THEORY OF COMPUTING, Volume 16 (4), 2020, pp. 1–50
www.theoryofcomputing.org

SPECIAL ISSUE: CCC 2018

On The Hardness of Approximate and Exact (Bichromatic) Maximum Inner Product

Lijie Chen[*]

Received July 1, 2018; Revised May 28, 2019; Published September 23, 2020

Abstract. In this paper we study the (Bichromatic) Maximum Inner Product Problem (Max-IP), in which we are given sets A and B of vectors, and the goal is to find $a \in A$ and $b \in B$ maximizing inner product $a \cdot b$. Max-IP is a basic question and serves as the base problem in the recent breakthrough of [Abboud et al., FOCS 2017] on hardness of approximation for polynomial-time problems. It is also used (implicitly) in the argument for hardness of exact ℓ_2-Furthest Pair (and other important problems in computational geometry) in poly-loglog dimensions in [Williams, SODA 2018]. We have three main results regarding this problem.

- **Characterization of Multiplicative Approximation**. First, we study the best multiplicative approximation ratio for Boolean Max-IP in subquadratic time. We show that, for Max-IP with two sets each consisting of n vectors from $\{0,1\}^d$, there is an $n^{2-\Omega(1)}$-time multiplicative t-approximation algorithm when $t = (d/\log n)^{\Omega(1)}$. We also show this is conditionally optimal, as a $(d/\log n)^{o(1)}$-approximation algorithm would refute SETH. Similar characterization is also achieved for additive approximation for Max-IP.

[*] Supported by an Akamai Fellowship.

ACM Classification: F.1.3

AMS Classification: 68Q17

Key words and phrases: Maximum Inner Product, SETH, hardness of approximation in P, fine-grained complexity, Hopcroft's problem, Chinese Remainder Theorem

DOI: 10.4086/toc.2019.v016a004

- **$2^{O(\log^* n)}$-dimensional Hardness for Exact Max-IP Over The Integers.** Second, we revisit the hardness of solving Max-IP exactly for vectors with integer entries. We show that, under SETH, for Max-IP with sets of n vectors from $\mathbb{Z}^d$ for some $d = 2^{O(\log^* n)}$, every exact algorithm requires $n^{2-o(1)}$ time. With the reduction from [Williams, SODA 2018], it follows that ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair in dimension $2^{O(\log^* n)}$ require $n^{2-o(1)}$ time.
- **Connection with $\mathsf{NP} \cdot \mathsf{UPP}$ Communication Protocols.** Last, we establish a connection between conditional lower bounds for exact Max-IP with integer entries and $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocols for Set-Disjointness, parallel to the connection between conditional lower bounds for approximate Max-IP and MA communication protocols for Set-Disjointness.

The lower bound in our first result is a direct corollary of the new MA protocol for Set-Disjointness introduced in [Rubinstein, STOC 2018], and our algorithms utilize the polynomial method and simple random sampling. Our second result follows from a new dimensionality self reduction from the Orthogonal Vectors problem for n vectors from $\{0,1\}^d$ to n vectors from $\mathbb{Z}^\ell$ where $\ell = 2^{O(\log^* d)}$, dramatically improving the previous reduction in [Williams, SODA 2018]. The key technical ingredient is a recursive application of *Chinese Remainder Theorem*.

As a by-product we obtain an MA communication protocol for Set-Disjointness with complexity $O(\sqrt{n \log n \log\log n})$, slightly improving the $O(\sqrt{n}\log n)$ bound [Aaronson and Wigderson, TOCT 2009], and approaching the $\Omega(\sqrt{n})$ lower bound [Klauck, CCC 2003].

Moreover, we show that (under SETH) one can apply the $O(\sqrt{n})$ BQP communication protocol for Set-Disjointness to prove near-optimal hardness for approximation to Max-IP with vectors in $\{-1,1\}^d$. This answers a question from [Abboud et al., FOCS 2017] in the affirmative.

1 Introduction

Maximum Inner Product Search is a fundamental similarity search problem in which you want to maintain a collection of vectors S, and answer queries of the form that given a new vector q, find the vector in S which is the most correlated to q (or an approximation to it). This problem is closely related to another fundamental problem called nearest neighbor search, in which one needs to maintain a collection of points, and find the nearest neighbor for the query points (or an approximation to it).

In this paper we consider a natural offline version of Maximum Inner Product Search, in which you are only required to compute the maximum correlation between S and all queries, and queries are given in advance.[1]

Definition 1.1 (Bichromatic Maximum Inner Product (Max-IP))**.** For $n, d \in \mathbb{N}$, the $\mathsf{Max\text{-}IP}_{n,d}$ problem is

[1] Clearly, this offline version is weaker than the online similarity search counterpart, so lower bounds for it automatically imply lower bounds for its online version.

defined as: *given two sets A,B each consisting of n vectors from $\{0,1\}^d$ compute*

$$\mathsf{OPT}(A,B) := \max_{a\in A, b\in B} a\cdot b.$$

We use $\mathbb{Z}$-Max-$\mathsf{IP}_{n,d}$ ($\mathbb{R}$-Max-$\mathsf{IP}_{n,d}$) to denote the same problem, but with A,B being sets of vectors from $\mathbb{Z}^d$ ($\mathbb{R}^d$).

1.1 Motivation and background

1.1.1 Hardness of approximate Max-IP

A natural brute-force algorithm solves Max-IP in $O(n^2\cdot d)$ time. Assuming SETH[2], there is no $n^{2-\Omega(1)}$-time algorithm for Max-$\mathsf{IP}_{n,d}$ when $d=\omega(\log n)$ [73].

Despite being one of the most central problems in similarity search and having numerous applications [48, 43, 15, 64, 65, 68, 17, 16, 18, 60, 69, 71, 14, 50, 11, 70, 34, 33], until recently it was unclear whether there could be a near-linear time, 1.1-approximation algorithm, before the recent breakthrough by Abboud, Rubinstein and Williams [5].[3]

In [5], a framework for proving inapproximability results for problems in P is established (the distributed PCP framework), from which it follows:

Theorem 1.2 (Abbaud, Rubinstein, Williams 2017)**.** Assuming SETH, no $n^{2-\Omega(1)}$-time algorithm for Max-$\mathsf{IP}_{n,n^{o(1)}}$ can achieve multiplicative $2^{(\log n)^{1-o(1)}}$-approximation.

Theorem 1.2 is an exciting development for hardness of approximation in P, implying other important inapproximability results for a host of problems including Bichromatic LCS Closest Pair Over Permutations, Approximate Regular Expression Matching, and Diameter in Product Metrics [5]. However, we still do not have a complete understanding of the approximation hardness of Max-IP yet. For instance, consider the following two concrete questions:

Question 1. *Is there a multiplicative $(\log n)$-approximation $n^{2-\Omega(1)}$-time algorithm for Max-$\mathsf{IP}_{n,\log^2 n}$? What about a multiplicative 2-approximation algorithm for Max-$\mathsf{IP}_{n,\log^2 n}$?*

Question 2. *Is there an additive $(d/\log n)$-approximation $n^{2-\Omega(1)}$-time algorithm for Max-$\mathsf{IP}_{n,d}$?*

We note that the lower bound from [5] cannot answer Question 1. Tracing the details of their proofs, one can see that it only shows approximation hardness for dimension $d=\log^{\omega(1)} n$. Question 2 concerning additive approximation is not addressed at all by [5]. Given the importance of Max-IP, it is interesting to ask:

For what ratios r do $n^{2-\Omega(1)}$-time r-approximation algorithms exist for Max-IP?

Does the best-possible approximation ratio (in $n^{2-\Omega(1)}$ time) relate to the dimensionality, in some way?

[2] SETH (Strong Exponential Time Hypothesis) states that for every $\varepsilon>0$ there is a k such that k-SAT cannot be solved in $O((2-\varepsilon)^n)$ time [47].

[3] See [5] for a thorough discussion on the state of affairs on hardness of approximation in P before their work.

In an important recent work, Rubinstein [67] improved the distributed PCP construction in a crucial way, from which one can derive more refined lower bounds on approximate $\mathsf{Max\text{-}IP}$. He then used the refined lower bounds to establish the SETH-hardness of $(1+o(1))$-approximate nearest neighbor search.

Building on Rubinstein's technique in this paper we provide full *characterizations*, determining essentially optimal multiplicative approximations and additive approximations to $\mathsf{Max\text{-}IP}$, under SETH.

1.1.2 Hardness of exact $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$

Recall that from [73], there is no $n^{2-\Omega(1)}$-time algorithm for exact Boolean $\mathsf{Max\text{-}IP}_{n,\omega(\log n)}$. Since in real-life applications of similarity search, one often deals with real-valued data instead of just Boolean data, it is natural to ask about $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ (which is certainly a special case of $\mathbb{R}$-$\mathsf{Max\text{-}IP}$): what is the maximum d such that $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,d}$ can be solved exactly in $n^{2-\Omega(1)}$ time?

Besides being interesting in its own right, $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ also has reductions to ℓ_2-Furthest Pair and to Bichromatic ℓ_2-Closest Pair. Hence, lower bounds for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ imply lower bounds for these two famous problems in computational geometry (see [75] for a discussion on this topic).

Prior to our work, it was implicitly shown in [75] that:

Theorem 1.3 ([75])**.** Assuming SETH, there is no $n^{2-\Omega(1)}$-time algorithm for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,\omega((\log\log n)^2)}$ with vectors of $O(\log n)$-bit entries.

However, the best known algorithm for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ runs in $n^{2-\Theta(1/d)}$ time [58, 10, 78][4], hence there is still a gap between the lower bound and the best known upper bounds. To confirm these algorithms are in fact optimal, we would like to prove a lower bound with $\omega(1)$ dimensions.

In this paper, we significantly strengthen the previous lower bound from $\omega((\log\log n)^2)$ dimensions to $2^{O(\log^* n)}$ dimensions. ($2^{O(\log^* n)}$ is an *extremely slowly growing* function. See the Preliminaries for its formal definition.)

1.1.3 Fine-grained complexity and communication complexity

One intriguing aspect of the distributed PCP framework is that it makes use of the $\widetilde{O}(\sqrt{n})$ MA communication protocol for Set-Disjointness [1]. Several follow-up works [51, 67] explored this connection further, and settled the hardness of approximation to several fundamental problems (under SETH).

Given the success of the interplay between these two seemingly unrelated fields, it is natural to seek more results from it. In particular, it is asked in [5] whether the $O(\sqrt{n})$ BQP communication protocol for Set-Disjointness can be utilized.

In this paper, we answer the question affirmatively by showing that BQP communication protocol implies hardness for approximation to $\{-1,1\}$-$\mathsf{Max\text{-}IP}$.[5] Moreover, we also establish a connection between $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ lower bounds and $\mathsf{NP}\cdot\mathsf{UPP}$ communication protocols for Set-Disjointness, which suggests a new perspective on our results on $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$.

[4] [10, 78] are for ℓ_2-Furthest Pair or Bichromatic ℓ_2-Closest Pair. They also work for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ as there are reductions from $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ to these two problems, see [75] or Lemma 4.6 and Lemma 4.7.

[5] That is, $\mathsf{Max\text{-}IP}$ with sets A and B each consisting of n vectors from $\{-1,1\}^d$.

1.2 Hypothesis assumed in this paper

We use $\mathsf{OV}_{n,d}$ to denote the Orthogonal Vectors problem: given two sets of vectors A,B each consisting of n vectors from $\{0,1\}^d$, determine whether there is an $a \in A$ and $b \in B$ such that $a \cdot b = 0$.[6] Similarly, we use $\mathbb{Z}$-$\mathsf{OV}_{n,d}$ to denote the same problem except for that A and B consist of vectors from $\mathbb{Z}^d$ (which is also called Hopcroft's problem).

All our results are based on the following widely used conjecture about OV:

Conjecture 1.4 (Orthogonal Vectors Conjecture (OVC) [73, 8]). For every $\varepsilon > 0$, there exists a $c \geq 1$ such that $\mathsf{OV}_{n,d}$ requires $n^{2-\varepsilon}$ time when $d = c \log n$.

OVC is a plausible conjecture as it is implied by the popular Strong Exponential Time Hypothesis [47, 29] on the time complexity of solving k-SAT [73, 76].

1.3 Our results on $\{0,1\}$-Max-IP

The first main result of our paper characterizes when there is a truly subquadratic-time ($n^{2-\Omega(1)}$-time, for some universal constant hidden in the big-Ω) multiplicative t-approximation algorithm for Max-IP, and characterizes the best-possible additive approximations as well. Let $\mathbb{P}$ be an optimization problem (can be either minimization or maximization). We begin with formal definitions of these two standard types of approximation:

- We say $\mathbb{A}$ is a multiplicative t-approximation algorithm for $\mathbb{P}$, if for all instances I, $\mathbb{A}$ outputs a value $\widetilde{\mathsf{OPT}}(I)$ such that $\widetilde{\mathsf{OPT}}(I) \in [\mathsf{OPT}(I), \mathsf{OPT}(I) \cdot t]$, where $\mathsf{OPT}(I)$ is the answer to $\mathbb{P}$ on instance I.

- We say $\mathbb{A}$ is an additive t-approximation algorithm for $\mathbb{P}$, if for all instances I, $\mathbb{A}$ outputs a value $\widetilde{\mathsf{OPT}}(I)$ such that $|\widetilde{\mathsf{OPT}}(I) - \mathsf{OPT}(I)| \leq t$.

- To avoid ambiguity, we call an algorithm computing $\mathsf{OPT}(I)$ exactly an *exact* algorithm for $\mathbb{P}$.

1.3.1 Characterizations of hardness of multiplicative approximations to Max-IP

In the multiplicative case, our characterization (formally stated below) basically says that there is a multiplicative t-approximation $n^{2-\Omega(1)}$-time algorithm for Max-$\mathsf{IP}_{n,d}$ if and only if $t = (d/\log n)^{\Omega(1)}$. Note that in the following theorem we require $d = \omega(\log n)$, since in the case of $d = O(\log n)$, there *are* $n^{2-\varepsilon}$-time algorithms for exact Max-$\mathsf{IP}_{n,d}$ [14, 13].

Theorem 1.5. Letting $\omega(\log n) < d < n^{o(1)}$ and $t \geq 2$,[7] the following holds.

[6] Here we use the bichromatic version of OV instead of the monochromatic one for convenience, as they are equivalent.

[7] Note that t and d are both functions of n, we assume they are computable in $n^{o(1)}$ time throughout this paper for simplicity.

1. There is an $n^{2-\Omega(1)}$-time multiplicative t-approximation algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if
$$t = (d/\log n)^{\Omega(1)},$$
and under SETH (or OVC), there is no $n^{2-\Omega(1)}$-time multiplicative t-approximation algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if
$$t = (d/\log n)^{o(1)}.$$

2. Moreover, let
$$\varepsilon = \min\left(\frac{\log t}{\log(d/\log n)}, 1\right).$$
There are multiplicative t-approximation deterministic algorithms for $\mathsf{Max\text{-}IP}_{n,d}$ running in time
$$O\left(n^{2-\Omega(\varepsilon)} \cdot \mathrm{polylog}(n)\right).$$

Remark 1.6. *We remark that our algorithm still gets a non-trivial speed-up over the brute force algorithm as long as* $\varepsilon = \omega(\log\log n/\log n)$.

Comparison with [11]. We remark that in [11], subquadratic-time algorithms for multiplicative t-approximation Max-IP is given when $t = d^{\Omega(1)}$.[8] Our algorithms achieve a slightly better approximation ratio $t = (d/\log n)^{\Omega(1)}$ which matches the conditional lower bound. Moreover, our algorithms work for the following more general case which is not handled by [11].

The algorithms in Theorem 1.5 indeed work for the case where the sets consist of non-negative reals (i. e., $\mathbb{R}^+$-Max-IP).

Corollary 1.7. Assuming $\omega(\log n) < d < n^{o(1)}$ and letting
$$\varepsilon = \min\left(\frac{\log t}{\log(d/\log n)}, 1\right),$$
there is a multiplicative t-approximation deterministic algorithm for $\mathbb{R}^+$-$\mathsf{Max\text{-}IP}_{n,d}$ running in time
$$O\left(n^{2-\Omega(\varepsilon)} \cdot \mathrm{polylog}(n)\right).$$

The lower bound of Theorem 1.5 is a direct corollary of the new improved MA protocols for Set-Disjointness from [67], which is based on Algebraic Geometry codes. Together with the framework of [5], that MA-protocol implies a reduction from OV to approximate Max-IP.

Our upper bounds are applications of the polynomial method [74, 6]: defining appropriate sparse polynomials for approximating Max-IP on small groups of vectors, and using fast matrix multiplication to speed up the evaluation of these polynomials on many pairs of points.

Via the known reduction from Max-IP to LCS-Pair in [5], we also obtain a more refined lower bound for approximating the LCS Closest Pair problem (defined below).

[8] They in fact consider the data structure version of a more fine-grained version of this problem, where there is an additional parameter s such that you want to decide whether $\mathsf{OPT}(A,B) \le t \cdot s$ or $\ge s$.

Definition 1.8 (LCS Closest Pair)**.** The $\mathsf{LCS\text{-}Closest\text{-}Pair}_{n,d}$ problem is: *given two sets A,B each consisting of n strings from Σ^d (Σ is a finite alphabet), determine*

$$\max_{a\in A,b\in B} \mathsf{LCS}(a,b),$$

where $\mathsf{LCS}(a,b)$ is the length of the longest common subsequence of strings a and b.

Corollary 1.9 (Improved Inapproximability for $\mathsf{LCS\text{-}Closest\text{-}Pair}$)**.** Assuming SETH (or OVC), for every $t \ge 2$, computing a multiplicative t-approximation to $\mathsf{LCS\text{-}Closest\text{-}Pair}_{n,d}$ requires $n^{2-o(1)}$ time, if $d = t^{\omega(1)} \cdot \log^5 n$.

1.3.2 Characterizations of hardness of additive approximations to Max-IP

Our characterization for additive approximations to $\mathsf{Max\text{-}IP}$ says that there is an additive t-approximation $n^{2-\Omega(1)}$-time algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if and only if $t = \Omega(d)$.

Theorem 1.10. Letting $\omega(\log n) < d < n^{o(1)}$ and $0 \le t \le d$, the following holds:

1. There is an $n^{2-\Omega(1)}$-time additive t-approximation algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if
$$t = \Omega(d),$$
and under SETH (or OVC), there is no $n^{2-\Omega(1)}$-time additive t-approximation algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if
$$t = o(d).$$

2. Moreover, letting $\varepsilon = t/d$, there is an
$$O\left(n^{2-\Omega(\varepsilon^{1/3}/\log \varepsilon^{-1})}\right)$$
time, additive t-approximation randomized algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ when $\varepsilon = \omega(\log^6 \log n/\log^3 n)$.

The lower bound above is already established in [67], while the upper bound works by reducing the problem to the $d = O(\log n)$ case via random-sampling coordinates, and solving the reduced problem via known methods [14, 13].

Remark 1.11. *We want to remark here that the lower bounds for approximate $\mathsf{Max\text{-}IP}$ are direct corollaries of the new* MA *protocols for Set-Disjointness in [67]. Our main contribution is providing the complementary* upper bounds *to show that these lower bounds are indeed* tight *assuming* SETH.

All-Pair-Max-IP. Finally, we remark that our algorithms (with slight adaptations) also work for the following stronger problem[9]: $\mathsf{All\text{-}Pair\text{-}Max\text{-}IP}_{n,d}$, in which we are given two sets A and B each consisting of n vectors from $\{0,1\}^d$, and for each $x \in A$ we must compute

$$\mathsf{OPT}(x,B) := \max_{y\in B} x \cdot y.$$

We say an algorithm is a multiplicative t-approximation (additive t-approximation) algorithm for $\mathsf{All\text{-}Pair\text{-}Max\text{-}IP}$, if for all instances of $\mathsf{OPT}(x,B)$, it computes corresponding approximate answers.

[9] Since All-Pair-Max-IP is stronger than Max-IP, lower bounds for Max-IP automatically apply to All-Pair-Max-IP.

Corollary 1.12. Suppose $\omega(\log n) < d < n^{o(1)}$, and let

$$\varepsilon_M := \min\left(\frac{\log t}{\log(d/\log n)}, 1\right) \text{ and } \varepsilon_A := \frac{\min(t,d)}{d}.$$

There is an $n^{2-\Omega(\varepsilon_M)}$ polylog(n)-time multiplicative t-approximation algorithm and an $n^{2-\Omega(\varepsilon_A^{1/3}/\log \varepsilon_A^{-1})}$-time additive t-approximation algorithm for All-Pair-Max-IP$_{n,d}$, when $\varepsilon_A = \omega(\log^6 \log n / \log^3 n)$.

1.4 BQP communication protocols and approximate $\{-1,1\}$-Max-IP

Making use of the $O(\sqrt{n})$-degree approximate polynomial for OR [27, 77], we also give a completely different proof for the hardness of multiplicative approximations to $\{-1,1\}$-Max-IP. Note that the answer to $\{-1,1\}$-Max-IP can be negative, so we consider the variant of maximizing *unsigned* inner product here; that is, we want to approximate $\max_{a\in A, b\in B} |a\cdot b|$ (this variant is also studied in [11]). Lower bound from that approach is inferior to Theorem 1.5: in particular, *it cannot achieve a characterization.*[10]

It is asked in [5] that whether we can make use of the $O(\sqrt{n})$ BQP communication protocol for Set-Disjointness [28] to prove conditional lower bounds. Indeed, that quantum communication protocol is based on the $O(\sqrt{n})$-time quantum query algorithm for OR (Grover's algorithm [42]), which induces the needed approximate polynomial for OR. Hence, the following theorem in some sense answers their question in the affirmative.

Theorem 1.13 (Informal)**.** Assuming SETH (or OVC), no $n^{2-\Omega(1)}$-time algorithm for

$$\{-1,1\}\text{-Max-IP}_{n,n^{o(1)}}$$

can achieve multiplicative $n^{o(1)}$-approximation.

The full statement can be found in Theorem 7.1 and Theorem 7.2.

1.5 Our results on $\mathbb{Z}$-Max-IP

1.5.1 Hardness of exact $\mathbb{Z}$-Max-IP in dimension $2^{O(\log^* n)}$

Now we turn to discussing our results on $\mathbb{Z}$-Max-IP. We show that $\mathbb{Z}$-Max-IP is hard to solve in $n^{2-\Omega(1)}$ time, even with $2^{O(\log^* n)}$-dimensional vectors.

Theorem 1.14. Assuming SETH (or OVC), there is a constant c such that any exact algorithm for $\mathbb{Z}$-Max-IP$_{n,d}$ in dimension $d = c^{\log^* n}$ requires $n^{2-o(1)}$ time, with vectors of $O(\log n)$-bit entries.

As direct corollaries of the above theorem, using reductions implicit in [75], we also conclude hardness for ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair under SETH (or OVC) in dimension $2^{O(\log^* n)}$.

[10] We also remark that in particular, the hardness of approximate $\{-1,1\}$-Max-IP follows from the hardness of approximate $\{0,1\}$-Max-IP. We include the proof via BQP communication protocols here as it is an interesting alternative proof.

Theorem 1.15 (Hardness of ℓ_2-Furthest Pair in Dimension $c^{\log^* n}$)**.** Assuming SETH (or OVC), there is a constant c such that ℓ_2-Furthest Pair in dimension $c^{\log^* n}$ requires $n^{2-o(1)}$ time, with vectors of $O(\log n)$-bit entries.

Theorem 1.16 (Hardness of Bichromatic ℓ_2-Closest Pair in Dimension $c^{\log^* n}$)**.** Assuming SETH (or OVC), there is a constant c such that Bichromatic ℓ_2-Closest Pair in dimension $c^{\log^* n}$ requires $n^{2-o(1)}$ time, with vectors of $O(\log n)$-bit entries.

The above lower bounds on ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair are in sharp contrast with the case of ℓ_2-*Closest Pair*, which can be solved in $2^{O(d)} \cdot n \log^{O(1)} n$ time [23, 53, 38].

1.5.2 Improved dimensionality reduction for OV and Hopcroft's problem

Our hardness of $\mathbb{Z}$-Max-IP is established by a reduction from Hopcroft's problem, whose hardness is in turn derived from the following significantly improved dimensionality reduction for OV.

Lemma 1.17 (Improved Dimensionality Reduction for OV)**.** Let $1 \le \ell \le d$. There is an

$$O\left(n \cdot \ell^{O(6^{\log^* d} \cdot (d/\ell))} \cdot \operatorname{poly}(d)\right)\text{-time}$$

reduction from $\mathsf{OV}_{n,d}$ to $\ell^{O(6^{\log^* d} \cdot (d/\ell))}$ instances of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$, with vectors of entries with bit-length $O(d/\ell \cdot \log \ell \cdot 6^{\log^* d})$.

Comparison with [75]. Comparing to the old construction in [75], our reduction here is more efficient when ℓ is much smaller than d (which is the case we care about). That is, in [75], $\mathsf{OV}_{n,d}$ can be reduced to $d^{d/\ell}$ instances of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$, while we get $\{\ell^{6^{\log^* d}}\}^{d/\ell}$ instances in our improved one. So, for example, when $\ell = 7^{\log^* d}$, the old reduction yields $d^{d/7^{\log^* d}} = n^{\omega(1)}$ instances (recall that $d = c \log n$ for an arbitrary constant c), while our improved one yields only $n^{o(1)}$ instances, each with $2^{O(\log^* n)}$ dimensions.

From Lemma 1.17, the following theorem follows in the same way as in [75].

Theorem 1.18 (Hardness of Hopcroft's Problem in Dimension $c^{\log^* n}$)**.** Assuming SETH (or OVC), there is a constant c such that $\mathbb{Z}$-$\mathsf{OV}_{n,c^{\log^* n}}$ with vectors of $O(\log n)$-bit entries requires $n^{2-o(1)}$ time.

1.5.3 Connection between $\mathbb{Z}$-Max-IP lower bounds and $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocols

We also show a new connection between $\mathbb{Z}$-Max-IP and a special type of communication protocol. Let us first recall the Set-Disjointness problem.

Definition 1.19 (Set-Disjointness)**.** Let $n \in \mathbb{N}$. In Set-Disjointness (DISJ_n), Alice holds a vector $X \in \{0,1\}^n$, Bob holds a vector $Y \in \{0,1\}^n$, and they want to determine whether $X \cdot Y = 0$.

In [5], the hardness of approximate Max-IP is established via a connection to MA communication protocols (in particular, an MA communication protocol with small communication complexity for Set-Disjointness). Our lower bound for (exact) $\mathbb{Z}$-Max-IP can also be connected to similar $\mathsf{NP} \cdot \mathsf{UPP}$ protocols (note that $\mathsf{MA} = \mathsf{NP} \cdot \mathsf{promiseBPP}$).

Formally, we define $\mathsf{NP} \cdot \mathsf{UPP}$ protocols as follows.[11]

Definition 1.20. For a problem Π with two n-bit strings x, y as inputs (Alice holds x and Bob holds y), we say a communication protocol is an (m,ℓ)*-efficient* $\mathsf{NP} \cdot \mathsf{UPP}$ *communication protocol* if the following holds:

- There are three parties Alice, Bob and Merlin in the protocol.
- Merlin sends Alice and Bob an advice string z of length m, which is a function of x and y.
- Given y and z, Bob sends Alice ℓ bits, and Alice decides to accept or not.[12] They have an unlimited supply of private random coins (not public, which is important) during their conversation. The following conditions hold:
    - If $\Pi(x,y) = 1$, then there is an advice z from Merlin such that Alice accepts with probability $\geq 1/2$.
    - Otherwise, for all possible advice strings from Merlin, Alice accepts with probability $< 1/2$.

Moreover, we say the protocol is (m,ℓ)-computational-efficient, if in addition the probability distributions of both Alice's and Bob's behavior can be computed in $\mathrm{poly}(n)$ time given their input and the advice.

Our new reduction from OV to $\mathsf{Max\text{-}IP}$ actually implies an efficient $\mathsf{NP} \cdot \mathsf{UPP}$ protocol for Set-Disjointness.

Theorem 1.21. For every $1 \leq \alpha \leq n$, there is an

$$\left(\alpha \cdot 6^{\log^* n} \cdot (n/2^{\alpha}), O(\alpha)\right)\text{-computational-efficient}$$

$\mathsf{NP} \cdot \mathsf{UPP}$ communication protocol for DISJ_n.

For example, when $\alpha = 3\log^* n$, Theorem 1.21 implies there is an $O(o(n), O(\log^* n))$-computational-efficient $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocol for DISJ_n. Moreover, we show that if the protocol of Theorem 1.21 can be improved a little bit (like removing the $6^{\log^* n}$ term), we would obtain the desired hardness for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$ in dimension $\omega(1)$.

Theorem 1.22. Assuming SETH (or OVC), if there is an increasing and unbounded function f such that for every $1 \leq \alpha \leq n$, there is an

$$(n/f(\alpha), \alpha)\text{-computational-efficient}$$

$\mathsf{NP} \cdot \mathsf{UPP}$ communication protocol for DISJ_n, then $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,\omega(1)}$ requires $n^{2-o(1)}$ time with vectors of $\mathrm{polylog}(n)$-bit entries. The same holds for ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair.

[11] Here are some comments on the name $\mathsf{NP} \cdot \mathsf{UPP}$. Roughly speaking, a UPP protocol Π is a private-coin randomized communication protocol in which Alice and Bob accept with probability $> 1/2$ if and only if the answer is yes. For a communication protocol class $\mathcal{D}$, $\mathsf{NP} \cdot \mathcal{D}$ denotes a new class of protocol resembling a Merlin-Arthur game: A prover (who knows the inputs of Alice and Bob) sends a proof to both Alice and Bob first; then Alice and Bob run a prescribed $\mathcal{D}$ protocol on their inputs and the proof to decide whether they accept the proof or not. The protocol needs to satisfy the soundness and completeness conditions similar to an original MA protocol.

[12] Our simplification here is justified by the fact that one-way UPP communication protocols are equivalent to their seemingly more powerful two-way analogs [63].

1.6 Improved MA protocols for Set-Disjointness

Finally, we also obtain a new MA protocol for Set-Disjointness, which improves on the previous $O(\sqrt{n}\log n)$ protocol in [1], and is closer to the $\Omega(\sqrt{n})$ lower bound by [54]. Like the protocol in [1], our new protocol also works for the following slightly harder problem Inner Product.

Definition 1.23 (Inner Product)**.** Let $n \in \mathbb{N}$. In Inner Product (IP_n), Alice holds a vector $X \in \{0,1\}^n$, Bob holds a vector $Y \in \{0,1\}^n$, and they want to compute $X \cdot Y$.

Theorem 1.24. *There is an* MA *protocol for* DISJ_n *and* IP_n *with communication complexity*

$$O\left(\sqrt{n\log n\log\log n}\right).$$

In [67], the author asked whether the MA communication complexity of DISJ_n (IP_n) is $\Theta(\sqrt{n})$ or $\Theta(\sqrt{n}\log n)$. Our result makes progress on that question by showing that the true complexity lies between $\Theta(\sqrt{n})$ and $\Theta(\sqrt{n\log n\log\log n})$.

1.7 Intuition for dimensionality self-reduction for OV

The $2^{O(\log^* n)}$ factor in Lemma 1.17 is not common in theoretical computer science,[13] and our new reduction for OV is considerably more complicated than the polynomial-based construction from [75]. Hence, it is worth discussing the intuition behind Lemma 1.17, and the reason why we get a factor of $2^{O(\log^* n)}$.

A direct approach based on the Chinese Remainder Theorem. We first discuss a direct reduction based on the *Chinese Remainder Theorem* (CRT) (see Theorem 2.5 for a formal definition). CRT says that given a collection of distinct primes $q_1,\dots,q_b$, and a collection of integers $r_1,\dots,r_b$, there exists a unique integer $t = \mathsf{CRR}(\{r_j\};\{q_j\})$ such that $t \equiv r_j \pmod{q_j}$ for each $j \in [b]$ and $0 \le t < \prod_{j=1}^{b} q_j$ (CRR stands for *Chinese Remainder Representation*).

Now, let $b,\ell \in \mathbb{N}$. Suppose we would like to have a dimensionality reduction φ from $\{0,1\}^{b\cdot\ell}$ to $\mathbb{Z}^{\ell}$. We can partition an input $x \in \{0,1\}^{b\cdot\ell}$ into ℓ blocks, each of length b, and represent each block via CRT: that is, for a block $z \in \{0,1\}^b$, we map it into a single integer $\varphi_{\mathsf{block}}(z) := \mathsf{CRR}(\{z_j\};\{q_j\})$, and the concatenations of φ_{block} over all blocks of x is $\varphi(x) \in \mathbb{Z}^{\ell}$.

The key idea here is that, for $z,z' \in \{0,1\}^b$, $\varphi_{\mathsf{block}}(z)\cdot\varphi_{\mathsf{block}}(z') \pmod{q_j}$ is simply $z_j \cdot z'_j$. That is, the multiplication between two *integers* $\varphi_{\mathsf{block}}(z)\cdot\varphi_{\mathsf{block}}(z')$ simulates the coordinate-wise multiplication between two *vectors* z and z'!

Therefore, if we make all primes q_j larger than ℓ, we can in fact determine $x \cdot y$ from $\varphi(x)\cdot\varphi(y)$, by looking at $\varphi(x)\cdot\varphi(y) \pmod{q_j}$ for each j. That is,

$$x\cdot y = 0 \Leftrightarrow \varphi(x)\cdot\varphi(y) \equiv 0 \pmod{q_j} \quad \text{for every } j \in [b].$$

[13] Other examples include an $O\big(2^{O(\log^* n)}n^{4/3}\big)$-time algorithm for $\mathbb{Z}$-$\mathsf{OV}_{n,3}$ [59], $O\big(2^{O(\log^* n)}n\log n\big)$-time algorithms (Fürer's algorithm with its modifications) for Fast Integer Multiplication [39, 36, 44] and an old $O(n^{d/2}2^{O(\log^* n)})$-time algorithm for Klee's measure problem [30].

Hence, let V be the set of all integers$0 \le v \le \ell \cdot \left(\prod_{j=1}^{b} q_j\right)^2$ that $v \equiv 0 \pmod{q_j}$ for every $j \in [b]$. We have

$$x \cdot y = 0 \Leftrightarrow \varphi(x) \cdot \varphi(y) \in V.$$

The reduction is completed by constructing a $\mathbb{Z}$-OV instance for each $v \in V$: we append corresponding values to make $\varphi_A(x) = [\varphi(x), -1]$ and $\varphi_B(y) = [\varphi(y), v]$ (this step is from [75]).

Note that a nice property for φ is that each $\varphi(x)_i$ only depends on the i-th block of x, and the mapping is the same on each block (φ_{block}); we call this the *block mapping property*.

Analysis of the direct reduction. To continue building intuition, let us analyze the above reduction. The size of V is the number of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$ instances we create, and $|V| \ge \prod_{j=1}^{b} q_j$. These primes q_j have to be all distinct, and it follows that $\prod_{j=1}^{b} q_j$ is $b^{\Theta(b)}$. Since we want to create at most $n^{o(1)}$ instances (or n^{ε} for arbitrarily small ε), we need to set $b \le \log n / \log\log n$. Moreover, to base our hardness on OVC which deals with $c \log n$-dimensional vectors, we need to set $b \cdot \ell = d = c \cdot \log n$ for an arbitrary constant c. Therefore, we must have $\ell \ge \log\log n$, and the above reduction only obtains the same hardness result as [75].

Key observation: "Most space modulo q_j" is actually wasted. To improve the above reduction, we need to make $|V|$ smaller. Our key observation about φ is that, for the primes q_j, they are mostly larger than $b \gg \ell$, but $\varphi(x) \cdot \varphi(y) \in \{0, 1, \ldots, \ell\} \pmod{q_j}$ for all these q_j. Hence, *"most space modulo q_j" is actually wasted.*

Make more "efficient" use of the "space": recursive reduction. Based on the previous observation, we want to use the "space modulo q_j" more efficiently. It is natural to consider a *recursive reduction*. We will require all our primes q_j to be larger than b. Let b_{m} be a very small integer compared to b, and $\psi : \{0,1\}^{b_{\mathsf{m}} \cdot \ell} \to \mathbb{Z}^{\ell}$ with a set V_{ψ} and a block mapping ψ_{block} be a similar reduction on much smaller inputs: for $x, y \in \{0,1\}^{b_{\mathsf{m}} \cdot \ell}$, $x \cdot y = 0 \Leftrightarrow \psi(x) \cdot \psi(y) \in V_{\psi}$. We also require here that $\psi(x) \cdot \psi(y) \le b$ for every x and y.

For an input $x \in \{0,1\}^{b \cdot \ell}$ and a block $z \in \{0,1\}^{b}$ of x, our key idea is to partition z again into b/b_{m} "micro" blocks each of size b_{m}. And for a block z in x, let $z^1, \ldots, z^{b/b_{\mathsf{m}}}$ be its b/b_{m} micro blocks. We map z into an integer $\varphi_{\mathsf{block}}(z) := \mathsf{CRR}(\{\psi_{\mathsf{block}}(z^j)\}_{j=1}^{b/b_{\mathsf{m}}}; \{q_j\}_{j=1}^{b/b_{\mathsf{m}}})$.

Now, given two blocks $z, w \in \{0,1\}^{b}$, we can see that

$$\varphi_{\mathsf{block}}(z) \cdot \varphi_{\mathsf{block}}(w) \equiv \psi_{\mathsf{block}}(z^j) \cdot \psi_{\mathsf{block}}(w^j) \pmod{q_j}.$$

For two inputs $x, y \in \{0,1\}^{b \cdot \ell}$, for $(i, j) \in [\ell] \times [b/b_{\mathsf{m}}]$, we use $x^{i,j} \in \{0,1\}^{b_{\mathsf{m}}}$ ($y^{i,j}$) to denote the j-th micro block in the i-th block of x (y, respectively). We also define $x^{[j]} \in \{0,1\}^{b_{\mathsf{m}} \cdot \ell}$ as the concatenation of $x^{1,j}, x^{2,j}, \ldots, x^{\ell,j}$ ($y^{[j]}$ is defined similarly). Then we have

$$\begin{aligned} \varphi(x) \cdot \varphi(y) &\equiv \sum_{i=1}^{\ell} \psi_{\mathsf{block}}(x^{i,j}) \cdot \psi_{\mathsf{block}}(y^{i,j}) \pmod{q_j} \\ &= \psi(x^{[j]}) \cdot \psi(y^{[j]}). && (\psi(x^{[j]}) \cdot \psi(y^{[j]}) \le b < q_j) \end{aligned}$$

Hence, for every $j \in [b/b_{\mathsf{m}}]$, we can determine whether $x^{[j]} \cdot y^{[j]} = 0$ from whether $\varphi(x) \cdot \varphi(y) \pmod{q_j} \in V_{\varphi}$, and therefore also determine whether $x \cdot y = 0$ from $\varphi(x) \cdot \varphi(y)$.

We can now observe that $|V| \le b^{\Theta(b/b_{\mathsf{m}})}$, smaller than before; thus we get an improvement, depending on how large can b_{m} be. Clearly, the reduction ψ can also be constructed from even smaller reductions, and after recursing $\Theta(\log^* n)$ times, we can switch to the direct construction discussed before. By a straightforward (but tedious) calculation, we can derive Lemma 1.17.

High-level explanation on the $2^{O(\log^* n)}$ factor. Ideally, we want to have a reduction from OV to $\mathbb{Z}$-OV with only $\ell^{O(b)}$ instances, in other words, we want $|V| = \ell^{O(b)}$. The reason we need to pay an extra $2^{O(\log^* n)}$ factor in the exponent is as follows:

In our reduction, $|V|$ is at least $\prod_{j=1}^{b/b_{\mathsf{m}}} q_j$, which is also the bound on each coordinate of the reduction: $\varphi(x)_i$ equals to a CRR encoding of a vector with $\{q_j\}_{j=1}^{b/b_{\mathsf{m}}}$, whose value can be as large as $\prod_{j=1}^{b/b_{\mathsf{m}}} q_j - 1$. That is, all we want is to control the upper bound on the coordinates of the reduction.

Suppose we are constructing an "outer" reduction $\varphi : \{0,1\}^{b\cdot\ell} \to \mathbb{Z}^{\ell}$ from the "micro" reduction $\psi : \{0,1\}^{b_{\mathsf{m}}\cdot\ell} \to \mathbb{Z}^{\ell}$ with coordinate upper bound L_ψ ($\psi(x)_i \le L_\psi$), and let $L_\psi = \ell^{\kappa\cdot b_{\mathsf{m}}}$. (That is, κ is the extra factor comparing to the ideal case.) Recall that we have to ensure $q_j > \psi(x)\cdot\psi(y)$ to make our construction work, and therefore we have to set q_j larger than L_ψ^2.

Then the coordinate upper bound for φ becomes $L_\varphi = \prod_{j=1}^{b/b_{\mathsf{m}}} q_j \ge (L_\psi)^{2\cdot b/b_{\mathsf{m}}} = \ell^{2\kappa\cdot b}$. Therefore, we can see that after one recursion, the "extra factor" κ at least doubles. Since our recursion proceeds in $\Theta(\log^* n)$ rounds, we have to pay an extra $2^{O(\log^* n)}$ factor on the exponent.

1.8 Related work

SETH-based conditional lower bound. SETH is one of the most fruitful conjectures in the Fine-Grained Complexity. There are numerous conditional lower bounds based on it for problems in P among different areas, including: dynamic data structures [61, 7, 45, 55, 3, 46, 41], computational geometry [24, 37, 75, 67], pattern matching [8, 22, 21, 25, 26], graph algorithms [66, 40, 9, 56]. See [72] for a recent survey on SETH-based lower bounds (and more).

Communication Complexity and conditional hardness. The connection between communication protocols (in various model) for Set-Disjointness and SETH dates back at least to [62], in which it is shown that a sublinear computational efficient protocol for 3-party Number-On-Forehead Set-Disjointness problem would refute SETH. And it is worth mentioning that Abboud and Rubinstein's result [4] builds on the $\widetilde{O}(\log n)$ IP communication protocol for Set-Disjointness in [1]. Making use of the IP communication protocol for low-space computation, [31] establish an equivalence class for LCS-$\mathsf{Closest}$-Pair.

In [32], Σ_2 communication protocols are utilized to show the subquadratic-time equivalence between $\mathsf{OV}_{n,O(\log n)}$, Max-$\mathsf{IP}_{n,O(\log n)}$, Approximate Bichromatic Closest Pair and several other problems.

Distributed PCP. Using Algebraic Geometry codes (AG codes), [67] obtains a better MA protocol, which in turn improves the efficiency of the previous distributed PCP construction of [5]. He then shows the $n^{2-o(1)}$-time hardness for $1+o(1)$-approximation to Bichromatic Closest Pair and $o(d)$-additive approximation to Max-$\mathsf{IP}_{n,d}$ with this new technique.

[51] use the Distributed PCP framework to derive inapproximability results for k-Dominating Set under various assumptions. In particular, building on the techniques of [67], it is shown that under SETH,

k-Dominating Set has no $(\log n)^{1/\operatorname{poly}(k,e(\varepsilon))}$ approximation in $n^{k-\varepsilon}$ time.[14]

[52] also utilize AG codes and polynomial method to show hardness results for Exact and Approximate Monochromatic Closest Pair and Approximate Monochromatic Maximum Inner Product.

Hardness of approximation in P. Making use of Chebyshev embeddings, [11] prove a $2^{\Omega\left(\frac{\sqrt{\log n}}{\log\log n}\right)}$ inapproximability lower bound on $\{-1,1\}$-Max-IP. [2] take an approach different from Distributed PCP, and shows that under certain complexity assumptions, LCS does not have a *deterministic* $1+o(1)$-approximation in $n^{2-\varepsilon}$ time. They also establish a connection with circuit lower bounds and show that the existence of such a *deterministic* algorithm implies E^{NP} does not have non-uniform linear-size Valiant Series Parallel circuits. In [4], it is improved to that any constant factor approximation deterministic algorithm for LCS in $n^{2-\varepsilon}$ time implies that E^{NP} does not have non-uniform linear-size NC^1 circuits. See [5] for more related results in hardness of approximation in P.

Organization of the paper

In Section 2, we introduce the needed preliminaries for this paper. In Section 3, we prove our characterizations for approximate Max-IP and other related results. In Section 4, we prove $2^{O(\log^* n)}$ dimensional hardness for $\mathbb{Z}$-Max-IP and other related problems. In Section 5, we establish the connection between $\mathsf{NP}\cdot\mathsf{UPP}$ communication protocols and SETH-based lower bounds for exact $\mathbb{Z}$-Max-IP. In Section 6, we present the $O\left(\sqrt{n\log n\log\log n}\right)$ MA protocol for Set-Disjointness.

2 Preliminaries

We begin by introducing some notation. For an integer d, we use $[d]$ to denote the set of integers from 1 to d. For a vector u, we use u_i to denote the i-th element of u.

We use $\log(x)$ to denote the logarithm of x with respect to base 2 and $\ln(x)$ to denote the natural logarithm of x.

In our arguments, we use the iterated logarithm function $\log^*(n)$, which is defined recursively as follows:

$$\log^*(n) := \begin{cases} 0 & n \le 1; \\ \log^*(\log n) + 1 & n > 1. \end{cases}$$

2.1 Fast rectangular matrix multiplication

Similarly to previous algorithms using the polynomial method, our algorithms make use of the algorithms for fast rectangular matrix multiplication.

Theorem 2.1 ([57])**.** There is an $N^{2+o(1)}$-time algorithm for multiplying two matrices A and B with size $N \times N^{\alpha}$ and $N^{\alpha} \times N$, where $\alpha > 0.31389$.

Theorem 2.2 ([35])**.** There is an $N^2 \cdot \operatorname{polylog}(N)$-time algorithm for multiplying two matrices A and B with size $N \times N^{\alpha}$ and $N^{\alpha} \times N$, where $\alpha > 0.172$.

[14] where $e : \mathbb{R}^+ \to \mathbb{N}$ is some function

2.2 Number theory

Here we recall some facts from number theory. In our reduction from OV to $\mathbb{Z}$-OV, we will apply the famous prime number theorem, which supplies a good estimate of the number of primes smaller than a certain number. See e.g. [19] for a reference on this.

Theorem 2.3 (Prime Number Theorem)**.** Let $\pi(n)$ be the number of primes $\leq n$. We have

$$\lim_{n\to\infty} \frac{\pi(n)}{n/\ln n} = 1.$$

From a simple calculation, we have the following lemma.

Lemma 2.4. There are $10n$ distinct primes in $[n+1, n^2]$ for all sufficiently large n.

Proof. For a sufficiently large n, from the prime number theorem, the number of primes in $[n+1, n^2]$ is equal to

$$\pi(n^2) - \pi(n) \sim n^2/2\ln n - n/\ln n \gg 10n. \qquad \square$$

Next we recall the Chinese remainder theorem, and Chinese remainder representation.

Theorem 2.5. Given d pairwise co-prime integers $q_1, q_2, \ldots, q_d$, and d integers $r_1, r_2, \ldots, r_d$, there is exactly one integer $0 \leq t < \prod_{i=1}^{d} q_i$ such that

$$t \equiv r_i \pmod{q_j} \quad \text{for every } i \in [d].$$

We call this t the Chinese remainder representation (or the CRR encoding) of the r_i (with respect to these q_i). We also denote

$$t = \mathsf{CRR}(\{r_i\}; \{q_i\})$$

for convenience. We sometimes omit the sequence $\{q_i\}$ for simplicity, when it is clear from the context.

Moreover, t can be computed in polynomial time with respect to the total bits of all the given integers.

2.3 Communication complexity

In our paper we will make use of a certain kind of MA protocol, we call them (m, r, ℓ, s)-efficient protocols.[15]

Definition 2.6. We say an MA Protocol is (m, r, ℓ, s)-efficient for a communication problem, if in the protocol:

- There are three parties Alice, Bob and Merlin in the protocol, Alice holds input x and Bob holds input y.
- Merlin sends an advice string z of length m to Alice, which is a function of x and y.

[15] Our notation here is adopted from [51]. They also defined similar k-party communication protocols, while we only discuss 2-party protocols in this paper.

- Alice and Bob jointly toss r coins to obtain a random string w of length r.
- Given y and w, Bob sends Alice a message of length ℓ.
- After that, Alice decides whether to accept or not.
  - When the answer is yes, Merlin has exactly one advice such that Alice always accept.
  - When the answer is no, or Merlin sends the wrong advice, Alice accepts with probability at most s.

2.4 Derandomization

We make use of expander graphs to reduce the amount of random coins needed in one of our communication protocols. We abstract the following result for our use here.

Theorem 2.7 (see, e. g., Theorem 21.12 and Theorem 21.19 in [20])**.** Let m be an integer. There is a universal constant c_1 such that for every $\varepsilon < 1/2$, there is a $\operatorname{poly}(\log m, \log \varepsilon^{-1})$-time computable function $\mathcal{F} : \{0,1\}^{\log m + c_1 \cdot \log \varepsilon^{-1}} \to [m]^{c_1 \cdot \log \varepsilon^{-1}}$, such that for every set $B \subseteq [m]$ of size at least $m/2$,

$$\Pr_{w \in \{0,1\}^{\log m + c_1 \cdot \log \varepsilon^{-1}}} [a \notin B \text{ for every } a \in \mathcal{F}(w)] \leq \varepsilon,$$

here $a \in \mathcal{F}(w)$ means a is one of the elements in the sequence $\mathcal{F}(w)$.

3 Hardness of Approximate Max-IP

In this section we prove our characterizations of approximate Max-IP.

3.1 The multiplicative case

We begin with the proof of Theorem 1.5. In Lemma 3.2, we construct the desired approximation algorithm and in Corollary 3.4 we prove the lower bound.

First we need the following simple lemma, which says that the k-th root of the sum of the k-th powers of non-negative reals gives a good approximation to their maximum.

Lemma 3.1. Let S be a set of non-negative real numbers, k be an integer, and $x_{\max} := \max_{x \in S} x$. We have

$$\left(\sum_{x \in S} x^k\right)^{1/k} \in \left[x_{\max}, x_{\max} \cdot |S|^{1/k}\right].$$

Proof. Since

$$\left(\sum_{x \in S} x^k\right) \in \left[x_{\max}^k, |S| \cdot x_{\max}^k\right],$$

the lemma follows directly by taking the k-th root of both sides.

□

Lemma 3.2. Assuming $\omega(\log n) < d < n^{o(1)}$ and letting

$$\varepsilon = \min\left(\frac{\log t}{\log(d/\log n)}, 1\right),$$

there are multiplicative t-approximation deterministic algorithms for $\mathbb{R}^{+}$-Max-$\mathsf{IP}_{n,d}$ running in time[16]

$$O\left(n^{2+o(1)-0.31\cdot\frac{1}{\varepsilon^{-1}+\frac{0.31}{2}}}\right) = O\left(n^{2+o(1)-\Omega(\varepsilon)}\right)$$

or time

$$O\left(n^{2-0.17\cdot\frac{1}{\varepsilon^{-1}+\frac{0.17}{2}}}\cdot\mathrm{polylog}(n)\right) = O\left(n^{2-\Omega(\varepsilon)}\cdot\mathrm{polylog}(n)\right).$$

Proof. Let $d = c\cdot\log n$. From the assumption, we have $c = \omega(1)$, and $\varepsilon = \min\left(\log(t)/\log(c), 1\right)$. When $\log t > \log c$, we simply use a multiplicative c-approximation algorithm instead, hence in the following we assume $\log t \leq \log c$. We begin with the first algorithm here.

Construction and analysis of the Power of Sum Polynomial $P_r(z)$. Let r be a parameter to be specified later and z be a vector from $(\mathbb{R}^{+})^{d}$. We define a polynomial $P_r(z)$ as

$$P_r(z) := \left(\sum_{i=1}^{d} z_i\right)^{r}.$$

Let $E := \{(e_1, e_2, \ldots, e_d) \mid \sum_{i=1}^{d} e_i = r, \text{the } e_i \text{ are non-negative integers}\}$.
We have

$$|E| = \binom{r+d-1}{d-1} = \binom{r+d-1}{r}.$$

For each $e \in E$, we define $z^{e} := \prod_{i=1}^{d} z_i^{e_i}$. Now, by expanding out the polynomial, we can write $P_r(z)$ as

$$P_r(z) = \sum_{e\in E} c_e \cdot z^{e},$$

where the c_e are the corresponding coefficients. Then consider $P_r(x,y) := P_r(x_1\cdot y_1, x_2\cdot y_2, \ldots, x_d\cdot y_d)$, plugging in $z_i := x_i\cdot y_i$, it can be written as

$$P_r(x,y) := \sum_{e\in E} c_e \cdot x^{e}\cdot y^{e},$$

where x^{e} and y^{e} are defined similarly as z^{e}.

Construction and analysis of the Batch Evaluation Polynomial $P_r(X,Y)$. Now, let X and Y be two sets each consisting of $b = t^{r/2}$ vectors from $\{0,1\}^{d}$. We define[17]

$$P_r(X,Y) := \sum_{x\in X, y\in Y} P_r(x,y) = \sum_{x\in X, y\in Y} (x\cdot y)^{r}.$$

[16] In the following we assume a real RAM model of computation for simplicity.

[17] We remark that similar polynomials are also used in [12] to give a simple algorithm for solving the light bulb problem.

By Lemma 3.1, we have

$$P_r(X,Y)^{1/r} \in [\mathsf{OPT}(X,Y), \mathsf{OPT}(X,Y) \cdot t],$$

recall that $\mathsf{OPT}(X,Y) := \max_{x \in X, y \in Y} x \cdot y$.

Embedding into Rectangular Matrix Multiplication. Now, for $x, y \in \{0,1\}^d$, we define the mappings $\phi_x(x)$ and $\phi_y(y)$ as,

$$\phi_x(x) := (c_{e_1} \cdot x^{e_1}, c_{e_2} \cdot x^{e_2}, \dots, c_{e_m} \cdot x^{e_m})$$

and

$$\phi_y(y) := (y^{e_1}, y^{e_2}, \dots, y^{e_m}),$$

where $m = |E|$ and $e_1, e_2, \dots, e_m$ is an enumeration of all vectors in E.

From the definition, it follows that

$$\phi_x(x) \cdot \phi_y(y) = P_r(x,y)$$

for every $x, y \in \{0,1\}^d$.

Then for each X and Y, we map them into m-dimensional vectors $\phi_X(X)$ and $\phi_Y(Y)$ simply by a summation:

$$\phi_X(X) := \sum_{x \in X} \phi_x(x) \quad \text{and} \quad \phi_Y(Y) := \sum_{y \in Y} \phi_y(y).$$

We can see

$$\phi_X(X) \cdot \phi_Y(Y) = \sum_{x \in X} \phi_x(x) \cdot \sum_{y \in Y} \phi_y(y) = \sum_{x \in X} \sum_{y \in Y} P_r(x,y) = P_r(X,Y).$$

Given two sets A, B each consisting of n vectors from $\{0,1\}^d$, we split A into n/b sets $A_1, A_2, \dots, A_{n/b}$ of size b, and split B in the same way as well. Then we construct a matrix $M_A(M_B)$ of size $n/b \times m$, such that the i-th row of $M_A(M_B)$ is the vector $\phi_X(A_i)(\phi_Y(B_i))$. After that, the evaluation of $P_r(A_i, B_j)$ for all integers $i, j \in [n/b]$ can be reduced to computing the matrix product $M_A \cdot M_B^T$. After knowing all the $P_r(A_i, B_j)$, we simply compute their maximum, whose r-th root gives us a t-approximate answer of the original problem.

Analysis of the running time. Finally, we are going to specify the parameter r and analyze the time complexity. In order to utilize the fast matrix multiplication algorithm from Theorem 2.1, we need to have

$$m \le (n/b)^{0.313},$$

then our running time is simply $(n/b)^{2+o(1)} = n^{2+o(1)}/b^2$.

We are going to set $r = k \cdot \log n / \log c$, and our choice of k will satisfy $k = \Theta(1)$ and $r \le d$. We have

$$m \le \left(\frac{\mathrm{e} \cdot (r+d)}{r} \right)^r \le \left(\frac{\mathrm{e} \cdot 2d}{r} \right)^r \le \left(\frac{2c \log n \cdot \mathrm{e}}{k \cdot \log n / \log c} \right)^{k \cdot \log n / \log c},$$

and therefore

$$\log m \le k \cdot \log n \left[\log \frac{2c \log c}{k} + \log \mathrm{e} \right] \Big/ \log c.$$

Since $c = \omega(1)$ and $k = \Theta(1)$, we have

$$\log m \leq (1+o(1)) \cdot k \log n = k \log n + o(\log n).$$

Plugging in, we have

$$\begin{aligned}
& m \leq (n/b)^{0.313} \\
\Longleftarrow\ & \log m \leq 0.313 \cdot (\log n - \log b) \\
\Longleftarrow\ & k \log n \leq 0.31 \cdot (\log n - \log b) \\
\Longleftarrow\ & 0.31 \cdot (r/2) \cdot \log t + k \log n \leq 0.31 \log n && (b = t^{r/2}) \\
\Longleftarrow\ & \frac{\log n}{\log c} \cdot k \cdot \log t \cdot \frac{0.31}{2} + k \log n \leq 0.31 \log n && (r = k \cdot \log n / \log c) \\
\Longleftarrow\ & k \cdot \left\{ 1 + \frac{\log t}{\log c} \cdot \frac{0.31}{2} \right\} \leq 0.31 \\
\Longleftarrow\ & k = \frac{0.31}{1 + \frac{\log t}{\log c} \cdot \frac{0.31}{2}} = \frac{0.31}{1 + \frac{0.31}{2} \cdot \varepsilon}.
\end{aligned}$$

Note since $\varepsilon \in [0,1]$, k is indeed $\Theta(1)$.

Finally, with our choice of k specified, our running time is $n^{2+o(1)}/b^2 = n^{2+o(1)}/t^r$.

By a simple calculation,

$$\begin{aligned}
\log t^r &= r \cdot \log t \\
&= k \cdot \log n / \log c \cdot \log t \\
&= \log n \cdot \left\{ \frac{\log t}{\log c} \cdot \frac{0.31}{1 + \frac{0.31}{2} \cdot \varepsilon} \right\} \\
&= \log n \cdot \frac{0.31\varepsilon}{1 + \frac{0.31}{2} \cdot \varepsilon} \\
&= \log n \cdot \frac{0.31}{\varepsilon^{-1} + \frac{0.31}{2}}.
\end{aligned}$$

Hence, our running time is

$$\frac{n^{2+o(1)}}{t^r} = n^{2+o(1) - \frac{0.31}{\varepsilon^{-1} + \frac{0.31}{2}}}$$

as stated.

The second algorithm. The second algorithm follows exactly the same except for that we apply Theorem 2.2 instead, hence the constant 0.31 is replaced by 0.17. □

The lower bound follows directly from the new MA protocol for Set-Disjointness in [67]. We present an explicit proof here for completeness.

Before proving the lower bound, we need the following reduction from OV to approximate Max-IP.

Lemma 3.3 (Implicit in Theorem 4.1 of [67]). There is a universal constant c_1 such that, for every integer c, real$\varepsilon \in (0,1]$ and $\tau \geq 2$, $\mathsf{OV}_{n,c\log n}$ can be reduced to n^{ε} $\mathsf{Max\text{-}IP}_{n,d}$ instances (A_i, B_i) for $i \in [n^{\varepsilon}]$, such that:

- $d = \tau^{\mathrm{poly}(c/\varepsilon)} \cdot \log n$.
- Letting $T = c \log n \cdot \tau^{c_1}$, if there is an $a \in A$ and $b \in B$ such that $a \cdot b = 0$, then there exists an i such that $\mathsf{OPT}(A_i, B_i) \geq T$.
- Otherwise, for every $i \in [n^{\varepsilon}]$ we must have $\mathsf{OPT}(A_i, B_i) \leq T/\tau$.

The reduction above follows directly from the new MA communication protocols in [67] together with the use of expander graphs to reduce the amount of random coins. A proof for the lemma is deferred to Section 3.5.

Now we are ready to show the lower bound on approximate $\mathsf{Max\text{-}IP}$.

Corollary 3.4. Assuming SETH (or OVC) and letting $d = \omega(\log n)$ and $t \geq 2$, no $n^{2-\Omega(1)}$-time algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ can achieve multiplicative t-approximation if

$$t = (d/\log n)^{o(1)}.$$

Proof. Let $c = d/\log n$. Note that $t = c^{o(1)}$ (recall that t and d are two functions of n).

Suppose for contradiction that there is an $n^{2-\varepsilon'}$-time multiplicative $t(n)$-approximation algorithm $\mathbb{A}$ for $\mathsf{Max\text{-}IP}(n,d)$ for some $\varepsilon' > 0$.

Let $\varepsilon = \varepsilon'/2$. Now, for every constant c_2, we apply the reduction in Lemma 3.3 with $\tau = t$ to reduce an $\mathsf{OV}_{n,c_2\log n}$ instance to n^{ε} $\mathsf{Max\text{-}IP}_{n,t^{\mathrm{poly}(c_2/\varepsilon)}\cdot\log n}$ instances. Since $t^{\mathrm{poly}(c_2/\varepsilon)} = t^{O(1)}$ and $t = c^{o(1)}$, it follows that for sufficiently large n, $t^{O(1)} \cdot \log n = c^{o(1)} \cdot \log n = o(d)$. It in turn implies that for sufficiently large n, n^{ε} calls to $\mathbb{A}$ are enough to solve the $\mathsf{OV}_{n,c_2\log n}$ instance.

Therefore, we can solve $\mathsf{OV}_{n,c_2\log n}$ in $n^{2-\varepsilon'} \cdot n^{\varepsilon} = n^{2-\varepsilon}$ time for all constants c_2. Contradiction to OVC. □

Finally, the correctness of Theorem 1.5 follows directly from Lemma 3.2 and Corollary 3.4.

3.2 The additive case

In this subsection we prove Theorem 1.10. We proceed similarly as in the multiplicative case by establishing the algorithm first.

The algorithm is actually very easy, we simply apply the following algorithm from [13].

Lemma 3.5 (Implicit in Theorem 5.1 in [13]). Assuming $\varepsilon = \omega(\log^6 \log(d \log n)/\log^3 n)$, there is an

$$n^{2-\Omega\left(\varepsilon^{1/3}/\log\left(\frac{d}{\varepsilon \log n}\right)\right)}\text{-time}$$

additive $\varepsilon \cdot d$-approximation randomized algorithm for $\mathsf{Max\text{-}IP}_{n,d}$.

Lemma 3.6. Let $\varepsilon = \min(t,d)/d$. There is an

$$O\left(n^{2-\Omega(\varepsilon^{1/3}/\log\varepsilon^{-1})}\right)$$

time, additive t-approximation randomized algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ when $\varepsilon = \omega(\log^6 \log n/\log^3 n)$.

Proof. When $t > d$ the problem becomes trivial, so we can assume $t \le d$, and now $t = \varepsilon \cdot d$.

Let $\varepsilon_1 = \varepsilon/2$ and c_1 be a constant to be specified later. Given a $\mathsf{Max\text{-}IP}_{n,d}$ instance with two sets A and B each consisting of n vectors from $\{0,1\}^d$, we create another $\mathsf{Max\text{-}IP}_{n,d_1}$ instance with sets $\widetilde{A}$, $\widetilde{B}$ and $d_1 = c_1 \cdot \varepsilon_1^{-2} \cdot \log n$ as follows:

- Pick d_1 uniform random indices $i_1, i_2, i_3, \dots, i_{d_1} \in [d]$, each i_k for $k \in [d_1]$ is an independent uniform random number in $[d]$.
- Then we construct $\widetilde{A}$ from A by reducing each $a \in A$ into $\tilde{a} = (a_{i_1}, a_{i_2}, \dots, a_{i_{d_1}}) \in \{0,1\}^{d_1}$ and $\widetilde{B}$ from B in the same way.

Note for each $a \in A$ and $b \in B$, by a Chernoff bound, we have

$$\Pr\left[\left|\frac{\tilde{a}\cdot\tilde{b}}{d_1} - \frac{a \cdot b}{d}\right| \ge \varepsilon_1\right] < 2\mathrm{e}^{-2d_1\varepsilon_1^2} = 2n^{-2\cdot c_1}.$$

By setting $c_1 = 2$, the above probability is smaller than $1/n^3$.

Hence, by a simple union bound, with probability at least $1 - 1/n$, we have

$$\left|\frac{\widetilde{a}\cdot\widetilde{b}}{d_1} - \frac{a\cdot b}{d}\right| \le \varepsilon_1$$

for every $a \in A$ and $b \in B$. Hence, it means that this reduction only changes the "relative inner product"($a \cdot b/d$ or $\widetilde{a}\cdot\widetilde{b}/d_1$) of each pair by at most ε_1. Hence, the maximum of the "relative inner product" also changes by at most ε_1, and we have $|\mathsf{OPT}(A,B)/d - \mathsf{OPT}(\widetilde{A},\widetilde{B})/d_1| \le \varepsilon_1$.

Then we apply the algorithm in Lemma 3.5 on the instance with sets $\widetilde{A}$ and $\widetilde{B}$ with error $\varepsilon = \varepsilon_1$ to obtain an estimate $\widetilde{O}$, and our final answer is simply $(\widetilde{O}/d_1) \cdot d$.

From the guarantee from Lemma 3.5, we have $|\mathsf{OPT}(\widetilde{A},\widetilde{B})/d_1 - \widetilde{O}/d_1| \le \varepsilon_1$, and therefore

$$|\mathsf{OPT}(A,B)/d - \widetilde{O}/d_1| \le 2\varepsilon_1 = \varepsilon,$$

from which the correctness of our algorithm follows directly.

For the running time, note that the reduction part runs in linear time $O(n \cdot d)$, and the rest takes

$$n^{2-\Omega\left(\varepsilon^{1/3}/\log(\frac{d_1}{\varepsilon_1 \log n})\right)} = n^{2-\Omega(\varepsilon^{1/3}/\log\varepsilon^{-1})}$$

time. □

The lower bound is already established in [67], we show it follows from Lemma 3.3 here for completeness.

Lemma 3.7 (Theorem 4.1 of [67])**.** Assuming SETH (or OVC), and letting $d = \omega(\log n)$ and $t > 0$, there is no $n^{2-\Omega(1)}$-time additive t-approximation randomized algorithm for $\mathsf{Max\text{-}IP}_{n,d}$ if

$$t = o(d).$$

Proof. Recall that t and d are all functions of n. Suppose for contradiction that there is an $n^{2-\varepsilon'}$-time additive $t(n)$-approximation algorithm $\mathbb{A}$ for $\mathsf{Max\text{-}IP}(n,d)$ for some $\varepsilon' > 0$.

Let $\varepsilon = \varepsilon'/2$. Now, for every constant c_2, we apply the reduction in Lemma 3.3 with $\tau = 2$ to reduce an $\mathsf{OV}_{n,c_2 \log n}$ instance to n^{ε} $\mathsf{Max\text{-}IP}_{n,d_1}$ instances, where $d_1 = 2^{\mathrm{poly}(c_2/\varepsilon)} \cdot \log n = O(1) \cdot \log n$. In addition, from Lemma 3.3, to solve the $\mathsf{OV}_{n,c_2 \log n}$ instance, we only need to compute additive $T/6 = \Omega(\log n) = \Omega(d_1)$-approximations to these $\mathsf{Max\text{-}IP}$ instances obtained via the reduction.

This can be solved, via n^{ε} calls to $\mathbb{A}$ as follows: for each $\mathsf{Max\text{-}IP}_{n,d_1}$ instance $\mathcal{I}$ we get, we duplicate each coordinate d/d_1 times (note that $d = \omega(\log n) \gg d_1 = O(\log n)$, and for simplicity we assume $d_1 \mid d$), to obtain a $\mathsf{Max\text{-}IP}_{n,d}$ instance $\mathcal{I}^{\mathsf{new}}$, such that $\mathsf{OPT}(\mathcal{I}^{\mathsf{new}}) = d/d_1 \cdot \mathsf{OPT}(\mathcal{I})$. Then $\mathbb{A}$ can be used to estimate $\mathsf{OPT}(\mathcal{I}^{\mathsf{new}})$ within an additive error $t = o(d)$. Scaling its estimate by d_1/d, it can also be used to estimate $\mathsf{OPT}(\mathcal{I})$ within an additive error $o(d_1) = o(\log n) \le T/6$ for sufficiently large n.

Therefore, we can solve $\mathsf{OV}_{n,c_2 \log n}$ in $n^{2-\varepsilon'} \cdot n^{\varepsilon} = n^{2-\varepsilon}$ time for all constants c_2. Contradiction to OVC. □

Finally, the correctness of Theorem 1.10 follows directly from Lemma 3.6 and Lemma 3.7.

3.3 Adaptation to All-Pair-Max-IP

Now we sketch the adaptation of our algorithms to work for the All-Pair-Max-IP problem.

Reminder of Corollary 1.12 *Suppose $\omega(\log n) < d < n^{o(1)}$, and let*

$$\varepsilon_M := \min\left(\frac{\log t}{\log(d/\log n)}, 1\right) \text{ and } \varepsilon_A := \frac{\min(t,d)}{d}.$$

There is an $n^{2-\Omega(\varepsilon_M)}\,\mathrm{polylog}(n)$-time multiplicative t-approximation algorithm and an $n^{2-\Omega(\varepsilon_A^{1/3}/\log \varepsilon_A^{-1})}$-time additive t-approximation algorithm for All-Pair-Max-IP$_{n,d}$, when $\varepsilon_A = \omega(\log^6 \log n/\log^3 n)$.

Proof sketch. Note that the algorithm in Lemma 3.5 from [13] actually works for the All-Pair-Max-IP$_{n,d}$. Hence, we can simply apply that algorithm after the coordinate sampling phase, and obtain an additive t-approximation algorithm for All-Pair-Max-IP$_{n,d}$.

For multiplicative t-approximation algorithm, suppose we are given with two sets, A and B, of n vectors each, from $\{0,1\}^d$. Instead of partitioning each of them into n/b subsets (the notation used here is the same as in the proof of Lemma 3.2), we only partition B into n/b subsets, $B_1, B_2, \ldots, B_{n/b}$, of size b each, and calculate $P_r(x, B_i) := \sum_{y \in B_i} P_r(x,y)$ for every $x \in A$ and $i \in [n/b]$ using similar reduction to rectangular matrix multiplication as in Lemma 3.2. Note that here we are multiplying an $n \times m$ matrix and an $m \times (n/b)$ matrix, and this can be reduced to b instances of multiplication of an $(n/b) \times m$ matrix and an $m \times (n/b)$ matrix, and now our running time becomes $n^2/b \cdot \mathrm{polylog}(n)$ instead of $n^2/b^2 \cdot \mathrm{polylog}(n)$.

By a similar analysis, these can be done in $n^{2-\Omega(\varepsilon_M)} \cdot \mathrm{polylog}(n)$ time, and then we can compute the multiplicative t-approximate answers for the given All-Pair-Max-IP$_{n,d}$ instance. □

3.4 Improved hardness for LCS-Closest Pair problem

We finish this section with the proof of Corollary 1.9. First we abstract the reduction from Max-IP to LCS-Closest-Pair in [5] here.

Lemma 3.8 (Implicit in Theorem I.10 in [5])**.** For every real $t \geq 2$ and integer n, computing a multiplicative t-approximation to $\mathsf{Max\text{-}IP}_{n,d}$ reduces to computing a multiplicative $t/2$-approximation to $\mathsf{LCS\text{-}Closest\text{-}Pair}_{n,O(d^3 \log^2 n)}$ in $O(n \operatorname{poly}(d, \log n))$ time.

Now we are ready to prove Corollary 1.9 (restated below for convenience).

Reminder of Corollary 1.9 *Assuming SETH (or OVC), for every $t \geq 2$, computing a multiplicative t-approximation to $\mathsf{LCS\text{-}Closest\text{-}Pair}_{n,d}$ requires $n^{2-o(1)}$ time, if $d = t^{\omega(1)} \cdot \log^5 n$.*

Proof. From Corollary 3.4, assuming SETH (or OVC), for every $t \geq 2$, we have that computing a multiplicative $2t$-approximation to $\mathsf{Max\text{-}IP}_{n,d}$ requires $n^{2-o(1)}$ time if $d = t^{\omega(1)} \cdot \log n$. Then from Lemma 3.8, we immediately have that computing a multiplicative t-approximation to $\mathsf{LCS\text{-}Closest\text{-}Pair}_{n,d^3 \cdot \log^2 n} = \mathsf{LCS\text{-}Closest\text{-}Pair}_{n,t^{\omega(1)} \cdot \log^5 n}$ requires $n^{2-o(1)}$ time. □

3.5 A proof of Lemma 3.3

Finally, we present a proof of Lemma 3.3, which is implicit in [67].

We need the following efficient MA protocol for Set-Disjointness from [67], which is also used in [51].[18]

Lemma 3.9 (Theorem 3.2 of [67])**.** For every α and m, there is an $(m/\alpha, \log_2 m + O(1), \operatorname{poly}(\alpha), 1/2)$-efficient MA protocol for DISJ_m.

We want to reduce the error probability while keeping the number of total random coins relatively low. To achieves this, we can use an expander graph (Theorem 2.7) to prove the following theorem.

Lemma 3.10. For every α, m and $\varepsilon < 1/2$, there is an $(m/\alpha, \log_2 m + O(\log \varepsilon^{-1}), \operatorname{poly}(\alpha) \cdot \log \varepsilon^{-1}, \varepsilon)$-efficient MA protocol for DISJ_m.

Proof. Let c_1 and $\mathcal{F} : \{0,1\}^{\log m + c_1 \cdot \log \varepsilon^{-1}} \to [m]^{c_1 \cdot \log \varepsilon^{-1}}$ be the corresponding constant and function as in Theorem 2.7, and Π be the $(m/\alpha, \log_2 m, \operatorname{poly}(\alpha), 1/2)$-efficient MA protocol for DISJ_m in Lemma 3.9. Set $q = c_1 \cdot \log \varepsilon^{-1}$ and our new protocol Π_{new} works as follows:

- Merlin still sends the same advice to Alice as in Π.
- Alice and Bob jointly toss $r = \log m + q$ coins to get a string $w \in \{0,1\}^r$. Then we let $w_1, w_2, \ldots, w_q$ be the sequence corresponding to $\mathcal{F}(w)$. Each of them can be interpreted as $\log m$ bits.
- Bob sends Alice q messages, the i-th message m_i corresponds to Bob's message in Π when the random bits is w_i.

[18] The protocol in [51] also works for the k-party number-in-hand model.

- After that, Alice decides whether to accept or not as follows:
  - If for every $i \in [q]$, Alice would accept Bob's message m_i with random bits w_i in Π, then Alice accepts.
  - Otherwise, Alice rejects.

It is easy to verify that the advice length, message length and number of random coins satisfy our requirements.

For the error probability, note that when these two sets are disjoint, the same advice in Π leads to acceptance of Alice. Otherwise, suppose the advice from Merlin is either wrong or these two sets are intersecting, then half of the random bits in $\{0,1\}^{\log m}$ leads to the rejection of Alice in Π. Hence, from Theorem 2.7, with probability at least $1-\varepsilon$, at least one of the random bits w_i would lead to the rejection of Alice, which completes the proof. □

Finally, we prove Lemma 3.3 (restated below).

Reminder of Lemma 3.3 *There is a universal constant c_1 such that, for every integer c, real$\varepsilon \in (0,1]$ and $\tau \geq 2$, $\mathsf{OV}_{n,c\log n}$ can be reduced to n^{ε} $\mathsf{Max\text{-}IP}_{n,d}$ instances (A_i,B_i) for $i \in [n^{\varepsilon}]$, such that:*

- *$d = \tau^{\mathrm{poly}(c/\varepsilon)} \cdot \log n$.*
- *Letting $T = c\log n \cdot \tau^{c_1}$, if there is $a \in A$ and $b \in B$ such that $a \cdot b = 0$, then there exists an i such that $\mathsf{OPT}(A_i,B_i) \geq T$.1*
- *Otherwise, for every i we must have $\mathsf{OPT}(A_i,B_i) \leq T/\tau$.*

Proof. The reduction follows exactly the same as in [5], we recap here for completeness.

Set $\alpha = c/\varepsilon$, $m = c \cdot \log n$ and $\varepsilon = 1/\tau$, and let Π be the $(m/\alpha, \log_2 m + O(\log \varepsilon^{-1}), \mathrm{poly}(\alpha) \cdot \log \varepsilon^{-1}, \varepsilon)$-efficient MA protocol for Set-Disjointness as in Lemma 3.10.

Now, we first enumerate all of $2^{m/\alpha} = 2^{\varepsilon \cdot \log n} = n^{\varepsilon}$ possible advice strings, and create a $\mathsf{Max\text{-}IP}$ instance for each of the advice strings.

For a fixed advice $\psi \in \{0,1\}^{\varepsilon \cdot \log n}$, we create a $\mathsf{Max\text{-}IP}$ instance with sets A_ψ and B_ψ as follows. We use $a \circ b$ to denote the concatenation of the strings a and b.

Let $r = \log_2 m + c_1 \cdot \log \varepsilon^{-1}$, where c_1 is the constant hidden in the big O notation in Lemma 3.10, and $\ell = \mathrm{poly}(\alpha) \cdot \log \varepsilon^{-1}$. Let $m_1, m_2, \ldots, m_{2^\ell}$ be an enumeration of all strings in $\{0,1\}^{\ell}$.

- For each $a \in A$, and for each string $w \in \{0,1\}^r$, we create a vector $a^w \in \{0,1\}^{2^\ell}$, such that a_i^w indicates that given advice ψ and randomness w, whether Alice accepts message m_i or not (1 for acceptance, 0 for rejection). Let the concatenation of all these a^w be a_ψ. Then A_ψ is the set of all these a_ψ for $a \in A$.
- For each $b \in B$, and for each string $w \in \{0,1\}^r$, we create a vector $b^w \in \{0,1\}^{2^\ell}$, such that $b_i^w = 1$ if Bob sends the message m_i given advice ψ and randomness w, and $= 0$ otherwise. Let the concatenation of all these b^w be b_ψ. Then B_ψ is the set of all these b_ψ for $b \in B$.

We can see that for $a \in A$ and $b \in B$, $a_\psi \cdot b_\psi$ is precisely the number of random coins leading Alice to accept the message from Bob given advice ψ when Alice and Bob holds a and b correspondingly. Therefore, let $T = 2^r = c\log n \cdot \tau^{c_1}$. From the properties of the protocol Π, we can see that:

- If there is an $a \in A$ and $b \in B$ such that $a \cdot b = 0$, then there is a $\psi \in \{0,1\}^{\varepsilon \cdot \log n}$ such that $a_\psi \cdot b_\psi \geq T$.
- Otherwise, for every $a \in A$, $b \in B$ and advice $\psi\{0,1\}^{\varepsilon \cdot \log n}$, $a_\psi \cdot b_\psi \leq T/\tau$.

And this completes the proof. □

4 Hardness of Exact $\mathbb{Z}$-Max-IP, Hopcroft's problem and more

In this section we show hardness of Hopcroft's problem, exact $\mathbb{Z}$-Max-IP, ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair. Essentially our results follow from the framework of [75], in which it is shown that hardness of Hopcroft's problem implies hardness of other three problems, and is implied by dimensionality reduction for OV.

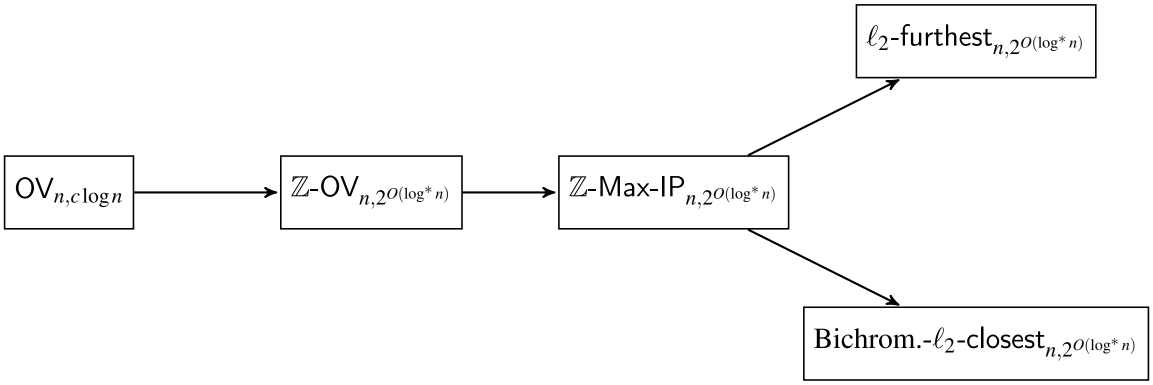
ℓ_2-furthest$_{n,2^{O(\log^* n)}}$
$\mathsf{OV}_{n,c\log n}$
$\mathbb{Z}$-$\mathsf{OV}_{n,2^{O(\log^* n)}}$
$\mathbb{Z}$-Max-IP$_{n,2^{O(\log^* n)}}$
Bichrom.-ℓ_2-closest$_{n,2^{O(\log^* n)}}$

Figure 1: A diagram for all reductions in this section.

The organization of this section

In Section 4.1, we prove the improved dimensionality reduction for OV. In Section 4.2, we establish the hardness of Hopcroft's problem in dimension $2^{O(\log^* n)}$ with the improved reduction. In Section 4.3, we show Hopcroft's problem can be reduced to $\mathbb{Z}$-Max-IP and thus establish the hardness for the later one. In Section 4.4, we show $\mathbb{Z}$-Max-IP can be reduced to ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair, therefore the hardness for the last two problems follow. In Section 4.5, we show that to construct better dimensionality reduction for OV, it suffices to show the existence of such reductions, instead of constructing them algorithmically. See Figure 1 for a diagram of all reductions covered in this section.

The reductions in Section 4.2, Section 4.3 and Section 4.4 are all from [75] (either explicitly or implicitly), we make them explicit here for our ease of exposition and for making the paper self-contained.

4.1 Improved dimensionality reduction for OV

We begin with the improved dimensionality reduction for OV. The following theorem is one of the technical cores of this paper, which makes use of the CRR encoding (see Theorem 2.5) recursively.

Theorem 4.1. Let b,ℓ be two sufficiently large integers. There is a reduction $\psi_{b,\ell} : \{0,1\}^{b\cdot\ell} \to \mathbb{Z}^{\ell}$ and a set $V_{b,\ell} \subseteq \mathbb{Z}$, such that for every $x,y \in \{0,1\}^{b\cdot\ell}$,

$$x\cdot y = 0 \Leftrightarrow \psi_{b,\ell}(x)\cdot\psi_{b,\ell}(y) \in V_{b,\ell}$$

and

$$0 \le \psi_{b,\ell}(x)_i < \ell^{6^{\log^*(b)}\cdot b}$$

for every $x \in \{0,1\}^{b\cdot\ell}$ and $i \in [\ell]$. Moreover, the computation of $\psi_{b,\ell}(x)$ takes $\operatorname{poly}(b\cdot\ell)$ time, and the set $V_{b,\ell}$ can be constructed in

$$O\left(\ell^{O(6^{\log^*(b)}\cdot b)}\cdot\operatorname{poly}(b\cdot\ell)\right)$$

time.

Remark 4.2. *We didn't make much effort to minimize the base* 6 *above to keep the calculation clean, it can be replaced by any constant* > 2 *with a tighter calculation.*

Proof. We are going to construct our reduction in a recursive way. ℓ will be the same throughout the proof, hence in the following we use ψ_b (V_b) instead of $\psi_{b,\ell}$ ($V_{b,\ell}$) for simplicity.

Direct CRR for small b: When $b < \ell$, we use a direct Chinese remainder representation of numbers. We pick b distinct primes $q_1,q_2,\ldots,q_b$ in $[\ell+1,\ell^2]$ (they are guaranteed to exist by Lemma 2.4), and use them for our CRR encoding.

For $x \in \{0,1\}^{b\cdot\ell}$, we partition it into ℓ equal-sized subvectors, and use x^i to denote the i-th subvector. That is, x^i is the subvector of x from the $((i-1)\cdot b+1)$-th bit to the $(i\cdot b)$-th bit.

Then we define $\psi_b(x)$ as

$$\psi_b(x) := \left(\mathsf{CRR}\left(\{x^1_j\}_{j=1}^{b}\right), \mathsf{CRR}\left(\{x^2_j\}_{j=1}^{b}\right), \ldots, \mathsf{CRR}\left(\left\{x^{\ell}_j\right\}_{j=1}^{b}\right)\right).$$

That is, the i-th coordinate of $\psi_b(x)$ is the CRR encoding of the i-th subvector x^i with respect to the primes q_j.

Now, for $x,y \in \{0,1\}^{b\cdot\ell}$, note that for $j \in [b]$,

$$\begin{aligned}
&\psi_b(x)\cdot\psi_b(y) \pmod{q_j}\\
\equiv&\sum_{i=1}^{\ell}\mathsf{CRR}\left(\{x^i_j\}_{j=1}^{b}\right)\cdot\mathsf{CRR}\left(\{y^i_j\}_{j=1}^{b}\right) \pmod{q_j}\\
\equiv&\sum_{i=1}^{\ell}x^i_j\cdot y^i_j \pmod{q_j}.
\end{aligned}$$

Since the sum $\sum_{i=1}^{\ell} x_j^i \cdot y_j^i$ is in $[0,\ell]$, and $q_j > \ell$, we can see

$$\sum_{i=1}^{\ell} x_j^i \cdot y_j^i = 0 \Leftrightarrow \psi_b(x) \cdot \psi_b(y) \equiv 0 \pmod{q_j}.$$

Therefore, $x \cdot y = \sum_{j=1}^{b} \sum_{i=1}^{\ell} x_j^i \cdot y_j^i = 0$ is equivalent to that

$$\psi_b(x) \cdot \psi_b(y) \equiv 0 \pmod{q_j}$$

for every $j \in [b]$.

Finally, we have $0 \le \psi_b(x)_i < \prod_{j=1}^{b} q_j < \ell^{2 \cdot b} \le \ell^{6^{\log^*(b)} \cdot b}$. Therefore

$$\psi_b(x) \cdot \psi_b(y) < \ell^{6^{\log^*(b)} \cdot 2b+1},$$

and we can set V_b to be the set of all integers in $[0, \ell^{6^{\log^*(b)} \cdot 2b+1}]$ that is 0 modulo each q_j, and it is easy to see that

$$x \cdot y \Leftrightarrow \psi_b(x) \cdot \psi_b(y) \in V_b$$

for every $x, y \in \{0,1\}^{b \cdot \ell}$.

Recursive construction for larger b: When $b \ge \ell$, suppose the theorem holds for all integers $b' < b$. Let b_{m} be the number such that (we ignore the rounding issue here and pretend that b_{m} is an integer for simplicity),

$$\ell^{6^{\log^*(b_{\mathsf{m}})} \cdot b_{\mathsf{m}}} = b.$$

Then we pick b/b_{m} primes $p_1, p_2, \dots, p_{b/b_{\mathsf{m}}}$ in $[(b^2\ell), (b^2\ell)^2]$, and use them as our reference primes in the CRR encodings.

For $x \in \{0,1\}^{b \cdot \ell}$, as before, we partition x into ℓ equal-sized subvectors $x^1, x^2, \dots, x^{\ell}$, where x^i consists of the $((i-1) \cdot b + 1)$-th bit of x to the $(i \cdot b)$-th bit of x. Then we partition each x^i again into b/b_{m} micro groups, each of size b_{m}. We use $x^{i,j}$ to denote the j-th micro group of x^i after the partition.

Now, we use $x^{[j]}$ to denote the concatenation of the vectors $x^{1,j}, x^{2,j}, \dots, x^{\ell,j}$. That is, $x^{[j]}$ is the concatenation of the j-th micro group in each of the ℓ subvectors. Note that $x^{[j]} \in \{0,1\}^{b_{\mathsf{m}} \cdot \ell}$, and can be seen as a smaller instance, on which we can apply $\psi_{b_{\mathsf{m}}}$.

Our recursive construction then goes in two steps. In the first step, we make use of $\psi_{b_{\mathsf{m}}}$, and transform each b_{m}-size micro group into a single number in $[0,b)$. This step transforms x from a vector in $\{0,1\}^{b \cdot \ell}$ into a vector $S(x)$ in $\mathbb{Z}^{(b/b_{\mathsf{m}}) \cdot \ell}$. And in the second step, we use a similar CRR encoding as in the base case to encode $S(x)$, to get our final reduced vector in $\mathbb{Z}^{\ell}$.

$S(x)$ is simply

$$\begin{aligned} S(x) := \Big(& \psi_{b_{\mathsf{m}}}(x^{[1]})_1, \psi_{b_{\mathsf{m}}}(x^{[2]})_1, \dots, \psi_{b_{\mathsf{m}}}(x^{[b/b_{\mathsf{m}}]})_1, \\ & \psi_{b_{\mathsf{m}}}(x^{[1]})_2, \psi_{b_{\mathsf{m}}}(x^{[2]})_2, \dots, \psi_{b_{\mathsf{m}}}(x^{[b/b_{\mathsf{m}}]})_2, \\ & \dots, \dots, \dots \\ & \psi_{b_{\mathsf{m}}}(x^{[1]})_{\ell}, \psi_{b_{\mathsf{m}}}(x^{[2]})_{\ell}, \dots, \psi_{b_{\mathsf{m}}}(x^{[b/b_{\mathsf{m}}]})_{\ell} \Big). \end{aligned}$$

That is, we apply $\psi_{b_{\mathsf{m}}}$ to all the $x^{[j]}$, and shrink all the corresponding micro groups in x into integers. Again, we partition $S = S(x)$ into ℓ equal size groups $S^1, S^2, \ldots, S^\ell$.

Then we define $\psi_b(x)$ as

$$\psi_b(x) := \left(\mathsf{CRR}\left(\left\{S^1_j\right\}_{j=1}^{b/b_{\mathsf{m}}}\right), \mathsf{CRR}\left(\left\{S^2_j\right\}_{j=1}^{b/b_{\mathsf{m}}}\right), \ldots, \mathsf{CRR}\left(\left\{S^\ell_j\right\}_{j=1}^{b/b_{\mathsf{m}}}\right)\right).$$

In other words, the i-th coordinate of $\psi_b(x)$ is the CRR representation of the number sequence S^i, with respect to our primes $\{q_j\}_{j=1}^{b/b_{\mathsf{m}}}$.

Now, note that for $x, y \in \{0,1\}^{b\cdot\ell}$, $x \cdot y = 0$ is equivalent to $x^{[j]} \cdot y^{[j]} = 0$ for every $j \in [b/b_{\mathsf{m}}]$, which is further equivalent to

$$\psi_{b_{\mathsf{m}}}(x^{[j]}) \cdot \psi_{b_{\mathsf{m}}}(y^{[j]}) \in V_{b_{\mathsf{m}}}$$

for every $j \in [b/b_{\mathsf{m}}]$, by our assumption on $\psi_{b_{\mathsf{m}}}$.

Since $0 \le \psi_{b_{\mathsf{m}}}(x^{[j]})_i, \psi_{b_{\mathsf{m}}}(y^{[j]})_i < b$ for every $x, y \in \{0,1\}^{b\cdot\ell}$, $i \in [\ell]$ and $j \in [b/b_{\mathsf{m}}]$, we also have $\psi_{b_{\mathsf{m}}}(x^{[j]}) \cdot \psi_{b_{\mathsf{m}}}(y^{[j]}) < b^2 \cdot \ell$, therefore we can assume that $V_{b_{\mathsf{m}}} \subseteq [0, b^2\ell)$.

For every $x, y \in \{0,1\}^{b\cdot\ell}$ and $j \in [b/b_{\mathsf{m}}]$, we have

$$\begin{aligned}
&\psi_b(x) \cdot \psi_b(y)\\
\equiv &\sum_{i=1}^{\ell} \mathsf{CRR}\left(\left\{S(x)^i_j\right\}_{j=1}^{b/b_{\mathsf{m}}}\right) \cdot \mathsf{CRR}\left(\left\{S(y)^i_j\right\}_{j=1}^{b/b_{\mathsf{m}}}\right) \pmod{p_j}\\
\equiv &\sum_{i=1}^{\ell} S(x)^i_j \cdot S(y)^i_j \pmod{p_j}\\
\equiv &\sum_{i=1}^{\ell} \psi_{b_{\mathsf{m}}}(x^{[j]})_i \cdot \psi_{b_{\mathsf{m}}}(y^{[j]})_i \pmod{p_j}\\
\equiv &\psi_{b_{\mathsf{m}}}(x^{[j]}) \cdot \psi_{b_{\mathsf{m}}}(y^{[j]}) \pmod{p_j}.
\end{aligned}$$

Since $p_j \ge b^2 \cdot \ell$, we can determine $\psi_{b_{\mathsf{m}}}(x^{[j]}) \cdot \psi_{b_{\mathsf{m}}}(y^{[j]})$ from $\psi_b(x) \cdot \psi_b(y)$ by taking modulo p_j. Therefore $x \cdot y = 0$ is equivalent to

$$(\psi_b(x) \cdot \psi_b(y) \mod p_j) \in V_{b_{\mathsf{m}}} \text{ for every } j \in [b/b_{\mathsf{m}}].$$

Finally, recall that we have

$$\ell^{6^{\log^*(b_{\mathsf{m}})} \cdot b_{\mathsf{m}}} = b.$$

Taking logarithm of both sides, we have

$$6^{\log^*(b_{\mathsf{m}})} \cdot b_{\mathsf{m}} \cdot \log \ell = \log b.$$

Then we can upper bound $\psi_b(x)_i$ by

$$\begin{aligned}
\psi_b(x)_i &< \prod_{j=1}^{b/b_{\mathsf{m}}} p_j \\
&< (b^2\ell)^{2\cdot(b/b_{\mathsf{m}})} && (b \geq \ell.) \\
&\leq 2^{6\cdot b/b_{\mathsf{m}}\cdot\log b} \\
&\leq 2^{6\cdot b/b_{\mathsf{m}}\cdot 6^{\log^*(b_{\mathsf{m}})}\cdot b_{\mathsf{m}}\cdot\log\ell} \\
&\leq \ell^{6\cdot 6^{\log^*(b_{\mathsf{m}})}\cdot b} \\
&\leq \ell^{6^{\log^*(b)}\cdot b} && (b_{\mathsf{m}} \leq \log b, \log^*(b_{\mathsf{m}})+1 \leq \log^*(\log b)+1 = \log^*(b).)
\end{aligned}$$

Therefore, we can set V_b as the set of all integers t in $[0, \ell^{6^{\log^*(b)}\cdot 2b+1})$ such that

$$(t \mod p_j) \in V_{b_{\mathsf{m}}} \text{ for every } j \in [b/b_{\mathsf{m}}].$$

And it is easy to see this V_b satisfies our requirement.

Finally, it is easy to see that the straightforward way of constructing $\psi_b(x)$ takes $O(\mathrm{poly}(b\cdot\ell))$ time, and we can construct V_b by enumerating all possible values of $\psi_b(x)\cdot\psi_b(y)$ and checking each of them in $O(\mathrm{poly}(b\cdot\ell))$ time. Since there are at most $\ell^{O(6^{\log^*(b)}\cdot b)}$ such values, V_b can be constructed in

$$O\left(\ell^{O(6^{\log^*(b)}\cdot b)}\cdot\mathrm{poly}(b\cdot\ell)\right)$$

time, which completes the proof. □

Now we prove Lemma 1.17, we recap its statement here for convenience.

Reminder of Lemma 1.17 *Let $1 \leq \ell \leq d$. There is an*

$$O\left(n\cdot\ell^{O(6^{\log^* d}\cdot(d/\ell))}\cdot\mathrm{poly}(d)\right)\text{-time}$$

reduction from $\mathsf{OV}_{n,d}$ to $\ell^{O(6^{\log^* d}\cdot(d/\ell))}$ instances of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$, with vectors of entries with bit-length $O(d/\ell\cdot\log\ell\cdot 6^{\log^* d})$.

Proof. The proof is exactly the same as the proof for Lemma 1.1 in [75] with different parameters. We recap it here for convenience.

Given two sets A' and B' each consisting of n vectors from $\{0,1\}^d$, we apply $\psi_{d/\ell,\ell}$ to each of the vectors in A' (B') to obtain a set A (B) of vectors from $\mathbb{Z}^\ell$. From Theorem 4.1, there is a $(u,v) \in A' \times B'$ such that $u\cdot v = 0$ if and only if there is a $(u,v) \in A \times B$ such that $u\cdot v \in V_{d/\ell,\ell}$.

Now, for each element $t \in V_{d/\ell,\ell}$, we are going to construct two sets A_t and B_t of vectors from $\mathbb{Z}^{\ell+1}$ such that there is a $(u,v) \in A \times B$ with $u\cdot v = t$ if and only if there is a $(u,v) \in A_t \times B_t$ with $u\cdot v = 0$. We construct a set A_t as the collection of all vectors $u_A = [u,1]$ for $u \in A$, and a set B_t as the collection of all vectors $v_B = [v,-t]$ for $v \in B$. It is easy to verify this reduction has the properties we want.

Note that there are at most $\ell^{O(6^{\log^* d}\cdot(d/\ell))}$ numbers in $V_{d/\ell,\ell}$, so we have such a number of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$ instances. And from Theorem 4.1, the reduction takes

$$O\left(n\cdot\ell^{O(6^{\log^* d}\cdot(d/\ell))}\cdot\mathrm{poly}(d)\right)$$

time.

Finally, the bit-length of reduced vectors is bounded by

$$\log\left(\ell^{O(6^{\log^* d}\cdot(d/\ell))}\right)=O\left(d/\ell\cdot\log\ell\cdot 6^{\log^* d}\right),$$

which completes the proof. □

A transformation from nonuniform construction to uniform construction. The proof for Theorem 4.1 works recursively. In one recursive step, we reduce the construction of $\psi_{b,\ell}$ to the construction of $\psi_{b_{\mathsf{m}},\ell}$, where $b_{\mathsf{m}}\le\log b$. Applying this reduction $\log^* n$ times, we get a sufficiently small instance that we can switch to a direct CRR construction.

An interesting observation here is that after applying the reduction only three times, the block length parameter becomes $b'\le\log\log\log b$. Such a b' is so small that we can actually use brute force to find the "optimal" construction $\psi_{b',\ell}$ in $b^{o(1)}$ time instead of recursing deeper. Hence, to find a construction better than Theorem 4.1, we only need to prove the existence of such a construction. See Section 4.5 for details.

4.2 Improved hardness for Hopcroft's problem

In this subsection we are going to prove Theorem 1.18 using our new dimensionality reduction from Lemma 1.17, we recap its statement here for completeness.

Reminder of Theorem 1.18 *[Hardness of Hopcroft's problem in $c^{\log^* n}$ dimension] Assuming SETH (or OVC), there is a constant c such that $\mathbb{Z}$-$\mathsf{OV}_{n,c^{\log^* n}}$ with vectors of $O(\log n)$-bit entries requires $n^{2-o(1)}$ time.*

Proof. The proof here follows roughly the same as the proof for Theorem 1.1 in [75].

Let c be an arbitrary constant and $d:=c\cdot\log n$. We show that an algorithm $\mathbb{A}$ solving $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$ where $\ell=7^{\log^* n}$ in $O(n^{2-\delta})$ time for some $\delta>0$ can be used to construct an $O(n^{2-\delta+o(1)})$-time algorithm for $\mathsf{OV}_{n,d}$, and therefore contradicts the OVC.

We simply invoke Lemma 1.17, note that we have

$$\begin{aligned}\log\left\{\ell^{O\left(6^{\log^* d}\cdot(d/\ell)\right)}\right\}&=\log\ell\cdot O\left(6^{\log^* d}\cdot(d/\ell)\right)\\&=O\left(\log^* n\cdot 6^{\log^* n}\cdot c\cdot\log n/7^{\log^* n}\right)\\&=O\left(\log^* n\cdot(6/7)^{\log^* n}\cdot c\cdot\log n\right)\\&=o(\log n).\end{aligned}$$

Therefore, the reduction takes $O(n \cdot \ell^{O\left(6^{\log^* d} \cdot (d/\ell)\right)} \cdot \mathrm{poly}(d)) = n^{1+o(1)}$ time. It reduces an $\mathsf{OV}_{n,d}$ instance to $n^{o(1)}$ instances of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$, whose vectors have bit length $o(\log n)$ as calculated above. We simply solve all these $n^{o(1)}$ instances using $\mathbb{A}$, and this gives us an $O(n^{2-\delta+o(1)})$-time algorithm for $\mathsf{OV}_{n,d}$, which completes the proof. □

4.3 Hardness for $\mathbb{Z}$-Max-IP

Now we move to hardness of exact $\mathbb{Z}$-$\mathsf{Max\text{-}IP}$.

Theorem 4.3 (Implicit in Theorem 1.2 [75])**.** There is an $O(\mathrm{poly}(d) \cdot n)$-time algorithm which reduces a $\mathbb{Z}$-$\mathsf{OV}_{n,d}$ instance into a $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,d^2}$ instance.

Proof. We remark here that this reduction is implicitly used in the proof of Theorem 1.2 in [75], we abstract it here only for our exposition.

Given a $\mathbb{Z}$-$\mathsf{OV}_{n,d}$ instance with sets A, B. Consider the following polynomial $P(x,y)$, where $x, y \in \mathbb{Z}^d$.

$$P(x,y) = -(x \cdot y)^2 = \sum_{i,j \in [d]} -(x_i \cdot y_i) \cdot (x_j \cdot y_j) = \sum_{i,j \in [d]} -(x_i \cdot x_j) \cdot (y_i \cdot y_j).$$

It is easy to see that whether there is an $(x,y) \in A \times B$ such that $x \cdot y = 0$ is equivalent to whether the maximum value of $P(x,y)$ is 0.

Now, for each $x \in A$ and $y \in B$, we identify $[d^2]$ with $[d] \times [d]$ and construct $\widetilde{x}, \widetilde{y} \in \mathbb{Z}^{d^2}$ such that

$$\widetilde{x}_{(i,j)} = x_i \cdot x_j \quad \text{and} \quad \widetilde{y}_{(i,j)} = -y_i \cdot y_j.$$

Then we have $\widetilde{x} \cdot \widetilde{y} = P(x,y)$. Hence, let $\widetilde{A}$ be the set of all these vectors $\widetilde{x}$, and $\widetilde{B}$ be the set of all these vectors $\widetilde{y}$. Whether there is an $(x,y) \in A \times B$ such that $x \cdot y = 0$ is equivalent to whether $\mathsf{OPT}(\widetilde{A}, \widetilde{B}) = 0$, and our reduction is completed. □

Now, Theorem 1.14 (restated below) is just a simple corollary of Theorem 4.3 and Theorem 1.18.

Reminder of Theorem 1.14 *Assuming SETH (or OVC), there is a constant c such that every exact algorithm for $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,d}$ for $d = c^{\log^* n}$ requires $n^{2-o(1)}$ time, with vectors of $O(\log n)$-bit entries.*

4.3.1 A dimensionality reduction for Max-IP

The reduction $\psi_{b,\ell}$ from Theorem 4.1 actually does more: For $x, y \in \{0,1\}^{b \cdot \ell}$, from $\psi_{b,\ell}(x) \cdot \psi_{b,\ell}(y)$ we can in fact determine the inner product $x \cdot y$ itself, not only whether $x \cdot y = 0$. Formally, we have the following corollary.

Corollary 4.4. Let b, ℓ be two sufficiently large integers. There is a reduction $\psi_{b,\ell} : \{0,1\}^{b \cdot \ell} \to \mathbb{Z}^{\ell}$ and $b \cdot \ell + 1$ sets $V_{b,\ell}^0, V_{b,\ell}^1, \ldots, V_{b,\ell}^{b \cdot \ell} \subseteq \mathbb{Z}$, such that for every $x, y \in \{0,1\}^{b \cdot \ell}$,

$$x \cdot y = k \Leftrightarrow \psi_{b,\ell}(x) \cdot \psi_{b,\ell}(y) \in V_{b,\ell}^k \quad \text{for every } 0 \le k \le b \cdot \ell,$$

and

$$0 \leq \psi_{b,\ell}(x)_i < \ell^{6^{\log^*(b)} \cdot b}$$

for for every $x \in \{0,1\}^{b \cdot \ell}$ and $i \in [\ell]$. Moreover, the computation of $\psi_{b,\ell}(x)$ takes $\mathrm{poly}(b \cdot \ell)$ time, and the sets $V_{b,\ell}^k$ can be constructed in $O(\ell^{O(6^{\log^*(b)} \cdot b)} \cdot \mathrm{poly}(b \cdot \ell))$ time.

Combining Corollary 4.4 with Theorem 4.3, we can in fact derive a similar dimensionality self-reduction from Max-IP to $\mathbb{Z}$-Max-IP.

Corollary 4.5. Let $1 \leq \ell \leq d$. There is an

$$O\left(n \cdot \ell^{O(6^{\log^* d} \cdot (d/\ell))} \cdot \mathrm{poly}(d)\right)\text{-time}$$

reduction from $\mathsf{Max\text{-}IP}_{n,d}$ to $d \cdot \ell^{O(6^{\log^* d} \cdot (d/\ell))}$ instances of $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,(\ell+1)^2}$, with vectors of entries with bit-length $O\left(d/\ell \cdot \log \ell \cdot 6^{\log^* d}\right)$.

Proof sketch. Let $b = d/\ell$ (assume ℓ divides d here for simplicity), A and B be the sets in the given $\mathsf{Max\text{-}IP}_{n,d}$ instance. We proceed similarly as the case for OV.

For each k from 0 to d, we construct the set $V_{b,\ell}^k$ as specified in Corollary 4.4. Then there is an $(x,y) \in A \times B$ such that $x \cdot y = k$ if and only if there is an $(x,y) \in A \times B$ such that $\psi_{b,\ell}(x) \cdot \psi_{b,\ell}(y) \in V_{b,\ell}^k$. Using exactly the same reduction as in Lemma 1.17, we can in turn reduce this into $\ell^{O(6^{\log^*(b)} \cdot b)}$ instances of $\mathbb{Z}$-$\mathsf{OV}_{n,\ell+1}$.

Applying Theorem 4.3, by solving all these $(d+1) \cdot \ell^{O(6^{\log^*(b)} \cdot b)}$ $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,(\ell+1)^2}$ instances, we can determine whether there is an $(x,y) \in A \times B$ such that $x \cdot y = k$ for every k, from which we can compute the answer to the $\mathsf{Max\text{-}IP}_{n,d}$ instance. □

4.4 Hardness for ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair

Now we turn to the proof of hardness of ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair. The two reductions below are slight adaptations of the reductions in the proofs of Theorem 1.2 and Corollary 2.1 in [75].

Lemma 4.6. Assuming $d = n^{o(1)}$, there is an $O(\mathrm{poly}(d) \cdot n)$-time algorithm which reduces a $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,d}$ instance into an instance of ℓ_2-Furthest Pair on $2n$ points in $\mathbb{R}^{d+2}$. Moreover, if the $\mathbb{Z}$-Max-IP instance consists of vectors of $O(\log n)$-bit entries, so does the ℓ_2-Furthest Pair instance.

Proof. Let A, B be the sets in the $\mathbb{Z}$-$\mathsf{Max\text{-}IP}_{n,d}$ instance, and k be the smallest integer such that all vectors from A and B consist of $(k \cdot \log n)$-bit entries.

Let W be $n^{C \cdot k}$ where C is a large enough constant. Given $x \in A$ and $y \in B$, we construct point

$$\widetilde{x} = \left(x, \sqrt{W - \|x\|^2}, 0\right) \quad \text{and} \quad \widetilde{y} = \left(-y, 0, \sqrt{W - \|y\|^2}\right).$$

That is, we append two corresponding values into the end of vectors x and $-y$.

Now, we can see that for $x_1, x_2 \in A$, the squared distance between their reduced points is

$$\|\widetilde{x_1} - \widetilde{x_2}\|^2 \le \|x_1 - x_2\|^2 + W \le 4 \cdot d \cdot n^{2k} + W.$$

Similarly we have

$$\|\widetilde{y_1} - \widetilde{y_2}\|^2 \le 4 \cdot d \cdot n^{2k} + W$$

for $y_1, y_2 \in B$.

Next, for $x \in A$ and $y \in B$, we have

$$\|\widetilde{x} - \widetilde{y}\|^2 = \|\widetilde{x}\|^2 + \|\widetilde{y}\|^2 - 2 \cdot \widetilde{x} \cdot \widetilde{y} = 2 \cdot W + 2 \cdot (x \cdot y) \ge 2 \cdot W - d \cdot n^{2k} > 4 \cdot d \cdot n^{2k} + W,$$

the last inequality holds when we set C to be 5.

Putting everything together, we can see the ℓ_2-farthest pair among all points $\widetilde{x}$ and $\widetilde{y}$ must be a pair $(\widetilde{x}, \widetilde{y})$ with $x \in A$ and $y \in B$. And maximizing $\|\widetilde{x} - \widetilde{y}\|$ is equivalent to maximizing $x \cdot y$, which proves the correctness of our reduction. Furthermore, when k is a constant, the reduced instance only needs vectors with $O(k) \cdot \log n = O(\log n)$-bit entries. □

Lemma 4.7. Assuming $d = n^{o(1)}$, there is an $O(\mathrm{poly}(d) \cdot n)$-time algorithm which reduces a $\mathbb{Z}$-Max-$\mathsf{IP}_{n,d}$ instance into an instance of Bichromatic ℓ_2-Closest Pair on $2n$ points in $\mathbb{R}^{d+2}$. Moreover, if the $\mathbb{Z}$-Max-IP instance consists of vectors of $O(\log n)$-bit entries, so does the Bichromatic ℓ_2-Closest Pair instance.

Proof. Let A, B be the sets in the $\mathbb{Z}$-Max-$\mathsf{IP}_{n,d}$ instance, and k be the smallest integer such that all vectors from A and B consist of $(k \cdot \log n)$-bit entries.

Let W be $n^{C \cdot k}$ where C is a large enough constant. Given $x \in A$ and $y \in B$, we construct point

$$\widetilde{x} = \left(x, \sqrt{W - \|x\|^2}, 0\right) \quad \text{and} \quad \widetilde{y} = \left(y, 0, \sqrt{W - \|y\|^2}\right).$$

That is, we append two corresponding values into the end of vectors x and y. And our reduced instance is to find the closest point between the set $\widetilde{A}$ (consisting of all these $\widetilde{x}$ where $x \in A$) and the set $\widetilde{B}$ (consisting of all these $\widetilde{y}$ where $y \in B$).

Next, for $x \in A$ and $y \in B$, we have

$$\|\widetilde{x} - \widetilde{y}\|^2 = \|\widetilde{x}\|^2 + \|\widetilde{y}\|^2 - 2 \cdot \widetilde{x} \cdot \widetilde{y} = 2 \cdot W - 2 \cdot (x \cdot y).$$

Hence minimizing $\|\widetilde{x} - \widetilde{y}\|$ where $x \in A$ and $y \in B$ is equivalent to maximize $x \cdot y$, which proves the correctness of our reduction. Furthermore, when k is a constant, the reduced instance only needs vectors with $O(k) \cdot \log n = O(\log n)$-bit entries. □

Now Theorem 1.15 and Theorem 1.16 are simple corollaries of Lemma 4.6, Lemma 4.7 and Theorem 1.14.

4.5 Nonuniform to uniform transformation for dimensionality reduction for OV

Finally, we discuss the transformation from nonuniform construction to uniform one for dimensionality reduction for OV. In order to state our result formally, we need to introduce some definitions.

Definition 4.8 (Nonuniform Reduction)**.** Let $b,\ell,\kappa \in \mathbb{N}$. We say a function $\varphi : \{0,1\}^{b\cdot\ell} \to \mathbb{Z}^{\ell}$ together with a set $V \subseteq \mathbb{Z}$ is a (b,ℓ,κ)-reduction, if the following holds:

- For every $x,y \in \{0,1\}^{b\cdot\ell}$,
$$x \cdot y = 0 \Leftrightarrow \varphi(x)\cdot\varphi(y) \in V.$$
- For every $x \in \{0,1\}^{b\cdot\ell}$ and $i \in [\ell]$,
$$0 \le \varphi(x)_i < \ell^{\kappa\cdot b}.$$

Similarly, let τ be an increasing function. We say a function family $\{\varphi_{b,\ell}\}_{b,\ell}$ together with a set family $\{V_{b,\ell}\}_{b,\ell}$ is a τ-reduction family, if for every integer b and ℓ, $(\varphi_{b,\ell}, V_{b,\ell})$ is a $(b,\ell,\tau(b))$-reduction.

Moreover, if for all integers b and $\ell \le \log\log\log b$, there is an algorithm $\mathbb{A}$ which computes $\varphi_{b,\ell}(x)$ in $\operatorname{poly}(b)$ time given b,ℓ and $x \in \{0,1\}^{b\cdot\ell}$, and constructs the set $V_{b,\ell}$ in $O\left(\ell^{O(\tau(b)\cdot b)} \cdot \operatorname{poly}(b)\right)$ time given b and ℓ, then we call $(\varphi_{b,\ell}, V_{b,\ell})$ a uniform-τ-reduction family.

Remark 4.9. *The reason we assume ℓ to be so small compared to b is that in our applications we only care about very small ℓ, and that greatly simplifies the notation. From Theorem 4.1, there is a uniform-$\left(6^{\log^* b}\right)$-reduction family, and a better uniform-reduction family implies better hardness for $\mathbb{Z}$-OV and other related problems as well (Lemma 1.17, Theorem 4.3, Lemma 4.7 and Lemma 4.6).*

Now we are ready to state our nonuniform to uniform transformation result formally.

Theorem 4.10. Letting τ be an increasing function such that $\tau(n) = O(\log\log\log n)$ and supposing there is a τ-reduction family, then there is a uniform-$O(\tau)$-reduction family.

Proof sketch. The construction in Theorem 4.1 is recursive, it constructs the reduction $\psi_{b,\ell}$ from a much smaller reduction $\psi_{b_{\mathsf{m}},\ell}$, where $b_{\mathsf{m}} \le \log b$. In the original construction, it takes $\log^* b$ recursions to make the problem instance sufficiently small so that a direct construction can be used. Here we only apply the reduction three times. First let us abstract the following lemma from the proof of Theorem 4.1.

Lemma 4.11 (Implicit in Theorem 4.1)**.** Letting $b,\ell,b_{\mathsf{m}},\kappa \in \mathbb{N}$ and supposing $\ell^{\kappa\cdot b_{\mathsf{m}}} = b$ and there is a $(b_{\mathsf{m}},\ell,\kappa)$-reduction (φ, V'), the following holds:

- There is a $(b,\ell,6\cdot\kappa)$-reduction (ψ, V).
- Given (φ, V'), for every $x \in \{0,1\}^{b\cdot\ell}$, $\psi(x)$ can be computed in $\operatorname{poly}(b\cdot\ell)$, and V can be constructed in $O\left(\ell^{O(\kappa\cdot b)} \cdot \operatorname{poly}(b\cdot\ell)\right)$ time.

Now, let $b,\ell \in \mathbb{N}$. We are going to construct our reduction as follows.

Let b_1 be the number such that
$$\ell^{\tau(b)\cdot 6^2\cdot b_1} = b.$$

Similarly, we set b_2 and b_3 so that

$$\ell^{\tau(b)\cdot 6\cdot b_2} = b_1 \quad \text{and} \quad \ell^{\tau(b)\cdot b_3} = b_2.$$

We can calculate from above that $b_3 \le \log\log\log b$.

From the assumption that there is a τ-reduction, there is a $(b_3, \ell, \tau(b_3))$-reduction $(\varphi_{b_3,\ell}, V_{b_3,\ell})$, which is also a $(b_3, \ell, \tau(b))$-reduction, as τ is increasing. Note that we can assume $\ell \le \log\log\log b$ and $\tau(b) \le \log\log\log b$ from assumption. Now we simply use a brute force algorithm to find $(\varphi_{b_3,\ell}, V_{b_3,\ell})$. There are

$$\ell^{\tau(b)\cdot b_3\cdot \ell\cdot 2^{b_3\cdot\ell}} = b^{o(1)}$$

possible functions from $\{0,1\}^{b_3\cdot\ell} \to \{0,\dots \ell^{\tau(b_3)\cdot b_3} - 1\}^{\ell}$. Given such a function φ, one can check in $\mathrm{poly}(2^{b_3\cdot\ell}) = b^{o(1)}$ time that whether one can construct a corresponding set V to obtain our $(b_3, \ell, \tau(b))$-reduction.

Applying Lemma 4.11 three times, one obtain a $(b, \ell, O(\tau(b)))$-reduction (ψ, V). And since $\varphi_{b_3,\ell}$ can be found in $b^{o(1)}$ time, together with Lemma 4.11, we obtain a uniform-τ-reduction family. □

Finally, we give a direct corollary of Theorem 4.10 that the existence of an $O(1)$-reduction family implies hardness of $\mathbb{Z}$-OV, $\mathbb{Z}$-Max-IP, ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair in dimension $\omega(1)$.

Corollary 4.12. If there is an $O(1)$-reduction family, then for every $\varepsilon > 0$, there exists a $c \ge 1$ such that $\mathbb{Z}$-OV, $\mathbb{Z}$-Max-IP, ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair in dimension c with $O(\log n)$-bit entries require $n^{2-\varepsilon}$ time.

Proof sketch. Note that since the hardness of $\mathbb{Z}$-OV implies the harnesses of other three, we only need to consider $\mathbb{Z}$-OV here.

From Theorem 4.10 and the assumption, there exists a uniform-$O(1)$-reduction. Proceeding similar as in Lemma 1.17 with the uniform-$O(1)$-reduction, we obtain a better dimensionality self reduction from OV to $\mathbb{Z}$-OV. Then exactly the same argument as in Theorem 1.18 with different parameters gives us the lower bound required. □

5 NP · UPP communication protocol and exact hardness for $\mathbb{Z}$-Max-IP

We note that the inapproximability results for (Boolean) Max-IP is established via a connection to the MA communication complexity protocol of Set-Disjointness [5]. In the light of this, in this section we view our reduction from OV to $\mathbb{Z}$-Max-IP (Lemma 1.17 and Theorem 4.3) in the perspective of communication complexity.

We observe that in fact, our reduction can be understood as an NP · UPP communication protocol for Set Disjointness. Moreover, we show that if we can get a slightly better NP · UPP communication protocol for Set-Disjointness, then we would be able to prove $\mathbb{Z}$-Max-IP is hard even for $\omega(1)$ dimensions (and also ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair).

5.1 $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocol for Set-Disjointness

First, we rephrase the results of Lemma 1.17 and Theorem 4.3 in a more convenient way for our use here.

Lemma 5.1 (Rephrasing of Lemma 1.17 and Theorem 4.3)**.** Let $1 \le \ell \le d$, and $m = \ell^{O(6^{\log^* d} \cdot (d/\ell))}$. There exists a family of functions

$$\psi^i_{\mathsf{Alice}}, \psi^i_{\mathsf{Bob}} : \{0,1\}^d \to \mathbb{R}^{(\ell+1)^2}$$

for $i \in [m]$ such that:

- when $x \cdot y = 0$, there is an i such that $\psi^i_{\mathsf{Alice}}(x) \cdot \psi^i_{\mathsf{Bob}}(y) \ge 0$;
- when $x \cdot y > 0$, for every i $\psi^i_{\mathsf{Alice}}(x) \cdot \psi^i_{\mathsf{Bob}}(y) < 0$;
- all $\psi^i_{\mathsf{Alice}}(x)$ and $\psi^i_{\mathsf{Bob}}(y)$ can be computed in $\mathrm{poly}(d)$ time.

We also need the standard connection between UPP communication protocols and sign-rank [63] (see also Chapter 4.11 of [49]).

Recall that for a function $F : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$, a UPP protocol for F is a private-coin randomized communication protocol such that: Alice and Bob hold $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ respectively; $F(x,y) = 1$ if and only if Alice and Bob accepts with probability $> 1/2$. The cost of the protocol is the maximum bits communicated over all pairs $(x,y) \in \mathcal{X} \times \mathcal{Y}$ and Alice's and Bob's corresponding private random coins.

Lemma 5.2 (Equivalence of sign-rank and UPP communication protocol (Theorem 3 of [63]))**.** The following holds.

- There is a d-cost UPP protocol for F implies that for $\ell = d+1$, there are mappings $\psi^{\mathcal{X}} : \mathcal{X} \to \mathbb{R}^{2^\ell}$ and $\psi^{\mathcal{Y}} : \mathcal{Y} \to \mathbb{R}^{2^\ell}$ such that for every $(x,y) \in \mathcal{X} \times \mathcal{Y}$:
  - if $F(x,y) = 1$, $\psi^{\mathcal{X}}(x) \cdot \psi^{\mathcal{Y}}(y) \ge 0$;
  - otherwise, $\psi^{\mathcal{X}}(x) \cdot \psi^{\mathcal{Y}}(y) < 0$.
- There are mappings $\psi^{\mathcal{X}} : \mathcal{X} \to \mathbb{R}^{2^\ell}$ and $\psi^{\mathcal{Y}} : \mathcal{Y} \to \mathbb{R}^{2^\ell}$ satisfying the above conditions implies that for $d = \ell + 1$, there is a d-cost UPP protocol for F.

From the above lemmas, we immediately get the needed communication protocol and prove Theorem 1.21 (restated below for convenience).

Reminder of Theorem 1.21 *For every $1 \le \alpha \le n$, there is an*

$$\left(\alpha \cdot 6^{\log^* n} \cdot (n/2^\alpha), O(\alpha)\right)\textit{-computational-efficient}$$

$\mathsf{NP} \cdot \mathsf{UPP}$ *communication protocol for DISJ_n.*

Proof sketch. We set $\alpha = \log \ell$ here. Given the function families $\{\psi^i_{\mathsf{Alice}}\}, \{\psi^i_{\mathsf{Bob}}\}$ from Lemma 5.1, Merlin just sends the index $i \in [m]$ and the rest follows from Lemma 5.2. □

5.2 Slightly better protocols imply hardness in dimension $\omega(1)$

Finally, we show that if we have a slightly better $\mathsf{NP} \cdot \mathsf{UPP}$ protocol for Set-Disjointness, then we can show $\mathbb{Z}$-Max-IP requires $n^{2-o(1)}$ time even for $\omega(1)$ dimensions (and so do ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair). We restate Theorem 1.22 here for convenience.

Reminder of Theorem 1.22 *Assuming SETH (or OVC), if there is an increasing and unbounded function f such that for every $1 \le \alpha \le n$, there is an*

$$(n/f(\alpha), \alpha)\text{-computational-efficient}$$

$\mathsf{NP} \cdot \mathsf{UPP}$ *communication protocol for DISJ_n, then $\mathbb{Z}$-Max-$\mathsf{IP}_{n,\omega(1)}$ requires $n^{2-o(1)}$ time with vectors of* $\operatorname{polylog}(n)$*-bit entries. The same holds for ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair.*

Proof. Suppose otherwise, there is a constant $\varepsilon_1 > 0$ and an algorithm $\mathbb{A}$ for $\mathbb{Z}$-Max-$\mathsf{IP}_{n,d}$ running in $n^{2-\varepsilon_1}$ time for all constants d. (Note from Lemma 4.6 and Lemma 4.7, we only need to consider $\mathbb{Z}$-Max-IP here.)

Now, letting c be an arbitrary constant, we are going to construct an $n^{2-\Omega(1)}$-time algorithm for $\mathsf{OV}_{n,c\log n}$, contradicting OVC.

Let $\varepsilon = \varepsilon_1/2$, and α be the first number such that $c/f(\alpha) < \varepsilon$. Note that α is also a constant. Consider the $(c\log n/f(\alpha), \alpha)$-computational-efficient $\mathsf{NP} \cdot \mathsf{UPP}$ protocol Π for $\mathsf{DISJ}_{c\log n}$, and let A, B be the two sets in the $\mathsf{OV}_{n,c\log n}$ instance. Our algorithm via reduction works as follows:

- There are 2^α possible messages in $\{0,1\}^\alpha$, let $m_1, m_2, \dots, m_{2^\alpha}$ be an enumeration of them.
- We first enumerate all possible advice strings from Merlin in Π. There are $2^{c\log n/f(\alpha)} \le 2^{\varepsilon \cdot \log n} = n^\varepsilon$ such strings, let $\phi \in \{0,1\}^{\varepsilon \cdot \log n}$ be such an advice string.
  - For each $x \in A$, let $\psi_{\mathsf{Alice}}(x) \in \mathbb{R}^{2^\alpha}$ be the probabilities that Alice accepts each message from Bob. That is, $\psi_{\mathsf{Alice}}(x)_i$ is the probability that Alice accepts the message m_i, given its input x and the advice ϕ.
  - Similarly, for each $y \in B$, let $\psi_{\mathsf{Bob}}(y) \in \mathbb{R}^{2^\alpha}$ be the probabilities that Bob sends each message. That is, $\psi_{\mathsf{Bob}}(y)_i$ is the probability that Bob sends the message m_i, give its input y and the advice ϕ.
  - Then, for each $x \in A$ and $y \in B$, $\psi_{\mathsf{Alice}}(x) \cdot \psi_{\mathsf{Bob}}(y)$ is precisely the probability that Alice accepts at the end when Alice and Bob hold x and y respectively and the advice is ϕ. Now we let A_ϕ be the set of all the vectors $\psi_{\mathsf{Alice}}(x)$, and B_ϕ be the set of all the vectors $\psi_{\mathsf{Bob}}(y)$.
- If there is a ϕ such that $\mathsf{OPT}(A_\phi, B_\phi) \ge 1/2$, then we output yes, and otherwise output no.

From the definition of Π, it is straightforward to see that the above algorithm solves $\mathsf{OV}_{n,c\cdot\log n}$. Moreover, notice that from the computational-efficient property of Π, the reduction itself works in $n^{1+\varepsilon} \cdot \operatorname{polylog}(n)$ time, and all the vectors in A_ϕ and B_ϕ have at most $\operatorname{polylog}(n)$ bit precision, which means $\mathsf{OPT}(A_\phi, B_\phi)$ can be solved by a call to $\mathbb{Z}$-Max-$\mathsf{IP}_{n,2^\alpha}$ with vectors of $\operatorname{polylog}(n)$-bit entries.

Hence, the final running time for the above algorithm is bounded by $n^\varepsilon \cdot n^{2-\varepsilon_1} = n^{2-\varepsilon}$ (2^α is still a constant), which contradicts the OVC. □

6 Improved MA protocols

In this section we prove Theorem 1.24 (restated below for convenience).

Reminder of Theorem 1.24 *There is an* MA *protocol for* DISJ_n *and* IP_n *with communication complexity*

$$O\left(\sqrt{n\log n\log\log n}\right).$$

To prove Theorem 1.24, we need the following intermediate problem.

Definition 6.1 (The Inner Product Modulo p Problem (IP_n^p))**.** Let p and n be two positive integers. In IP_n^p, Alice and Bob are given vectors X and Y in $\{0,1\}^n$ respectively and they want to compute $X \cdot Y \pmod{p}$.

Note that IP_n and IP_n^p are not Boolean functions, so we need to generalize the definition of an MA protocol. In an MA protocol for IP_n, Merlin sends the answer directly to Alice together with a proof to convince Alice and Bob. The correctness condition becomes that for the right answer $X \cdot Y$, Merlin has a proof such that Alice and Bob will accept with high probability (like $2/3$). And the soundness condition becomes that for the wrong answers, every proof from Merlin will be rejected with high probability.

We are going to use the following MA protocol for IP_n^p, which is a slight adaptation of the protocol from [67].

Lemma 6.2 (Implicit in Theorem 3.1 of [67])**.** For a sufficiently large prime q and integers T and n, there is an

$$\Big(O\left(n/T\cdot\log q\right), \log n + O(1), O\left(T\cdot\log q\right), 1/2\Big)\text{-efficient}$$

MA protocol for IP_n^q.

Now we ready to prove Theorem 1.24.

Proof of Theorem 1.24. Since an IP_n protocol trivially implies a DISJ_n protocol, we only need to consider IP_n in the following.

Now, let x be the number such that $x^x = n$. For convenience we are going to pretend that x is an integer. It is easy to see that $x = \Theta(\log n/\log\log n)$. Then we pick $10x$ distinct primes $p_1, p_2, \ldots, p_{10x}$ in $[x+1, x^2]$. (We can assume that n is large enough to make x satisfy the requirement of Lemma 2.4.) Let T be a parameter. We use Π_{p_i} to denote the $\big(O\left(n/T\cdot\log p_i\right), \log n + O(1), O\left(T\cdot\log p_i\right), 1/2\big)$-efficient MA protocol for $\mathsf{IP}_n^{p_i}$.

Our protocol for IP_n works as follows:

- Merlin sends Alice all the advice strings from the protocols $\Pi_{p_1}, \Pi_{p_2}, \ldots, \Pi_{p_{10x}}$, together with a presumed inner product $0 \le z \le n$.
- Note that Π_{p_i} contains the presumed value of $X \cdot Y \pmod{p_i}$, Alice first checks whether z is consistent with all these Π_{p_i}, and rejects immediately if it does not.
- Alice and Bob jointly toss $O(\log(10x))$ coins, to pick a uniform random number $i^\star \in [10x]$, and then they simulate $\Pi_{p_{i^\star}}$. That is, they pretend they are the Alice and Bob in the protocol $\Pi_{p_{i^\star}}$ with the advice from Merlin in $\Pi_{p_{i^\star}}$ (which Alice does have).

Correctness. Let $X,Y \in \{0,1\}^n$ be the vectors of Alice and Bob. If $X \cdot Y = z$, then by the definition of these protocols Π_{p_i}, Alice always accepts with the correct advice from Merlin. Otherwise, let $d = X \cdot Y \neq z$. We are going to analyze the probability that we pick a "good" $p_{i^\star}$ such that $p_{i^\star}$ does not divide $|d-z|$. Since $p_i > x$ for all the p_i and $x^x > n \geq |d-z|$, $|d-z|$ cannot be a multiplier of more than x primes among the p_i. Therefore, with probability at least 0.9, our pick of $p_{i^\star}$ is good. And in this case, from the definition of the protocols Π_{p_i}, Alice and Bob would reject afterward with probability at least $1/2$, as $d \pmod{p_{i^\star}}$ differs from $z \pmod{p_{i^\star}}$. In summary, when $X \cdot Y \neq z$, Alice rejects with probability at least $0.9/2 = 0.45$, which finishes the proof for the correctness.

Complexity. Now, note that the total advice length is

$$O\left(n/T \cdot \sum_{i=1}^{10x} \log p_i\right) = O\left(n/T \cdot \log \prod_{i=1}^{10x} x^2\right) = O\left(n/T \cdot \log x^{20x}\right) = O\left(n/T \cdot \log n\right).$$

And the communication complexity between Alice and Bob is bounded by

$$O\left(T \cdot \log x^2\right) = O\left(T \cdot \log\log n\right).$$

Setting $T = \sqrt{n \log n / \log\log n}$ balances the above two quantities, and we obtain the needed MA-protocol for IP_n. □

7 Hardness of Approximate $\{-1,1\}$-Max-IP via approximate polynomial for OR

In this section we apply the $O(\sqrt{n})$-degree approximate polynomial for OR [27, 77] to show hardness of approximate $\{-1,1\}$-Max-IP. We first give a reduction from OV to approximate $\{-1,1\}$-Max-IP.

Theorem 7.1. Letting $\varepsilon \in (0,1)$, an $\mathsf{OV}_{n,d}$ instance with sets A,B reduces to a $\{-1,1\}$-$\mathsf{Max\text{-}IP}_{n,d_1}$ instance with sets $\widetilde{A}$ and $\widetilde{B}$, such that:

- $d_1 = \binom{d}{\leq O\left(\sqrt{d \log 1/\varepsilon}\right)}^3 \cdot 2^{O\left(\sqrt{d \log 1/\varepsilon}\right)} \cdot \varepsilon^{-1}$, in which the notation $\binom{n}{\leq m}$ denotes $\sum_{i=0}^{m} \binom{n}{i}$.
- There is an integer $T > \varepsilon^{-1}$ such that if there is an $(a,b) \in A \times B$ such that $a \cdot b = 0$, then $\mathsf{OPT}(\widetilde{A},\widetilde{B}) \geq T$.
- Otherwise, $|\mathsf{OPT}(\widetilde{A},\widetilde{B})| \leq T \cdot \varepsilon$.
- Moreover, the reduction takes $n \cdot \mathrm{poly}(d_1)$ time.

We remark here that the above reduction fails to achieve a characterization: setting $\varepsilon = 1/2$ and $d = c \log n$ for an arbitrary constant c, we have $d_1 = 2^{\widetilde{O}(\sqrt{\log n})}$, much larger than $\log n$. Another interesting difference between the above theorem and Lemma 3.3 (the reduction from OV to approximate Max-IP) is that Lemma 3.3 reduces one OV instance to many Max-IP instances, while the above reduction only reduces it to one $\{-1,1\}$-Max-IP instance.

Proof of Theorem 7.1.

Construction and analysis of the polynomial $P_\varepsilon(z)$. By [27, 77], there is a polynomial $P_\varepsilon : \{0,1\}^d \to \mathbb{R}$ such that:

- P_ε is of degree $D = O\left(\sqrt{d \log 1/\varepsilon}\right)$.
- For every $z \in \{0,1\}^d$, $P_\varepsilon(z) \in [0,1]$.
- Given $z \in \{0,1\}^d$, if $\mathsf{OR}(z) = 0$, then $P_\varepsilon(z) \geq 1 - \varepsilon$, otherwise $P_\varepsilon(z) \leq \varepsilon$.
- P_ε can be constructed in time polynomial in its description size.

Now, let us analyze P_ε further. For a set $S \subseteq [d]$, let $\chi_S : \{0,1\}^d \to \mathbb{R}$ be $\chi_S(z) := \prod_{i \in S}(-1)^{z_i}$. Then we can write P_ε as

$$P_\varepsilon := \sum_{S \subseteq [d], |S| \leq D} \chi_S \cdot \langle \chi_S, P_\varepsilon \rangle,$$

where $\langle \chi_S, P_\varepsilon \rangle$ is the inner product of χ_S and P_ε, defined as $\langle \chi_S, P_\varepsilon \rangle := \mathbb{E}_{x \in \{0,1\}^d} \chi_S(x) \cdot P_\varepsilon(x)$.

Let $c_S = \langle \chi_S, P_\varepsilon \rangle$. From the definition it is easy to see that $c_S \in [-1,1]$.

Discretization of polynomial P_ε. Note that $P_\varepsilon(z)$ has real coefficients, we need to turn it into another polynomial with integer coefficients first.

Let $M := \binom{d}{\leq D}$. Consider the following polynomial $\widehat{P_\varepsilon}$:

$$\widehat{P_\varepsilon} := \sum_{S \subseteq [d], |S| \leq D} \lfloor c_S \cdot 2M/\varepsilon \rfloor \cdot \chi_S.$$

We can see that $|\widehat{P_\varepsilon}(z)/(2M/\varepsilon) - P_\varepsilon(z)| \leq \varepsilon$ for every $z \in \{0,1\}^d$, and we let $\hat{c}_S := \lfloor c_S \cdot M \cdot 2/\varepsilon \rfloor$ for convenience.

Simplification of the polynomial $\widehat{P_\varepsilon}$. $\widehat{P_\varepsilon}(z)$ is expressed over the basis consisting of the χ_S. We need to turn it into a polynomial over the standard basis.

For each $S \subseteq [d]$, consider χ_S, it can also be written as

$$\chi_S(z) = \prod_{i \in S}(-1)^{z_i} := \prod_{i \in S}(1 - 2z_i) = \sum_{T \subseteq S} (-2)^{|T|} z_T,$$

where $z_T := \prod_{i \in T} z_i$. Plugging it into the expression of $\widehat{P_\varepsilon}$, we have

$$\widehat{P_\varepsilon}(z) := \sum_{T \subseteq [d], |T| \leq D} \left(\sum_{S \subseteq [d], |S| \leq D, T \subseteq S} \hat{c}_S \right) \cdot (-2)^{|T|} z_T.$$

Set

$$\tilde{c}_T := \left(\sum_{S \subseteq [d], |S| \leq D, T \subseteq S} \hat{c}_S \right) \cdot (-2)^{|T|},$$

the above simplifies to

$$\widehat{P}_{\varepsilon}(z) := \sum_{T \subseteq [d], |T| \le D} \tilde{c}_T \cdot z_T .$$

Properties of the polynomial $\widehat{P}_{\varepsilon}$. Let us summarize some properties of $\widehat{P}_{\varepsilon}$ for now. First we need a bound on $|\tilde{c}_T|$. We can see $|\hat{c}_S| \le M \cdot 2/\varepsilon$, and by a simple calculation we have

$$|\tilde{c}_T| \le M^2 \cdot 2^D \cdot 2/\varepsilon .$$

Let $B = M^2 \cdot 2^D \cdot 2/\varepsilon$ for convenience. For $x, y \in \{0,1\}^d$, consider $\widehat{P}_{\varepsilon}(x,y) := \widehat{P}_{\varepsilon}(x_1y_1, x_2y_2, \dots, x_dy_d)$ (that is, plugging in $z_i = x_iy_i$), we have

$$\widehat{P}_{\varepsilon}(x,y) := \sum_{T \subseteq [d], |T| \le D} \tilde{c}_T \cdot x_T \cdot y_T ,$$

where $x_T := \prod_{i \in T} x_i$ and y_T is defined similarly. Moreover, we have

- If $x \cdot y = 0$, then $\widehat{P}_{\varepsilon}(x,y) \ge (2M/\varepsilon) \cdot (1 - 2\varepsilon)$.
- If $x \cdot y \ne 0$, then $|\widehat{P}_{\varepsilon}(x,y)| \le (2M/\varepsilon) \cdot 2\varepsilon$.

The reduction. Now, let us construct the reduction, we begin with some notation. For two vectors a, b, we use $a \circ b$ to denote their concatenation. For a vector a and a real τ, we use $a \cdot \tau$ to denote the vector resulting from multiplying each coordinate of a by τ. Let $\operatorname{sgn}(\tau)$ be the sign function that outputs 1 when $\tau > 0$, -1 when $\tau < 0$, and 0 when $\tau = 0$. For $\tau \in \{-B, -B+1, \dots, B\}$, we use $e_\tau \in \{-1,0,1\}^B$ to denote the vector whose first $|\tau|$ elements are $\operatorname{sgn}(\tau)$ and the rest are zeros. We also use $\mathbf{1}$ to denote the all-1 vector with length B.

Let $T_1, T_2, \dots, T_M$ be an enumeration of all subsets $T \subseteq [d]$ such that $|T| \le D$. We define

$$\varphi_x(x) := \circ_{i=1}^{M} (e_{\tilde{c}_{T_i}} \cdot x_{T_i}) \text{ and } \varphi_y(y) := \circ_{i=1}^{M} (\mathbf{1} \cdot y_{T_i}).$$

And we have

$$\varphi_x(x) \cdot \varphi_y(y) = \sum_{i=1}^{M} (e_{\tilde{c}_{T_i}} \cdot \mathbf{1}) \cdot (x_{T_i} \cdot y_{T_i}) = \sum_{i=1}^{M} \tilde{c}_{T_i} \cdot x_{T_i} \cdot y_{T_i} = \widehat{P}_{\varepsilon}(x,y).$$

To move from $\{-1,0,1\}$ to $\{-1,1\}$, we use the following carefully designed reductions $\psi_x, \psi_y : \{-1,0,1\} \to \{-1,1\}^2$, such that

$$\psi_x(-1) = \psi_y(-1) = (-1,-1,-1,-1), \quad \psi_x(0) = (-1,-1,1,1),$$

$$\psi_y(0) := (1,-1,1,-1), \quad \text{and} \quad \psi_x(1) = \psi_y(1) = (1,1,1,1).$$

It is easy to check that for $a, b \in \{-1,0,1\}$, we have $\psi_x(a) \cdot \psi_y(b) = 4 \cdot (a \cdot b)$.

Hence, composing the above two reductions, we get our desired reductions $\phi_x = \psi_x^{\otimes(B \cdot M)} \circ \varphi_x$ and $\phi_y = \psi_y^{\otimes(B \cdot M)} \circ \varphi_y$ such that for $x, y \in \{0,1\}^d$, $\phi_x(x), \phi_y(y) \in \{-1,1\}^{4B \cdot M}$ and $\phi_x(x) \cdot \phi_y(y) = 4 \cdot \widehat{P}_{\varepsilon}(x,y)$.

Finally, given an $\mathsf{OV}_{n,d}$ instance with two sets A and B, we construct two sets $\widetilde{A}$ and $\widetilde{B}$, such that $\widetilde{A}$ consists of all the vectors $\phi_x(x)$ for $x \in A$, and $\widetilde{B}$ consists of all the vectors $\phi_y(y)$ for $y \in B$.

Then we can see $\widetilde{A}$ and $\widetilde{B}$ consist of n vectors from $\{-1,1\}^{d_1}$, where

$$d_1 = 4B \cdot M = M^3 \cdot 2^D \cdot 8/\varepsilon = \left(\frac{d}{\leq O\left(\sqrt{d \log 1/\varepsilon}\right)} \right)^3 \cdot 2^{O\left(\sqrt{d \log 1/\varepsilon}\right)} \cdot \varepsilon^{-1}$$

as stated.

It is not hard to see the above reduction takes $n \cdot \mathrm{poly}(d_1)$ time. Moreover, if there is an $(x,y) \in A \times B$ such that $x \cdot y = 0$, then $\mathsf{OPT}(\widetilde{A},\widetilde{B}) \geq (8M/\varepsilon) \cdot (1-2\varepsilon)$, otherwise, $\mathsf{OPT}(\widetilde{A},\widetilde{B}) \leq (8M/\varepsilon) \cdot 2\varepsilon$. Setting ε above to be $1/3$ times the ε in the statement finishes the proof. □

With Theorem 7.1, we are ready to prove our hardness results on $\{-1,1\}$-Max-IP.

Theorem 7.2. Assume SETH (or OVC). Letting $\alpha : \mathbb{N} \to \mathbb{R}$ be any function of n such that $\alpha(n) = n^{o(1)}$, there is another function β satisfying $\beta(n) = n^{o(1)}$ and an integer $T > \alpha$ (β and T depend on α), such that there is no $n^{2-\Omega(1)}$-time algorithm for $\{-1,1\}$-Max-$\mathsf{IP}_{n,\beta(n)}$ distinguishing the following two cases:

- $\mathsf{OPT}(A,B) \geq T$ (A and B are the sets in the $\{-1,1\}$-Max-IP instance).
- $|\mathsf{OPT}(A,B)| \leq T/\alpha(n)$.

Proof. Letting $\alpha = n^{o(1)}$ and $k = \log\alpha/\log n$, we have $k = o(1)$. Setting $d = c\log n$ where c is an arbitrary constant and $\varepsilon = \alpha^{-1}$ in Theorem 7.1, we have that an $\mathsf{OV}_{c\log n}$ reduces to a certain $\alpha(n)$-approximation to a $\{-1,1\}$-Max-IP_{n,d_1} instance with sets A and B, where

$$d_1 = \left(\frac{c\log n}{\leq O(\sqrt{ck}\log n)} \right)^3 \cdot 2^{O(\sqrt{ck}\log n)} \leq \left(\frac{\sqrt{c}}{\sqrt{k}} \right)^{O(\sqrt{ck}\log n)} \cdot 2^{O(\sqrt{ck}\log n)} = n^{O(\log(c/k)\cdot\sqrt{ck})}.$$

Now let $\beta = n^{k^{1/3}}$ and T be the integer specified by Theorem 7.1. Since $k = o(1)$, $\beta = n^{o(1)}$. Suppose otherwise there is an $n^{2-\Omega(1)}$-time algorithm for distinguishing whether $\mathsf{OPT}(A,B) \geq T$ or $|\mathsf{OPT}(A,B)| \leq T/\alpha(n)$. Then for any constant c, $O(\log(c/k)\sqrt{ck}) \leq k^{1/3}$ for sufficiently large n, which means $d_1 \leq \beta(n)$ for a sufficiently large n, and there is an $n^{2-\Omega(1)}$-time algorithm for $\mathsf{OV}_{c\log n}$ by Theorem 7.1, contradiction to OVC. □

8 Future work

We end our paper by discussing a few interesting research directions.

1. The most important future direction from this work is to further improve the dimensionality reduction for OV. It is certainly weird to consider $2^{O(\log^* n)}$ to be the right answer for the limit of the dimensionality reduction. This term seems to follow from the limitation of our recursive number-theoretical construction, and not from the nature of the problem itself. We conjecture that there should be an $\omega(1)$ dimensional reduction with a more direct construction.

One possible direction is to combine the original polynomial-based construction from [75] together with our new number-theoretical one. These two approaches seem completely different, hence a clever combination of them may prove our conjecture.

2. In order to prove $\omega(1)$ dimensional hardness for ℓ_2-Furthest Pair and Bichromatic ℓ_2-Closest Pair, we can also bypass the OV dimensionality reduction approach by proving $\omega(1)$ dimensional hardness for $\mathbb{Z}$-Max-IP directly. One possible way to approach this question is to start from the $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocol connection as in Section 5 (apply Theorem 1.22), and (potentially) draw some connections from some known UPP communication protocols.

3. We have seen an efficient reduction from $\mathbb{Z}$-OV to $\mathbb{Z}$-Max-IP which only blows up the dimension quadratically, is there a similar reduction from $\mathbb{Z}$-Max-IP back to $\mathbb{Z}$-OV? Are $\mathbb{Z}$-Max-IP and $\mathbb{Z}$-OV equivalent?

4. By making use of the new AG-code based MA protocols, we can shave a $\widetilde{O}(\sqrt{\log n})$ factor from the communication complexity, can we obtain an $O(\sqrt{n})$ MA communication protocol matching the lower bound for DISJ_n? It seems new ideas are required.

   Since our MA protocol works for both DISJ and IP, and IP does seem to be a harder problem. It may be better to find an MA protocol only works for DISJ. It is worth noting that an $O(\sqrt{n})$ AMA communication protocol for DISJ is given by [67], which doesn't work for IP.

5. Can the dependence on ε in the algorithms from Theorem 1.5 be further improved? Is it possible to apply ideas in the $n^{2-1/\widetilde{\Omega}(\sqrt{c})}$ algorithm for $\mathsf{Max\text{-}IP}_{n,c\log n}$ from [13]?

6. For the complexity of computing 2-multiplicative-approximation to $\mathsf{Max\text{-}IP}_{n,c\log n}$, Theorem 1.5 implies that there is an algorithm running in $n^{2-1/O(\log c)}$ time, the same as the best algorithm for $\mathsf{OV}_{n,c\log n}$ [6]. Is this just a coincidence? Or are there some connections between these two problems?

7. We obtain a connection between hardness of $\mathbb{Z}$-Max-IP and $\mathsf{NP} \cdot \mathsf{UPP}$ communication protocols for Set-Disjointness. Can we get similar connections from other $\mathsf{NP} \cdot \mathcal{C}$ type communication protocols for Set-Disjointness? Some candidates include $\mathsf{NP} \cdot \mathsf{SBP}$ and $\mathsf{NP} \cdot \mathsf{promiseBQP}$ (QCMA).

Acknowledgment

I would like to thank Ryan Williams for introducing the problem to me, countless encouragement and helpful discussions during this work, and also many comments on a draft of this paper. In particular, the idea of improving OV dimensionality self-reduction using CRT (the direct CRT based approach) is introduced to me by Ryan Williams.

I am grateful to Virginia Vassilevska Williams, Kaifeng Lyu and Peilin Zhong for helpful discussions and suggestions. I would like to thank Aviad Rubinstein for sharing a manuscript of his paper, and pointing out that the $O(\sqrt{n \log n \log\log n})$ MA protocol also works for Inner Product.

References

[1] SCOTT AARONSON AND AVI WIGDERSON: Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory*, 1(1):2:1–2:54, 2009. Preliminary version in STOC'08. [doi:10.1145/1490270.1490272] 4, 11, 13

[2] AMIR ABBOUD AND ARTURS BACKURS: Towards hardness of approximation for polynomial time problems. In *Proc. 8th Symp. Innovations in Theoretical Computer Science (ITCS'17)*, pp. 11:1–11:26. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. [doi:10.4230/LIPIcs.ITCS.2017.11] 14

[3] AMIR ABBOUD AND SØREN DAHLGAARD: Popular conjectures as a barrier for dynamic planar graph algorithms. In *Proc. 57th FOCS*, pp. 477–486. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.58, arXiv:1605.03797] 13

[4] AMIR ABBOUD AND AVIAD RUBINSTEIN: Fast and deterministic constant factor approximation algorithms for LCS imply new circuit lower bounds. In *Proc. 9th Symp. Innovations in Theoretical Computer Science (ITCS'18)*, pp. 35:1–35:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018. [doi:10.4230/LIPIcs.ITCS.2018.35] 13, 14

[5] AMIR ABBOUD, AVIAD RUBINSTEIN, AND R. RYAN WILLIAMS: Distributed PCP theorems for hardness of approximation in P. In *Proc. 58th FOCS*, pp. 25–36. IEEE Comp. Soc., 2017. [doi:10.1109/FOCS.2017.12, arXiv:1706.06407] 3, 4, 6, 8, 9, 13, 14, 23, 24, 35

[6] AMIR ABBOUD, RICHARD R. RYAN WILLIAMS, AND HUACHENG YU: More applications of the polynomial method to algorithm design. In *Proc. 26th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'15)*, pp. 218–230. ACM Press, 2015. [doi:10.1137/1.9781611973730.17] 6, 43

[7] AMIR ABBOUD AND VIRGINIA VASSILEVSKA WILLIAMS: Popular conjectures imply strong lower bounds for dynamic problems. In *Proc. 55th FOCS*, pp. 434–443. IEEE Comp. Soc., 2014. [doi:10.1109/FOCS.2014.53, arXiv:1402.0054] 13

[8] AMIR ABBOUD, VIRGINIA VASSILEVSKA WILLIAMS, AND OREN WEIMANN: Consequences of faster alignment of sequences. In *Proc. 41st Internat. Colloq. on Automata, Languages and Programming (ICALP'14)*, pp. 39–51. Springer, 2014. [doi:10.1007/978-3-662-43948-7_4] 5, 13

[9] AMIR ABBOUD, VIRGINIA VASSILEVSKA WILLIAMS, AND HUACHENG YU: Matching triangles and basing hardness on an extremely popular conjecture. *SIAM J. Comput.*, 47(3):1098–1122, 2018. Preliminary version in STOC'15. [doi:10.1137/15M1050987] 13

[10] PANKAJ K. AGARWAL, HERBERT EDELSBRUNNER, OTFRIED SCHWARZKOPF, AND EMO WELZL: Euclidean minimum spanning trees and bichromatic closest pairs. *Discrete Comput. Geom.*, 6(3):407–422, 1991. Preliminary version in SoCG'90. [doi:10.1007/BF02574698] 4

[11] THOMAS DYBDAHL AHLE, RASMUS PAGH, ILYA P. RAZENSHTEYN, AND FRANCESCO SILVESTRI: On the complexity of inner product similarity join. In *Proc. 35th ACM Symp. Principles

of Database Sys. (PODS'16), pp. 151–164. ACM Press, 2016. [doi:10.1145/2902251.2902285, arXiv:1510.02824] 3, 6, 8, 14

[12] JOSH ALMAN: An illuminating algorithm for the light bulb problem. In *2nd Symp. on Simplicity in Algorithms (SOSA'19)*, pp. 2:1–2:11, 2019. [doi:10.4230/OASIcs.SOSA.2019.2, arXiv:1810.06740] 17

[13] JOSH ALMAN, TIMOTHY M. CHAN, AND R. RYAN WILLIAMS: Polynomial representations of threshold functions and algorithmic applications. In *Proc. 57th FOCS*, pp. 467–476. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.57, arXiv:1608.04355] 5, 7, 20, 22, 43

[14] JOSH ALMAN AND R. RYAN WILLIAMS: Probabilistic polynomials and Hamming nearest neighbors. In *Proc. 56th FOCS*, pp. 136–150. IEEE Comp. Soc., 2015. [doi:10.1109/FOCS.2015.18, arXiv:1507.05106] 3, 5, 7

[15] ALEXANDR ANDONI AND PIOTR INDYK: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Comm. ACM*, 51(1):117–122, 2008. Conference version in FOCS'06. [doi:10.1145/1327452.1327494] 3

[16] ALEXANDR ANDONI, PIOTR INDYK, THIJS LAARHOVEN, ILYA P. RAZENSHTEYN, AND LUDWIG SCHMIDT: Practical and optimal LSH for angular distance. In *Adv. in Neural Inform. Proc. Systems (NIPS'15)*, pp. 1225–1233. MIT Press, 2015. NIPS. [arXiv:1509.02897] 3

[17] ALEXANDR ANDONI, PIOTR INDYK, HUY LÊ NGUYỄN, AND ILYA P. RAZENSHTEYN: Beyond locality-sensitive hashing. In *Proc. 25th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'14)*, pp. 1018–1028. ACM Press, 2014. [doi:10.1137/1.9781611973402.76, arXiv:1306.1547] 3

[18] ALEXANDR ANDONI AND ILYA P. RAZENSHTEYN: Optimal data-dependent hashing for approximate near neighbors. In *Proc. 47th STOC*, pp. 793–801. ACM Press, 2015. [doi:10.1145/2746539.2746553, arXiv:1501.01062] 3

[19] TOM M. APOSTOL: *Introduction to Analytic Number Theory*. Springer, 2013. [doi:10.1007/978-1-4757-5579-4] 15

[20] SANJEEV ARORA AND BOAZ BARAK: *Computational Complexity – A Modern Approach*. Cambridge Univ. Press, 2009. 16

[21] ARTURS BACKURS AND PIOTR INDYK: Which regular expression patterns are hard to match? In *Proc. 57th FOCS*, pp. 457–466. IEEE Comp. Soc., 2016. [doi:10.1109/FOCS.2016.56, arXiv:1511.07070] 13

[22] ARTURS BACKURS AND PIOTR INDYK: Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM J. Comput.*, 47(3):1087–1097, 2018. Preliminary version in STOC'15. [doi:10.1137/15M1053128, arXiv:1412.0348] 13

[23] JON LOUIS BENTLEY AND MICHAEL IAN SHAMOS: Divide-and-conquer in multidimensional space. In *Proc. 8th STOC*, pp. 220–230. ACM Press, 1976. [doi:10.1145/800113.803652] 9

[24] KARL BRINGMANN: Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Proc. 55th FOCS*, pp. 661–670. IEEE Comp. Soc., 2014. [doi:10.1109/FOCS.2014.76, arXiv:1404.1448] 13

[25] KARL BRINGMANN, ALLAN GRØNLUND, AND KASPER GREEN LARSEN: A dichotomy for regular expression membership testing. In *Proc. 58th FOCS*, pp. 307–318. IEEE Comp. Soc., 2017. [doi:10.1109/FOCS.2017.36, arXiv:1611.00918] 13

[26] KARL BRINGMANN AND MARVIN KÜNNEMANN: Multivariate fine-grained complexity of longest common subsequence. In *Proc. 29th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'18)*, pp. 1216–1235. ACM Press, 2018. [doi:10.1137/1.9781611975031.79, arXiv:1803.00938] 13

[27] HARRY BUHRMAN, RICHARD CLEVE, RONALD DE WOLF, AND CHRISTOF ZALKA: Bounds for small-error and zero-error quantum algorithms. In *Proc. 40th FOCS*, pp. 358–368. IEEE Comp. Soc., 1999. [doi:10.1109/SFFCS.1999.814607, arXiv:cs/9904019] 8, 39, 40

[28] HARRY BUHRMAN, RICHARD CLEVE, AND AVI WIGDERSON: Quantum vs. classical communication and computation. In *Proc. 30th STOC*, pp. 63–68. ACM Press, 1998. [doi:10.1145/276698.276713, arXiv:quant-ph/9802040] 8

[29] CHRIS CALABRO, RUSSELL IMPAGLIAZZO, AND RAMAMOHAN PATURI: The complexity of satisfiability of small depth circuits. In *4th Internat. Workshop on Parameterized and Exact Computation (IWPEC'09)*, pp. 75–85. Springer, 2009. [doi:10.1007/978-3-642-11269-0_6] 5

[30] TIMOTHY M. CHAN: A (slightly) faster algorithm for Klee's measure problem. *Comput. Geom.*, 43(3):243–250, 2010. Preliminary version in SoCG'08. [doi:10.1016/j.comgeo.2009.01.007] 11

[31] LIJIE CHEN, SHAFI GOLDWASSER, KAIFENG LYU, GUY N. ROTHBLUM, AND AVIAD RUBINSTEIN: Fine-grained complexity meets IP = PSPACE. In *Proc. 30th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'19)*, pp. 1–20. ACM Press, 2019. [doi:10.1137/1.9781611975482.1, arXiv:1805.02351] 13

[32] LIJIE CHEN AND R. RYAN WILLIAMS: An equivalence class for orthogonal vectors. In *Proc. 30th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'19)*, pp. 21–40. ACM Press, 2019. [doi:10.1137/1.9781611975482.2, arXiv:1811.12017] 13

[33] TOBIAS CHRISTIANI: A framework for similarity search with space-time tradeoffs using locality-sensitive filtering. In *Proc. 28th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'17)*, pp. 31–46. ACM Press, 2017. [doi:10.1137/1.9781611974782.3, arXiv:1605.02687] 3

[34] TOBIAS CHRISTIANI AND RASMUS PAGH: Set similarity search beyond minhash. In *Proc. 49th STOC*, pp. 1094–1107. ACM Press, 2017. [doi:10.1145/3055399.3055443, arXiv:1612.07710] 3

[35] DON COPPERSMITH: Rapid multiplication of rectangular matrices. *SIAM J. Comput.*, 11(3):467–471, 1982. [doi:10.1137/0211037] 14

[36] SVYATOSLAV COVANOV AND EMMANUEL THOMÉ: Fast integer multiplication using generalized Fermat primes. *Math. Comput.*, 88(317):1449–1477, 2019. [doi:10.1090/mcom/3367] 11

[37] ROEE DAVID, KARTHIK C. S., AND BUNDIT LAEKHANUKIT: On the complexity of closest pair via polar-pair of point-sets. *SIAM J. Discrete Math.*, 33(1):509–527, 2019. Preliminary version in SoCG'18. [doi:10.1137/17M1128733, arXiv:1608.03245] 13

[38] MARTIN DIETZFELBINGER, TORBEN HAGERUP, JYRKI KATAJAINEN, AND MARTTI PENTTONEN: A reliable randomized algorithm for the closest-pair problem. *J. Algorithms*, 25(1):19–51, 1997. [doi:10.1006/jagm.1997.0873] 9

[39] MARTIN FÜRER: Faster integer multiplication. *SIAM J. Comput.*, 39(3):979–1005, 2009. Preliminary version in STOC'07. [doi:10.1137/070711761] 11

[40] JIAWEI GAO, RUSSELL IMPAGLIAZZO, ANTONINA KOLOKOLOVA, AND R. RYAN WILLIAMS: Completeness for first-order properties on sparse structures with algorithmic applications. *ACM Trans. Algorithms*, 15(3):23:1–23:35, 2018. Preliminary version in SODA'17. [doi:10.1145/3196275] 13

[41] ISAAC GOLDSTEIN, TSVI KOPELOWITZ, MOSHE LEWENSTEIN, AND ELY PORAT: Conditional lower bounds for space/time tradeoffs. In *Proc. 15th Internat. Workshop on Algorithms and Data Structures (WADS'17)*, pp. 421–436. Springer, 2017. [doi:10.1007/978-3-319-62127-2_36, arXiv:1706.05847] 13

[42] LOV K. GROVER: A fast quantum mechanical algorithm for database search. In *Proc. 28th STOC*, pp. 212–219. ACM Press, 1996. [doi:10.1145/237814.237866, arXiv:quant-ph/9605043] 8

[43] SARIEL HAR-PELED, PIOTR INDYK, AND RAJEEV MOTWANI: Approximate nearest neighbors: Towards removing the curse of dimensionality. *Theory of Computing*, 8(14):321–350, 2012. Preliminary versions in STOC'98 and FOCS'01. [doi:10.4086/toc.2012.v008a014] 3

[44] DAVID HARVEY, JORIS VAN DER HOEVEN, AND GRÉGOIRE LECERF: Even faster integer multiplication. *J. Complexity*, 36(C):1–30, 2016. [doi:10.1016/j.jco.2016.03.001, arXiv:1407.3360] 11

[45] MONIKA HENZINGER, SEBASTIAN KRINNINGER, DANUPON NANONGKAI, AND THATCHAPHOL SARANURAK: Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proc. 47th STOC*, pp. 21–30. ACM Press, 2015. [doi:10.1145/2746539.2746609, arXiv:1511.06773] 13

[46] MONIKA HENZINGER, ANDREA LINCOLN, STEFAN NEUMANN, AND VIRGINIA VASSILEVSKA WILLIAMS: Conditional hardness for sensitivity problems. In *Proc. 8th Symp. Innovations in Theoretical Computer Science (ITCS'17)*, pp. 26:1–26:31. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. [doi:10.4230/LIPIcs.ITCS.2017.26, arXiv:1703.01638] 13

[47] RUSSELL IMPAGLIAZZO AND RAMAMOHAN PATURI: On the complexity of k-SAT. *J. Comput. System Sci.*, 62(2):367–375, 2001. Preliminary version in CCC'99. [doi:10.1006/jcss.2000.1727] 3, 5

[48] PIOTR INDYK AND RAJEEV MOTWANI: Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th STOC*, pp. 604–613. ACM Press, 1998. [doi:10.1145/276698.276876] 3

[49] STASYS JUKNA: *Boolean Function Complexity: Advances and Frontiers*. Volume 27. Springer, 2012. [doi:10.1007/978-3-642-24508-4] 36

[50] MATTI KARPPA, PETTERI KASKI, AND JUKKA KOHONEN: A faster subquadratic algorithm for finding outlier correlations. *ACM Trans. Algorithms*, 14(3):31:1–31:26, 2018. Preliminary version in SODA'16. [doi:10.1145/3174804, arXiv:1510.03895] 3

[51] KARTHIK C. S., BUNDIT LAEKHANUKIT, AND PASIN MANURANGSI: On the parameterized complexity of approximating dominating set. *J. ACM*, 66(5):33:1–33:38, 2019. Preliminary version in STOC'18. [doi:10.1145/3325116, arXiv:1711.11029] 4, 13, 15, 23

[52] KARTHIK C. S. AND PASIN MANURANGSI: On closest pair in Euclidean metric: Monochromatic is as hard as bichromatic. *Combinatorica*, 2020. Preliminary version in ITCS'19. [doi:10.1007/s00493-019-4113-1, arXiv:1812.00901] 14

[53] SAMIR KHULLER AND YOSSI MATIAS: A simple randomized sieve algorithm for the closest-pair problem. *Inform. and Comput.*, 118(1):34–37, 1995. [doi:10.1006/inco.1995.1049] 9

[54] HARTMUT KLAUCK: Rectangle size bounds and threshold covers in communication complexity. In *Proc. 18th Conf. Computational Complexity (CCC'03)*, pp. 118–134. IEEE Comp. Soc., 2003. [doi:10.1109/CCC.2003.1214415, arXiv:cs/0208006] 11

[55] TSVI KOPELOWITZ, SETH PETTIE, AND ELY PORAT: Higher lower bounds from the 3SUM conjecture. In *Proc. 27th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'16)*, pp. 1272–1287. ACM Press, 2016. [doi:10.1137/1.9781611974331.ch89, arXiv:1407.6756] 13

[56] ROBERT KRAUTHGAMER AND OHAD TRABELSI: Conditional lower bounds for all-pairs max-flow. *ACM Trans. Algorithms*, 14(4):42:1–42:15, 2018. Preliminary version in ICALP'17. [doi:10.1145/3212510, arXiv:1702.05805] 13

[57] FRANÇOIS LE GALL AND FLORENT URRUTIA: Improved rectangular matrix multiplication using powers of the Coppersmith-Winograd tensor. In *Proc. 29th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'18)*, pp. 1029–1046. ACM Press, 2018. [doi:10.1137/1.9781611975031.67, arXiv:1708.05622] 14

[58] JIŘÍ MATOUŠEK: Efficient partition trees. *Discrete Comput. Geom.*, 8(3):315–334, 1992. Preliminary version in SoCG'91. [doi:10.1007/BF02293051] 4

[59] JIŘÍ MATOUŠEK: Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993. Preliminary version in SoCG'92. [doi:10.1007/BF02573972] 11

[60] BEHNAM NEYSHABUR AND NATHAN SREBRO: On symmetric and asymmetric LSHs for inner product search. In *Proc. 32nd Int. Conf. on Machine Learning (ICML'15)*, pp. 1926–1934, 2015. MLR Press. [arXiv:1410.5518] 3

[61] MIHAI PĂTRAŞCU: Towards polynomial lower bounds for dynamic problems. In *Proc. 42nd STOC*, pp. 603–610. ACM Press, 2010. [doi:10.1145/1806689.1806772] 13

[62] MIHAI PĂTRAŞCU AND R. RYAN WILLIAMS: On the possibility of faster SAT algorithms. In *Proc. 21st Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'10)*, pp. 1065–1075. ACM Press, 2010. [doi:10.1137/1.9781611973075.86] 13

[63] RAMAMOHAN PATURI AND JANOS SIMON: Probabilistic communication complexity. *J. Comput. System Sci.*, 33(1):106–123, 1986. Preliminary version in FOCS'84. [doi:10.1016/0022-0000(86)90046-2] 10, 36

[64] ALI RAHIMI AND BENJAMIN RECHT: Random features for large-scale kernel machines. In *Adv. in Neural Inform. Proc. Systems (NIPS'07)*, pp. 1177–1184. MIT Press, 2007. NIPS. 3

[65] PARIKSHIT RAM AND ALEXANDER G. GRAY: Maximum inner-product search using cone trees. In *Proc. 18th Internat. Conf. on Knowledge Discovery and Data Mining (KDD'12)*, pp. 931–939. ACM Press, 2012. [doi:10.1145/2339530.2339677] 3

[66] LIAM RODITTY AND VIRGINIA VASSILEVSKA WILLIAMS: Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proc. 45th STOC*, pp. 515–524. ACM Press, 2013. [doi:10.1145/2488608.2488673] 13

[67] AVIAD RUBINSTEIN: Hardness of approximate nearest neighbor search. In *Proc. 50th STOC*, pp. 1260–1268. ACM Press, 2018. [doi:10.1145/3188745.3188916, arXiv:1803.00904] 4, 6, 7, 11, 13, 19, 20, 21, 22, 23, 38, 43

[68] ANSHUMALI SHRIVASTAVA AND PING LI: Asymmetric LSH (ALSH) for sublinear time maximum inner product search (MIPS). In *Adv. in Neural Inform. Proc. Systems (NIPS'14)*, pp. 2321–2329. MIT Press, 2014. NIPS. [arXiv:1405.5869] 3

[69] ANSHUMALI SHRIVASTAVA AND PING LI: Asymmetric minwise hashing for indexing binary inner products and set containment. In *Proc. 24th Int. World Wide Web Conf. (WWW'15)*, pp. 981–991, 2015. [doi:10.1145/2736277.2741285] 3

[70] CHRISTINA TEFLIOUDI AND RAINER GEMULLA: Exact and approximate maximum inner product search with LEMP. *ACM Trans. Database Syst.*, 42(1):5:1–5:49, 2016. [doi:10.1145/2996452] 3

[71] GREGORY VALIANT: Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. *J. ACM*, 62(2):13:1–13:45, 2015. Preliminary version in FOCS'12. [doi:10.1145/2728167] 3

[72] VIRGINIA VASSILEVSKA WILLIAMS: On some fine-grained questions in algorithms and complexity. In *Proc. Internat. Congr. of Mathematicians (ICM 2018)*, volume 4, pp. 3447–3487. World Scientific, 2019. [doi:10.1142/9789813272880_0188] 13

[73] R. RYAN WILLIAMS: A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoret. Comput. Sci.*, 348(2–3):357–365, 2005. Preliminary version in ICALP'04. [doi:10.1016/j.tcs.2005.09.023] 3, 4, 5

[74] R. RYAN WILLIAMS: Faster all-pairs shortest paths via circuit complexity. *SIAM J. Comput.*, 47(5):1965–1985, 2018. Preliminary version in STOC'14. [doi:10.1137/15M1024524, arXiv:1312.6680] 6

[75] R. RYAN WILLIAMS: On the difference between closest, furthest, and orthogonal pairs: Nearly-linear vs barely-subquadratic complexity. In *Proc. 29th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'18)*, pp. 1207–1215. ACM Press, 2018. [doi:10.1137/1.9781611975031.78, arXiv:1709.05282] 4, 8, 9, 11, 12, 13, 25, 29, 30, 31, 32, 43

[76] R. RYAN WILLIAMS AND HUACHENG YU: Finding orthogonal vectors in discrete structures. In *Proc. 25th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'14)*, pp. 1867–1877. ACM Press, 2014. [doi:10.1137/1.9781611973402.135] 5

[77] RONALD DE WOLF: A note on quantum algorithms and the minimal degree of ε-error polynomials for symmetric functions. *Quantum Info. Comput.*, 8(10):943–950, 2008. [doi:10.26421/QIC8.10, arXiv:0802.1816] 8, 39, 40

[78] ANDREW CHI-CHIH YAO: On constructing minimum spanning trees in k-dimensional spaces and related problems. *SIAM J. Comput.*, 11(4):721–736, 1982. [doi:10.1137/0211059] 4

AUTHOR

Lijie Chen
Ph. D. student
MIT, Cambridge, MA
lijieche@mit.edu
http://www.mit.edu/~lijieche/

ABOUT THE AUTHOR

LIJIE CHEN got an undergraduate degree from Tsinghua University in 2017, and is now a Ph. D. student at MIT. His advisor is Ryan Williams. He likes turning red bull into theorems. While not doing that, he enjoys playing music games.