




Bellman–Ford Is Optimal for Shortest Hop-Bounded Paths

Tomasz Kociumaka   

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Adam Polak   

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Abstract

This paper is about the problem of finding a shortest s - t path using at most h edges in edge-weighted graphs. The Bellman–Ford algorithm solves this problem in $O(hm)$ time, where m is the number of edges. We show that this running time is optimal, up to subpolynomial factors, under popular fine-grained complexity assumptions.

More specifically, we show that under the APSP Hypothesis the problem cannot be solved faster already in undirected graphs with nonnegative edge weights. This lower bound holds even restricted to graphs of arbitrary density and for arbitrary $h \in O(\sqrt{m})$. Moreover, under a stronger assumption, namely the Min-Plus Convolution Hypothesis, we can eliminate the restriction $h \in O(\sqrt{m})$. In other words, the $O(hm)$ bound is tight for the entire space of parameters h , m , and n , where n is the number of nodes.

Our lower bounds can be contrasted with the recent near-linear time algorithm for the negative-weight Single-Source Shortest Paths problem, which is the textbook application of the Bellman–Ford algorithm.

2012 ACM Subject Classification Theory of computation → Shortest paths

Keywords and phrases Fine-grained complexity, graph algorithms, lower bounds, shortest paths

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.72

Funding *Adam Polak*: Part of this work was done at École Polytechnique Fédérale de Lausanne, supported by the Swiss National Science Foundation projects *Lattice Algorithms and Integer Programming* (185030) and *Complexity of Integer Programming* (CRFS-2_207365).

Acknowledgements The second author would like to thank Danupon Nanongkai and Luca Trevisan for bringing his attention to the problem discussed in this paper, Alexandra Lassota – for useful feedback on an early draft of the manuscript, and Imbir – a ginger tabby, who supervised initial stages of this work.

1 Introduction

The Bellman–Ford algorithm [24, 14, 4] is the textbook solution for the Single-Source Shortest Paths (SSSP) problem in graphs with negative edge weights. It runs in $O(nm)$ time, where n denotes the number of nodes and m is the number of edges. If we limit the outer for-loop (see Algorithm 1) to only $h \leq n - 1$ iterations, the algorithm computes single-source shortest paths that use at most h edges (or *hops*) and runs in $O(hm)$ time.

Algorithm 1 The Bellman–Ford algorithm.

```
 $d^{(0)} \leftarrow [+\infty, +\infty, \dots, +\infty];$   
 $d^{(0)}[s] \leftarrow 0;$   
for  $i$  from 1 to  $n - 1$  do  
   $d^{(i)} \leftarrow d^{(i-1)};$   
  foreach edge  $(u, v) \in E$  do  
     $d^{(i)}[v] \leftarrow \min\{d^{(i)}[v], d^{(i-1)}[u] + w(u, v)\};$ 
```



© Tomasz Kociumaka and Adam Polak;
licensed under Creative Commons License CC-BY 4.0
31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 72;
pp. 72:1–72:10



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Negative-weight SSSP has seen a lot of improvements over Bellman–Ford’s running time: scaling algorithms [15, 16, 18], which eventually led to $O(n^{1/2}m \log W)$ running time, where W denotes the maximum absolute value of a negative edge weight; interior-point methods for the more general Minimum-Cost Flow problem, which recently led to an almost-linear $O(m^{1+o(1)} \log W)$ time algorithm [8]; and finally, the recent combinatorial near-linear $O(m \log^8(n) \log W)$ time algorithm [5], subsequently improved to run in $O(m \log^2(n) \log(nW) \log \log n)$ time [7].

Can we get similar improvements for the problem of finding shortest hop-bounded paths?

This basic question stays embarrassingly open. Even in undirected graphs with only nonnegative edge weights, Bellman–Ford remains the fastest known algorithm for that problem. In this paper, we give a negative answer to the above question, up to subpolynomial factors, under popular fine-grained complexity assumptions.

1.1 Our results

Let us begin with formally stating the computational problem that we study: Given a graph $G = (V, E)$ with edge weights $w : E \rightarrow \mathbb{Z}$, two distinguished nodes $s, t \in V$, and a nonnegative integer $h \in \mathbb{Z}_{\geq 0}$, find (the length of) a shortest path from s to t that uses at most¹ h edges. We call such paths *h-hop-bounded*, or simply *hop-bounded* when h is implicit in the context.

In this paper, we give two fine-grained reductions, each proving that the $O(hm)$ running time of the Bellman–Ford algorithm is conditionally optimal for the problem, up to subpolynomial factors. Our two hardness results differ from each other in (1) how the parameters n, m, h of the hard instances relate to each other, and (2) which hardness assumption is required. Table 1 summarizes these differences.

Our first result holds under the APSP Hypothesis.

► **Theorem 1.** *Unless the APSP Hypothesis fails, there is neither an $O(h^{1-\varepsilon}m)$ nor an $O(hm^{1-\varepsilon})$ time algorithm for finding the length of a shortest h -hop-bounded s - t path in undirected graphs with nonnegative edge weights, for any constant $\varepsilon > 0$.*

This holds even restricted to instances with density $n = \Theta(\sqrt{m})$ and hop bound $h = \Theta(m^\eta)$ for arbitrarily chosen $\eta \in (0, 1/2]$.

Although the hard instances in Theorem 1 are dense, one can trivially obtain sparser instances by adding isolated nodes. Indeed, such nodes influence neither the length of a shortest h -hop-bounded s - t path nor the running time bounds as functions of h and m .

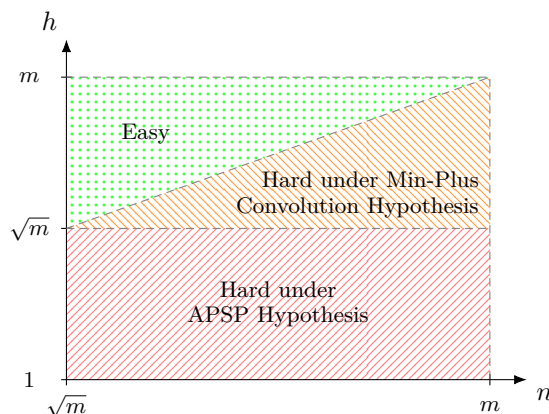
► **Corollary 2.** *The result of Theorem 1 holds even restricted to instances with density $n = \Theta(m^\nu)$ and hop bound $h = \Theta(m^\eta)$ for arbitrarily chosen $\nu \in [1/2, 1]$ and $\eta \in (0, 1/2]$.*

We remark that it is not very surprising that our reduction from APSP can only produce instances with $h \in O(\sqrt{m})$. The conjectured time complexity of APSP in n -node graphs is $n^{3-o(1)} = |\text{input}|^{3/2-o(1)}$. For $h = \Theta(m^\eta)$, the $O(hm)$ time bound is actually $O(|\text{input}|^{1+\eta})$. Fine-grained reductions from problems with smaller complexity to problems with larger complexity are possible (see, e.g., [22]) but rare, and to our best knowledge no such reduction from APSP is known. If Theorem 1 worked for $\eta > 1/2$, this would be the first such example.

¹ We could also consider a variant of the problem asking for a *walk* with *exactly* h edges. It is the harder of the two variants (adding a length-0 self-loop to node s reduces the “at most h ” variant to the “exactly h ” variant), and we prove the hardness of the easier one already.

■ **Table 1** Summary and comparison of our conditional hardness results for the problem of finding (the length of) a shortest hop-bounded path between two nodes.

	Density	Hops	Hypothesis
Corollary 2	$n = \Theta(m^\nu)$, $\nu \in [1/2, 1]$	$h = \Theta(m^\eta)$, $\eta \in (0, 1/2]$	APSP
Corollary 4	$n = \Theta(m^\nu)$, $\nu \in [1/2, 1]$	$h = \Theta(m^\eta)$, $\eta \in [1/2, \nu]$	Min-Plus Convolution



■ **Figure 1** Parameter space. The upper-left triangle represents the case of $h \geq n$, where the problem degenerates to the standard Shortest Path problem, without hop bound, which can be solved in $\tilde{O}(m)$ time.

In order to cover the remaining combinations of parameters, we use a stronger hypothesis, concerning a problem with conjectured quadratic time complexity, namely the Min-Plus Convolution Hypothesis. Since this hypothesis implies the APSP Hypothesis, it is also a sufficient condition for Theorem 1 and thus gives hardness for the entire parameter space.

► **Theorem 3.** *Unless the Min-Plus Convolution Hypothesis fails, there is neither an $O(h^{1-\varepsilon}m)$ nor an $O(hm^{1-\varepsilon})$ time algorithm for finding the length of a shortest h -hop-bounded s - t path problem in undirected graphs with nonnegative edge weights, for any constant $\varepsilon > 0$.*

This holds even restricted to instances with density $n = \Theta(m^\nu)$ and hop bound $h = \Theta(m^\eta)$ for arbitrarily chosen $\eta \in [1/2, 1]$.

Just like before, we can sparsify the hard instances by adding isolated nodes.

► **Corollary 4.** *The result of Theorem 3 holds even restricted to instances with density $n = \Theta(m^\nu)$ and hop bound $h = \Theta(m^\eta)$ for arbitrarily chosen $\nu \in [1/2, 1]$ and $\eta \in [1/2, \nu]$.*

Combining Corollaries 2 and 4, we cover the entire range of parameters $\nu \in [1/2, 1]$ and $\eta \in (0, \nu]$ for which the $O(hm)$ running time is optimal; see Figure 1.

Let us point out that, even though the Bellman–Ford algorithm finds paths from a single source s to all the nodes in the graph, the above two hardness results hold even for the easier problem of finding a single path between two distinguished nodes s and t . Moreover, note that any (shortest path) algorithm for directed graphs could also be used for undirected graphs (but not necessarily vice versa). Bellman–Ford works in directed graphs with possibly negative edge weights, while our hardness results already hold for undirected graphs with nonnegative edge weights.

1.2 Hardness assumptions

In this section, we briefly recall the two hypotheses that we use in our theorems.

The APSP Hypothesis is the assertion that the All-Pairs Shortest Paths (APSP) problem in n -node graphs cannot be solved in truly subcubic $O(n^{3-\varepsilon})$ time for any constant $\varepsilon > 0$. It is one of the three main hypotheses in fine-grained complexity, the other two being the 3SUM Hypothesis and Strong Exponential Time Hypothesis (SETH) [26]. The APSP Hypothesis is strengthened by the existence of a large class of problems that are equivalent to the APSP problem under subcubic reductions [27, 26], and by the lack of a truly subcubic algorithm for any of these problems.

In the Min-Plus Convolution problem, we are given two sequences $(a[i])_{i=0}^{n-1}$, $(b[i])_{i=0}^{n-1}$, and the goal is to output sequence $(c[i])_{i=0}^{n-1}$ defined as $c[k] = \min_{i+j=k} (a[i] + b[j])$. So far, only subpolynomial $2^{O(\sqrt{\log n})}$ -factor improvements [6, 28] over the naive quadratic running time are known. The Min-Plus Convolution Hypothesis [21, 11] states that the problem cannot be solved in truly subquadratic time, that is, $O(n^{2-\varepsilon})$ for any constant $\varepsilon > 0$. Similarly to APSP, there is also a class (albeit smaller) of problems equivalent to Min-Plus Convolution under subquadratic reductions [11].²

The two hypotheses are closely related because APSP is runtime-equivalent (up to constant factors) to the Min-Plus Product problem [13], which is the matrix product analogue of Min-Plus Convolution. There is a reduction from the convolution to the product problem [6], which entails that the Min-Plus Convolution Hypothesis implies the APSP Hypothesis, and the former is therefore a stronger assumption.

As is customary in fine-grained complexity, these hypotheses, as well as all the results in this paper, are stated in the word RAM model of computation with $O(\log n)$ -bit machine words. We assume all input numbers fit into single machine words. Alternatively, the APSP Hypothesis is sometimes stated as follows [26]: For every $\varepsilon > 0$ there is a constant c such that APSP cannot be solved in $O(N^{3-\varepsilon})$ time in N -node graphs with edge weights in $\{-N^c, \dots, N^c\}$. Our results could also be stated this way because our reductions do not increase weights by more than polynomial factors.

1.3 Related work

Hop-bounded paths are studied in various areas related to graph algorithms, e.g., distributed algorithms [17], dynamic algorithms [25, 23], or even polyhedral combinatorics [12]. Problems of finding shortest hop-bounded paths appear, e.g., in the context of quality-of-service (QoS) routing in networks [2, 9].

Guérin and Orda [19] and Cheng and Ansari [10] studied the problem of finding shortest h -hop-bounded paths from single source s to all nodes in the graph and for all hop bounds $h \leq H$. They proved an $\Omega(Hm)$ lower bound for that problem against so-called *path-comparison-based algorithms*, i.e., algorithms that only access the edge weights by comparing the lengths of two paths. Although Dijkstra and Bellman–Ford are known to be path-comparison-based, algebraic algorithms relying on fast matrix multiplication are not.

² One of the problems equivalent to Min-Plus Convolution is the Knapsack problem on instances with target value $t = \Theta(n)$. Recall that Knapsack can be solved in $O(nt)$ time by a dynamic programming algorithm due to Bellman [3]. Hence, we can half-jokingly rephrase Theorem 3 and say that if one Bellman’s algorithm is optimal for Knapsack, then the other Bellman’s algorithm is optimal for shortest hop-bounded paths.

Bicriteria Path. In the Bicriteria Path problem, we are given a graph $G = (V, E)$ with two types of nonnegative edge weights $l, c : E \rightarrow \mathbb{Z}$, called *lengths* and *costs*, respectively; two budgets $L, C \in \mathbb{Z}$; and two distinguished nodes $s, t \in V$. The goal is to find a path from s to t with the total length at most L and the total cost at most C . Joksč's algorithm [20] solves the problem in pseudopolynomial $O(\min(L, C) \cdot m)$ time. For the special case of all edge costs equal to 1, the Bicriteria Path problem is equivalent to the problem we study in this paper and, moreover, Joksč's algorithm runs in the same time as the Bellman–Ford algorithm.

Aboud, Bringmann, Hermelin, and Shabtay [1] proved that, unless SETH fails, there is no algorithm solving the Bicriteria Path problem on sparse graphs (with $m = \Theta(n)$ edges) with budgets $L, C = \Theta(n^\gamma)$ in time $O(n^{1+\gamma-\varepsilon})$ for any $\varepsilon > 0$ and $\gamma > 0$. In other words, they proved that Joksč's algorithm is conditionally optimal, up to subpolynomial factors. Their reduction, however, heavily uses both types of edge weights, and hence it does not imply any lower bound for the special case with unit costs, i.e., for our problem of interest.

2 Hardness under APSP Hypothesis


Preliminaries. Given a complete tripartite graph $G = (A \cup B \cup C, E)$ with edge weights $w : E \rightarrow \mathbb{Z}$, the Negative Triangle problem asks to find three nodes $a \in A$, $b \in B$, and $c \in C$ with $w(a, b) + w(b, c) + w(c, a) < 0$. Vassilevska Williams and Williams [27] proved that APSP and Negative Triangle are equivalent under subcubic reductions. In particular, unless the APSP Hypothesis fails, there is no $O(N^{3-\varepsilon})$ time algorithm for Negative Triangle with $|A| = |B| = |C| = N$, for any $\varepsilon > 0$.

Via a by now standard argument, under the same assumption, for any $\varepsilon > 0$, there is no $O(N^{\alpha+2-\varepsilon})$ time algorithm for the problem restricted to instances with $|A| = \Theta(N^\alpha)$ and $|B| = |C| = N$ for arbitrarily chosen $\alpha \in (0, 1]$. Specifically, the reduction partitions the original set A into $\Theta(N^{1-\alpha})$ subsets of size $\Theta(N^\alpha)$ each; the sets B and C are copied to all $\Theta(N^{1-\alpha})$ produced instances. A negative triangle exists in the original instance if and only if it exists in at least one of the produced instances. Thus, if each of the obtained instances could be solved in $O(N^{\alpha+2-\varepsilon})$ time, then the original instance could be solved in $O(N^{1-\alpha} \cdot (N^2 + N^{\alpha+2-\varepsilon})) = O(N^{3-\alpha} + N^{3-\varepsilon})$ time³, violating the APSP Hypothesis.

Reduction. We show how to reduce an instance of Negative Triangle, with $|A| = \Theta(N^\alpha)$ and $|B| = |C| = N$, to finding the minimum length of an h -hop-bounded s - t path in an undirected graph with $n = \Theta(N)$ nodes, $m = \Theta(N^2)$ edges, and the hop bound $h = \Theta(N^\alpha)$. In order to prove Theorem 1, we set $\alpha = 2\eta$ so that $n = \Theta(N) = \Theta(\sqrt{m})$ and $h = \Theta(N^\alpha) = \Theta(N^{2\eta}) = \Theta(m^\eta)$. An $O(h^{1-\varepsilon}m)$ time (or $O(hm^{1-\varepsilon})$ time) algorithm finding a shortest h -hop-bounded s - t path would thus yield an $O(N^{\alpha(1-\varepsilon)+2}) = O(N^{\alpha+2-\alpha\varepsilon})$ time (respectively, $O(N^{\alpha+2-2\varepsilon})$ time) algorithm for the original instance of the Negative Triangle problem. As argued above, no such algorithm exists unless the APSP Hypothesis fails.

Let $P = |A|$. Suppose that $A = \{a_1, \dots, a_P\}$, $B = \{b_1, \dots, b_N\}$, and $C = \{c_1, \dots, c_N\}$. Moreover, let W denote the maximum absolute value of an edge weight. Create an undirected graph (see Figure 2) with node set $A \cup B \cup C \cup \bar{A}$, where $\bar{A} = \{\bar{a}_1, \dots, \bar{a}_P\}$ is a disjoint copy of A .

³ The N^2 term corresponds to the time it takes to construct each instance, which consists of $O(N^2)$ edges.



■ **Figure 2** The graph created in the reduction from Negative Triangle.

- For every $a_i \in A$ and $b_j \in B$, add edge $\{a_i, b_j\}$ with weight $w(a_i, b_j) + 3(P + 1 - i)W$.
- For every $b_j \in B$ and $c_k \in C$, add edge $\{b_j, c_k\}$ with weight $w(b_j, c_k) + 3W$.
- For every $c_k \in C$ and $a_i \in A$, add edge $\{c_k, \bar{a}_i\}$ with weight $w(c_k, \bar{a}_i) + 3iW$.
- For every $i \in \{1, \dots, P - 1\}$, add edges $\{a_i, a_{i+1}\}$ and $\{\bar{a}_i, \bar{a}_{i+1}\}$ with weights 0.

Consider a shortest path in this graph from $s \stackrel{\text{def}}{=} a_1$ to $t \stackrel{\text{def}}{=} \bar{a}_P$ using at most $h \stackrel{\text{def}}{=} P + 2$ hops. We claim that its total length is less than $(3P + 2)W$ if and only if there is a negative triangle in the initial graph. Indeed, each triple $(a_i, b_j, c_k) \in A \times B \times C$ corresponds to path $a_1 - a_2 - \dots - a_i - b_j - c_k - \bar{a}_i - \bar{a}_{i+1} - \dots - \bar{a}_P$, which uses exactly $P + 2$ hops and has total length

$$\begin{aligned} w(a_i, b_j) + 3(P + 1 - i)W + w(b_j, c_k) + 3W + w(c_k, \bar{a}_i) + 3iW \\ = (w(a_i, b_j) + w(b_j, c_k) + w(c_k, \bar{a}_i)) + (3P + 2)W. \end{aligned}$$

Hence, the “if” direction follows. For the “only if” direction, fix an s - t path with at most $P + 2$ hops and a total length strictly less than $(3P + 2)W$. The path must be of the form $a_1 - a_2 - \dots - a_i - b_j - \dots - c_k - \bar{a}_{i'} - \bar{a}_{i'+1} - \dots - \bar{a}_P$, where $\{a_i, b_j\}$ is the first edge that leaves A and $\{c_k, \bar{a}_{i'}\}$ is the last edge that enters \bar{A} . Every edge has weight at least 0, every edge incident to b_j or c_k has weight at least $-W + 3W = 2W$, and the direct edge between b_j and c_k has weight at most $W + 3W = 2 \cdot 2W$. Thus, the direct edge is the cheapest walk from b_j to c_k both in terms of the length and the number of hops. Consequently, we may assume without loss of generality that our s - t path proceeds directly from b_j to c_k . This means that the number of hops is $i + 3 + (P - 1 - i') = P + 2 + i - i'$, whereas the total length is

$$\begin{aligned} w(a_i, b_j) + 3(P + 1 - i)W + w(b_j, c_k) + 3W + w(c_k, \bar{a}_{i'}) + 3i'W \\ = (w(a_i, b_j) + w(b_j, c_k) + w(c_k, \bar{a}_{i'})) + (3P + 2)W + 3(i' - i)W. \end{aligned}$$

If $i' < i$, then the number of hops is at least $P + 3$, which is larger than assumed. If $i' > i$, then the path length is at least $-3W + (3P + 2)W + 3W \geq (3P + 2)W$, a contradiction. Therefore, $i = i'$ holds. Since the path length is less than $(3P + 2)W$, we conclude that $w(a_i, b_j) + w(b_j, c_k) + w(c_k, \bar{a}_i) < 0$, i.e., (a_i, b_j, c_k) is a negative triangle in the initial graph.

3 Hardness under Min-Plus Convolution Hypothesis

Preliminaries. In the Max-Plus Convolution Upper Bound problem, we are given three sequences $(a[i])_{i=0}^{n-1}$, $(b[i])_{i=0}^{n-1}$, and $(c[i])_{i=0}^{n-1}$ of n numbers each, and the goal is to decide whether $c[k] \geq \max_{i+j=k} (a[i] + b[j])$ holds for all k . In other words, we want to find i, j, k such that $i + j = k$ and $c[k] < a[i] + b[j]$. The Max-Plus Convolution Upper Bound and Min-Plus Convolution problems are equivalent under subquadratic reductions [11, Theorem 3.1]; thus, in particular, unless the Min-Plus Convolution Hypothesis fails, there is no $O(n^{2-\varepsilon})$ time algorithm for Max-Plus Convolution Upper Bound, for any $\varepsilon > 0$.

Let us introduce an intermediate problem, which we call Common Max-Plus Convolution Upper Bound: Given M pairs of sequences

$$((a_1[i])_{i=0}^{N-1}, (b_1[i])_{i=0}^{N-1}), \dots, ((a_M[i])_{i=0}^{N-1}, (b_M[i])_{i=0}^{N-1}),$$


and one sequence $(c[i])_{i=0}^{N-1}$, decide if there exist i, j, k, ℓ such that $i + j = k$ and $c[k] < a_\ell[i] + b_\ell[j]$. We call such (i, j, k, ℓ) a *violating quadruple*. First, we show that the naive running time of $O(MN^2)$ is conditionally optimal for this problem.

► **Lemma 5.** *Unless the Min-Plus Convolution Hypothesis fails, there is no $O(MN^{2-\varepsilon})$ time algorithm for Common Max-Plus Convolution Upper Bound, for any $\varepsilon > 0$, even when restricted to instances with $M = \Theta(N^\alpha)$ for an arbitrarily chosen constant $\alpha \geq 0$.*

Proof. The argument is based on a self-reduction of Min-Plus Convolution (see [11, proof of Theorem 5.5]). Let $\beta = \alpha/(1+\alpha) \in [0, 1)$. We start with an instance of Max-Plus Convolution Upper Bound with three sequences a, b, c , each of length n . We split a and b into $\Theta(n^\beta)$ blocks of consecutive elements, each block of length $\lceil n^{1-\beta} \rceil$ (the last block can be shorter). For every pair of blocks, one from a and the other from b , we want to check if the corresponding fragment of c is an upper bound of their max-plus convolution. Similarly to [11], we add suitable padding so that all three sequences are of the same length. This way, we end up with $\Theta(n^{2\beta})$ smaller instances of Max-Plus Convolution Upper Bound. The key step is to classify these instances according to the third sequence, which results in $\Theta(n^\beta)$ groups of size $\Theta(n^\beta)$ each (the instances in any single group share the same third sequence). Each such group becomes a single instance of Common Max-Plus Convolution Upper Bound, with $M = \Theta(n^\beta)$ and $N = \Theta(n^{1-\beta})$. If each of these instances can be solved in $O(MN^{2-\varepsilon}) = O(n^{\beta+(1-\beta)(2-\varepsilon)})$ time, then the original instance can be solved in $O(n^{2\beta+(1-\beta)(2-\varepsilon)}) = O(n^{2-(1-\beta)\varepsilon})$ time, and the Min-Plus Convolution Hypothesis fails. We conclude the proof by observing that $n^\beta = n^{\alpha/(1+\alpha)} = (n^{1/(1+\alpha)})^\alpha = (n^{1-\beta})^\alpha$, and thus $M = \Theta(N^\alpha)$ holds as desired. ◀

Cygan, Mucha, Węgrzycki, and Włodarczyk [11, proof of Theorem 5.4] showed that, without loss of generality, the input sequences to Max-Plus Convolution Upper Bound are nonnegative and strictly increasing. The same argument applies to Common Max-Plus Convolution Upper Bound. Using the additional structure, we can replace the condition $i + j = k$ with $i + j \leq k$. Indeed, suppose we find (i, j, k, ℓ) with $i + j \leq k$ and $c[k] < a_\ell[i] + b_\ell[j]$; then, by monotonicity, $c[i + j] \leq c[k]$, and hence $(i, j, i + j, \ell)$ is a quadruple satisfying the original condition.

► **Observation 6.** *The lower bound of Lemma 5 holds even restricted to instances with all the sequences nonnegative and strictly increasing. For such instances, the condition $i + j = k$ can be equivalently replaced by $i + j \leq k$.*



■ **Figure 3** The graph created in the reduction from Common Max-Plus Convolution Upper Bound.

Reduction. We show how to reduce an instance of Common Max-Plus Convolution Upper Bound, with M pairs of length- N sequences, to finding the length of a shortest hop-bounded s - t path in an undirected graph with $n = \Theta(N + M)$ nodes, $m = \Theta(NM)$ edges, and the hop bound $h = \Theta(N)$. This will let us conclude that an $O(h^{1-\varepsilon}m)$ time (or $O(hm^{1-\varepsilon})$ time) shortest path algorithm would give an $O(MN^{2-\varepsilon})$ time (respectively, $O(M^{1-\varepsilon}N^{2-\varepsilon})$ time) algorithm for the Common Max-Plus Convolution Upper Bound problem and thus violate the Min-Plus Convolution Hypothesis.

Let W denote the maximum value of any input sequence element. Create an undirected graph (see Figure 3) composed of three paths $u_0 - u_1 - \dots - u_{N-1}$, $v_0 - v_1 - \dots - v_{N-1}$, $w_0 - w_1 - \dots - w_{N-1}$, and an independent set of M nodes x_1, x_2, \dots, x_M . Set the weights of all the path edges to 0. For every $i \in \{0, 1, \dots, N-1\}$ and $\ell \in \{1, 2, \dots, M\}$, add an edge between u_i and x_ℓ with weight $5W - a_\ell[i]$. Then, for every $j \in \{0, 1, \dots, N-1\}$ and $\ell \in \{1, 2, \dots, M\}$, add an edge between x_ℓ and v_j with weight $5W - b_\ell[j]$. Finally, for every $k \in \{0, 1, \dots, N-1\}$, add an edge between v_0 and w_k with weight $5W + c[k]$.

Consider a shortest path in this graph from $s \stackrel{\text{def}}{=} u_0$ to $t \stackrel{\text{def}}{=} w_{N-1}$ using at most $h \stackrel{\text{def}}{=} N + 2$ hops. We claim that its total length is less than $15W$ if and only if there is a quadruple (i, j, k, ℓ) with $i + j \leq k$ and $c[k] < a_\ell[i] + b_\ell[j]$. Indeed, each quadruple (i, j, k, ℓ) corresponds to path

$$u_0 - u_1 - \dots - u_i - x_\ell - v_j - v_{j-1} - \dots - v_0 - w_k - w_{k+1} - \dots - w_{N-1}. \quad (\star)$$

Such a path uses $i + 1 + 1 + j + 1 + (N - 1 - k) = (N + 2) + (i + j - k)$ hops and has total length $15W - a_\ell[i] - b_\ell[j] + c[k]$. Hence, the “if” direction follows. For the “only if” direction, since every non-zero edge weight is at least $4W$, an s - t path of total length less than $15W$ can use at most three such edges, and therefore it must be of the form (\star) . The hop bound implies $i + j - k \leq 0$ and the total length bound implies $c[k] - a_\ell[i] - b_\ell[j] < 0$.

Recall that $n = \Theta(N + M)$, $m = \Theta(NM)$, and $h = \Theta(N)$. For $\eta \in [1/2, 1]$, in order to get hard shortest hop-bounded path instances with density $\Theta(m^\eta)$ and hop bound $h = \Theta(m^\eta)$, we start with the Common Max-Plus Convolution Upper Bound problem restricted to instances with $M = \Theta(N^{(1-\eta)/\eta})$. This implies $h = \Theta(N) = \Theta((NM)^\eta) = \Theta(m^\eta)$. Moreover, due to $\eta \geq 1/2$, we have $M \leq O(N)$, and thus $n = \Theta(N + M) = \Theta(N) = \Theta(m^\eta)$ holds as desired. This concludes the proof of Theorem 3.

References

- 1 Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. SETH-based lower bounds for subset sum and bicriteria path. *ACM Trans. Algorithms*, 18(1):6:1–6:22, 2022. doi:10.1145/3450524.
- 2 Anantaram Balakrishnan and Kemal Altinkemer. Using a hop-constrained model to generate alternative communication network design. *INFORMS J. Comput.*, 4(2):192–205, 1992. doi:10.1287/ijoc.4.2.192.
- 3 Richard Bellman. Notes on the theory of dynamic programming IV - maximization over discrete sets. *Naval Research Logistics Quarterly*, 3(1-2):67–70, March 1956. doi:10.1002/nav.3800030107.
- 4 Richard Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958. doi:10.1090/qam/102435.
- 5 Aaron Bernstein, Danupon Nanongkai, and Christian Wulff-Nilsen. Negative-weight single-source shortest paths in near-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 600–611. IEEE, 2022. doi:10.1109/FOCS54457.2022.00063.
- 6 David Bremner, Timothy M. Chan, Erik D. Demaine, Jeff Erickson, Ferran Hurtado, John Iacono, Stefan Langerman, Mihai Patrascu, and Perouz Taslakian. Necklaces, convolutions, and X+Y. *Algorithmica*, 69(2):294–314, 2014. doi:10.1007/s00453-012-9734-3.
- 7 Karl Bringmann, Alejandro Cassis, and Nick Fischer. Negative-weight single-source shortest paths in near-linear time: Now faster! In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*. IEEE, 2023. arXiv:2304.05279.
- 8 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.
- 9 Shigang Chen and Klara Nahrstedt. An overview of quality of service routing for next-generation high-speed networks: problems and solutions. *IEEE Netw.*, 12(6):64–79, 1998. doi:10.1109/65.752646.
- 10 Gang Cheng and Nirwan Ansari. Finding all hops shortest paths. *IEEE Commun. Lett.*, 8(2):122–124, 2004. doi:10.1109/LCOMM.2004.823365.
- 11 Marek Cygan, Marcin Mucha, Karol Węgrzycki, and Michał Włodarczyk. On problems equivalent to (min, +)-convolution. *ACM Trans. Algorithms*, 15(1):14:1–14:25, 2019. doi:10.1145/3293465.
- 12 Geir Dahl and Luis Eduardo Neves Gouveia. On the directed hop-constrained shortest path problem. *Oper. Res. Lett.*, 32(1):15–22, 2004. doi:10.1016/S0167-6377(03)00026-9.
- 13 Michael J. Fischer and Albert R. Meyer. Boolean matrix multiplication and transitive closure. In *12th Annual Symposium on Switching and Automata Theory*, pages 129–131. IEEE Computer Society, 1971. doi:10.1109/SWAT.1971.4.
- 14 Lester Randolph Ford. Network flow theory. Technical report, RAND Corporation, 1956. URL: <https://www.rand.org/pubs/papers/P923.html>.
- 15 Harold N. Gabow. Scaling algorithms for network problems. *J. Comput. Syst. Sci.*, 31(2):148–168, 1985. doi:10.1016/0022-0000(85)90039-X.
- 16 Harold N. Gabow and Robert Endre Tarjan. Faster scaling algorithms for network problems. *SIAM J. Comput.*, 18(5):1013–1036, 1989. doi:10.1137/0218069.
- 17 Mohsen Ghaffari, Bernhard Haeupler, and Goran Zuzic. Hop-constrained oblivious routing. In *53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2021*, pages 1208–1220. ACM, 2021. doi:10.1145/3406325.3451098.
- 18 Andrew V. Goldberg. Scaling algorithms for the shortest paths problem. *SIAM J. Comput.*, 24(3):494–504, 1995. doi:10.1137/S0097539792231179.
- 19 Roch Guérin and Ariel Orda. Computing shortest paths for any number of hops. *IEEE/ACM Trans. Netw.*, 10(5):613–620, 2002. doi:10.1109/TNET.2002.803917.

- 20 H. C. Jochsch. The shortest route problem with constraints. *Journal of Mathematical Analysis and Applications*, 14(2):191–197, 1966. doi:10.1016/0022-247X(66)90020-5.
- 21 Marvin Künnemann, Ramamohan Paturi, and Stefan Schneider. On the fine-grained complexity of one-dimensional dynamic programming. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, volume 80 of *LIPICs*, pages 21:1–21:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.ICALP.2017.21.
- 22 Andrea Lincoln, Adam Polak, and Virginia Vassilevska Williams. Monochromatic triangles, intermediate matrix products, and convolutions. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020*, volume 151 of *LIPICs*, pages 53:1–53:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ITCS.2020.53.
- 23 Jakub Łącki and Yasamin Nazari. Near-optimal decremental hopsets with applications. In *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022*, volume 229 of *LIPICs*, pages 86:1–86:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.ICALP.2022.86.
- 24 Alfonso Shimbel. Structure in communication nets. In *Symposium on Information Networks, 1954*, pages 199–203. Polytechnic Press of the Polytechnic Institute of Brooklyn, 1955.
- 25 Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *37th Annual ACM Symposium on Theory of Computing, STOC 2005*, pages 112–119. ACM, 2005. doi:10.1145/1060590.1060607.
- 26 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *International Congress of Mathematicians, ICM 2018*, pages 3447–3487, 2018. doi:10.1142/9789813272880_0188.
- 27 Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5):27:1–27:38, 2018. doi:10.1145/3186893.
- 28 R. Ryan Williams. Faster all-pairs shortest paths via circuit complexity. *SIAM J. Comput.*, 47(5):1965–1985, 2018. doi:10.1137/15M1024524.