Approximating APSP without Scaling: Equivalence of Approximate Min-Plus and Exact Min-Max

Karl Bringmann^{*} Marvin Künnemann[†]

Karol Węgrzycki[‡]

Abstract

Zwick's $(1+\varepsilon)$ -approximation algorithm for the All Pairs Shortest Path (APSP) problem runs in time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$, where $\omega \leq 2.373$ is the exponent of matrix multiplication and W denotes the largest weight. This can be used to approximate several graph characteristics including the diameter, radius, median, minimum-weight triangle, and minimum-weight cycle in the same time bound.

Since Zwick's algorithm uses the scaling technique, it has a factor $\log W$ in the running time. In this paper, we study whether APSP and related problems admit approximation schemes avoiding the scaling technique. That is, the number of arithmetic operations should be independent of W; this is called *strongly polynomial*. Our main results are as follows.

- We design approximation schemes in strongly polynomial time $\mathcal{O}(\frac{n^{\omega}}{\varepsilon} \operatorname{polylog}(\frac{n}{\varepsilon}))$ for APSP on undirected graphs as well as for the graph characteristics diameter, radius, median, minimum-weight triangle, and minimum-weight cycle on directed or undirected graphs.
- For APSP on directed graphs we design an approximation scheme in strongly polynomial time $\mathcal{O}(n^{\frac{\omega+3}{2}}\varepsilon^{-1}\operatorname{polylog}(\frac{n}{\varepsilon}))$. This is significantly faster than the best exact algorithm.
- We explain why our approximation scheme for APSP on directed graphs has a worse exponent than ω : Any improvement over our exponent $\frac{\omega+3}{2}$ would improve the best known algorithm for Min-Max Product. In fact, we prove that approximating directed APSP and exactly computing the Min-Max Product are equivalent.

Our techniques yield a framework for approximation problems over the (min,+)-semiring that can be applied more generally. In particular, we obtain the first strongly polynomial approximation scheme for Min-Plus Convolution in strongly subquadratic time, and we prove an equivalence of approximate Min-Plus Convolution and exact Min-Max Convolution.

^{*}Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, kbringma@mpi-inf.mpg.de

[†]Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, marvin@mpi-inf.mpg.de

[‡]Institute of Informatics, University of Warsaw, Warsaw, Poland, k.wegrzycki@mimuw.edu.pl. Supported by the grants 2016/21/N/ST6/01468 and 2018/28/T/ST6/00084 of the Polish National Science Center and project TOTAL that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 677651).

1 Introduction

Scaling is one of the most fundamental algorithmic techniques. On a problem involving weights from a range $\{1, \ldots, W\}$, the main idea of scaling is to consider each of the log W bits one-by-one. Roughly speaking, in each phase we only consider the current bit, which simplifies the weighted problem to an unweighted problem. The scaling technique was particularly successful for graph problems (e.g., [27, 22, 17, 38, 21]). For instance, a scaling-based algorithm solves maximum weighted matching in time $\mathcal{O}(m\sqrt{n}\log(nW))$ [22], which was recently improved to time $\mathcal{O}(m\sqrt{n}\log W)$ [17].

However, in some situations scaling-based algorithms may be slower than alternative approaches, since they naturally require a factor log W in the running time. In particular, in practice weights are often given as floating-point numbers, and thus log W can easily be as large as n, rendering most scaling-based algorithms inferior to naive approaches. For this reason as well as for the genuinely theoretical interest, research on *strongly polynomial* algorithms received major attention (e.g., [49, 37, 56, 44]). We say that an algorithm runs in strongly polynomial time if its number of arithmetic operations does not depend on W. For example, the fastest strongly polynomial algorithms for maximum weight matching run in time $\tilde{\mathcal{O}}(nm)$ [18, 50]. (Here and throughout the paper we let $\tilde{\mathcal{O}}(T) = \mathcal{O}(T \operatorname{polylog} T)$, in particular, $\tilde{\mathcal{O}}(n^c)$ never hides a factor log W.)

The resulting challenge is to design improved strongly polynomial algorithms whose running times come as close as possible to the best known scaling-based algorithms, but without any log W-factors. In this paper, we tackle this challenge for a large class of approximation algorithms. This is achieved in part by an algorithmic framework that allows us to switch between approximate problems over the (min,+)-semiring and exact problems over the (min,max)-semiring.

1.1 Approximating APSP, Matrix Products, and Graph Characteristics

In this paper, we study the following problems (see Appendix A for formal problem definitions).

- Shortest path problems: The *All-Pairs Shortest Path* problem (APSP) asks to compute, given a directed graph with positive edge weights, the length of the shortest path between any two vertices.
- Matrix products: Given matrices $A, B \in \mathbb{R}^{n \times n}_+$, their product over the (\oplus, \otimes) -semiring is the matrix $C \in \mathbb{R}^{n \times n}_+$ with $C[i, j] = \bigoplus_{1 \le k \le n} (A[i, k] \otimes B[k, j])$. In general, the product can be computed using $\mathcal{O}(n^3)$ semiring operations. Over the $(+, \cdot)$ -ring, the problem is standard matrix multiplication and can be solved in time $\mathcal{O}(n^{\omega}) \le \mathcal{O}(n^{2.373})$ [25]. *Min-Plus Product* is the problem of computing the matrix product over the $(\min, +)$ -semiring.
- Graph characteristics: Specifically, we study the graph characteristics Diameter, Radius, Median, Minimum-Weight Triangle, and Minimum-Weight Cycle.

These graph characteristics and Min-Plus Product can be reduced to APSP, and thus all of these problems can be solved in time $\mathcal{O}(n^3)$ [19, 57] and using a recent algorithm by Williams [58] in time $n^3/2^{\Omega(\sqrt{\log n})}$. Moreover, with the exception of Diameter, an $\mathcal{O}(n^{3-\delta})$ -algorithm for one of these graph characteristics, or for Min-Plus Product, or for APSP would yield an $\mathcal{O}(n^{3-\delta'})$ -algorithm for all of these problems [59, 2]. It is therefore conjectured that none of them can be solved in truly subcubic time [59, 2]. Zwick designed a $(1 + \varepsilon)$ -approximation algorithm for APSP running in time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ [62]. This yields approximation schemes with the same guarantees for Min-Plus Product and the mentioned graph characteristics [2, 40]. Zwick's running time is close to optimal¹, except that it is open whether the factor log W is necessary. To the best of our knowledge, no strongly polynomial approximation scheme is known for any of the mentioned problems. This leads to our main question:

Do APSP, Min-Plus Product, and the mentioned graph characteristics have strongly polynomial approximation schemes running in time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$? Or at least in time $\widetilde{\mathcal{O}}(\frac{n^{3-\delta}}{\varepsilon})$ for some $\delta > 0$?

Note that in the setting of strongly polynomial algorithms, by *time* we mean the number of arithmetic operations. However, there is also a corresponding question that considers the bit complexity. In fact, variants of our main question are reasonable and open in at least three different settings:

- Number of arithmetic operations: When we only count arithmetic operations, then in particular we can add/multiply two log W-bit input integers in constant time. Thus, it is not clear why the running time of an algorithm should depend on log W at all. Nevertheless, Zwick's algorithm requires $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ arithmetic operations. It is open whether this can be reduced to $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$ (or even to $\widetilde{\mathcal{O}}(\frac{n^{3-\delta}}{\varepsilon})$ for any $\delta > 0$).
- Bit complexity with integers: In bit complexity, an arithmetic operation on b-bit integers has cost $\widetilde{\mathcal{O}}(b)$. Note that the input to APSP consists of n^2 many log W-bit integers, and suppose that we keep this number format throughout the algorithm. Running Zwick's algorithm in this setting results in a bit complexity of $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon}\log^2 W)$, since each arithmetic operation has bit complexity $\widetilde{\mathcal{O}}(\log(nW))$. One log W-factor is natural, since we operate on log W-bit integers. The question thus becomes whether the second log W-factor of Zwick's algorithm is necessary, or whether it can be improved to bit complexity $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon}\log W)$.
- Bit complexity with floating point approximations: One can improve upon the bit complexity of Zwick's algorithm as described above by changing the number format to floating point. Note that changing any input number by a factor in $[1, 1 + \varepsilon]$ changes the resulting distances by at most $1 + \varepsilon$ and thus still yields a $(1 + \mathcal{O}(\varepsilon))$ -approximation. We can therefore round any input integer in the range $\{1, \ldots, W\}$ to a floating point number with an $\mathcal{O}(\log \frac{1}{\varepsilon})$ -bit mantissa and an $\mathcal{O}(\log \log W)$ -bit exponent. We argue that this is the natural input format of Approximate APSP in Section 2.1. In this format, arithmetic operations on input numbers have bit complexity $\widetilde{\mathcal{O}}(\log \frac{1}{\varepsilon} + \log \log W)$, and thus a factor $\log \log W$ in the bit complexity would be natural. However, implementing Zwick's algorithm in this setting yields bit complexity $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$. The question now becomes whether this can be improved to $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log \log W)$, after converting the input numbers to floating point in time $\widetilde{\mathcal{O}}(n^2 \log W)$ (we will ignore this conversion time throughout the paper since it is near-linear in the input size).

Note that in all three settings potentially Zwick's algorithm could be improved by a factor up to $\widetilde{\mathcal{O}}(\log W)$. We focus on the first setting in this paper, where our goal is to design algorithms whose number of arithmetic operations is independent of W. However, our algorithms also yield improvements in the other two settings, which we will briefly mention below.

¹For APSP and Min-Plus Product, any $(1 + \varepsilon)$ -approximation can be used to compute the Boolean matrix product [13], and thus requires time $\Omega(n^{\omega})$. Moreover, the dependence on $\frac{1}{\varepsilon}$ should be at least polynomial, since the hardness conjecture for APSP is stated for W = poly(n) [59], and a setting of $\varepsilon = 1/W$ yields an exact algorithm.

1.2 Our Results

In this paper, we answer our main question affirmatively for all listed problems (for Directed APSP we need the relaxed form of the question). Our results hold on the Word RAM, see Section 2.1 for details of the machine model.

For the mentioned graph characteristics, obtaining time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$ is an easy exercise. Since the result is a single number, we can first compute a $\operatorname{poly}(n)$ -approximation, round edge weights to obtain $W = \operatorname{poly}(n/\varepsilon)$, and then use the $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ -time approximation scheme as a black box.

Theorem 1.1. Diameter, Radius, Median, Minimum-Weight Triangle, and Minimum-Weight Cycle on directed and undirected graphs have approximation schemes in strongly polynomial time² $\widetilde{O}(\frac{n^{\omega}}{\varepsilon})$.

For APSP restricted to *undirected* graphs, we also obtain time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$. We augment an essentially standard scaling-based algorithm for APSP by contracting light edges. This is more involved than our solution for graph characteristics, and is inspired by an iterative algorithm of Tardos [49]. Similar edge contraction arguments have been used in the context of parallel algorithms for approximate APSP on undirected graphs [30, 12].

Theorem 1.2. $(1 + \varepsilon)$ -Approximate Undirected APSP is in strongly polynomial time² $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$.

For APSP on *directed* graphs the ideas used above fail, since there are n^2 output numbers and we cannot contract directed edges. As our most involved result in this paper, we obtain a *truly subcubic* strongly polynomial approximation scheme for APSP; no such algorithm was known before.

Theorem 1.3. $(1 + \varepsilon)$ -Approximate Directed APSP is in strongly polynomial time² $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}}/\varepsilon)$.

Our approximation scheme for (directed) APSP is, in fact, a reduction from approximate APSP to the exact problem Min-Max Product, i.e., the problem of computing the matrix product over the (min, max)-semiring. This problem is closely related to the All-Pairs Bottleneck Path problem.³ Min-Max Product and All-Pairs Bottleneck Path can be solved in time $\tilde{\mathcal{O}}(n^{\frac{\omega+3}{2}})$ [15], which is why this term appears in our approximation scheme for APSP.

Furthermore, our reduction also works in the other direction, which yields an *equivalence* of approximation schemes for APSP and exact algorithms for Min-Max Product. In particular, for readers willing to believe that the best known running time for Min-Max Product is essentially optimal, this can be seen as a conditional lower bound for approximate APSP, showing that any improvements upon our approximation scheme in terms of the exponent of n is unlikely.

Theorem 1.4. For any $c \ge 2$, if one of the following statements is true, then all are:

- $(1 + \varepsilon)$ -Approximate Directed APSP can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\text{poly}(\varepsilon))$,
- $(1+\varepsilon)$ -Approximate Min-Plus Product can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\text{poly}(\varepsilon))$,
- exact Min-Max Product can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c)$,
- exact All-Pairs Bottleneck Path can be solved in strongly polynomial time $\mathcal{O}(n^c)$.

²Here, by time we mean the number of *arithmetic operations* performed on a RAM machine. The *bit complexity* of the algorithm is bounded by the number of arithmetic operations times $\log \log W$ (up to terms hidden by $\tilde{\mathcal{O}}$).

³In All-Pairs Bottleneck Path we are given a directed graph with capacities on its edges, and want to determine for all vertices u, v the capacity of a single path for which a maximum amount of flow can be routed from u to v.

Our techniques also transfer to other problems over the $(\min, +)$ -semiring. In particular, we design the first strongly polynomial-time approximation scheme for Min-Plus Convolution. Analogously to Theorem 1.4, we complement this result by an equivalence of approximating Min-Plus Convolution and exactly solving Min-Max Convolution.

Theorem 1.5. $(1+\varepsilon)$ -Approximate Min-Plus Convolution is in strongly polynomial time² $\widetilde{\mathcal{O}}(n^{3/2}/\sqrt{\varepsilon})$. Furthermore, for any $c \geq 1$, if one of the following statements is true, then both are:

- $(1+\varepsilon)$ -Approximate Min-Plus Convolution can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\operatorname{poly}(\varepsilon))$,
- exact Min-Max Product can be solved in strongly polynomial time $\mathcal{O}(n^c)$.

As an application, we obtain an approximation scheme with the same guarantees for the related Tree Sparsity problem.

Theorem 1.6. $(1 + \varepsilon)$ -Approximate Tree Sparsity is in strongly polynomial time² $\widetilde{O}(\frac{n^{3/2}}{\sqrt{\varepsilon}})$.

We prove these results related to Min-Plus Convolution in Section 8 as Theorems 8.3, 8.2, and Corollary 8.8.

Technical Overview 1.3

Our main technical contribution is the following Sum-to-Max-Covering, which yields a framework for reducing approximate problems over the $(\min, +)$ -semiring to exact or approximate problems over the (min,max)-semiring. The most intriguing results of this paper (the approximation scheme for directed APSP as well as the equivalence with Min-Max Product) are essentially immediate consequences of Sum-to-Max-Covering, see Sections 3 and 4. Here we denote $[n] = \{1, \ldots, n\}$.

Theorem 1.7 (Sum-to-Max-Covering). Given vectors $A, B \in \mathbb{R}^n_+$ and $\varepsilon > 0$, in linear time in the output size we can compute vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}((\frac{1}{\varepsilon} + \log n) \log \frac{1}{\varepsilon})$ such that for all $i, j \in [n]$:

$$A[i] + B[j] \le \min_{\ell \in [s]} \max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \le (1 + \varepsilon)(A[i] + B[j]).$$

There are two main issues that make the proof of this statement non-trivial. For close pairs i, j, meaning $\frac{A[i]}{B[j]} \in [\varepsilon, \frac{1}{\varepsilon}]$, the sum A[i] + B[j] and the maximum max $\{A[i], B[j]\}$ differ significantly. It is thus necessary to change the values of the vectors A, B. Roughly speaking, we handle this issue by splitting A into vectors $A^{(\ell)}$ such that all entries $A^{(\ell)}[i], A^{(\ell)}[i']$ differ by either less than a factor $1 + \varepsilon$ or by more than a factor $poly(1/\varepsilon)$. Then we can choose $B^{(\ell)}$ such that $B^{(\ell)}[j]$ is approximately A[i] + B[j] for all close pairs i, j. This ensures that for close pairs $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\}\$ is approximately A[i] + B[j]. For details see Close Covering (Lemma 5.2).

For the distant pairs i, j, with $\frac{A[i]}{B[j]} \notin [\varepsilon, \frac{1}{\varepsilon}]$, the sum A[i] + B[j] and the maximum max $\{A[i], B[j]\}$ differ by less than a factor $1 + \varepsilon$, so we do not have to change any values. However, we need to remove some entries (by setting them to ∞) in order to not interfere with close pairs. We show how to cover all distant pairs but no too-close pairs, via a recursive splitting into $\log n$ levels of chunks and treating boundaries between chunks by introducing several shifts of restricted areas. For details see Distant Covering (Lemma 5.3).

1.4 Further Related Work

It is known that in general not every scaling-based algorithm can be made strongly polynomial, see, e.g., Hochbaum's work on the allocation problem [28].

APSP and Min-Plus Product For undirected graphs with weights in $\{-W, \ldots, W\}$, APSP can be solved exactly in time $\tilde{\mathcal{O}}(Wn^{\omega})$ [45, 23, 5, 48], where $\omega < 2.373$ is the matrix multiplication exponent [25]. For directed graphs with weights in $\{-W, \ldots, W\}$, Zwick [62] presented an $\mathcal{O}(W^{0.68}n^{2.575})$ -time algorithm that also uses fast matrix multiplication (in fact, recent advances for rectangular matrix multiplication yield slightly stronger bounds [24, 26]).

The closest related work to our paper is by Vassilevska and Williams [54], who considered the real-valued Min-Plus Product. They proposed a method to compute the k most significant bits of each entry of the Min-Plus Product in time $\mathcal{O}(2^k n^{2.687} \log n)$, in the traditional comparisonaddition model of computation. This is similar to an additive $W/2^k$ -approximation. However, it is incomparable to a $(1 + \varepsilon)$ -approximation algorithm for Min-Plus Product, since (1) the k most significant bits might all be 0, in which case they do not provide a multiplicative approximation, and (2) a $(1 + \varepsilon)$ -approximation not necessarily allows to determine any particular bit of the result, e.g., if a number is very close to being a power of 2. Subsequently, their dependence on n was improved to $\mathcal{O}(2^k n^{2.684})$ [60], which was further refined to $\mathcal{O}(2^{0.96k} n^{2.687})$ and to $\mathcal{O}(2^{ck} n^{2.684})$ for some c < 1 [33].

For approximate APSP for real-valued graphs with weights in $[-n^{o(1)}, n^{o(1)}]$, Yuster [61] presented an additive ε -approximation in time $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}})$. More recently, among other results, Roditty and Shapira [39] gave an algorithm computing every distance $d_G(u, v)$ up to an additive error of $d_G(u, v)^p$ in time $\widetilde{\mathcal{O}}(Wn^{2.575-p/(7.4-2.3p)})$. For very small W, this interpolates between Zwick's fastest exact algorithm and his approximation algorithm [62].

In this paper we will focus on the problem of $(1 + \varepsilon)$ -approximating APSP when ε is close to 0. For $\varepsilon < 1$ the problem is at least as hard as Boolean matrix multiplication [13] and thus requires time $\Omega(n^{\omega})$. However, there are more efficient algorithms in the regime $\varepsilon \ge 1$ for undirected graphs, using spanners and distance oracles [52].

All-Pairs Bottleneck Path and Min-Max Product The All-Pairs Bottleneck Path (APBP) problem is, given an edge-weighted directed graph G, to determine for all vertices u, v the maximal weight w such that there is a path from u to v using only edges of weight at least w. It is known that APBP is equivalent to Min-Max Product, up to lower order factors in running time. The first truly subcubic algorithm for Min-Max Product was given by Vassilevska et al. [55], which was improved to time $O(n^{\frac{\omega+3}{2}})$ by Duan and Pettie [15].

Shapira et al. [47] proposed an $\mathcal{O}(n^{2.575})$ -time algorithm for a vertex-weighted variant of APBP. Duan and Ren [16] introduced the problem All-Pairs Shortest Path for All Flows (APSP-AF) and provided an approximation algorithm in time $\tilde{\mathcal{O}}(n^{\frac{\omega+3}{2}}\varepsilon^{-3/2}\log W)$. They also proved an equivalence with Min-Max Product. However, in contrast to the equivalences presented in this paper, their equivalence loses a factor log W, and thus does not work for strongly polynomial algorithms.

APSP and APBP can be easily computed in time $\mathcal{O}(n^{2.5})$ on quantum computers [35]. Le Gall and Nishimura [33] designed the first quantum algorithm for computing Min-Max Product in time $\mathcal{O}(n^{2.473})$, and noted that every problem equivalent to APBP admits a nontrivial $\mathcal{O}(n^{2.5-\varepsilon})$ -time algorithm in the quantum realm. It is also worth mentioning that there are efficient algorithms for products in other algebraic structures, e.g., dominance product, $(+, \min)$ -product, (\min, \leq) -product (see, e.g., [53]).

Hardness of Approximation in P There is a growing literature on hardness of approximation in P (see, e.g., [3, 41, 1, 10, 29, 11]), building on recent progress in fine-grained complexity theory. For readers that are willing to believe that the current algorithms for Min-Max Product are close to optimal, our equivalence of approximating APSP and exactly computing Min-Max Product is a hardness of approximation result, and in fact it is one of the first tight lower bounds for approximation algorithms for problems in P (cf. [11]).

1.5 Organization

After preliminaries on the machine model in Section 2, we present our approximation scheme for APSP in Section 3 and the equivalence with Min-Max Product in Section 4. The main technical result, Max-to-Sum-Covering, is proved in Section 5. In Section 6 we discuss Undirected APSP, and in Section 7 we discuss certain graph characteristics. Finally, in Section 8 we present our approximation scheme for Min-Plus Convolution and prove the equivalence with Min-Max Convolution. Formal problem definitions can be found in Appendix A.

2 Preliminaries

Notation By W we denote the largest input weight. We use $\tilde{\mathcal{O}}$ -notation to suppress polylogarithmic factors in n and ε , but never in W. By $\omega < 2.373$ we denote the exponent of matrix multiplication [25]. Observe that $\frac{\omega+3}{2} < 2.687$. We use [n] to denote the set $\{1, 2, \ldots, n\}$. We assume all input weights to be positive real numbers – in particular after scaling we can assume all numbers to be at least 1, so the input range is $\mathbb{R}_+ = \{a \ge 1 \mid a \in \mathbb{R}\} \cup \{\infty\}$. We denote by W the largest finite input number.

We will state our results for both directed and undirected graphs. By default, G denotes a graph, V the set of its vertices and E set of edges. In most cases the graph is weighed with a function $w : E \to \mathbb{R}_+$. When we talk about graph algorithms, n denotes the number of vertices and m the number of edges. We consider multiplicative $(1 + \varepsilon)$ -approximation algorithms, where the deviation from the exact value is always one-sided. We assume that $\varepsilon > 0$ is sufficiently small ($\varepsilon < 1/10$). For formal definitions of the problems considered in this paper, see Appendix A.

2.1 Machine Model and Input Format

Throughout the paper, we will assume that for all approximate problems input numbers are represented in floating-point, while for all exact problems input numbers are integers represented in usual bit representation. This choice of representation is not necessary for our new approximation algorithms (they would also work on the Word RAM with input in bit representation or on the Real RAM allowing only additions and comparisons); however, it is necessary for our equivalences between approximate and exact problems, as we discuss at the end of this section. We first describe the details of these formats as well as why this choice is well-motivated and natural.

The reader is invited to skip over the machine model details and consider an unrealistic, but significantly simpler model of computation throughout the paper: A Real RAM model where all logical and arithmetic operations on real numbers have unit cost, including rounding operations. This model is too powerful to be a realistic model of computation [43], but considering our algorithms in this model captures the main ideas.

Floating-Point Representation for Approximate Problems For all approximate problems considered in this paper, we can change every input weight by a factor $1 + \varepsilon$ in a preprocessing step; this changes the result by at most a factor $1 + \varepsilon$. It therefore suffices to store for each input weight w its rounded logarithm $e = \lfloor \log_2 w \rfloor$, which requires only $\mathcal{O}(\log \log W)$ bits, and a $(1 + \varepsilon)$ -approximation of $w/2^e \in [1, 2]$, which requires only $\mathcal{O}(\log 1/\varepsilon)$ bits. Note that this is floating-point representation. Hence, floating-point is the natural input format for the approximate problems studied in this paper!

The necessity for rigorous models for floating-point numbers in theoretical computer science was observed in [4, 8, 51]. Here we follow the format proposed by Thorup [50], except that we slightly simplify it, since we only want to represent positive reals. In floating-point representation, a positive real number is given as a pair x = (e, f), where the *exponent* e is a κ -bit integer and the *mantissa* f is a γ -bit string f_1, \ldots, f_{γ} . The pair x represents the real number

$$2^e \cdot \left(1 + \sum_{i=1}^{\gamma} f_i / 2^i\right).$$

Here γ, κ are parameters of the model. Moreover, we assume that all arithmetic operations on floating-point numbers can be performed in constant time.

For all approximate problems considered in this paper, we assume the input weights to be given in floating-point format. In particular, if the input weights are in the range [1, W], we assume floating-point representation with $\Theta(\log n)$ -bit mantissa and $\Theta(\log n + \log \log W)$ -bit exponent. The unit-cost assumption (that all arithmetic operations on floating-point numbers take constant time) thus hides at most a factor $\widetilde{\mathcal{O}}(\log n + \log \log W)$ compared to, e.g., the complexity of performing these operations by a device operating on bits. Note that many other formats can be efficiently converted into floating-point, and thus our algorithms also work in other settings.

Note that using a fixed floating-point precision introduces inherent inaccuracies when performing arithmetic operations. For simplicity of presentation, however, we shall assume that all arithmetic operations yield an exact result. For the algorithms in this paper, it is easy to see that this assumption can be removed by increasing the precision slightly.

Bit Representation for Exact Problems The only two exact problems that we consider in this problem are Min-Max Product and Min-Max Convolution. Since both problems are of the Min-Max type, it is easy to see that we can replace all input numbers by their ranks, i.e., their index in the sorted ordering of all input numbers. Solving the problem on the ranks, we can then infer the result. Hence, up to additional near-linear time in the input size to determine the ranks, we can assume that all input numbers are integers in the range $\{1, \ldots, \text{poly}(n)\}$, and thus all input numbers are $\mathcal{O}(\log n)$ -bit integers. This is the reason why for the exact problems studied in this paper, bit representation is the natural input format, and not floating-point! As usual for the Word RAM, we assume that each memory cell stores $\Omega(\log n)$ -bit integers, and thus operations on input numbers can be performed in constant time. Necessity of our Choice of Input Representation We crucially use our choice of input formats in our equivalences of approximate Min-Plus and exact Min-Max problems (see Theorem 1.4, Theorem 8.2): In the reduction from exact Min-Max to approximate Min-Plus we need to exponentiate some numbers. In usual bit representation, this would translate $\mathcal{O}(\log n)$ -bit integers to poly(n)-bit integers and thus not be efficient enough. However, if m is an $\mathcal{O}(\log n)$ -bit integer in standard bit-representation, then we can store 2^m in floating-point representation by storing m as the exponent; the resulting floating-point number has an $\mathcal{O}(\log n)$ -bit exponent (and an $\mathcal{O}(1)$ -bit mantissa).

For the other direction, from approximate problems in floating-point to exact problems in bit representation, we use that for Min-Max problems we can replace input numbers by their ranks, which converts floating-point numbers to $\mathcal{O}(\log n)$ -bit integers in bit representation.

3 Strongly Polynomial Approximation for Directed APSP

We present a strongly polynomial $(1 + \varepsilon)$ -approximation algorithm for APSP with running time $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}}\varepsilon^{-1})$, proving Theorem 1.3. To this end, we first recall the reduction from approximate APSP to approximate Min-Plus Product from [62] (see Theorem 3.1). Then we observe that Sum-To-Max-Covering yields a reduction from approximate Min-Plus Product to Min-Max Product. Using the known $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}})$ -time algorithm for the latter shows the result (see Theorem 3.2).

Theorem 3.1 (Implicit in [62]). If $(1 + \varepsilon)$ -Approximate Min-Plus Product can be solved in time $T(n, \varepsilon)$, then $(1 + \varepsilon)$ -Approximate APSP can be solved in time $\mathcal{O}(T(n, \varepsilon/\log n) \cdot \log n)$.

Proof. For the sake of completeness, we repeat the argument of Zwick [62, Theorem 8.1]. Let A be the adjacency matrix of a given edge-weighted directed graph G, i.e., if there is an edge $(i, j) \in E$ of weight w(i, j) then A[i, j] = w(i, j), and $A[i, j] = \infty$ otherwise. We also add self-loops of weight 0, i.e., we set A[i, i] = 0 for all $i \in [n]$. Given $\varepsilon > 0$, we set $\varepsilon' := \ln(1 + \varepsilon)/\lceil \log n \rceil$ (where ln is the natural logarithm and log is base 2). We will perform $\lceil \log n \rceil$ iterations of repeated squaring. In each iteration, we execute $(1 + \varepsilon')$ -Approximate Min-Plus Product on the current matrix A with itself, i.e., we square the current matrix A. An easy inductive proof shows that after r iterations each entry A[i, j] is bounded from below by the distance from i to j in G, and bounded from above by $(1 + \varepsilon')^r$ times the length of the shortest 2^r -hop path from i to j. Since any shortest path uses at most n edges, after $\lceil \log n \rceil$ iterations each entry A[i, j] is an approximation of the distance from i to j in G, by a multiplicative factor of

$$(1+\varepsilon')^{\lceil \log n \rceil} = \left(1 + \frac{\ln(1+\varepsilon)}{\lceil \log n \rceil}\right)^{\lceil \log n \rceil} \le 1+\varepsilon.$$

The direct running time of the reduction is $\mathcal{O}(n^2 \log n)$ and there are $\mathcal{O}(\log n)$ calls to $(1 + \varepsilon')$ -Approximate Min-Plus Product with $\varepsilon' = \Theta(\frac{\varepsilon}{\log n})$.

Theorem 3.2. $(1 + \varepsilon)$ -Approximate Min-Plus Product can be solved in time $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}}\varepsilon^{-1})$.

Proof. We use Sum-To-Max-Covering to reduce approximate Min-Plus Product to exact Min-Max Product and then use a known algorithm for the latter; the pseudocode is shown in Algorithm 1.

Consider input matrices $A, B \in \mathbb{R}^{n \times n}_+$ on which we want to compute $C \in \mathbb{R}^{n \times n}_+$ with $C[i, j] = \min_{k \in [n]} \{A[i, k] + B[k, j]\}$ for all $i, j \in [n]$. We view the matrices A, B as vectors in $\mathbb{R}^{n^2}_+$, in order to

apply Sum-To-Max-Covering (Lemma 5.1). This yields vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^{n^2}_+$, which we re-interpret as matrices in $\mathbb{R}^{n \times n}_+$. We compute the Min-Max Product of every layer $A^{(\ell)}, B^{(\ell)}$ and return the entry-wise minimum of the results, see Algorithm 1. (Note that we can replace the entries of $A^{(\ell)}, B^{(\ell)}$ by their ranks before computing the Min-Max Product and then infer the actual result — this is necessary since our input format for approximate Min-Plus Product is floating-point, but for Min-Max Product our input format is standard bit representation.)

Algorithm 1 APPROXIMATEMINPROD (A, B, ε) . 1: $\{(A^{(1)}, B^{(1)}), \dots, (A^{(s)}, B^{(s)})\}$ = SUMTOMAXCOVERING (A, B, ε) 2: $C^{(\ell)} := \text{MINMAXPROD}(A^{(\ell)}, B^{(\ell)})$ for all $\ell \in [s]$ 3: $\tilde{C}[i, j] := \min_{\ell \in [s]} C^{(\ell)}[i, j]$ for all $i, j \in [n]$ 4: return \tilde{C}

Let us prove that the output matrix \tilde{C} is a $(1 + \varepsilon)$ -approximation of C. Sum-To-Max-Covering yields that for any i, j, k we have

$$A[i,k] + B[k,j] \le \min_{\ell \in [s]} \max\{A^{(\ell)}[i,k], B^{(\ell)}[k,j]\} \le (1+\varepsilon)(A[i,k] + B[k,j]).$$

In particular, since $C[i, j] = \min_{\ell \in [s]} C^{(\ell)}[i, j] = \min_{\ell \in [s]} \min_{k \in [n]} \max\{A^{(\ell)}[i, k], B^{(\ell)}[k, j]\}$, and $C[i, j] = \min_{k \in [n]} (A[i, k] + B[k, j])$, we obtain

$$C[i,j] \le \tilde{C}[i,j] \le (1+\varepsilon)C[i,j].$$

Sum-To-Max-Covering runs in time $\widetilde{\mathcal{O}}(n^2/\varepsilon)$. Computing *s* times the Min-Max Product runs in time $\widetilde{\mathcal{O}}(sn^{\frac{\omega+3}{2}})$. We conclude the proof by noting that Sum-To-Max-Covering yields $s = \mathcal{O}(\frac{1}{\varepsilon} \operatorname{polylog}(n/\varepsilon))$.

Combining Theorems 3.1 and 3.2 yields a $(1 + \varepsilon)$ -approximation for APSP in time $\widetilde{\mathcal{O}}(n^{\frac{\omega+3}{2}}\varepsilon^{-1})$.

4 Equivalence of Approximate APSP and Min-Max Product

We next prove our equivalence of approximating APSP, exactly computing the Min-Max Product, and other problems. The theorem is restated here for convenience.

Theorem 1.4. For any $c \ge 2$, if one of the following statements is true, then all are:

- $(1 + \varepsilon)$ -Approximate Directed APSP can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\text{poly}(\varepsilon))$,
- $(1+\varepsilon)$ -Approximate Min-Plus Product can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\operatorname{poly}(\varepsilon))$,
- exact Min-Max Product can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c)$,
- exact All-Pairs Bottleneck Path can be solved in strongly polynomial time $\mathcal{O}(n^c)$.

Proof. Equivalence of $(1 + \varepsilon)$ -Approximate APSP and $(1 + \varepsilon)$ -Approximate Min-Plus Product is essentially known. One direction is given by Theorem 3.1. For the other direction, given matrices A, B we build a 3-layered graph, with edge weights between the first two layers as in A, edge weights between the last two layers as in B, and all edges directed from left to right. Then we observe that

the pairwise distances between the first and third layers are in one-to-one correspondence to Min-Plus Product on A, B, also in an approximate setting.

Equivalence of Min-Max Product and All-Pairs Bottleneck Path is folklore (see, e.g., [15]). Both directions of this equivalence work exactly as for (approximate) Min-Plus Product vs. APSP.

Our main contribution is the equivalence of $(1 + \varepsilon)$ -Approximate Min-Plus Product and exact Min-Max Product. Observe that if Min-Max Product can be solved in time T(n) then the algorithm from Theorem 3.2 runs in time $\widetilde{\mathcal{O}}(T(n)/\varepsilon)$.

It remains to show a reduction from Min-Max Product to $(1+\varepsilon)$ -Approximate Min-Plus Product. Fix any constant $\varepsilon > 0$. Given matrices A, B, denote their Min-Max Product by C. Let r be the value of $4(1+\varepsilon)^2$ rounded up to the next power of 2, and consider the matrices A', B' with $A'[i, j] := r^{A[i,j]}$ and $B'[i,j] := r^{B[i,j]}$. (Recall that the input A, B for Min-Max Product is in standard bit representation, so in constant time we can compute $r^{A[i,j]}$ in floating-point representation, by writing $A[i,j] \cdot \log r$ into the exponent.) Let C' be the result of $(1 + \varepsilon)$ -Approximate Min-Plus Product on A', B'.

Claim 4.1. We have $r^{C[i,j]} \leq C'[i,j] \leq r^{C[i,j]+1/2}$ for all i, j.

Using this claim, we can infer C from C' by computing $C[i, j] = \lfloor \log_r C'[i, j] \rfloor$ (i.e., we simply read the most significant bits of the exponent of the floating-point number C'[i, j]). If $(1 + \varepsilon)$ -Approximate Min-Plus Product can be solved in time T(n) (recall that ε is fixed), then this yields an algorithm for Min-Max Product running in time $\widetilde{\mathcal{O}}(T(n))$.

Proof of Claim 4.1. We will use $\min_k(A'[i,k] + B'[k,j]) \leq C'[i,j] \leq (1+\varepsilon) \min_k(A'[i,k] + B'[k,j])$ for all $i, j \in [n]$. For any i, j there exists k with $C[i,j] = \max\{A[i,k], B[k,j]\}$. Hence,

$$C'[i,j] \le (1+\varepsilon)(A'[i,k] + B'[k,j]) = (1+\varepsilon)(r^{A[i,k]} + r^{B[k,j]}) \le 2(1+\varepsilon)r^{\max\{A[i,k],B[k,j]\}} = 2(1+\varepsilon)r^{C[i,j]} + 2(1+\varepsilon)r^{C[i,j]} \le 2(1+\varepsilon)r^{C[i,j]} + 2(1+\varepsilon)r^{C[i,j]} \le 2(1+\varepsilon)r^{C[i,j]} + 2(1+\varepsilon)r^{C[i,j]} \le 2(1+\varepsilon)r$$

and by $r \ge 4(1+\varepsilon)^2$ we obtain $C'[i,j] \le r^{C[i,j]+1/2}$. Moreover, for any i,j there exists k with $C'[i,j] \ge A'[i,k] + B'[k,j]$. We thus obtain

$$C'[i,j] \ge A'[i,k] + B'[k,j] = r^{A[i,k]} + r^{B[k,j]} \ge r^{\max\{A[i,k],B[k,j]\}} \ge r^{C[i,j]}.$$

We remark that for scaling algorithms this proof shows an equivalence of the $\widetilde{\mathcal{O}}(Wn^{\omega})$ -time exact algorithm for Min-Max Product and the $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\operatorname{poly}(\varepsilon)}\log W)$ -time approximation scheme for Min-Plus Product.

5 Sum-To-Max-Covering

In this section, we prove the main technical result of this paper, which we slightly reformulate here.

Theorem 5.1 (Sum-to-Max-Covering, Reformulated). Given vectors $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, there are vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon} + \log n \log \frac{1}{\varepsilon})$ and:

(i) for all
$$i, j \in [n]$$
 and all $\ell \in [s]$:

$$\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge A[i] + B[j],$$

(ii) for all $i, j \in [n]$ there exists $\ell \in [s]$:

 $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \le (1+\varepsilon)(A[i]+B[j]).$

We can compute such vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)}$ in time $\mathcal{O}(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon} + n \log n \log \frac{1}{\varepsilon})$.

We split the construction into two parts, covering the pairs i, j with $\frac{A[i]}{B[j]} \in [\varepsilon, 1/\varepsilon]$ (Close Covering Lemma, Section 5.1) and covering the remaining pairs (Distant Covering Lemma, Section 5.2). We show how to combine both cases in Section 5.3.

5.1 Close Covering

We first want to cover all pairs i, j with $\frac{A[i]}{B[j]} \in [\varepsilon, 1/\varepsilon]$. To get an intuition, let $d \in \mathbb{Z}$ and consider only the entries A[i] in the range $[(1 + \varepsilon)^{d-1}, (1 + \varepsilon)^d)$. Remove all other entries of A by setting them to ∞ , obtaining a vector A'. Since we consider the close case, we are only interested in entries B[j] that differ by at most a factor $1/\varepsilon$ from A[i], so consider the entries B[j] in the range $[\varepsilon(1+\varepsilon)^{d-1}, \frac{1}{\varepsilon}(1+\varepsilon)^d)$. Add $(1+\varepsilon)^d$ to all such entries B[j] and remove all other entries of B by setting them to ∞ , obtaining a vector B'. Then for the considered entries we have $\max\{A'[i], B'[j]\} =$ $B'[j] = B[j] + (1 + \varepsilon)^d$, which is between A[i] + B[j] and $(1 + \varepsilon)(A[i] + B[j])$. This covers all considered pairs in the sense of Max-to-Sum-Covering.

However, naively we would need to repeat this construction for too many values of d. The main observation of our construction is that we can perform this construction in parallel for all values $d \in D = \{s, 2s, 3s, \ldots\}$. That is, we only remove an entry of A if it is irrelevant for all $d \in D$, and similarly for the entries of B. For a sufficiently large integer $s = \Theta(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$, it turns out that the considered entries for different d's do not interfere. Performing this construction for all shifts $D + 1, D + 2, \ldots, D + s$ covers all close pairs. See Figure 1 for an illustration.



Figure 1: An illustration for Algorithm 2. Shown is the positive real line in log-scale. Entries of A that lie outside the red/dark-shaded areas are set to ∞ . Entries of B that lie outside the blue/light-shaded areas are set to ∞ . We set $A^{(\ell)}[i] := (1 + \varepsilon)^d$ and $B^{(\ell)}[j] := B[j] + (1 + \varepsilon)^d$. This guarantees the approximation $(1 - \varepsilon)(A[i] + B[j]) \leq \max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \leq (1 + \varepsilon)(A[i] + B[j])$ for close pairs. Numbers in non-overlapping parts differ by so much that their sum and their max are equal up to a factor $1 + \varepsilon$. This ensures that they do not interfere with the close pairs.

Lemma 5.2 (Close Covering). Given vectors $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, there are vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ such that:

(i) for all
$$i, j \in [n]$$
 and all $\ell \in [s]$: $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge (1 - \varepsilon)(A[i] + B[j]), and$

(ii) for all
$$i, j \in [n]$$
 if $\frac{A[i]}{B[i]} \in [\varepsilon, 1/\varepsilon]$ then $\exists \ell \in [s]$: $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \le (1+\varepsilon)(A[i]+B[j])$.

We can compute such vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)}$ in time $\mathcal{O}(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon})$.

Proof. We choose $s = \Theta(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ with sufficiently large hidden constant, and for any $\ell \in \{1, \ldots, s\}$ construct vectors $A^{(\ell)}, B^{(\ell)}$ as described in Algorithm 2.

Algorithm 2 CLOSECOVERING $(A,$	$(B, \varepsilon).$
1: Set $s := 1 + \lceil 2 \log_{1+\varepsilon}(1/\varepsilon) \rceil$	\triangleright Note that $s = \Theta(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}).$
2: for $\ell = 1,, s$ do	\triangleright Take care of $A[i] \approx (1 + \varepsilon)^{ks+\ell}$ for any k .
3: $D_{\ell} := \{ks + \ell \mid k \in \mathbb{Z}\}$	
4: $A^{(\ell)}[i] := \begin{cases} (1+\varepsilon)^d & \text{if } A[i] \\ \infty & \text{otherw} \end{cases}$	$\in \left[(1+\varepsilon)^{d-1}, (1+\varepsilon)^d \right)$ for some $d \in D_\ell$ vise
5: $B^{(\ell)}[j] := \begin{cases} B[j] + (1+\varepsilon)^d \\ \infty \end{cases}$	if $B[j] \in \left[\varepsilon(1+\varepsilon)^{d-1}, \frac{1}{\varepsilon}(1+\varepsilon)^d\right)$ for some $d \in D_\ell$ otherwise
6: end for 7: return $\{(A^{(1)}, B^{(1)}), \dots, (A^{(s)})\}$	$B^{(s)})\}$

Note that the condition for B[j] is well-defined in the sense that it applies for at most one $d \in D_{\ell}$. To see this, since two consecutive values in D_{ℓ} differ by s, we only need to show the inequality $\frac{1}{\varepsilon}(1+\varepsilon)^d \leq \varepsilon(1+\varepsilon)^{d+s-1}$, which holds since $s \geq 1 + \log_{1+\varepsilon}(1/\varepsilon^2)$. The same can be immediately seen to hold for A[i].

The size and time bounds are immediate. It remains to prove correctness.

For property (ii), consider any i, j with $\frac{A[i]}{B[j]} \in [\varepsilon, 1/\varepsilon]$. Note that there is a unique $\ell \in \{1, \ldots, s\}$ such that $A^{(\ell)}[i] \neq \infty$. For this ℓ , we have $A^{(\ell)}[i] = (1 + \varepsilon)^d$ with $(1 + \varepsilon)^{d-1} \leq A[i] < (1 + \varepsilon)^d$, for some $d \in D_{\ell}$. By the assumption $\frac{A[i]}{B[j]} \in [\varepsilon, 1/\varepsilon]$, we obtain $\varepsilon(1 + \varepsilon)^{d-1} \leq B[j] \leq \frac{1}{\varepsilon}(1 + \varepsilon)^d$, and thus $B^{(\ell)}[j]$ is not set to ∞ , and we have $B^{(\ell)}[j] = B[j] + (1 + \varepsilon)^d$. We conclude by observing that $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} = B^{(\ell)}[j] = B[j] + (1 + \varepsilon)^d \leq (1 + \varepsilon)(B[j] + A[j])$.

For property (i), consider any i, j and ℓ . If one of $A^{(\ell)}[i], B^{(\ell)}[j]$ is set to ∞ , then the property holds trivially. Otherwise, we have $A^{(\ell)}[i] = (1 + \varepsilon)^d$ for some $d \in D_\ell$ and $B^{(\ell)}[j] = B[j] + (1 + \varepsilon)^{d'}$ for some $d' \in D_\ell$. We consider two cases.

Case 1: $d \leq d'$. Then $A[i] \leq (1+\varepsilon)^d \leq (1+\varepsilon)^{d'}$, and thus $B^{(\ell)}[j] = B[j] + (1+\varepsilon)^{d'} \geq A[i] + B[j]$. Case 2: d > d'. Then by definition of D_{ℓ} we have $d \geq d' + s$. We bound

$$B[j] \le \frac{1}{\varepsilon} (1+\varepsilon)^{d'} \le \frac{1}{\varepsilon} (1+\varepsilon)^{d-s} \le \frac{1}{\varepsilon} (1+\varepsilon)^{1-s} A[i] \le \varepsilon A[i],$$
(1)

where the last inequality uses $s \ge 1 + \log_{1+\varepsilon}(1/\varepsilon^2)$. This yields

$$A^{(\ell)}[i] \ge A[i] \stackrel{(1)}{\ge} (1-\varepsilon)A[i] + B[j] \ge (1-\varepsilon)(A[i] + B[j]).$$

In both cases we have $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge (1-\varepsilon)(A[i]+B[j])$, which proves property (i).

5.2 Distant Covering

We now want to cover all pairs i, j with $\frac{A[i]}{B[j]} \notin [\varepsilon, 1/\varepsilon]$. Our solution for this case is similar to the well-known *Well-Separated Pair Decomposition* (see [9, 6]), which we use in a one-dimensional setting and in log-scale. The main difference is that we unite sufficiently distant pairs of the decomposition that lie on the same level.

Our constructed vectors $A^{(\ell)}$ will correspond to subsets of the entries of A, i.e., we have $A^{(\ell)}[i] \in \{A[i], \infty\}$, and similarly for B. For this reason, we switch to subset notation for the majority of this section, and then return to our usual notation of vectors $A^{(\ell)}, B^{(\ell)}$ in Corollary 5.10.

For $x, y \in \mathbb{R}_+$, we define their distance as $d(x, y) := \max\{\frac{x}{y}, \frac{y}{x}\}$ if $x, y < \infty$ and $d(x, \infty) = d(\infty, x) = \infty$ otherwise. For sets $X, Y \subset \mathbb{R}_+$, we define their distance as $d(X, Y) := \min_{x \in X, y \in Y} d(x, y)$.

Lemma 5.3 (Distant Covering, Set Variant). Given a set $Z \subset \mathbb{R}_+$ of size n and a parameter $\varepsilon > 0$, there are sets $X_1, \ldots, X_s \subseteq Z$ and $Y_1, \ldots, Y_s \subseteq Z$ with $s = \mathcal{O}(\log n \log \frac{1}{\varepsilon})$ such that:

- (i) for any $\ell \in [s]$ we have $d(X_{\ell}, Y_{\ell}) > \frac{1}{\epsilon}$, and
- (ii) for any $x, y \in Z$ with $d(x, y) \geq \frac{2}{\varepsilon}$ and x < y there is $\ell \in [s]$ such that $x \in X_{\ell}$ and $y \in Y_{\ell}$.

We can compute sets X_1, \ldots, X_s and Y_1, \ldots, Y_s satisfying (1) and (2) time $\mathcal{O}(n \log n \log \frac{1}{\varepsilon})$.

We will later use Z as the set of all entries of vectors A and B. Regarding (i), observe that if $d(x,y) > \frac{1}{\varepsilon}$, then the sum x + y and the maximum $\max\{x, y\}$ differ by less than a factor $1 + \varepsilon$. This allows us to ensure point (i) of Sum-to-Max-Covering. Property (ii) ensures that we cover all distant pairs and thus corresponds to point (ii) of Sum-to-Max-Covering.

The proof outline is as follows, see also Algorithm 3 for pseudocode. To simplify notation we assume n to be a power of 2 (this is without loss of generality since we can fill up Z with arbitrary numbers). We first sort Z, so from now on we assume that $Z = \{z_1, \ldots, z_n\}$ with $z_1 \leq \ldots \leq z_n$. The algorithm performs $\log n$ iterations. In iteration r, we split Z into chunks of size $n/2^r$, and we remove some chunks that are irrelevant for covering distant pairs, see procedure SPLITCHUNKS and Figure 2. Then we separate the resulting list of chunks into two sub-lists, see procedure SEPARATECHUNKS and Figure 3. Finally, we handle the transition between any two chunks by introducing a restricted area at their boundary, applied with $\mathcal{O}(\log \frac{1}{\varepsilon})$ different shifts, see procedure SHIFTEDTRANSITIONS and Figure 4. In the following subsections we describe the individual procedures in detail.

Algorithm 3 DISTANTCOVERINGLEMMA (Z, ε)

1: sort(Z) $\triangleright Z = \{z_1 \leq z_2 \leq \ldots \leq z_n\}$ 2: Set T_0 as a list containing one element, $T_0[1] := Z$ 3: for $r = 1, 2, \ldots, \lceil \log n \rceil$ do 4: $T_r := \text{SPLITCHUNKS}(T_{r-1}, \varepsilon)$ 5: $T_{r,1}, T_{r,2} := \text{SEPARATECHUNKS}(T_r)$ 6: $S_{r,1} := \text{SHIFTEDTRANSITIONS}(T_{r,1}, \varepsilon)$ 7: $S_{r,2} := \text{SHIFTEDTRANSITIONS}(T_{r,2}, \varepsilon)$ 8: end for 9: return $\bigcup_r S_{r,1} \cup S_{r,2}$

5.2.1 SplitChunks

Algorithm 4 describes the procedure of selecting chunks T_r in every level, see also Figure 2 for an illustration. We start with a big chunk $T_0[1] = Z$, containing the whole input. Then we iterate over all levels $r = 1, 2, \ldots, \log n$ and construct refined chunks as follows. In iteration r, we iterate over all previous chunks $T_{r-1}[i]$. If $T_{r-1}[i]$ does not contain any two numbers in distance greater than

 $\frac{1}{\varepsilon}$, then we can ignore it. Otherwise, we split $T_{r-1}[i]$ at the middle into two chunks of half the size and add them to the list of chunks T_r . For any r, this yields a list of chunks T_r such that

(P1) every chunk $T_r[i]$ is a subset of Z of the form $\{z_a, z_{a+1}, \ldots, z_b\}$ and of size $|T_r[i]| = n/2^r$, and

(P2) every $x \in T_r[i]$ is smaller than every $y \in T_r[j]$, for any i < j.

Note that at the bottom level, chunks have size 1. Moreover, for any r > 0 the list T_r contains an even number of chunks; this will also hold for all lists of chunks constructed later.

Algorithm 4 SplitChunks $(T_{r-1}[1...\ell], \varepsilon)$ 1: Initialize T_r as an empty list, and k := 12: for $i = 1, 2, ..., \ell$ do By construction, $T_{r-1}[i]$ is of the form $\{z_a, z_{a+1}, \ldots, z_b\}$ for some $a \leq b$ 3: if $z_a < \varepsilon \cdot z_b$ then 4: $T_r[2k-1] := \{z_a, \ldots, z_{(a+b-1)/2}\}$ 5: \triangleright Split $T_{r-1}[i]$ in the middle 6: $T_r[2k] := \{z_{(a+b+1)/2}, \dots, z_b\}$ k := k+17: 8: end if 9: 10: end for 11: return T_r



Figure 2: Illustration of the procedure SPLITCHUNKS, which splits and selects chunks of the input numbers on different levels r.

Claim 5.4. We have $d(T_r[2k], T_r[2k+3]) > \frac{1}{\epsilon}$ for any level r and any index k.

Proof. Write the parent chunk $T_r[2k+1] \cup T_r[2k+2]$ in the form $\{z_a, z_{a+1}, \ldots, z_b\}$. Since $T_r[2k+1], T_r[2k+2]$ have been added, we have $z_a < \varepsilon \cdot z_b$. Since every number in T[2k] is smaller than z_a and every number in T[2k+3] is larger than z_b , the distance of T[2k], T[2k+3] is greater than $\frac{1}{\varepsilon}$.

The main property of our splitting procedure is that all $x, y \in Z$ with $d(x, y) > \frac{1}{\varepsilon}$ eventually are contained in consecutive chunks (see Figure 2) – note that we will only make use of consecutive chunks with indices 2k - 1, 2k for some k (as opposed to 2k, 2k + 1).

Claim 5.5. For any $x, y \in Z$, if $d(x, y) > \frac{1}{\varepsilon}$ and x < y, then there exist a level r and index k such that $x \in T_r[2k-1]$ and $y \in T_r[2k]$.

Proof. Consider the largest r such that x, y are contained in the same chunk $T_r[k']$. The chunk $T_r[k']$ contains at least two elements, so $r < \log n$. Since $d(x, y) > \frac{1}{\varepsilon}$, Algorithm 4 splits $T_r[k']$ into chunks $T_{r+1}[2k-1]$ and $T_{r+1}[2k]$. By maximality of r and by x < y, it follows that $x \in T_{r+1}[2k-1]$ and $y \in T_{r+1}[2k]$. This proves the claim.

5.2.2 SEPARATECHUNKS

The procedure SEPARATECHUNKS is given a list T_r of chunks and separates it into two subsequences $T_{r,1}$ and $T_{r,2}$, where $T_{r,1}$ contains all chunks T[i] with $(i \mod 4) \in \{1,2\}$, and $T_{r,2}$ contains the remaining chunks in T. See Algorithm 5 for pseudocode and Figure 3 for an illustration.

Algorithm 5 SEPARATECHUNKS $(T_r[1...2\ell])$ 1: Initialize $T_{r,1}, T_{r,2}$ as empty lists, and b := 12: for $k = 1, 2, ..., \ell$ do 3: Append $T_r[2k-1]$ and $T_r[2k]$ to $T_{r,b}$ 4: b := 3 - b5: end for 6: return $T_{r,1}, T_{r,2}$



Figure 3: Illustration of SEPARATECHUNKS, which separates the list of chunks T_r on some level r into two sub-lists $T_{r,1}$ and $T_{r,2}$. The selected chunks are marked in red/light-shaded and blue/dark-shaded. (In the next step, the red/light-shaded chunks will form a set X_{ℓ} , and the blue/dark-shaded ones will form a set Y_{ℓ} . Note that within $T_{r,1}$ every red chunk is ε -distant from every blue chunk, except for its right neighbor. This will be used by the procedure SHIFTEDTRANSITIONS.)

This construction ensures that consecutive chunks with indices 2k and 2k + 1 have distance at least $\frac{1}{\varepsilon}$, as shown in the following claim.

Claim 5.6. We have $d(T_{r,b}[2k], T_{r,b}[2k+1]) > \frac{1}{\epsilon}$ for any level r, index k, and $b \in \{1, 2\}$.

Proof. Because of how $T_{r,1}, T_{r,2}$ are constructed, the chunks $T_{r,b}[2k]$ and $T_{r,b}[2k+1]$ correspond to chunks $T_r[2k']$ and $T_r[2k'+3]$ for some k'. The statement now follows from Claim 5.4.

The following analogue of Claim 5.5 is immediate.

Claim 5.7. For any $x, y \in Z$, if $d(x, y) > \frac{1}{\varepsilon}$ and x < y, then there exist a level r, index k, and $b \in \{1, 2\}$ such that $x \in T_{r,b}[2k-1]$ and $y \in T_{r,b}[2k]$.

Proof. Consecutive chunks $T_r[2k-1]$ and $T_r[2k]$ are either both added to $T_{r,1}$ or both added to $T_{r,2}$. The statement thus follows from Claim 5.5.

5.2.3 ShiftedTransitions

The procedure SHIFTEDTRANSITIONS is given a list of chunks $T = T_{r,b}$ and returns $\mathcal{O}(\log \frac{1}{\varepsilon})$ many pairs (X_t, Y_t) of the final covering (recall the statement of Lemma 5.3). Naively, we would like to assign every odd chunk to X_t and every even chunk to Y_t , i.e., $X_t = \bigcup_k T[2k-1]$ and $Y_t = \bigcup_k T[2k]$. From Claim 5.6 we know that even chunks are distant from their right neighbors, i.e., $d(T[2k], T[2k+1]) > \frac{1}{\varepsilon}$. This is not necessarily true for d(T[2k-1], T[2k]), and therefore we introduce a restricted area at their boundary, applied with $\mathcal{O}(\log \frac{1}{\varepsilon})$ different shifts, as illustrated in Figure 4. See Algorithm 6.

Algorithm 6 SHIFTEDTRANSITIONS $(T[1, \dots 2\ell], \varepsilon)$

1: for $t = 0, 1, \ldots, \left\lceil \log_2 \frac{1}{\epsilon} \right\rceil$ do for $k \in \{1, \ldots, \ell\}$ do 2: let z_{\min} be the minimal number in T[2k]3: $T'[2k-1] := \{ z \in T[2k-1] \mid z \le \varepsilon 2^t z_{\min} \}$ 4: $T'[2k] := \{ z \in T[2k] \mid z > 2^t z_{\min} \}$ 5:end for 6: $X_t := \bigcup_k T'[2k-1]$ 7: $Y_t := \bigcup_k T'[2k]$ 8: 9: end for 10: return $\{(X_t, Y_t) \mid 0 \le t \le \lceil \log_2 \frac{1}{\varepsilon} \rceil\}$

Claim 5.8. For any output (X_t, Y_t) by SHIFTEDTRANSITIONS $(T_{r,b}, \varepsilon)$ we have $d(X_t, Y_t) > \frac{1}{\varepsilon}$.

Proof. Let $T = T_{r,b}$. By Claim 5.6 and sortedness (see property (P2)), any chunks T[i] and T[j] with $j \ge i+2$ have distance greater than $\frac{1}{\varepsilon}$. Since X_t only contains numbers from odd chunks T[2k-1], and Y_t only contains numbers from even chunks T[2k], we obtain that any $x \in X_t, y \in Y_t$ within distance $\frac{1}{\varepsilon}$ satisfy $x \in T[2k-1], y \in T[2k]$ for some index k. However, in any iteration t the subsets $T'[2k-1] \subseteq T[2k-1]$ and $T'[2k] \subseteq T[2k]$ are chosen to have distance greater than $\frac{1}{\varepsilon}$, and hence $d(X_t, Y_t) > \frac{1}{\varepsilon}$.

Claim 5.9. For any $x, y \in Z$, if $d(x, y) \geq \frac{2}{\varepsilon}$ and x < y, then there exist a level r and $b \in \{1, 2\}$ such that SHIFTEDTRANSITIONS $(T_{r,b}, \varepsilon)$ outputs a pair (X_t, Y_t) with $x \in X_t, y \in Y_t$.

Proof. By Claim 5.7, there are r, k, b with $x \in T_{r,b}[2k-1]$ and $y \in T_{r,b}[2k]$. Let $T = T_{r,b}$ and let z_{\min} be the minimal number in T[2k]. If $x \leq \varepsilon \cdot z_{\min}$, then in iteration t = 0 we construct T'[2k-1] containing x, and T'[2k] = T[2k] contains y, so $x \in X_0$ and $y \in Y_0$. Otherwise, let $t \in \mathbb{N}$ be minimal with $x \leq \varepsilon 2^t \cdot z_{\min}$. By sortedness (see property (P2)) we have $x < z_{\min}$ and thus $t \leq \lceil \log_2 \frac{1}{\varepsilon} \rceil$. Hence, in iteration t the set T'[2k-1] contains x. Moreover, by minimality of tand $d(x, y) \geq \frac{2}{\varepsilon}$ we have $\varepsilon 2^{t-1} \cdot z_{\min} < x \leq \frac{\varepsilon}{2}y$, and thus $y > 2^t \cdot z_{\min}$, so y is contained in T'[2k], yielding $x \in X_t, y \in Y_t$.

5.2.4 Proof of Distant Covering

Proof of Lemma 5.3. Properties (i) and (ii) follow immediately from Claims 5.8 and 5.9. The bound $s = \mathcal{O}(\log n \log \frac{1}{\varepsilon})$ on the number of constructed pairs (X_{ℓ}, Y_{ℓ}) is immediate from the loops in



Figure 4: Illustration of SHIFTEDTRANSITIONS in log-scale. Dashed areas represent the added/removed numbers from iteration t to t + 1. In each iteration, we shift to the right by a factor 2, resulting in at most $\lceil \log_2 \frac{1}{\varepsilon} \rceil$ iterations. Note that in each iteration the red/light-shaded numbers are in distance greater than $\frac{1}{\varepsilon}$ from the blue/dark-shaded numbers. Moreover, the distance between any two numbers in the dashed area is less than $\frac{2}{\varepsilon}$.

Algorithms 3 and 6. Finally, the running time of $\mathcal{O}(n \log n \log \frac{1}{\varepsilon})$ is immediate from inspecting the pseudocode.

It remains to translate Lemma 5.3 to the vector notation of Sum-to-Max-Covering.

Corollary 5.10 (Distant Covering, Vector Variant). Given vectors $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, there are vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\log n \log \frac{1}{\varepsilon})$ such that:

(i) for all
$$i, j \in [n]$$
 and all $\ell \in [s]$: $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge (1 - 2\varepsilon)(A[i] + B[j]), and$

(ii) for all
$$i, j \in [n]$$
 if $\frac{A[i]}{B[j]} \notin [\varepsilon, 1/\varepsilon]$ then $\exists \ell \in [s]$: $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \leq A[i] + B[j]$.

We can compute such vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)}$ in time $\mathcal{O}(n \log n \log \frac{1}{\epsilon})$.

Proof. Set $Z := \{A[i], B[i] \mid i \in [n]\}$ and run Lemma 5.3 on (Z, ε') with $\varepsilon' := 2\varepsilon$ to obtain subsets $X_1, \ldots, X_s \subseteq Z$ and $Y_1, \ldots, Y_s \subseteq Z$ with $s = \mathcal{O}(\log n \log \frac{1}{\varepsilon})$. We only double the number of subsets by considering $(X'_1, \ldots, X'_{2s}) := (X_1, \ldots, X_s, Y_1, \ldots, Y_s)$ and $(Y'_1, \ldots, Y'_{2s}) := (Y_1, \ldots, Y_s, X_1, \ldots, X_s)$. We construct vectors $A^{(\ell)}, B^{(\ell)}$ with $\ell \in [2s]$ by setting

$$A^{(\ell)}[i] := \begin{cases} A[i] & \text{if } A[i] \in X'_{\ell}, \\ \infty & \text{otherwise,} \end{cases} \qquad B^{(\ell)}[i] := \begin{cases} B[i] & \text{if } B[i] \in Y'_{\ell}, \\ \infty & \text{otherwise.} \end{cases}$$

The size and time bounds are immediate.

For any numbers $x, y \in \mathbb{R}_+$ with $d(x, y) > \frac{1}{\varepsilon}$ we claim that $\max\{x, y\} \ge (1 - \varepsilon)(x + y)$. Indeed, assume without loss of generality $x < \varepsilon y$, then we have $\max\{x, y\} = y > (1 - \varepsilon)(x + y) \ge (1 - \varepsilon)(x + y)$.

Property (i) of Lemma 5.3 yields that $d(A^{(\ell)}[i], B^{(\ell)}[j]) > \frac{1}{\epsilon'}$ for any i, j, ℓ . Hence, we have

$$\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge (1 - \varepsilon')(A^{(\ell)}[i] + B^{(\ell)}[j]) \ge (1 - \varepsilon')(A[i] + B[j]) = (1 - 2\varepsilon)(A[i] + B[j]),$$

proving property (i) of the corollary.

Note that by switching from X_1, \ldots, X_s and Y_1, \ldots, Y_s to X'_1, \ldots, X'_{2s} and Y'_1, \ldots, Y'_{2s} we made Lemma 5.3 symmetric, and thus the requirement x < y of its property (ii) is removed. Thus, property (ii) of Lemma 5.3 yields that for any i, j with $d(A[i], B[j]) \geq \frac{2}{\varepsilon'} = \frac{1}{\varepsilon}$ there exists $\ell \in [2s]$ with $A^{(\ell)}[i] = A[i]$ and $B^{(\ell)}[j] = B[j]$, and thus

$$\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} = \max\{A[i], B[j]\} \le A[i] + B[j], B^{(\ell)}[j]\} \le A[i] + B[j], B^{(\ell)}[j]$$

proving property (ii) of the corollary.

5.3 Proof of Sum-To-Max-Covering

The following variant of Sum-To-Max-Covering follows immediately from combining Close Covering (Lemma 5.2) and Distant Covering (Lemma 5.3). Note that compared to Lemma 5.1 there is an additional factor $1 - 2\varepsilon$ on the right hand side of property (i).

Lemma 5.11 (Weaker Sum-to-Max-Covering). Given vectors $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, there are vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon} + \log n \log \frac{1}{\varepsilon})$ such that:

(i) for all $i, j \in [n]$ and all $\ell \in [s]$:

$$\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge (1 - 2\varepsilon)(A[i] + B[j]),$$

(ii) for all $i, j \in [n]$ there exists $\ell \in [s]$:

$$\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \le (1+\varepsilon)(A[i]+B[j]).$$

We can compute such vectors $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)}$ in time $\mathcal{O}(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon} + n \log n \log \frac{1}{\varepsilon})$.

Proof of Lemma 5.11. Given vectors $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, we simply run Close Covering and Distant Covering on (A, B, ε) and concatenate the results, see Algorithm 7. Correctness as well as size and time bounds are immediate consequences of Lemma 5.2 and Corollary 5.10.

Algorithm 7 WEAKERSUMTOMAXCOVERING (A, B, ε) .					
1: return DISTANTCOVERING $(A, B, \varepsilon) \cup$ CLOSECOVERING (A, B, ε)					

Proof of Lemma 5.1. To prove the stronger variant given in the beginning of Section 5, we have to remove the factor $1 - 2\varepsilon$ on the right hand side of property (i) in Lemma 5.11. To this end, given A, B, ε , we run the construction from Lemma 5.11 on A, B, ε' with $\varepsilon' := \frac{\varepsilon}{5}$, and we divide every entry of the resulting vectors by $1 - 2\varepsilon'$, see Algorithm 8. For correctness, note that the division by $1 - 2\varepsilon'$ removes the factor $1 - 2\varepsilon'$ from property (i) in Lemma 5.11, i.e., we obtain the claimed $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \ge A[i] + B[j]$ for all i, j, ℓ . For property (ii), note that the division by $1 - 2\varepsilon'$ leaves us with $\max\{A^{(\ell)}[i], B^{(\ell)}[j]\} \le \frac{1+\varepsilon'}{1-2\varepsilon'}(A[i] + B[j])$ for all i, j and some ℓ . Using $\varepsilon' = \frac{\varepsilon}{5}$ and $\frac{1+\varepsilon/5}{1-2\varepsilon/5} \le 1+\varepsilon$ for any $\varepsilon \in (0, 1]$ finishes the proof.

Algorithm 8 SUMTOMAXCOVERING (A, B, ε) .

```
1: return \frac{1}{1-2\varepsilon/5} · WEAKERSUMTOMAXCOVERING(A, B, \frac{\varepsilon}{5})
```

6 Strongly Polynomial Approximation for Undirected APSP

In this section, we present a strongly polynomial $(1 + \varepsilon)$ -approximation for APSP on undirected graphs that runs in time $\tilde{\mathcal{O}}(n^{\omega}/\varepsilon)$. Contrary to the title of this paper our algorithm will use the scaling technique, but we combine it with edge contractions and amortized analysis to avoid the factor log W; this is inspired by Tardos [49] and uses similar edge contraction arguments as [30, 12]. We start by describing the previously fastest approximation algorithm for APSP (Section 6.1), and then present our adaptations (Section 6.2).

6.1 Zwick's Approximation for APSP

Zwick [62] obtained an $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ -time $(1 + \varepsilon)$ -approximation algorithm for (directed or undirected) APSP as follows. The first step is to reduce approximate APSP to approximate Min-Plus Product, see Theorem 3.1. It is well-known that Min-Plus Product on $n \times n$ -matrices with entries in $\{1, \ldots, W\}$ can be solved exactly in time $\widetilde{\mathcal{O}}(Wn^{\omega})$. Zwick utilized this fact via *adaptive scaling* to realize his second step, as shown in Algorithm 10.

Algorithm 9 SCALE (A, q, ε) . 1: $A'[i, j] = \begin{cases} \lceil A[i, j]/(\varepsilon 2^q) \rceil & \text{if } A[i, j] \le 2^q \\ \infty & \text{otherwise} \end{cases}$

2: return A'

Algorithm 10 ZWICK-APX-MINPROD (A, B, ε) .

1: Initialize $\tilde{C}[i, j] := \infty$ for all i, j2: Let W be the largest entry of A, B3: for $q = 0, 1, 2, \dots, \lceil \log W \rceil + 1$ do 4: $A' = \text{SCALE}(A, q, \varepsilon) \qquad \triangleright$ Scale matrix A so entries are from $\{0, \dots, \lceil \frac{1}{\varepsilon} \rceil\}$ 5: $B' = \text{SCALE}(B, q, \varepsilon)$ 6: $C' = \text{MINPLUSPROD}(A', B') \qquad \triangleright$ This works in time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$ 7: $\tilde{C}[i, j] = \min\{C[i, j], \varepsilon 2^q C'[i, j]\}$ for all i, j8: end for 9: return \tilde{C}

In each iteration q of Algorithm 10 the entries that are greater than 2^q are ignored (they are replaced by ∞). The remaining entries are scaled and rounded to lie in the range $\{1, \ldots, \lceil \frac{1}{\varepsilon} \rceil\}$ by Algorithm 9. This yields scaled matrices A', B' with integer entries bounded by $\mathcal{O}(\frac{1}{\varepsilon})$, so their Min-Plus Product C' can be computed in time $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$. The output \tilde{C} of the algorithm is the entry-wise minimum of all computed products C' over all q. The total running time is $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$.

Since we always round up, one can see that each entry of \tilde{C} is at least the corresponding entry of the correct Min-Plus Product C of A, B. For the other direction, for all i, j there is an iteration q such

that $2^{q-1} < C[i, j] \le 2^q$. Consider any k with C[i, j] = A[i, k] + B[k, j]. Then $A[i, k], B[k, j] \le 2^q$, so they are not set to ∞ . We obtain that $C'[i, j] \le A'[i, k] + B'[k, j] \le \frac{A[i, k] + B[k, j]}{\varepsilon^{2q}} + 2$, and thus $\tilde{C}[i, j] \le \varepsilon 2^q C'[i, j] \le C[i, j] + \varepsilon 2^{q+1} \le (1 + 4\varepsilon)C[i, j]$. Replacing ε by $\varepsilon/4$ yields the claimed approximation.

6.2 Undirected APSP in Strongly Polynomial Matrix-Multiplication Time

We now essentially remove the $\log W$ -factor from Zwick's algorithm for undirected graphs, proving the following restated theorem from the introduction.

Theorem 1.2. $(1 + \varepsilon)$ -Approximate Undirected APSP is in strongly polynomial time⁴ $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$.

Proof of Theorem 1.2. The algorithm proceeds in iterations $q = 1, 2, 4, \ldots$ up to the largest power of 2 bounded by nW. In iteration q, the goal is to find all shortest paths of length in [q, 2q). For this, all edges of G of weight at least 2q are irrelevant. Therefore, in the first iteration we start with an empty graph H, and in iteration q we add all edges of G with weight in [q, 2q) to H. We then round down all edge weights to multiples of $q\varepsilon/n$. This may result in edges of weight 0, which we contract (this crucially uses that we consider undirected graphs). Finally, we run Zwick's algorithm on H and update the corresponding distances. Specifically, if we compute a distance $\tilde{d}_H(u,v) \in [(1-\varepsilon)q, (1+\varepsilon)2q)$ in H, then we iterate over all vertices i in G that were contracted to u in H, and similarly over all j that were contracted to v, and we update our estimated distance D[i, j]. See Algorithm 11.

Algorithm 11 APPROXIMATEUNDIRECTEDAPSP (G, ε) . 1: Initialize H to be the graph with n isolated nodes, i.e., $H = (V(G), \emptyset)$ 2: Initialize $D[i, j] := \infty$ for all $i, j \in V(G)$ 3: for $q = 1, 2, 4, ..., 2^{\lfloor \log(nW) \rfloor}$ do \triangleright Find all shortest paths of length in [q, 2q): Add all edges of G with weight in [q, 2q) to H 4: Round down all edge weights of H to multiples of $\frac{q\varepsilon}{n}$ 5:Contract all edges of H with weight 0 6: Run Zwick's $(1 + \varepsilon)$ -approximation for APSP on H, obtaining distances $\tilde{d}_H(u, v)$ 7: for all nodes u, v of H with $d_H(u, v) \in [(1 - \varepsilon)q, (1 + \varepsilon)2q)$ do 8: $D[i,j] := \min\{D[i,j], d_H(u,v)\}$ for every node i(j) of G that was contracted to u(v), 9: end for 10: 11: end for 12: Return D

Correctness Denote by $d_G(i, j)$ the correct distance between i and j in G, and by $d_H(.,.)$ the correct distance in H. The computed approximation satisfies $d_H(u, v) \leq \tilde{d}_H(u, v) \leq (1 + \varepsilon)d_H(u, v)$ for any $u, v \in V(H)$.

Claim 6.1. Consider $i, j \in V(G)$ that have been contracted to u, v in H, respectively. If we have $\tilde{d}_H(u, v) \ge (1 - \varepsilon)q$ then $\tilde{d}_H(u, v) \ge (1 - \varepsilon)d_G(i, j)$.

⁴Here, by time we mean the number of *arithmetic operations* performed on a RAM machine. The *bit complexity* of the algorithm is bounded by the number of arithmetic operations times $\log \log W$ (up to terms hidden by $\tilde{\mathcal{O}}$).

Since we only update distances when $d_H(u, v) \in [(1 - \varepsilon)q, (1 + \varepsilon)2q)$, by this claim the output of Algorithm 11 satisfies $D[i, j] \ge (1 - \varepsilon)d_G(i, j)$ for all $i, j \in V(G)$.

Proof of Claim 6.1. Consider a path P from u to v in H realizing $d_H(u, v)$. Uncontract all contracted edges of H to obtain a graph H'. Since we only contracted edges of (rounded) weight 0, the path P corresponds to a path P' from i to j in H' of (rounded) length $d_H(u, v)$. Since we can assume P' to be a simple path and we rounded down edge weights to multiples of $q\varepsilon/n$, in the graph G path P' has length at most $d_H(u, v) + q\varepsilon$. Hence, we have $d_G(i, j) \leq d_H(u, v) + q\varepsilon$.

Note that the claim is trivial if $d_G(i,j) \leq q$, so assume $d_G(i,j) > q$. Then we conclude by bounding $\tilde{d}_H(u,v) \geq d_H(u,v) \geq d_G(i,j) - q\varepsilon > (1-\varepsilon)d_G(i,j)$.

It remains to show $D[i, j] \leq (1 + \varepsilon)d_G(i, j)$ for all $i, j \in V(G)$. Consider the iteration q with $d_G(i, j) \in [q, 2q)$. Note that all edges of any shortest path between i and j are added in or before iteration q. Let u, v be the nodes that i, j have been contracted to in H in iteration q, respectively. Since rounding down to multiples of $q\varepsilon/n$ reduces the distance by at most $q\varepsilon$, and since contracting edges of rounded weight 0 does not change any distances, we have $d_H(u, v) \in [(1 - \varepsilon)d_G(i, j), d_G(i, j)]$. This yields $d_H(u, v) \in [(1 - \varepsilon)q, 2q)$, and in particular we have $u \neq v$. By the properties of the approximation (i.e., $d_H(u, v) \leq \tilde{d}_H(u, v) \leq (1 + \varepsilon)d_H(u, v)$), we obtain $\tilde{d}_H(u, v) \in [(1 - \varepsilon)q, (1 + \varepsilon)2q)$. This triggers the update of D[i, j], which yields $D[i, j] \leq \tilde{d}_H(u, v) \leq (1 + \varepsilon)d_H(u, v) \leq (1 + \varepsilon)d_H(u, v)$.

In total, we obtain that $(1-\varepsilon)d_G(i,j) \leq D[i,j] \leq (1+\varepsilon)d_G(i,j)$ for all $i,j \in V(G)$. By dividing all output entries D[i,j] by $(1-\varepsilon)$, we may instead obtain a "one-sided" approximation of the form $d_G(i,j) \leq D'[i,j] \leq (1+\mathcal{O}(\varepsilon))d_G(i,j)$.

Running Time We call an iteration q void if G has no edge with weight in [q, 2q) and H is empty (i.e., H consists of isolated vertices). Observe that void iterations do not change H or D. In order to avoid the log W-factor of a naive implementation of Algorithm 11, we skip over all void iterations.

In iteration q, let $C_{q,1}, \ldots, C_{q,k(q)}$ be all connected components of H of size at least 2, and let $n_{q,1}, \ldots, n_{q,k(q)}$ be their sizes (i.e., number of vertices). Note that instead of running Zwick's algorithm on H, it suffices to run it on each $C_{q,i}$. Moreover, after rounding, the ratio of the largest to smallest weight in H is $\mathcal{O}(n/\varepsilon)$, and thus the log W factor of Zwick's algorithm is $\mathcal{O}(\log(n/\varepsilon))$. Hence, we can bound the contribution of Zwick's algorithm to our running time by $\widetilde{\mathcal{O}}(\sum_{q,i}(n_{q,i}^{\omega}/\varepsilon)\log(n/\varepsilon))$.

The life-cycle of every pair of vertices $u, v \in V(G)$ can be described by the following states: (1) u and v are not connected in H; (2) u and v are connected in H; (3) u and v are contracted into the same vertex of H. Observe that each pair $u, v \in V(G)$ can be in state (2) for at most $\mathcal{O}(\log(n/\varepsilon))$ iterations. Indeed, if u and v are connected in iteration q, then in iteration $q' > \frac{n}{\varepsilon}q$ they have been contracted. It follows that $\sum_{q,i} n_{q,i}^2 = \mathcal{O}(n^2 \log(n/\varepsilon))$, since the former counts the pairs of connected nodes in H, summed over all iterations q.

Combining this with our running time bound of $\widetilde{\mathcal{O}}(\sum_{q,i}(n_{q,i}^{\omega}/\varepsilon)\log(n/\varepsilon))$, the fact $\omega \geq 2$, and Jensen's inequality, yields a time bound of $\widetilde{\mathcal{O}}(n^{\omega}/\varepsilon)$. This bounds the running time spent in calls to Zwick's algorithm. Most other parts of our algorithm can be seen to take time $\widetilde{\mathcal{O}}(n^2)$ in total.

A subtle point is the update of matrix entries D[i, j] in line 9 of Algorithm 11. Consider any i, j and let u, v be the vertices that i, j have been contracted to in H in iteration q, respectively. Note that D[i, j] can only change if $u \neq v$ and u and v are connected in H. This situation can only

happen for $\mathcal{O}(\log(n/\varepsilon))$ iterations, since then u and v will be contracted to the same vertex. Hence, in total updating D takes time $\mathcal{O}(n^2 \log(n/\varepsilon))$. This finishes the proof.

7 Strongly Polynomial Approximation for Graph Characteristics

One of the fundamental challenges in network science is the identification of "important" or "central" nodes in a network. Different graph characteristics have been proposed to capture this notion [20]. For example the Median of a graph is a node that minimizes the sum of the distances to all other nodes in graph, the Center of a graph is a node that minimizes the maximum distance to any other node (this distance is called Radius) and the Diameter of a graph is the distance of the furthest pair of vertices in the graph. Centrality measures are actively generalized to weighted graphs [36]. In this section, we present a simple argument that yields strongly polynomial approximation schemes for these problems. The following theorem is restated from the introduction.

Theorem 1.1. Diameter, Radius, Median, Minimum-Weight Triangle, and Minimum-Weight Cycle on directed and undirected graphs have approximation schemes in strongly polynomial time⁵ $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon})$.

Abboud et al. [2] observed that *Diameter*, *Radius* and *Median* admit $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ -time approximation schemes via Zwick's approximation of APSP. Similarly, Roditty and Williams [40] observed that *Minimum Weight Triangle* admits an $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ -time approximation scheme. They used this as a black-box to show that *Minimum Weight Cycle* (both in directed and undirected graphs) admits an $\widetilde{\mathcal{O}}(\frac{n^{\omega}}{\varepsilon} \log W)$ -time approximation scheme.

Proof of Theorem 1.1. Let G be a given graph. For any number $w \in \mathbb{R}_+$, define G_w as the graph G where we remove all edges of weight > w and change the weight of all remaining edges to w. On G_w we can compute a 2-approximation for each of the considered problems in time $\widetilde{\mathcal{O}}(n^{\omega})$ (since $\varepsilon = 1$ is constant and there are only two different edge weights, so also W is constant). Note that if the result on G_w is infinite, then the solution value on G is greater than w, as we need to include at least one edge of weight greater than w. Moreover, if the solution value on G_w is finite, then it is at most $w \cdot n^2$, since this is the total weight of all edges in G_w . In particular, this means that the solution value of G is at most wn^2 .

We use this as follows. First, we sort all edge weights of G and perform binary search to determine the smallest edge weight w of G such that the solution value on G_w is finite. It follows that the solution value on G is in $[w, wn^2]$, so we have an $\mathcal{O}(n^2)$ -approximation. Now we round up all edge weights of G to multiples of $w\varepsilon/n^2$. This changes the total edge weight of G by at most εw , and thus also the weight of an optimal solution by at most a multiplicative factor $1 + \varepsilon$. The ratio between the largest and smallest weight in the resulting graph G' is at most $W \leq \frac{n^4}{\varepsilon}$. Hence the $\widetilde{\mathcal{O}}(\frac{n^\omega}{\varepsilon} \log W)$ -time approximation scheme runs in time $\widetilde{\mathcal{O}}(\frac{n^\omega}{\varepsilon})$ on G'.

8 Strongly Polynomial Approximation for Min-Plus Convolution

In this section, we consider sequences $A \in \mathbb{R}^n_+$ and index their entries by $A[0], \ldots, A[n-1]$. Given two sequences $A, B \in \mathbb{R}^n_+$, the convolution problem in the (\otimes, \oplus) -semiring is to compute the sequence

⁵Here, by time we mean the number of *arithmetic operations* performed on a RAM machine. The *bit complexity* of the algorithm is bounded by the number of arithmetic operations times $\log \log W$ (up to terms hidden by $\tilde{\mathcal{O}}$).

 $C \in \mathbb{R}^n_+$ with $C[k] = \bigoplus_{i+j=k} (A[i] \otimes B[j])$ for all $0 \leq k < n$. Clearly, the problem can be solved using $\mathcal{O}(n^2)$ ring operations. In the standard $(\cdot, +)$ -ring, the problem can be solved in time $\mathcal{O}(n \log n)$ by Fast Fourier Transform (FFT).

Min-Max Convolution is the problem of computing convolution in the (min,max)-semiring. Kosaraju [31] was the first to design a subquadratic-time algorithm for this problem, obtaining time $\mathcal{O}(n^{3/2}\sqrt{\log n})$. He conjectured that his algorithm can be improved to time $\tilde{\mathcal{O}}(n)$ but so far no improvement has been found.

Min-Plus Convolution is the problem of computing convolution in the (min,+)-semiring. Computing Min-Plus Convolution in time $\mathcal{O}(n^{2-\delta})$ for any $\delta > 0$ is a major open problem [14, 32]. Backurs et al. [7] were the first to study $(1+\varepsilon)$ -approximation algorithms for Min-Plus Convolution. They obtained an $\mathcal{O}(\frac{n}{\varepsilon^2} \log n \log^2 W)$ -time algorithm, which they used to design an approximation algorithm for Tree Sparsity. Subsequently, their running time was improved to $\mathcal{O}(\frac{n}{\varepsilon} \log(n/\varepsilon) \log W)$ [34].

In this section, we start with a simple $(1+\varepsilon)$ -approximation algorithm for Min-Plus Convolution that directly follows from our Sum-To-Max-Covering and runs in time $\tilde{\mathcal{O}}(n^{3/2}/\varepsilon)$, see Theorem 8.1. This is the first strongly polynomial $(1 + \varepsilon)$ -approximation algorithm for this problem (with a running time of $\mathcal{O}(n^{2-\delta})$ for any $\delta > 0$). We then prove an equivalence of approximate Min-Plus Convolution and exact Min-Max Convolution, see Theorem 8.2. Finally, we use more problemspecific arguments to obtain an improved approximation algorithm running in time $\tilde{\mathcal{O}}(n^{3/2}/\varepsilon^{1/2})$, see Theorem 8.3.

8.1 Simple Approximation Algorithm

Direct application of our Sum-to-Max-Covering yields the following strongly polynomial approximation algorithm for Min-Plus Convolution, similarly as for APSP.

Theorem 8.1. $(1+\varepsilon)$ -Approximate Min-Plus Convolution can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^{3/2}/\varepsilon)$.

Proof. Consider input sequences $(A[i])_{i=0}^{n-1}$, $(B[i])_{i=0}^{n-1}$ on which we want to compute $(C[k])_{k=0}^{n-1}$ with $C[k] = \min_{i+j=k}(A[i] + B[j])$. We view the sequences A, B as vectors in \mathbb{R}^n_+ , in order to apply Sum-To-Max-Covering (Lemma 5.1). This yields sequences $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\varepsilon^{-1} \operatorname{polylog}(n/\varepsilon))$. We compute the Min-Max Convolution of every layer $A^{(\ell)}, B^{(\ell)}$. (Note that we can first replace the entries of $A^{(\ell)}, B^{(\ell)}$ by their ranks; this is necessary since the input format for approximate Min-Plus Convolution is floating-point, but Min-Max Convolution requires standard bit representation of integers.) We then return the entry-wise minimum of the results, see Algorithm 12. The output sequence \tilde{C} is a $(1 + \varepsilon)$ -approximation of C; this follows from the properties of Sum-To-Max-Covering, analogously to the proof of Theorem 3.2. Since Sum-To-Max-Covering takes time $\tilde{\mathcal{O}}(n/\varepsilon)$, the running time is dominated by computing s times a Min-Max Convolution, resulting in $\tilde{\mathcal{O}}(sn^{3/2}) = \tilde{\mathcal{O}}(n^{3/2}/\varepsilon)$. □

Algorithm 12 APPROXIMATEMINCONV (A, B, ε)

- 1: $\{(A^{(1)}, B^{(1)}), \dots, (A^{(s)}, B^{(s)})\} = \text{SUMTOMAXCOVERING}(A, B, \varepsilon)$
- 2: $C^{(\ell)} := \text{MINMAXCONV}(A^{(\ell)}, B^{(\ell)})$ for any $\ell \in [s]$
- 3: $\tilde{C}[k] := \min_{\ell \in [s]} \{ C^{(\ell)}[k] \}$ for any $0 \le k < n$
- 4: return \tilde{C}

8.2 Equivalence of Approximate Min-Plus and Exact Min-Max Convolution

We next show an equivalence of approximate Min-Plus Convolution and exact Min-Max Convolution, similarly to the equivalence for matrix products in Theorem 1.4.

Theorem 8.2. For any $c \ge 1$, if one of the following statements is true, then both are:

- $(1+\varepsilon)$ -Approximate Min-Plus Convolution can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c/\text{poly}(\varepsilon))$,
- exact Min-Max Product can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^c)$.

In particular, any further improvement on the exponent of n in Theorem 8.1 would yield an improved algorithm for Min-Max Convolution.

Proof. For one direction, observe that if Min-Max Convolution can be solved in time T(n) then the algorithm from Theorem 8.1 runs in time $\widetilde{\mathcal{O}}(n/\varepsilon + T(n)/\varepsilon)$, which is $\widetilde{\mathcal{O}}(T(n))$ for constant $\varepsilon > 0$.

For the other direction, on input A, B denote the result of Min-Max Convolution by C. Set $r := \lceil 4(1 + \varepsilon)^2 \rceil$ and consider the sequences A', B' with $A'[i] := r^{A[i]}$ and $B'[j] := r^{B[j]}$. (Note that the integers A[i], B[j] are in standard bit representation, so we can compute floating-point representations of $r^{A[i]}, r^{B[j]}$ in constant time, essentially by writing A[i], B[j] into the exponent.) Let C' be the result of $(1 + \varepsilon)$ -Approximate Min-Plus Convolution on A', B'. Then as in the proof of Theorem 1.4, we obtain $r^{C[k]} \leq C'[k] \leq r^{C[k]+1/2}$. Hence, we can infer the Min-Max Convolution C of A, B by setting $C[k] = \lfloor \log_r C'[k] \rfloor$ (i.e., we essentially only read the exponent of the floating-point number C'[k]). If $(1 + \varepsilon)$ -Approximate Min-Plus Convolution is in time T(n), this yields an algorithm for Min-Max Convolution running in time $\widetilde{O}(n + T(n)) = \widetilde{O}(T(n))$.

8.3 Improved Approximation Algorithm

In the remainder of this section, we improve the simple approximation algorithm of Theorem 8.1.

Theorem 8.3. $(1+\varepsilon)$ -Approximate Min-Plus Convolution can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(n^{3/2}/\sqrt{\varepsilon})$.

We will divide the algorithm for Min-Plus Convolution into two parts: the first part will handle the case when summands are *close* and the second will handle the *distant* case.

8.3.1 Approximating Min-Plus Convolution for Distant Summands

First, we will simply use Distant Covering (Corollary 5.10) to correctly compute Min-Plus Convolution for summands that differ by at least a factor $\frac{1}{\epsilon}$.

Lemma 8.4. Given sequences $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, let C be the result of Min-Plus Convolution on A, B. In strongly polynomial $\mathcal{O}(n^{3/2} \operatorname{polylog}(\frac{n}{\varepsilon}))$ time we can compute a sequence \tilde{C} such that:

- (i) for any $k \in [n]$ we have $\tilde{C}[k] \geq C[k]$, and
- (ii) if there are i + j = k with C[k] = A[i] + B[j] and $\frac{A[i]}{B[j]} \notin [\frac{\varepsilon}{4}, \frac{4}{\varepsilon}]$ then $\tilde{C}[k] \le (1 + \varepsilon)C[k]$.

Algorithm 13 DISTANTMINCONV (A, B, ε)

1: $\{(A^{(1)}, B^{(1)}), \ldots, (A^{(s)}, B^{(s)})\}$ = DISTANTCOVERING $(A, B, \varepsilon/4)$ 2: $C^{(\ell)} := \operatorname{MINMAXCONV}(A^{(\ell)}, B^{(\ell)})$ for any $\ell \in [s]$ 3: $\tilde{C}[k] := \frac{1}{1-\varepsilon/2} \cdot \min_{\ell \in [s]} \{C^{(\ell)}[k]\}$ for any $0 \le k < n$ 4: return \tilde{C}

Proof of Lemma 8.4. We view the sequences A, B as vectors in \mathbb{R}^n_+ , and apply Distant Covering (Corollary 5.10) with $\varepsilon' := \frac{\varepsilon}{4}$. This yields sequences $A^{(1)}, \ldots, A^{(s)}, B^{(1)}, \ldots, B^{(s)} \in \mathbb{R}^n_+$ with $s = \mathcal{O}(\operatorname{polylog}(\frac{n}{\varepsilon}))$. We compute the Min-Max Convolution of every layer $A^{(\ell)}, B^{(\ell)}$. Then we compute the entry-wise minimum of the results, scale it by a factor $\frac{1}{1-2\varepsilon'}$, and return the resulting sequence \tilde{C} , see Algorithm 13.

For correctness, note that the scaling factor $\frac{1}{1-2\varepsilon'}$ removes the factor $1-2\varepsilon'$ from the right hand side of property (i) in Corollary 5.10. This yields $\tilde{C}[k] \geq C[k]$. Moreover, by property (ii) of Corollary 5.10, for any indices i + j = k with $\frac{A[i]}{B[j]} \notin [\varepsilon', \frac{1}{\varepsilon'}] = [\frac{\varepsilon}{4}, \frac{4}{\varepsilon}]$ there is an ℓ such that $C^{(\ell)}[k] \leq A[i] + B[j]$. Minimizing over all ℓ and multiplying by $\frac{1}{1-2\varepsilon'} = \frac{1}{1-\varepsilon/2} < 1+\varepsilon$ yields the claimed property (ii).

Since Distant Covering takes time $\mathcal{O}(n \operatorname{polylog}(\frac{n}{\varepsilon}))$, the running time is dominated by computing s times Min-Max Convolution. Using the fastest known algorithm for Min-Max Convolution [31], we obtain time $\widetilde{\mathcal{O}}(n^{3/2}s) = \mathcal{O}(n^{3/2}\operatorname{polylog}(\frac{n}{\varepsilon}))$.

8.3.2 Approximating Min-Plus Convolution for Close Summands

We now use a variant of a known scaling-based approximation scheme for Min-Plus Convolution to handle the close summands.

Lemma 8.5. Given sequences $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, let C be the result of Min-Plus Convolution on A, B. In strongly polynomial $\widetilde{\mathcal{O}}(n^{3/2}/\sqrt{\varepsilon})$ time we can compute a sequence \tilde{C} such that:

- (i) for any $k \in [n]$ we have $\tilde{C}[k] \ge C[k]$, and
- (ii) if there are i + j = k with C[k] = A[i] + B[j] and $\frac{A[i]}{B[j]} \in [\frac{\varepsilon}{4}, \frac{4}{\varepsilon}]$ then $\tilde{C}[k] \le (1 + \varepsilon)C[k]$.

Algorithm 14 $SCALE(A, q, \varepsilon)$.						
1: $A'[i] = \begin{cases} \left\lceil \frac{4}{\varepsilon q} \cdot A[i] \right\rceil \\ \infty \end{cases}$	$ if \frac{\varepsilon q}{16} \le A[i] \le q $ otherwise					
2: return A'						

Proof. Consider the procedures SCALE and CLOSEMINCONV (Algorithms 14 and 15), which are modifications of the known $\tilde{\mathcal{O}}(\frac{n}{\varepsilon} \log W)$ -time approximation scheme for Min-Plus Convolution [34]. The main difference to [34] is that the entries of A that are set to ∞ in the procedure SCALE are not only the ones that are too large (greater than q) but also the ones that are too small (smaller than $\frac{\varepsilon q}{16}$). We claim that this algorithm proves Lemma 8.5. Correctness is based on the following rounding lemma.

Algorithm 15 CLOSEMINCONV (A, B, ε) .

1: Initialize $\tilde{C}[k] := \infty$ for all k2: for $q = 1, 2, 4, \dots, 2^{\lceil \log 2W \rceil}$ do 3: $A' := \text{SCALE}(A, q, \varepsilon)$ 4: $B' := \text{SCALE}(B, q, \varepsilon)$ 5: V := EXACTMINCONV(A', B')6: $\tilde{C}[k] := \min\{\tilde{C}[k], V[k] \cdot \frac{q\varepsilon}{4}\}$ for all k7: end for 8: return \tilde{C}

Lemma 8.6 (cf. Lemma B.2 in [34]). For any $x, y, q, \varepsilon \in \mathbb{R}_+$ with $x + y \ge q/2$ and $0 < \varepsilon < 1$ we have:

$$x + y \le \left(\left\lceil \frac{4x}{q\varepsilon} \right\rceil + \left\lceil \frac{4y}{q\varepsilon} \right\rceil \right) \frac{q\varepsilon}{4} \le (1 + \varepsilon)(x + y).$$

Proof. We repeat the proof for completeness. The lower bound is immediate. For the upper bound, note that

$$\left(\left\lceil\frac{4x}{q\varepsilon}\right\rceil + \left\lceil\frac{4y}{q\varepsilon}\right\rceil\right)\frac{q\varepsilon}{4} \le x + y + 2\frac{q\varepsilon}{4} = x + y + \varepsilon \cdot \frac{q}{2} \le (1+\varepsilon)(x+y).$$

Correctness Correctness of CLOSEMINCONV can now be shown similarly as in [34]. Regarding property (i), the lower bound of Lemma 8.6 ensures that we have $\tilde{C}[k] \ge C[k]$ throughout the run of the algorithm. Regarding property (ii), for any i + j = k with C[k] = A[i] + B[j] there exists a precision parameter q with $q/2 \le A[i] + B[j] \le q$. In particular, we have $A[i], B[j] \le q$ and $\max\{A[i], B[j]\} \ge \frac{q}{4}$. If we additionally have $\frac{A[i]}{B[j]} \in [\frac{\varepsilon}{4}, \frac{4}{\varepsilon}]$, then

$$\min\{A[i], B[j]\} \ge \frac{\varepsilon}{4} \cdot \max\{A[i], B[j]\} \ge \frac{\varepsilon q}{16},$$

and thus A'[i] and B'[j] both are not set to ∞ by the procedure SCALE. The upper bound of Lemma 8.6 now implies $\tilde{C}[k] \leq (1 + \varepsilon)(A[i] + B[j]) = (1 + \varepsilon)C[k]$.

Running time For the running time analysis, denote by α_q the number of entries of A' that are not set to ∞ in iteration q. Note that if an entry A[i] is not set to ∞ in iteration q, then we have $\frac{\varepsilon q}{16} \leq A[i] \leq q$, or, equivalently, $A[i] \leq q \leq \frac{16}{\varepsilon}A[i]$. Since q grows geometrically, there are $\mathcal{O}(\log \frac{1}{\varepsilon})$ iterations q in which entry A[i] is not set to ∞ . Hence, we obtain $\sum_q \alpha_q = \mathcal{O}(n \log \frac{1}{\varepsilon})$. We similarly define β_q as the number of non- ∞ entries of B' in iteration q, and obtain

$$\sum_{q} \beta_q = \mathcal{O}\left(n \log \frac{1}{\varepsilon}\right). \tag{2}$$

We argue in the following that procedure CLOSEMINCONV can be implemented in such a way that the running time for iteration q is $\mathcal{O}(\min\{\alpha_q\beta_q, \frac{n}{\varepsilon}\}\operatorname{polylog}(\frac{n}{\varepsilon}))$.

To this end, we use an event queue to be able to skip all iterations with $\alpha_q = 0$ or $\beta_q = 0$. Moreover, we can maintain all non- ∞ entries of A', B' in time $\mathcal{O}(\alpha_q + \beta_q)$ per iteration q. Note that $\mathcal{O}(\alpha_q + \beta_q) \leq \mathcal{O}(\min\{\alpha_q\beta_q, \frac{n}{\varepsilon}\})$. This yields the claimed time bound for lines 3 and 4 of procedure CLOSEMINCONV.

For line 5 of procedure CLOSEMINCONV, note that naively the exact Min-Plus Convolution of A'and B' can be computed in time $\mathcal{O}(\alpha_q \beta_q)$. Moreover, since A', B' have entries in $\{1, \ldots, W\} \cup \{\infty\}$ for $W = \left\lceil \frac{4}{\varepsilon} \right\rceil$, by Fast Fourier Transform their Min-Plus Convolution can be computed in time $\widetilde{\mathcal{O}}(Wn) = \widetilde{\mathcal{O}}(\frac{n}{\varepsilon})$. Using the better of the two yields the claimed time bound.

Finally, line 6 of procedure CLOSEMINCONV can be implemented in time $\mathcal{O}(\min\{\alpha_q\beta_q,n\})$, since this is an upper bound on the number of non- ∞ entries of V.

Altogether, we obtain time $\mathcal{O}(\min\{\alpha_q\beta_q, \frac{n}{\varepsilon}\} \operatorname{polylog}(\frac{n}{\varepsilon}))$ per iteration q. Hence, the total running time of procedure CLOSEMINCONV is bounded by

$$\sum_{q} \min\left\{\alpha_{q}\beta_{q}, \frac{n}{\varepsilon}\right\} \operatorname{polylog}(\frac{n}{\varepsilon}).$$

We split this sum into the cases $\beta_q \leq \lambda$ and $\beta_q > \lambda$, and note that the second case can occur at most $\mathcal{O}(\frac{n}{\lambda}\log\frac{1}{\varepsilon})$ times due to (2). This yields a total running time of at most

$$\left(\left(\sum_{q} \alpha_{q} \lambda\right) + \frac{n}{\varepsilon} \cdot \frac{n}{\lambda} \log \frac{1}{\varepsilon}\right) \operatorname{polylog}(\frac{n}{\varepsilon}) \le \left(n\lambda + \frac{n^{2}}{\varepsilon\lambda}\right) \operatorname{polylog}(\frac{n}{\varepsilon}).$$

Setting $\lambda := (n/\varepsilon)^{1/2}$ yields the claimed total running time bound $\widetilde{\mathcal{O}}(n^{3/2}/\varepsilon^{1/2})$.

8.3.3 Proof of Theorem 8.3

\mathbf{A}	lgorithm	16	Apx	V	lin	C	ONV	V (A	, I	З	$,\varepsilon)$).
--------------	----------	----	-----	---	-----	---	-----	-----	---	-----	---	-----------------	----

1: $\tilde{C}_1 := \text{DISTANTMINCONV}(A, B, \varepsilon)$ 2: $\tilde{C}_2 := \text{CLOSEMINCONV}(A, B, \varepsilon)$ 3: $\tilde{C}[k] := \min{\{\tilde{C}_1[k], \tilde{C}_2[k]\}}$ for any $0 \le k < n$ 4: return \tilde{C}

Proof of Theorem 8.3. Given sequences $A, B \in \mathbb{R}^n_+$ and a parameter $\varepsilon > 0$, we simply run our algorithms for approximating Min-Plus Convolution on distant and close summands, and compute their entry-wise minimum (see Algorithm 16). Correctness as well as size and time bounds are immediate consequences of Lemmas 8.4 and 8.5.

8.4 Applications for Tree Sparsity

A direct consequence of the strongly polynomial $(1 + \varepsilon)$ -approximation for Min-Plus Convolution is an analogous strongly polynomial $(1 + \varepsilon)$ -approximation for Tree Sparsity. We will make use of the following strongly polynomial reduction that uses an approximation algorithm for Min-Plus Convolution as a black-box.

Theorem 8.7 (cf. Theorem 7.1 [34], reformulation of [7]). If $(1 + \varepsilon)$ -Approximate Min-Plus Convolution can be solved in time $T(n,\varepsilon)$, then $(1 + \varepsilon)$ -Approximate Tree Sparsity can be solved in time $\mathcal{O}((n + T(n,\varepsilon/\log^2 n))\log n)$.

Note, that this reduction runs in strongly polynomial time, and the log W-factors in the running time of [7, 34] come exclusively from the use of scaling for approximating Min-Plus Convolution. In consequence, our strongly polynomial $(1 + \varepsilon)$ -approximation for Min-Plus Convolution yields a strongly polynomial $(1 + \varepsilon)$ -approximation for Tree Sparsity.

Corollary 8.8. $(1+\varepsilon)$ -Approximate Tree Sparsity can be solved in strongly polynomial time $\widetilde{\mathcal{O}}(\frac{n^{3/2}}{\sqrt{\varepsilon}})$.

Tree Sparsity has applications in image processing, computational biology [46] and machine learning [7]. The main issue of previous approximation schemes for this problem was that in applications the input consists of real numbers and thus $\log W$ -factors significantly influence the running time [42].

References

- A. Abboud and A. Rubinstein. Fast and deterministic constant factor approximation algorithms for LCS imply new circuit lower bounds. In Proc. 9th Innovations in Theoretical Computer Science Conference (ITCS'18), volume 94, pages 35:1–35:14, 2018.
- [2] A. Abboud, F. Grandoni, and V. V. Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In Proc. 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15), pages 1681–1697, 2015.
- [3] A. Abboud, A. Rubinstein, and R. R. Williams. Distributed PCP theorems for hardness of approximation in P. In Proc. 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS'17), pages 25–36, 2017.
- [4] E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. SIAM J. Comput., 38(5):1987–2006, 2009.
- [5] N. Alon, Z. Galil, and O. Margalit. On the exponent of the all pairs shortest path problem. J. Comput. Syst. Sci., 54(2):255-262, 1997.
- [6] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In Proc. 27th Annual ACM Symposium on Theory of Computing (STOC'95), pages 489–498, 1995.
- [7] A. Backurs, P. Indyk, and L. Schmidt. Better approximations for tree sparsity in nearly-linear time. In Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'17), pages 2215–2229, 2017.
- [8] L. Blum, F. Cucker, M. Shub, and S. Smale. Complexity and real computation. Springer Science & Business Media, 2012.
- [9] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. J. ACM, 42(1):67–90, 1995.
- [10] P. Chalermsook, M. Cygan, G. Kortsarz, B. Laekhanukit, P. Manurangsi, D. Nanongkai, and L. Trevisan. From Gap-ETH to FPT-inapproximability: Clique, dominating set, and more. In Proc. 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS'17), pages 743–754, 2017.
- [11] L. Chen. On the hardness of approximate and exact (bichromatic) maximum inner product. In Proc. 33rd Computational Complexity Conference (CCC'18), volume 102, pages 14:1–14:45, 2018.

- [12] E. Cohen. Using selective path-doubling for parallel shortest-path computations. J. Algorithms, 22(1):30–56, 1997.
- [13] E. Cohen and U. Zwick. All-pairs small-stretch paths. J. Algorithms, 38(2):335–353, 2001.
- [14] M. Cygan, M. Mucha, K. Wegrzycki, and M. Włodarczyk. On problems equivalent to (min, +)-convolution. ACM Trans. Algorithms, 15(1):14:1–14:25, 2019.
- [15] R. Duan and S. Pettie. Fast algorithms for (max,min)-matrix multiplication and bottleneck shortest paths. In Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09), pages 384–391, 2009.
- [16] R. Duan and H. Ren. Approximating all-pair bounded-leg shortest path and APSP-AF in truly-subcubic time. In Proc. 45th International Colloquium on Automata, Languages, and Programming (ICALP'18), pages 42:1–42:12, 2018.
- [17] R. Duan, S. Pettie, and H. Su. Scaling algorithms for weighted matching in general graphs. In Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'17), pages 781–800, 2017.
- [18] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. J. ACM, 19(2):248–264, 1972.
- [19] R. W. Floyd. Algorithm 97: Shortest path. Commun. ACM, 5(6):345, 1962.
- [20] L. C. Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3): 215–239, 1978.
- [21] H. N. Gabow and R. E. Tarjan. Faster scaling algorithms for network problems. SIAM J. Comput., 18(5):1013–1036, 1989.
- [22] H. N. Gabow and R. E. Tarjan. Faster scaling algorithms for general graph-matching problems. J. ACM, 38(4):815–853, 1991.
- [23] Z. Galil and O. Margalit. All pairs shortest paths for graphs with small integer length edges. J. Comput. Syst. Sci., 54(2):243–254, 1997.
- [24] F. L. Gall. Faster algorithms for rectangular matrix multiplication. In Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science, (FOCS 2012), pages 514–523, 2012.
- [25] F. L. Gall. Powers of tensors and fast matrix multiplication. In Proc. 39th International Symposium on Symbolic and Algebraic Computation (ISSAC'14), pages 296–303, 2014.
- [26] F. L. Gall and F. Urrutia. Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18), pages 1029–1046, 2018.
- [27] A. V. Goldberg. Scaling algorithms for the shortest paths problem. SIAM J. Comput., 24(3): 494–504, 1995.
- [28] D. S. Hochbaum. Lower and upper bounds for the allocation problem and other nonlinear optimization problems. *Math. Oper. Res.*, 19(2):390–409, 1994.

- [29] Karthik C. S., B. Laekhanukit, and P. Manurangsi. On the parameterized complexity of approximating dominating set. In Proc. 50th Annual Symposium on Theory of Computing (STOC'18), pages 1283–1296, 2018.
- [30] P. N. Klein and S. Sairam. A parallel randomized approximation scheme for shortest paths. In Proc. 24th Annual ACM Symposium on Theory of Computing (STOC'92), pages 750–758, 1992.
- [31] S. R. Kosaraju. Efficient tree pattern matching (preliminary version). In Proc. 30th Annual Symposium on Foundations of Computer Science (FOCS'89), pages 178–183, 1989.
- [32] M. Künnemann, R. Paturi, and S. Schneider. On the fine-grained complexity of one-dimensional dynamic programming. In Proc. 44th International Colloquium on Automata, Languages, and Programming (ICALP'17), pages 21:1–21:15, 2017.
- [33] F. Le Gall and H. Nishimura. Quantum algorithms for matrix products over semirings. *Chicago J. Theor. Comput. Sci.*, 2017, 2017.
- [34] M. Mucha, K. Wegrzycki, and M. Włodarczyk. A subquadratic approximation scheme for partition. In Proc. 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19), pages 70–88, 2019.
- [35] A. Nayebi and V. V. Williams. Quantum algorithms for shortest paths problems in structured instances. CoRR, 2014. URL http://arxiv.org/abs/1410.6220.
- [36] T. Opsahl, F. Agneessens, and J. Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251, 2010.
- [37] J. B. Orlin. A faster strongly polynomial minimum cost flow algorithm. Operations Research, 41(2):338–350, 1993.
- [38] J. B. Orlin and R. K. Ahuja. New scaling algorithms for the assignment and minimum mean cycle problems. *Math. Program.*, 54:41–56, 1992.
- [39] L. Roditty and A. Shapira. All-pairs shortest paths with a sublinear additive error. In Proc. 35th International Colloquium on Automata, Languages and Programming (ICALP'08), pages 622–633, 2008.
- [40] L. Roditty and V. V. Williams. Minimum weight cycles and triangles: Equivalences and algorithms. In Proc. 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'11), pages 180–189, 2011.
- [41] A. Rubinstein. Hardness of approximate nearest neighbor search. In Proc. 50th Annual Symposium on Theory of Computing (STOC'18), pages 1260–1268, 2018.
- [42] L. Schmidt. personal communication, 2017.
- [43] A. Schönhage. On the power of random access machines. In Proc. 6th International Colloquium on Automata, Languages, and Programming (ICALP'79), volume 71, pages 520–529, 1979.

- [44] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. J. Comb. Theory, Ser. B, 80(2):346–355, 2000.
- [45] R. Seidel. On the all-pairs-shortest-path problem. In Proc. 24th Annual ACM Symposium on Theory of Computing (STOC'92), pages 745–749, 1992.
- [46] O. Serang. A fast numerical method for max-convolution and the application to efficient maxproduct inference in Bayesian networks. *Journal of Computational Biology*, 22(8):770–783, 2015.
- [47] A. Shapira, R. Yuster, and U. Zwick. All-pairs bottleneck paths in vertex weighted graphs. In Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'07), pages 978–985, 2007.
- [48] A. Shoshan and U. Zwick. All pairs shortest paths in undirected graphs with integer weights. In Proc. 40th Annual IEE Symposium on Foundations of Computer Science (FOCS'99), pages 605–615, 1999.
- [49] E. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3): 247–256, 1985.
- [50] M. Thorup. Floats, integers, and single source shortest paths. J. Algorithms, 35(2):189–201, 2000.
- [51] M. Thorup. Equivalence between priority queues and sorting. In Proc. 43rd Symposium on Foundations of Computer Science (FOCS'02), pages 125–134, 2002.
- [52] M. Thorup and U. Zwick. Approximate distance oracles. J. ACM, 52(1):1–24, 2005.
- [53] V. Vassilevska. Efficient algorithms for path problems in weighted graphs. Carnegie Mellon University, 2008. PhD Thesis.
- [54] V. Vassilevska and R. Williams. Finding a maximum weight triangle in $O(n^{3-\delta})$ time, with applications. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC'06)*, pages 225–231, 2006.
- [55] V. Vassilevska, R. Williams, and R. Yuster. All pairs bottleneck paths and max-min matrix products in truly subcubic time. *Theory of Computing*, 5(1):173–189, 2009.
- [56] L. A. Végh. A strongly polynomial algorithm for generalized flow maximization. Math. Oper. Res., 42(1):179–211, 2017.
- [57] S. Warshall. A theorem on Boolean matrices. J. ACM, 9(1):11–12, 1962.
- [58] R. R. Williams. Faster all-pairs shortest paths via circuit complexity. SIAM J. Comput., 47 (5):1965–1985, 2018.
- [59] V. V. Williams and R. R. Williams. Subcubic equivalences between path, matrix, and triangle problems. J. ACM, 65(5):27:1–27:38, 2018.

- [60] R. Yuster. Efficient algorithms on sets of permutations, dominance, and real-weighted APSP. In Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09), pages 950–957, 2009.
- [61] R. Yuster. Approximate shortest paths in weighted graphs. J. Comput. Syst. Sci., 78(2): 632–637, 2012.
- [62] U. Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. J. ACM, 49(3):289–317, 2002.

A Problem Definitions

For an edge-weighted (directed or undirected) graph G, $d_G(u, v)$ denotes the minimal total weight of any path from u to v in G.

Directed/Undirected All-Pairs Shortest Path (APSP) **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute $d_G(u, v)$ for any $u, v \in [n]$

 $(1 + \varepsilon)$ -Approximate Directed/Undirected All-Pairs Shortest Path (APSP) **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute values $\tilde{d}(u, v)$ for any $u, v \in [n]$ such that $d_G(u, v) \leq \tilde{d}(u, v) \leq (1 + \varepsilon)d_G(u, v)$

All-Pairs Bottleneck Path

Input: An edge-weighed directed graph G on n nodes, with weights in \mathbb{R}_+ **Task:** For every two nodes u, v determine the maximum over all paths from u to v of the minimum edge weight along the path

Min-Plus Product **Input:** Matrices $A, B \in \mathbb{R}^{n \times n}_+$ **Task:** Compute matrix $C \in \mathbb{R}^{n \times n}_+$ with $C[i, j] = \min_{k \in [n]} (A[i, k] + B[k, j])$ for any $i, j \in [n]$

Min-Max Product **Input:** Matrices $A, B \in \mathbb{R}^{n \times n}_+$ **Task:** Compute matrix $C \in \mathbb{R}^{n \times n}_+$ with $C[i, j] = \min_{k \in [n]} \max\{A[i, k], B[k, j]\}$ for any $i, j \in [n]$

 $(1 + \varepsilon)$ -Approximate Min-Plus Product **Input:** Matrices $A, B \in \mathbb{R}^{n \times n}_+$ **Task:** Compute a matrix $\tilde{C} \in \mathbb{R}^{n \times n}_+$ with $C[i, j] \leq \tilde{C}[i, j] \leq (1 + \varepsilon)C[i, j]$ for any $i, j \in [n]$, where C is the correct output of Min-Plus Product Min-Plus Convolution **Input:** Sequences $(A[i])_{i=0}^{n-1}, (B[i])_{i=0}^{n-1} \in \mathbb{R}^n_+$ **Task:** Compute sequence $(C[i])_{i=0}^{n-1}$ with $C[k] = \min_{i+j=k}(A[i] + B[j])$

Min-Max Convolution **Input:** Sequences $(A[i])_{i=0}^{n-1}, (B[i])_{i=0}^{n-1} \in \mathbb{R}^n_+$ **Task:** Compute sequence $(C[i])_{i=0}^{n-1}$ with $C[k] = \min_{i+j=k} \max\{A[i], B[j]\}$

 $(1 + \varepsilon)$ -Approximate Min-Plus Convolution **Input:** Sequences $(A[i])_{i=0}^{n-1}$, $(B[i])_{i=0}^{n-1} \in \mathbb{R}^n_+$ **Task:** Compute a sequence $(\tilde{C}[i])_{i=0}^{n-1}$ with $C[k] \leq \tilde{C}[k] \leq (1+\varepsilon)C[k]$ for any $0 \leq k < n$, where C denotes the correct output of Min-Plus Convolution

Tree Sparsity

Input: A node-weighted rooted tree T, with weights in \mathbb{R}_+ , and an integer k**Task:** Find the maximal total weight of any rooted subtree of T consisting of k vertices

Directed/Undirected Minimum Weight Triangle Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ Task: Compute the minimal total weight of any triangle in G

 $(1 + \varepsilon)$ -Approximate Directed/Undirected Minimum Weight Triangle **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute a number $\tilde{T} \in [T, (1+\varepsilon)T]$, where T is the minimal total weight of any triangle in G

Directed/Undirected Minimum Weight Cycle **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute the minimal total weight of any cycle in G

 $(1 + \varepsilon)$ -Approximate Directed/Undirected Minimum Weight Cycle **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute a number $\tilde{C} \in [C, (1 + \varepsilon)C]$, where C is the minimal total weight of any cycle in G

Radius

Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute $\min_{s \in V(G)} \max_{v \in V(G)} d_G(s, v)$

 $(1 + \varepsilon)$ -Approximate Radius

Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute a number $\tilde{R} \in [R, (1 + \varepsilon)R]$, where R is the radius of G Diameter

Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute $\max_{u,v \in V(G)} d_G(u, v)$

 $(1 + \varepsilon)$ -Approximate Diameter **Input:** An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute a number $\tilde{D} \in [D, (1 + \varepsilon)D]$, where D is the diameter of G

Median

Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute $\min_{u \in V(G)} \sum_{v \in V(G)} d_G(u, v)$

 $(1 + \varepsilon)$ -Approximate Median

Input: An edge-weighed directed/undirected graph G on n nodes, with weights in \mathbb{R}_+ **Task:** Compute a number $\tilde{M} \in [M, (1 + \varepsilon)M]$, where M is the median of G