On the Parameterized Complexity of Approximating Dominating Set

Karthik C. S.[*]   Bundit Laekhanukit[†]   Pasin Manurangsi[‡]

April 12, 2018

Abstract

We study the parameterized complexity of approximating the k-Dominating Set (DomSet) problem where an integer k and a graph G on n vertices are given as input, and the goal is to find a dominating set of size at most $F(k) \cdot k$ whenever the graph G has a dominating set of size k. When such an algorithm runs in time $T(k) \cdot \text{poly}(n)$ (i.e., FPT-time) for some computable function T, it is said to be an $F(k)$-*FPT-approximation algorithm* for k-DomSet. Whether such an algorithm exists is listed in the seminal book of Downey and Fellows (2013) as one of the "most infamous" open problems in Parameterized Complexity. This work gives an almost complete answer to this question by showing the non-existence of such an algorithm under $\mathsf{W}[1] \neq \mathsf{FPT}$ and further providing tighter running time lower bounds under stronger hypotheses. Specifically, we prove the following for every computable functions T, F and every constant $\varepsilon > 0$:

- Assuming $\mathsf{W}[1] \neq \mathsf{FPT}$, there is no $F(k)$-*FPT-approximation algorithm* for k-DomSet.
- Assuming the Exponential Time Hypothesis (ETH), there is no $F(k)$-approximation algorithm for k-DomSet that runs in $T(k) \cdot n^{o(k)}$ time.
- Assuming the Strong Exponential Time Hypothesis (SETH), for every integer $k \geq 2$, there is no $F(k)$-approximation algorithm for k-DomSet that runs in $T(k) \cdot n^{k-\varepsilon}$ time.
- Assuming the k-SUM Hypothesis, for every integer $k \geq 3$, there is no $F(k)$-approximation algorithm for k-DomSet that runs in $T(k) \cdot n^{\lceil k/2 \rceil - \varepsilon}$ time.

Previously, only constant ratio FPT-approximation algorithms were ruled out under $\mathsf{W}[1] \neq \mathsf{FPT}$ and $(\log^{1/4-\varepsilon} k)$-FPT-approximation algorithms were ruled out under ETH [Chen and Lin, FOCS 2016]. Recently, the non-existence of an $F(k)$-FPT-approximation algorithm for any function F was shown under Gap-ETH [Chalermsook et al., FOCS 2017]. Note that, to the best of our knowledge, no running time lower bound of the form $n^{\delta k}$ for any absolute constant $\delta > 0$ was known before even for any constant factor inapproximation ratio.

Our results are obtained by establishing a connection between communication complexity and hardness of approximation, generalizing the ideas from a recent breakthrough work of Abboud et al. [FOCS 2017]. Specifically, we show that to prove hardness of approximation of a certain parameterized variant of the label cover problem, it suffices to devise a specific protocol for a communication problem that depends on which hypothesis we rely on. Each of these communication problems turns out to be either a well studied problem or a variant of one; this allows us to easily apply known techniques to solve them.

---

[*]Weizmann Institute of Science, Israel. Email: `karthik.srikanta@weizmann.ac.il`

[†]Shanghai University of Finance and Economics, China & Simons Institute for the Theory of Computing, USA. Email: `blaekhan@mpi-inf.mpg.de`

[‡]University of California, Berkeley, USA. Email: `pasin@berkeley.edu`

1 Introduction

In the *dominating set* problem (DomSet), we are given an undirected graph G on n vertices and an integer k, and the goal is to decide whether there is a subset of vertices $S \subseteq V(G)$ of size k such that every vertex outside S has a neighbor in S (i.e., S *dominates* every vertex in G and is thus called a *dominating set*). Often regarded as one of the classical problems in computational complexity, DomSet was first shown to be NP-complete in the seminal work of Karp [Kar72][1]. Thus, its optimization variant, namely the *minimum dominating set*, where the goal is to find a dominating set of smallest possible size, is also NP-hard. To circumvent this apparent intractability of the problem, the study of an approximate version was initiated. The quality of an approximation algorithm is measured by the *approximation ratio*, which is the ratio between the size of the solution output by an algorithm and the size of the minimum dominating set. A simple greedy heuristic for the problem, which has by now become one of the first approximation algorithms taught in undergraduate and graduate algorithm courses, was intensively studied and was shown to yield a $(\ln n - \ln \ln n + \Theta(1))$-approximation for the problem [Joh74, Chv79, Lov75, Sri95, Sla96]. On the opposite side, a long line of works in hardness of approximation [LY94, RS97, Fei98, AMS06, Mos15] culminated in the work of Dinur and Steurer [DS14], who showed that obtaining a $(1 - \varepsilon) \ln n$-approximation for the problem is NP-hard for every $\varepsilon > 0$. This essentially settles the approximability of the problem.

Besides approximation, another widely-used technique to cope with NP-hardness is *parameterization*. The parameterized version of DomSet, which we will refer to simply as *k*-DomSet, is exactly the same as the original decision version of the problem except that now we are not looking for a polynomial time algorithm but rather a *fixed parameter tractable* (FPT) algorithm – one that runs in time $T(k) \cdot \text{poly}(n)$ for some computable function T (e.g., $T(k) = 2^k$ or 2^{2^k}). Such running time will henceforth be referred to as *FPT time*. Alas, even with this relaxed requirement, *k*-DomSet still remains intractable: in the same work that introduced the W-hierarchy, Downey and Fellows [DF95a] showed that *k*-DomSet is complete for the class W[2], which is generally believed to not be contained in FPT, the class of fixed parameter tractable problems. In the ensuing years, stronger running time lower bounds have been shown for *k*-DomSet under strengthened assumptions. Specifically, Chen et al. [CHKX06] ruled out $T(k) \cdot n^{o(k)}$-time algorithm for *k*-DomSet assuming the *Exponential Time Hypothesis* (ETH)[2]. Furthermore, Pătrașcu and Williams [PW10] proved, for every $k \geq 2$, that, under the *Strong Exponential Time Hypothesis* (SETH)[3], not even $O(n^{k-\varepsilon})$ algorithm exists for *k*-DomSet for any $\varepsilon > 0$. Note that the trivial algorithm that enumerates through every *k*-size subset and checks whether it forms a dominating set runs in $O(n^{k+1})$ time. It is possible to speed up this running time using fast matrix multiplication [EG04, PW10]. In particular, Pătrașcu and Williams [PW10] themselves also gave an $n^{k+o(1)}$-time algorithm for every $k \geq 7$, painting an almost complete picture of the complexity of the problem.

Given the strong negative results for *k*-DomSet discussed in the previous paragraph, it is natural to ask whether we can somehow incorporate the ideas from the area of approximation algorithms to come up with a *fix parameter approximation (FPT-approximation) algorithm* for *k*-DomSet.

---

[1]To be precise, Karp showed NP-completeness of Set Cover, which is well-known to be equivalent to DomSet.

[2]ETH [IP01] states that no subexponential time algorithm can solve 3-CNF-SAT; see Hypothesis 2.

[3]SETH [IP01], a strengthening of ETH, states that, for every $\varepsilon > 0$, there exists $k = k(\varepsilon) \in \mathbb{N}$ such that no $O(2^{(1-\varepsilon)n})$-time algorithm can solve *k*-CNF-SAT; see Hypothesis 3.

To motivate the notion of FPT-approximation algorithms, first notice that the seemingly reasonable $O(\log n)$-approximation given by the greedy algorithm can become unjustifiable when the optimum k is small since it is even possible that the overhead paid is unbounded in terms of k. As a result, FPT-approximation algorithms require the approximation ratios to be bounded only in terms of k; specifically, for any computable function F, we say that an algorithm is an $F(k)$-FPT-approximation algorithm for k-DomSet if it runs in FPT time and, on any input (G, k) such that the minimum dominating set of G is of size at most k, it outputs a dominating set of size at most $F(k) \cdot k$.

This brings us to the main question addressed in our work: *Is there an $F(k)$-FPT-approximation algorithm for k-DomSet for some computable function F?* This question, which dates back to late 1990s (see, e.g., [DFM06]), has attracted significant attention in literature [DFM06, CGG06, DFMR08, CH10, DF13, HKK13, CHK13, BEKP15, CL16, CCK+17]. In fact, it is even listed in the seminal textbook of Downey and Fellows [DF13] as one of the six "most infamous" open questions[4] in the area of Parameterized Complexity. While earlier attempts fell short of ruling out either $F(k)$ that is super constant or all FPT algorithms (see Section 1.2 for more details), the last couple of years have seen significant progresses on the problem. In a remarkable result of Chen and Lin [CL16], it was shown that no FPT-approximation for k-DomSet exists for any constant ratio unless $\mathsf{W}[1] = \mathsf{FPT}$. They also proved that, assuming ETH, the inapproximability ratio can be improved to $\log^{1/4-\varepsilon} k$ for any constant $\varepsilon > 0$. Very recently, Chalermsook et al. [CCK+17] proved, under the *Gap Exponential Time Hypothesis* (Gap-ETH)[5], that no $F(k)$-approximation algorithm for k-DomSet exists for any computable function F. Such non-existence of FPT-approximation algorithms is referred to in literature as the *total FPT-inapproximability* of k-DomSet.

Although Chalermsook et al.'s result on the surface seems to settle the parameterized complexity of approximating dominating set, several aspects of the result are somewhat unsatisfactory. First, while Gap-ETH may be plausible, it is quite strong and, in a sense, does much of the work in the proof. Specifically, Gap-ETH itself already gives the gap in hardness of approximation; once there is such a gap, it is not hard[6] to build on it and prove other inapproximability results. As an example, in the analogous situation in NP-hardness of approximation, once one inapproximability result can be shown, others follow via relatively simple gap-preserving reductions (see, e.g., [PY91]). On the other hand, creating a gap in the first place requires the PCP Theorem [AS98, ALM+98], which involves several new technical ideas such as local checkability of codes and proof composition[7]. Hence, it is desirable to bypass Gap-ETH and prove total FPT-inapproximability under assumptions that do not involve hardness of approximation in the first place. Drawing a parallel to the theory of NP-hardness of approximation once again, it is imaginable that a success in bypassing Gap-ETH may also reveal a "PCP-like Theorem" for parameterized complexity.

An additional reason one may wish to bypass Gap-ETH for the total FPT-inapproximability of k-

---

[4]Since its publication, one of the questions, the parameterized complexity of biclique, has been resolved [Lin15].

[5]Gap-ETH [Din16, MR16], another strengthening of ETH, states that no subexponential time algorithm can distinguish satisfiable 3-CNF formulae from ones that are not even $(1 - \varepsilon)$-satisfiable for some $\varepsilon > 0$.

[6]One issue grossed over in this discussion is that of *gap amplification*. While Gap-ETH gives some constant gap, Chalermsook et al. still needed to amplify the gap to arrive at total FPT-inapproximability. Fortunately, unlike the NP-hardness regime that requires Raz's parallel repetition theorem [Raz98], the gap amplification step in [CCK+17], while non-trivial, only involved relatively simple combinatorial arguments. (See [CCK+17, Theorem 4.3].)

[7]Even in the "combinatorial proof" of the PCP Theorem [Din07], many of these tools still remain in use, specifically in the alphabet reduction step of the proof.

DomSet is that the latter is a statement purely about parameterized complexity, so one expects it to hold under a standard parameterized complexity assumption. Given that Chen and Lin [CL16] proved W[1]-hardness of approximating k-DomSet to within any constant factor, a concrete question here is whether we can show W[1]-hardness of approximation for every function $F(k)$:

Open Question 1.1. *Can we base the total FPT-inapproximability of k-DomSet on* W[1] $\neq$ FPT*?*

Another issue not completely resolved by [CCK+17] is the running time lower bound. While the work gives a quite strong running time lower bound that rules out any $T(k) \cdot n^{o(k)}$-time $F(k)$-approximation algorithm for any computable functions T and F, it is still possible that, say, an $O(n^{0.5k})$-time algorithm can provide a very good (even constant ratio) approximation for k-DomSet. Given the aforementioned $O(n^{k-\varepsilon})$ running time lower bound for exact algorithms of k-DomSet by Pătraşcu and Williams [PW10], it seems reasonable to ask whether such a lower bound can also be established for approximation algorithms:

Open Question 1.2. *Is it hard to approximate k-DomSet in $O(n^{k-\varepsilon})$-time?*

This question has perplexed researchers, as even with the running time of, say, $O(n^{k-0.1})$, no $F(k)$-approximation algorithm is known for k-DomSet for any computable function F.

1.1   Our Contributions

Our contributions are twofold. Firstly, at a higher level, we prove parameterized inapproximabilty results for k-DomSet, answering the two aforementioned open questions (and more). Secondly, at a lower level, we demonstrate a connection between communication complexity and parameterized inapproximability, allowing us to translate running time lower bounds for parameterized problems into parameterized hardness of approximation. This latter part of the contribution extends ideas from a recent breakthrough of Abboud et al. [ARW17], who discovered similar connections and used them to establish inapproximability for problems in P. In this subsection, we only focus on the first part of our contributions. The second part will be discussed in detail in Section 2.

Parameterized Inapproximability of Dominating Set

Our first batch of results are the inapproximability results for k-DomSet under various standard assumptions in parameterized complexity and fine-grained complexity: W[1] $\neq$ FPT, ETH, SETH and the k-SUM Hypothesis. First, we show total inapproximability of k-DomSet under W[1] $\neq$ FPT. In fact, we show an even stronger[8] inapproximation ratio of $(\log n)^{1/\mathsf{poly}(k)}$:

Theorem 1.3. *Assuming* W[1] $\neq$ FPT*, no FPT time algorithm can approximate k-DomSet to within a factor of $(\log n)^{1/\mathsf{poly}(k)}$.*

---

[8]Note that the factor of the form $(\log n)^{1/\mathsf{poly}(k)}$ is stronger than that of the form $F(k)$. To see this, assume that we have an $F(k)$-FPT-approximation algorithm for some computable function F. We can turn this into a $(\log n)^{1/\mathsf{poly}(k)}$-approximation algorithm by first checking which of the two ratios is smaller. If $F(k)$ is smaller, then just run the $F(k)$-FPT-approximation algorithm. Otherwise, use brute force search to solve the problem. Since the latter case can only occur when $n \leq \exp(F(k)^{\mathsf{poly}(k)})$, we have that the running time remains FPT.

4

Our result above improves upon the constant factor inapproximability result of Chen and Lin [CL16] and resolves the question of whether we can base total FPT inapproximability of k-DomSet on a purely parameterized complexity assumption. Furthermore, if we are willing to assume the stronger ETH, we can even rule out all $T(k) \cdot n^{o(k)}$-time algorithms:

Theorem 1.4. *Assuming* ETH, *no* $T(k) \cdot n^{o(k)}$-*time algorithm can approximate k-DomSet to within a factor of* $(\log n)^{1/\text{poly}(k)}$.

Note that the running time lower bound and approximation ratio ruled out by the above theorem are exactly the same as those of Charlermsook et al.'s result based on Gap-ETH [CCK$^+$17]. In other words, we successfully bypass Gap-ETH from their result completely. Prior to this, the best known ETH-based inapproximability result for k-DomSet due to Chen and Lin [CL16] ruled out only $(\log^{1/4+\varepsilon} k)$-approximation for $T(k) \cdot n^{o(\sqrt{k})}$-time algorithms.

Assuming the even stronger hypothesis, SETH, we can rule out $O(n^{k-\varepsilon})$-time approximation algorithms for k-DomSet, matching the running time lower bound from [PW10] while excluding not only exact but also approximation algorithms. We note, however, that the approximation ratio we get in this case is not $(\log n)^{1/\text{poly}(k)}$ anymore, but rather $(\log n)^{1/\text{poly}(k,e(\varepsilon))}$ for some function e, which arises from SETH and the Sparsification Lemma [IPZ01].

Theorem 1.5. *There is a function $e : \mathbb{R}^+ \to \mathbb{N}$ such that, assuming* SETH, *for every integer $k \geq 2$ and for every $\varepsilon > 0$, no $O(n^{k-\varepsilon})$-time algorithm can approximate k-DomSet to within a factor of* $(\log n)^{1/\text{poly}(k,e(\varepsilon))}$.

Finally, to demonstrate the flexibility of our proof techniques (which will be discussed at length in the next section), we apply the framework to the k-SUM Hypothesis[9] which yields an $n^{\lceil k/2 \rceil - \varepsilon}$ running time lower bound for approximating k-DomSet as stated below.

Theorem 1.6. *Assuming the k-SUM Hypothesis, for every integer $k \geq 3$ and for every $\varepsilon > 0$, no $O(n^{\lceil k/2 \rceil - \varepsilon})$-time algorithm can approximate k-DomSet to within a factor of* $(\log n)^{1/\text{poly}(k)}$.

We remark here that the k-SUM problem is known to be W[1]-hard [DF95a, ALW14] and our proof of Theorem 1.6 indeed yields an alternative proof of W[1]-hardness of approximating k-DomSet (Theorem 1.3). Nevertheless, we provide a different self-contained W[1]-hardness reduction directly from Clique since the ideas there are also useful for our ETH-hardness result (Theorem 1.4).

The summary of our results and those from previous works are shown in Table 1.

1.2 Comparison to Previous Works

In addition to the lower bounds previously mentioned, the parameterized inapproximability of k-DomSet has also been investigated in several other works [DFMR08, CHK13, HKK13, BEKP15]. Specifically, Downey et al. [DFMR08] showed that obtaining an *additive* constant approximation for k-DomSet is W[2]-hard. On the other hand, in [HKK13, CHK13], the authors ruled

---

[9]The k-SUM Hypothesis states that, for every $\varepsilon > 0, k \in \mathbb{N}$ such that $k \geq 3$, no $O(n^{\lceil k/2 \rceil - \varepsilon})$-time algorithm solves the k-SUM problem; see Hypothesis 4.

Summary of Previous Works and The Results in This Paper

Complexity Assumption	Inapproximability Ratio	Running Time Lower Bound	Reference
W[1] $\neq$ FPT	Any constant	$T(k) \cdot \text{poly}(n)$	[CL16]
	$(\log n)^{1/\text{poly}(k)}$	$T(k) \cdot \text{poly}(n)$	This paper
ETH	$(\log k)^{1/4+\epsilon}$	$T(k) \cdot n^{o(\sqrt{k})}$	[CL16]
	$(\log n)^{1/\text{poly}(k)}$	$T(k) \cdot n^{o(k)}$	This paper
Gap-ETH	$(\log n)^{1/\text{poly}(k)}$	$T(k) \cdot n^{o(k)}$	[CCK$^+$17]
SETH	Exact	$O(n^{k-\epsilon})$	[PW10]
	$(\log n)^{1/\text{poly}(k,e(\varepsilon))}$	$O(n^{k-\epsilon})$	This paper
k-SUM Hypothesis	$(\log n)^{1/\text{poly}(k)}$	$O(n^{\lceil k/2 \rceil -\epsilon})$	This paper

Table 1: Summary of our and previous results on k-DomSet. We only show those whose inapproximability ratios are at least some constant greater than one (i.e., we exclude additive inapproximability results). Here $e : \mathbb{R}^+ \to \mathbb{N}$ is some function, $T : \mathbb{N} \to \mathbb{N}$ can be any computable function and ε can be any positive constant. The Gap-ETH has been bypassed in this paper, and prior to this paper, the k-SUM Hypothesis had never been used in proving inapproximability of k-DomSet.

out $(\log k)^{1+\varepsilon}$-approximation in time $\exp(\exp((\log k)^{1+\varepsilon})) \cdot \text{poly}(n)$ for some fixed constant $\varepsilon > 0$ by assuming ETH and the *projection game conjecture* proposed in [Mos15]. Further, Bonnet et al. [BEKP15] ruled out $(1 + \varepsilon)$-FPT-approximation, for some fixed constant $\varepsilon > 0$, assuming Gap-ETH[10]. We note that, with the exception of W[2]-hardness results [DF95a, DFMR08], our results subsume all other aforementioned lower bounds regarding k-DomSet, both for approximation [CHK13, HKK13, BEKP15, CL16, CCK$^+$17] and exact algorithms [CHKX06, PW10].

While our techniques will be discussed at a much greater length in the next section (in particular we compare our technique with [ARW17] in Section 2.2), we note that our general approach is to first show inapproximability of a parameterized variant of the *Label Cover* problem called MaxCover and then reduce MaxCover to k-DomSet. The first step employs the connection between communication complexity and inapproximability of MaxCover, whereas the second step follows directly from the reduction in [CCK$^+$17] (which is in turn based on [Fei98]). While MaxCover was not explicitly defined until [CCK$^+$17], its connection to k-DomSet had been implicitly used both in the work of Pătrașcu and Williams [PW10] and that of Chen and Lin [CL16].

From this perspective, the main difference between our work and [PW10, CL16, CCK$^+$17] is the source of hardness for MaxCover. Recall that Pătrașcu and Williams [PW10] ruled out only exact algorithms; in this case, a relatively simple reduction gave hardness for the exact version of MaxCover. On the other hand, both Chalermsook et al. [CCK$^+$17] and Chen and Lin [CL16] ruled out approximation algorithms, meaning that they needed gaps in their hardness results

---

[10]The authors assume the same statement as Gap-ETH (albeit, with imperfect completeness) but have an additional assertion that it is implied by ETH (see Hypothesis 1 in [BEKP15]). It is not hard to see that their assumption can be replaced by Gap-ETH.

for MaxCover. Chalermsook et al. obtained their initial gap from their assumption (Gap-ETH), after which they amplified it to arrive at an arbitrarily large gap for MaxCover. On the other hand, [CL16] derived their gap from the hardness of approximating Maximum *k*-Intersection shown in Lin's earlier breakthrough work [Lin15]. Lin's proof [Lin15] made use of certain combinatorial objects called *threshold graphs* to prove inapproximability of Maximum *k*-Intersection. Unfortunately, this construction was not very flexible, in the sense that it produced MaxCover instances with parameters that were not sufficient for proving total-FPT-inapproximability for *k*-DomSet. Moreover, his technique (i.e., threshold graphs) was limited to reductions from *k*-Clique and was unable to provide a tight running time lower bound under ETH. By resorting to the connection between MaxCover and communication complexity, we can generate MaxCover instances with wider ranges of parameters from much more general assumptions, allowing us to overcome the aforementioned barriers.

Organization.   In the next section, we give an overview of our lower level contributions; for readers interested in the general ideas without too much notational overhead, this section covers most of the main ideas from our paper through a proof sketch of our W[1]-hardness of approximation result (Theorem 1.3). After that, in Section 3, we define additional notations and preliminaries needed to formalize our proofs. Section 4 provides a definition for Product Space Problems (PSP) and rewrites the hypotheses in these terms. Next, in Section 5, we establish a general theorem converting communication protocols to a reduction from PSP to MaxCover. Sections 6, 7 and 8 provide communication protocols for our problems of interest: Set Disjointness, Multi-Equality and Sum-Zero. Finally, in Section 9, we conclude with a few open questions and research directions.

2   Connecting Communication Complexity and Parameterized Inapproximability: An Overview

This section is devoted to presenting our connection between communication complexity and parameterized inapproximability (which is one of our main contributions as discussed in the introduction) and serves as an overview for all the proofs in this paper. As mentioned previously, our discovery of this connection is inspired by the work of Abboud et al. [ARW17] who showed the connection between the communication protocols and hardness of approximation for problems in P. More specifically, they showed how a Merlin-Arthur protocol for *Set Disjointness* with certain parameters implies the SETH-hardness of approximation for a problem called *PCP-Vectors* and used it as the starting point to prove inapproximability of other problems in P. We extend this idea by identifying a communication problem associated with each of the complexity assumptions (W[1] $\neq$ FPT, ETH, SETH and *k*-SUM Hypothesis) and then prove a generic theorem that translates communication protocols for these problems to conditional hardness of approximation for a parameterized variant of the *Label Cover* problem called MaxCover [CCK$^+$17]. Since the hardness of MaxCover is known to imply the hardness of *k*-DomSet [CCK$^+$17] (see Section 3.3), we have arrived at our inapproximability results for *k*-DomSet. As the latter part is not the contribution of this paper, we will focus on explaining the connection between communication complexity and the hardness of approximating MaxCover. We start by defining MaxCover:

Definition 2.1. *The input for* MaxCover *is a* label cover *instance* Γ, *which consists of a bipartite graph* $G = (U, W; E)$ *such that U is partitioned into $U_1 \cup \cdots \cup U_q$ and W is partitioned into $W_1 \cup \cdots \cup W_h$. We sometimes refer to U_i's and W_j's as left and right* super-nodes *of Γ, respectively.*

A solution to MaxCover *is called a* labeling, *which is a subset of vertices $S \subseteq W$ formed by picking a vertex w_j from each W_j (i.e., $|S \cap W_j| = 1$ for all $j \in [h]$). We say that a labeling S covers a left super-node U_i if there exists a vertex $u_i \in U_i$ such that u_i is a neighbor of every vertex in S. The goal in* MaxCover *is to find a labeling that covers the maximum fraction of left super-nodes.*

For concreteness, we focus on the W[1]-hardness proof (Theorem 1.3); at the end of this subsection, we will discuss how this fits into a larger framework that encapsulates other hypotheses as well.

For the purpose of our current discussion, it suffices to think of MaxCover as being parameterized by h, the number of right super-nodes; from this viewpoint, we would like to show that it is W[1]-hard to approximate MaxCover to within $(\log n)^{1/\mathsf{poly}(h)}$ factor. For simplicity, we shall be somewhat imprecise in our overview below, all proofs will be formalized later in the paper.

We reduce from the k-Clique problem, which is well-known to be W[1]-hard [DF95a]. The input to k-Clique is an integer k and a graph which we will call $G' = (V', E')$ to avoid confusion with the label cover graph. The goal is to determine whether G' contains a clique of size k. Recall that, to prove the desired W[1]-hardness, it suffices to provide an *FPT-reduction* from any k-Clique instance (G', k) to approximate MaxCover instance $G = (U, W; E)$; this is an FPT-time reduction such that the new parameter h is bounded by a function of the original parameter k. Furthermore, since we want a hardness of approximation result for the latter, we will also show that, when (G', k) is a YES instance of k-Clique, there is a labeling of G that covers all the left super-nodes. On the other hand, when (G', k) is a NO instance of k-Clique, we wish to show that every labeling of G will cover at most $1/(\log n)^{1/\mathsf{poly}(h)}$ fraction of the left super-nodes. If we had such a reduction, then we would have arrived at the total FPT-inapproximability of MaxCover under W[1] $\neq$ FPT. But, how would we come up with such a reduction? We will do this by devising a specific kind of protocol for a communication problem!

2.1 A Communication Problem for k-Clique

The communication problem related to k-Clique we consider is a multi-party problem where there are $h = \binom{k}{2}$ players, each associated with a two-element subset $\{i, j\}$ of $[k]$. The players cannot communicate with each other. Rather, there is a referee that they can send messages to. Each player $\{i, j\}$ is given two vertices $u_i^{\{i,j\}}$ and $u_j^{\{i,j\}}$ such that $\{u_i^{\{i,j\}}, u_j^{\{i,j\}}\}$ forms an edge in G'. The vertices $u_i^{\{i,j\}}$ and $u_j^{\{i,j\}}$ are allegedly the i-th and j-th vertices of a clique respectively. The goal is to determine whether there is indeed a k-clique in G' such that, for every $\{i, j\} \subseteq [k]$, $u_i^{\{i,j\}}$ and $u_j^{\{i,j\}}$ are the i-th and j-th vertices of the clique.

The communication protocol that we are looking for is a one-round protocol with public randomness and by the end of which the referee is the one who outputs the answer. Specifically, the protocol proceeds as follows. First, the players and the referee together toss r random coins. Then, each player sends an ℓ-bit message to the referee. Finally, the referee decides, based on the messages received and the randomness, either to accept or reject. The protocol is said to have perfect

completeness and soundness s if (1) when there is a desired clique, the referee always accepts and (2) when there is no such clique, the referee accepts with probability at most s. The model described here is referred to in the literature as the multi-party *Simultaneous Message Passing* (SMP) model [Yao79, BGKL03, FOZ16]. We refer to a protocol in the SMP model as an SMP protocol.

From Communication Protocol to MaxCover**.** Before providing a protocol for the previously described communication problem, let us describe how to turn the protocol into a label cover instance $G = (U = U_1 \cup \cdots \cup U_q, W = W_1 \cup \cdots \cup W_h; E)$.

- Let $h = \binom{k}{2}$. Again, we associate elements in $[h]$ with two-element subsets of $[k]$. Each right super-node $W_{\{i,j\}}$ represents Player $\{i,j\}$. Each vertex in $W_{\{i,j\}}$ represents a possible input to the player, i.e., we have one vertex $a_{\{u,v\}} \in W_{\{i,j\}}$ for each edge $\{u,v\} \in E'$ in the graph G'. Assume w.l.o.g. that $i < j$ and $u < v$. This vertex $a_{\{u,v\}}$ represents player $\{i,j\}$ receiving u and v as the alleged i-th and j-th vertices of the clique respectively.

- Let $q = 2^r$. We associate each element in $[q]$ with an r-bit string. For each $\gamma \in \{0,1\}^r$, the left super-node U_γ contains one node corresponding to each *accepting configuration* on randomness γ; that is, for each h-tuple of ℓ-bit strings $(m_{\{1,2\}}, \ldots, m_{\{k-1,k\}}) \in (\{0,1\}^\ell)^h$, there is a node $(m_{\{1,2\}}, \ldots, m_{\{k-1,k\}})$ in U_γ iff the referee on randomness γ and message $m_{\{1,2\}}, \ldots, m_{\{k-1,k\}}$ from all the players accepts.

- The edges in E are defined as follows. Recall that each node a in a right super-node $W_{\{i,j\}}$ corresponds to an input that each player receives in the protocol. For each $\gamma \in \{0,1\}^r$, suppose that the message produced on this randomness by the $\{i,j\}$-th player on the input corresponding to a is $m^{a,\gamma}$. We connect a to every accepting configuration on randomness γ that agrees with the message $m^{a,\gamma}$. More specifically, for every $\gamma \in \{0,1\}^r$, a is connected to every vertex $(m_{\{1,2\}}, \ldots, m_{\{k-1,k\}}) \in U_\gamma$ iff $m_{\{i,j\}} = m^{a,\gamma}$.

Consider any labeling $S \subseteq W$. It is not hard to see that, if we run the protocol where the $\{i,j\}$-th player is given the edge corresponding to the unique element in $S \cap W_{\{i,j\}}$ as an input, then the referee accepts a random string $\gamma \in \{0,1\}^r$ if and only if the left super-node U_γ is covered by the labeling S. In other words, the fraction of the left super-nodes covered by S is exactly equal to the acceptance probability of the protocol. This means that if (G', k) is a YES-instance of k-Clique, then we can select S corresponding to the edges of a k-clique and every left super-node will be covered. On the other hand, if (G', k) is a NO-instance of k-Clique, there is no subset S that corresponds to a valid k-clique, meaning that every labeling S covers at most s fraction of the edges. Hence, we have indeed arrived at hardness of approximation for MaxCover. Before we move on to describe the protocol, let us note that the running time of the reduction is $\text{poly}(2^{r+\ell h}, |E'|)$, which also gives an upper bound on the number of vertices in the label cover graph G.

SMP Protocol. Observe first that the trivial protocol, one where every player sends the whole input to the referee, does not suffice for us; this is because the message length ℓ is $\Omega(\log n)$, meaning that the running time of the reduction is $n^{\Omega(h)} = n^{\Omega(k^2)}$ which is not FPT time.

Nevertheless, there still is a simple protocol that does the job. Notice that the input vertices $u_i^{\{i,j\}}$ and $u_j^{\{i,j\}}$ given to Player $\{i,j\}$ are already promised to form an edge. Hence, the only thing the referee needs to check is whether each alleged vertex of the clique sent to different players are the same; namely, he only needs to verify that, for every $i \in [k]$, we have $u_i^{\{i,1\}} = u_i^{\{i,2\}} = \cdots = u_i^{\{i,i-1\}} = u_i^{\{i,i+1\}} = \cdots = u_i^{\{i,k\}}$. In other words, he only needs to check equalities for each of the k unknowns. The equality problem and its variants are extensively studied in communication complexity (see, e.g., [Yao79, KN97]). In our case, the protocol can be easily derived using any error-correcting code. Specifically, for an outcome $\gamma \in \{0,1\}^r$ of the random coin tosses, every Player $\{i,j\}$ encodes each of his input ($u_i^{\{i,j\}}$ and $u_j^{\{i,j\}}$) using a binary error-correcting code and sends only the γ-th bit of each encoded word to the referee. The referee then checks whether, for every $i \in [k]$, the received γ-th bits of the encodings of $u_i^{\{i,1\}}, u_i^{\{i,2\}}, \ldots, u_i^{\{i,k\}}$ are equal.

In the protocol described above, the message length ℓ is now only two bits (one bit per vertex), the randomness r used is logarithmic in the block length of the code, the soundness s is one minus the relative distance of the code. If we use a binary code with constant rate and constant relative distance (aka *good codes*), then r will be simply $O(\log \log n)$; this means that the running time of the reduction is $\text{poly}(n, \exp(O(k^2)))$ as desired. While the soundness in this case will just be some constant less than one, we can amplify the soundness by repeating the protocol multiple times independently; this increases the randomness and message length, but it is still not hard to see that, with the right number of repetitions, all parameters lie within the desired ranges. With this, we have completed our sketch for the proof of W[1]-hardness of approximating MaxCover.

2.2   A Framework for Parameterized Hardness of Approximation

The W[1]-hardness proof sketched above is an example of a much more general connection between communication protocol and the hardness of approximating MaxCover. To gain insight on this, consider any function $f : X_1 \times \cdots \times X_k \to \{0,1\}$. This function naturally induces both a communication problem and a computational problem. The communication problem for f is one where there are k players, each player i receives an input $a_i \in X_i$, and they together wish to compute $f(a_1, \ldots, a_k)$. The computational problem for f, which we call the *Product Space Problem*[11] of f (abbreviated as PSP(f)), is one where the input consists of subsets $A_1 \subseteq X_1, \ldots, A_k \subseteq X_k$ and the goal is to determine whether there exists $(a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k$ such that $f(a_1, \ldots, a_k) = 1$. The sketch reduction to MaxCover above in fact not only applies to the specific communication problem of k-Clique: the analogous construction is a generic way to translate any SMP protocol for the communication problem of any function f to a reduction from PSP(f) to MaxCover. To phrase it somewhat differently, if we have an SMP protocol for f with certain parameters and PSP(f) is hard to solve, then MaxCover is hard to approximate.

This brings us to the framework we use in this paper. It consists of only two steps. First, we rewrite the problem in the hypotheses as Product Space Problems of some family of functions $\mathcal{F}$. This gives us the conditional hardness for solving PSP($\mathcal{F}$). Second, we devise an SMP protocol for every function $f \in \mathcal{F}$. Given the connection outlined in the previous paragraph, this automatically yields the parameterized hardness of approximating MaxCover.

---

[11]The naming comes from the product structure of the domain of f.

To gain more intuition into the framework, note that in the case of k-Clique above, the function $f \in \mathcal{F}$ we consider is just the function $f : X_{\{1,2\}} \times \cdots \times X_{\{k,k-1\}}$ where each of $X_{\{1,2\}}, \cdots, X_{\{k,k-1\}}$ is a copy of the edge set. The function f "checks" that the edges selected form a clique, i.e., that, for every $i \in [k]$, the alleged i-th vertex of the clique specified in the $\{i,j\}$-coordinate is equal for every $j \neq i$. Since this is a generalization of the equality function, we call such a class of functions "multi-equality". It turns out that 3-CNF-SAT can also be written as PSP of multi-equality; each X_i contains assignments to $1/k$ fraction of the clauses and the function f checks that each variable is assigned the same value across all X_i's they appear in. A protocol essentially the same as the one given above also works in this setting and immediately gives our ETH-hardness result (Theorem 1.4)! Unfortunately, this does not suffice for our SETH-hardness. In that case, the function used is the k-way set disjointness; this interpretation of SETH is well-known (see, e.g., [Wil05]) and is also used in [ARW17]. Lastly, the k-SUM problem is already written in PSP form where f is just the Sum-Zero function that checks whether the sum of k specified numbers equals to zero.

Let us note that in the actual proof, we have to be more careful about the parameters than in the above sketch. Specifically, the reduction from MaxCover to k-DomSet from [CCK+17] incurs a blow-up in size that is exponential in terms of the number of vertices in each left super-node (i.e., exponential in $|U_\gamma|$). This means that we need $|U_1|, \ldots, |U_r| = o(\log n)$. In the context of communication protocol, this translates to keeping the message length $O(\log \log n)$ where $O(\cdot)$ hides a sufficiently small constant. Nevertheless, for the protocol for k-Clique reduction (and more generally for multi-equality), this does not pose a problem for us since the message length before repetitions is $O(1)$ bits; we can make sure that we apply only $O(\log \log n)$ repetitions to the protocol.

For Sum-Zero, known protocols either violate the above requirement on message length [Nis94] or use too much randomness [Vio15]. Nonetheless, a simple observation allows us to compose Nisan's protocol [Nis94] and Viola's protocol [Vio15] and arrive at a protocol with the best of both parameters. This new protocol may also be of independent interest beyond the scope of our work.

On the other hand, well-known communication complexity lower bounds on set disjointness [Raz92, KS92, BJKS04] rule out the existence of protocols with parameters we wish to have! [ARW17] also ran into this issue; in our language, they got around this problem by allowing the referee to receive an advice. This will also be the route we take. Even with advice, however, devising a protocol with the desired parameters is a technically challenging issue. In particular, until very recently, no protocol for set disjointness with $O(\log \log n)$ message length (and $o(n)$ advice length) was known. This was overcome in the work of Rubinstein [Rub18] who used algebraic geometric codes to give such a protocol for the two-player case. We extend his protocol in a straightforward manner to the k-player case; this extension was also suggested to us by Rubinstein [Rub17].

A diagram illustrating the overview of our approach can be found in Figure 1.

Comparison to Abboud et al. The main result of Abboud et al. [ARW17] is their SETH-hardness of the gap label cover problem which they refer to as the PCP-Vectors problem. In fact, PCP-Vectors is equivalent to MaxCover when $h = 2$ (i.e., the number of right super nodes is two). However, formulating the label cover problem as MaxCover instead of PCP-Vectors is beneficial

Figure 1: Overview of Our Framework. The first step is to reformulate each hypothesis in terms of hardness of a PSP problem, which is done in Section 4. Using the connection between SMP protocols and MaxCover outlined earlier (and formalized in Section 5), our task is now to devise SMP protocols with certain parameters for the corresponding communication problems; these are taken care of in Sections 6, 7 and 8. For completeness, the final reduction from MaxCover to k-DomSet which was shown in [CCK+17] is included in Appendix A.

for us, as our goal it to reduce to graph problems.

In their work, they merge the roles of the referee and the first player as it is necessary to achieve the goal of proving hardness of approximation for important problems in P (which are usually defined on one or two sets of vectors). However, by doing this the details of the proof become a little convoluted. On the contrary, our framework with the SMP model is arguably a cleaner framework to work with and it works well for our goal of proving hardness of approximation for parameterized problems.

Finally, we note that our observation that the hardness of approximating MaxCover can be obtained from any arbitrary hypothesis as long as there is an underlying product structure (as formalized via PSPs) is a new contribution of this paper.

3 Preliminaries and Backgrounds

We use standard graph terminology. Let G be any graph. The vertex and edge sets of G are denoted by $V(G)$ and $E(G)$, respectively. We say that a subset of vertices $S \subseteq V(G)$ *dominates* G if every vertex $v \in S \setminus V(G)$ has a neighbor in S, and we call S a *dominating set* of G. A k-*dominating set* of G is a dominating set of G with cardinality k. The *domination number*, denoted by (DomSet (G)) of G is the size of the smallest dominating set of G. A *clique H* in G is a complete subgraph of G, and we say that H is a k-*clique* if H contains k vertices. The *clique number* of G is the size of the largest clique in G. Sometime we abuse notation and call a subset of vertices $S \subseteq V(G)$, a clique, if S induces a complete subgraph in G.

3.1 Problem Definitions

Below are the list of problems considered in this paper.

- **Dominating Set.** In the *k-Dominating Set* problem (*k*-DomSet), we are given a graph G, and the goal is to decide whether G has a dominating set of size k. In the minimization version, called *Minimum Dominating Set* (DomSet, for short), the goal is to find a dominating set in G of minimum size.

- **Clique.** In the *k-Clique* problem (*k*-Clique), we are given a graph G, and the goal is to decide whether G has a clique of size k. In the maximization version, called *Maximum Clique* (Clique, for short), the goal is to find a clique in G of maximum size.

- ***k*-SAT.** In the *k-SAT* problem (*k*-CNF-SAT), we are given a CNF formula Φ with **at most** k literals in a clause and the goal is to decide whether Φ is satisfiable.

- ***k*-Sum.** In *k*-SUM, we are given k subsets $S_1, \ldots, S_k \subseteq ([-M, M] \cap \mathbb{Z})$ of integers between $-M$ and M (inclusive), and the goal is to determine whether there exist $x_1 \in S_1, \ldots, x_k \in S_k$ such that $x_1 + \cdots + x_k = 0$. That is, we wish to pick one integer from each subset so that they sum to zero.

In addition to the above problems, we devote the next section to define and discuss a variant of the *label cover* problem, namely MaxCover.

3.2 MaxCover – A Variant of Label Cover

We now define a variant of the label cover problem called MaxCover, which was introduced by Chalermsook et al. [CCK+17] to capture the parameterized inapproximability of *k*-Clique and *k*-DomSet. The approximation hardness of MaxCover will be the basis of our hardness results.

The input of MaxCover is a *label cover instance*; a label cover instance Γ consists of a bipartite graph $G = (U, W; E)$ such that U is partitioned into $U = U_1 \cup \cdots \cup U_\ell$ and W is partitioned into $W = W_1 \cup \cdots \cup W_h$. We sometimes refer to U_i's and W_j's as *left super-nodes* and *right super-nodes* of Γ, respectively. Another parameter we will be interested in is the maximum size of left super nodes, i.e., $\max_{i \in [\ell]} |U_i|$; we refer to this quantity as the *left alphabet size* of the instance.

A solution to MaxCover is called a *labeling*, which is a subset of vertices $S \subseteq W$ formed by picking a vertex w_j from each W_j (i.e., $|S \cap W_j| = 1$ for all $j \in [h]$). We say that a labeling S *covers* a left super-node U_i if there exists a vertex $u_i \in U_i$ such that u_i is a neighbor of every vertex in S. The goal in MaxCover is to find a labeling that covers the maximum fraction of left super-nodes. We abuse the notation MaxCover and also use it for the optimum as well, i.e.,

$$\mathsf{MaxCover}(\Gamma) = \frac{1}{\ell} \left(\max_{\text{labeling } S} |\{i \in [\ell] \mid U_i \text{ is covered by } S\}| \right).$$

The above terminologies for MaxCover are from [CCK+17]. Note, however, that our definitions are phrased somewhat different than theirs; in our definitions, the input graphs are the so-called

label-extended graphs whereas in their definitions, the input graphs are the *constraint graphs*. Nevertheless, it is not hard to see that the two versions are in fact equivalent. Another difference is that we use MaxCover to denote the *fraction* of left super-nodes covered by the optimal labeling whereas [CCK$^+$17] uses the notion for the *number* of covered left super-nodes. The former notation is somewhat more convenient for us as the value is between zero and one.

3.3   Inapproximability of DomSet from MaxCover

The relation between MaxCover and DomSet has been observed in literature. The *k-prover system* introduced by Feige in [Fei98] can be casted as a special case of MaxCover with *projection property*, and it has been shown that this proof system can be transformed into an instance of DomSet. We note, however, that the optimal value of the DomSet instance produced by Feige's *k*-prover system has size dependent on the number of left super-nodes rather than k, the number of right super-nodes. Recently, Chalermsook et al. [CCK$^+$17] observed that even without the projection property, the relation between MaxCover and DomSet still holds, and the value of the optimal solution can be reduced to k. This is stated formally below.

Theorem 3.1 (Reduction from MaxCover to DomSet [CCK$^+$17]). *There is an algorithm that, given a* MaxCover *instance* $\Gamma = (U = \bigcup_{j=1}^{q} U_j, W = \bigcup_{i=1}^{k} W_i, E)$, *outputs a* k-DomSet *instance G such that*

- *If* $\mathsf{MaxCover}(\Gamma) = 1$, *then* $\mathsf{DomSet}(G) = k$.

- *If* $\mathsf{MaxCover}(\Gamma) \leq \varepsilon$, *then* $\mathsf{DomSet}(G) \geq (1/\varepsilon)^{1/k} \cdot k$.

- $|V(G)| = |W| + \sum_{j \in [q]} k^{|U_j|}$.

- *The reduction runs in time* $O\left(|W| \left(\sum_{j \in [q]} k^{|U_j|}\right)\right)$.

For the sake of self-containedness, we provide the proof of Theorem 3.1 in Appendix A.

3.4   The Hypotheses

We now list and discuss the computational complexity hypotheses on which our results are based.

3.4.1 $\mathsf{W}[1] \neq \mathsf{FPT}$ **Hypothesis**

The first hypothesis is $\mathsf{W}[1] \neq \mathsf{FPT}$, which is one of the most popular hypotheses used in the area of parameterized complexity since many fundamental parameterized problems turn out to be $\mathsf{W}[1]$-hard. For the interest of space, we do not give the full definition of W-hierarchy; we refer the readers to standard textbook in the field (e.g. [DF13, CFK$^+$15]) for the definition and discussions regarding the hierarchy. Rather, since it is well-know that k-Clique is $\mathsf{W}[1]$-complete, we will use a more convenient formulation of $\mathsf{W}[1] \neq \mathsf{FPT}$, which simply states that k-Clique is not in FPT:

Hypothesis 1 ($\mathsf{W}[1] \neq \mathsf{FPT}$ Hypothesis). *For any computable function $T : \mathbb{N} \to \mathbb{N}$, no algorithm can solve k-Clique in $T(k) \cdot \mathrm{poly}(n)$ time where n denotes the number of vertices in the input graph.*

3.4.2 Exponential Time Hypothesis and Strong Exponential Time Hypothesis

Our second hypothesis is the Exponential Time Hypothesis (ETH), which can be stated as follows.

Hypothesis 2 (Exponential Time Hypothesis (ETH) [IP01, IPZ01, Tov84]). *There exists $\delta > 0$ such that no algorithm can solve 3-CNF-SAT in $O(2^{\delta n})$ time where n is the number of variables. Moreover, this holds even when restricted to formulae in which each variable appears in at most three clauses.*

Note that the original version of the hypothesis from [IP01] does not enforce the requirement that each variable appears in at most three clauses. To arrive at the above formulation, we first apply the Sparsification Lemma of [IPZ01], which implies that we can assume without loss of generality that the number of clauses m is $O(n)$. We then apply Tovey's reduction [Tov84] which produces a 3-CNF instance with at most $3m + n = O(n)$ variables and every variable occurs in at most three clauses. This means that the bounded occurrence restriction is also without loss of generality.

We will also use a stronger hypothesis called the Strong Exponential Time Hypothesis (SETH):

Hypothesis 3 (Strong Exponential Time Hypothesis (SETH) [IP01, IPZ01]). *For every $\varepsilon > 0$, there exists $k = k(\varepsilon) \in \mathbb{N}$ such that no algorithm can solve k-CNF-SAT in $O(2^{(1-\varepsilon)n})$ time where n is the number of variables. Moreover, this holds even when the number of clauses m is at most $c(\varepsilon) \cdot n$ where $c(\varepsilon)$ denotes a constant that depends only on ε.*

Again, we note that, in the original form [IP01], the bound on the number of clauses is not enforced. However, the Sparsification Lemma [IPZ01] allows us to do so without loss of generality.

3.4.3 k-SUM Hypothesis

Our final hypothesis is the k-SUM Hypothesis, which can be stated as follows.

Hypothesis 4 (k-SUM Hypothesis [AL13]). *For every integer $k \geq 3$ and every $\varepsilon > 0$, no $O(n^{\lceil k/2 \rceil - \varepsilon})$ time algorithm can solve k-SUM where n denotes the total number of input integers, i.e., $n = |S_1| + \cdots + |S_k|$. Moreover, this holds even when $M = n^{2k}$.*

The above hypothesis is a natural extension of the more well-known 3-SUM Hypothesis [GO95, Pat10], which states that 3-SUM cannot be solved in $O(n^{2-\varepsilon})$ time for any $\varepsilon > 0$. Moreover, the k-SUM Hypothesis is closely related to the question of whether SUBSET-SUM can be solved in $O(2^{(1/2-\varepsilon)n})$ time; if the answer to this question is negative, then k-SUM cannot be solved in $O(n^{k/2-\varepsilon})$ time for every $\varepsilon > 0, k \in \mathbb{N}$. We remark that, if one is only willing to assume this latter weaker lower bound of $O(n^{k/2-\varepsilon})$ instead of $O(n^{\lceil k/2 \rceil - \varepsilon})$, our reduction would give an $O(n^{k/2-\varepsilon})$ running time lower bound for approximating k-DomSet. Finally, we note that the assumption that $M = n^{2k}$ can be made without loss of generality since there is a randomized reduction from the general version of the problem (where M is, say, 2^n) to this version of the problem and this reduction can be derandomized under a certain circuit complexity assumption [ALW14].

3.5 Error-Correcting Codes

An error correcting code C over alphabet Σ is a function $C : \Sigma^m \to \Sigma^d$ where m and d are positive integers which are referred to as the *message length* and *block length* of C respectively. Intuitively, the function C encodes an original message of length m to an encoded message of length d. Since we will also deal with communication protocols, for which "message length" has another meaning, we will sometimes refer to the message length of codes as *code message length* whenever there is an ambiguity. The *rate* of a code $\rho(C)$ is defined as the ratio between its message length and its block length, i.e., $\rho(C) = m/d$. The *relative distance* of a code, denoted by $\delta(C)$, is defined as $\min_{x \neq y \in \Sigma^m} \delta(C(x), C(y))$ where $\delta(C(x), C(y))$ is the *relative Hamming distance* between $C(x)$ and $C(y)$, i.e., the fraction of coordinates on which $C(x)$ and $C(y)$ disagree.

3.5.1 Good Codes

In the construction of our communication protocol in Section 7, we require our codes to have constant rate and constant relative distance (referred to as *good codes*). It is not hard to see that random codes, ones where each codeword $C(x)$ is randomly selected from Σ^d independently from each other, satisfy these properties. For binary codes (i.e., $|\Sigma| = 2$), one can explicitly construct such codes using expander graphs (so called *Expander Codes* [SS96]); alternatively *Justesen Code* [Jus72] also have the same property (see Appendix E.1.2.5 from [Gol08] for an excellent exposition).

Fact 3.2. *For some absolute constant $\delta, \rho > 0$, there exists a family of codes $\mathcal{C} := \{C_m : \{0,1\}^m \to \{0,1\}^{d(m)}\}_{m \in \mathbb{N}}$ such that for every $m \in \mathbb{N}$ the rate of C_m is at least ρ and the relative distance of C_m is at least δ. Moreover, any codeword of C_m can be computed in time $\mathrm{poly}(m)$.*

3.5.2 Algebraic Geometric Codes

In the construction of our communication protocol in Section 6, we require our codes to have some special algebraic properties which have been shown to be present in algebraic geometric codes [GS96]. First, we will introduce a couple of additional definitions.

Definition 3.3 (Systematicity)**.** *Given $s \in \mathbb{N}$, a code $C : \Sigma^m \to \Sigma^d$ is s-systematic if there exists a size-s subset of $[d]$, which for convenience we identify with $[s]$, such that for every $x \in \Sigma^s$ there exists $w \in \Sigma^m$ in which $x = C(w)\mid_{[s]}$.*

Definition 3.4 (Degree-t Closure)**.** *Let Σ be a finite field. Given two codes $C : \Sigma^m \to \Sigma^d, C' : \Sigma^{m'} \to \Sigma^d$ and positive integer t, we say that C' is a degree-t closure of C if, for every $w_1, \ldots, w_r \in \Sigma^m$ and $P \in \mathbb{F}[X_1, \ldots, X_r]$ of total degree at most t, it holds that $\omega := P(C(w_1), \ldots, C(w_r))$ is in the range of C', where $\omega \in \Sigma^d$ is defined coordinate-wise by the equation $\omega_i := P(C(w_1)_i, \ldots, C(w_r)_i)$.*

Below we provide a self-contained statement of the result we rely on in Section 6; it follows from Theorem 7 of [SAK$^+$01], which gives an efficient construction of the algebraic geometric codes based on [GS96]'s explicit towers of function fields.

Theorem 3.5 ([GS96, SAK$^+$01]). *There are two polynomial functions $\hat{r}, \hat{q} : \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathbb{N}$ and any prime $q > \hat{q}(k)$, there are two code families $\mathcal{A} = \{A_n\}_{n \in \mathbb{N}}$, $\mathcal{B} = \{B_n\}_{n \in \mathbb{N}}$ such that the following holds for all $n \in \mathbb{N}$,*

- *A_n and B_n are n-systematic code with alphabet $\mathbb{F}_{q^2}$,*

- *A_n and B_n have block length less than $n \cdot \hat{r}(k)$.*

- *B_n has relative distance $\geq 1/2$,*

- *B_n is a degree-k closure of A_n, and,*

- *Any codeword in A_n or B_n can be computed in poly(n) time .*

We remark here that variants of the above theorem have previously found applications in the construction of special kinds of PCPs [BCG$^+$16, BKK$^+$16]. In these works, the theorems are also stated in a language similar to Theorem 3.5 above.

4 Product Space Problems and Popular Hypotheses

In this section, we define a class of computational problems called Product Space Problems (PSP). As the name suggests, a problem in this class is defined on a class of functions whose domain is a k-ary Cartesian Product, i.e., $f : X_1 \times \cdots \times X_k \to \{0, 1\}$. The input of the problem are subsets[12] $A_1 \subseteq X_1, \ldots, A_k \subseteq X_k$, and the goal is to determine whether there exists $(a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k$ such that $f(a_1, \ldots, a_k) = 1$. The size of the problem is determined by $\max_{i \in [k]} |A_i|$. A formal definition of PSP can be found below.

Definition 4.1 (Product Space Problem). *Let $m : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be any function and $\mathcal{F} := \{f_{N,k} : \{0, 1\}^{m(N,k) \times k} \to \{0, 1\}\}_{N,k \in \mathbb{N}}$ be a family of Boolean functions indexed by N and k. For each $k \in \mathbb{N}$, the* product space problem PSP$(k, \mathcal{F})$ *of order N is defined as follows: given k subsets $A_1, \ldots, A_k$ of $\{0, 1\}^{m(N,k)}$ each of cardinality at most N as input, determine if there exists $(a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k$ such that $f_{N,k}(a_1, \ldots, a_k) = 1$. We use the following shorthand PSP$(k, \mathcal{F}, N)$ to describe PSP$(k, \mathcal{F})$ of order N.*

In all the PSPs considered in this paper, the input length $m(N, k)$ is always at most $\text{poly}(k) \cdot \log N$ and $f_{N,k}$ is always computable in time $\text{poly}(m(N, k))$. In such a case, there is a trivial $N^{k + o_k(1)}$-time algorithm to solve PSP$(k, \mathcal{F}, N)$: enumerating all $(a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k$ and check whether $f_{N,k}(a_1, \ldots, a_k) = 1$. The rest of this section is devoted to rephrasing the hypotheses (SETH, ETH, W[1] $\neq$ FPT and the k-SUM Hypothesis) in terms of lower bounds for PSPs. The function families $\mathcal{F}$'s, and running time lower bounds will depend on the hypotheses. For example, SETH will corresponds to set disjointness whereas W[1] $\neq$ FPT will correspond to a generalization of equality called "multi-equality"; the former will give an $N^{k(1-o(1))}$ running time lower bound whereas the latter only rules out FPT time algorithms.

---

[12]Each A_i will be explicitly given as part of the input through the elements that it contains.

We would like to remark that the class of problems called 'locally-characterizable sets' introduced by Goldreich and Rothblum [GR18b] are closely related to PSPs. Elaborating, we may interpret locally-characterizable sets as the negation of PSPs, i.e., for any PSP$(k, \mathcal{F}, N)$, we may define the corresponding locally-characterizable set $\mathcal{S}$ as follows:

$$\mathcal{S} = \{(A_1, \ldots, A_k) \mid \text{for all } (a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k \text{ we have } f_{N,k}(a_1, \ldots, a_k) = 0\}.$$

Finally, we note that the class of problems called 'counting local patterns' introduced in [GR18a] are the counting counterpart of PSPs, i.e., for any instance $(A_1, \ldots, A_k)$ of PSP$(k, \mathcal{F}, N)$, we may define the corresponding counting local pattern solution to be the number of distinct $(a_1, \ldots, a_k) \in A_1 \times \cdots \times A_k$ such that $f_{N,k}(a_1, \ldots, a_k) = 1$.

4.1 k-SUM **Hypothesis**

To familiarize the readers with our notations, we will start with the k-SUM Hypothesis, which is readily in the PSP form. Namely, the functions in the family are the Sum-Zero functions that checks if the sum of k integers is zero:

Definition 4.2 (Sum-Zero). *Let $k, m \in \mathbb{N}$. SUMZERO$_{m,k} : (\{0,1\}^m)^k \to \{0,1\}$ is defined by*

$$\text{SUMZERO}_{m,k}(x_1, \ldots, x_k) = \begin{cases} 1 \text{ if } \sum_{i \in [k]} x_i = 0, \\ 0 \text{ otherwise,} \end{cases}$$

where we think of each x_i as a number in $[-2^{m-1}, 2^{m-1} - 1]$, and the addition is over $\mathbb{Z}$.

The function family $\mathcal{F}^{\text{SUMZERO}}$ can now be defined as follows.

Definition 4.3 (Sum-Zero Function Family). *Let $m : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be a function defined by $m(N, k) = 2k\lceil \log N \rceil$. $\mathcal{F}^{\text{SUMZERO}}$ is defined as $\{\text{SUMZERO}_{m(N,k),k}\}_{N \in \mathbb{N}, k \in N}$.*

The following proposition is immediate from the definition of the k-SUM Hypothesis.

Proposition 4.4. *Assuming the k-SUM Hypothesis, for every integer $k \geq 3$ and every $\varepsilon > 0$, no $O(N^{\lceil k/2 \rceil - \varepsilon})$-time algorithm can solve PSP$(k, \mathcal{F}^{\text{SUMZERO}}, N)$ for all $N \in \mathbb{N}$.*

4.2 **Set Disjointness and** SETH

We recall the k-way disjointness function, which has been studied extensively in literature (see, e.g., [LS09] and references therein).

Definition 4.5 (Set Disjointness). *Let $k, m \in \mathbb{N}$. DISJ$_{m,k} : (\{0,1\}^m)^k \to \{0,1\}$ is defined by*

$$\text{DISJ}_{m,k}(x_1, \ldots, x_k) = \neg \left(\bigvee_{i \in [m]} \left(\bigwedge_{j \in [k]} (x_j)_i \right) \right).$$

The function family $\mathcal{F}_c^{\text{DISJ}}$ can now be defined as follows.

Definition 4.6 (Set Disjointness Function Family). *For every $c \in \mathbb{N}$, let $m_c : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be a function defined by $m_c(N, k) = c\lceil k \log N\rceil$. $\mathcal{F}_c^{\text{DISJ}}$ is defined as $\{\text{DISJ}_{m_c(N,k),k}\}_{N \in \mathbb{N}, k \in N}$.*

We have the following proposition which follows easily from the definition of SETH and its well-known connection to the Orthogonal Vectors Hypothesis [Wil05].

Proposition 4.7. *Let $k \in \mathbb{N}$ such that $k > 1$. Assuming SETH, for every $\varepsilon > 0$ there exists $c := c_\varepsilon \in \mathbb{N}$ such that no $O(N^{k(1-\varepsilon)})$-time algorithm can solve $\text{PSP}(k, \mathcal{F}_c^{\text{DISJ}}, N)$ for all $N \in \mathbb{N}$.*

Proof. Fix $\varepsilon > 0$ and $k > 1$. By SETH, there exists $w := w(\varepsilon) \in \mathbb{N}$ and $c := c(\varepsilon) \in \mathbb{N}$ such that no algorithm can solve w-CNF-SAT in $O(2^{(1-\varepsilon)n})$ time where n is the number of variables and $m \le cn$ is the number of clauses. For every w-CNF-SAT formula ϕ, we will build $A_1^\phi, \ldots, A_k^\phi \subseteq \{0,1\}^m$ each of cardinality $N := 2^{n/k}$ such that there exists $(a_1, \ldots, a_k) \in A_1^\phi \times \cdots \times A_k^\phi$ such that $\text{DISJ}_{m,k}(a_1, \ldots, a_k) = 1$ if and only if ϕ is satisfiable. Thus, if there was an $O(N^{k(1-\varepsilon)})$-time algorithm that can solve $\text{PSP}(k, \mathcal{F}_c^{\text{DISJ}}, N)$ for all $N \in \mathbb{N}$, then it would violate SETH.

All that remains is to show the construction of $A_1^\phi, \ldots, A_k^\phi$ from ϕ. Fix $i \in [k]$. For every partial assignment σ to the variables $x_{(i-1)*(n/k)+1}, \ldots, x_{i*(n/k)}$ we build an m-bit vector $a_\sigma \in A_i^\phi$ as follows: $\forall j \in [m]$, we have $a_\sigma(j) = 0$ is σ satisfies the j^{th} clause, and $a_\sigma(j) = 1$ otherwise (i.e., the clause is not satisfied, or its satisfiability is indeterminate). It is easy to verify that there exists $(a_1, \ldots, a_k) \in A_1^\phi \times \cdots \times A_k^\phi$ such that $\text{DISJ}_{m,k}(a_1, \ldots, a_k) = 1$ if and only if ϕ is satisfiable. $\square$

We remark that we can prove a similar statement as that of Proposition 4.7 for ETH: assuming ETH, there exists k_0 such that for every $k > k_0$ there exists $c := c_{k_0} \in \mathbb{N}$ such that no $O(N^{o(k)})$-time algorithm can solve $\text{PSP}(k, \mathcal{F}_c^{\text{DISJ}}, N)$ for all $N \in \mathbb{N}$. However, instead of associating ETH with DISJ, we will associate with the Boolean function MULTEQ (which will be defined in the next subsection) and its corresponding PSP. This is because, associating ETH with MULTEQ provides a more elementary proof of Theorem 1.4 (in particular we will not need to use algebraic geometric codes – which are essentially inevitable if we associate ETH with DISJ).

4.3 $\text{W}[1] \ne \text{FPT}$ **Hypothesis and** ETH

Again, we recall the k-way Equality function which has been studied extensively in literature (see, e.g., [AMS12, ABC09, CRR14, CMY08, LV11, PVZ12] and references therein).

Definition 4.8 (Equality). *Let $k, m \in \mathbb{N}$. $\text{EQ}_{m,k} : (\{0,1\}^m)^k \to \{0,1\}$ is defined by*

$$\text{EQ}_{m,k}(x_1, \ldots, x_k) = \bigwedge_{i,j \in [k]} (x_i = x_j)$$

where $x_i = x_j$ is a shorthand for $\bigwedge_{p \in [m]} (x_i)_p = (x_j)_p$.

Unfortunately, the PSP associated with EQ is in fact not hard: given sets $A_1, \ldots, A_k$, it is easy to find whether they share an element by just sorting the combined list of $A_1 \cup \cdots \cup A_k$. Hence, we will need a generalization of the equality function to state our hard problem. Before we do so, let us first state an intermediate helper function, which is a variant of the usual equality function where some of the k inputs may be designed as "null" and the function only checks the equality over the non-null inputs. We call this function the Selective-Equality (SELEQ) function. For notational convenience, in the definition below, each of the k inputs is now viewed as $(x_i, y_i) \in \{0,1\}^{m-1} \times \{\bot, \top\}$; if $y_i = \bot$, then (x_i, y_i) represents the "null" input.

Definition 4.9 (Selective-Equality). *Let $k, m \in \mathbb{N}$. $\text{SELEQ}_{m,k} : (\{0,1\}^{m-1} \times \{\bot, \top\})^k \to \{0,1\}$ is defined by*

$$\text{SELEQ}_{m,k}((x_1, y_1), \ldots, (x_k, y_k)) = \bigwedge_{i,j \in [k]} \left((y_i = \bot) \vee (y_j = \bot) \vee (x_i = x_j) \right).$$

Next, we introduce the variant of EQ whose associated PSP is hard under $W[1] \neq$ FPT and ETH. In the settings of both Equality and Selective-Equality defined above, there is only one unknown that is given in each of the k inputs $a_1 \in A_1, \ldots, a_k \in A_k$ and the functions check whether they are equal. The following function, which we name Multi-Equality, is the t-unknown version of Selective-Equality. Specifically, the i^{th} part of the input is now a tuple $((x_{i,1}, y_{i,1}), \ldots, (x_{i,t}, y_{i,t}))$ where $x_{i,1}, \ldots, x_{i,t}$ are bit strings representing the supposed values of the t unknowns while, similar to Selective-Equality, each $y_{i,q} \in \{\bot, \top\}$ is a symbol indicating whether $(x_{i,q}, y_{i,q})$ is the "null" input. Below is the formal definition of MULTEQ; note that for convenience, we use $(x_{i,q}, y_{i,q})_{q \in [t]}$ as a shorthand for $((x_{i,1}, y_{i,1}), \ldots, (x_{i,t}, y_{i,t}))$, i.e., the i^{th} part of the input.

Definition 4.10 (Multi-Equality). *Let $k, t \in \mathbb{N}$ and let $m \in \mathbb{N}$ be any positive integer such that m is divisible by t. Let $m' = m/t$. $\text{MULTEQ}_{m,k,t} : ((\{0,1\}^{m'-1} \times \{\bot, \top\})^t)^k \to \{0,1\}$ is defined by*

$$\text{MULTEQ}_{m,k,t}((x_{1,q}, y_{1,q})_{q \in [t]}, \ldots, (x_{k,q}, y_{k,q})_{q \in [t]}) = \bigwedge_{q \in [t]} \text{SELEQ}_{m',k}((x_{1,q}, y_{1,q}), \ldots, (x_{k,q}, y_{k,q})).$$

Next, we define the family $\mathcal{F}^{\text{MULTEQ}}$; note that in the definition below, we simply choose $t(k)$, the number of unknowns, to be $k + \binom{k}{2} + \binom{k}{3}$. As we will see later, this is needed for ETH-hardness. For $W[1]$-hardness, it suffices to use a smaller number of variables. However, we choose to define $t(k)$ in such a way so that we can conveniently use one family for both ETH and $W[1]$-hardness.

Definition 4.11. *Let $t : \mathbb{N} \to \mathbb{N}$ be defined by $t(k) = k + \binom{k}{2} + \binom{k}{3}$. Let $m : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be defined by $m(N, k) = t(k)(1 + k\lceil \log N \rceil)$. We define $\mathcal{F}^{\text{MULTEQ}}$ as $\{\text{MULTEQ}_{m(N,k),k,t(k)}\}_{N \in \mathbb{N}, k \in \mathbb{N}}$.*

We next show a reduction from k-Clique to $\text{PSP}(k', \mathcal{F}^{\text{MULTEQ}})$ where $k' = \binom{k}{2}$. The overall idea of the reduction is simple. First, we associate the integers in $[k']$ naturally with the elements of $\binom{[k]}{2}$. We then create the sets $\left(A_{\{i,j\}} \right)_{\{i,j\} \subseteq [k], i \neq j}$ in such a way that each element of the set $A_{\{i,j\}}$ corresponds to picking an edge between the i-th and the j-th vertices in the supposed k-clique. Then, MULTEQ is used to check that these edges are consistent, i.e., that, for every $i \in [k]$, $a_{\{i,j\}}$ and $a_{\{i,j'\}}$ pick the same vertex to be the i^{th} vertex in the clique for all $j, j' \in [k] \setminus \{i\}$. This idea is formalized in the following proposition and its proof.

Proposition 4.12. *Let $k \in \mathbb{N}$ and $k' = \binom{k}{2}$. There exists a $\mathrm{poly}(N,k)$-time reduction from any instance (G,k) of* Clique *to an instance $(A_1, \ldots, A_{k'})$ of the $\mathsf{PSP}(k', \mathcal{F}^{\mathrm{MULTEQ}}, N')$ where N denotes the number of vertices of G and $N' = \binom{N}{2}$.*

Proof. Given a Clique instance[13] (G,k), the reduction proceeds as follows. For convenience, we assume that the vertex set $V(G)$ is $[N]$. Furthermore, we associate the elements of $[k']$ naturally with the elements of $\binom{[k]}{2}$. For the sake of conciseness, we sometimes abuse notation and think of $\{i,j\}$ as an ordered pair (i,j) where $i < j$. For every $\{i,j\} \in \binom{[k]}{2}$ such that $i < j$, the set $A_{\{i,j\}}$ contains one element $a_{\{i,j\}}^{\{u,v\}} = \left(a_{\{i,j\},1}^{\{u,v\}}, \ldots, a_{\{i,j\},t(k')}^{\{u,v\}} \right)$ for each edge $\{u,v\} \in E(G)$ such that $u < v$, where

$$a_{\{i,j\},q}^{\{u,v\}} = \begin{cases} (u, \top) & \text{if } q = i, \\ (v, \top) & \text{if } q = j, \\ (0, \bot) & \text{otherwise.} \end{cases}$$

Note that in the definition above, we view u, v and 0 as $\left(m(N',k')/t(k') - 1 \right)$-bit strings, where $m : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ is as in Definition 4.11. Also note that each set $A_{\{i,j\}}$ has size at most $\binom{N}{2} = N'$, meaning that $(A_{\{i,j\}})_{\{i,j\} \subseteq [k]}$ is indeed a valid instance of $\mathsf{PSP}(k', \mathcal{F}^{\mathrm{MULTEQ}}, N')$. For brevity, below we will use f as a shorthand for $\mathrm{MULTEQ}_{m(N',k'),k',t(k')}$.

($\Rightarrow$) Suppose that (G,k) is a YES instance for Clique, i.e., there exists a k-clique $\{u_1, \ldots, u_k\}$ in G. Assume without loss of generality that $u_1 < \cdots < u_k$. We claim that,

$$f\left(\left(a_{\{i,j\}}^{\{u_i,u_j\}} \right)_{i,j \in [k], i < j} \right) = 1.$$

To see that this is the case, observe that for every $q \in [t(k')]$ and for every $\{i,j\} \subseteq [k]$ such that $i < j$, we have either $a_{\{i,j\},q}^{\{u_i,u_j\}} = (0, \bot)$ or $a_{\{i,j\},q}^{\{u_i,u_j\}} = (u_q, \top)$. This means that, $\mathrm{SELEQ}\left(\left(a_{\{i,j\},q}^{\{u_i,u_j\}} \right)_{i,j \in [k], i < j} \right) = 1$ for every $q \in [t(k')]$.

($\Leftarrow$) Suppose that $(A_{\{i,j\}})_{\{i,j\} \subseteq [k]}$ is a YES instance for $\mathsf{PSP}(k', \mathcal{F}^{\mathrm{MULTEQ}}, N')$, i.e., there exists $a_{\{i,j\}}^* \in A_{\{i,j\}}$ for every $\{i,j\} \subseteq [k]$ such that $f((a_{\{i,j\}}^*)_{\{i,j\} \subseteq [k]}) = 1$. Suppose that $a_{\{i,j\}}^* = (x_{\{i,j\},1}^*, y_{\{i,j\},1}^*, \ldots, x_{\{i,j\},t(k')}^*, y_{\{i,j\},t(k')}^*)$. From this solution $\{a_{\{i,j\}}^*\}_{\{i,j\} \subseteq [k]}$, we can recover the k-clique as follows. For each $i \in [k]$, pick an arbitrary $j(i) \in [k]$ that is not equal to i. Let u_i be $x_{\{i,j\},i}^*$. We claim that $u_1, \ldots, u_k$ forms a k-clique in G. To show this, it suffices to argue that, for every distinct $i, i' \in [k]$, there is an edge between u_i and $u_{i'}$ in G. To see that this holds, consider $a_{\{i,i'\}}^*$. Since $y_{\{i,i'\},i}^* = y_{\{i,j(i)\},i}^* = \top$, we have $x_{\{i,i'\},i}^* = x_{\{i,j(i)\},i}^* = u_i$. Similarly, we have $x_{\{i,i'\},i'}^* = u_{i'}$. Since $a_{\{i,i'\}}^* \in A_{\{i,i'\}}$ and from how the set $A_{\{i,i'\}}$ is defined, we have $\{u_i, u_{i'}\} \in E(G)$, which concludes our proof. $\square$

Lemma 4.13. *Assuming $\mathsf{W}[1] \neq \mathsf{FPT}$, for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot \mathrm{poly}(N)$ time algorithm that can solve $\mathsf{PSP}(k, \mathcal{F}^{\mathrm{MULTEQ}}, N)$ for every $N, k \in \mathbb{N}$.*

---

[13]We assume without loss of generality that G does not contain any self-loop.

Proof. Suppose for the sake of contradiction that, for some computable function $T : \mathbb{N} \to \mathbb{N}$, there is a $T(k) \cdot \text{poly}(N)$ time algorithm $\mathcal{A}$ that can solve $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ for every $N, k \in \mathbb{N}$. We will show that this algorithm can also be used to solve k-Clique parameterized by k in FPT time.

Given an instance (G, k) of k-Clique, we first run the reduction from Proposition 4.12 to produce an instance $(A_1, \ldots, A_{k'})$ of $\text{PSP}(k', \mathcal{F}^{\text{MULTEQ}}, N')$ in $\text{poly}(N, k)$ time where $N = |V(G)|, N' = \binom{N}{2}$ and $k' = \binom{k}{2}$. We then run $\mathcal{A}$ on $(A_1, \ldots, A_{k'})$, which takes time $T(k') \cdot \text{poly}(N')$. This means that we can also solve our k-Clique instance (G, k) in time $\text{poly}(N, k) + T(k') \cdot \text{poly}(N') = \text{poly}(N, k) + T\left(\binom{k}{2}\right) \cdot \text{poly}(N)$, which is FPT time. Since k-Clique is $\text{W}[1]$-complete, this contradicts with $\text{W}[1] \neq \text{FPT}$. $\square$

Next, we will prove ETH-hardness of $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}})$. Specifically, we will reduce a 3-CNF-SAT instance ϕ where each variable appears in at most three clauses to an instance of $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ where $N = 2^{O(n/k)}$ and n denotes the number of variables in ϕ. The overall idea is to partition the set of clauses into k parts of equal size and use each element in A_j to represent a partial assignment that satisfies all the clauses in the j^{th} partition. This indeed means that each group has size $2^{O(n/k)}$ as intended. However, choosing the unknowns are not as straightforward as in the reduction from k-Clique above; in particular, if we view each variable by itself as an unknown, then we would have n unknowns, which is much more than the designated $t(k) = k + \binom{k}{2} + \binom{k}{3}$ unknowns! This is where we use the fact that each variable appears in at most three clauses: we group the variables of ϕ together based on which partitions they appear in and view each group as a single variable. Since each variable appears in at most three clauses, the number of ways they can appear in the k partitions is $k + \binom{k}{2} + \binom{k}{3}$ which is indeed equal to $t(k)$. The ideas are formalized below.

Proposition 4.14. *Let $k \in \mathbb{N}$. There exists a $\text{poly}(N, k)$-time reduction from any instance ϕ of 3-CNF-SAT such that each variable appears in at most three clauses in an instance $(A_1, \ldots, A_k)$ of the $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ where $N = 2^{3\lceil m/k \rceil}$ and m denotes the number of clauses in ϕ.*

Proof. Given a 3-CNF-SAT formula ϕ such that each variable appears in at most three clauses. Let the variable set of ϕ be $\mathcal{Z} = \{z_1, \ldots, z_n\}$ and the clauses of ϕ be $\mathcal{C} = \{C_1, \ldots, C_m\}$. Then for every $k \in \mathbb{N}$, we produce an instance $(A_1, \ldots, A_k)$ of $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ where $N = 2^{3\lceil m/k \rceil}$ as follows.

First, we partition the clause set $\mathcal{C}$ into k parts $\mathcal{C}_1, \ldots, \mathcal{C}_k$ each of size at most $\lceil m/k \rceil$. For each variable z_i, let S_i denote $\{j \in [k] \mid \exists C_h \in \mathcal{C}_j \text{ such that } z_i \in C_h \text{ or } \bar{z}_i \in C_h\}$. Since every z_i appears in at most three clauses, we have $S_i \in \binom{[k]}{\leq 3}$. For each $S \in \binom{[k]}{\leq 3}$, let $\nu(S)$ denote the set of all variables z_i's such that $S_i = S$ (i.e. S is exactly equal to the set of all partitions that z_i appears in). The general idea of the reduction is that we will view a partial assignment to the variables in $\nu(S)$ as an unknown for MULTEQ; let us call this unknown X_S (hence there are $k + \binom{k}{2} + \binom{k}{3} = t(k)$ unknowns). For each $j \in [k]$, A_j contains one element for each partial assignment to the variables that appear in the clauses in $\mathcal{C}_j$ and that satisfies all the clauses in $\mathcal{C}_j$. Such a partial assignment specifies $\left(1 + k + \binom{k}{2}\right)$ unknowns: all the X_S such that $j \in S$. The MULTEQ function is then used to check the consistency between the partial assignments to the variables from different A_j's.

To formalize this intuition, we first define more notations. Let $\tilde{m} = 3k\lceil m/k \rceil$. For every subsets $T \subseteq T' \subseteq \mathcal{Z}$ and every partial assignment $\alpha : T' \to \{0, 1\}$, the restriction of α to T, denoted by $\alpha|_T$ is the function from T to $\{0, 1\}$ where $\alpha|_T(z) = \alpha(z)$ for every $z \in T$. Furthermore, we define the

operator $\text{ext}(\alpha)$, which "extends" α to $\widetilde{m}$ bits, i.e., the i-th bit of $\text{ext}(\alpha)$ is $\alpha(z_i)$ if $z_i \in T$ and is zero otherwise. Finally, we use $\text{var}(\mathcal{C}_j)$ to denote the set of all variables that appear in at least one of the clauses from $\mathcal{C}_j$, i.e., $\text{var}(\mathcal{C}_j) = \bigcup_{C \in \mathcal{C}_j} \text{var}(C)$ where $\text{var}(C)$ denotes $\{z_i \in \mathcal{Z} \mid z_i \in C \text{ or } \bar{z}_i \in C\}$.

Now, since our $t(k)$ is exactly $\left| \binom{[k]}{\leq 3} \right|$, we can associate each element of $[t]$ with a subset $S \in \binom{[k]}{\leq 3}$. Specifically, for each partial assignment $\alpha : \text{var}(\mathcal{C}_j) \to \{0, 1\}$ such that α satisfies all the clauses in $\mathcal{C}_j$, the set A_j contains an element $a_j^\alpha = (a_{j,S}^\alpha)_{S \in \binom{[k]}{\leq 3}}$ where, for every $S \in \binom{[k]}{\leq 3}$,

$$a_{j,S}^\alpha = \begin{cases} \left(\text{ext}\left(\alpha|_{\nu(S)} \right), \top \right) & \text{if } j \in S, \\ (0^{\widetilde{m}}, \bot) & \text{otherwise.} \end{cases}$$

Fix $S \in \binom{[k]}{\leq 3}$. For every $j \in S$, observe that $\nu(S) \subseteq \text{var}(\mathcal{C}_j)$. Moreover, since each $\mathcal{C}_j$ contains at most $\lceil m/k \rceil$ clauses, there are at most $3\lceil m/k \rceil$ variables in $\text{var}(\mathcal{C}_j)$. This means that A_j has size at most $2^{3\lceil m/k \rceil}$. Hence, $(A_1, \ldots, A_k)$ is indeed a valid instance of PSP $(k, \mathcal{F}^{\text{MULTEQ}}, N)$ where $N = 2^{3\lceil m/k \rceil}$. For brevity, below we will use f as a shorthand for $\text{MULTEQ}_{m(N,k),k,t(k)}$.

($\Rightarrow$) Suppose that ϕ is satisfiable. Let $\alpha : \mathcal{C} \to \{0, 1\}$ be an assignment that satisfies all the clauses. Let $a_j^* = a_j^{\alpha|_{\text{var}(\mathcal{C}_j)}} \in A_j$ for every $j \in [k]$. Observe that, for every $S \in \binom{[k]}{\leq 3}$ and every $j \in [k]$, we either have $a_{j,S}^* = (0^{\widetilde{m}}, \bot)$ or $a_{j,S}^* = (\text{ext}(\alpha|_{\nu(S)}), \top)$. This indeed implies that $f(a_1^*, \ldots, a_k^*) = 1$.

($\Leftarrow$) Suppose that there exists $(a_1^{\alpha_1}, \ldots, a_k^{\alpha_k}) \in A_1 \times \cdots \times A_k$ such that $f(a_1^{\alpha_1}, \ldots, a_k^{\alpha_k}) = 1$. We construct an assignment $\alpha : \mathcal{Z} \to \{0, 1\}$ as follows. For each $i \in [n]$, pick an arbitrary $j(i) \in [k]$ such that $z_i \in \text{var}(\mathcal{C}_{j(i)})$ and let $\alpha(z_i) = \alpha_{j(i)}(z_i)$. We claim that α satisfies every clause. To see this, consider any clause $C \in \mathcal{C}$. Suppose that C is in the partition $\mathcal{C}_j$. It is easy to check that $f(a_1^{\alpha_1}, \ldots, a_k^{\alpha_k}) = 1$ implies that $\alpha|_{\text{var}(C)} = \alpha_j|_{\text{var}(C)}$. Since α_j is a partial assignment that satisfies C, α must also satisfy C. In other words, α satisfies all clauses of ϕ. $\qquad \square$

Lemma 4.15. *Assuming* ETH, *for any computable function* $T : \mathbb{N} \to \mathbb{N}$, *there is no* $T(k) \cdot N^{o(k)}$ *time algorithm that can solve* $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ *for every* $N, k \in \mathbb{N}$.

Proof. Let $\delta > 0$ be the constant in the running time lower bound in ETH. Suppose for the sake of contradiction that ETH holds but, for some function T, there is a $T(k) \cdot N^{o(k)}$ time algorithm $\mathcal{A}$ that can solve $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ for every $N, k \in \mathbb{N}$. Thus, there exists a sufficiently large k such that the running time of $\mathcal{A}$ for solving $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ is at most $O(N^{\delta k/10})$ for every $N \in \mathbb{N}$.

Given a 3-CNF formula ϕ such that each variable appears in at most three clauses. Let n, m denote the number of variables and the number of clauses of ϕ, respectively. We first run the reduction from Proposition 4.14 on ϕ with this value of k. This produces an instance $(A_1, \ldots, A_k)$ of PSP $(k, \mathcal{F}^{\text{MULTEQ}}, N)$ where $N = 2^{3\lceil m/k \rceil}$. Since each variable appears in at most three clauses, we have $m \leq 3n$, meaning that $N = O(2^{9n/k})$. By running $\mathcal{A}$ on this instance, we can decide whether ϕ is satisfiable in time $O(N^{\delta k/10}) = O(2^{0.9\delta n})$, contradicting ETH. $\qquad \square$

5 Communication Protocols and Reduction to Gap Label Cover

In this section, we first introduce a communication model for multiparty communication known in literature as the Simultaneous Message Passing model. Then, we introduce a notion of "efficient" communication protocols, and connect the existence of such protocols to a reduction from PSP to a gap version of MaxCover.

5.1 Efficient Protocols in Simultaneous Message Passing Model

The two-player Simultaneous Message Passing (SMP) model was introduced by Yao [Yao79] and has been extensively studied in literature [KN97]. In the multiparty setting, the SMP model is considered popularly with the number-on-forehead model, where each player can see the input of all the other players but not his own [CFL83, BGKL03]. In this paper, we consider the multiparty SMP model where the inputs are given as in the number-in-hand model (like in [FOZ16, WW15]).

Simultaneous Message Passing Model. Let $f : \{0,1\}^{m \times k} \to \{0,1\}$. In the k-player simultaneous message passing communication model, we have k players each with an input $x_i \in \{0,1\}^m$ and a referee who is given an advice $\mu \in \{0,1\}^*$ (at the same time when the players are given the input). The communication task is for the referee to determine if $f(x_1, \ldots, x_k) = 1$. The players are allowed to only send messages to the referee. In the randomized setting, we allow the players *and* the referee to jointly toss some random coins before sending messages, i.e., we allow public randomness.

Next, we introduce the notion of *efficient* protocols, which are in a nutshell one-round randomized protocols where the players and the referee are in a computationally bounded setting.

Efficient Protocols. Let π be a communication protocol for a problem in the SMP model. We say that π is a (w, r, ℓ, s)-efficient protocol if the following holds:

- The referee receives w bits of advice.

- The protocol is one-round with public randomness, i.e., the following actions happen sequentially:

    1. The players receive their inputs and the referee receives his advice.
    2. The players and the referee jointly toss r random coins.
    3. Each player on seeing the randomness (i.e. results of r coin tosses) deterministically sends an ℓ-bit message to the referee.
    4. Based on the advice, the randomness, and the total $\ell \cdot k$ bits sent from the players, the referee outputs accept or reject.

- The protocol has completeness 1 and soundness s, i.e.,

    - If $f(x_1, \ldots, x_k) = 1$, then there exists an advice on which the referee always accepts.
    - If $f(x_1, \ldots, x_k) = 0$, then, on any advice, the referee accepts with probability at most s.

- The players and the referee are computationally bounded, i.e., all of them perform all their computations in poly(m)-time.

The following proposition follows immediately from the definition of an efficient protocol and will be very useful in later sections for gap amplification.

Proposition 5.1. *Let $z \in \mathbb{N}$ and π be a communication protocol for a problem in the SMP model. Suppose π is a (w, r, ℓ, s)-efficient protocol. Then there exists a $(w, z \cdot r, z \cdot \ell, s^z)$-efficient protocol for the same problem.*

Proof. The proof follows by a simple repetition argument. More precisely, we repeat steps 2-4 in the protocol z times, each time using fresh randomness, but note that the z steps of drawing random coins can be clubbed into one step, and the decision by the referee can be reserved till the end of the entire protocol, wherein he accepts if and only if he would accept in each of the individual repetitions. □

5.2   Lower Bounds on Gap-MaxCover

The following theorem is the main conceptual contribution of the paper: we show below that the existence of efficient protocols can translate (exact) hardness of PSPs to hardness of approximating MaxCover.

Theorem 5.2. *Let $m : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be any function. Let $\mathcal{F} := \{f_{N,k} : \{0,1\}^{m(N,k) \times k} \to \{0,1\}\}_{N,k \in \mathbb{N}}$ be a family of Boolean functions indexed by N, k. Suppose there exists a (w, r, ℓ, s)-efficient protocol[14] for $f_{N,k}$ in the k-player SMP model for every $N, k \in \mathbb{N}$. Then, there is a reduction from any instance $(A_1, \ldots, A_k)$ of $\mathsf{PSP}(k, \mathcal{F}, N)$ to 2^w label cover instances $\{\Gamma^\mu\}_{\mu \in \{0,1\}^w}$ such that*

- *The running time of the reduction is $2^{w+r+\ell k} \mathsf{poly}(m(N, k))$.*

- *Each $\Gamma^\mu = (U^\mu = U_1^\mu \cup \cdots \cup U_q^\mu, W^\mu = W_1^\mu \cup \cdots \cup W_h^\mu; E^\mu)$ has the following parameters:*

  - *Γ^μ has at most Nk right nodes, i.e., $|W^\mu| \leq Nk$,*
  - *Γ^μ has k right super nodes, i.e., $h = k$,*
  - *Γ^μ has 2^r left super nodes, i.e., $q = 2^r$,*
  - *Γ^μ's left alphabet size is at most $2^{\ell k}$, i.e., $|U_1^\mu|, \ldots, |U_q^\mu| \leq 2^{\ell k}$.*

- *If $(A_1, \ldots, A_k)$ is a YES instance of $\mathsf{PSP}(k, \mathcal{F}, N)$, then $\mathsf{MaxCover}(\Gamma^\mu) = 1$ for some $\mu \in \{0,1\}^w$.*

- *If $(A_1, \ldots, A_k)$ is a NO instance of $\mathsf{PSP}(k, \mathcal{F}, N)$, then $\mathsf{MaxCover}(\Gamma^\mu) \leq s$ for every $\mu \in \{0,1\}^w$.*

Proof. Given a (w, r, ℓ, s)-efficient protocol π of $f_{N,k}$ and an instance $(A_1, \ldots, A_k)$ of $\mathsf{PSP}(k, \mathcal{F}, N)$, we will generate 2^w instances of MaxCover. Specifically, for each $\mu \in \{0,1\}^w$, we construct an instance $\Gamma^\mu = (U^\mu = U_1^\mu \cup \cdots \cup U_q^\mu, W^\mu = W_1^\mu \cup \cdots \cup W_h^\mu; E^\mu)$ of MaxCover as follows.

---

[14]w, r, ℓ and s can depend on N and k.

- Let $h = k$. For each $j \in [h]$, the right super-node W_j^μ contains one node for each $x_j \in A_j$.

- Let $q = 2^r$. For each random string $\gamma \in \{0,1\}^r$, the left-super node U_γ^μ contains one node for each of the possible accepting messages from the k players, i.e., each vertex in U_γ^μ corresponds to $(m_1, \ldots, m_k) \in (\{0,1\}^\ell)^k$ where in the protocol π the referee, on an advice μ and a random string γ, accepts if the messages he received from the k players are $m_1, \ldots, m_k$.

- We add an edge between $x_j \in W_j^\mu$ and $(m_1, \ldots, m_k) \in U_\gamma^\mu$ if m_j is equal to the message that j sends on an input x_j and a random string γ in the protocol π.

Observe that there is a bijection between labelings of Γ^μ and elements of $A_1 \times \cdots \times A_k$.

Now consider a labeling $S \subseteq W^\mu$ of Γ^μ and the corresponding $(x_1, \ldots, x_k) \in A_1 \times \cdots \times A_k$. For each random string $\gamma \in \{0,1\}^r$, observe that the referee accepts on an input $(x_1, \ldots, x_k)$, an advice μ, and a random string γ if and only if there is a vertex $u \in U_\gamma$ (corresponding to the messages sent by the players) that has an edge to every vertex in S. Therefore, the acceptance probability of the protocol on advice μ is the same as the fraction of left super-nodes covered by S. The completeness and soundness then easily follows:

Completeness. If there exists $(x_1, \ldots, x_k) \in A_1 \times \cdots \times A_k$ such that $f_{N,k}(x_1, \ldots, x_k) = 1$, then there is an advice $\mu \in \{0,1\}^w$ on which the referee always accepts for this input $(x_1, \ldots, x_k)$, meaning that the corresponding labeling covers every left super-node of Γ^μ, i.e., $\mathsf{MaxCover}(\Gamma^\mu) = 1$.

Soundness. If $f_{N,k}(x_1, \ldots, x_k) = 0$ for every $(x_1, \ldots, x_k) \in A_1 \times \cdots \times A_k$, then, for any advice $\mu \in \{0,1\}^w$, the referee accepts with probability at most s on every input $(x_1, \ldots, x_k) \in A_1 \times \cdots \times A_k$. This means that, for any $\mu \in \{0,1\}^w$, no labeling covers more than s fraction of left the super-nodes. In other words, $\mathsf{MaxCover}(\Gamma^\mu) \le s$ for all $\mu \in \{0,1\}^w$. $\square$

For the rest of this subsection, we will use the following shorthand. Let $\Gamma = (U = U_1 \cup \cdots \cup U_q, W = W_1 \cup \cdots \cup W_h; E)$ be a label cover instance, and we use the shorthand $\Gamma(N, k, r, \ell)$ to say that the label cover instance has the following parameters:

- Γ has at most Nk right nodes, i.e., $|W| \le Nk$,

- Γ has k right super nodes, i.e., $h = k$,

- Γ has 2^r left super nodes, i.e., $q = 2^r$,

- Γ has left alphabet size of at most $2^{\ell k}$, i.e., $|U_1|, \ldots, |U_q| \le 2^{\ell k}$.

The rest of this section is devoted to combining Theorem 5.2 with the results in Section 4 to obtain conditional hardness for the gap-MaxCover problem, assuming that we have efficient protocols with certain parameters. These protocols will be devised in the three subsequent sections.

Understanding the Parameters. Before we state the exact dependency of parameters, let us first discuss some intuition behind it. First of all, if we start with an instance of $\mathsf{PSP}(k, \mathcal{F}, N)$, Theorem 5.2 will produce 2^w instances of $\Gamma(N, k, r, \ell)$. Roughly speaking, since we want the lower bounds from PSP to translate to MaxCover, we would like the number of instances to be $N^{o(1)}$, meaning that we want $w = o(\log N)$. Recall that in all function families we consider $m = \Theta_k(\log N)$. Hence, this requirement is the same as $w = o_k(m)$. Moreover, we would like the instance size of $\Gamma(N, k, r, \ell)$ to also be $O_k(N)$, meaning that the number of left vertices, $2^{r+\ell k}$ has to be $O_k(N)$. Thus, it suffices to have a protocol where $r + \ell k = o_k(m)$.

If we additionally want the hardness to translate also to k-DomSet, the parameter dependencies become more subtle. Specifically, applying Theorem 3.1 to the MaxCover instances results in a blow-up of $\sum_{j \in [q]} k^{|U_j|} = 2^r \cdot k^{2^{\ell k}} = 2^{r + (\log k) \cdot 2^{\ell k}}$. We also want this to be at most $N^{o(1)}$, meaning that we need $r + (\log k) \cdot 2^{\ell k} = o(\log N) = o_k(m)$. In other words, it suffices for us to require that $\ell k < (\log m)/\beta$ for some constant $\beta > 1$. The exact parameter dependencies are formalized below.

SETH

Corollary 5.3. *For any $c \in \mathbb{N}$, let $\mathcal{F}_c^{\mathrm{DISJ}}$ be the family of Boolean functions as defined in Definition 4.6. For every $\delta > 0$, suppose there exists a (w, r, ℓ, s)-efficient protocol for $\mathrm{DISJ}_{m,k}$ in the k-player SMP model for every $k \in \mathbb{N}$ and every $m \in \mathbb{N}$, such that $w \leq \delta m$ and $r + \ell k = o_k(m)$. Then, assuming SETH, for every $\varepsilon > 0$ and integer $k > 1$, no $O(N^{k(1-\varepsilon)})$-time algorithm can distinguish between $\mathsf{MaxCover}(\Gamma) = 1$ and $\mathsf{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N \in \mathbb{N}$. Moreover, if $\ell < (\log m)/\beta \cdot k$ for some constant $\beta > 1$, then assuming SETH, for every $\varepsilon > 0$ and integer $k > 1$, no $O(N^{k(1-\varepsilon)})$-time algorithm can distinguish between $\mathsf{DomSet}(G) = k$ and $\mathsf{DomSet}(G) \geq \left(\frac{1}{s}\right)^{\frac{1}{k}} \cdot k$ for any graph G with at most $O_k(N)$ vertices, for all $N \in \mathbb{N}$.*

Proof. The proof of the first part of the theorem statement is by contradiction. Suppose there is an $O(N^{k(1-\varepsilon)})$-time algorithm $\mathcal{A}$ for some fixed constant $\varepsilon > 0$ and integer $k > 1$ which can distinguish between $\mathsf{MaxCover}(\Gamma) = 1$ and $\mathsf{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N \in \mathbb{N}$. From Proposition 4.7, we have that there exists $c_\varepsilon \in \mathbb{N}$ such that no $O(N^{k(1-\varepsilon/2)})$-time algorithm can solve $\mathsf{PSP}(k, \mathcal{F}_{c_\varepsilon}^{\mathrm{DISJ}}, N)$ for all $N \in \mathbb{N}$. Fix $\delta = \varepsilon/3c_\varepsilon$. Next, by considering Theorem 5.2 for the case of (w, r, ℓ, s)-efficient protocols, we have that there are 2^w label cover instances $\{\Gamma^\mu\}_{\mu \in \{0,1\}^w}$ which can be constructed in $2^{\delta m(1+o_k(1))}$ time. Note that $2^{\delta m(1+o_k(1))} = N^{k\varepsilon/3(1+o_k(1))}$ by our choice of δ. Thus, we can run $\mathcal{A}$ on each Γ^μ and solve $\mathsf{PSP}(k, \mathcal{F}_c^{\mathrm{DISJ}}, N)$ for all $N, k \in \mathbb{N}$ in time less than $N^{k(1-\varepsilon/2)}$. This contradicts Proposition 4.7.

To prove the second part of the theorem statement, we apply the reduction described in Theorem 3.1 and note that $2^r = N^{o(1)}$ and $k^{2^{\ell k}} = 2^{(\log_2 k) \cdot (m(N,k))^{1/\beta}} = N^{o(1)}$. $\qquad\square$

The proof of Theorem 1.5 follows by plugging in the parameters of the protocol described in Corollary 6.2 to the above corollary.

ETH

Corollary 5.4. *Let $\mathcal{F}^{\text{MULTEQ}}$ be the family of Boolean functions as defined in Definition 4.11. Suppose there exists a (w, r, ℓ, s)-efficient protocol for $\text{MULTEQ}_{m,k,t}$ in the k-player SMP model for every $k, t, m \in \mathbb{N}$ such that $w + r + \ell k = o_k(m)$. Then, assuming ETH, for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot N^{o(k)}$ time algorithm that can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N, k \in \mathbb{N}$. Moreover, if $\ell < (\log m)/\beta \cdot k$ for some constant $\beta > 1$, then assuming ETH, for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot N^{o(k)}$ time algorithm that can distinguish between $\text{DomSet}(G) = k$ and $\text{DomSet}(G) \geq \left(\frac{1}{s}\right)^{\frac{1}{k}} \cdot k$ for any graph G with at most $O_k(N)$ vertices, for all $N, k \in \mathbb{N}$.*

Proof. The proof of the first part of the theorem statement is by contradiction. Suppose there is an algorithm $\mathcal{A}$ running in time $\widetilde{T}(k) \cdot N^{o(k)}$ for some computable function $\widetilde{T} : \mathbb{N} \to \mathbb{N}$ that can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N, k \in \mathbb{N}$. From Lemma 4.15, we have that for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot N^{o(k)}$ time algorithm that can solve $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ for every $N, k \in \mathbb{N}$. Next, by considering Theorem 5.2 for the case of (w, r, ℓ, s)-efficient protocols, we have that there are 2^w label cover instances $\{\Gamma^\mu\}_{\mu \in \{0,1\}^w}$ which can be constructed in $2^{o_k(m)}$ time. Note that $2^{o_k(m)} = O_k(N^{o(1)})$ by the choice of $m(N, k)$ in Definition 4.11. Thus, we can run $\mathcal{A}$ on each Γ^μ and solve $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ for all $N, k \in \mathbb{N}$ in time $\widetilde{T}(k) \cdot N^{o(k)}$. This contradicts Lemma 4.15.

To prove the second part of the theorem statement, we apply the reduction described in Theorem 3.1 and note that $2^r = N^{o(1)}$ and $k2^{\ell k} = 2^{(m(N,k))^{1/\beta} \cdot \log_2 k} = N^{o(1)}$. $\qquad \square$

The proof of Theorem 1.4 follows by plugging in the parameters of the protocol described in Corollary 7.2 to the above corollary.

W[1] $\neq$ FPT

Corollary 5.5. *Let $\mathcal{F}^{\text{MULTEQ}}$ be the family of Boolean functions as defined in Definition 4.11. Suppose there exists a (w, r, ℓ, s)-efficient protocol for $\text{MULTEQ}_{m,k,t}$ in the k-player SMP model for every $k, t, m \in \mathbb{N}$ such that $w + r + \ell k < m/tk$. Then, assuming $\text{W}[1] \neq \text{FPT}$, for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot \text{poly}(N)$-time algorithm that can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N, k \in \mathbb{N}$. Moreover, if $r < m/2tk$ and $\ell < (\log m)/\beta \cdot k$ for some constant $\beta > 1$, then assuming $\text{W}[1] \neq \text{FPT}$, for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot \text{poly}(N)$-time algorithm that can distinguish between $\text{DomSet}(G) = k$ and $\text{DomSet}(G) \geq \left(\frac{1}{s}\right)^{\frac{1}{k}} \cdot k$ for any graph G with at most $O_k(N)$ vertices, for all $N, k \in \mathbb{N}$.*

Proof. The proof of the first part of the theorem statement is by contradiction. Suppose there is an algorithm $\mathcal{A}$ running in time $\widetilde{T}(k) \cdot \text{poly}(N)$ for some computable function $\widetilde{T} : \mathbb{N} \to \mathbb{N}$ that can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N, k \in \mathbb{N}$. From Lemma 4.13, we have that for any computable function $T : \mathbb{N} \to \mathbb{N}$, there is no $T(k) \cdot \text{poly}(N)$ time algorithm that can solve $\text{PSP}(k, \mathcal{F}^{\text{MULTEQ}}, N)$ for every $N, k \in \mathbb{N}$. Next, by considering Theorem 5.2 for the case of (w, r, ℓ, s)-efficient protocols, we have that there

are 2^w label cover instances $\{\Gamma^\mu\}_{\mu \in \{0,1\}^w}$ which can be constructed in $2^{m(N,k)/k \cdot t(k)} \cdot \text{poly}(m(N,k))$ time. Note that $2^{m(N,k)/k \cdot t(k)} = O(N)$ and $\text{poly}(m(N,k)) = N^{o(1)}$ by the choice of $m(N,k)$ in Definition 4.11. Thus, we can run $\mathcal{A}$ on each Γ^μ and solve $\text{PSP}(k, \mathcal{F}^{\text{MultEq}}, N)$ for all $N, k \in \mathbb{N}$ in time less than $\widetilde{T}(k) \cdot \text{poly}(N)$. This contradicts Lemma 4.13.

To prove the second part of the theorem statement, we apply the reduction described in Theorem 3.1 and note that $2^r = O(\sqrt{N})$ and $k^{2^{\ell k}} = 2^{(m(N,k))^{1/\beta} \cdot \log_2 k} = N^{o(1)}$. $\qquad\square$

k-SUM **Hypothesis**

Corollary 5.6. *Let $\mathcal{F}^{\text{SumZero}}$ be the family of Boolean functions as defined in Definition 4.3. Suppose there exists a (w, r, ℓ, s)-efficient protocol for $\text{SumZero}_{m,k}$ in the k-player SMP model for every $m, k \in \mathbb{N}$, such that $w + r + \ell k = o_k(m)$. Then assuming the k-SUM Hypothesis, for every integer $k \geq 3$ and every $\varepsilon > 0$, no $O(N^{\lceil k/2 \rceil - \varepsilon})$-time algorithm can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N \in \mathbb{N}$. Moreover, if $\ell < {}^{(\log m)}/_{\beta \cdot k}$ for some constant $\beta > 1$, then assuming the k-SUM Hypothesis, for every $\varepsilon > 0$ no $O(N^{\lceil k/2 \rceil - \varepsilon})$-time algorithm can distinguish between $\text{DomSet}(G) = k$ and $\text{DomSet}(G) \geq \left(\frac{1}{s}\right)^{\frac{1}{k}} \cdot k$ for any graph G with at most $O_k(N)$ vertices for all $N \in \mathbb{N}$.*

Proof. The proof of the first part of the theorem statement is by contradiction. Suppose there is an algorithm $\mathcal{A}$ running in time $O(N^{\lceil k/2 \rceil - \varepsilon})$ for some fixed constant $\varepsilon > 0$ and some integer $k \geq 3$ that can distinguish between $\text{MaxCover}(\Gamma) = 1$ and $\text{MaxCover}(\Gamma) \leq s$ for any label cover instance $\Gamma(N, k, r, \ell)$ for all $N \in \mathbb{N}$. From Proposition 4.4, we have that no $O(N^{\lceil k/2 \rceil - \varepsilon/2})$-time algorithm can solve $\text{PSP}(k, \mathcal{F}^{\text{SumZero}}, N)$ for all $N \in \mathbb{N}$. Next, by considering Theorem 5.2 for the case of (w, r, ℓ, s)-efficient protocols, we have that there are 2^w label cover instances $\{\Gamma^\mu\}_{\mu \in \{0,1\}^w}$ which can be constructed in $2^{o_k(m)}$ time. Note that $2^{o_k(m)} = O_k(N^{o(1)})$ by the choice of $m(N,k)$ in Definition 4.3. Thus, we can run $\mathcal{A}$ on each Γ^μ and solve $\text{PSP}(k, \mathcal{F}^{\text{SumZero}}, N)$ for all $N \in \mathbb{N}$ in time $O(N^{\lceil k/2 \rceil - \varepsilon})$. This contradicts Proposition 4.4.

To prove the second part of the theorem statement, we apply the reduction described in Theorem 3.1 and note that $2^r = N^{o(1)}$ and $k^{2^{\ell k}} = 2^{(m(N,k))^{1/\beta} \cdot \log_2 k} = N^{o(1)}$. $\qquad\square$

The proof of Theorem 1.6 follows by plugging in the parameters of the protocol described in Corollary 8.6 to the above corollary.

6 An Efficient Protocol for Set Disjointness

Set Disjointness has been extensively studied primarily in the two-player setting (i.e., $k = 2$). In that setting, we know that the randomized communication complexity is $\Omega(m)$ [KS92, Raz92, BJKS04], where m is the input size of each player. Surprisingly, [AW09] showed that the MA-complexity of two-player set disjointness is $\tilde{O}(\sqrt{m})$. Their protocol was indeed an

$(\tilde{O}(\sqrt{m}), O(\log m), \tilde{O}(\sqrt{m}), 1/2)$-efficient protocol for the case when $k = 2$. Recently, [ARW17] improved (in terms of the message size) the protocol to be an $(m/\log m, O(\log m), O((\log m)^3), 1/2)$-efficient protocol for the case when $k = 2$. Both these results can be extended naturally for all $k > 1$, to give an $(mk/\log m, O_k(\log m), O_k((\log m)^3), 1/2)$-efficient protocol. However, this does not suffice for proving Theorem 1.5 since we need a (w, r, ℓ, s)-efficient protocol with $w = o(m)$ and $\ell = o(\log m)$. Fortunately, Rubinstein [Rub18] recently showed that the exact framework of the MA-protocols as in [AW09, ARW17] but with the use of algebraic geometric codes instead of Reed Muller or Reed Solomon codes gives the desired parameters in the two-player case. Below we naturally extend Rubinstein's protocol to the k-player setting. This extension was suggested to us by Rubinstein [Rub17].

Theorem 6.1. *There is a polynomial function $\hat{\ell} : \mathbb{N} \times \mathbb{N} \to [1, \infty)$ such that for every $k \in \mathbb{N}$ and every $\alpha \in \mathbb{N}$, there is a protocol for k-player $\mathrm{DISJ}_{m,k}$ in the SMP model which is an $(m/\alpha, \log_2 m, \hat{\ell}(k, \alpha), 1/2)$-efficient protocol, where each player is given m bits as input, and the referee is given at most m/α bits of advice.*

Proof. Fix $k, m, \alpha \in \mathbb{N}$. Let q be the smallest prime greater than $\hat{q}(k)$ such that $q > (2\alpha\hat{r}(k))^2$, where the functions $\hat{r}$ and $\hat{q}$ are as defined in Theorem 3.5. Let $\mathcal{G} = \mathbb{F}_{q^2}$. Let $T = 2\alpha\hat{r}(k)\log_2 q$.

We associate $[m]$ with $[T] \times [m/T]$ and write the input $\mathbf{x}_j \in \{0,1\}^m$ as vectors $\mathbf{x}_j^1, \ldots, \mathbf{x}_j^T$ where $\mathbf{x}_j^t \in \{0,1\}^{m/T}$. For every $j \in [k]$, Player j computes $A_{m/T}(\mathbf{x}_j^t)$ for every $t \in [T]$. We denote the block length of $A_{m/T}$ by d. From the systematicity guaranteed by Theorem 3.5 we have that $A_{m/T}(\mathbf{x}_j^t) \mid_{[m/T]} = \mathbf{x}_j^t$. Also, notice that for all $t \in [T]$, we have $\bigwedge_{j \in [k]} \mathbf{x}_j^t = 0^{m/T}$ if and only if $\prod_{j \in [k]} A_{m/T}(\mathbf{x}_j^t) = 0^{m/T}$.

With the above observation in mind, we define the marginal sum $\Gamma \in \mathcal{G}^d$ as follows:

$$\forall i \in [d], \ \Gamma_i = \sum_{t \in [T]} \prod_{j \in [k]} A_{m/T}(\mathbf{x}_j^t)_i.$$

Again, notice that for all $t \in [T]$, we have $\bigwedge_{j \in [k]} \mathbf{x}_j^t = 0^{m/T}$ if and only if $\Gamma_i = 0$ for all $i \in [m/T]$. This follows from the following:

- For all $i \in [m/T]$, we have $A_{m/T}(\mathbf{x}_j^t)_i \in \{0,1\}$ and thus $\prod_{j \in [k]} A_{m/T}(\mathbf{x}_j^t)_i \in \{0,1\}$.

- The characteristic of $\mathcal{G}$ being greater than $(2\alpha\hat{r}(k)) \cdot \sqrt{q} \geq (2\alpha\hat{r}(k)) \cdot \log_2 q = T$.

More importantly, we remark that Γ is a codeword[15] in $B_{m/T}$. This follows because $B_{m/T}$ is a degree k closure code of $A_{m/T}$. To see this, in Definition 3.4, set $t = k$, $r = k \cdot T$, and $P[x_{1,1}, \ldots, x_{k,T}] = \sum_{i \in [T]} \prod_{j \in [k]} x_{i,j}$.

---

[15]We would like to remark that we use the multiplicity of Algebraic Geometric codes to find a non-trivial advice. On a related note, Meir [Mei13] had previously shown that error correcting codes with the multiplicity property suffice to show the IP theorem (i.e., the IP=PSPACE result).

The protocol

1. Merlin sends the referee Φ which is allegedly equal to the marginal sums codeword Γ defined above.

2. All players jointly draw $r \in [d]$ uniformly at random.

3. For every $j \in \{1, \ldots, k\}$, Player j sends to the referee $A_{m/T}(\mathbf{x}_j^t)_r$, $\forall t \in [T]$.

4. The referee accepts if and only if both of the following hold:

$$\forall i \in [m/T], \quad \Phi_i = 0 \tag{1}$$

$$\Phi_r = \sum_{t \in [T]} \prod_{j \in [k]} A_{m/T}(\mathbf{x}_j^t)_r. \tag{2}$$

Analysis

Advice Length. To send the advice, Merlin only needs to send a codeword in $B_{m/T}$ to the verifier. This means that the advice length (in bits) is no more than $\log_2 q^2$ times the block length, which is $d \leq (m/T)\hat{r}(k) = m/2\alpha \log_2 q$ where the equality comes from our choice of T and the inequality from Theorem 3.5.

Message Length. Each player sends T elements of $\mathcal{G}$. Hence, the message length is $T \log_2 q^2 \leq \alpha \hat{r}(k)(2 \log_2 q)^2$. Recall that q can be upper bounded by a polynomial in k and α. Hence, the message length is upper bounded entirely as a polynomial in k and α as desired.

Randomness. The number of coin tosses is $\log_2(d) \leq \log_2\left(m\hat{r}(k)/T\right) = \log_2\left(m/2\alpha \log_2 q\right) < \log_2 m$.

Completeness. If the k sets are disjoint, Merlin can send the true Γ, and the verifier always accepts.

Soundness. If the k sets are not disjoint and Φ is actually Γ, then (1) is false and the verifier always rejects. On the other hand, if $\Phi \neq \Gamma$, then, since both are codewords of $B_{m/T}$, from Theorem 3.5 their relative distance must be at least $1/2$. As a result, with probability at least $1/2$, $\Phi_r \neq \Gamma_r$. Since the right hand side of (2) is simply Γ_r, the verifier will reject for such r. Hence, the rejection probability is at least $1/2$. $\qquad\square$

The following corollary follows immediately by applying Proposition 5.1 with $z = (\log_2 m)/2k \cdot \ell(k,\alpha)$ to the above theorem.

Corollary 6.2. *There is a polynomial function $\hat{s} : \mathbb{N} \times \mathbb{N} \rightarrow [1, \infty)$ such that for every $k \in \mathbb{N}$ and every $\alpha \in \mathbb{N}$ there is a protocol for k-player $\mathrm{DISJ}_{m,k}$ in the SMP model which is an $\left(m/\alpha, O\left((\log_2 m)^2\right), (\log_2 m)/2k, (1/m)^{1/\hat{s}(k,\alpha)}\right)$-efficient protocol, where each player is given m bits as input, and the referee is given at most m/α bits of advice.*

7 An Efficient Protocol for Multi-Equality

Equality has been extensively studied, primarily in the two-player setting (i.e., $k = 2$). In that setting, when public randomness is allowed, we know that the randomized communication complexity is $O(1)$ [Yao79, KN97], and the protocols can be naturally extended to the k-player SMP model that yields a randomized communication complexity of $O(k)$. There are many protocols which achieve this complexity bound but for the purposes of proving Theorems 1.3 and 1.4, we will use the protocol where the players encode their input using a fixed good binary code and then send a jointly agreed random location of the encoded input to the referee who checks if all the messages he received are equal. Below we extend that protocol for Multi-Equality.

Theorem 7.1. *For some absolute constant $\delta > 0$, for every $t, k \in \mathbb{N}$ and every $m \in \mathbb{N}$ such that m is divisible by t, there is a $(0, \log m + O(1), 2t, 1 - \delta)$-efficient protocol for $\text{MULTEQ}_{m,k,t}$ in the k-player SMP model.*

Proof. Let $\mathcal{C} = \{C_m : \{0,1\}^m \to \{0,1\}^{d(m)}\}_{m \in \mathbb{N}}$ be the family of good codes with rate at least ρ and relative distance at least δ as guaranteed by Fact 3.2. Fix $m, k, t \in \mathbb{N}$ as in the theorem statement. Let $m' := m/t$.

The protocol

1. All players jointly draw $i^\star \in [d(m')]$ uniformly at random.

2. If player j's input is $x_j = (x_{j,1}, y_{j,1}, \ldots, x_{j,t}, y_{j,t})$, then he sends $(C(x_{j,1})_{i^*}, y_{j,1}, \ldots, C(x_{j,t})_{i^*}, y_{j,t})$ to the referee.

3. The referee accepts if and only if the following holds:

$$\text{MULTEQ}_{2t,k,t}((C(x_{1,1})_{i^*}, y_{1,1}, \ldots, C(x_{1,t})_{i^*}, y_{1,t}), \ldots, (C(x_{k,1})_{i^*}, y_{k,1}, \ldots, C(x_{k,t})_{i^*}, y_{k,t})) = 1.$$

Analysis

Parameters of the Protocol. In the first step of the protocol all the players jointly draw $i^\star \in [d(m')]$ uniformly, which requires $\lceil \log d(m') \rceil \leq \log m' + \log(1/\rho) \leq \log m + O(1)$ random bits. Then, for every $j \in [k]$, player j sends the referee $2t$ bits. Finally, since the code $C_{m'}$ is efficient, it is easy to see that the players and referee run in $\text{poly}(m)$-time.

Completeness. If $\text{MULTEQ}_{m,k,t}(x_1, \ldots, x_k) = 1$, then, for every $q \in [t]$ and $i, j \in [k]$, we have $(y_{i,q} = \bot) \vee (y_{j,q} = \bot) \vee (x_{i,q} = x_{j,q}) = 1$. This implies that, for every $q \in [t], i, j \in [k]$ and $i^* \in [d(m)]$, we have $(y_{i,q} = \bot) \vee (y_{j,q} = \bot) \vee (C(x_{i,q})_{i^*} = C(x_{j,q})_{i^*})$. In other words, $\text{MULTEQ}_{2t,k}((C(x_{1,1})_{i^*}, y_{1,1}, \ldots, C(x_{1,t})_{i^*}, y_{1,t}), \ldots, (C(x_{k,1})_{i^*}, y_{k,1}, \ldots, C(x_{k,t})_{i^*}, y_{k,t})) = 1$ for every $i^* \in [d(m)]$, meaning that the referee always accepts.

Soundness. Suppose that $\text{MULTEQ}_{m,k,t}(x_1, \ldots, x_k) = 0$. Then, there exists some $q \in [t]$ and $i, j \in [k]$ such that $y_{i,q} = \top, y_{j,q} = \top$ and $x_{i,q} \neq x_{j,q}$. Since the code C has relative distance at least δ, $C(x_{i,q})$ and $C(x_{j,q})$ must differ on at least δ fraction

of the coordinates. When the randomly selected i^* is such a coordinate, we have that $\text{MULTEQ}_{2t,k}((C(x_{1,1})_{i^*}, y_{1,1}, \ldots, C(x_{1,t})_{i^*}, y_{1,t}), \ldots, (C(x_{k,1})_{i^*}, y_{k,1}, \ldots, C(x_{k,t})_{i^*}, y_{k,t})) = 0$, i.e., the referee rejects. Hence, the referee rejects with probability at least δ. $\qquad\square$

The following corollary follows immediately by applying Proposition 5.1 to the above theorem with $z = \log_2 m / 4kt$.

Corollary 7.2. *For every $t, k \in \mathbb{N}$ and every $m \in \mathbb{N}$ such that m is divisible by t, there is a $(0, O((\log m)^2), (\log_2 m)/2k, (1/m)^{1/O(kt)})$-efficient protocol for $\text{MULTEQ}_{m,k,t}$ in the k-player SMP model.*

8 An Efficient Protocol for Sum-Zero

The Sum-Zero problem has been studied in the SMP model and efficient protocols with the following parameters have been obtained.

Theorem 8.1 ([Nis94]). *For every $k \in \mathbb{N}$ and $m \in \mathbb{N}$ there is a $(0, O(\log(m + \log k)), O(\log(m + \log k)), 1/2)$-efficient protocol for $\text{SUMZERO}_{m,k}$ in the k-player SMP model.*

The above protocol is based on a simple (yet powerful) idea of picking a small random prime p and checking if the numbers sum to zero modulo p. Viola put forth a protocol with better parameters than the above protocol (i.e., smaller message length) using specialized hash functions and obtained the following:

Theorem 8.2 ([Vio15]). *For every $k \in \mathbb{N}$ and $m \in \mathbb{N}$ there is a $(0, O(m), O(\log k), 1/2)$-efficient protocol for $\text{SUMZERO}_{m,k}$ in the k-player SMP model.*

In order to prove Theorem 1.6, we need a protocol with $o(m)$ randomness *and* $o(\log m)$ message length, and both the protocols described above do not meet these conditions. We show below that the above two results can be composed to get a protocol with $O(\log k)$ communication complexity and $O_k(\log m)$ randomness. In fact, we will use a slightly different protocol from [Vio15]: namely, the protocol for the Sum-Zero($\mathbb{Z}_p$) problem as stated below.

Definition 8.3 (Sum-Zero($\mathbb{Z}_p$) Problem). *Let $k, m, p \in \mathbb{N}$. $\mathbb{Z}_p\text{-SUMZERO}_{m,k} : (\{0,1\}^m)^k \to \{0,1\}$ is defined by*

$$\mathbb{Z}_p\text{-SUMZERO}_{m,k}(x_1, \ldots, x_k) = \begin{cases} 1 \text{ if } \sum_{i \in [k]} x_i = 0 \mod p, \\ 0 \text{ otherwise,} \end{cases}$$

where we think of each x_i as a number in $[-2^{m-1}, 2^{m-1} - 1]$.

Theorem 8.4 ([Vio15]). *For every $p, k \in \mathbb{N}$ and $m \in \mathbb{N}$, there is a $(0, O(\log p), O(\log k), 1/2)$-efficient protocol[16] for $\mathbb{Z}_p\text{-SUMZERO}_{m,k}$ in the k-player SMP model.*

---

[16] As written in [Vio15], the protocol has the following steps. First, the players send some messages to the referee. Then the players and the referee jointly draw some random coins, and finally the players send some more messages to the referee. However, we note that the first and second steps of the protocol can be swapped in [Vio15] to obtain an efficient protocol as the drawing of randomness do not depend on the messages sent by the players to the referee.

Below is our theorem which essentially combines Theorem 8.1 and Theorem 8.4.

Theorem 8.5. *For every $k \in \mathbb{N}$ and $m \in \mathbb{N}$ there is a $(0, O(\log(m + \log k)), O(\log k), 3/4)$-efficient protocol for $\text{SUMZERO}_{m,k}$ in the k-player SMP model.*

Proof. Let π^* be the protocol of Viola from Theorem 8.4.

The protocol

Let $t = 2(m - 1 + \log k)$. Let $p_1, \ldots, p_t$ be the first t primes.

1. All players jointly draw $i^* \in [t]$ uniformly at random.

2. The players and the referee run π^* where each player now has input $y_i = x_i \mod p_{i^*}$, and the referee accepts if and only if $\sum_{i \in [k]} y_i = 0 \mod p_{i^*}$.

Notice that the above protocol is still an efficient protocol as the first step of the above protocol can be combined with the draw of random coins in the first step of π^* to form a single step in which the players and the referee jointly draw random coins from the public randomness.

Analysis

Randomness. In the first step of the protocol all the players jointly draw $i^* \in [t]$ uniformly, which requires $\lceil \log_2 t \rceil$ random bits. Then, the players draw $O(\log p_{i^*})$ additional random coins for π^*. The bound on the randomness follows by noting that $p_{i^*} \leq p_t = O(t \log t)$.

Message Length. For every $j \in [k]$, Player j sends the referee $O(\log k)$ bits as per π^*.

Completeness. If the k numbers sum to zero, then for any $p \in \mathbb{N}$, they sum to zero mod p and thus the referee always accepts.

Soundness. If the k numbers do not sum to zero, then let $s(\pi^*)$ be the soundness of π^* and let S be the subset of the first t primes defined as follows:

$$S = \left\{ p_i \middle| i \in [t], \sum_{j \in [k]} x_j = 0 \bmod p_i \right\}$$

It is clear that the referee rejects with probability at least $(1 - |S|/t) \cdot (1 - s(\pi^*))$. Therefore, it suffices to show that $|S| < t/2$ as $s(\pi^*) \leq 1/2$. Let $x := \sum_{j \in [k]} x_j$. We have that $x \in [-k \cdot 2^{m-1}, k \cdot 2^{m-1}]$ and that, for every $p \in S$, p divides x. Since $x \neq 0$, we know that $|x| \geq \prod_{p \in S} p \geq \prod_{i=1}^{|S|} p_i \geq 2^{|S|}$.

Since $|x| \leq k \cdot 2^{m-1}$, we have that $|S| < m - 1 + \log k = t$, and the proof follows. $\square$

The following corollary follows immediately by applying Proposition 5.1 with $z = \log_2 m / 2k \cdot c \log k$ to the above theorem, where c is some constant such that the message length of the protocol in Theorem 8.5 is at most $c \log k$.

Corollary 8.6. *For every* $k, m \in \mathbb{N}$ *there is a* $\left(0, O((\log(m + \log k))^2), {}^{(\log_2 m)}/{2k}, \left({}^{1}/{m}\right)^{1/O(k \log k)}\right)$- *efficient protocol for* SUMZERO$_{m,k}$ *in the k-player SMP model.*

9 Conclusion and Open Questions

We showed the parameterized inapproximability results for k-DomSet under $\mathsf{W}[1] \neq \mathsf{FPT}$, ETH, SETH and k-SUM Hypothesis, which almost resolve the complexity status of approximating parameterized k-DomSet. Although we showed the $\mathsf{W}[1]$-hardness of the problem, the exact version of k-DomSet is $\mathsf{W}[2]$-complete. Thus, a remaining question is whether approximating k-DomSet is $\mathsf{W}[2]$-hard:

Open Question 9.1. *Can we base total inapproximability of k-DomSet on* $\mathsf{W}[2] \neq \mathsf{FPT}$?

We note that even 1.01-approximation of k-DomSet is not known to be $\mathsf{W}[2]$-hard.

Another direction is to look beyond k-DomSet and try to prove inapproximability of other parameterized problems. Since k-DomSet is $\mathsf{W}[2]$-complete, there are already known reductions from it to other $\mathsf{W}[2]$-hard problems. Some of these reductions such as those for k-Set Cover, k-Hitting Set, k-DomSet on tournaments [DF95b] and k-Contiguous Set [CNW16] are gap-preserving reductions (in the sense of [CCK+17, Definition 3.4]) and total inapproximability under $\mathsf{W}[1] \neq \mathsf{FPT}$ translate directly to those problems. In this sense, one may wish to go beyond $\mathsf{W}[2]$-hard problems and prove hardness of approximation for problems in $\mathsf{W}[1]$. One of the most well-studied $\mathsf{W}[1]$-complete problems is k-Clique which we saw earlier in our proof; in its optimization variant, one wishes to find a clique of maximum size in the graph. Note that in this case a k-approximation is trivial since we can just output one vertex. It was shown in [CCK+17] that no $o(k)$-approximation is possible in FPT time assuming Gap-ETH. Since this rules out all non-trivial approximation algorithms, such a non-existential result is also referred to as the *total FPT-inapproximability* of k-Clique. Hence, the question is whether one can bypass Gap-ETH for this result as well:

Open Question 9.2. *Can we base total inapproximability of k-Clique on* $\mathsf{W}[1] \neq \mathsf{FPT}$ *or* ETH?

Again, we note that the current state-of-the-art results do not even rule out 1.01-FPT-approximation algorithms for k-Clique under $\mathsf{W}[1] \neq \mathsf{FPT}$ or ETH.

Another direction is to look beyond parameterized complexity questions. As mentioned earlier, Abboud et al. [ARW17] used the hardness of approximating of PCP-Vectors as a starting point of their inapproximability results of problems in P. Since MaxCover is equivalent to PCP-Vectors when the number of right super-nodes is two, it may be possible that MaxCover for larger number of right super-nodes can also be used to prove hardness of problems in P as well. At the moment, however, we do not have any natural candidate in this direction (see Appendix B for further discussions).

References

[ABC09] Chrisil Arackaparambil, Joshua Brody, and Amit Chakrabarti. Functional monitoring without monotonicity. In *ICALP*, pages 95–106, 2009.

[AL13] Amir Abboud and Kevin Lewi. Exact weight subgraphs and the k-sum conjecture. In *ICALP*, pages 1–12, 2013.

[ALM⁺98] 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.

[ALW14] Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In *ESA*, pages 1–12, 2014.

[AMS06] Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k-restrictions. *ACM Trans. Algorithms*, 2(2):153–177, 2006.

[AMS12] Noga Alon, Ankur Moitra, and Benny Sudakov. Nearly complete graphs decomposable into large induced matchings and their applications. In *STOC*, pages 1079–1090, 2012.

[ARW17] Amir Abboud, Aviad Rubinstein, and Ryan Williams. Distributed PCP theorems for hardness of approximation in P. In *FOCS*, pages 25–36, 2017.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[AW09] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *TOCT*, 1(1):2:1–2:54, 2009.

[BCG⁺16] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. Short interactive oracle proofs with constant query complexity, via composition and sumcheck. *ECCC*, 23:46, 2016.

[BEKP15] Edouard Bonnet, Bruno Escoffier, Eun Jung Kim, and Vangelis Th. Paschos. On subexponential and FPT-time inapproximability. *Algorithmica*, 71(3):541–565, 2015.

[BGKL03] László Babai, Anna Gál, Peter G. Kimmel, and Satyanarayana V. Lokam. Communication complexity of simultaneous messages. *SIAM J. Comput.*, 33(1):137–166, 2003.

[BJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.

[BKK+16] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. Constant rate PCPs for circuit-sat with sublinear query complexity. *J. ACM*, 63(4):32:1–32:57, 2016.

[CCK+17] Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From Gap-ETH to FPT-inapproximability: Clique, Dominating Set, and more. In *FOCS*, pages 743–754, 2017.

[CFK+15] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

[CFL83] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In *STOC*, pages 94–99, 1983.

[CGG06] Yijia Chen, Martin Grohe, and Magdalena Grüber. On parameterized approximability. In *IWPEC*, pages 109–120, 2006.

[CH10] Liming Cai and Xiuzhen Huang. Fixed-parameter approximation: Conceptual framework and approximability results. *Algorithmica*, 57(2):398–412, 2010.

[CHK13] Rajesh Hemant Chitnis, MohammadTaghi Hajiaghayi, and Guy Kortsarz. Fixed-parameter and approximation algorithms: A new look. In *IPEC*, pages 110–122, 2013.

[CHKX06] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006.

[Chv79] Vasek Chvátal. A greedy heuristic for the set-covering problem. *Math. Oper. Res.*, 4(3):233–235, 1979.

[CL16] Yijia Chen and Bingkai Lin. The constant inapproximability of the parameterized dominating set problem. In *FOCS*, pages 505–514, 2016.

[CMY08] Graham Cormode, S. Muthukrishnan, and Ke Yi. Algorithms for distributed functional monitoring. In *SODA*, pages 1076–1085, 2008.

[CNW16] Moses Charikar, Yonatan Naamad, and Anthony Wirth. On approximating target set selection. In *APPROX*, pages 4:1–4:16, 2016.

[CRR14] Arkadev Chattopadhyay, Jaikumar Radhakrishnan, and Atri Rudra. Topology matters in communication. In *FOCS*, pages 631–640, 2014.

[DF95a]   Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: on completeness for W[1]. *Theor. Comput. Sci.*, 141(1&2):109–131, 1995.

[DF95b]   Rodney G. Downey and Michael R. Fellows. *Parameterized Computational Feasibility*, pages 219–244. Birkhäuser Boston, Boston, MA, 1995.

[DF13]   Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.

[DFM06]   Rodney G. Downey, Michael R. Fellows, and Catherine McCartin. Parameterized approximation problems. In *IWPEC*, pages 121–129, 2006.

[DFMR08]   Rodney G. Downey, Michael R. Fellows, Catherine McCartin, and Frances A. Rosamond. Parameterized approximation of dominating set problems. *Inf. Process. Lett.*, 109(1):68–70, 2008.

[Din07]   Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.

[Din16]   Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *ECCC*, 23:128, 2016.

[DS14]   Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *STOC*, pages 624–633, 2014.

[EG04]   Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theor. Comput. Sci.*, 326(1-3):57–67, 2004.

[Fei98]   Uriel Feige. A threshold of ln n for approximating set cover. *J. ACM*, 45(4):634–652, 1998.

[FOZ16]   Orr Fischer, Rotem Oshman, and Uri Zwick. Public vs. private randomness in simultaneous multi-party communication complexity. In *SIROCCO*, pages 60–74, 2016.

[GO95]   Anka Gajentaan and Mark H. Overmars. On a class of o(n2) problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.

[Gol08]   Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, New York, NY, USA, 1 edition, 2008.

[GR18a]   Oded Goldreich and Guy N. Rothblum. Counting t-cliques: Worst-case to average-case reductions and direct interactive proof systems. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:46, 2018.

[GR18b]   Oded Goldreich and Guy N. Rothblum. Simple doubly-efficient interactive proof systems for locally-characterizable sets. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, pages 18:1–18:19, 2018.

[GS96]   Arnaldo Garcia and Henning Stichtenoth. On the asymptotic behaviour of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248 – 273, 1996.

[HKK13]   Mohammad Taghi Hajiaghayi, Rohit Khandekar, and Guy Kortsarz. The foundations of fixed parameter inapproximability. *CoRR*, abs/1310.2711, 2013.

[IP01]   Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

[IPZ01]   Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.

[Joh74]   David S. Johnson. Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.*, 9(3):256–278, 1974.

[Jus72]   Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Trans. Information Theory*, 18(5):652–656, 1972.

[Kar72]   Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, pages 85–103, 1972.

[KN97]   Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.

[KS92]   Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.

[Lin15]   Bingkai Lin. The parameterized complexity of k-biclique. In *SODA*, pages 605–615, 2015.

[Lov75]   L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975.

[LS09]   Troy Lee and Adi Shraibman. Lower bounds in communication complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):263–398, 2009.

[LV11]   Guanfeng Liang and Nitin H. Vaidya. Multiparty equality function computation in networks with point-to-point links. In *SIROCCO*, pages 258–269, 2011.

[LY94]   Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *J. ACM*, 41(5):960–981, 1994.

[Mei13]   Or Meir. IP = PSPACE using error-correcting codes. *SIAM J. Comput.*, 42(1):380–403, 2013.

[Mos15]   Dana Moshkovitz. The projection games conjecture and the NP-hardness of ln n-approximating set-cover. *Theory of Computing*, 11:221–235, 2015.

[MR16]   Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense CSPs. *CoRR*, abs/1607.02986, 2016.

[Nis94]   Noam Nisan. The communication complexity of threshold gates. In *Proceedings of Combinatorics, Paul Erdos is Eighty*, pages 301–315, 1994.

[Pat10]   Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *STOC*, pages 603–610, 2010.

[PVZ12] Jeff M. Phillips, Elad Verbin, and Qin Zhang. Lower bounds for number-in-hand multiparty communication complexity, made easy. In *SODA*, pages 486–501, 2012.

[PW10] Mihai Patrascu and Ryan Williams. On the possibility of faster SAT algorithms. In *SODA*, pages 1065–1075, 2010.

[PY91] Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991.

[Raz92] Alexander A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.

[RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *STOC*, pages 475–484, 1997.

[Rub17] Aviad Rubinstein. Personal communication, 2017.

[Rub18] Aviad Rubinstein. Hardness of approximate nearest neighbor search. In *STOC*, 2018. To appear.

[SAK+01] Kenneth W. Shum, Ilia Aleshnikov, P. Vijay Kumar, Henning Stichtenoth, and Vinay Deolalikar. A low-complexity algorithm for the construction of algebraic-geometric codes better than the Gilbert-Varshamov bound. *IEEE Trans. Information Theory*, 47(6):2225–2241, 2001.

[Sla96] Petr Slavík. A tight analysis of the greedy algorithm for set cover. In *STOC*, pages 435–441, 1996.

[Sri95] Aravind Srinivasan. Improved approximations of packing and covering problems. In *STOC*, pages 268–276, 1995.

[SS96] Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Trans. Information Theory*, 42(6):1710–1722, 1996.

[Tov84] Craig A. Tovey. A simplified NP-complete satisfiability problem. *Discrete Applied Mathematics*, 8(1):85–89, 1984.

[Vio15] Emanuele Viola. The communication complexity of addition. *Combinatorica*, 35(6):703–747, 2015.

[Wil05] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.

[WW15] Omri Weinstein and David P. Woodruff. The simultaneous communication of disjointness with applications to data streams. In *ICALP*, pages 1082–1093, 2015.

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC*, pages 209–213, 1979.

A   A Reduction From MaxCover to DomSet

In this section, we provide the reduction from MaxCover to DomSet (i.e. Theorem 3.1). The reduction and its proof presented here are quoted almost directly from Chalermsook et al. [CCK$^+$17] except with appropriate notational adjustments. We include it here only for the sake of self-containedness.

To ease the presentation, we will present the reduction in terms of a different variant of the label cover problem called MinLabel. The input to MinLabel is the same as that of MaxCover, i.e., it is a label cover instance $\Gamma = (U = \bigcup_{i \in [q]} U_i, W = \bigcup_{j \in [k]} W_j; E)$ as defined in Subsection 3.2.

However, the solution and the objective of the problem will be different from MaxCover. Specifically, a solution of MinLabel is called a *multi-labeling*, which is simply a subset of vertices $\widehat{S} \subseteq W$. We say that a multi-labeling $\widehat{S}$ *covers* a left super-node U_i if there exists a vertex $u_i \in U_i$ that has a neighbor in $\widehat{S} \cap W_j$ for every $j \in [k]$. The goal in MinLabel is to find a minimum-size multi-labeling that covers all the left super-nodes U_i's. Again, we overload the notation and use MinLabel to also denote the optimum, i.e.,

$$\mathsf{MinLabel}(\Gamma) = \min_{\substack{\text{multi-labeling } \widehat{S} \\ \text{that covers every } U_i}} |\widehat{S}|.$$

The following lemma demonstrates relations between MaxCover and MinLabel on the same input label cover instance. These relations in turns allow us to deduce inapproximability results of MinLabel from that of MaxCover.

Proposition A.1. *Let Γ be any label cover instance with k right super-nodes. Then, we have*

- *If $\mathsf{MaxCover}(\Gamma) = 1$, then $\mathsf{MinLabel}(\Gamma) = k$.*

- *If $\mathsf{MaxCover}(\Gamma) \leq \epsilon$ for some $0 < \epsilon$, then $\mathsf{MinLabel}(\Gamma) \geq (1/\epsilon)^{1/k} \cdot k$.*

Before proceeding to prove Proposition A.1, let us state the desired reduction from label cover to DomSet (Theorem 3.1) in terms of MinLabel instead of MaxCover:

Lemma A.2 (Reduction from MinLabel to DomSet [CCK$^+$17, Theorem 5.4])). *There is an algorithm that, given a label cover instance $\Gamma = (U = \bigcup_{j=1}^{q} U_j, W = \bigcup_{i=1}^{k} W_i, E)$, outputs a k-DomSet instance G such that*

- $\mathsf{MinLabel}(\Gamma) = \mathsf{DomSet}(G).$

- $|V(G)| = |W| + \sum_{i \in [q]} k^{|U_q|}.$

- *The reduction runs in time $O\left(|W| \left(\sum_{i \in [q]} k^{|U_q|}\right)\right).$*

It is obvious that Proposition A.1 and Lemma A.2 together imply Theorem 3.1. The rest of this section is devoted to proving Proposition A.1 and Lemma A.2. Their proofs can be found in Subsections A.1 and A.2, respectively.

A.1 From MaxCover to MinLabel: Proof of Proposition A.1

In this section, we prove Proposition A.1. The proof presented here is due to [CCK$^+$17] with slight changes to match our notations.

Proof of Proposition A.1. (1) Suppose MaxCover(Γ) = 1. Then there exists a labeling $S \subseteq W$ that covers every left super-node U_i. Hence, S is also a multi-labeling that covers every left super-node, which implies that MinLabel(Γ) = k.

(2) We prove by contrapositive. Assume that MinLabel(Γ) < $(1/\varepsilon)^{1/k}k$. Then there exists a multi-labeling $\widehat{S} \subseteq W$ of size less than $(1/\varepsilon)^{1/k}k$ that covers every left super-node. Observe that we can construct a feasible S' solution to MaxCover by uniformly and independently choosing one node from $\widehat{S} \cap W_j$ at random, for each of the right super-node W_j.

Thus, the expected number of left super-nodes covered by S' is

$$\mathbb{E}_{S'}\left[|\{i \in [q] : S' \text{ covers } U_i\}|\right] \geq \sum_{i \in [q]} \prod_{j \in [k]} |\widehat{S} \cap W_j|^{-1}$$

$$(\text{By AM-GM inequality}) \geq \sum_{i \in [q]} \left(\frac{1}{k} \sum_{j \in [k]} |\widehat{S} \cap W_j|\right)^{-k}$$

$$> q \cdot \left(\frac{1}{k} \cdot \left(\left(\frac{1}{\varepsilon}\right)^{1/k} k\right)\right)^{-k}$$

$$= q \cdot \varepsilon$$

Hence, there is a labeling that covers $> \varepsilon q$ left super-nodes, i.e., MaxCover(Γ) $> \varepsilon$. $\qquad\square$

A.2 From MinLabel to DomSet: Proof of Lemma A.2

We devote this subsection to show a reduction from MinLabel to DomSet as stated in Lemma A.2.

Proof of Lemma A.2. We show the reduction from MinLabel to the bipartite version of DomSet, where the input graph $G = (A, B; E)$ is bipartite, and we are asked to find a subset of (left) vertices of A that dominates all the (right) vertices in B. This variant of DomSet is known as Red-Blue DomSet, which is equivalent to the *Set Cover* problem. One can transform the bipartite version of DomSet into the general case just by forming a clique on A. It is not hard to see that the optimum remains unchanged.

We apply the reduction from [CCK$^+$17], which is in turn taken from [Fei98] (and [LY94]) with small adjustments. Our construction requires the *hypercube partition system*.

Hypercube Partition System. The (κ, ρ)-hypercube partition system consists of the universe $\mathcal{M}$ and a collection of subset $\{P_{x,y}\}_{x \in [\rho], y \in [\kappa]}$ where

$$\mathcal{M} = [\kappa]^\rho \quad \text{and} \quad P_{x,y} = \{z \in \mathcal{M} : z_x = y\}$$

The following are properties that can be observed from the hypercube partition system:

- **Partition Cover:** For each row $x \in [\rho]$, $\mathcal{P}_x = (P_{x,1}, \ldots, P_{x,\kappa})$ is a partition of $\mathcal{M}$. Thus, $\bigcup_{y \in [\kappa]} P_{x,y} = \mathcal{M}$ and $P_{x,y} \cap P_{x,y'} = \emptyset$ for $y \neq y'$.

  That is, ones can cover the universe $\mathcal{M}$ by picking all κ subsets from some row $x \in [\kappa]$.

- **Robust Property:** For any $y_1^*, \ldots, y_\rho^* \in [\kappa]$, $\bigcup_{x \in [\alpha], y \in [\kappa] \setminus \{y_x^*\}} P_{x,y} \neq \mathcal{M}$.

  That is, the only way to cover the universe $\mathcal{M}$ is to pick all the κ subsets from the same row. Otherwise, even if we include $\kappa - 1$ subsets from every row $x \in [\rho]$, it is not possible to cover the universe $\mathcal{M}$.

- **Size:** The number of elements in the partition is κ^ρ.

The Reduction. Now we present our reduction. Let $\Gamma = (U = \bigcup_{j=1}^q U_j, W = \bigcup_{i=1}^k W_i; E')$ be an instance of MinLabel. We will construct a bipartite graph $G = (A, B; E)$ of Red-Blue DomSet from Γ. The vertices of A are taken from the right vertices of Γ, i.e., $A = W$.

For each super-node U_i, we create the $(k, |U_i|)$-hypercube partition system $(\mathcal{M}^i, \{P_{x,y}^i\}_{x \in [|U_q|], y \in [k]})$ (i.e., $\rho = |U_q|$ and $\kappa = k$); then we take the elements of $\mathcal{M}^j$ as left vertices of G.

We name vertices in U_i by $u_1^i, u_2^i, \ldots, u_{|U_i|}^i$ for $i \in [q]$. For each $j \in k$ and each $w^j \in W_j$, we join w_j to every vertex $z \in P_{x,j}^i$ if $\{u_x^i, w^j\}$ is an edge in Γ. More precisely, the bipartite graph $G = (A, B; E)$ is defined such that

$$A := W = \bigcup_j W_j, B := \bigcup_{i=1}^q \mathcal{M}^i, E := \{\{w^j, z\} : i \in [q], j \in [k], x \in [|U_i|], w^j \in W_j, z \in P_{x,j}^i, \{u_x^i, w^j\} \in E'\}.$$

Size of The Construction. It is clear from that the numbers of left and right vertices in our construction is as follows.

$$\text{The Number of Left veritces of } G = |W|$$
$$\text{The Number of Right vertices of } G = \sum_{i \in [q]} k^{|U_i|}$$

Analysis.   Clearly, the size translation satisfies the lemma statement. For the claim $\mathsf{MinLabel}(\Gamma) = \mathsf{DomSet}(G)$, observe that any solution S to the MinLabel instance corresponds to a solution to the Red-Blue DomSet instance. We claim that S is feasible to the MinLabel instance if and only if it is feasible to the DomSet instance. This will prove the claim and thus imply the lemma.

First, suppose S is feasible to MinLabel. Then, for every U_i, there is a vertex $u_x^i \in U_i$ that has an edge to some vertex $w_y^j \in W_j \cap S$, for all $j \in [k]$. This means that we have an edge from w_y^j to every vertex $z \in P_{x,y}^i$. It then follows by the Partition Cover property of the hypercube partition system that S dominates every vertex in the universe $\mathcal{M}^i$. Consequently, we conclude that S dominates $B = \bigcup_{i \in [q]} \mathcal{M}^i$.

Conversely, suppose S is feasible to DomSet. Then S has an edge to every vertex in the universe $\mathcal{M}^i \subseteq B$. We know by construction and the Robustness Property of the hypercube partition system that, for some $x \in [|U_i|]$, S contains vertices $w^1 \in S \cap W_1, w^2 \in S \cap W_2, \ldots, w^k \in S \cap W_k$, which correspond to the subsets $P_{x,y}^i$'s in the partition system, and that $w^1, w^2, \ldots, w^k$ are incident to the vertex $u_x^i \in U_i$. This means that S covers the left super-node U_i. Consequently, we conclude that S is a feasible multi-labeling, thus completing the proof. □

B   Connection to Fine-Grained Complexity

In this section, we will demonstrate the conditional hardness of problems in P by basing them on the conditional hardness of PSP. First, we define the problem in P of interest to this section.

Definition B.1 (*k*-linear form inner product). *Let $k, m \in \mathbb{N}$. Given $x_1, \ldots, x_k \in \mathbb{R}^m$, we define the inner product of these k vectors as follows:*

$$\langle x_1, \ldots, x_k \rangle = \sum_{i \in [m]} \prod_{j \in [k]} x_i(j),$$

where $x_i(j)$ denotes the j^{th} coordinate of the vector x_i.

Definition B.2 (*k*-chromatic Maximum Inner Product). *Let $k \in \mathbb{N}$. Given k collections $A_1, A_2, \ldots, A_k$ each of N vectors in $\{0,1\}^D$, where $D = N^{o(1)}$, and an integer s, the k-chromatic Maximum Inner Product (MIP) problem is to determine if there exists $a_i \in A_i$ for all $i \in [k]$ such that $\langle a_1, \ldots, a_k \rangle \geq s$.*

We continue to use the shorthand $\Gamma(N, k, r, \ell)$ introduced in Section 5.2. Moreover, while using the shorthand $\Gamma(N, k, r, \ell)$, we will assume that for all $i \in [h]$, $|W_i| \leq N$. We define Unique MaxCover to be the MaxCover problem with the following additional structure: for every labeling $S \subseteq W$ (see Section 3.2 to recall the definition of a labeling) and any left super-node U_i, there is at most one node in U_i which is a neighbor to all the nodes in S. We remark that the reduction in Theorem 5.2 already produces instances of Unique MaxCover. This follows from the proof of Theorem 5.2 by noting that on every random string, each player sends a message to the referee in a deterministic way. Finally, we have the following connection between unique MaxCover and MIP.

Theorem B.3. *Let $N, k, r, \ell \in \mathbb{N}$. There is a reduction from any* Unique MaxCover *instance $\Gamma(N, k, r, \ell)$ to k-chromatic* MIP *instance $(A_1, \ldots, A_k, s)$ such that*

- *For all $i \in [k]$, $|A_i| \leq N$, $s = 2^r$, and $D \leq 2^{r+\ell k}$.*

- *The running time of the reduction is $O(Nk \cdot 2^{r+\ell k})$.*

- *For any integer s^*, there exists $a_i \in A_i$ for all $i \in [k]$ such that $\langle a_1, \ldots, a_k \rangle \geq s^*$ if and only if $\mathsf{MaxCover}(\Gamma) \geq s^*/2^r$.*

Proof. For every $i \in [k]$, we associate each A_i with the right super-node W_i. Each node in W_i corresponds to a vector in A_i. Given a node w in W_i, the corresponding vector $a \in A_i$ is constructed as follows. Each coordinate of a corresponds to a left node. Let the j^{th} coordinate of a correspond to a node $u \in U$. The j^{th} coordinate of a is 1 if there is an edge between u and w in the graph G of the instance Γ; otherwise, the j^{th} coordinate of a is 0. It is easy to see that $|A_i| \leq N$ and $D \leq 2^{r+\ell k}$, and the running time of the reduction is $O(Nk \cdot 2^{r+\ell k})$. It is also easy to see that if $\mathsf{MaxCover}(\Gamma) \geq s^*/2^r$ then the vectors a_i corresponding to the nodes in the labeling resulting in the maximum cover, have inner product at least s^*.

We will now show that for any integer s^*, if there exists $a_i \in A_i$ for all $i \in [k]$ such that $\langle a_1, \ldots, a_k \rangle \geq s^*$ then $\mathsf{MaxCover}(\Gamma) \geq s^*/2^r$. Fix an integer s^*. Suppose that there exists $a_i \in A_i$ for all $i \in [k]$ such that $\langle a_1, \ldots, a_k \rangle \geq s^*$. The k-tuple $(a_1, \ldots, a_k)$ corresponds to a labeling S of unique MaxCover. For each coordinate such that all the $a_i = 1$ on that coordinate, we note that there is a left node (corresponding to the coordinate) which is a common neighbor of all the nodes in the labeling. Moreover, since in the unique MaxCover we can have at most one node in each left super-node which is a common neighbor of the nodes in the labeling, we have that:

$$\mathsf{MaxCover}(\Gamma) \geq \frac{1}{2^r} \cdot |\{i \in [2^r] \mid U_i \text{ is covered by } S\}| = \langle a_1, \ldots, a_k \rangle \geq s^*. \qquad \square$$

The results for hardness of approximation for problems in P is obtained by simply fixing the value of k in the above theorem to some universal constant (such as $k = 2$). For example, by fixing $k = 2$ and applying the above reduction to Corollary 5.3, we recover the main result of [ARW17] on bichromatic MIP. This is not surprising as under SETH, our framework is just a generalization of the distributed PCP framework of [ARW17].

Furthermore, we demonstrate the flexibility of our framework by fixing $k = 3$ and applying the above reduction to Corollary 5.6, to establish a hardness of approximation result of trichromatic MIP under the 3-SUM Hypothesis as stated below. We note here that the running time lower bound is only $N^{2-o(1)}$, which is likely not tight since the lower bound from SETH is $N^{3-o(1)}$.

Theorem B.4. *Assuming the 3-SUM Hypothesis, for every $\varepsilon > 0$, no $O(N^{2-\varepsilon})$ time algorithm can, given three collections A, B, and C each of N vectors in $\{0,1\}^D$, where $D = N^{o(1)}$, and an integer s, distinguish between the following two cases:*

Completeness. *There exists $a \in A, b \in B, c \in C$ such that $\langle a, b, c \rangle \geq s$.*

Soundness. *For every $a \in A, b \in B, c \in C$, $\langle a, b, c \rangle \leq s/2^{(\log N)^{1-o(1)}}$.*

Proof. We apply Proposition 5.1 with $z = m/(\log_2 m)^2$ and $k = 3$ to Theorem 8.5, to obtain a $\left(0, O(m/\log m), O(m/(\log m)^2), (1/2)^{m^{1-o(1)}}\right)$-efficient protocol for $\textsc{SumZero}_{m,3}$ in the 3-player SMP

model. By plugging in the parameters of the above protocol to Corollary 5.6, we obtain that assuming the 3-SUM hypothesis, for every $\varepsilon > 0$, no $O(N^{2-\varepsilon})$-time algorithm can distinguish between $\mathsf{MaxCover}(\Gamma) = 1$ and $\mathsf{MaxCover}(\Gamma) \leq (1/2)^{(\log N)^{1-o(1)}}$ for any label cover instance $\Gamma(N, 3, r, \ell)$ for all $N \in \mathbb{N}$. The proof of the theorem concludes by applying Theorem B.3 to the above hardness of MaxCover (Note that Theorem 5.2 provides a reduction from $\mathsf{PSP}(k, \mathcal{F}, N)$ to Unique MaxCover). $\square$