On Average Baby PIH and Its Applications

Yuwei Liu
Shanghai Jiao Tong University, China

Yijia Chen
Shanghai Jiao Tong University, China

Shuangle Li
Nanjing University, China

Bingkai Lin
Nanjing University, China

Xin Zheng
Nanjing University, China

Abstract

The *Parameterized Inapproximability Hypothesis* (PIH) asserts that no FPT algorithm can decide whether a given 2CSP instance parameterized by the number of variables is satisfiable, or at most a constant fraction of its constraints can be satisfied simultaneously. In a recent breakthrough, Guruswami, Lin, Ren, Sun, and Wu (STOC 2024) proved the PIH under the Exponential Time Hypothesis (ETH). However, it remains a major open problem whether the PIH can be established assuming only $\mathsf{W}[1] \neq \mathsf{FPT}$. Towards this goal, Guruswami, Ren, and Sandeep (CCC 2024) showed a weaker version of the PIH called the *Baby PIH* under $\mathsf{W}[1] \neq \mathsf{FPT}$. In addition, they proposed one more intermediate assumption known as the *Average Baby PIH*, which might lead to further progress on the PIH. As the main contribution of this paper, we prove that the Average Baby PIH holds assuming $\mathsf{W}[1] \neq \mathsf{FPT}$.

Given a 2CSP instance where the number of its variables is the parameter, the Average Baby PIH states that no FPT algorithm can decide whether (i) it is satisfiable or (ii) any *multi-assignment* that satisfies all constraints must assign each variable more than r values on *average* for any fixed constant $r > 1$. So there is a *gap* between (i) and (ii) on the average number of values that are assigned to a variable, i.e., 1 vs. r. If this gap occurs in *each* variable instead of on average, we get the original Baby PIH. So central to our paper is an FPT self-reduction for 2CSP instances that turns the above gap for each variable into a gap on average. By the known W[1]-hardness for the Baby PIH, this proves that the Average Baby PIH holds under $\mathsf{W}[1] \neq \mathsf{FPT}$.

As applications, we obtain (i) for the first time, the W[1]-hardness of constant approximating k-ExactCover, and (ii) a tight relationship between running time lower bounds in the Average Baby PIH and approximating the parameterized Nearest Codeword Problem (k-NCP).

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Average Baby PIH, Parameterized Inapproximability, Constraint Satisfaction Problem, Exact Set Cover, W[1]-hardness

Digital Object Identifier 10.4230/LIPIcs.STACS.2025.65

Funding Yuwei Liu and Yijia Chen are supported by the National Natural Science Foundation of China (Project 62372291).

Acknowledgements The authors want to thank Guohang Liu, Mingjun Liu, Yangluo Zheng for discussions in the early stage of this work. The comments and suggestions from anonymous reviewers also help to improve the paper significantly. In particular, Theorem 2 is due to one of them.

42nd International Symposium on Theoretical Aspects of Computer Science (STACS 2025).
Editors: Olaf Beyersdorff, Michał Pilipczuk, Elaine Pimentel, and Nguyễn Kim Thắng;
Article No. 65; pp. 65:1–65:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In classical complexity theory, the PCP theorem [3, 2, 8] serves as an essential tool for proving most of the existing results in the hardness of approximation. As a variant, the Multi-Assignment PCP theorem [1, Lemma 11] states that, for any constant $r > 1$ and $0 < \varepsilon < 1$, it is even NP-hard to decide whether a CSP instance is satisfiable, or any multi-assignment (see Definition 7.) that assigns each variable no more than r values cannot satisfy a $(1-\varepsilon)$-fraction of constraints. Among others, the Multi-Assignment PCP theorem was used to show the NP-hardness of approximating SetCover [1]. It turns out that the Multi-Assignment PCP theorem is a simple consequence of the PCP theorem by a straightforward probabilistic argument. Nevertheless, Barto and Kozik [4] gave a direct and purely combinatorial proof for the simple case of $\varepsilon = 0$. Observe that it means that the CSP instances under consideration are either satisfiable or cannot be satisfied by a desired multi-assignment. This restricted version of the Multi-Assignment PCP theorem is termed the *Baby PCP Theorem* in [4].

As an analog of the PCP theorem in parameterized complexity theory, the Parameterized Inapproximability Hypothesis [21], PIH for short, is an important assumption from which we can prove many FPT inapproximability results, including the inapproximability of k-Clique, k-ExactCover [13], and Direct Odd Cycle Transversal [21], etc. It claims that for some constant $0 < \varepsilon < 1$, no $f(k) \cdot n^{O(1)}$-time (i.e., FPT) algorithm can distinguish a satisfiable 2CSP instance with k variables from one where less than $(1-\varepsilon)$-fraction of constraints can be satisfied simultaneously [21]. Unlike the PCP theorem, the PIH is still a major open problem in parameterized complexity. The current state of the art is that the PIH holds under the Exponential Time Hypothesis (ETH) [12, 11], and a proof of the PIH under the minimum assumption $\mathsf{W}[1] \neq \mathsf{FPT}$ remains elusive. Toward this goal, studying some consequences of the PIH and proving them under $\mathsf{W}[1] \neq \mathsf{FPT}$ might provide new insights and valuable lessons.

Recently, Guruswami, Ren, and Sandeep [13] proved a parameterized version of the Baby PCP Theorem, appropriately coined as the *Baby PIH*, under $\mathsf{W}[1] \neq \mathsf{FPT}$. The Baby PIH states that for any constant $r > 1$, no FPT algorithm can distinguish a satisfiable 2CSP instance from one with no satisfying multi-assignment which assigns each variable no more than r values. Just like the relationship between the PCP theorem and the Baby PCP theorem, the Baby PIH is a direct consequence of the PIH. As a next step, a further complexity assumption is suggested, i.e., the *Average Baby PIH* [13, Conjecture 3], which seems to be sandwiched between the PIH and the Baby PIH. It postulates the W[1]-hardness of the problem known as Avg-r-Gap-2CSP (see Definition 8) which asks for distinguishing a satisfiable 2CSP instance from one without satisfying multi-assignment which assigns each variable no more than r values *on average*. The authors of [13] also demonstrated the difference between the Baby PIH and the Average Baby PIH. In fact, for all $r > 1$ and $\delta > 0$, they constructed 2CSP instances with variable set X that cannot be satisfied by any multi-assignment assigning each variable in X no more than r values, but can be satisfied by a multi-assignment that assigns in total $(1+\delta)|X|$ values to all the variables in X, that is, every variable is assigned $1+\delta$ values on average. Compared to proving W[1]-hardness[1] of the PIH, it is apparently easier to show the W[1]-hardness of the Average Baby PIH, and

[1] Strictly speaking, the PIH is not a computational problem and we cannot directly define its hardness. The formal statement should be "proving W[1]-hardness of the problem described in the PIH", and we use "W[1]-hardness for the PIH" for short in the introduction.

studying the Average Baby PIH might bring us further closer to a proof of the W[1]-hardness for the PIH. Moreover, the Average Baby PIH is already sufficient for proving some non-trivial inapproximability results such as k-ExactCover [13].

1.1 Main Results

Let $\Pi = (X, \Sigma, \Phi)$ be a 2CSP instance with a set X of variables, an alphabet Σ, and a set Φ of constraints. A multi-assignment $\hat{\sigma} : X \to 2^{\Sigma}$ relaxes the standard notion of assignments by assigning each variable $x \in X$ a *set* of values in Σ, i.e., $\hat{\sigma}(x) \subseteq \Sigma$. Thereby, Π is said to be *satisfied* by $\hat{\sigma}$ if for every constraint $\varphi \in \Phi$, one can pick for each variable x of φ a value from the set $\hat{\sigma}(x)$ assigned to this variable to satisfy φ. We say that $\hat{\sigma}$ assigns $\sum_{x \in X} |\hat{\sigma}(x)|$ values to X in total, or equivalently, each variable in X is assigned $\frac{\sum_{x \in X} |\hat{\sigma}(x)|}{|X|}$ values *on average* (see Definition 7). Our main result is the following theorem stating that the Average Baby PIH holds under $\mathsf{W}[1] \neq \mathsf{FPT}$.

▶ **Theorem 1** (Informal, see Theorem 16)**.** *Assume* $\mathsf{W}[1] \neq \mathsf{FPT}$*. Then for any constant* $r > 1$*, given a 2CSP instance* $\Pi = (X, \Sigma, \Phi)$ *parameterized by* $|X|$*, no* FPT *time algorithm can distinguish between:*
- Π *is satisfiable.*
- *Any multi-assignment assigning no more than* $r|X|$ *values to* X *does not satisfy* Π*.*

Clearly any standard assignment $\sigma : X \to \Sigma$ can be identified with a multi-assignment $\hat{\sigma}$ that assigns each variable $x \in X$ a set of a single value, i.e., $\hat{\sigma}(x) = \{\sigma(x)\}$. Hence, there is a constant r gap in Theorem 1 between YES and NO instances on the average number of values assigned to each variable, which gives us the aforementioned Avg-r-Gap-2CSP problem. On the other hand, the constant gap for the PIH is on the fraction of constraints that can be satisfied by an assignment. That is, a YES instance is a 2CSP instance whose *all* constraints can be satisfied by an assignment, while any assignment can only satisfy at most *a constant fraction* of constraints in a NO instance. So, perhaps surprisingly, the difference between the Average Baby PIH and the PIH can be pinpointed within the Avg-r-Gap-2CSP problem precisely in terms of whether a given instance contains a "dense" or "sparse" set of constraints.[2] More precisely:

▶ **Theorem 2.**
- *Under* $\mathsf{W}[1] \neq \mathsf{FPT}$*, the Average Baby PIH holds for* Avg-r-Gap-2CSP *instances* $\Pi = (X, \Sigma, \Phi)$ *with* $|\Phi| = \omega(|X|)$*.*
- *If the Average Baby PIH holds for* Avg-r-Gap-2CSP *instances* $\Pi = (X, \Sigma, \Phi)$ *with* $|\Phi| = O(|X|)$*, then the PIH holds as well.*

As a first application of the Average Baby PIH under $\mathsf{W}[1] \neq \mathsf{FPT}$, using a reduction in [13], we obtain the W[1]-hardness of constant approximating the k-ExactCover problem (see Definition 9), improving its previous approxmation lower bound under a stronger assumption, i.e., the Gap-ETH [23].

▶ **Theorem 3** (Theorem 27 restated)**.** *For any constant* $r > 1$*,* r*-approximating* k*-*ExactCover *is* W[1]*-hard.*

[2] This is pointed out by an anonymous reviewer.

We remark that the W[1]-hardness of approximating k-ExactCover has been a long-standing open problem in parameterized complexity. Although the W[1], W[2], ETH-hardness of approximating the k-SetCover problem has been established in [6, 15, 18, 20], as a special case of k-SetCover, the hardness of approximating k-ExactCover was only known under the PIH [23] prior to our work.

The second application is a close relationship between running time lower bounds for constant approximating the parameterized Nearest Codeword Problem γ-Gap-k-NCP$_p$ (see Definition 10) and Avg-r-Gap-2CSP. Its proof is a straightforward composition of two known reductions in [23, 13] .

▶ **Theorem 4.** *For any prime p, computable function g, and constant r, if no $f(k) \cdot n^{o(g(k))}$-time algorithm can decide Avg-r-Gap-2CSP with k variables for any computable function f, then r-Gap-k-NCP$_p$ cannot be solved in time $f(k) \cdot n^{o(g(k))}$ for any computable function f.*

Proof Sketch. Theorem 4 follows from the gap-preserving reduction [13] from Avg-r-Gap-2CSP to k-ExactCover (see also Section A), and the gap-preserving reduction from k-ExactCover to r-Gap-k-NCP$_p$ [23, Theorem 28]. Note that in both reductions the parameter k is preserved. ◀

1.2 Technical Overview: Local-to-Global Reduction For 2CSP

To prove Theorem 1, we show that the W[1]-hardness for the Baby PIH implies the W[1]-hardness for the Average Baby PIH. Here, the "W[1]-hardness for the Baby PIH" means that, for all $r > 1$, it's W[1]-hard to decide whether (i) a 2CSP instance is satisfiable, or (ii) it cannot be satisfied by any multi-assignment assigning each variable no more than r values. Thus there is a constant r gap between (i) and (ii) in the number of values assigned to *each* variable. Thereby, the gap is "local." As already mentioned, for the Average Baby PIH, the gap is on the average number of values, or equivalently, the *total* number of values assigned to *all* the variables. Hence, the gap is "global." Our reduction from the Baby PIH to the Average Baby PIH is thus said to be "local-to-global."

Technically, our reduction relies on a simple but crucial property of high-distance error-correcting codes (ECC) shown in [16, 20]. In particular, we need an ECC $\mathcal{C} \subseteq \mathbb{F}_p^m$ with relative distance $1-\delta$ such that any two distinct codewords $x, y \in \mathcal{C}$ can agree on at most δm entries. So if we have a set S of codewords and an ε-fraction of entries (denoted by $I \subseteq [m]$) with $\varepsilon \gg \delta$ such that for each $i \in I$, we can find distinct $x, y \in S$ that agree on their i-th entry, then the size of S must be large. The lower bound of $|S|$ is called the *collision number* of $\mathcal{C}$, denote by $\mathrm{Col}_\varepsilon(\mathcal{C})$. A simple counting argument in [20] shows that $\mathrm{Col}_\varepsilon(\mathcal{C}) \geq \sqrt{2\varepsilon/\delta}$.

Now given a 2CSP instance $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$, let $n = |\Pi_0| > |\Sigma_0|$, parameter $k = |X_0|$, and $k' = |\Phi_0|$. We fix some ECC with very large relative distance (e.g., a Reed-Solomon code) $\mathcal{C} : \mathbb{F}_p^k \to \mathbb{F}_p^{k''}$ for prime $n^{1/k} \leq p < 2n^{1/k}$ and $k'' = \Theta(k^5)$. Then we have $\mathrm{Col}_\varepsilon(\mathcal{C}) > 2rk$. We construct a new 2CSP instance $\Pi = (X_1 \dot\cup X_2, \Sigma, \Phi)$ as:

- $X_1 = \{u_1, \cdots, u_{k'}\}$, $X_2 = \{v_1, \cdots, v_{k''}\}$.
- Each u_j takes value from the (encoding of) satisfying assignments of $\varphi_j = (x_{i_1} x_{i_2}, C_j) \in \Phi_0$, i.e., $\{(\mathcal{C}(a_1), \mathcal{C}(a_2)) : (a_1, a_2) \in C_j\}$. Each v_ℓ takes value from $\mathbb{F}_p^k$.
- For each $u_j \in X_1$ and $v_\ell \in X_2$, there is a constraint that checks whether u_j's assigned value (w_1, w_2) and v_ℓ's assigned value s satisfy $w_1[\ell] = s[i_1]$ and $w_2[\ell] = s[i_2]$.

See Figure 1 for an illustration. Finally, we duplicate X_1 and X_2 each an appropriate number of times to make them equal in size, finishing our reduction.

In general, an assignment to X_1 should correspond to a satisfying multi-assignment to the original instance Π_0, and the value assigned to each $v_\ell \in X_2$ is a "guess" of the ℓ-th entry of the encoding of every x_i. It is easy to see that if the input instance Π_0 is satisfiable, then so

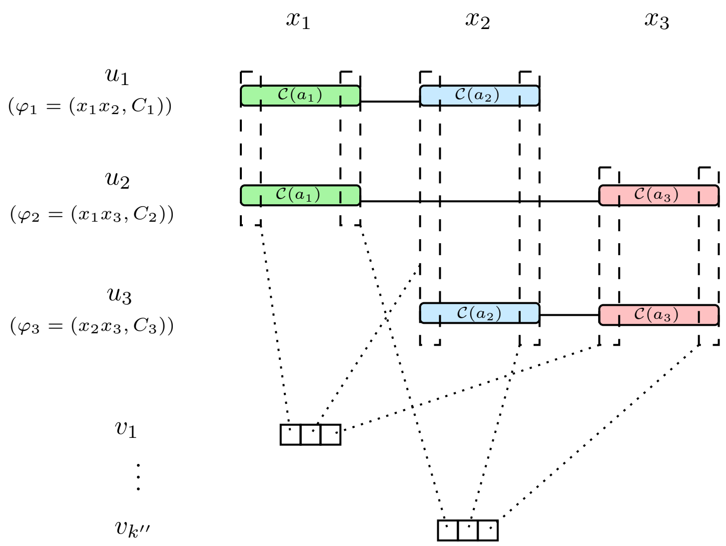

Figure 1 An illustration of our construction for an input instance $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$ with $|X_0| = |\Phi_0| = 3$.

does the new 2CSP instance Π, since for each $u_j \in X_1$ corresponding to $\varphi_j = (x_{i_1}x_{i_2}, C_j)$, we can simply assign it the value $(\mathcal{C}(\sigma(x_{i_1})), \mathcal{C}(\sigma(x_{i_2})))$, where σ is a satisfying assignment for Π_0. At the same time, the value assigned to each $v_\ell \in X_2$ is $(\mathcal{C}(\sigma(x_1))[\ell], \cdots, \mathcal{C}(\sigma(x_k))[\ell]) \in \mathbb{F}_p^k$.

For soundness, suppose Π_0 has no satisfying multi-assignment assigning at most r values to each variable, we need to argue that Π has no satisfying multi-assignment that assigns $r(1-\varepsilon)(|X_1| + |X_2|)/2$ values in total. To that end, we exploit the collision number of the code $\mathcal{C}$. Fix any satisfying multi-assignment $\hat{\sigma}$ to Π. Recall that each variable $u_j \in X_1$ is assigned some value from a satisfying partial assignment to $\varphi_j \in \Phi_0$. Then, $\hat{\sigma}(X_1)$ naturally gives a satisfying multi-assignment to Π_0, which implies that there exists a variable $x_i \in X_0$ such that more than r different values are assigned by $\hat{\sigma}$ to u_j, for which the corresponding constraint φ_j contains x_i.

Now we have two possible cases for $\hat{\sigma}$. In the first case, for $(1-\varepsilon)$-fraction of variables in X_2, the multi-assignment $\hat{\sigma}$ assigns each of them more than r values. We are done, since this implies that the total number of assigned values by $\hat{\sigma}$ is more than $(1-\varepsilon)r|X_2|$. For the second case, there exists an ε-fraction of variables in X_2, each of which is assigned by $\hat{\sigma}$ at most r values. Given such a variable v_ℓ, we have more than r different *codewords* assigned to X_1 in x_i's position which has at most r possible values in the ℓ-th entry. This entails the existence of two different codewords with the same ℓ-th entry. Since there are εm such entries, the assignment to X_1 must contain at least $\mathrm{Col}_\varepsilon(\mathcal{C})$ different codewords. Each assignment to $u_j \in X_1$ contributes two codewords, so the total number of assigned values by $\hat{\sigma}$ is at least $\mathrm{Col}_\varepsilon(\mathcal{C})/2 > rk$. In summary, both cases guarantee a constant gap on either X_1 or X_2, showing that any satisfying multi-assignment to Π must assign $r(1-\varepsilon)(|X_1| + |X_2|)/2$ values to $X_1 \cup X_2$ in total.

More details are referred to Section 3.

1.3 Discussions

For minimization problems, the technique of constructing two parts of variables and arguing that at least one part has a large gap seems quite general, as exhibited by the previous works showing the W[1]-hardness of approximating k-SetCover [6, 20] and k-NCP [17].

The use of the collision number of error-correcting codes in the context of parameterized inapproximability was first introduced in [16] and further developed in [20, 17]. We ask whether these techniques can be unified.

▶ **Question 1.** *Is there a general framework for proving parameterized inapproximability of minimization problems?*

We also suggest two new variants of the Average Baby PIH, which might serve as a next step towards proving the PIH under $\mathsf{W}[1] \neq \mathsf{FPT}$. On closer inspection of our construction, particularly Case 1 in the proof of Lemma 20, it only guarantees the total number of values assigned to X_1 (i.e., $\Sigma_{x \in X_1} |\hat{\sigma}(x)|$) is large. This could happen if an $o(1)$-fraction of x's in X_1 is assigned super-constant number of values. We ask if a larger gap can be achieved. So the first one asks whether the Average Baby PIH can be strengthened by requiring a constant fraction of variables to be assigned multiple values.

▶ **Question 2.** *Under* $\mathsf{W}[1] \neq \mathsf{FPT}$*, can we prove that for any constant* $r > 2$*, there exists a constant* $c > 0$ *such that no* FPT *algorithm can decide whether a 2CSP instance is satisfiable, or any satisfying multi-assignment must have at least a c-fraction of variables assigned r values?*

We remark that the case of $r = 2$ follows from the inapproximability of k-Clique [19, 14, 5].

The second variant is already contained in Theorem 2, thus equivalent to PIH.

▶ **Question 3.** *Under* $\mathsf{W}[1] \neq \mathsf{FPT}$*, can we prove that the Average PIH holds even for* Avg-r-Gap-2CSP *instances with the number of constraints being linear in the number of variables?*

It is also interesting to consider whether the inapproximability factor in the Average Baby PIH can be improved to $\omega(1)$, since this would directly lead to better lower bounds for approximating k-ExactCover. The current obstacle is that, although the running time of our reduction does not depend on the approximation factor, our reduction relies on the gap created in the Baby PIH [13]. In order to achieve an r-gap in the Baby PIH, the reduction in [13] runs in time $\Omega(n^{(2r)^r})$, consequently, the existing FPT reduction cannot create a super-constant $r = \omega(1)$ gap .

▶ **Question 4.** *Under* $\mathsf{W}[1] \neq \mathsf{FPT}$*, can we prove that the Average PIH (Theorem 1) holds for inapproximability factor* $r = \omega(1)$*, hence giving better inapproximability result for* k-ExactCover*?*

1.4 Organization

In Section 2, we introduce the main computational problems and complexity assumptions studied in this paper. As the central contribution, Section 3 explains our reduction from the Baby PIH to the Average Baby PIH. This, in fact, establishes the Average Baby PIH under $\mathsf{W}[1] \neq \mathsf{FPT}$. In Section A, we present a reduction from Avg-r-Gap-2CSP to the constant approximation of k-ExactCover, which slightly differs from the construction in [13].

2 Preliminaries

For a positive integer n, we use $[n]$ to denote the set $\{1, 2, \cdots, n\}$. $S_1 \dot\cup \cdots \dot\cup S_k$ is the *disjoint union* of sets $S_1, \ldots, S_k$, where we tacitly assume that $S_1, \ldots, S_k$ are pairwise disjoint. We use log (without subscript) to denote the logarithm number with base 2. For any prime number p, we write $\mathbb{F}_p$ for the (unique) finite field of size p. The asymptotic notations, i.e., O, Ω, ω, and Θ, are used following the general convention. The reader is assumed to be familiar with basic notions in parameterized complexity theory, in particular FPT and W[1]. Otherwise, the standard references are, e.g., [10, 9, 7].

2.1 Problems

▶ **Definition 5** (Parameterized 2CSP). *A* 2CSP *instance is defined as a triple* $\Pi = (X, \Sigma, \Phi)$ *where:*

- X *is a set of variable.*
- $\Sigma = \dot{\bigcup}_{x \in X} \Sigma_x$*, where each* Σ_x *contains values that the variable* $x \in X$ *can be assigned. Often, we assume that there exists an* $n \in \mathbb{N}$ *such that* $|\Sigma_x| \leq n$ *for all* $x \in X$*.*
- $\Phi = \{\varphi_1, \cdots, \varphi_{k'}\}$*, where each* $\varphi_j = (x_{i_1} x_{i_2}, C_j)$ *for some* $x_{i_1}, x_{i_2} \in X$*, and* C_i *is a subset of* $\Sigma_{x_{i_1}} \times \Sigma_{x_{i_2}}$*.*

The problem is to decide whether there exists an assignment $\sigma : X \to \Sigma$ *that satisfies:*

- *For all* $x \in X$*,* $\sigma(x) \in \Sigma_x$*.*
- *For all* $\varphi_j = (x_{i_1} x_{i_2}, C_j) \in \Phi$*,* $(\sigma(x_{i_1}), \sigma(x_{i_2})) \in C_j$*.*

The parameter for this problem is $k = |X|$*, the number of variables. Each pair of variables has at most one constraint, so* $|\Phi| \leq \binom{k}{2}$*. Without loss of generality, each variable is related to some constraint in* Φ*. The size of instance* Π *is defined as* $|\Pi| = |\Sigma| + |\Phi|$*, where the size of each* φ_j *is defined as* $|\varphi_j| = |C_j|$*.*

The approximation of parameterized 2CSP refers to the following problem.

▶ **Definition 6** (ε-Gap-2CSP). *Given a 2CSP instance* $\Pi = (X, \Sigma, \Phi)$ *with parameter* $k = |X|$*, we want to distinguish between:*

- Π *is satisfiable;*
- *any assignment can satisfy at most an* ε*-fraction of constraints in* Φ*.*

As already mentioned, the notion of multi-assignment extends the usual assignment in such a way that each variable can be assigned multiple values.

▶ **Definition 7** (Multi-assignment). *A multi-assignment of a 2CSP instance* $\Pi = (X, \Sigma, \Phi)$ *is a function* $\hat{\sigma} : X \to 2^{\Sigma}$*,* [3] *such that for all* $x \in X$ *we have* $\hat{\sigma}(x) \subseteq \Sigma_x$*. Furthermore, we say that* $\hat{\sigma}$ *satisfies* Π *if:*

- *For all* $\varphi_j = (x_{i_1} x_{i_2}, C_j) \in \Phi$*, there exist* $c_1 \in \hat{\sigma}(x_{i_1})$ *and* $c_2 \in \hat{\sigma}(x_{i_2})$ *with* $(c_1, c_2) \in C_j$*.*

The individual size *of* $\hat{\sigma}$ *is defined as* $\max_{x \in X} |\hat{\sigma}(x)|$*, and the* total size *of* $\hat{\sigma}$ *is* $\sum_{x \in X} |\hat{\sigma}(x)|$*.*

Let $r \geq 1$. We say that a 2CSP instance $\Pi = (X, \Sigma, \Phi)$ is *r-list satisfiable* if there exists a multi-assignment $\hat{\sigma}$ with individual size no more than r which satisfies Π, and Π is *r-average list satisfiable* if there exists a multi-assignment $\hat{\sigma}$ with total size no more than $r|X|$ which satisfies Π.

▶ **Definition 8** (Avg-r-Gap-2CSP). *Given a* 2CSP *instance* Π*, the goal is to distinguish between the following two cases:*

- Π *is satisfiable.*
- Π *is not* r*-average list satisfiable.*

[3] Here we use 2^{Σ} to denote the power set of Σ.

We also consider the k-ExactCover problem (aka, the k-UniqueSetCover problem) and the k-NCP problem (aka, k-MLD, for the parameterized Maximum Likelihood Decoding problem) as defined below.

▶ **Definition 9** (k-ExactCover). *Given a set U (which we call universe) and a collection of U's subsets $\mathcal{S}$, the goal is to distinguish between the following two cases:*
- *there exist at most k disjoint sets in $\mathcal{S}$ that form a partition of U,*
- *or U is not the union of any k sets in $\mathcal{S}$.*

▶ **Definition 10** (k-NCP). *For prime p, integer $d > 0$, given a (multi-)set V of vectors in $\mathbb{F}_p^d$, and a target vector $\vec{t} \in \mathbb{F}_p^d$, the k-NCP$_p$ problem asks for distinguishing between:*
- *the Hamming distance between $\vec{t}$ and the vector space spanned by V is at most k,*
- *or the Hamming distance between $\vec{t}$ and the vector space spanned by V is at least $k+1$.*

2.2 Hypotheses

▶ **Hypothesis 11** (PIH [21]). *For every constant $0 < \varepsilon < 1$, there is no* FPT *algorithm solving ε-Gap-2CSP.*

The Baby PIH, a hypothesis implied by PIH, asserts the hardness of approximating individual size of a satisfying multi-assignment. Formally,

▶ **Hypothesis 12** (Baby PIH [13]). *For any constant $r > 0$, no* FPT *algorithm can on input a 2CSP instance, distinguish whether it is satisfiable, or cannot be satisfied by any multi-assignment with* individual *size at most r.*

We emphasize that the Baby PIH is a hardness hypothesis with a **local** condition, i.e., the individual size of satisfying assignments. It is shown that the standard assumption W[1] $\neq$ FPT implies the Baby PIH:

▶ **Theorem 13** ([13]). *The Baby PIH holds under* W[1] $\neq$ FPT*.*

In contrast, the Average Baby PIH is defined on a **global** condition concerning the total size of satisfying assignments. The precise statement of this complexity assumption contains a technical property on the "shape" of the constraints in a 2CSP instance.

▶ **Definition 14** (Rectangular relation). *A 2CSP instance $\Pi = (X, \Sigma, \Phi)$ is said to have* rectangular relations *if for each $\varphi_j = (x_{i_1}x_{i_2}, C_j) \in \Phi$, there exist a set Q_j and mappings $\pi_j, \rho_j : \Sigma \to Q_j$, such that $(a, b) \in C_j$ iff $\pi_j(a) = \rho_j(b)$. We call Q_j the* underlying set *of φ_j.*

Some explanation for "rectangular" might be in order. Recall that a subset $S \subseteq \Sigma^2$ is a (combinatorial) *rectangle* if and only if there exist $A, B \subseteq \Sigma$ such that $S = A \times B$. It is easy to verify that $R \subseteq \Sigma^2$ is rectangular if and only if R is the union of a set of pairwise disjoint rectangles.

▶ **Hypothesis 15** (Average Baby PIH). *For any constant $r > 0$, there exists no* FPT *algorithm solving the Avg-r-Gap-2CSP problem, even when the instance contains only rectangular relations.*

3 Average Baby PIH from Baby PIH

In this section, we show that the Average Baby PIH even for instances with only rectangular relations, is implied by the Baby PIH.

3.1 Proofs of Main Results

We employ a local-to-global reduction developed in [17] to amplify the local gap for one variable (Theorem 13) into a global gap for all variables, thus proving the Average Baby PIH from the Baby PIH.

▶ **Theorem 16.** *Under* W[1] $\neq$ FPT*, for any constant* $r > 0$*, no* FPT *algorithm can distinguish a given* 2*CSP instance with rectangular relation is satisfiable, or cannot be satisfied by any multi-assignment with total size no more than* r*.*

To show Theorem 16, we first introduce some tools from coding theory. The *collision number* of an error-correcting code characterizes the number of codewords needed to find "collision" on a constant fraction of coordinates. We use the definition in [17]:

▶ **Definition 17** (ε-Collision Number)**.** *Let* $m \geq 1$ *and* $x, y \in \Sigma^m$ *with* $x \neq y$*. For every* $i \in [m]$ *we say that* x *and* y collide *on position* i *if* $x[i] = y[i]$*. Furthermore, a subset* $S \subseteq \Sigma^m$ collides *on position* i *if there exist distinct* $x, y \in S$ *with* $x[i] = y[i]$*. We define the* collision set *of* S *as*

$$\mathrm{ColSet}(S) = \big\{ i \in [m] \;\big|\; S \text{ collides on position } i \big\}.$$

Observe that if $|S| \leq 1$*, then* $\mathrm{ColSet}(S) = \varnothing$*.*

Now for every $C \subseteq \Sigma^m$ *and* $0 < \varepsilon < 1$ *the* ε-collision number of C*, denoted by* $\mathrm{Col}_\varepsilon(C)$*, is the* maximum $s \leq |C| + 1$ *such that for* all $S \in \binom{C}{s-1}$ *we have*

$$\big|\mathrm{ColSet}(S)\big| \leq \varepsilon m.$$

For Reed-Solomon codes, we have the following lower bounds on their collision number.

▶ **Theorem 18** (Theorem 10 in [20], see also [16])**.** *For any* $0 < \varepsilon < 1$*, any Reed-Solomon code* $\mathcal{C}^{RS} : \mathbb{F}_p^k \to \mathbb{F}_p^m$ *with sufficiently large* $k < m \leq p$*,* $\mathrm{Col}_\varepsilon(\mathcal{C}^{RS}) \geq \sqrt{\frac{2\varepsilon m}{k}}$*.*

Our proof of Theorem 16 consists of two reductions. The first one (Lemma 20) reduces 2CSP instances from the Baby PIH, i.e., with a "local" gap as explained in Section 1.2, to a new instance whose constraints are between two disjoint groups of variables. The new instance has a different "global" gap on each group of variables. As the sizes of the two groups might not be balanced, we do not necessarily have a "global" gap on all the variables. But this is easily remedied by the second reduction (Lemma 22) which makes an appropriate number of copies of the two groups.

▶ **Definition 19** ((r, s)-Average Multi-Assignment)**.** *Let* $\Pi = (X, \Sigma, \Phi)$ *be a* bipartite 2*CSP instance, in particular* $X = X_1 \dot\cup X_2$ *and every* $\varphi = (x_1 x_2, C) \in \Phi$ *has* $x_1 \in X_1$ *and* $x_2 \in X_2$*. Then for* $r_1, r_2 \geq 1$ *an* (r_1, r_2)-average multi-assignment *of* Π *is a multi-assignment* $\hat\sigma : X \to 2^\Sigma$ *such that*

$$\frac{\sum_{x \in X_1} |\hat\sigma(x)|}{|X_1|} \leq r_1 \quad \text{and} \quad \frac{\sum_{x \in X_2} |\hat\sigma(x)|}{|X_2|} \leq r_2.$$

That is, the total size of $\hat\sigma$ *restricted to* X_1 *is at most* $r_1|X_1|$*, and the total size of* $\hat\sigma$ *restricted to* X_2 *is at most* $r_2|X_2|$ *(cf. Definition 7). We say* Π *is* (r_1, r_2)-average list satisfiable *if there is an* (r_1, r_2)*-average multi-assignment which satisfies* Π*.*

▶ **Lemma 20.** *There is an algorithm* $\mathcal{A}$ *which on input a* 2*CSP instance* $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$*,* $\varepsilon > 0$*, and* $r \geq 1$ *computes a bipartite* 2*CSP instance* $\Pi = (X_1 \dot\cup X_2, \Sigma, \Phi)$ *with the following properties.*

Completeness. *If* Π_0 *is satisfiable, then so is* Π*,*

Soundness. *For every $r \geq 1$ if Π_0 is not $2r$-list satisfiable, then Π is not (r_1, r_2)-average list satisfiable for every $r_1, r_2 \in \mathbb{N}$ with*

$$r_1 + r_2 \leq 2(1-\varepsilon)r.$$

Rectangularity. *All constraints in Φ are rectangular.*

In addition, there exists a computable function f upper bounding the running time of $\mathcal{A}$ as

$$f(|X_0| + |\Phi_0| + 1/\varepsilon + r)|\Sigma_0|^{O(1)}. \tag{1}$$

And the number of variables $|X_1| + |X_2|$ and the number of constraints $|\Pi|$ in Π can also be upper bounded by $f(|X_0| + |\Phi_0| + 1/\varepsilon + r)$.

Proof. For the given 2CSP instance $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$ we let

$$k = |X_0| \quad \text{and} \quad k' = |\Phi_0|.$$

Thereby we fix some enumerations of the variables in X_0 and the constraints in Φ_0 as

$$X_0 = \{x_1, \ldots, x_k\} \quad \text{and} \quad \Phi_0 = \{\varphi_1, \ldots, \varphi_{k'}\}.$$

Let $\mathcal{C} : \mathbb{F}_p^k \to \mathbb{F}_p^{k''}$ be a Reed-Solomon code with

$$2|\Sigma_0|^{1/k} > p \geq |\Sigma_0|^{1/k} \quad \text{and} \quad k'' = \left\lfloor \frac{8(1-\varepsilon)^2 r^2}{\varepsilon} k(k')^2 \right\rfloor + 1.$$

Clearly $|\Sigma_0| \leq p^k$, and therefore we can assume without loss of generality

$$\Sigma_0 \subseteq \mathbb{F}_p^k.$$

Moreover, we only consider the case that

$$k'' \leq p \left(= |\mathbb{F}_p| \right) < 2|\Sigma_0|^{1/k},$$

i.e., Σ_0 is sufficiently larger than k and k'.[4] Hence we can invoke Theorem 18 on $\Sigma \leftarrow \mathbb{F}_p$, $k \leftarrow k$, $m \leftarrow k''$, and $\varepsilon \leftarrow \varepsilon$ to obtain

$$\mathrm{Col}_\varepsilon(\mathcal{C}(\mathbb{F}_p^k)) \geq \sqrt{\frac{2\varepsilon k''}{k}} > 4(1-\varepsilon)rk', \tag{2}$$

where the second inequality is by our choice of k''.

Now the algorithm $\mathcal{A}$ constructs the following bipartite 2CSP instance $\Pi = (X, \Sigma, \Phi)$.

Variables. $X = X_1 \dot\cup X_2$ with

$$X_1 = \{u_1, \ldots, u_{k'}\} \quad \text{and} \quad X_2 = \{v_1, \ldots, v_{k''}\}.$$

Alphabets. $\Sigma = \bigcup_{u \in X_1} \Sigma_u \cup \bigcup_{v \in X_2} \Sigma_v$ where:

- For every $j \in [k']$ the alphabet of the variable $u_j \in X_1$ is

$$\Sigma_{u_j} = \left\{ \left(\mathcal{C}(a_1), \mathcal{C}(a_2)\right) \;\middle|\; \varphi_j = (x_{i_1} x_{i_2}, C) \text{ and } (a_1, a_2) \in C \right\} \subseteq \left(\mathcal{C}(\mathbb{F}_p^k)\right)^2 \subseteq \left(\mathbb{F}_p^{k''}\right)^2. \tag{3}$$

  That is, Σ_{u_j} contains all the (partial) satisfying assignments of φ_j encoded by $\mathcal{C} : \mathbb{F}_p^k \to \mathbb{F}_p^{k''}$ as pairs of vectors in $\mathbb{F}_p^{k''}$. (Recall $\Sigma_0 \subseteq \mathbb{F}_p^k$.)
- For every $\ell \in [k'']$ we have $\Sigma_{v_\ell} = \mathbb{F}_p^k$. Since $p < 2|\Sigma_0|^{1/k}$, we have $|\Sigma_{v_\ell}| \leq 2^k |\Sigma_0|$.

[4] Otherwise, the original instance Π_0 can be solved in time of the form (1), and we can then output some predetermined Π depending on whether Π_0 is satisfiable.

Constraints. Let $j \in [k']$ and $\varphi_j = (x_{i_1} x_{i_2}, C)$. Then for every $\ell \in [k'']$ we have a constraint between the variable $u_j \in X_1$ and $v_\ell \in X_2$ which checks whether u_j is assigned to $(w_1, w_2) \in \left(\mathbb{F}_p^{k''}\right)^2$ and v_ℓ to $s \in \mathbb{F}_q^{\ell}$ such that

$$w_1[\ell] = s[i_1] \quad \text{and} \quad w_2[\ell] = s[i_2]. \tag{4}$$

Consequently (4) implies that the constraint is rectangular.[5] Moreover, the number of constraints in Π is

$$k'k'' = k' \left\lfloor \frac{2(1-\varepsilon)^2 r^2}{\varepsilon} k(k')^2 \right\rfloor + k'.$$

The completeness of our reduction is straightforward. So we turn to the soundness. In particular, we assume that the given 2CSP instance Π_0 is *not* $2r$*-list satisfiable*. Furthermore, let $\hat{\sigma} : X \to 2^{\Sigma}$ be a satisfying multi-assignment for Π. We need to show that, for any $r_1, r_2 \in \mathbb{N}$ if there is a satisfying (r_1, r_2)-average multi-assignment $\hat{\sigma}$, then

$$r_1 + r_2 > 2(1-\varepsilon)r. \tag{5}$$

To that end, let

$$\mathsf{Word}_{\hat{\sigma}} = \bigcup_{u_j \in X_1} \bigcup_{(w_1,w_2) \in \hat{\sigma}(u_j)} \{w_1, w_2\} \subseteq \mathbb{F}_p^{k''}. \tag{6}$$

That is, $\mathsf{Word}_{\hat{\sigma}}$ is the set of all codewords in $\mathbb{F}_p^{k''}$ that $\hat{\sigma}$ uses for the variables in X_1.

▷ Claim 21. Let $\ell \in [k'']$ with $|\hat{\sigma}(v_\ell)| \le 2r$. Then $\mathsf{Word}_{\hat{\sigma}}$ collides on position ℓ.

Proof of Claim 21. Let $\ell \in [k'']$ be fixed with $|\hat{\sigma}(v_\ell)| \le 2r$.

Consider an arbitrary constraint $\varphi_j = (x_{i_1} x_{i_2}, C) \in \Phi_0$ $\big($i.e., $j \in [k']\big)$. Since $\hat{\sigma}$ is a satisfying multi-assignment for Π, there exist

$$(w_1, w_2) \in \hat{\sigma}(u_j) \subseteq \Sigma_{u_j} \subseteq \left(\mathbb{F}_p^{k''}\right)^2 \quad \text{and} \quad s \in \hat{\sigma}(v_\ell) \subseteq \mathbb{F}_p^{k}$$

such that $u_j = (w_1, w_2)$ and $v_\ell = s$ satisfy the constraint between u_j and v_ℓ in Π. By $(w_1, w_2) \in \Sigma_{u_j}$ and (3) there are $a_1, a_2 \in \Sigma_0$ with $w_1 = \mathcal{C}(a_1)$ and $w_2 = \mathcal{C}(a_2)$ such that

$$x_{i_1} = a_1 \text{ and } x_{i_2} = a_2 \text{ satisfy } \varphi_j. \tag{7}$$

Then we say that a_1 is $(\hat{\sigma}, \varphi_j)$*-suitable* for x_{i_1} *with respect to* s, and similarly a_2 is $(\hat{\sigma}, \varphi_j)$*-suitable* for x_{i_2} *with respect to* s.

In addition, by (4)

$$\mathcal{C}(a_1)[\ell] = s[i_1] \quad \text{and} \quad \mathcal{C}(a_2)[\ell] = s[i_2]. \tag{8}$$

Now we define a *multi-assignment* $\hat{\sigma}_0 : X_0 \to 2^{\Sigma_0}$ for the original instance $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$ as follows. For every $x \in X_0$ let

$$\hat{\sigma}_0(x) = \bigcup_{s \in \hat{\sigma}(v_\ell)} \big\{a \in \Sigma_0 \bigm| j \in [k'] \text{ and } a \text{ is } (\hat{\sigma}, \varphi_j)\text{-suitable for } x \text{ with respect to } s\big\}. \tag{9}$$

(Recall that we have fixed an $\ell \in [k'']$ and hence $\hat{\sigma}(v_\ell)$.) Since every variable x must appear in at least one constraint $\varphi_j \in \Phi_0$ (cf. Definition 5), it is easy to see that $\hat{\sigma}_0$ is a satisfying multi-assignment for Π_0 by (7).

[5] To see this, we take $\pi(u_j) = \pi(w_1, w_2) = (w_1[\ell], w_2[\ell])$ and $\rho(v_\ell) = (v_\ell[i_1], v_\ell[i_2])$. Then equation (4) is precisely $\pi(u_j) = \rho(v_\ell)$ as in Definition 14.

As Π_0 is *not* $2r$-list satisfiable, there exists an $x_{i^*} \in X_0$ (i.e., $i^* \in [k]$) with

$$\left|\hat{\sigma}_0(x_{i^*})\right| \geq 2r+1.$$

We have assumed that

$$|\hat{\sigma}(v_\ell)| \leq 2r,$$

so by (9) there is an $s \in \hat{\sigma}(v_\ell)$ such that

$$\left|\left\{a \in \Sigma_0 \;\middle|\; j \in [k'] \text{ and } a \text{ is } (\hat{\sigma}, \varphi_j)\text{-suitable for } x_{i^*} \text{ with respect to } s\right\}\right| \geq 2.$$

Hence there are $a_1, a_2 \in \Sigma_0$ with $a_1 \neq a_2$ and $j_1, j_2 \in [k']$ such that
- a_1 is $(\hat{\sigma}, \varphi_{j_1})$-suitable for x_{i^*} with respect to s,
- and a_2 is $(\hat{\sigma}, \varphi_{j_2})$-suitable for x_{i^*} with respect to s.

Then (8) implies that

$$\mathcal{C}(a_1)[\ell] = s[i^*] = \mathcal{C}(a_2)[\ell].$$

In other words, $\mathcal{C}(a_1)$ and $\mathcal{C}(a_2)$ collide on position ℓ. Clearly $\mathcal{C}(a_1), \mathcal{C}(a_2) \in \mathsf{Word}_{\hat{\sigma}}$, so this finishes the proof of the claim. ◁

Let

$$r_1 = \frac{\sum_{x \in X_1} |\hat{\sigma}(x)|}{|X_1|} = \frac{\sum_{j \in [k']} \left|\hat{\sigma}(u_j)\right|}{k'} \quad \text{and} \quad r_2 = \frac{\sum_{x \in X_2} |\hat{\sigma}(x)|}{|X_2|} = \frac{\sum_{\ell \in [k'']} \left|\hat{\sigma}(v_\ell)\right|}{k''}.$$

Now we distinguish two cases.

1. There are more than ε fraction of $\ell \in [k'']$ such that $|\hat{\sigma}(v_\ell)| \leq 2r$, then Claim 21 implies that $\mathsf{Word}_{\hat{\sigma}}$ collides on more than ε fraction of positions $\ell \in [k'']$. Recall (2), i.e.,

$$\mathrm{Col}_\varepsilon(\mathcal{C}(\mathbb{F}_p^k)) \geq \sqrt{\frac{2\varepsilon k''}{k}} > 4(1-\varepsilon) r k'.$$

Hence,

$$\left|\mathsf{Word}_{\hat{\sigma}}\right| \geq \mathrm{Col}_\varepsilon(\mathcal{C}(\mathbb{F}_p^k)) > 4(1-\varepsilon) r k'.$$

By the definition (6) of $\mathsf{Word}_{\hat{\sigma}}$ we deduce

$$\begin{aligned}\left|\mathsf{Word}_{\hat{\sigma}}\right| &= \left|\bigcup_{u_j \in X_1} \bigcup_{(w_1,w_2) \in \hat{\sigma}(u_i)} \{w_1, w_2\}\right| \\ &\leq \sum_{u_j \in X_1} \left|\bigcup_{(w_1,w_2) \in \hat{\sigma}(u_i)} \{w_1, w_2\}\right| \leq \sum_{u_j \in X_1} 2\left|\hat{\sigma}(u_i)\right|\end{aligned}$$

It follows that

$$r_1 = \frac{\sum_{j \in [k']} \left|\hat{\sigma}(u_j)\right|}{k'} > \frac{4(1-\varepsilon) r k'}{2k'} = 2(1-\varepsilon) r.$$

2. There are at most ε fraction of $\ell \in [k'']$ with $|\hat{\sigma}(v_\ell)| \leq 2r$. Or equivalently, there are at least $(1-\varepsilon)$ fraction of $\ell \in [k'']$ with $|\hat{\sigma}(v_\ell)| \geq 2r+1$. Then

$$r_2 = \frac{\sum_{\ell \in [k'']} \left|\hat{\sigma}(v_\ell)\right|}{k''} \geq \frac{(1-\varepsilon) k''(2r+1) + \varepsilon k''}{k''} > 2(1-\varepsilon) r.$$

So both cases lead to (5) as desired. ◀

With some proper replication, the unbalanced (r_1, r_2)-gap can be turned into a balanced one, and yield the desired r-average list unsatisfiability.

▶ **Lemma 22.** *For any bipartite 2CSP instance $\Pi = (X_1 \dot\cup X_2, \Sigma, \Phi)$ and $r > 1$ we can compute in polynomial time a 2CSP instance $\Pi' = (X', \Sigma', \Phi')$ with*

$$|X| = 2|X_1||X_2|$$

such that

Completeness. *If Π is satisfiable, then so is Π',*

Soundness. *Let $r \geq 1$. If Π is not (r_1, r_2)-average list satisfiable for* every $r_1, r_2 \geq 1$ *with $r_1 + r_2 \leq 2r$, then Π' is not r-average list satisfiable. Or equivalently, if Π' is r-average list satisfiable, then for* some $r_1, r_2 \in \mathbb{N}$ *with $r_1 + r_2 \leq 2r$ the bipartite Π is (r_1, r_2)-average list satisfiable.*

Furthermore, if Π is rectangular, then so is Π'.

Proof. Let

$$k_1 = |X_1| \quad \text{and} \quad k_2 = |X_2|.$$

The desired $\Pi' = (X', \Sigma', \Phi')$ is constructed as below.

Variables. X' consists of k_2 copies of X_1 and k_1 copies of X_2, i.e., $X' = X_1' \dot\cup X_2'$ where

$$X_1' = \left\{x^{(i)} \;\middle|\; x \in X_1 \text{ and } i \in [k_2]\right\} \quad \text{and} \quad X_2' = \left\{x^{(i)} \;\middle|\; x \in X_2 \text{ and } i \in [k_1]\right\}.$$

Note, $|X_1'| = |X_2'| = k_1 k_2$, therefore Π' contains $2k_1k_2$ many variables.

Alphabets. $\Sigma' = \bigcup_{x \in X'} \Sigma'_x$ where:

- For every $x \in X_1$ and $i \in [k_2]$, let $\Sigma'_{x^{(i)}} = \Sigma_x$. Recall that $\Sigma_x \subseteq \Sigma$ is the alphabet for the variable x in the original 2CSP instance Π.
- Similarly, for every $x \in X_2$ and $i \in [k_1]$ let $\Sigma'_{x^{(i)}} = \Sigma_x$.

Constraints. For every constraint $\varphi = (x_1 x_2, C) \in \Phi$ with $x_1 \in X_1$ and $x_2 \in X_2$, $i_1 \in [k_2]$, and $i_2 \in [k_1]$ we have a constraint

$$\varphi^{i_1, i_2} = \left(x_1^{(i_1)} x_2^{(i_2)}, C\right) \in \Phi'.$$

That is, φ^{i_1,i_2} is a copy of φ where the variable x_1 is replaced by its i_1-th copy $x_1^{(i_1)}$ and x_2 by its i_2-th copy $x_2^{(i_2)}$. It immediately implies that if Π is rectangular, then Π' is rectangular too.

Again the completeness is immediate. Towards the soundness, let $\hat\sigma' : X' \to 2^{\Sigma'}$ be a satisfying r-average multi-assignment for Π'. In particular,

$$r = \frac{\sum_{x \in X'} \left|\hat\sigma'(x)\right|}{|X'|} = \frac{\sum_{x \in X_1'} \left|\hat\sigma'(x)\right| + \sum_{x \in X_2'} \left|\hat\sigma'(x)\right|}{|X_1'| + |X_2'|} = \frac{\sum_{x \in X_1'} \left|\hat\sigma'(x)\right| + \sum_{x \in X_2'} \left|\hat\sigma'(x)\right|}{2k_1k_2}.$$

We set

$$r_1 = \frac{\sum_{x \in X_1'} \left|\hat\sigma'(x)\right|}{|X_1'|} = \frac{\sum_{x \in X_1'} \left|\hat\sigma'(x)\right|}{k_1k_2} \quad \text{and} \quad r_2 = \frac{\sum_{x \in X_2'} \left|\hat\sigma'(x)\right|}{|X_2'|} = \frac{\sum_{x \in X_2'} \left|\hat\sigma'(x)\right|}{k_1k_2} \tag{10}$$

It follows that

$$r_1 + r_2 = \frac{\sum_{x \in X_1'} \left|\hat\sigma'(x)\right| + \sum_{x \in X_2'} \left|\hat\sigma'(x)\right|}{k_1k_2} = 2r.$$

Note that

$$X_1' = \dot{\bigcup}_{i\in[k_2]} \{x^{(i)} \mid x \in X_1\}. \tag{11}$$

Therefore,

$$\begin{aligned} r_1k_1k_2 &= r_1|X_1'| && \big(\text{by } |X_1'| = k_1k_2\big) \\ &= \sum_{x\in X_1'} \big|\hat{\sigma}'(x)\big| && \big(\text{by (10)}\big) \\ &= \sum_{i\in[k_2]}\sum_{x\in X_1} \big|\hat{\sigma}'(x^{(i)})\big|. && \big(\text{by (11)}\big) \end{aligned}$$

Hence, there exists an $i_1 \in [k_2]$ such that

$$\sum_{x\in X_1} \big|\hat{\sigma}'(x^{(i_1)})\big| \le r_1k_1, \quad \text{or equivalently} \quad \frac{\sum_{x\in X_1} \big|\hat{\sigma}'(x^{(i_1)})\big|}{|X_1|} \le r_1$$

by $|X_1| = k_1$. Arguing similarly for X_2 we get an $i_2 \in [k_1]$ such that

$$\frac{\sum_{x\in X_2} \big|\hat{\sigma}'(x^{(i_2)})\big|}{|X_2|} \le r_2$$

Finally we define a multi-assignment $\hat{\sigma}$ for the original instance Π by

$$\hat{\sigma}(x) = \begin{cases} \hat{\sigma}'(x^{(i_1)}) & \text{if } x \in X_1 \\ \hat{\sigma}'(x^{(i_2)}) & \text{if } x \in X_2. \end{cases}$$

By the above argument, $\hat{\sigma}$ is (r_1, r_2)-average. Moreover, it satisfies Π, since $\hat{\sigma}'$ satisfies Π'. ◀

Putting all pieces together, we have Theorem 16.

Proof of Theorem 16. We give an FPT reduction from instances in the Baby PIH (Theorem 13) to AVG-r-GAP-2CSP. Then, since the Baby PIH holds under W[1] $\neq$ FPT, we deduce that the Average Baby PIH also holds under W[1] $\neq$ FPT.

For any 2CSP instance $\Pi_0 = (X_0, \Sigma_0, \Phi_0)$, we can construct a bipartite 2CSP instance $\Pi_1 = (X_1, \Sigma_1, \Phi_1)$ by Lemma 20, and then construct an AVG-r-GAP-2CSP instance $\Pi = (X, \Sigma, \Phi)$ from Π_1 by Lemma 22. Trivially, Π is satisfiable when Π_0 is satisfiable. When Π_0 is not r-list satisfiable, Π_1 is not (r_1, r_2)-average list satisfiable for all constants r_1, r_2 with $r_1 + r_2 \ge 2(1-\varepsilon)r$, and thus Π is not $(1-\varepsilon)r$-average list satisfiable. Furthermore, Π has rectangular relations because Π_1 has rectangular relations.

Moreover, the running time of this reduction can be bounded by

$$f(|X_0| + |\Phi_0| + 1/\varepsilon + r)|\Sigma_0|^{O(1)}$$

for a computable function f, and

$$|X| + |\Phi| \le f(|X_0| + |\Phi_0| + 1/\varepsilon + r)|\Sigma_0|^{O(1)}$$

as well, so the reduction is an FPT reduction. ◀

3.2 Average Baby PIH on Dense and Sparse Instances

In this section we prove Theorem 2, which is divided into two separate lemmas to help with readability. For our purposes, a 2CSP instance $\Pi = (X, \Sigma, \Phi)$ is *dense* if $|\Phi| = \omega(|X|)$; or it is *sparse*, if $|\Phi| = O(|X|)$.

▶ **Lemma 23.** *Under* W[1] $\neq$ FPT*, the Average Baby PIH holds for all* Avg-*r*-Gap-2CSP *instances that are dense.*

Proof. The reduction in the proof of Lemma 20 yields *complete* bipartite 2CSP instances $\Pi_1 = (X_1 \dot\cup X_2, \Sigma_1, \Phi_1)$, i.e., for each $x_1 \in X_1$ and $x_2 \in X_2$, there exists a constraint $\varphi = (x_1x_2, C) \in \Phi$. Then the reduction in the proof of Lemma 22 makes $|X_2|$ copies of X_1 and $|X_1|$ copies of $|X_2|$, while keeping the constraints in each pair of copies. So in the final instance $\Pi = (X, \Sigma, \Phi)$ from the proof of Theorem 16, the number of constraints is

$$|\Phi| = |X_1|^2|X_2|^2 = \frac{|X|^2}{4}.$$

Now consider any function $h \in \omega(1)$. We produce a new instance $\Pi' = (X', \Sigma', \Phi')$ by simply copying Π for t times, where t is chosen as the minimum number satisfying

$$h(t|X|) \geq \frac{|X|}{4}.$$

Note that there is no constraint between different copies. Then, the new parameter is $|X'| = t|X|$, and

$$|\Phi'| = t|\Phi| = \frac{t|X|^2}{4} = \frac{|X|}{4}|X'| \leq h(|X'|)|X'|.$$

It's clear that this reduction runs in FPT time. Also, if Π is satisfiable, then Π' is satisfiable. If any satisfying multi-assignment to Π must have total size more than $r|X|$, then any satisfying multi-assignment to Π' must assign each copy of Π more than $r|X|$ values, so in total more than $r|X'|$ values, preserving the gap. ◀

▶ **Lemma 24.** *Let $r > 1$. If there exists a constant $c > 0$ such that no* FPT *algorithm can solve* Avg-*r*-Gap-2CSP *on instance $\Pi = (X, \Sigma, \Phi)$ with $|\Phi| \leq c \cdot r|X|$, i.e., Π is sparse, then the PIH holds.*

Proof Sketch. Let $\Pi = (X, \Sigma, \Phi)$ be a NO instance of Avg-r-Gap-2CSP. For any (standard) assignment $\sigma : X \to \Sigma$, assume that σ violates t constraints, then one can simply add at most $2t$ values to σ and obtain a satisfying multi-assignment $\hat{\sigma}$ with total size $|X| + 2t$. Since Π is a NO instance, we have $|X| + 2t > r|X|$. Thus,

$$t > \frac{r-1}{2}|X| = \frac{r-1}{2c \cdot r} \cdot c \cdot r|X| \geq \frac{r-1}{2c \cdot r}|\Phi|.$$

In other words, any assignment to Π must violate a constant fraction of the constraints in Π. This gives a reduction from Avg-r-Gap-2CSP to PIH. ◀

Putting Lemma 23 and Lemma 24 together, we obtain Theorem 2. As already mentioned in the introduction, this result indicates that the current barrier to the W[1]-hardness of the PIH is the lack of reduction for Avg-r-Gap-2CSP on sparse instances, i.e., instances with linearly many constraints.

References

1 Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k-restrictions. *ACM Trans. Algorithms*, 2(2):153–177, 2006. doi:10.1145/1150334.1150336.

2 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. doi:10.1145/278298.278306.

3 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998. doi:10.1145/273865.273901.

4 Libor Barto and Marcin Kozik. Combinatorial gap theorem and reductions between promise CSPs. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1204–1220. SIAM, 2022. doi:10.1137/1.9781611977073.50.

5 Yijia Chen, Yi Feng, Bundit Laekhanukit, and Yanlin Liu. Simple combinatorial construction of the $k^{o(1)}$-lower bound for approximating the parameterized k-Clique. In *2025 Symposium on Simplicity in Algorithms (SOSA)*, pages 263–280. Society for Industrial and Applied Mathematics, 2025. doi:10.1137/1.9781611978315.21.

6 Yijia Chen and Bingkai Lin. The constant inapproximability of the parameterized Dominating Set problem. *SIAM J. Comput.*, 48(2):513–533, 2019. doi:10.1137/17M1127211.

7 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.

8 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007. doi:10.1145/1236457.1236459.

9 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.

10 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006. doi:10.1007/3-540-29953-X.

11 Venkatesan Guruswami, Bingkai Lin, Xuandi Ren, Yican Sun, and Kewen Wu. Almost optimal time lower bound for approximating parameterized Clique, CSP, and more, under ETH. *CoRR*, abs/2404.08870, 2024. doi:10.48550/arXiv.2404.08870.

12 Venkatesan Guruswami, Bingkai Lin, Xuandi Ren, Yican Sun, and Kewen Wu. Parameterized Inapproximability Hypothesis under Exponential Time Hypothesis. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 24–35. ACM, 2024. doi:10.1145/3618260.3649771.

13 Venkatesan Guruswami, Xuandi Ren, and Sai Sandeep. Baby PIH: parameterized inapproximability of min CSP. In Rahul Santhanam, editor, *39th Computational Complexity Conference, CCC 2024, July 22-25, 2024, Ann Arbor, MI, USA*, volume 300 of *LIPIcs*, pages 27:1–27:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.CCC.2024.27.

14 Karthik C. S. and Subhash Khot. Almost polynomial factor inapproximability for parameterized k-Clique. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPIcs*, pages 6:1–6:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.CCC.2022.6.

15 Karthik C. S., Bundit Laekhanukit, and Pasin Manurangsi. On the parameterized complexity of approximating Dominating Set. *J. ACM*, 66(5):33:1–33:38, 2019. doi:10.1145/3325116.

16 Karthik C. S. and Inbal Livni Navon. On hardness of approximation of parameterized Set Cover and Label Cover: Threshold graphs from error correcting codes. In Hung Viet Le and Valerie King, editors, *4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11-12, 2021*, pages 210–223. SIAM, 2021. doi:10.1137/1.9781611976496.24.

17 Shuangle Li, Bingkai Lin, and Yuwei Liu. Improved lower bounds for approximating parameterized nearest codeword and related problems under ETH. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia*, volume 297 of *LIPIcs*, pages 107:1–107:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ICALP.2024.107.

18 Bingkai Lin. A simple gap-producing reduction for the parameterized Set Cover problem. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 81:1–81:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.ICALP.2019.81.

19 Bingkai Lin. Constant approximating k-Clique is W[1]-hard. 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 1749–1756. ACM, 2021. doi:10.1145/3406325.3451016.

20 Bingkai Lin, Xuandi Ren, Yican Sun, and Xiuhan Wang. Constant approximating parameterized k-SetCover is W[2]-hard. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3305–3316. SIAM, 2023. doi:10.1137/1.9781611977554.CH126.

21 Daniel Lokshtanov, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of Directed Odd Cycle Transversal. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2181–2200. SIAM, 2020. doi:10.1137/1.9781611975994.134.

22 Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *J. ACM*, 41(5):960–981, 1994. doi:10.1145/185675.306789.

23 Pasin Manurangsi. Tight running time lower bounds for strong inapproximability of maximum k-Coverage, Unique Set Cover and related problems (via t-wise agreement testing theorem). In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 62–81. SIAM, 2020. doi:10.1137/1.9781611975994.5.

A From Average Baby PIH to Inapproximability of k-ExactCover

We present a proof relying on a construction that slightly differs from the one in [13]. Their proof makes use of the (T, m)-set gadget [22, 1] that was previously used to show the hardness of approximating SetCover problem. On the other hand, our proof develops a novel composition of the collision number of ECCs (recall Definition 17) with the following well-known combinatorial object.

▶ **Definition 25** (Hypercube Partition System). *Let A, B be two sets. Then the (A, B)-*hypercube partition system *is defined by*

- *the universe $\mathcal{M} = A^B \left(= \{z \mid \text{a function } z : B \to A\}\right)$, and*
- *a collection of subsets $\{P_{x,y}\}_{x \in B, y \in A}$ where each $P_{x,y} = \{z \in \mathcal{M} \mid z(x) = y\}$.*

▶ **Theorem 26** (cf. Theorem 21 in [13]). *Assume that the Average Baby PIH holds on all 2CSP instances with rectangular relations. Then k-ExactCover cannot be approximated in* FPT *time within any constant factor. More precisely, for every constant $r > 1$ no* FPT *algorithm, on a given k-SetCover instance $\Pi = (S, U)$ with size n and $k \geq 1$, can distinguish between the following two cases:*

- *We can choose k* disjoint *sets in S whose union is U.*
- *U is not the union of any rk sets in S.*

Proof. Let $\Pi = (X, \Sigma, \Phi)$ be an AVG-r-GAP-2CSP instance with rectangular relations. We set $k = |X|$. Moreover, for each rectangular constraint $\varphi_j = (x_{i_1}x_{i_2}, C_j) \in \Phi$ we use Q_j to denote the underlying set and $\pi_j, \rho_j : \Sigma \to Q_j$ the associated mappings as in Definition 14. That is, for every $a, b \in \Sigma$, it holds that $(a, b) \in C_j$ if and only if $\pi_j(a) = \rho_j(b)$. Then we set

$$t = \max_{\varphi_j \in \Phi} |Q_j|. \tag{12}$$

Clearly, we can assume without loss of generality

$$t \leq |\Pi|.$$

Now we reduce Π to a k-EXACTCOVER instance. To that end, we choose a further alphabet Δ whose size is a prime and satisfies

$$\max\left\{\lceil \log t \rceil, 2^{2r^2k^2}\right\} \leq |\Delta| \leq 2\max\left\{\lceil \log t \rceil, 2^{2r^2k^2}\right\}.$$

Moreover,let

$$d = \left\lceil \frac{2r^2k^2 \log t}{\log |\Delta|} \right\rceil.$$

This leads to the following code with very large distance (here we simply use Reed-Solomon code again)

$$\text{Enc} : \Delta^{\left\lceil \frac{\log t}{\log |\Delta|} \right\rceil} \to \Delta^d.$$

Plugging

$$k \leftarrow \left\lfloor \frac{\log t}{\log |\Delta|} \right\rfloor, \ m \leftarrow d, \ p \leftarrow |\Delta|,^{6} \text{ and } \varepsilon \leftarrow 1/2$$

in Theorem 18 we conclude that the 1/2-collision number of Enc is

$$\text{Col}_{1/2}(\text{Enc}) \geq \sqrt{\frac{d}{\log t / \log |\Delta|}} > rk.$$

Observe that (12) implies that every tuple in C_i can be identified with a string in $\Delta^{\left\lceil \frac{\log m}{\log |\Delta|} \right\rceil}$, i.e., the domain of Enc.

Then, for each variable $x \in X$ and every its possible value $a \in \Sigma$, we define a set $S_{x,a}$ as follows. For each constraint $\varphi_j = (x_{i_1}x_{i_2}, C_j) \in \Phi$ with associated set T_j and mappings $\pi_j, \rho_j : \Sigma \to Q_j$, and for each $\ell \in [d]$, we construct a $([2], \Delta)$-hypercube partition system

$$\left(\mathcal{M}^{(j,\ell)}, \{P_{u,v}^{(j,\ell)}\}_{u \in \Delta, v \in [2]}\right). \tag{13}$$

[6] Observe that

$$\left\lfloor \frac{\log t}{\log |\Delta|} \right\rfloor < \left\lceil \frac{2r^2k^2 \log t}{\log |\Delta|} \right\rceil \leq \lceil \log t \rceil \leq |\Delta|,$$

hence, the condition $k < m \leq p$ in Theorem 18 is satisfied.

Then for each $(a,b) \in C_j$ we add $P^{(j,\ell)}_{\mathrm{Enc}(\pi_j(a))[\ell],1}$ to $S_{x_{i_1},a}$ and similarly $P^{(j,\ell)}_{\mathrm{Enc}(\rho_j(b))[\ell],2}$ to $S_{x_{i_2},b}$. Finally, let the universe be

$$U = \dot{\bigcup}_{\varphi_j \in \Phi, \ell \in [d]} \mathcal{M}^{(j,\ell)}, \quad \text{and} \quad \mathcal{S} = \left\{ S_{x,a} \;\middle|\; x \in X \text{ and } a \in \Sigma \right\}.$$

For the completeness, let $\sigma : X \to \Sigma$ be a satisfying assignment of Π, it is routine to check that $\{S_{x,\sigma(x)}\}_{x \in X}$ is a partition of U.

For the soundness, assume that every satisfying multi-assignment of Π has total size at least rk (cf. Definition 7). Let $\mathcal{S}' \subseteq \mathcal{S}$ be a cover of U. Consider the multi-assignment that maps every variable $x \in X$ to $\left\{ a \in \Sigma \;\middle|\; S_{x,a} \in \mathcal{S}' \right\}$. If this multi-assignment satisfies Π, the our assumption implies $|\mathcal{S}'| \geq rk$. Otherwise, assume that there exists some constraint $\varphi_j = (x_{i_1} x_{i_2}, C_j) \in \Phi$ which is not satisfied. Note that the above multi-assignment assigns x_{i_1} to $E_1 = \left\{ a \in \Sigma \;\middle|\; S_{x_{i_1},a} \in \mathcal{S}' \right\}$ and x_{i_2} to $E_2 = \left\{ b \in \Sigma \;\middle|\; S_{x_{i_2},b} \in \mathcal{S}' \right\}$. Since φ_j is not satisfied, for all $(a,b) \in E_1 \times E_2$ we have $\mathrm{Enc}(\pi_j(a)) \neq \mathrm{Enc}(\rho_j(b))$. However, for each $\ell \in [d]$, since $\mathcal{M}^{(j,\ell)}$ is covered by $\mathcal{S}'$, there must exist $a \in E_1$ and $b \in E_2$ with $\mathrm{Enc}(\pi_j(a))[\ell] = \mathrm{Enc}(\rho_j(b))[\ell]$. Therefore, the set $\{\pi_j(a)\}_{a \in E_1} \cup \{\rho_j(b)\}_{b \in E_2}$ collides on all coordinates $\ell \in [d]$, hence it must have size at least $\mathrm{Col}_{1/2}(\mathrm{Enc})$. We deduce

$$|\mathcal{S}'| \geq |E_1| + |E_2| \geq \left| \{\pi_j(a)\}_{a \in E_1} \cup \{\rho_j(b)\}_{b \in E_2} \right| \geq \mathrm{Col}_{1/2}(\mathrm{Enc}) > rk.$$

Finally, in each hypercube partition system (13) it holds that

$$\left| \mathcal{M}^{(j,\ell)} \right| = 2^{|\Delta|} \leq 4^{\lceil \log t \rceil} + 4^{2^{2r^2k^2}} \leq |\Pi|^2 + 4^{2^{2r^2k^2}},$$

and there are at most $\binom{k}{2} d \leq k^2 r^2 k^2 \log t \leq r^2 k^4 \log |\Pi|$ such systems. The size of the universe U is thus at most $g(r,k)|\Pi|^3$ for some appropriate computable function $g : \mathbb{N}^2 \to \mathbb{N}$, while the parameter of the k-ExactCover instance remains $k = |X|$. It follows easily that the running time of this reduction is FPT. ◀

Combining Theorem 26 and Theorem 16, we obtain:

▶ **Theorem 27.** *For any constant $r > 1$, r-approximating k-ExactCover is* W[1]*-hard.*