. volume 41, pages 1694–1703, 2012. doi:10.1137/090772721.
- [Erd64] P. Erdős. On extremal problems of graphs and generalized graphs. *Israel Journal of Mathematics*, 2(3):183–190, September 1964. doi:10.1007/BF02759942.
- [ERS05] Guy Even, Dror Rawitz, and Shimon (Moni) Shahar. Hitting sets when the VC-dimension is small. *Information Processing Letters*, 95(2):358–362, July 2005. doi:10.1016/j.ipl.2005.03.010.
- [FHHP19] Wai Shing Fung, Ramesh Hariharan, Nicholas J. A. Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. *SIAM J. Comput.*, 48(4):1196–1223, 2019. doi:10.1137/16M1091666.
- [FK17] Arnold Filtser and Robert Krauthgamer. Sparsification of Two-Variable Valued Constraint Satisfaction Problems. *SIAM Journal on Discrete Mathematics*, 31(2):1263–1276, January 2017. doi:10.1137/15M1046186.
- [Fra95] P. Frankl. Extremal set systems. In *Handbook of combinatorics, Vol. 1, 2*, pages 1293–1329. Elsevier Sci. B. V., Amsterdam, 1995.
- [FW02] Eugene C. Freuder and Richard J. Wallace. Suggestion strategies for constraint-based matchmaker agents. *International Journal on Artificial Intelligence Tools*, 11(01):3–18, March 2002. doi:10.1142/S0218213002000769.
- [GBY<sup>+</sup>24] Yuxin Guo, Deyu Bo, Cheng Yang, Zhiyuan Lu, Zhongjian Zhang, Jixi Liu, Yufei Peng, and Chuan Shi. Data-centric Graph Learning: A Survey. (arXiv:2310.04987), January 2024. arXiv:2310.04987, doi:10.48550/arXiv.2310.04987.
- [GHSY12] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the Locality of Codeword Symbols. *IEEE Transactions on Information Theory*, 58(11):6925–6934, November 2012. doi:10.1109/TIT.2012.2208937.
- [Gil22] Justin Gilmer. A constant lower bound for the union-closed sets conjecture. (arXiv:2211.09055), November 2022. arXiv:2211.09055, doi:10.48550/arXiv.2211.09055.
- [Gop09] Parikshit Gopalan. A note on Efremenko’s locally decodable codes. *Electron. Colloquium Comput. Complex.*, TR09, 2009. URL: <https://eccc.weizmann.ac.il/report/2009/069>, arXiv:TR09–069.
- [HKTH16] Alireza Hadjighasem, Daniel Karrasch, Hiroshi Teramoto, and George Haller. Spectral-clustering approach to Lagrangian vortex detection. *Physical Review E*, 93(6):063107, June 2016. doi:10.1103/PhysRevE.93.063107.

- [HLM<sup>+</sup>26] Jun-Ting Hsieh, Daniel Z. Lee, Sidhanth Mohanty, Aaron Putterman, and Rachel Yun Zhang. Sparsifying Cayley graphs on every group. In Kasper Green Larsen and Barna Saha, editors, *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2026, Vancouver, BC, Canada, January 11-14, 2026*, pages 6029–6041. SIAM, 2026. doi:10.1137/1.9781611978971.215.
- [HN10] Danny Harnik and Moni Naor. On the compressibility of *NP* instances and cryptographic applications. *SIAM J. Comput.*, 39(5):1667–1713, 2010. doi:10.1137/060668092.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- [Jan20] Bart M. P. Jansen. Crossing Paths with Hans Bodlaender: A Personal View on Cross-Composition for Sparsification Lower Bounds. In Fedor V. Fomin, Stefan Kratsch, and Erik Jan van Leeuwen, editors, *Treewidth, Kernels, and Algorithms: Essays Dedicated to Hans L. Bodlaender on the Occasion of His 60th Birthday*, pages 89–111. Springer International Publishing, Cham, 2020. doi:10.1007/978-3-030-42071-0\_8.
- [JLR17] Peter Jonsson, Victor Lagerkvist, and Biman Roy. Time Complexity of Constraint Satisfaction via Universal Algebra. *LIPICs, Volume 83, MFCS 2017*, 83:17:1–17:15, 2017. doi:10.4230/LIPICs.MFCS.2017.17.
- [JP19] Bart M. P. Jansen and Astrid Pieterse. Optimal Sparsification for Some Binary CSPs Using Low-Degree Polynomials. *ACM Trans. Comput. Theory*, 11(4):28:1–28:26, August 2019. doi:10.1145/3349618.
- [JW24] Bart M. P. Jansen and Michal Włodarczyk. Optimal polynomial-time compression for boolean max CSP. *ACM Trans. Comput. Theory*, 16(1):4:1–4:20, 2024. doi:10.1145/3624704.
- [Kar93] David R. Karger. Global min-cuts in rnc, and other ramifications of a simple min-cut algorithm. In Vijaya Ramachandran, editor, *Proceedings of the Fourth Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, 25-27 January 1993, Austin, Texas, USA*, pages 21–30. ACM/SIAM, 1993. URL: <http://dl.acm.org/citation.cfm?id=313559.313605>.
- [Kee11] Peter Keevash. Hypergraph Turán problems. In Robin Chapman, editor, *Surveys in Combinatorics 2011*, pages 83–140. Cambridge University Press, 1 edition, June 2011. doi:10.1017/CB09781139004114.004.
- [KK15] Dmitry Kogan and Robert Krauthgamer. Sketching Cuts in Graphs and Hypergraphs. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15*, pages 367–376, New York, NY, USA, January 2015. Association for Computing Machinery. doi:10.1145/2688073.2688093.
- [KK24] Yotam Kenneth and Robert Krauthgamer. Cut sparsification and succinct representation of submodular hypergraphs. In Karl Bringmann, Martin Grohe, Gabriele

- Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, Tallinn, Estonia, July 8-12, 2024*, LIPIcs, pages 97:1–97:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ICALP.2024.97.
- [KKTY21] Michael Kapralov, Robert Krauthgamer, Jakab Tardos, and Yuichi Yoshida. Towards tight bounds for spectral sparsification of hypergraphs. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 598–611. ACM, 2021. doi:10.1145/3406325.3451061.
- [KLM<sup>+</sup>17] Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. *SIAM J. Comput.*, 46(1):456–477, 2017. doi:10.1137/141002281.
- [KLOS14] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multi-commodity generalizations. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 217–226. ACM, New York, 2014. doi:10.1137/1.9781611973402.16.
- [KMP14] Ioannis Koutis, Gary L. Miller, and Richard Peng. Approaching optimality for solving SDD linear systems. *SIAM J. Comput.*, 43(1):337–354, 2014. doi:10.1137/110845914.
- [Kni94] Emanuel Knill. Graph generated union-closed families of sets, 1994. URL: <https://arxiv.org/abs/math/9409215>, arXiv:math/9409215.
- [KPS24a] Sanjeev Khanna, Aaron L. Putterman, and Madhu Sudan. Code Sparsification and its Applications. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 5145–5168. Society for Industrial and Applied Mathematics, January 2024. doi:10.1137/1.9781611977912.185.
- [KPS24b] Sanjeev Khanna, Aaron L. Putterman, and Madhu Sudan. Near-optimal Size Linear Sketches for Hypergraph Cut Sparsifiers. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, October 2024. doi:10.48550/arXiv.2407.03934.
- [KPS25a] Sanjeev Khanna, Aaron Putterman, and Madhu Sudan. Efficient algorithms and new characterizations for CSP sparsification. pages 407–416, 2025. doi:10.1145/3717823.3718205.
- [KPS25b] Sanjeev Khanna, Aaron Putterman, and Madhu Sudan. A theory of spectral CSP sparsification. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, Aarhus, Denmark, July 8-11, 2025*, LIPIcs, pages 107:1–107:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ICALP.2025.107.
- [KVV04] Ravi Kannan, Santosh Vempala, and Adrian Vetta. On clusterings: Good, bad and spectral. *J. ACM*, 51(3):497–515, May 2004. doi:10.1145/990308.990313.

- [LCY<sup>+</sup>21] Dongsheng Luo, Wei Cheng, Wenchao Yu, Bo Zong, Jingchao Ni, Haifeng Chen, and Xiang Zhang. Learning to Drop: Robust Graph Neural Network via Topological Denoising. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 779–787, Virtual Event Israel, March 2021. ACM. doi:10.1145/3437963.3441734.
- [LLMV10] Arnaud Lallouet, Matthieu Lopez, Lionel Martin, and Christel Vrain. On Learning Constraint Problems. In *2010 22nd IEEE International Conference on Tools with Artificial Intelligence*, volume 1, pages 45–52, October 2010. doi:10.1109/ICTAI.2010.16.
- [LMW26] Shachar Lovett, Raghu Meka, and Yimeng Wang. Moonflowers and efficient code sparsification, 2026. URL: <https://arxiv.org/abs/2605.08676>, arXiv:2605.08676.
- [LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999. doi:10.1145/331524.331526.
- [LS14] Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in  $\tilde{O}(\text{rank})$  iterations and faster algorithms for maximum flow. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 424–433. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.52.
- [LS18] Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almost-linear time. *SIAM J. Comput.*, 47(6):2315–2336, 2018. doi:10.1137/16M1061850.
- [LW17] Victor Lagerkvist and Magnus Wahlström. Kernelization of Constraint Satisfaction Problems: A Study Through Universal Algebra. In J. Christopher Beck, editor, *Principles and Practice of Constraint Programming*, pages 157–171, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-66158-2\_11.
- [LW18] Victor Lagerkvist and Magnus Wahlström. Which NP-Hard SAT and CSP Problems Admit Exponentially Improved Algorithms? (arXiv:1801.09488), January 2018. arXiv:1801.09488, doi:10.48550/arXiv.1801.09488.
- [LW20] Victor Lagerkvist and Magnus Wahlström. Sparsification of SAT and CSP Problems via Tractable Extensions. *ACM Transactions on Computation Theory*, 12(2):1–29, June 2020. doi:10.1145/3389411.
- [Mal54] A. I. Malcev. On the general theory of algebraic systems. *Mat. Sb. (N.S.)*, 35(77):3–20, 1954.
- [McG14] Andrew McGregor. Graph stream algorithms: a survey. *SIGMOD Rec.*, 43(1):9–20, 2014. doi:10.1145/2627692.2627694.
- [MR95] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995. URL: <https://doi.org/10.1017/cbo9780511814075>, doi:10.1017/CB09780511814075.

- [MU05] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. doi:10.1017/CB09780511813603.
- [PBS08] Mathias Paulin, Christian Bessiere, and Jean Sallantin. Automatic Design of Robot Behaviors through Constraint Network Acquisition. In *2008 20th IEEE International Conference on Tools with Artificial Intelligence*, volume 1, pages 275–282, November 2008. doi:10.1109/ICTAI.2008.83.
- [Peb22] Luke Pebody. Extension of a Method of Gilmer. (arXiv:2211.13139), November 2022. arXiv:2211.13139, doi:10.48550/arXiv.2211.13139.
- [Qua24] Kent Quanrud. Quotient sparsification for submodular functions. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 5209–5248. Society for Industrial and Applied Mathematics, January 2024. doi:10.1137/1.9781611977912.187.
- [Raf24] Akbar Rafiey. Decomposable submodular maximization in federated setting. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 41841–41866. PMLR / OpenReview.net, 2024. URL: <https://proceedings.mlr.press/v235/rafiey24a.html>.
- [Rag07] Prasad Raghavendra. A note on Yekhanin’s locally decodable codes. *Electron. Colloquium Comput. Complex.*, TR07-016, 2007. URL: <https://eccc.weizmann.ac.il/eccc-reports/2007/TR07-016/index.html>.
- [RS78] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976)*, Vol. II, volume 18 of *Colloq. Math. Soc. János Bolyai*, pages 939–945. North-Holland, Amsterdam-New York, 1978.
- [Sau72] N Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, July 1972. doi:10.1016/0097-3165(72)90019-2.
- [Saw23] Will Sawin. An improved lower bound for the union-closed set conjecture. (arXiv:2211.11504), June 2023. arXiv:2211.11504, doi:10.48550/arXiv.2211.11504.
- [Sch24] Gregory Schwartzman. Mini-batch submodular maximization. *CoRR*, abs/2401.12478, 2024. arXiv:2401.12478, doi:10.48550/arXiv.2401.12478.
- [SEB73] V. T. Sós, P. Erdős, and W. G. Brown. On the existence of triangulated spheres in 3-graphs, and related problems. *Periodica Mathematica Hungarica*, 3(3-4):221–228, September 1973. doi:10.1007/BF02018585.
- [She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J. Math.*, 41:247–261, 1972. URL: <http://projecteuclid.org/euclid.pjm/1102968432>.

- [SPR11] Venu Satuluri, Srinivasan Parthasarathy, and Yiye Ruan. Local graph sparsification for scalable clustering. In Timos K. Sellis, Renée J. Miller, Anastasios Kementsietsidis, and Yannis Velegarakis, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 721–732. ACM, 2011. doi:10.1145/1989323.1989399.
- [ST04] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.
- [ST11] Daniel A. Spielman and Shang-Hua Teng. Spectral Sparsification of Graphs. *SIAM Journal on Computing*, 40(4):981–1025, January 2011. doi:10.1137/08074489X.
- [SV26a] Amatya Sharma and Santhoshini Velusamy. Characterizing streaming decidability of csps via non-redundancy, 2026. URL: <https://arxiv.org/abs/2604.21922>, arXiv:2604.21922.
- [SV26b] Amatya Sharma and Santhoshini Velusamy. Non-redundancy of low-arity symmetric boolean csps, 2026. URL: <https://arxiv.org/abs/2605.14007>, arXiv:2605.14007.
- [Tak23] Rustem Takhanov. The algebraic structure of the densification and the sparsification tasks for CSPs. *Constraints*, 28(1):13–44, March 2023. doi:10.1007/s10601-022-09340-1.
- [TN22] Kunihiko Taira and Aditya G. Nair. Network-based analysis of fluid flows: Progress and outlook. *Progress in Aerospace Sciences*, 131:100823, May 2022. doi:10.1016/j.paerosci.2022.100823.
- [Tod16] Michael J. Todd. *Minimum volume ellipsoids - theory and algorithms*, volume 23 of *MOS-SIAM Series on Optimization*. SIAM, 2016.
- [TŽ15] Johan Thapper and Stanislav Živný. Necessary Conditions for Tractability of Valued CSPs. *SIAM Journal on Discrete Mathematics*, 29(4):2361–2384, January 2015. doi:10.1137/140990346.
- [TŽ17] Johan Thapper and Stanislav Živný. The Power of Sherali–Adams Relaxations for General-Valued CSPs. *SIAM Journal on Computing*, 46(4):1241–1279, January 2017. doi:10.1137/16M1079245.
- [Van18] Vitaly Vanchurin. Information Graph Flow: A Geometric Approximation of Quantum and Statistical Systems. *Foundations of Physics*, 48(6):636–653, June 2018. doi:10.1007/s10701-018-0166-z.
- [VC15] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, Cham, 2015. Reprint of Theor. Probability Appl. **16** (1971), 264–280.
- [Vis13] Nisheeth K. Vishnoi.  $Lx = b$ . *Foundations and Trends® in Theoretical Computer Science*, 8(1–2):1–141, May 2013. doi:10.1561/04000000054.

- [vN28] J. v. Neumann. Zur Theorie der Gesellschaftsspiele. *Math. Ann.*, 100(1):295–320, 1928. doi:10.1007/BF01448847.
- [Wak24] Tanay Wakhare. Two studies of constraints in high dimensions: Entropy inequalities and the randomized symmetric binary perceptron. Master’s thesis, Massachusetts Institute of Technology, 2024.
- [Wój99] Piotr Wójcik. Union-closed families of sets. *Discrete Mathematics*, 199(1):173–182, March 1999. doi:10.1016/S0012-365X(98)00208-8.
- [Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *J. ACM*, 55(1):1:1–1:16, February 2008. doi:10.1145/1326554.1326555.
- [Yek12] Sergey Yekhanin. Locally Decodable Codes. *Foundations and Trends® in Theoretical Computer Science*, 6(3):139–255, October 2012. doi:10.1561/04000000030.
- [Yos26] Yuichi Yoshida. Solving hypergraph laplacian systems in almost-linear time, 2026. URL: <https://arxiv.org/abs/2604.27651>, arXiv:2604.27651.
- [ZSW<sup>+</sup>23] Shichang Zhang, Atefeh Sohrabizadeh, Cheng Wan, Zijie Huang, Ziniu Hu, Yewen Wang, Yingyan, Lin, Jason Cong, and Yizhou Sun. A Survey on Graph Neural Network Acceleration: Algorithms, Systems, and Customized Hardware. (arXiv:2306.14052), June 2023. arXiv:2306.14052, doi:10.48550/arXiv.2306.14052.