

Fine-Grained Optimality of Partially Dynamic Shortest Paths and More

Barna Saha*, Virginia Vassilevska Williams†, Yinzhan Xu‡, Christopher Ye§

July 16, 2024

Abstract

Single Source Shortest Paths (SSSP) is among the most well-studied problems in computer science. In the incremental (resp. decremental) setting, the goal is to maintain distances from a fixed source in a graph undergoing edge insertions (resp. deletions). A long line of research culminated in a near-optimal deterministic $(1 + \varepsilon)$ -approximate data structure with $m^{1+o(1)}$ total update time over all m updates by Bernstein, Probst Gutenberg and Saranurak [FOCS 2021]. However, there has been remarkably little progress on the exact SSSP problem beyond Even and Shiloach’s algorithm [J. ACM 1981] for unweighted graphs. For weighted graphs, there are no exact algorithms beyond recomputing SSSP from scratch in $\tilde{O}(m^2)$ total update time, even for the simpler Single-Source Single-Target Shortest Path problem (stSP). Despite this lack of progress, known (conditional) lower bounds only rule out algorithms with amortized update time better than $m^{1/2-o(1)}$ in dense graphs.

In this paper, we give a tight (conditional) lower bound: any partially dynamic exact stSP algorithm requires $m^{2-o(1)}$ total update time for any sparsity m . We thus resolve the complexity of partially dynamic shortest paths, and separate the hardness of exact and approximate shortest paths, giving evidence as to why no non-trivial exact algorithms have been obtained while fast approximation algorithms are known.

Moreover, we give tight bounds on the complexity of combinatorial algorithms for several path problems that have been studied in the static setting since early sixties: Node-weighted shortest paths (studied alongside edge-weighted shortest paths), bottleneck paths (early work dates back to 1960), and earliest arrivals (early work dates back to 1958). These bounds rule out any nontrivial combinatorial algorithms for these problems in the partially dynamic setting. Interestingly, for all of the above path-variant problems, we obtain faster partially dynamic algorithms using fast matrix multiplication.

*University of California San Diego. bsaha@ucsd.edu. Supported by NSF grants 1652303, 1909046, 2112533, and HDR TRIPODS Phase II grant 2217058. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

†Massachusetts Institute of Technology. virgi@mit.edu. Supported by NSF Grant CCF-2330048, BSF Grant 2020356 and a Simons Investigator Award. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

‡Massachusetts Institute of Technology. xyzhan@mit.edu. Partially supported by NSF Grant CCF-2330048, BSF Grant 2020356 and a Simons Investigator Award. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, and the Institute for Emerging CORE Methods in Data Science (EnCORE).

§University of California San Diego. czye@ucsd.edu. Partially supported by NSF grants 1652303, 1909046, 2112533, and HDR TRIPODS Phase II grant 2217058. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

1 Introduction

Computing shortest paths is one of the most fundamental tasks in computer science. Given an n -node, m -edge *directed* graph¹ $G = (V, E)$ with nonnegative integer edge weights, the Single Source Shortest Paths (SSSP) problem asks to compute the distance from a given node s to all $v \in V$. Dijkstra's classical algorithm [Dij59, FT87] solves this problem in $O(m + n \log n)$ time. SSSP (and its variants for undirected graphs and for graphs with negative weights) has been extensively studied [Shi53, For56, Bel58, Moo59, Gab83, GT89, FW93, FW94, Gol95, Ram96, Tho96, Ram97, Tho99, Hag00, Tho00, PR05, Tho04, vdBLN⁺20, AMV20, BNW22, CKL⁺22, Fin24]. The current fastest SSSP algorithm runs in $O(m + n \log \log n)$ time [Tho04].

For the all-pairs version of the problem, All-Pairs Shortest Paths (APSP), one needs to compute the distance between every pair of nodes in the given graph. One could run Dijkstra's algorithm n times to obtain an $O(mn + n^2 \log n)$ running time. For a dense graph where $m = \Theta(n^2)$, this running time becomes $O(n^3)$. An alternative classical algorithm, the Floyd–Warshall algorithm, also achieves an $O(n^3)$ running time. A long series of works (e.g. [Fre76, Dob90, Tak92, Tak98, Han04, Pet04, Zwi04, Cha05, Han06, Cha10a, HT12, OV22]) improved poly-logarithmic factors over the classic $O(n^3)$ running time, culminating in the current fastest algorithm by Williams [Wil18] with $n^3/2^{\Theta(\sqrt{\log n})}$ running time, a super poly-logarithmic improvement over the classic $O(n^3)$ running time. Due to the lack of polynomial improvements, the APSP hypothesis, stating that there is no $O(n^{3-\varepsilon})$ time algorithm for any $\varepsilon > 0$ that computes APSP for n -node graphs, is among the most popular hypotheses in Fine-Grained Complexity.

Shortest paths problems have also been extensively studied in dynamic graphs. The main focus of our paper is the *partially dynamic* setting, where the dynamic graph only undergoes edge insertions (incremental) or only undergoes edge deletions (decremental). Partially dynamic problems are often easier than their fully-dynamic counterparts where both edge insertions and deletions are allowed. We are interested in the total running time of the algorithms over all the operations (starting from an empty graph in the incremental case, or deleting edges until the graph is empty in the decremental case). A seminal algorithm of Demetrescu and Italiano [DI04] shows how to maintain APSP in amortized $\tilde{O}(n^2)$ time,² even for fully dynamic graphs. This amortized running time is essentially optimal if our goal is to minimize the total running time, as the output size after each operation is already n^2 for APSP.³ Therefore, we will focus on partially dynamic SSSP.

For unweighted graphs, Even and Shiloach [ES81] gave a partially dynamic SSSP algorithm with total time $O(mn)$, a bound that has stood for more than 40 years. In the meantime, a long line of work has studied $(1 + \varepsilon)$ -approximate SSSP, where the algorithm must maintain approximate shortest paths after each update [BR11, HKN14a, HKN14b, BC16, Ber16, BC17, Ber17, HKN18, CK19, BGS20, GVW20, GW20a, GW20b, BGS21].

However, there has been little progress on maintaining *exact* SSSP in partially dynamic weighted graphs. The naive algorithm for partially dynamic SSSP is to rerun Dijkstra's algorithm after every update, which results in an $\tilde{O}(m^2)$ total running time, or $\tilde{O}(m)$ time per-update, even when amortized over all m updates. For dense graphs, this running time is $\tilde{O}(n^4)$. This trivial recomputation is essentially the best known for partially dynamic SSSP. Even for the simpler partially dynamic stSP problem where we only need to maintain the distance between one source node s and one sink node t , there are no better upper bounds.

¹Throughout the paper, we will assume $m \geq n - 1$.

² \tilde{O} hides $\text{polylog}(n)$ factors.

³It is still an active line of research to improve the worst-case time guarantee for fully dynamic APSP [Tho05, ACK17, GW20c, CZ23, Mao24], and the current fastest algorithm is designed by Mao [Mao24] which runs in $\tilde{O}(n^{2.5})$ time per operation.

Open Question 1. *Can we solve partially dynamic stSP in $O(m^{2-\varepsilon})$ total time for some $\varepsilon > 0$, for any sparsity m ? In particular, can we solve partially dynamic stSP in $O(n^{4-\varepsilon})$ total time for some $\varepsilon > 0$?*

There has been partial progress towards answering Open Question 1. Roditty and Zwick [RZ11] proved that under the APSP hypothesis, any partially dynamic SSSP algorithm handling n updates requires total time $n^{3-o(1)}$, which is a factor of n away from the naive algorithm for dense graphs. This lower bound was later strengthened by Abboud and Vassilevska Williams [AV14] to hold also for partially dynamic stSP. While this lower bound rules out any algorithm improving upon recomputation from scratch when amortized over only n updates, it only rules out algorithms with $m^{1/2-\varepsilon}$ amortized update time when amortizing over all m updates. In particular, previous lower bounds do not rule out the possibility of an algorithm which has only n costly updates with the rest of the updates being relatively fast, resulting in $m^{1/2-o(1)}$ amortized time when amortizing over all m updates.

More partial progress was made by Gutenberg, Vassilevska Williams and Wein [GVW20], who showed that in the special case of sparse graphs where $m = n^{1+o(1)}$, any algorithm solving partially dynamic stSP requires total time $m^{2-o(1)}$, under the k -Cycle hypothesis from Fine-Grained Complexity (see [LVW18, AHR⁺19]) which postulates that any algorithm that can detect whether a graph contains a k -cycle (for every constant $k \geq 3$) requires $m^{2-o(1)}$ time. In terms of n , detecting a k -cycle can be solved in $O(n^\omega)$ time⁴ for all constants k (see e.g., [IR78]). When $m = \Omega(n^{1+c})$ for some $c > 0$, $O(n^\omega)$ can be faster than m^2 , e.g. in the optimistic case where $\omega = 2$. Therefore, the k -Cycle hypothesis is most reasonable for graphs with $m = n^{1+o(1)}$. Thus, the lower bound in [GVW20] does not help with Open Question 1 for graphs with $m = \Omega(n^{1+c})$ for any $c > 0$. Furthermore, their result is not meaningful for partially dynamic SSSP with $m = n^{1+o(1)}$, because an $\Omega(mn) = m^{2-o(1)}$ lower bound is trivial, as mn is the total output size (n distances after each update).

We remark that *fully dynamic* stSP was shown to require $n^{2-o(1)}$ amortized update time over any (polynomial) number of updates for combinatorial⁵ algorithms by Jin and Xu [JX22]. Indeed, any combinatorial algorithm requires $n^{2-o(1)}$ amortized update time even for the simpler (s, t) -reachability problem in the fully dynamic setting.

Roditty and Zwick [RZ11] considered partially dynamic SSSP where instead of outputting the distances from s to all nodes in V after each update, the data structure only needs to support distance queries between s and any queried node v . They explicitly stated the following open question:

Open Question 2. *Is partially dynamic SSSP with m updates and n^2 queries in $\tilde{O}(mn)$ time?*

In particular, if the answer to Open Question 2 is affirmative, then one can solve partially dynamic stSP in $\tilde{O}(mn)$ time, as the total number of queries for partially dynamic stSP is $m \leq n^2$. Therefore, if the answer to Open Question 2 is affirmative, then the answer to Open Question 1 is also affirmative for $m = \Omega(n^{1+c})$ for $c > 0$. Neither of the previously mentioned lower bounds [AV14, GVW20] can resolve Open Question 2. In the contrapositive, if one can resolve Open Question 1 in the negative, then it would also imply a resolution of Open Question 2 in the negative.

⁴ $\omega < 2.372$ [DWZ23, VXXZ24] is the square matrix multiplication exponent. We will also use $\omega(a, b, c)$ to denote the exponent for multiplying an $n^a \times n^b$ matrix and an $n^b \times n^c$ matrix.

⁵“Combinatorial” algorithms are not well-defined, but the term has become synonymous with algorithms that do not use the heavy algebraic techniques used in fast matrix multiplication algorithms. Hardness under the BMM hypothesis essentially implies that fast algorithms can only exist if fast matrix multiplication is used.

Strong Hardness for Exact Partially Dynamic stSP. As the first main result of our paper, we fully resolve Open Question 1 in the negative based on a popular fine-grained hypothesis. Thus, the answer to Open Question 2 is (conditionally) negative as well, based on the above discussion. In particular, we resolve the complexity of partially dynamic shortest paths, and separate the hardness of exact and approximate shortest paths, giving evidence as to why no non-trivial exact algorithms have been obtained while fast approximation algorithms are known.

Our lower bound is based on the Minimum-Weight 4-Clique hypothesis, which states that finding the 4-clique with minimum total edge weights in an n -node graph requires $n^{4-o(1)}$ time. This hypothesis is a special case of the Minimum-Weight k -Clique hypothesis, which has been used in, e.g., [AVW14, BDT16, BT17, LVW18, ABDN18, BGMW20, GV21, BFvdH⁺24]. In particular, the Minimum-Weight 3-Clique hypothesis is equivalent to the APSP hypothesis [VW18]. Our result can be formally stated as follows:

Theorem 1.1. *Under the Minimum Weight 4-Clique hypothesis, any algorithm computing incremental/decremental stSP on n -node undirected graphs requires $n^{4-o(1)}$ total time.*

Through a simple graph transformation, Theorem 1.1 implies that any partially dynamic stSP algorithm requires $m^{2-o(1)}$ total time over m updates in an n -node m -edge undirected graph for *any sparsity* m . In our work, we mostly consider directed graphs, but notably, Theorem 1.1 even holds for undirected graphs.

1.1 Variants of Shortest Paths in the Dynamic Setting

In the static setting, there are several variants of shortest paths that have been studied extensively. Examples include shortest paths on structured graphs such as node-weighted graphs [Cha10a], bottleneck paths [Pol60, Hu61] (the bottleneck path between two nodes is the path whose minimum weight edge is maximized), and earliest arrivals [Min58] (the earliest arrival between two nodes is the minimum weight of the last edge among all paths whose edges are increasing in weight; earliest arrivals is also known as nondecreasing paths). All these problems are extremely basic and well-motivated: Node-weighted graphs naturally arise in applications for many problems related to shortest paths [VW06, VWY06, CL09, SYZ07, DP09, VW13]. Bottleneck paths [GT88, CKT⁺16] have many applications, from max-flow algorithms ([EK72], Chapter 7.3 in [AMO93]), to voting protocols [Sch11], to metabolic pathway analysis [ULH09] and more. Earliest arrivals has applications to public transportation scheduling such as train and planes (see e.g. [Min58, Vas10]).

Modifications of Dijkstra's algorithm can solve the single-source version of all three variants in $\tilde{O}(m)$ time. Faster than Dijkstra's algorithms exist for all of them: Single Source Shortest Paths on node-weighted graphs (nw-SSSP) can be solved in $O(m + n \log \log n)$ time [Tho04] similar to SSSP, Single Source Earliest Arrival (SSEA) can be solved in $O(m \log \log n)$ time in the addition-comparison model and $O(m + n)$ time on the word-RAM [Vas10], and Single Source Bottleneck Paths (SSBP) can be solved in $O(m\sqrt{\log n})$ time [DLX18] which is faster than Dijkstra's algorithm when $m = o(n\sqrt{\log n})$.

In contrast to APSP, the all-pairs versions of these problems all have *truly subcubic* time algorithms as first shown by [VWY09, Vas10, Cha10a]. The fastest algorithm to date for All-Pairs Shortest Paths on node-weighted graphs (nw-APSP) runs in $\tilde{O}(n^{(9+\omega)/4}) = O(n^{2.843})$ time [Cha10a]. Using rectangular matrix multiplication, this running time can be improved to $O(n^{2.830})$ [Yus09]. The fastest algorithms for All-Pairs Bottleneck Paths (APBP) [DP09] and All-Pairs Earliest Arrivals (APEA) [DJW19] both run in $\tilde{O}(n^{(3+\omega)/2}) = O(n^{2.686})$ time.

Notably, all of the aforementioned subcubic algorithms use fast matrix multiplication, and that is for a good reason. Under the Boolean Matrix Multiplication (BMM) hypothesis, no combinatorial

algorithm for these problems can achieve $O(n^{3-\varepsilon})$ for $\varepsilon > 0$ time.

To the best of our knowledge, the dynamic versions of these natural variants of shortest paths have not been explicitly studied before, despite the extensive study of their static versions. In this paper, we initiate the study of the dynamic versions of node-weighted shortest paths, bottleneck paths, and earliest arrivals, obtaining algorithms and conditional lower bounds.

New Partially Dynamic Algorithms. First, we give sub-quartic total time algorithms for the partially dynamic versions of all three variants, setting them apart from partially dynamic SSSP. In particular, for partially dynamic SSEA, our running time is near-linear. For node-weighted s - t shortest paths (nw-stSP), we also give an algorithm that is faster than nw-SSSP. We give similar faster algorithms for Single Source Bottleneck Paths (SSBP) and s - t Bottleneck Paths (stBP).

Theorem 1.2. *For n -node m -edge graphs, there are algorithms for*

- *incremental nw-stSP in $\tilde{O}(n^{4-\gamma/3}) = O(n^{3.887})$ total time;*
- *incremental nw-SSSP in $\tilde{O}(n^{4-\gamma/4}) = O(n^{3.915})$ total time;*
- *incremental/decremental stBP in $\tilde{O}(n^{2+\varepsilon_1+\varepsilon_2} + n^{2+\omega(1,\varepsilon_1,\varepsilon_2)-\varepsilon_1} + n^{2+\omega(1,1,\varepsilon_2)-\varepsilon_2}) = O(n^{3.405})$ total time;*
- *incremental SSBP in $\tilde{O}(n^{(5+\omega)/2}) = O(n^{3.686})$ time;*
- *incremental or decremental SSEA in $\tilde{O}(m)$ total time,*

where γ is the solution to the equation $\omega(1,1,1+\gamma) = 3-\gamma$ and $\varepsilon_1, \varepsilon_2$ are parameters to be optimized.

New Hardness Results. All algorithms except the one for SSEA in Theorem 1.2 use fast matrix multiplication, and it is natural to ask whether fast matrix multiplication is necessary in order to achieve $O(n^{4-\varepsilon})$ running time for these problems. We show that this is indeed the case, under the Combinatorial 4-Clique hypothesis, which states that any combinatorial algorithm detecting whether a graph contains a 4-clique requires $n^{4-o(1)}$ time. This hypothesis is a natural generalization of the BMM hypothesis, as the latter has been shown to be equivalent to the hypothesis that any combinatorial algorithm for Triangle Detection in n -node graphs requires $n^{3-o(1)}$ time [VW18]. The Combinatorial k -Clique hypothesis has been widely used in prior works (e.g, [Cha10b, BGL17, Li19, AGI⁺19, GVW20, JX22, HLSW23]). In fact, the reductions for this result are adapted from the reduction for Theorem 1.1, showcasing the versatility of our technique.

Theorem 1.3. *Under the Combinatorial 4-Clique hypothesis, any combinatorial algorithm computing incremental/decremental stBP or nw-stSP requires $n^{4-o(1)}$ total time.*

Similar to before, Theorem 1.3 implies an $m^{2-o(1)}$ lower bound for any sparsity m .

Even with fast matrix multiplication, the current best running time for detecting whether a graph contains a 4-clique is $O(n^{\omega(1,2,1)})$ [EG04]. Notably, the current best bound [VXXZ24] on $\omega(1,2,1)$ is $3.251 > 3$; if $\omega = 2$, then $\omega(1,2,1) = 3$.

Our method implies that, unless we can improve this 4-clique running time, any algorithm (even non-combinatorial ones) computing incremental/decremental stBP or nw-stSP requires $n^{\omega(1,2,1)-o(1)}$ total time. We can further improve the exponent for the lower bound, under the hardness of the OMv3 problem, proposed by [GVW20] as a generalization of the OMv problem [HKNS15], and the Minimum-Witness 3-Product problem, which is a generalization of the Minimum-Witness Product problem proposed in [KL22].

In the OMv3 problem, we need to preprocess an $n \times n$ Boolean matrix A , so that we can support the following types of queries: Given three n -dimensional Boolean vectors u, v, w , determine whether there exist $i, j, k \in [n]$, so that $u_i \wedge v_j \wedge w_k \wedge A_{ij} \wedge A_{jk} \wedge A_{ki}$. A naive algorithm solves each query

in $O(n^\omega)$ time, and it is hypothesized in [GVW20] that this is essentially the best algorithm among all algorithms with polynomial preprocessing time.

Hypothesis 1.4 (OMv3 hypothesis [GVW20]). *There is no algorithm with polynomial preprocessing time and total update and query time $O(n^{\omega+1-\varepsilon})$ solving the OMv3 problem, for any $\varepsilon > 0$.*

We show the following results under the OMv3 hypothesis:

Theorem 1.5. *Under the OMv3 hypothesis, any algorithm computing incremental/decremental stBP or nw-stSP requires $n^{\omega+1-o(1)}$ total time.*

For nw-SSSP, we are able to show a higher lower bound, based on the hardness of a variant of the Minimum-Witness Product problem.

In the Minimum-Witness Product problem, we are given two Boolean matrices A and B , and we need to compute $\min\{k \in [n] : A_{i,k} \wedge B_{k,j}\}$ for every $(i, j) \in [n] \times [n]$. The current best algorithm for this problem runs in $O(n^{2+\lambda})$ time [CKL07], where $\omega(1, \lambda, 1) = 1+2\lambda$. If $\omega = 2$, this running time is essentially $O(n^{2.5})$; the hypothesis that there is no $O(n^{2.5-\varepsilon})$ time algorithm for Minimum-Witness Product has been considered before [LPV20].

We consider the following natural generalization of Minimum-Witness Product to three matrices, first proposed by [KL22], which we call Minimum-Witness 3-Product. In this problem, we are given three $n \times n$ Boolean matrices A, B, C , and for every $i_1, i_2, i_3 \in [n]$, we need to find the minimum value of $j \in [n]$ such that $A_{i_1,j} \wedge B_{i_2,j} \wedge C_{i_3,j}$.

[KL22] gave an $O(n^{3.5})$ time algorithm (when $\omega = 2$) for Minimum-Witness 3-Product, which is an adaptation of the existing algorithm for Min-Witness Product [CKL07]. As $O(n^{3.5})$ is the best running time even when $\omega = 2$, the following hypothesis is plausible. See Section 3.2.2 for more discussion on why it is plausible.

Hypothesis 1.6 (Minimum-Witness 3-Product hypothesis). *There is no $O(n^{3.5-\varepsilon})$ time algorithm for the Minimum-Witness 3-Product problem for $\varepsilon > 0$.*

We obtain a conditional lower bound based on the Minimum-Witness 3-Product hypothesis.

Theorem 1.7. *Under the Minimum-Witness 3-Product hypothesis, any algorithm computing incremental/decremental nw-SSSP requires $n^{3.5-o(1)}$ total time.*

Weight-Dynamic Algorithms Typically, updates to dynamic graphs are considered in the context of edge insertion and deletion. [HPS21] introduced *weight-dynamic* graphs, on which updates can additionally modify the weight of an existing edge. Note that this generalizes the insertion/deletion model, as edge insertions/deletions can be modeled by setting weights to ∞ or $-\infty$ depending on the setting of the problem. In the partially dynamic versions of this model, one can either only increase the edges weights, or only decrease the edge weights. In this more general model, all our lower bounds hold, and we additionally show that combinatorial algorithms for stEA require $m^{2-o(1)}$ time under the Combinatorial 4-Clique hypothesis, whereas in the insertion/deletion model we obtain an $\tilde{O}(m)$ time algorithm for SSEA, exhibiting a separation between the two dynamic models.

Fully Dynamic Algorithm The main focus of our paper is partially dynamic algorithms, but we also show a fully dynamic algorithm for APBP. Namely, we show that fully dynamic APBP can be solved in $\tilde{O}(n^2)$ amortized update time.

Problem	Upper Bound	Lower Bound	Combinatorial Lower Bound	Known Lower Bound
Shortest Paths				
(s, t)	n^4	$n^{4-o(1)}$ (Thm. 1.1)	-	$n^{3-o(1)}$ [AV14]
single source	n^4	$n^{4-o(1)}$ (Thm. 1.1)	-	$n^{3-o(1)}$ [RZ11]
Node-Weighted Shortest Paths				
inc. (s, t)	$n^{3.887} [n^{3.834}]$ (Prop. 3.3)	$n^{\omega+1-o(1)} [n^{3-o(1)}]$ (Thm. 3.8)	$n^{4-o(1)}$ (Thm. 3.5)	n^2
inc. single source	$n^{3.915} [n^{3.875}]$ (Prop. 3.1)	$n^{3.5-o(1)}$ (Thm. 3.13)	$n^{4-o(1)}$ (Thm. 3.5)	$n^{2.75-o(1)}$ [RZ11]
Bottleneck Paths				
(s, t)	$n^{3.405} [n^{3.25}]$ (Thm. 4.2)	$n^{\omega+1-o(1)} [n^{3-o(1)}]$ (Thm. 4.18)	$n^{4-o(1)}$ (Thm. 4.16)	n^2
inc. single source	$n^{3.686} [n^{3.5}]$ (Thm. 4.11)	$n^{\omega+1-o(1)} [n^{3-o(1)}]$ (Thm. 4.18)	$n^{4-o(1)}$ (Thm. 4.16)	n^2
Earliest Arrivals				
single source	n^2 (Thm. 5.4)	n^2	n^2	n^2

Table 1: Summary of results for partially dynamic (s, t) and single source problems with respect to number of nodes n . The table shows total update and query time. Unless otherwise stated, our bounds hold for both incremental and decremental versions of the problem. Upper bounds are given ignoring polylogarithmic factors. Bounds are given with respect to the current best upper bound of (rectangular) matrix multiplication [VXXZ24] computed with [Bra]. Bounds within brackets assume $\omega = 2$.

Problem	Lower Bound	Combinatorial Lower Bound	Known Lower Bound
Shortest Paths			
(s, t)	$m^{2-o(1)}$ (Thm. 1.1)	-	$m^{2-o(1)}$ if $m = n^{1+o(1)}$ [GVW20]
Node-Weighted Shortest Paths			
(s, t)	$m^{\frac{\omega+1}{2}-o(1)} [m^{1.5-o(1)}]$ (Thm. 3.8)	$m^{2-o(1)}$ (Thm. 3.5)	m
single source	$n^{1.75-o(1)}$ (Thm. 3.13)	$m^{2-o(1)}$ (Thm. 3.5)	$m^{1.375-o(1)}$ [RZ11]
Bottleneck Paths			
(s, t)	$m^{\frac{\omega+1}{2}-o(1)} [m^{1.5-o(1)}]$ (Thm. 4.18)	$m^{2-o(1)}$ (Thm. 4.16)	m

Table 2: Summary of lower bounds for partially dynamic (s, t) and single source problems with respect to number of edges m . The table shows total update and query time. Our bounds hold for both incremental and decremental versions of the problem and all values of $m \leq n^2$ (Prop A.2). Bounds are given with respect to the current best upper bound ω [VXXZ24]. Bounds within brackets assume $\omega = 2$.

1.2 Preliminaries

Let $[n] = \{1, 2, \dots, n\}$. Let $\vec{v}[i]$ denote the i -th entry of vector \vec{v} and $A[i, j]$ denote the $[i, j]$ -th entry of A . Let $N(v)$ denote the neighborhood of a vertex v . If P_1, P_2 are paths in G with $P_1 = (a, \dots, b), P_2 = (b, \dots, c)$, let $P_1 \circ P_2$ denote the path concatenation (a, \dots, b, \dots, c) .

In this work, we consider a dynamic edge update model. Concretely, an incremental algorithm must handle edge insertions, a decremental must handle deletions, while a fully dynamic algorithm must handle both insertions and deletions. Unless otherwise specified, our upper bounds obtain algorithms that answer queries in $O(1)$ time. When specified, we also consider the weight update model. In this case, an incremental algorithm must handle weight increases, a decremental algorithm must handle weight decreases, while a fully dynamic algorithm must handle arbitrary weight changes.

1.3 Organization

The remainder of the paper is organized as follows. Section 2 is dedicated to showing a tight lower bound for partially dynamic shortest paths. Section 3, Section 4, and Section 5 delve into node-weighted shortest paths, bottleneck paths, and the earliest-arrival problems respectively, giving both lower bounds and faster algorithms in the partially dynamic setting.

2 A Tight Lower Bound for Partially Dynamic Shortest Paths

We now present our main result: a tight conditional lower bound for partially dynamic stSP. As a warm-up, we prove a lower bound for graphs allowing parallel edges. Prior reductions have used APSP, triangle detection as well as 4-Clique instances for dynamic shortest path lower bounds e.g. as in [RZ11, AV14, JX22]. We extend these constructions with the addition of parallel edges with carefully chosen weights and insertion sequence that allow for a partially dynamic lower bound. This extension from previous works with parallel edges in the constructed graph is relatively simple and it was also observed by [Pol24]. Later we explain how to remove the parallel edges, which is the crux of our contribution.

Proposition 2.1. *Under the Minimum-Weight 4-Clique hypothesis, any combinatorial algorithm computing incremental/decremental stSP on undirected (not necessarily simple) graphs with n vertices and m edges requires $n^{4-o(1)}$ total time.*

Suppose we are given a Minimum-Weight 4-Clique instance. As is typical with k -clique problems,⁶ we may assume that the given graph is k -partite. Thus, let the Minimum-Weight 4-Clique instance be a 4-partite graph G with $V(G) = A \cup B \cup C \cup D$, each part consisting of n nodes. We identify the vertices of A, B, C, D with the integer set $[n]$. We can also assume that G is a complete 4-partite graph, as the non-edges can be replaced by edges of large enough weight (for example $6W + 1$ where W is the maximum weight of any edge). For any tuple $(a, b, c, d) \in A \times B \times C \times D$, let

$$\text{wt}(a, b, c, d) = \text{wt}(a, b) + \text{wt}(a, c) + \text{wt}(a, d) + \text{wt}(b, c) + \text{wt}(b, d) + \text{wt}(c, d)$$

denote the weight of the 4-clique (a, b, c, d) . The Minimum-Weight 4-Clique problem asks to compute

$$\min\{\text{wt}(a, b, c, d) \text{ s.t. } a \in A, b \in B, c \in C, d \in D\}.$$

⁶For any k -clique instance $G = (V, E)$ we can create a k -partite G' with vertex set containing k copies of V , V_1, \dots, V_k , each an independent set, and for every two copies V_i and V_j we add an edge (u_i, v_j) for every edge (u, v) in G , where $u_i \in V_i, v_j \in V_j$ are the copies of u and v in V_i and V_j respectively. If G had weights, the vertices and edges of G' inherit the weights of the vertices and edges in G that they represent.

A simple algorithm enumerates all cliques in $O(n^4)$ time. The Minimum-Weight 4-Clique hypothesis states that this is essentially optimal.

Hypothesis 2.2 (Minimum-Weight k -Clique hypothesis). *There is no $O(n^{k-\varepsilon})$ algorithm for Minimum-Weight k -Clique on n -node graphs with non-negative weights, for any $\varepsilon > 0$.*

Proof of Proposition 2.1. For simplicity, we consider the incremental case on directed graphs. Let W be larger than the weight of any 4-clique in the graph (say larger than 6 times the maximum weight in the graph). We construct an stSP instance as illustrated in Figures 1 and 2.

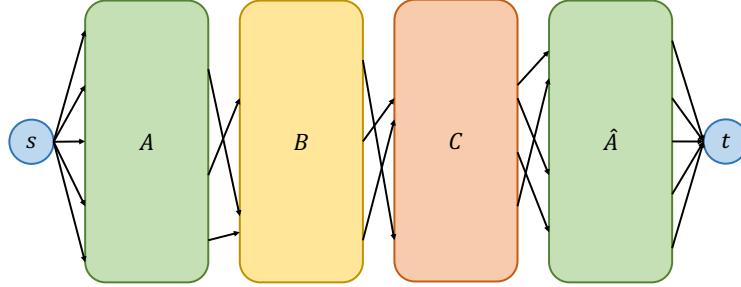


Figure 1: Overall structure of reduction of dynamic path problems. For each $d \in D$, the A, B, C and \hat{A} gadgets encode neighbors of d in A, B , and C . Edges between gadgets encode edges between A, B , and C in the Minimum-Weight 4-Clique instance.

We build gadgets of $O(n)$ vertices, with A, \hat{A} encoding the adjacency information between A and D , B encoding adjacency between B and D ; and C encoding adjacency between C and D . The gadget construction is described in Figure 2. We describe only the B gadget, noting that all gadgets are constructed analogously. The gadget consists of 2 layers B_1, B_2 , each a copy of part $B \subset V(G)$ in the Minimum-Weight 4-Clique instance. For each $b \in B$, we label the copy of b in B_i as $b^{(i)}$, $i = 1, 2$. Note that copies of a in \hat{A}_i are denoted $\hat{a}^{(i)}$. We have thus constructed a layered graph with vertex sets

$$\{s\} \cup A_1 \cup A_2 \cup B_1 \cup B_2 \cup C_1 \cup C_2 \cup \hat{A}_1 \cup \hat{A}_2 \cup \{t\}.$$

We use edges between gadgets to encode the adjacency information between A, B , and C . Formally, for every edge between A, B (resp. B, C and A, C) we insert an edge with weight $\text{wt}(a, b)$ between $(a^{(2)}, b^{(1)})$ (resp. $\text{wt}(b, c)$ between $(b^{(2)}, c^{(1)})$ and $\text{wt}(a, c)$ between $(c^{(2)}, \hat{a}^{(1)})$). Finally, the source s (resp. sink t) are connected to all nodes in the first layer A_1 (resp. last layer \hat{A}_2) with weight 0 edges. For the previous insertions, s cannot reach t and we do not need the stSP information in the reduction.

Next, we will iterate over node pairs (d, a) in decreasing lexicographic order, and use the dynamic stSP instance to compute the minimum weight 4-clique involving (d, a) . We begin with the iteration over a node d in the outer loop. For every neighbor b (resp. c) of d , we connect $(b^{(1)}, b^{(2)})$ (resp. $(c^{(1)}, c^{(2)})$) with an edge of weight $dW + \text{wt}(b, d)$ (resp. weight $dW + \text{wt}(c, d)$). Now, in the inner loop, we iterate over nodes $a \in A$ (also in decreasing order). We connect $(a^{(1)}, a^{(2)}), (\hat{a}^{(1)}, \hat{a}^{(2)})$ with edges of weight $dnW + aW + \frac{\text{wt}(a, d)}{2}$. Note that each gadget is a matching (with parallel edges). Then, we claim that the construction of the graph ensures that if the minimum-weight 4-clique involving (d, a) has weight w^* , then the shortest path from s to t has length,

$$\mathbf{d}(s, t) = 2dnW + (2a + 2d)W + w^*.$$

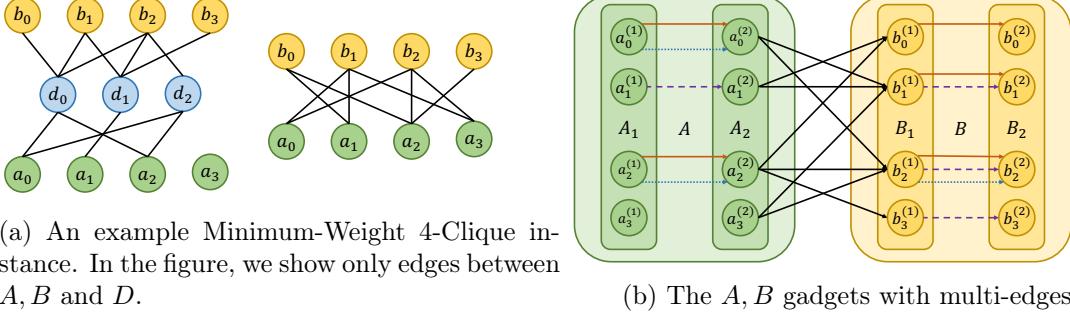


Figure 2: An example of encoding a Minimum-Weight 4-Clique instance into an incremental stSP instance with parallel edges. Solid black edges encode edges $A \times B$ with weight $\text{wt}(a, b)$. Edges within gadgets are inserted in decreasing order of weight. Let $n = 4$. Solid red edges join neighbors of d_2 with weight $2W + \text{wt}(b_j, d_2)$ between $(b_j^{(1)}, b_j^{(2)})$ in B and weight $2nW + kW + \text{wt}(a_k, d_2)$ between $(a_k^{(1)}, a_k^{(2)})$ in A . Dashed purple edges join neighbors of d_1 with weight $W + \text{wt}(b_j, d_1)$ between $(b_j^{(1)}, b_j^{(2)})$ in B and weight $nW + kW + \text{wt}(a_k, d_1)$ between $(a_k^{(1)}, a_k^{(2)})$ in A . Dotted blue edges join neighbors of d_0 with weight $\text{wt}(b_j, d_0)$ between $(b_j^{(1)}, b_j^{(2)})$ in B and weight $kW + \text{wt}(a_k, d_0)$ between $(a_k^{(1)}, a_k^{(2)})$ in A . The (s, t) -distance is maintained after each insertion.

Consider the minimum-weight 4-clique (a, b, c, d) involving (d, a) with weight w^* . Then by our construction there is a path

$$(s, a^{(1)}, a^{(2)}, b^{(1)}, b^{(2)}, c^{(1)}, c^{(2)}, \hat{a}^{(1)}, \hat{a}^{(2)}, t)$$

with length

$$2dnW + (2a + 2d)W + \text{wt}(a, b, c, d) = 2dnW + (2a + 2d)W + w^*.$$

We claim that there is no shorter path. Let P be the shortest path from s to t . Since each gadget is a matching, the edges in the gadgets correspond to some node in the 4-Clique instance and P has the form

$$P = (s, a_1^{(1)}, a_1^{(2)}, b^{(1)}, b^{(2)}, c^{(1)}, c^{(2)}, \hat{a}_2^{(1)}, \hat{a}_2^{(2)}, t).$$

Note that all paths from s to t in the graph must have length at least $2dnW + (2a + 2d)W$. Since we iterate over d in decreasing order, if we take an edge between $(b^{(1)}, b^{(2)})$ or $(c^{(1)}, c^{(2)})$ inserted by a previous iteration $d' > d$ over D , the path has length at least

$$2dnW + (2a + d + d')W \geq 2dnW + (2a + 2d + 1)W > 2dnW + (2a + 2d)W + w^*,$$

so P cannot be a shortest path. Similarly, if we take an edge between $(a^{(1)}, a^{(2)})$ or $(\hat{a}_2^{(1)}, \hat{a}_2^{(2)})$ inserted in a previous iteration $d' > d$, the path has length at least $(2d + 1)nW + (2a + 2d)W$. Finally, since we iterate over a in decreasing order, if $\max\{a_1, a_2\} > a$ then the path has length at least

$$2dnW + (a_1 + a_2 + 2d)W \geq 2dnW + (2a + 2d + 1)W.$$

Thus, if P is a shortest path, we may assume $a = a_1 = a_2$ and all edges in P within gadgets are inserted in the current iteration $d \in D$. Thus, if P is a shortest path, it must correspond to a 4-clique (a, b, c, d) involving (d, a) and furthermore the length of this path is

$$2dnW + (2a + 2d)W + \text{wt}(a, b, c, d).$$

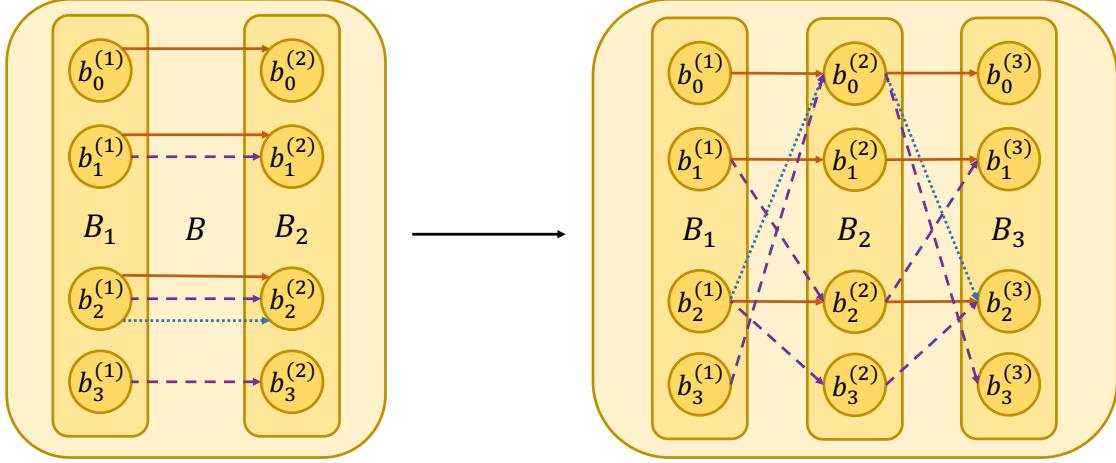


Figure 3: Removing parallel edges from the multi-edge gadget. On the left, solid red edges join neighbors of d_2 with weight $2W + \text{wt}(b_j, d_2)$. Dashed purple edges join neighbors of d_1 with weight $W + \text{wt}(b_j, d_1)$. Dotted blue edges join neighbors of d_0 with weight $\text{wt}(b_j, d_0)$. On the right, we insert an additional middle layer of $n = 4$ nodes. If b_j is a neighbor of d_i , two edges of weight $iW + \frac{\text{wt}(b_j, d_i)}{2}$ join $b_j^{(1)}, b_j^{(3)}$ via $b_{j+i \bmod n}^{(2)}$. Note that each $b_j^{(2)}$ participates in at most one path of weight $< (2i+1)W$, where d_i is the current smallest vertex of d . In particular, all paths with weight in this range are node disjoint, and must enter and leave the gadget from the copies of the same vertex $b \in B$. Furthermore, the simple graph only has n additional nodes.

Since we assumed that w^* is the weight of the minimum-weight 4-clique involving (d, a) , we conclude that $\mathbf{d}(P) = 2dnW + (2a + 2d)W + w^*$.

Thus, for each (d, a) we use the stSP distance to compute,

$$w^* = \mathbf{d}(s, t) - 2dnW - (2a + 2d)W.$$

The minimum-weight 4-clique of G is then computed by taking the minimum w^* over all (d, a) . Thus, if Minimum-Weight 4-Clique requires $n^{4-o(1)}$ time, then incremental stSP with $O(n^2)$ parallel edges requires $n^{4-o(1)}$ time. To conclude, we observe the proof generalizes to the decremental case by running the reduction in reverse and to undirected graphs by adding a sufficiently large weight to each edge. \square

Removing Parallel Edges. Proposition 2.1 has shown that partially dynamic stSP on graphs with parallel edges requires total time $n^{4-o(1)}$. If we naively turn this into a simple graph (say by turning each parallel edge into a two-edge path) or modify parallel edges into a path as in [GVW20], this drastically increases the number of vertices to n^2 . The key idea behind our reduction is that we can in fact modify Proposition 2.1 to a simple graph while only increasing the number of vertices by a constant factor. This idea also crucially underlies the lower bounds shown in later sections.

First, we give some intuition of how to avoid parallel edges. Note that parallel edges only occur within the gadgets A, B, C , and \hat{A} of the reduction. The parallel edges are inserted in the matching between the two layers of a single gadget to encode edge weights between D and vertex sets A, B , and C of G . Specifically, we use parallel edges to ensure that any path must enter and leave a gadget from copies of the *same* vertex. We now describe how to construct these gadgets without the use of parallel edges. Again, we describe only the B gadget, noting that the others are constructed analogously. Consider a 3-layer gadget $B_1 \cup B_2 \cup B_3$, with each copy of $b \in B$ denoted

by $b^{(1)}, b^{(2)}, b^{(3)}$. Fix an iteration over a vertex d in the outer loop. We join $b^{(1)}, b^{(3)}$ with a length two path through $((b+d) \bmod n)^{(2)}$ in the central layer with edges of weight $dW + \frac{\text{wt}(b,d)}{2}$. Thus, there is a path from $b^{(1)}$ to $b^{(3)}$ through the gadget of the form

$$(b^{(1)}, ((b+d) \bmod n)^{(2)}, b^{(3)}) \quad (1)$$

with weight $2dW + \text{wt}(b,d)$. We claim any other path through the B gadget has weight at least $(2d+1)W$. Consider a path,

$$(b_1^{(1)}, b_2^{(2)}, b_3^{(3)}).$$

Since we iterate over d in decreasing order, $b_2 \geq b_1 + d$ (and similarly $b_2 \geq b_3 + d$). If $b_2 > b_1 + d$ or $b_2 > b_3 + d$, then

$$\text{wt}(b_1^{(1)}, b_2^{(2)}) + \text{wt}(b_2^{(2)}, b_3^{(3)}) \geq (2b_2 - b_1 - b_3)W \geq (2d+1)W,$$

since $\text{wt}(b_1^{(1)}, b_2^{(2)}) \geq (b_2 - b_1)W$ and $\text{wt}(b_2^{(2)}, b_3^{(3)}) \geq (b_2 - b_3)W$. Thus, we conclude $b_1 = b_3 = b_2 - d$. Indeed, any shortest path through the B gadget must enter and leave from copies of the *same* $b \in B$ and this path has length $2dW + \text{wt}(b,d)$, just as in the parallel edge case.

Thus, we can construct a B gadget without parallel edges and the rest of the reduction follows as above. Figure 3 illustrates the gadget transformation. We give a full formal proof below.

Proof. We describe the reduction for incremental stSP. For the decremental setting, we can execute the reduction in reverse.

Suppose for contradiction there is an algorithm \mathcal{A} for incremental stSP with $O(n^{4-c})$ total time. Consider a Minimum Weight 4-Clique instance with vertex sets A, B, C, D of size n and integer weights in $[W-1]$. Throughout the reduction, assume that the vertex sets A, B, C, D are indexed from 0 to $n-1$. We design an algorithm for Minimum-Weight 4-Clique with total time $O(n^{4-c})$.

We construct a graph with vertices

$$\{s\} \cup (A_1 \cup A_2 \cup A_3) \cup (B_1 \cup B_2 \cup B_3) \cup (C_1 \cup C_2 \cup C_3) \cup (\hat{A}_1 \cup \hat{A}_2 \cup \hat{A}_3) \cup \{t\}.$$

Each vertex set $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, \hat{A}_1, \hat{A}_2$ and \hat{A}_3 has n vertices. For each $a \in A$, we create a copy $a^{(i)} \in A_i$ and a copy $\hat{a}^{(i)} \in \hat{A}_i$ for $i \in \{1, 2, 3\}$. Similarly, for each $b \in B$ and $c \in C$, we create a copy $b^{(i)} \in B_i$ and $c^{(i)} \in C_i$ for $i \in \{1, 2, 3\}$. Let $W_0 = 100n^2W$. Initially, insert edges $(s, a^{(1)})$ and $(\hat{a}^{(3)}, t)$ with weight W_0 for all $a \in A$. Furthermore, insert edges with weight $W_0 + \text{wt}(a, b)$ between $(a^{(3)}, b^{(1)})$ (resp. weight $W_0 + \text{wt}(b, c)$ between $(b^{(3)}, c^{(1)})$ and weight $W_0 + \text{wt}(c, a)$ between $(c^{(3)}, \hat{a}^{(1)})$) if and only if $(a, b) \in E$ (resp. if $(b, c) \in E$ and $(c, a) \in E$).

This creates a graph with $O(n)$ vertices. We now proceed to more edge insertions of the reduction. Initialize a variable $M \leftarrow \infty$. In the outer loop we iterate over $d_i \in D$ in decreasing order.

1. For each $b_j \in N(d_i)$, insert edges $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$ and $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ with weight $W_0 + i \cdot W + \frac{\text{wt}(b_j, d_i)}{2}$.
2. For each $c_j \in N(d_i)$, insert edges $(c_j^{(1)}, c_{j+i \bmod n}^{(2)})$ and $(c_{j+i \bmod n}^{(2)}, c_j^{(3)})$ with weight $W_0 + i \cdot W + \frac{\text{wt}(c_j, d_i)}{2}$.
3. In the inner loop, we iterate over $a_k \in A$ in decreasing order.
 - (a) If $a_k \in N(d_i)$, insert the following edges of weight $W_0 + i \cdot nW + k \cdot W + \frac{\text{wt}(a_k, d_i)}{4}$:
$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}).$$

(b) Use \mathcal{A} to compute $\mathbf{d}(s, t)$. If $\mathbf{d}(s, t) < 13 \cdot W_0 + 4i \cdot nW + 4(k + i + 1)W$ then update $M \leftarrow \min\{M, \mathbf{d}(s, t) - (13 \cdot W_0 + 4i \cdot nW + 4(k + i)W)\}$.

Finally, return M after all iterations. Clearly, the bottleneck of this algorithm is the total time of algorithm \mathcal{A} , so the running time of this algorithm is $O(n^{4-c})$.

To show the correctness of this algorithm, we first show the following lemma.

Lemma 2.3. *Consider the graph after adding all edges when iterating over d_i and a_k . There is a one-to-one correspondence between:*

1. paths P with length $\mathbf{d}(P) < 13 \cdot W_0 + 4i \cdot nW + 4(k + i + 1)W$
2. 4-cliques (a_k, b, c, d_i) containing a_k, d_i .

Furthermore, each such path P has length $\mathbf{d}(P) = 13 \cdot W_0 + 4i \cdot nW + 4(k + i)W + \mathbf{wt}(a_k, b, c, d_i)$.

Proof. Suppose d_i, a_k are in a 4-clique with vertices b_j, c_ℓ . Then, the following (s, t) -path exists in the graph and has weight $13 \cdot W_0 + 4i \cdot nW + 4(k + i)W + \mathbf{wt}(a_k, b_j, c_\ell, d_i)$:

$$(s, a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)}, b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)}, c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)}, \hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}, t).$$

Note that by our choice of W , $\mathbf{wt}(a_k, b_j, c_\ell, d_i) \leq 3W$ so this path has total weight less than the threshold $13 \cdot W_0 + 4i \cdot nW + 4(k + i + 1)W$.

We verify that this path exists and has the required weight. The edge $(s, a_k^{(1)})$ has weight W_0 . Each edge in the sub-path $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)})$ exists and has weight $W_0 + i \cdot nW + kW + \frac{\mathbf{wt}(a_k, d_i)}{4}$. Then, $(a_k, b_j) \in E$ implies that the edge $(a_k^{(3)}, b_j^{(1)})$ exists with weight $W_0 + \mathbf{wt}(a, c)$. A similar argument shows that the remaining edges in the path exist. Summing over all 13 edges in the path, we conclude the length of the path is $13 \cdot W_0 + 4i \cdot nW + 4(k + i)W + \mathbf{wt}(a_k, b_j, c_\ell, d_i) < 4i \cdot nW + 4(k + i + 1)W$.

Conversely, suppose there is some path of length $< 13 \cdot W_0 + 4i \cdot nW + 4(k + i + 1)W$. Recall that all weights in the 4-Clique instance are non-negative. Since each edge in the graph has weight at least W_0 and we have chosen W_0 so that

$$4i \cdot nW + 4(k + i + 1)W < W_0,$$

any (s, t) path with the required length has at most 13 edges. Since we have a 14-layered graph, with edges only between adjacent layers, any (s, t) -path with the required length must have the following structure:

$$P = (s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, b_{i_4}^{(1)}, b_{i_5}^{(2)}, b_{i_6}^{(3)}, c_{i_7}^{(1)}, c_{i_8}^{(2)}, c_{i_9}^{(3)}, \hat{a}_{i_{10}}^{(1)}, \hat{a}_{i_{11}}^{(2)}, \hat{a}_{i_{12}}^{(3)}, t).$$

At the time of the query in the i -th iteration over D and k -th iteration over A , the only edges in $(A_1 \times A_2) \cup (A_2 \times A_3) \cup (\hat{A}_1 \times \hat{A}_2) \cup (\hat{A}_2 \times \hat{A}_3)$ with weight less than $W_0 + i \cdot nW + (k + 1)W$ are $(a_k^{(1)}, a_{k+i \bmod n}^{(2)})$, $(a_{k+i \bmod n}^{(2)}, a_k^{(3)})$, $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)})$, and $(\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$ for $a_k \in N(d_i)$. If P uses these edges, then $k = i_1 = i_3 = i_{10} = i_{12}$. Similarly, the only edges in $B_1 \times B_2$ and $B_2 \times B_3$ with weight less than $W_0 + (i + 1) \cdot W$ are $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$, $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ for $b_j \in N(d_i)$. If P uses these edges, then $i_4 = i_6 = j$ for some $b_j \in N(d_i)$, since all paths of this form are vertex disjoint. Similarly, the only edges in $C_1 \times C_2$ and $C_2 \times C_3$ with weight less than $W_0 + (i + 1) \cdot W$ are $(c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)})$, $(c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)})$ for $c_\ell \in N(d_i)$. By a similar argument, if P uses these edges then $i_7 = i_9 = \ell$. If any of the above edges do not exist in P , the (s, t) -path in question has length at least $13 \cdot W_0 + 4i \cdot nW + 4(k + i + 1)W$, contradicting the assumption.

Finally, it remains to show a_k, b_j, c_ℓ form a triangle. Since the edge $(a_k^{(3)}, b_j^{(1)})$ exists, $(a_k, b_j) \in E$. A similar argument shows that $(a_k, c_\ell), (b_j, c_\ell) \in E$.

Thus, we have shown paths of length less than $13 \cdot W_0 + 4i \cdot nW + 4(k+i+1)W$ correspond to 4-cliques containing d_i, a_k . Furthermore, any such path has weight $4i \cdot nW + 4(k+i)W + \text{wt}(a_k, b, c, d_i)$ for b, c in the clique. \square

From Lemma 2.3, at the time of computing $\mathbf{d}(s, t)$, each path P with length $\mathbf{d}(P) < 13 \cdot W_0 + 4i \cdot nW + 4(k+i+1)W$ corresponds to a 4-clique (a_k, b, c, d_i) in the input graph containing d_i, a_k and furthermore

$$\mathbf{d}(P) = 13 \cdot W_0 + 4i \cdot nW + 4(k+i)W + \text{wt}(a_k, b, c, d_i),$$

where $\text{wt}(a_k, b, c, d_i)$ denotes the weight of the 4-clique. Thus, if $\mathbf{d}(s, t) < 13 \cdot W_0 + 4i \cdot nW + 4(k+i+1)W$ (which must happen if a_k and d_i are in any 4-clique), then $\mathbf{d}(s, t) - (13 \cdot W_0 + 4i \cdot nW + 4(k+i)W)$ is the minimum weight 4-clique containing a_k, d_i . Since we iterate over all a_k, d_i , the above procedure solves the Minimum Weight 4-Clique problem, contradicting the Minimum-Weight 4-Clique hypothesis. \square

3 Node-Weighted Shortest Paths

In this section, we give our algorithms and lower bounds for partially dynamic shortest paths on *node-weighted* graphs. We start the section with some simple algorithms that use fast matrix multiplication (FMM) to beat the naive algorithm recomputing shortest paths with each update. Next, we show that using FMM is in fact necessary by establishing a conditional lower bound that rules out any nontrivial combinatorial partially dynamic algorithms for node-weighted shortest paths. Finally, we show some nontrivial lower bounds that hold even for algebraic algorithms.

3.1 Faster Algebraic Algorithms for Partially Dynamic nw-SSSP

We prove the following result that utilizes fast matrix multiplication.

Proposition 3.1. *There is an algorithm solving incremental nw-SSSP in $\tilde{O}(n^{2-\gamma/4})$ amortized time where γ is the solution to the equation $\omega(1, 1, 1 + \gamma) = 3 - \gamma$.*

To prove Proposition 3.1, we require the static node-weighted APSP algorithm of Yuster [Yus09].

Lemma 3.2 (Theorem 1.1 of [Yus09]). *There is an algorithm computing nw-APSP in $O(n^{3-\gamma/2})$ -time, where γ is the solution to the equation $\omega(1, 1, 1 + \gamma) = 3 - \gamma$.*

With this in hand, we prove Proposition 3.1. At a high level, we handle the updates to the graph in batches. At the start of a batch, we compute nw-APSP on the current graph. Then, with each edge insertion, we construct a layered graph with $O(n^{1+t})$ edges, by using either the inserted edges or an edge encoding the shortest path computed by the nw-APSP instance, so that computing SSSP on the layered graph finds the shortest paths in the updated graph (since any paths not using edges inserted in this batch will use edges computed by the nw-APSP instance).

Proof of Proposition 3.1. Let t be a parameter to be fixed later. We will process updates in batches of size n^t . At the start of each batch, compute nw-APSP on the current graph using Lemma 3.2. Let $\mathbf{d}_0(u, v)$ denote the distances computed at the start of each batch. Then, after k updates $\{(u_i, v_i)\}_{i=1}^k$, construct the following 4-layer graph with vertices

$$\{s\} \cup \{u_i\}_{i=1}^k \cup \{v_i\}_{i=1}^k \cup V,$$

with edge weights $\text{wt}(s, u_i) = \mathbf{d}_0(s, u_i)$ for all $i \in [k]$, $\text{wt}(u_i, v_i) = 0$ (since the distances \mathbf{d}_0 include the weights on endpoints) for all $i \in [k]$, $\text{wt}(v_i, u_j) = \mathbf{d}_0(v_i, u_j)$ for all $i, j \in [k]$ and $\text{wt}(v_i, v) = \mathbf{d}_0(v_i, v)$ for all $i \in [k]$ and $v \in V$. We also insert edge weights $\text{wt}(s, v) = \mathbf{d}_0(s, v)$ for all $v \in V$. On this graph with $O(n^{1+t})$ edges, we compute SSSP in $\tilde{O}(n^{1+t})$ -time. The amortized running time is

$$\tilde{O}\left(n^{3-\gamma/2-t} + n^{1+t}\right),$$

where γ is the solution to the equation $\omega(1, 1, 1 + \gamma) = 3 - \gamma$. Optimizing for parameter $t = 1 - \frac{\gamma}{4}$ gives an amortized running time of

$$\tilde{O}\left(n^{2-\frac{\gamma}{4}}\right).$$

We claim this computes incremental nw-SSSP. First, since every distance in the constructed graph is realized by a path in the input graph, it suffices to show that the shortest path is encoded into the constructed graph. After i updates in the batch, let $(s = w_0, w_1, \dots, w_d = v)$ be a shortest path from s to v . If none of these edges are inserted during the batch, then they must have been present during the nw-APSP computation, and therefore the constructed graph has a path of length $\mathbf{d}_0(s, v)$. Otherwise, let $\{(w_{i_j}, w_{i_j+1})\}_j$ denote the edges inserted in this batch. Then, these edges are present in our graph, while the remaining portions of the path are encoded in the graph using the nw-APSP solution \mathbf{d}_0 , so a path with length at most the distance exists in the constructed graph. \square

Following a similar technique, we obtain an incremental algorithm for nw-stSP.

Proposition 3.3. *There is an algorithm solving incremental nw-stSP in $\tilde{O}(n^{2-\gamma/3})$ amortized time where γ is the solution to the equation $\omega(1, 1, 1 + \gamma) = 3 - \gamma$.*

Proof. Let t be a parameter to be fixed later. We will process updates in batches of size n^t . At the start of each batch, compute nw-APSP on the current graph using Lemma 3.2. Let $\mathbf{d}_0(u, v)$ denote the distances computed at the start of each batch. Then, after k updates $\{(u_i, v_i)\}_{i=1}^k$, construct the following 4-layer graph with vertices

$$\{s\} \cup \{u_i\}_{i=1}^k \cup \{v_i\}_{i=1}^k \cup \{t\},$$

with edge weights $\text{wt}(s, u_i) = \mathbf{d}_0(s, u_i)$ for all $i \in [k]$, $\text{wt}(u_i, v_i) = 0$ for all $i \in [k]$ (again the distance \mathbf{d}_0 include endpoints), $\text{wt}(v_i, u_j) = \mathbf{d}_0(v_i, u_j)$ for all $i, j \in [k]$ and $\text{wt}(v_i, t) = \mathbf{d}_0(v_i, t)$ for all $i \in [k]$. We additionally insert edge weight $\text{wt}(s, t) = \mathbf{d}_0(s, t)$. On this graph with $O(n^{2t})$ edges, we compute stSP in $\tilde{O}(n^{2t})$ -time. The amortized running time is

$$\tilde{O}\left(n^{3-\gamma/2-t} + n^{2t}\right),$$

where γ is the solution to the equation $\omega(1, 1, 1 + \gamma) = 3 - \gamma$. Optimizing for parameter $t = 1 - \frac{\gamma}{6}$ gives an amortized running time of

$$\tilde{O}\left(n^{2-\frac{\gamma}{3}}\right).$$

We claim this computes incremental nw-stSP. First, since every distance in the constructed graph is realized by a path in the input graph, it suffices to show that the shortest path is encoded into the constructed graph. After i updates in the batch, let $(s = w_0, w_1, \dots, w_d = t)$ be a shortest path from s to t . If none of these edges are inserted during the batch, then they must have been present during the nw-APSP computation, and therefore the constructed graph has a path of length $\mathbf{d}_0(s, t)$. Otherwise, let $\{(w_{i_j}, w_{i_j+1})\}_j$ denote the edges inserted in this batch. Then, these edges are present in our graph, while the remaining portions of the path are encoded in the graph using the nw-APSP solution \mathbf{d}_0 , so a path with length at most the distance exists in the constructed graph. \square

3.2 Conditional Lower Bounds for Partially Dynamic nw-SSSP

We have exploited FMM to design algorithms faster than trivial recomputation. We show that using FMM is in fact necessary: any combinatorial algorithm that can solve partially dynamic nw-stSP requires total time $n^{4-o(1)}$. However, our previous graph simplification technique (Figure 3) no longer suffices to produce vertex-disjoint paths. Instead, we construct a new gadget for each multi-edge vertex set that ensures all short paths are vertex disjoint. The resulting simple graph is node-weighted, and the number of vertices increases only by a constant factor.

Before stating the theorem, we state the relevant hardness conjecture.

Hypothesis 3.4 (Combinatorial k -Clique hypothesis). *There is no $O(n^{k-\varepsilon})$ combinatorial algorithm for k -Clique Detection on n -node graphs, for any $\varepsilon > 0$.*

Theorem 3.5. *Under the Combinatorial 4-Clique hypothesis, any combinatorial algorithm computing incremental/decremental nw-stSP on undirected graphs requires $n^{4-o(1)}$ total time.*

Proof. For simplicity, we describe only the reduction for incremental nw-stSP. For the decremental setting, we can execute the reduction in reverse.

Suppose for contradiction there is a combinatorial algorithm \mathcal{A} for partially dynamic nw-stSP with total time $O(n^{4-c})$ for some $c > 0$. Consider a 4-Clique instance with vertex sets A, B, C, D of size n . Throughout the reduction, assume that the vertex sets A, B, C, D are indexed from 0 to $n - 1$. We design an algorithm with total time $O(n^{4-c})$, thus contradicting the Combinatorial 4-Clique hypothesis. In the following, let $W = 100n$.

Vertex Set Gadget Before describing the construction of the full graph, we describe the key gadget in the reduction. First, consider vertex set B , and we construct the gadget $G_R(B)$. The gadget construction is shown in Figure 4.

In our previous lower bounds, we used a 3-layer graph to ensure that at each stage, there are *vertex disjoint* paths between $b_j^{(1)}, b_j^{(3)}$ for all $b_j \in N(d_i)$. This in turn ensured that the *same* vertex b_j was adjacent to both a_k, c_ℓ (and d_i) in the 4-Clique instance. Crucially, we used edge weights to ensure that through each middle vertex, there was only one path with edges of weight between $[W_0 + i \cdot W, W_0 + (i + 1) \cdot W]$. In a node-weighted graph, we describe the following modification.

$G_R(B)$ consists of 4 layers of vertices

$$B_1 \cup B_2 \cup B_3 \cup B_4,$$

where $|B_1| = |B_2| = |B_4| = n$ and $|B_3| = 2n$.⁷ For each $b \in B$, we create a copy $b^{(i)} \in B_i$ for $i \in \{1, 2, 4\}$. We index $B_3 = \{b_0^{(3)}, \dots, b_{2n-1}^{(3)}\}$ and assign weights as follows,

$$\begin{aligned} \text{wt}\left(b_j^{(2)}\right) &= \text{wt}\left(b_j^{(4)}\right) = W^4 + j \cdot W^3 \\ \text{wt}\left(b_j^{(3)}\right) &= W^4 + (4n - 2j) \cdot W^3 \\ \text{wt}\left(b_j^{(1)}\right) &= W^4 + j \cdot W \end{aligned}$$

We similarly define the vertex sets $G_R(C) = \bigcup_{i=1}^4 C_i$, $G_R(A) = \bigcup_{i=1}^4 A_i$ and $G_R(\hat{A}) = \bigcup_{i=1}^4 \hat{A}_i$ with the same vertex set sizes. The weights are defined as follows,

$$\text{wt}(s) = \text{wt}(t) = W^4,$$

⁷For the current reduction, it suffices to construct a 3-layer gadget $(B_2 \cup B_3 \cup B_4)$, but we require the 4-layer gadget for a later reduction (Theorem 3.13).

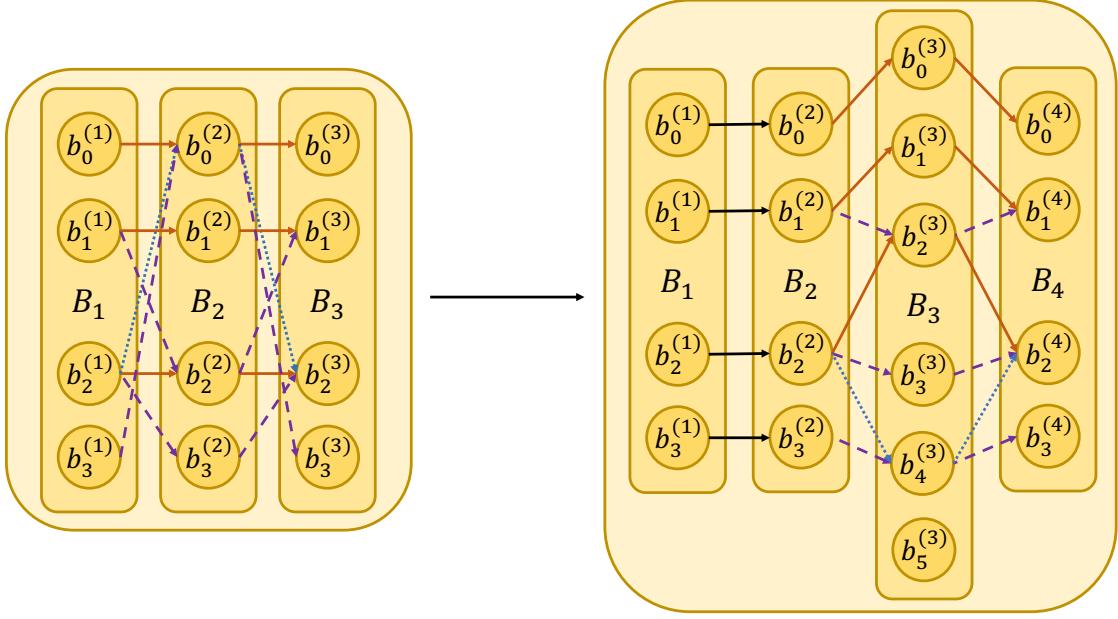


Figure 4: Comparison of node-weighted gadget with an edge-weighted gadget. The gadget on the left is the edge-weighted gadget from Theorem 1.1. We now create a 4-layer node-weighted graph where the weight of $b_j^{(i)}$ vertices in B_1, B_2, B_4 are $j \cdot W$ for $i \in \{1\}$ and $j \cdot W^3$ for $i \in \{2, 4\}$ and the weight of $b_j^{(3)}$ vertices are $(4n - 2j) \cdot W^3$ for some large constant W . While the edges are no longer weighted, we format the edges as before to illustrate corresponding paths. Note that each $b_j^{(3)}$ participates in at most one path of weight $< (4n - 2i + 1) \cdot W^3$ where d_i is the current smallest vertex of d . As in Figure 3, all paths with weight in this range are vertex disjoint, and must enter and leave the gadget from copies of the same vertex $b \in B$. Furthermore, the simple graph only has $4n$ additional vertices.

$$\begin{aligned}
\text{wt}(c_\ell^{(1)}) &= W^4 + \ell, \\
\text{wt}(a_k^{(1)}) &= \text{wt}(\hat{a}_k^{(1)}) = W^4 + k \cdot W^2, \\
\text{wt}(c_i^{(2)}) &= \text{wt}(a_i^{(2)}) = \text{wt}(\hat{a}_i^{(2)}) = W^4 + i \cdot W^3, \\
\text{wt}(c_i^{(3)}) &= \text{wt}(a_i^{(3)}) = \text{wt}(\hat{a}_i^{(3)}) = W^4 + (4n - 2i) \cdot W^3, \\
\text{wt}(c_i^{(4)}) &= \text{wt}(a_i^{(4)}) = \text{wt}(\hat{a}_i^{(4)}) = W^4 + i \cdot W^3.
\end{aligned}$$

Note that all vertex weights are nonnegative.

Lower Bound Construction We construct a graph with the following vertices:

$$\{s\} \cup G_R(A) \cup G_R(B) \cup G_R(C) \cup G_R(\hat{A}) \cup \{t\},$$

with vertex weights described above.

Initially, insert edges $(s, a^{(1)})$ and $(\hat{a}^{(4)}, t)$ for all $a \in A$. Furthermore, insert edges between $(a^{(4)}, b^{(1)})$ (resp. between $(b^{(4)}, c^{(1)})$, between $(c^{(4)}, \hat{a}^{(1)})$) if and only if $(a, b) \in E$ (resp. if $(b, c) \in E$, $(c, a) \in E$). Insert edges $(a^{(1)}, a^{(2)}), (b^{(1)}, b^{(2)}), (c^{(1)}, c^{(2)}), (\hat{a}^{(1)}, \hat{a}^{(2)})$ for all a, b, c .

We now proceed to more edge insertions of the reduction. In the outer loop we iterate over $d_i \in D$ in increasing order.

1. For each $b_j \in N(d_i)$, insert edges, $(b_j^{(2)}, b_{j+i}^{(3)})$ and $(b_{j+i}^{(3)}, b_j^{(4)})$.
2. For each $c_j \in N(d_i)$, insert edges $(c_j^{(2)}, c_{j+i}^{(3)})$ and $(c_{j+i}^{(3)}, c_j^{(4)})$.
3. In the inner loop, we iterate over $a_k \in A$ in decreasing order.
 - (a) If $a_k \in N(d_i)$, insert the following edges:

$$(a_k^{(2)}, a_{k+i}^{(3)}), (a_{k+i}^{(3)}, a_k^{(4)}), (\hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}), (\hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)}).$$

- (b) Query $\mathbf{d}(s, t)$ from \mathcal{A} . If $\mathbf{d}(s, t) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$ return TRUE.

Finally, after all iterations, return FALSE if no query satisfied the required distance constraint.

Clearly, the bottleneck of this algorithm is running \mathcal{A} , so the total running time is $O(n^{4-c})$. Here, emphasize that D is iterated over in increasing order while A is iterated over in decreasing order.

To show correctness, we first prove the following lemma.

Lemma 3.6. *Consider the query when iterating over d_i and a_k . Then, there is a one-to-one correspondence between:*

1. paths P with $\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$
2. 4-cliques (a_k, b, c, d_i) containing a_k, d_i .

Furthermore, each such path P has length

$$\mathbf{d}(P) = 18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + b \cdot W + c.$$

Proof. Suppose d_i, a_k are in a 4-clique with vertices b_j, c_ℓ . Then, the following (s, t) -path exists:

$$(s, a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)}, b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)}, c_\ell^{(1)}, c_\ell^{(2)}, c_{\ell+i}^{(3)}, c_\ell^{(4)}, \hat{a}_k^{(1)}, \hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)}, t).$$

Furthermore, this path has length

$$18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + j \cdot W + \ell.$$

Note that by our choice of W , $j \cdot W + \ell < W^2$, so this path satisfies the desired constraint. We verify that this path exists and has the required length. Recall that the edge $(s, a_k^{(1)})$ exists. Since $a_k \in N(d_i)$, each edge in the sub-path $(a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)})$ exists and the vertices have combined weight

$$4 \cdot W^4 + (4n - 2(k + i) + 2k) \cdot W^3 + k \cdot W^2 = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + k \cdot W^2.$$

Next, $(a_k, b_j) \in E$ implies that the edge $(a_k^{(4)}, b_j^{(1)})$ exists. Since $b_j \in N(d_i)$, each edge in the sub-path $(b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)})$ exists and the vertices have combined weight

$$4 \cdot W^4 + (4n - 2(j + i) + 2j) \cdot W^3 + j \cdot W = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + j \cdot W.$$

Then, $(b_j, c_\ell) \in E$ implies that the edge $(b_j^{(4)}, c_\ell^{(1)})$ exists. Since $c_\ell \in N(d_i)$, each edge in the sub-path $(c_\ell^{(1)}, c_\ell^{(2)}, c_{\ell+i}^{(3)}, c_\ell^{(4)})$ exists and the vertices have combined weight

$$4 \cdot W^4 + (4n - 2(\ell + i) + 2\ell) \cdot W^3 + \ell = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + \ell.$$

Then, $(c_\ell, a) \in E$ implies that the edge $(c_\ell^{(4)}, \hat{a}_k^{(1)})$ exists. Finally, as $a_k \in N(d_i)$, each edge in the sub-path $(\hat{a}_k^{(1)}, \hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)})$ exists and the vertices have combined weight

$$4 \cdot W^4 + (4n - 2(k + i) + 2k) \cdot W^3 + k \cdot W^2 = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + k \cdot W^2.$$

We conclude by observing that edge $(\hat{a}_k^{(4)}, t)$ exists. Summing over all vertices in the path, we conclude the length of the path is

$$18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + j \cdot W + \ell.$$

Conversely, suppose there is some path P of length

$$\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2.$$

Since all vertices have weight at least W^4 and

$$(16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2 < W^4,$$

any (s, t) -path P satisfying the length constraint has at most 18 vertices. Since we have a layered graph and edges only exist between adjacent layers, any such path P must have form

$$P = \left(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, a_{i_4}^{(4)}, b_{i_5}^{(1)}, b_{i_6}^{(2)}, b_{i_7}^{(3)}, b_{i_8}^{(4)}, c_{i_9}^{(1)}, c_{i_{10}}^{(2)}, c_{i_{11}}^{(3)}, c_{i_{12}}^{(4)}, \hat{a}_{i_{13}}^{(1)}, \hat{a}_{i_{14}}^{(2)}, \hat{a}_{i_{15}}^{(3)}, \hat{a}_{i_{16}}^{(4)}, t \right).$$

At the time of the query in the i -th iteration over D and k -th iteration over A in the inner-loop, the minimum weight of any path from A_2 to A_4 is $3 \cdot W^4 + (4n - 2i) \cdot W^3$, and this is attained only by paths of the form

$$\left(a_{k'}^{(2)}, a_{k'+i}^{(3)}, a_{k'}^{(4)} \right) \quad (2)$$

for $k' \geq k$ since we iterate over A in decreasing order. In fact, any other path from A_2, A_4 has weight at least $3 \cdot W^4 + (4n - 2i + 1) \cdot W^3$. This is because $a_{i_3}^{(3)}$ is adjacent only to vertices $a_{i_2}^{(2)}, a_{i_4}^{(4)}$ with $i_2, i_4 \geq i_3 - i$. Thus, if either i_2 or i_4 is $> i_3 - i$, the three vertices have weight at least

$$3 \cdot W^4 + (4n - 2i_3 + 2i_3 - 2i + 1) \cdot W^3 = 3 \cdot W^4 + (4n - 2i + 1) \cdot W^3.$$

By a similar argument, the minimum weight of any path from B_2 to B_4 , C_2 to C_4 , \hat{A}_2 to \hat{A}_4 is $3 \cdot W^4 + (4n - 2i) \cdot W^3$ and any path not satisfying the form of Eq. (2) has weight at least $3 \cdot W^4 + (4n - 2i + 1) \cdot W^3$. Since

$$\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2 < 18 \cdot W^4 + (16n - 7i) \cdot W^3,$$

we conclude

$$\begin{aligned} 0 &= i_2 - i_4 = i_6 - i_8 = i_{10} - i_{12} = i_{14} - i_{16}, \\ i &= i_3 - i_2 = i_7 - i_6 = i_{11} - i_{10} = i_{15} - i_{14}. \end{aligned}$$

Next, observe that for all $x \in \{a, b, c, \hat{a}\}$ and $X \in \{A, B, C, \hat{A}\}$, $x_j^{(2)}$ has a unique neighbor in X_1 , namely $x_j^{(1)}$. In particular, there are indices k', j', ℓ', \hat{k}' such that

$$P = \left(s, a_{k'}^{(1)}, a_{k'}^{(2)}, a_{k'+i}^{(3)}, a_{k'}^{(4)}, b_{j'}^{(1)}, b_{j'}^{(2)}, b_{j'+i}^{(3)}, b_{j'}^{(4)}, c_{\ell'}^{(1)}, c_{\ell'}^{(2)}, c_{\ell'+i}^{(3)}, c_{\ell'}^{(4)}, c_{\ell'}^{(5)}, \hat{a}_{k'}^{(1)}, \hat{a}_{k'}^{(2)}, \hat{a}_{k'+i}^{(3)}, \hat{a}_{k'}^{(4)}, t \right).$$

Next, we claim $k = k' = \hat{k}'$. Because we enumerate A in decreasing order, we know that $k', \hat{k}' \geq k$. If $k' > k$, then

$$\text{wt}\left(a_{k'}^{(1)}\right) > k \cdot W^2.$$

Similarly, if $\hat{k}' > k$, then

$$\text{wt}\left(\hat{a}_{\hat{k}'}^{(1)}\right) > k \cdot W^2.$$

If either condition holds, then $\mathbf{d}(P) \geq 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$, a contradiction. Thus, $k = k' = \hat{k}'$.

In particular, we have $a_k, b_{j'}, c_{\ell'} \in N(d_i)$ and furthermore $a_k, b_{j'}, c_{\ell'}$ form a triangle as edges $(a_k^{(4)}, b_{j'}^{(1)}), (b_{j'}^{(4)}, c_{\ell'}^{(1)}),$ and $(c_{\ell'}^{(4)}, \hat{a}_k^{(1)})$ exist, so that a_k, d_i form a 4-clique with $b_{j'}, c_{\ell'}$. Note that this path P has length

$$\mathbf{d}(P) = 18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + j' \cdot W + \ell'.$$

as desired. \square

From Lemma 3.6, at the time of the query, path P with length $\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$ corresponds to 4-clique (a_k, b, c, d_i) in the input graph containing d_i, a_k . Since we iterate over all a_k, d_i , the above procedure solves the 4-Clique Detection instance in total time $O(n^{4-c})$, contradicting the 4-Clique hypothesis. \square

3.2.1 Non-Combinatorial Lower Bounds

Using an essentially identical reduction, we show that under the OMv3 hypothesis, any algorithm computing partially dynamic nw-stSP with polynomial preprocessing requires $n^{\omega+1-o(1)}$ total update and query time. The OMv3 hypothesis, introduced by [GVW20] is a generalization of the OMv hypothesis [HKNS15].

OMv3 hypothesis. Beyond its inability to deal with non-combinatorial algorithms, the Combinatorial k -Clique hypothesis (Hypothesis 3.4) has one further weakness: It does not rule out algorithms with polynomially large preprocessing time. Conceptually, this may not be a large issue in the incremental setting (since in the preprocessing phase the algorithm only has access to an empty graph), it is not unreasonable to expect a decremental algorithm, given sufficient preprocessing time, to handle updates and queries efficiently. In the $k = 3$ case, the OMv hypothesis [HKNS15] strengthens conditional lower bounds based on the combinatorial BMM hypothesis [AV14], ruling out arbitrary algorithms with polynomial preprocessing. The OMv3 hypothesis was introduced as a natural generalization of the OMv hypothesis, ruling out arbitrary algorithms with polynomial preprocessing time [GVW20].

Definition 3.7 (OMv3 Problem [GVW20]). During preprocessing, the algorithm is given an $n \times n$ boolean matrix A . During the query phase, the algorithm receives n queries. Each query consist of three n -dimensional boolean vectors $\vec{u}, \vec{v}, \vec{w}$ and must output

$$\bigvee_{i,j,k} (\vec{u}_i \wedge \vec{v}_j \wedge \vec{w}_k \wedge A_{ij} \wedge A_{jk} \wedge A_{ki})$$

before the next query is received.

Naturally, each OMv3 query can be answered in $O(n^\omega)$ time using fast matrix multiplication. The OMv3 hypothesis states that this is essentially optimal.

Hypothesis 1.4 (OMv3 hypothesis [GVW20]). *There is no algorithm with polynomial preprocessing time and total update and query time $O(n^{\omega+1-\varepsilon})$ solving the OMv3 problem, for any $\varepsilon > 0$.*

Theorem 3.8. *Under the OMv3 hypothesis, any algorithm computing incremental/decremental nw-stSP on undirected graphs with polynomial preprocessing time requires $n^{\omega+1-o(1)}$ total update and query time.*

Proof. Suppose for contradiction there is an incremental stSP algorithm \mathcal{A} with polynomial preprocessing time and total update and query time $O(n^{\omega+1-c})$ for some $c > 0$. We will describe appropriate modifications for the decremental case where necessary. We design an efficient algorithm for OMv3.

In the preprocessing phase we receive a Boolean matrix A . We construct a graph with the same vertex set as in Theorem 3.5:

$$\{s\} \cup G_R(A) \cup G_R(B) \cup G_R(C) \cup \hat{G}_R(A) \cup \{t\}.$$

As before, insert edges $(s, a^{(1)})$ and $(\hat{a}^{(4)}, t)$ for all $a \in A$. Furthermore, insert edges between $(a^{(4)}, b^{(1)})$ (resp. $(b^{(4)}, c^{(1)})$, $(c^{(4)}, \hat{a}^{(1)})$) if and only if $A[a, b] = 1$ (resp. $A[b, c] = 1$, $A[a, c] = 1$). Insert edges $(a^{(1)}, a^{(2)})$, $(b^{(1)}, b^{(2)})$, $(c^{(1)}, c^{(2)})$, $(\hat{a}^{(1)}, \hat{a}^{(2)})$ for all a, b, c .

In the decremental case, we begin with a graph that additionally contains the following edges.

1. For all i, j , insert edges $(b_j^{(2)}, b_{j+i}^{(3)})$ and $(b_{j+i}^{(3)}, b_j^{(4)})$.
2. For all i, j , insert edges $(c_j^{(2)}, c_{j+i}^{(3)})$ and $(c_{j+i}^{(3)}, c_j^{(4)})$.
3. For all i, k , insert the following edges:

$$(a_k^{(2)}, a_{k+i}^{(3)}), (a_{k+i}^{(3)}, a_k^{(4)}), (\hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}), (\hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)}).$$

We then run the preprocessing step of \mathcal{A} on this $O(n)$ -vertex graph.

We now proceed to the dynamic phase of the reduction. Throughout, we assume that both queries and coordinates are indexed starting at 0 (just as vertices are indexed from 0 in Theorem 3.5). Recall that in the vertex gadgets $G_R(A), G_R(B)$ the third layer A_3, B_3 has $2n - 1$ vertices. Suppose we have received the i -th query, $\vec{u}_i, \vec{v}_i, \vec{w}_i$.

1. If $\vec{v}_i[j] = 1$, insert edges $(b_j^{(2)}, b_{j+i}^{(3)})$ and $(b_{j+i}^{(3)}, b_j^{(4)})$. In the decremental case, remove edges $(b_j^{(2)}, b_{j+n-i}^{(3)})$ and $(b_{j+n-i}^{(3)}, b_j^{(4)})$ if $\vec{v}_i[j] = 0$.
2. If $\vec{w}_i[j] = 1$, insert edges $(c_j^{(2)}, c_{j+i}^{(3)})$ and $(c_{j+i}^{(3)}, c_j^{(4)})$. In the decremental case, remove edges $(c_j^{(2)}, c_{j+n-i}^{(3)})$ and $(c_{j+n-i}^{(3)}, c_j^{(4)})$ if $\vec{w}_i[j] = 0$.
3. Now, we iterate over the coordinates k of vector \vec{v}_i in decreasing order. In the decremental case, we instead iterate over the coordinates k in increasing order.
 - (a) If $\vec{u}_i[k] = 1$, insert the edges

$$(a_k^{(2)}, a_{k+i}^{(3)}), (a_{k+i}^{(3)}, a_k^{(4)}), (\hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}), (\hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)}).$$

In the decremental case, remove the edges

$$(a_k^{(2)}, a_{k+n-i}^{(3)}), (a_{k+n-i}^{(3)}, a_k^{(4)}), (\hat{a}_k^{(2)}, \hat{a}_{k+n-i}^{(3)}), (\hat{a}_{k+n-i}^{(3)}, \hat{a}_k^{(4)}),$$

whenever $\vec{u}_i[k] = 0$.

- (b) Query $\mathbf{d}(s, t)$ from \mathcal{A} . If $\mathbf{d}(s, t) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$ return TRUE.
- (c) In the decremental case, we additionally remove all remaining edges

$$\left(a_k^{(2)}, a_{k+n-i}^{(3)} \right), \left(a_{k+n-i}^{(3)}, a_k^{(4)} \right), \left(\hat{a}_k^{(2)}, \hat{a}_{k+n-i}^{(3)} \right), \left(\hat{a}_{k+n-i}^{(3)}, \hat{a}_k^{(4)} \right).$$

Otherwise, return FALSE if none of the queries satisfied the required distance constraint.

- 4. In the decremental case, we additionally remove all remaining edges

$$\left(b_j^{(2)}, b_{j+n-i}^{(3)} \right), \left(b_{j+n-i}^{(3)}, b_j^{(4)} \right), \left(c_j^{(2)}, c_{j+n-i}^{(3)} \right), \left(c_{j+n-i}^{(3)}, c_j^{(4)} \right).$$

Clearly, the bottleneck of this algorithm is running \mathcal{A} , so the total time over all n queries is $O(n^{\omega+1-c})$.

Note that at the query after the i -th query and the k -th iteration of the inner loop, we have constructed the same graph as in Theorem 3.5. In particular, we have the following lemma, whose proof is deferred to Appendix A.1 as it is identical to the proof of Lemma 3.6.

The following lemma shows that stSP queries compute OMv3 correctly.

Lemma 3.9. *Consider the query $(\vec{u}_i, \vec{v}_i, \vec{w}_i)$ and the query after updating the k -th coordinate of \vec{u}_i . The following are equivalent:*

1. $d(s, t) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2$;
2. $\bigvee_{j, \ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k, j] \wedge A[j, \ell] \wedge A[\ell, k]) = \text{TRUE}$.

In particular, Lemma 3.9 shows that this query returns true if and only if

$$\bigvee_{j, \ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k, j] \wedge A[j, \ell] \wedge A[\ell, k]).$$

Thus, if the above is satisfied for any k , we return TRUE. Otherwise, we return FALSE, answering the query correctly in either case. Therefore, using \mathcal{A} we have obtained an algorithm for the OMv3 instance with polynomial preprocessing and total update and query time $O(n^{\omega+1-c})$, contradicting the OMv3 hypothesis. \square

We prove that the stSP instance computes OMv3.

Proof of Lemma 3.9. We prove the equivalence in the incremental case, leaving the modifications for the decremental case to the end of the proof. Suppose 2 is true and let j, ℓ index one clause that is TRUE. Then, the following (s, t) -path exists

$$(s, a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)}, b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)}, c_\ell^{(1)}, c_\ell^{(2)}, c_{\ell+i}^{(3)}, c_\ell^{(4)}, \hat{a}_k^{(1)}, \hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)}, t).$$

and has length

$$18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + j \cdot W + \ell.$$

Note that by our choice of W , $j \cdot W + \ell < W^2$ so this path satisfies the desired constraint. We verify that this path exists and has the required length. Recall that the edge $(s, a_k^{(1)})$ exists. Since $\vec{u}_i[k] = 1$, each edge in the sub-path $(a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)})$ exists and the vertices have combined weight,

$$4 \cdot W^4 + (4n - 2(k + i) + 2k) \cdot W^3 + k \cdot W^2 = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + k \cdot W^2$$

Next, $A[k, j] = 1$ implies that the edge $(a_k^{(4)}, b_j^{(1)})$ exists. Since $\vec{v}_i[k] = 1$, each edge in the sub-path $(b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)})$ exists and the vertices have combined weight,

$$4 \cdot W^4 + (4n - 2(j+i) + 2j) \cdot W^3 + j \cdot W = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + j \cdot W$$

Then, $A[j, \ell] = 1$ implies that the edge $(b_j^{(4)}, c_\ell^{(1)})$ exists. Since $\vec{w}_i[\ell] = 1$, each edge in the sub-path $(c_\ell^{(1)}, c_\ell^{(2)}, c_{\ell+i}^{(3)}, c_\ell^{(4)})$ exists and the vertices have combined weight,

$$4 \cdot W^4 + (4n - 2(\ell+i) + 2\ell) \cdot W^3 + \ell = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + \ell$$

Then, $A[\ell, k] = 1$ implies that the edge $(c_\ell^{(4)}, \hat{a}_k^{(1)})$ exists. Finally, as $\vec{u}_i[k] = 1$, each edge in the sub-path $(\hat{a}_k^{(1)}, \hat{a}_k^{(2)}, \hat{a}_{k+i}^{(3)}, \hat{a}_k^{(4)})$ exists and the vertices have combined weight,

$$4 \cdot W^4 + (4n - 2(k+i) + 2k) \cdot W^3 + k \cdot W^2 = 4 \cdot W^4 + (4n - 2i) \cdot W^3 + k \cdot W^2$$

We conclude by observing that edge $(\hat{a}_k^{(4)}, t)$ exists. Summing over all vertices in the path, we conclude the length of the path is,

$$18 \cdot W^4 + (16n - 8i) \cdot W^3 + 2k \cdot W^2 + j \cdot W + \ell$$

Conversely, suppose there is some path P of length

$$\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2.$$

Since all vertices have weight at least W^4 and

$$(16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2 < W^4,$$

any (s, t) -path P satisfying the length constraint has at most 18 vertices. Since we have a layered graph and only edges between adjacent layers, any such path P must have the form,

$$P = (s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, a_{i_4}^{(4)}, b_{i_5}^{(1)}, b_{i_6}^{(2)}, b_{i_7}^{(3)}, b_{i_8}^{(4)}, c_{i_9}^{(1)}, c_{i_{10}}^{(2)}, c_{i_{11}}^{(3)}, c_{i_{12}}^{(4)}, \hat{a}_{i_{13}}^{(1)}, \hat{a}_{i_{14}}^{(2)}, \hat{a}_{i_{15}}^{(3)}, \hat{a}_{i_{16}}^{(4)}, t)$$

At the time of the query after updating the k -th coordinate of \vec{u} in the OMv3 query $\vec{u}_i, \vec{v}_i, \vec{w}_i$, the minimum weight of any path from A_2 to A_4 is $3 \cdot W^4 + (4n - 2i) \cdot W^3$, and as argued in Lemma 3.6 this is attained only by paths of the form of Equation (2).

Similarly, the minimum weight of any path from B_2 to B_4 , C_2 to C_4 , \hat{A}_2 to \hat{A}_4 is $3 \cdot W^4 + (4n - 2i) \cdot W^3$ and any path not satisfying the form of Equation (2) has weight at least $3 \cdot W^4 + (4n - 2i + 1) \cdot W^3$. Since

$$\mathbf{d}(P) < 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k + 1) \cdot W^2 < 18 \cdot W^4 + (16n - 7i) \cdot W^3,$$

we conclude

$$\begin{aligned} 0 &= i_2 - i_4 = i_6 - i_8 = i_{10} - i_{12} = i_{14} - i_{16} \\ i &= i_3 - i_2 = i_7 - i_6 = i_{11} - i_{10} = i_{15} - i_{14}. \end{aligned}$$

Next, observe that for all $x \in \{a, b, c, \hat{a}\}$ and $X \in \{A, B, C, \hat{A}\}$, $x_j^{(2)}$ has a unique neighbor in X_1 , namely $x_j^{(1)}$. In particular, there are indices k', j', ℓ', \hat{k}' such that,

$$P = (s, a_{k'}^{(1)}, a_{k'}^{(2)}, a_{k'+i}^{(3)}, a_{k'}^{(4)}, b_{j'}^{(1)}, b_{j'}^{(2)}, b_{j'+i}^{(3)}, b_{j'}^{(4)}, c_{\ell'}^{(1)}, c_{\ell'}^{(2)}, c_{\ell'+i}^{(3)}, c_{\ell'}^{(4)}, \hat{a}_{\hat{k}'}^{(1)}, \hat{a}_{\hat{k}'}^{(2)}, \hat{a}_{\hat{k}'+i}^{(3)}, \hat{a}_{\hat{k}'}^{(4)}, t)$$

Next, we claim $k = k' = \hat{k}'$. As before, we have $k', \hat{k}' \geq k$ from Equation (4) since we iterate over k in decreasing order. If $k' > k$, then

$$\text{wt}\left(a_{k'}^{(1)}\right) > (k+1) \cdot W^3.$$

Similarly, if $\hat{k}' > k$, then

$$\text{wt}\left(\hat{a}_{\hat{k}'}^{(1)}\right) > (k+1) \cdot W^2.$$

If either condition fails, then $\mathbf{d}(P) \geq 18 \cdot W^4 + (16n - 8i) \cdot W^3 + (2k+1) \cdot W^2$, a contradiction. Thus, $k = k' = \hat{k}'$.

In particular, we have $\vec{u}_i[k] = \vec{v}_i[j'] = \vec{w}_i[\ell'] = 1$ and furthermore $A[k, j] = A[j, \ell] = A[\ell, k] = 1$ as edges $(a_k^{(4)}, b_{j'}^{(1)})$, $(b_{j'}^{(4)}, c_{\ell'}^{(1)})$, and $(c_{\ell'}^{(4)}, \hat{a}_k^{(1)})$ exist, so that

$$\vec{u}_i[k] = \vec{v}_i[j'] = \vec{w}_i[\ell'] = A[k, j] = A[j, \ell] = A[\ell, k] = \text{TRUE}.$$

Correctness of Decremental Reduction We now describe the modifications for the decremental case. In the decremental case, we instead have the path

$$\left(s, a_k^{(1)}, a_k^{(2)}, a_{k+n-i}^{(3)}, a_k^{(4)}, b_j^{(1)}, b_j^{(2)}, b_{j+n-i}^{(3)}, b_j^{(4)}, c_{\ell}^{(1)}, c_{\ell}^{(2)}, c_{\ell+n-i}^{(3)}, c_{\ell}^{(4)}, \hat{a}_k^{(1)}, \hat{a}_k^{(2)}, \hat{a}_{k+n-i}^{(3)}, \hat{a}_k^{(4)}, t\right),$$

which has length

$$18 \cdot W^4 + (8n + 8i) \cdot W^3 + 2k \cdot W^2 + j \cdot W + \ell.$$

We verify that this path exists and has the required length. Since $\vec{u}_i[k] = 1$, each edge in the sub-path $(a_k^{(1)}, a_k^{(2)}, a_{k+n-i}^{(3)}, a_k^{(4)})$ exists and the vertices have combined weight

$$4 \cdot W^4 + (4n - 2(k+n-i) + 2k) \cdot W^3 + k \cdot W^2 = 4 \cdot W^4 + (2n + 2i) \cdot W^3 + k \cdot W^2.$$

Following similar arguments as above the path exists with the required length.

Conversely, suppose there is some path P of length

$$\mathbf{d}(P) < 18 \cdot W^4 + (8n + 8i) \cdot W^3 + (2k+1) \cdot W^2.$$

As above, the path has at most 18 vertices and therefore has the form

$$P = \left(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, a_{i_4}^{(4)}, b_{i_5}^{(1)}, b_{i_6}^{(2)}, b_{i_7}^{(3)}, b_{i_8}^{(4)}, c_{i_9}^{(1)}, c_{i_{10}}^{(2)}, c_{i_{11}}^{(3)}, c_{i_{12}}^{(4)}, \hat{a}_{i_{13}}^{(1)}, \hat{a}_{i_{14}}^{(2)}, \hat{a}_{i_{15}}^{(3)}, \hat{a}_{i_{16}}^{(4)}, t\right).$$

At the time of the query after updating the k -th coordinate of \vec{u} in the OMv3 query $\vec{u}_i, \vec{v}_i, \vec{w}_i$, the minimum weight of any path from A_2 to A_4 is $3 \cdot W^4 + (2n + 2i) \cdot W^3$ and this is attained only by paths of the form of

$$(a_{k'}^{(2)}, a_{k'+n-i}^{(3)}, a_{k'}^{(4)}), \quad (3)$$

for $k' \geq k$ since we iterate over coordinates k in increasing order. In fact, any other path from A_2 to A_4 has weight at least $3 \cdot W^4 + (2n + 2i + 1) \cdot W^3$. This is because $a_{i_3}^{(3)}$ is adjacent only to vertices $a_{i_2}^{(2)}, a_{i_4}^{(4)}$ with $i_3 \leq i_2 + n - i$ or $i_2 \geq i_3 - n + i$. Thus, if either i_2 or i_4 is $> i_3 - n + i$, the three vertices have weight at least

$$3 \cdot W^4 + (4n - 2i_3 + 2i_3 - 2n + 2i + 1) \cdot W^3 = 3 \cdot W^4 + (2n + 2i + 1) \cdot W^3.$$

Similarly, the minimum weight of any path from B_2 to B_4 , C_2 to C_4 , \hat{A}_2 to \hat{A}_4 is $3 \cdot W^4 + (2n+2i) \cdot W^3$ and any path not satisfying the form of Equation (3) has weight at least $3 \cdot W^4 + (2n+2i+1) \cdot W^3$. Since

$$\mathbf{d}(P) < 18 \cdot W^4 + (8n + 8i) \cdot W^3 + (2k + 1) \cdot W^2 < 18 \cdot W^4 + (8n + 9i) \cdot W^3,$$

we conclude

$$\begin{aligned} 0 &= i_2 - i_4 = i_6 - i_8 = i_{10} - i_{12} = i_{14} - i_{16} \\ i &= i_3 - i_2 = i_7 - i_6 = i_{11} - i_{10} = i_{15} - i_{14}. \end{aligned}$$

Following previous arguments, there are indices k', j', ℓ', \hat{k}' such that,

$$P = \left(s, a_{k'}^{(1)}, a_{k'}^{(2)}, a_{k'+i}^{(3)}, a_{k'}^{(4)}, b_{j'}^{(1)}, b_{j'}^{(2)}, b_{j'+i}^{(3)}, b_{j'}^{(4)}, c_{\ell'}^{(1)}, c_{\ell'}^{(2)}, c_{\ell'+i}^{(3)}, c_{\ell'}^{(4)}, \hat{a}_{\hat{k}'}^{(1)}, \hat{a}_{\hat{k}'}^{(2)}, \hat{a}_{\hat{k}'+i}^{(3)}, \hat{a}_{\hat{k}'}^{(4)}, t \right)$$

Next, we claim $k = k' = \hat{k}'$. As before, we have $k', \hat{k}' \geq k$ from Equation (4) since we iterate over k in increasing order. Thus, $k = k' = \hat{k}'$ as in the incremental case. Following an identical proof as in the incremental case, we conclude

$$\vec{u}_i[k] = \vec{v}_i[j'] = \vec{w}_i[\ell'] = A[k, j] = A[j, \ell] = A[\ell, k] = \text{TRUE}.$$

□

3.2.2 Lower Bounds under Minimum Witness Hypothesis

Under the Minimum-Witness 3-Product hypothesis (Hypothesis 1.6), we show a stronger lower bound for the single source version of the problem. Before proving the lower bound, we provide some further discussion on the Minimum-Witness hypothesis. The Minimum-Witness Product problem is defined below.

Definition 3.10 (Min-Witness Product). Given two Boolean matrices A, B , their min-witness product C is defined by $C_{ij} = \min\{k \mid A_{ik} \wedge B_{kj}\}$.

[CKL07] computes the product in $O(n^{2+\lambda})$ time, where $\omega(1, \lambda, 1) = 1 + 2\lambda$. If $\omega = 2$, this running time is essentially $O(n^{2.5})$. [LPV20] conjectures that this is essentially optimal.

Hypothesis 3.11 (Minimum-Witness Product hypothesis). *There is no $O(n^{2.5-\varepsilon})$ algorithm for computing Minimum-Witness Product between two $n \times n$ Boolean matrices, for any $\varepsilon > 0$.*

Minimum-Witness Product can also be generalized to a problem between $k > 2$ input Boolean matrices [KL22]. In this work, we only need the following version where $k = 3$.

Definition 3.12 (Minimum-Witness 3-Product). Given three $n \times n$ Boolean matrices A, B, C , for every $i_1, i_2, i_3 \in [n]$, find the minimum value of $j \in [n]$ such that $A_{i_1, j} \wedge B_{i_2, j} \wedge C_{i_3, j}$.

The current best algorithm for Minimum-Witness 3-Product runs in $O(n^{3.5})$ time when $\omega = 2$ [KL22], and it seems that new ideas are required in order to break this $O(n^{3.5})$ running time. Below is some evidence, similar to [JX22]’s argument for why the so-called OuMv $_k$ hypothesis is plausible. First, we observe that we can use the Minimum-Witness 3-Product problem to solve 4-Clque Detection. Given a 4-partite graph $G = (V_1 \cup V_2 \cup V_3 \cup V_4, E)$, we can use the matrix A to encode the edges between V_1 and V_4 , matrix B to encode the edges between V_2 and V_4 and matrix C to encode the edges between V_3 and V_4 . Thus, the result of Minimum-Witness 3-Product can tell us, for every triple $(v_1, v_2, v_3) \in V_1 \times V_2 \times V_3$, whether there is some $v_4 \in V_4$ that is connected to all of

them. Then we can detect whether the graph has a 4-clique in $O(n^3)$ time. Currently, the all truly sub-quartic time algorithms for 4-Clique Detection essentially groups two node parts together, then solves a Triangle Detection instance on an unbalanced tripartite graph. Therefore, assuming there is no better framework for 4-Clique Detection, we also need to perform similar groupings for the Minimum-Witness 3-Product problem, as it can be used to solve 4-Clique Detection. If we group any two of the indices i_1, i_2, i_3, j together, we obtain a Min-Witness Product instance between an $n \times n^2$ matrix and an $n^2 \times n$ matrix, or between an $n^2 \times n$ matrix and an $n \times n$ matrix. Either way, the current fastest algorithm for them has running time $O(n^{3.5})$ if $\omega = 2$, which essentially follows the same framework for the square inputs case [CKL07]. Therefore, in order to refute the following hypothesis, it seems that one would need new ideas for either 4-Clique Detection or Min-Witness Product.

Hypothesis 1.6 (Minimum-Witness 3-Product hypothesis). *There is no $O(n^{3.5-\varepsilon})$ time algorithm for the Minimum-Witness 3-Product problem for $\varepsilon > 0$.*

3.2.3 Lower Bound for nw-SSSP

We now give our lower bound for nw-SSSP.

Theorem 3.13. *Under the Minimum-Witness 3-Product hypothesis, any algorithm computing incremental/decremental nw-SSSP on undirected graphs requires $n^{3.5-o(1)}$ total update and query time.*

Proof. As before, we describe only the reduction for incremental stSP. For the decremental setting, we can execute the reduction in reverse.

Suppose for contradiction there is an algorithm \mathcal{A} for incremental nw-SSSP in $O(n^{3.5-c})$ total time. Consider a Minimum-Witness 3-Product instance with matrices A, C, D of dimension $n \times n$. Throughout the reduction, assume that the matrices A, C, D are indexed from 0 to $n-1$. We design an algorithm computing min-witness 3-product with total time $O(n^{3.5-c})$, thus contradicting the Minimum-Witness 3-Product hypothesis.

We construct a graph consisting of the gadgets

$$\{s\} \cup G_R(A) \cup G_R(B) \cup C_1$$

of Theorem 3.5 with the following modifications. Whereas A was a vertex set of n nodes, think of A now as a collection of n rows of the matrix A . Whereas B was a vertex set of n nodes, think of B now as a collection of n columns. Whereas $C_1 = \{c^{(1)} : c \in \{0, \dots, n-1\}\}$ was a copy of the vertex set C in the 4-Clique instance, think of $c^{(1)}$ now as representing the c -th row of matrix C .

We describe the formal reduction below. Initially, insert edges $(s, a_i^{(1)})$ all $i \in \{0, 1, \dots, n-1\}$. Next, insert edges between $(a_i^{(4)}, b_j^{(1)})$ (respectively between $(b_j^{(4)}, c_i^{(1)})$) if and only if $A[i, j] = \text{TRUE}$ (respectively $C[i, j] = \text{TRUE}$). Insert edges $(a^{(1)}, a^{(2)}), (b^{(1)}, b^{(2)})$ for all a, b .

We now proceed to insert more edges in the reduction. In the outer loop we iterate over rows $i \in \{0, 1, \dots, n-1\}$ of D in increasing order.

1. For each $D[i, j] = \text{TRUE}$, insert edges $(b_j^{(2)}, b_{j+i}^{(3)}), (b_{j+i}^{(3)}, b_j^{(4)})$.
2. In the inner loop, we iterate over rows $k \in \{0, 1, \dots, n-1\}$ of A in decreasing order.
 - (a) Insert edges between $(a_k^{(2)}, a_{k+i}^{(3)})$ and $(a_{k+i}^{(3)}, a_k^{(4)})$.
 - (b) For each $c_\ell^{(1)} \in C_1$, query $\mathbf{d}(s, c_\ell^{(1)})$ from \mathcal{A} .
If $\mathbf{d}(s, c_\ell^{(1)}) < 10 \cdot W^4 + (8n - 4i) \cdot W^3 + (k+1) \cdot W^2$, let

$$\mathbf{d}(s, c_\ell^{(1)}) = 10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + b' \cdot W + \ell.$$

for some integer $0 \leq b' < W$. Then, assign the minimum witness of (k, ℓ, i) to be b' . Otherwise, if the $\mathbf{d}(s, c)$ does not satisfy the required constraint, assign the minimum witness to \perp , or in other words there is no witness j such that $A[k, j] = C[\ell, j] = D[i, j] = \text{TRUE}$.

Again, we emphasize that we iterate over D in increasing order and A in decreasing order. Clearly, the bottleneck is running \mathcal{A} , so the total time of this algorithm is $O(n^{3.5-c})$.

To show correctness, we first prove the following lemma.

Lemma 3.14. *Consider the query when iterating over the i -th row of D and the k -th row of A . Let ℓ be a row of C . Then, there is a one-to-one correspondence between:*

1. *Paths P between $s, c_\ell^{(1)}$ with $\mathbf{d}(P) < 10 \cdot W^4 + (8n - 4i) \cdot W^3 + (k + 1) \cdot W^2$;*
2. *Witnesses j such that $A[k, j] \wedge C[\ell, j] \wedge D[i, j]$.*

Furthermore, each such path P has length

$$\mathbf{d}(P) = 10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + j \cdot W + \ell.$$

Proof. Suppose there is a witness j such that $A[k, j] = C[\ell, j] = D[i, j] = 1$. Then, the following $(s, c_\ell^{(1)})$ -path exists :

$$(s, a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)}, b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)}, c_\ell^{(1)}).$$

Furthermore, this path has length

$$10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + j \cdot W + \ell.$$

Note that by our choice of W , $j \cdot W + \ell < W^2$ so this path satisfies the desired constraints. We verify that this path exists and has the claimed length. Recall that the edge $(s, a_k^{(1)})$ exists. Also, each edge in the sub-path $(a_k^{(1)}, a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)})$ exists and the vertices have combined weight $4 \cdot W^4 + (4n - 2i) \cdot W^3 + k \cdot W^2$. Next, $A[k, j] = \text{TRUE}$ implies that the edge $(a_k^{(4)}, b_j^{(1)})$ exists. Since $D[i, j] = \text{TRUE}$, each edge in the sub-path $(b_j^{(1)}, b_j^{(2)}, b_{j+i}^{(3)}, b_j^{(4)})$ exists and the vertices have combined weight $4 \cdot W^4 + (4n - 2i) \cdot W^3 + j \cdot W$. Then, $C[\ell, j] = \text{TRUE}$ implies that the edge $(b_j^{(4)}, c_\ell^{(1)})$ exists, thus verifying the existence of the path. Summing over all vertices in the path, we conclude the length of the path is

$$10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + j \cdot W + \ell.$$

Conversely, suppose there is some path P of length

$$\mathbf{d}(P) < 10 \cdot W^4 + (8n - 4i) \cdot W^3 + (k + 1) \cdot W^2.$$

Since all vertices have weight at least W^4 and

$$(8n - 4i) \cdot W^3 + (k + 1) \cdot W^2 < W^4,$$

any (s, t) -path P satisfying the length constraint has at most 10 vertices. Since we have a layered graph and edges only exist between adjacent layers, any such path P must have the form

$$P = (s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, a_{i_4}^{(4)}, b_{i_5}^{(1)}, b_{i_6}^{(2)}, b_{i_7}^{(3)}, b_{i_8}^{(4)}, c_\ell^{(1)}).$$

At the time of the query in the i -th iteration over D and k -th iteration over A in the inner-loop, the minimum weight of any path from A_2 to A_4 is $3W^4 + (4n - 2i) \cdot W^3$, and this is attained only by paths of the form

$$\left(a_k^{(2)}, a_{k+i}^{(3)}, a_k^{(4)} \right). \quad (4)$$

As before, any other path from A_2, A_4 has weight at least $3 \cdot W^4 + (4n - 2i + 1) \cdot W^3$. By a similar argument, the minimum weight of any path from B_2 to B_4 is $3 \cdot W^4 + (4n - 2i) \cdot W^3$ and any path not satisfying the form of Eq. (4) has weight at least $3 \cdot W^4 + (4n - 2i + 1) \cdot W^3$. Since

$$\mathbf{d}(P) < 10 \cdot W^4 + (8n - 4i) \cdot W^3 + (k + 1) \cdot W^2 < 12 \cdot W^4 + (8n - 3i) \cdot W^3,$$

we conclude

$$\begin{aligned} 0 &= i_2 - i_4 = i_6 - i_8, \\ i &= i_3 - i_2 = i_7 - i_6. \end{aligned}$$

Next, observe that for all $x \in \{a, b\}$ and $X \in \{A, B\}$, $x_j^{(2)}$ has a unique neighbor in X_1 , namely $x_j^{(1)}$. In particular, there are indices k', j' such that,

$$P = \left(s, a_{k'}^{(1)}, a_{k'}^{(2)}, a_{k'+i}^{(3)}, a_{k'}^{(4)}, b_{j'}^{(1)}, b_{j'}^{(2)}, b_{j'+i}^{(3)}, b_{j'}^{(4)}, c_\ell^{(1)} \right).$$

Following a similar argument as Lemma 3.6, we see that $k = k'$. In particular, we have $D[i, j'] = \text{TRUE}$ and $A[k, j'] = \text{TRUE}$ since $(a_k^{(4)}, b_{j'}^{(1)})$ is an edge, and $C[\ell, j'] = \text{TRUE}$ since $(b_{j'}^{(4)}, c_\ell^{(1)})$ is an edge. In particular, j' is a witness such that $A[k, j'] = C[\ell, j'] = D[i, j'] = \text{TRUE}$. Note that this path P has length

$$\mathbf{d}(P) = 10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + j' \cdot W + \ell$$

as desired. \square

From Lemma 3.14, at the time of the query, each path P from s to $c_\ell^{(1)}$ with length $\mathbf{d}(P) < 10 \cdot W^4 + (8n - 4i) \cdot W^3 + (k + 1) \cdot W^2$ corresponds to a witness j such that $A[k, j] = C[\ell, j] = D[i, j] = \text{TRUE}$. Furthermore, Lemma 3.14 shows that

$$\mathbf{d}(P) = 10 \cdot W^4 + (8n - 4i) \cdot W^3 + k \cdot W^2 + j \cdot W + \ell,$$

so that the shortest path returns the minimum j such that $A[k, j] \wedge C[\ell, j] \wedge D[i, j]$. In particular it computes the minimum-witness 3-product for this (k, ℓ, i) triplet. If there is no path, then there is no witness j such that $A[k, j] \wedge C[\ell, j] \wedge D[i, j]$. Thus, we compute the Minimum-Witness 3-Product instance in $O(n^{3.5-c})$ time, contradicting the hypothesis. \square

4 Bottleneck Paths

In this section, we give our algorithms and lower bounds for partially dynamic bottleneck paths. The bottleneck path problem is formally defined below.

Definition 4.1 (Bottleneck Paths). Let $G = (V, E, \text{wt})$ be a directed, weighted graph. For any path P , let $\mathbf{b}(P) = \min_{e \in P} \text{wt}(e)$ denote the **bottleneck capacity** of path P . For any pair of nodes u, v , let $\mathbf{bp}(u, v)$ denote the **bottleneck path** (breaking ties arbitrarily), the path between u, v maximizing $\mathbf{b}(P)$. Let $\mathbf{b}(u, v) = \mathbf{b}(\mathbf{bp}(u, v))$.

The (s, t) -**Bottleneck Path** problem (stBP) asks to compute $\mathbf{b}(s, t)$ for fixed nodes s, t . The **Single Source Bottleneck Paths** problem (SSBP) asks to compute $\mathbf{b}(s, v)$ for a single source s and for all $v \in V$. The **All Pairs Bottleneck Paths** problem (APBP) asks to compute $\mathbf{b}(u, v)$ for all $u, v \in V$.

4.1 Incremental/Decremental Algorithm for stBP

We begin with partially dynamic algorithms for bottleneck paths. First, we present a simple algorithm using a fully dynamic (s, t) -reachability data structure to compute partially dynamic stBP, in which we need to maintain an n -vertex graph undergoing arbitrary edge insertions and deletions, and we need to support querying whether there is a path from s to t .

Proposition 4.2. *The partially dynamic stBP can be solved in $O(m \cdot T(n))$ time, where $T(n) = \tilde{O}(n^{\varepsilon_1 + \varepsilon_2} + n^{\omega(1, \varepsilon_1, \varepsilon_2) - \varepsilon_1} + n^{\omega(1, 1, \varepsilon_2) - \varepsilon_2}) = O(n^{1.405})$ [vdBNS19] is the update/query time for fully dynamic (s, t) -reachability in an n -node graph.*

As it seems that [vdBNS19]'s algorithm works against oblivious, adaptive adversaries, our algorithm also does.

Proof. We first consider the decremental setting. We initialize a fully dynamic (s, t) -reachability data structure. We keep adding edges to the data structure from large weight to small, as long as s cannot reach t . Once we remove an edge, if s and t are not connected anymore, we again add more edges, from large weight to small, until s can reach t again. It is not difficult to see that the last edge we add to the data structure has the weight equal to $\mathbf{b}(s, t)$. To bound the running time, observe that each edge is inserted and deleted at most one time each.

The incremental setting can be solved similarly. \square

4.2 Incremental Algorithm for SSBP

Next, we will give an incremental SSBP algorithm. First, we formally define the relevant matrix products below, for which there are known subcubic algorithms using FMM.

Definition 4.3 (Dominance and Equality Product). Given two matrices A, B over a totally ordered set, their dominance product $C = A \leq_{\#} B$ is defined by $C_{ij} = |\{k \mid A_{ik} \leq B_{kj}\}|$.

Their equality product is defined by $C_{ij} = |\{k \mid A_{ik} = B_{kj}\}|$

Definition 4.4 ((min, \leq)-Product and (max, \leq)-Product). Given two matrices A, B over a totally ordered set, their (\min, \leq) product C is defined by $C_{ij} = \min_k \{B_{kj} \mid A_{ik} \leq B_{kj}\}$. The (\max, \leq) product C is defined by $C_{ij} = \max_k \{A_{ik} \mid A_{ik} \leq B_{kj}\}$.

Definition 4.5 ((max, min)-Product). Given two matrices A, B over a totally ordered set, their (\max, \min) product $C = A \otimes B$ is defined by $C_{ij} = \max_k \min \{A_{ik}, B_{kj}\}$.

As observed in [VWY07], the (\max, \min) product can be computed by taking the entry-wise maximum of two (\max, \leq) products, A, B and B^T, A^T .

Our SSBP algorithm requires the following lemma, which modified an existing algorithm for computing the (\max, \min) -product [DP09, GIL⁺21]. Formally, we split the computation of the (\max, \min) -product into a preprocessing phase and a query phase.

Lemma 4.6. *Given an $n \times n^b$ matrix A and an $n^b \times n$ matrix B , and a parameter $0 \leq g \leq b$, we can preprocess in $\tilde{O}(n^{\omega(1, b, 1) + b - g})$ time, so that for every given i , we can output the i -th row of the (\max, \min) -product between A and B in $\tilde{O}(n^{1+g})$ time.*

In particular, we can compute the (\max, \min) -product between A and B in $\tilde{O}(n^{(\omega(1, b, 1) + 2 + b)/2})$ time, by setting $g = (\omega(1, b, 1) + b - 2)/2$ and perform n queries.

We recall the definition of row (resp. column) balancing [DP09], which constructs a matrix A' given any matrix A ensuring that no row (resp. column) of A' has too many finite elements.

Let A be an $n \times p$ matrix with m finite elements (the others are either all ∞ or $-\infty$). The **row-balancing** of A , denoted $rb(A)$ is a pair of matrices (A', A'') each with at most $k = \lceil m/n \rceil$ finite elements in each row. The **column-balancing** of A , denoted $cb(A)$ is the pair of matrices (A'^T, A''^T) where $(A', A'') = rb(A^T)$.

The row-balancing of A is computed as follows. Assume that all non-finite elements are ∞ . For each row i , sort the finite elements of row i in increasing order, and partition this list into parts $T_i^1, \dots, T_i^{a_i}$ where the first $a_i - 1$ parts have exactly $k = \lceil m/n \rceil$ elements while the last part has at most k elements. Now, define $A'[i, j] = A[i, j]$ if $A[i, j] \in T_i^{a_i}$ and ∞ otherwise to be the matrix containing all elements in the last part. Now, since there are at most $\frac{m}{k} \leq n$ parts T_i^q , we choose an arbitrary injective mapping $\rho : (i, q) \mapsto [n]$ so that the part T_i^q is placed into the $\rho(i, q)$ -th row of A'' . In particular, $A''[i', j] = A[i, j]$ if $i' = \rho(i, q)$ and $A[i, j] \in T_i^q$. Otherwise, $A''[i, j] = \infty$.

Given row and column balanced matrices, [DP09] give an efficient algorithm for computing the dominance product (Definition 4.3) of sparse matrices. We repeat the proof for completeness.

Lemma 4.7 (Sparse Dominance Product [DP09]). *Suppose A is an $n^a \times n^b$ matrix with m_1 entries less than ∞ and B is an $n^b \times n^c$ matrix with m_2 entries greater than $-\infty$. The dominance product $A \leq_{\#} B$ can be computed in time $O\left(\frac{m_1 m_2}{n^b} + n^{\omega(a,b,c)}\right)$.*

Proof. Let $(A', A'') = cb(A)$ be the column-balancing of A . Note that A', A'' have at most $\lceil \frac{m_1}{n^b} \rceil$ finite entries per column. We construct the following two matrices

$$\hat{A}[i, k] = \begin{cases} 1 & A''[i, k] < \infty \\ 0 & \text{o/w} \end{cases}, \quad \hat{B}[k, j] = \begin{cases} 1 & B[k', j] \geq \max T_{k'}^{q'} \text{ and } \rho(k', q') = k \\ 0 & \text{o/w} \end{cases}.$$

We compute the matrix product $\hat{A}\hat{B}$. $\hat{A}[i, k]\hat{B}[k, j] = 1$ implies that $A''[i, k] < \infty$ and $B[k', j] \geq \max T_{k'}^{q'}$ where $\rho(k', q') = k$. In other words, $B[k', j]$ is larger than all elements in column k of A'' , which is a subset of column k' of A containing $A''[i, k] = A[i, k']$. Therefore, $(\hat{A}\hat{B})[i, j]$ counts the number of k' where $A[i, k'] \leq B[k', j]$ and $A[i, k'] \in T_{k'}^{q'}$ where $q' < a_k$ and $B[k', j]$ dominates all of $T_{k'}^{q'}$. It remains to check the following cases:

1. $A[i, k'] \in T_{k'}^{a_{k'}}$. For every $B[k', j]$, there are most $\lceil \frac{m_1}{n^b} \rceil$ rows i with $A[i, k'] \in T_{k'}^{a_{k'}}$. We compare $B[k', j]$ with elements $A[i, k'] \in T_{k'}^{a_{k'}}$ and increment $(A \leq_{\#} B)[i, j]$ if $A[i, k'] \leq B[k', j]$. This takes $O(m_2 \cdot \frac{m_1}{n^b})$ time.
2. $B[k', j]$ dominates some elements but not all in $T_{k'}^{q'}$. For every $B[k', j]$, this condition can only hold for one set $T_{k'}^{q'}$ as the parts are sorted in increasing order. We compare $B[k', j]$ with elements $A[i, k'] \in T_{k'}^{q'}$ and increment $(A \leq_{\#} B)[i, j]$ if $A[i, k'] \leq B[k', j]$. This takes $O(m_2 \cdot \frac{m_1}{n^b})$ time.

The overall time complexity is $O\left(\frac{m_1 m_2}{n^b} + n^{\omega(a,b,c)}\right)$. □

Given that we hope to construct a data structure for the (max, min)-product, we modify the computation of the sparse dominance product to compute a data structure separating the computation into preprocessing and query stages.

Lemma 4.8 (Sparse Dominance Product Data Structure). *Suppose A is an $n \times n^b$ matrix with m_1 entries less than ∞ and B is an $n^b \times n$ matrix with m_2 entries greater than $-\infty$. There is a data structure with $O(n^{\omega(1,b,1)})$ preprocessing time, answering column queries of the dominance product $A \leq_{\#} B$ in $O(c_j \frac{m_2}{n})$ time where the j -th column of B contains c_j finite elements.*

Proof. In the preprocessing phase, compute the row-balancing $(A', A'') = cb(A)$ as in Lemma 4.7. Similarly, construct matrices \hat{A}, \hat{B} , and compute the matrix product $\hat{A}\hat{B}$. This requires $n^{\omega(1,b,1)}$ time.

Now, suppose there is a query for the j -th column of $A \leq_{\#} B$. By assumption, the j -th column of B has c_j finite elements $< \infty$. Our goal is to answer the column query i.e. output $(A \leq_{\#} B)[i, j]$ for all i . As in Lemma 4.7, $(\hat{A}\hat{B})[i, j]$ counts the number of k' where $A[i, k'] \leq B[k', j]$ and $A[i, k'] \in T_{k'}^{q'}$ where $q' < a_k$ and $B[k', j]$ dominates all of $T_{k'}^{q'}$. It remains to check the following cases:

1. $A[i, k'] \in T_{k'}^{a_{k'}}$. For every finite $B[k', j]$, there are most $\lceil \frac{m_1}{n^b} \rceil$ rows i with $A[i, k'] \in T_{k'}^{a_{k'}}$. We compare $B[k', j]$ with elements $A[i, k'] \in T_{k'}^{a_{k'}}$ and increment $(A \leq_{\#} B)[i, j]$ if $A[i, k'] \leq B[k', j]$. This requires $O(c_j \cdot \frac{m_1}{n^b})$ time.
2. $B[k', j]$ dominates some elements but not all in $T_{k'}^{q'}$. For every finite $B[k', j]$, this condition can only hold for one set $T_{k'}^{q'}$ as the parts are sorted in increasing order. We compare $B[k', j]$ with elements $A[i, k'] \in T_{k'}^{q'}$ and increment $(A \leq_{\#} B)[i, j]$ if $A[i, k'] \leq B[k', j]$. This requires $O(c_j \cdot \frac{m_1}{n^b})$ time.

The query time complexity is therefore $O(c_j \cdot \frac{m_1}{n^b})$. \square

Using a similar argument, we can also construct a data structure answering row queries.

Corollary 4.9. *Suppose A is an $n \times n^b$ matrix with m_1 entries less than ∞ and B is an $n^b \times n$ matrix with m_2 entries greater than $-\infty$. There is a data structure with $O(n^{\omega(1,b,1)})$ preprocessing answering row queries of the dominance product $A \leq_{\#} B$ in $O(r_i \frac{m_2}{n})$ time where the i -th row of A contains r_i finite elements.*

Now, we prove Lemma 4.6.

Proof of Lemma 4.6. We modify the (\max, \leq) -Product algorithm of [DP09], since the (\max, \min) -Product can be reduced to two (\max, \leq) -Product instances. Fix a parameter $0 \leq g \leq b$. Let L be a sorted list of all entries in A, B and partition L into G groups, L_1, \dots, L_G such that the first $G-1$ parts have $\lceil \frac{2n^{1+b}}{G} \rceil$ elements and the final part has the remaining elements. For each $r \in [G]$, define the following matrices

$$A_r[i, j] = \begin{cases} A[i, j] & A[i, j] \in L_r \\ \infty & \text{o/w} \end{cases}, \quad B_r[i, j] = \begin{cases} B[i, j] & B[i, j] \in L_r \\ -\infty & \text{o/w} \end{cases}.$$

Compute the row-balancing $(A'_r, A''_r) = rb(A_r)$ and dominance products $A_r \leq_{\#} B, A'_r \leq_{\#} B, A''_r \leq_{\#} B$ as follows. For the dominance product, we do not fully compute them, only preprocessing the input as in Lemma 4.8.

1. To compute $A_r \leq_{\#} B$, construct Boolean matrices $\hat{A}_r[i, k] = 1$ if $A[i, k] \in L_r$ and $\hat{B}_r[k, j] = 1$ if $B[k, j] \in \bigcup_{i=r+1}^g L_i$. Compute $\hat{A}_r \hat{B}_r$ and the preprocess $A_r \leq_{\#} B_r$ according to Lemma 4.8. This takes $O(n^{\omega(1,b,1)})$ time.
2. To compute $A'_r \leq_{\#} B$, construct Boolean matrices $\hat{A}'_r[i, k] = 1$ if $A'[i, k] \in L_r$ and \hat{B}_r identically. Then, preprocess $A'_r \leq_{\#} B_r$ and compute $\hat{A}'_r \hat{B}_r$ in $O(n^{\omega(1,b,1)})$ time.
3. To compute $A''_r \leq_{\#} B$, construct Boolean matrices $\hat{A}''_r[i, k] = 1$ if $A''[i, k] \in L_r$ and \hat{B}_r identically. Then, preprocess $A''_r \leq_{\#} B_r$ and compute $\hat{A}''_r \hat{B}_r$ in $O(n^{\omega(1,b,1)})$ time.

We have computed $\hat{A}_r \hat{B}_r, \hat{A}'_r \hat{B}_r, \hat{A}''_r \hat{B}_r$ and preprocessed $A_r \leq_{\#} B_r$ and $A'_r \leq_{\#} B_r$ for all $r \in [G]$. This concludes the preprocessing phase which has overall time complexity

$$O\left(Gn^{\omega(1,b,1)}\right).$$

Then, in the query phase, we hope to compute the i -th row of the (\max, \leq) product. We begin by computing the i -th row of the preprocessed dominance products. For each $r \in [G]$, let c_r denote the number of finite elements in the i -th row of A_r . First, we observe

$$A_r \leq_{\#} B = A_r \leq_{\#} B_r + \hat{A}_r \hat{B}_r,$$

since the first term counts dominance pairs where both entries are in L_r while the second counts dominance pairs where $B[k, j] \in L_i$ for $i > r$. By Lemma 4.9, we can compute the i -th row of $A_r \leq_{\#} B_r$ in time $O\left(c_r \cdot \frac{n^{1+b}}{Gn}\right) = O\left(\frac{c_r n^b}{G}\right)$ since B_r has $O\left(\frac{n^{1+b}}{G}\right)$ finite elements. Adding this to $\hat{A}_r \hat{B}_r$ computes $A_r \leq_{\#} B$. Similarly, we use the following identities to compute $A'_r \leq_{\#} B, A''_r \leq_{\#} B_r$ in $O\left(\frac{c_r n^b}{G}\right)$ time:

$$\begin{aligned} A'_r \leq_{\#} B &= A'_r \leq_{\#} B_r + \hat{A}'_r \hat{B}_r, \\ A''_r \leq_{\#} B &= A''_r \leq_{\#} B_r + \hat{A}''_r \hat{B}_r. \end{aligned}$$

We will in fact need to compute all rows of $A''_r \leq_{\#} B_r$ indexed by $i' = \rho(i, q)$ for $q \in [a_i - 1]$, accounting for $\lceil \frac{c_r G}{n^b} \rceil$ rows in $(A''_r \leq_{\#} B_r)$. However, each of these rows contains at most $O\left(\frac{n^b}{G}\right)$ finite elements (and this is also true of the single i -th row in A'_r). Thus, querying all such rows takes time $O\left(c_r \frac{n^b}{G}\right)$ time. Overall computing the rows of these dominance products takes time

$$O\left(\sum_{r=1}^G \frac{c_r n^b}{G}\right) = O\left(\frac{n^{1+b}}{G}\right).$$

For each entry (i, j) we find the largest r such that $(A_r \leq_{\#} B)[i, j] > 0$. Let C denote the (\max, \leq) product of A, B . There are now two cases:

1. Suppose $(A'_r \leq_{\#} B)[i, j] > 0$. Then, since the i -th row of A'_r consists of the largest entries of the i -th row of A_r , we have $C[i, j] \in A'_r[i]$. By brute-force search over the elements of A'_r , we obtain $C[i, j]$ in $O\left(\frac{n^b}{G}\right)$ time.
2. Otherwise, find the largest q such that $(A''_r \leq_{\#} B)[\rho(i, q), j] > 0$ (this is in fact computed in the row-querying phase above) so that $C[i, j] \in T_i^q$. We compute $C[i, j]$ by brute force search in $O\left(\frac{n^b}{G}\right)$ time.

Since we compute this for all j , the row query requires $O(n^{1+b}G)$ time. To complete the proof, choose $G = \Theta(n^{b-g})$ so that $n^g = \Theta\left(\frac{n^b}{G}\right)$. \square

We now present our incremental algorithm for SSBP, which improves upon the simple framework used for our incremental nw-SSSP algorithms. Recall that we processed updates in batches, recomputing nw-APSP on the graph at the start of each batch. However, this is wasteful: We do not need *every* distance in the all-pairs solution, only the distances involving endpoints of the inserted edges. For bottleneck paths, we showed that the corresponding (\max, \min) -product is amenable to a data structure with a preprocessing phase and a query phase. Thus, at the end of each batch,

we preprocess the (max, min)-product on the current graph. Then, in future batches, whenever the capacity between a pair of vertices is required, we can simply query the relevant data structure. Since we do not need to query every single pair of vertices, this should save time over a complete APBP computation.

In the following, let $0 \leq g \leq t \leq 1$ be parameters. Consider the whole sequence of updates in incremental SSBP, and consider all intervals of the sequence of the form $(2^\ell \cdot (s-1) \cdot n^t, 2^\ell \cdot s \cdot n^t]$, for $\ell \geq 0, s > 0$. These are essentially dyadic intervals with endpoints multiplied by n^t .

Lemma 4.10. *For every interval L of length $n^{T'} = 2^\ell n^t$ described above, we can create a data structure in $\tilde{O}(n^{1+T+g} + n^{T \cdot (3+\omega)/2} + n^{(1+2T+\omega(1,T,T))/2} + n^{\omega(1,T,1)+T-g})$ preprocessing time where $T = \min\{1, T'\}$, so that given any vertex v , we can compute in $\tilde{O}(n^{1+g})$ time SSBP to/from v among paths whose latest updated edge is in L .*

Proof. We can prove the lemma by induction, and assume that the data structures have been created for all intervals strictly before L .

Let $U \subseteq V$ be the set of vertices that are among the endpoints of all updated edges in L . Clearly, $|U| = O(n^T)$. For each $u \in U$, we query the previous data structures to find out SSBP to/from u among paths whose latest updated edge is strictly before L . As the intervals are like dyadic intervals, we can find $O(\log n)$ intervals that cover all updates before L , so we only need to query $O(\log n)$ previous data structures. Therefore, the total time for this step is $\tilde{O}(n^{1+T+g})$. We use $\mathbf{b}_<(u, v)$ to denote bottleneck distances among paths whose latest updated edge is strictly before L .

We create a $V \times U$ matrix A , a $U \times U$ matrix S , and a $U \times V$ matrix B . All entries of A, S, B correspond to the $\mathbf{b}_<$ distances. Additionally, for entries of S , we also update it with the updated edge weights in the interval L . Then it is not difficult to see that $(A \oslash (S)^{\otimes n} \oslash B)_{i,j}$ is the bottleneck distance from i to j among paths whose latest updated edge is in L . Thus, we can do the following:

1. Compute $(S)^{\otimes n}$ in $\tilde{O}(n^{T \cdot (3+\omega)/2})$ time.
2. Compute $A \oslash (S)^{\otimes n}$ in $\tilde{O}(n^{(1+2T+\omega(1,T,T))/2})$ time using Lemma 4.6.
3. Finally, use the data structure part of 4.6 to preprocess between $(A \oslash (S)^{\otimes n})$ and B in $\tilde{O}(n^{\omega(1,T,1)+T-g})$ time, so that we can later support row-queries of $(A \oslash (S)^{\otimes n}) \oslash B$ in $\tilde{O}(n^{1+g})$ time.

Summing over all terms obtains the desired time complexity. \square

We are now ready to give the incremental SSBP algorithm.

Proposition 4.11. *Incremental SSBP can be solved in $\tilde{O}(n^{(1+\omega)/2})$ amortized time per update.*

Proof. Besides the data structures in Lemma 4.10, we also perform the followings after each edge insertion.

1. Let the new edge be (u, v) , and let τ be the previous timestamp that is a multiple of n^t .
2. Use the data structures in Lemma 4.10 to compute the SSBP to/from u, v and s among paths whose latest updated edge is on or before τ . This can be done in $\tilde{O}(n^{1+g})$ time. Now, it suffices to compute SSBP from s among paths whose latest updated edge is after τ .
3. Create the following “three-layer” graph. The first layer is $\{s\}$, the second layer consists of all vertices that are endpoints of some edge updated after τ , and the third layer is V . From the first layer to the second layer, the second layer to the third layer, and within the second layer, we add edges using distances computed in the previous bullet point (all these distances have been computed, either in this update or an earlier update). Between vertices in the second layer, we also add edges updated after τ (if there was an edge before, we keep the edge with larger weight). Then we compute SSBP of this graph in $\tilde{O}(n^{1+t})$ time (because the number

of updates after τ is $O(n^t)$, so the number of vertices in the second layer is $O(n^t)$, and the total number of edges is $O(n^{1+t})$.

Runtime analysis. The amortized running time of the preprocessing of the data structures in Lemma 4.10 is

$$\max_{T \in [t,1]} \tilde{O}(n^{1+g} + n^{T \cdot (1+\omega)/2} + n^{(1+\omega(1,T,T))/2} + n^{\omega(1,T,1)-g}) = \tilde{O}(n^{1+g} + n^{(1+\omega)/2} + n^{\omega-g})$$

and the running time of the remaining parts described above is

$$\tilde{O}(n^{1+g} + n^{1+t}).$$

Taking $t = g = \frac{\omega-1}{2}$ gives the claimed $\tilde{O}(n^{(1+\omega)/2})$ running time. \square

Remark 4.12 (Efficient Algorithm for Fully Dynamic Bottleneck Path in **Undirected** Graphs). Given an undirected graph G , it was known how to compute the weights of the bottleneck paths using a maximum spanning tree T of G [Hu61]: The weight of the bottleneck path between two vertices u and v is exactly the weight of the smallest edge on the tree path between u and v in T . Therefore, in order to maintain the bottleneck paths for a dynamic graph, it suffices to maintain its maximum spanning tree, and to support minimum weight queries on tree paths. For instance, we can use [HRWN15] to maintain the maximum spanning tree, and use link-cut tree [ST83] to support minimum weight queries on tree paths, to get a deterministic algorithm with $O(\log^4 n / \log \log n)$ amortized update time, and $O(\log n)$ amortized query time.

4.3 Partially Dynamic Algorithm for SSBP

Our lower bound in Theorems 4.16 and 4.18 crucially uses n^2 distinct weights in the constructed graph. The following proposition shows that this is necessary. In particular, by an appropriate modification of ES-Trees [ES81], partially dynamic SSBP can be computed efficiently on graphs with fewer distinct weights. Moreover, this algorithm is combinatorial.

Proposition 4.13. *There is an algorithm solving incremental/decremental SSBP in $\tilde{O}(mW)$ total-time update time, where W is the number of distinct weights. Furthermore, the algorithm answers queries $\mathbf{b}(s, v)$ in $O(\log W)$ time. This algorithm holds against oblivious, adaptive adversaries.*

Proof. We begin by describing the decremental setting. Let G denote the initial graph. For every distinct weight w , consider the graph $G_w = \{e \text{ s.t. } \text{wt}(e) \geq w\}$. Given graph G , we initialize a decremental SSR algorithm of [BPW19] on G_w for each w . Given an edge deletion, we delete e from every instance that contains w .

Given a query $\mathbf{b}(s, v)$, we proceed by binary search to find the maximum w where v is reachable from s on G_w and return the maximum such w as the bottleneck capacity between s, v .

Since the decremental SSR instance of [BPW19] answers queries in $O(1)$ time, we can answer queries in $O(\log W)$ time via binary search. Since there are at most W instances, the total time required by the algorithm is $\tilde{O}(mW)$.

Note that we can handle the incremental setting since incremental SSR can also be solved in total time $\tilde{O}(m)$. \square

4.4 Fully Dynamic All-Pairs Bottleneck Paths

We present our fully dynamic APBP algorithm.

Proposition 4.14. *Fully dynamic APBP can be solved in $\tilde{O}(n^2)$ amortized time per update deterministically.*

Proof. We first describe the intuition of our algorithm. Suppose all the edge weights are integers from $[1, 2n^2]$. Let M be some sufficiently large integer, and we replace the weight w of an edge by M^{2n^2-w} . Consider a path from u to v where the edge weights on the path are w_1, \dots, w_ℓ , then the total length of the path under the new weight is $\sum_i M^{2n^2-w_i}$, which is dominated by the edge with the smallest value of w_i . Thus, the shortest path from u to v under the new weight will maximize the minimum old edge weight, i.e., it will correspond to a bottleneck path. Intuitively, we could run Demetrescu and Italiano's dynamic APSP algorithm [DI04] under the new weight to achieve $\tilde{O}(n^2)$ amortized running time, assuming we can perform arithmetic operations between the new (huge) weights in constant time. Below, we describe algorithm in more details and the way to remove the assumptions.

We first use online list labelling to remove the assumption that all edge weights are from $[1, 2n^2]$.

Theorem 4.15 ([IKR81]). *We can deterministically store a set of t dynamically-changing items in an array of length $2t$, so that the sorted orders of the items are preserved, and we only move $\tilde{O}(1)$ items per operation in amortization.*

We use Theorem 4.15 to store the original edge weights of the graph as the items, and use the indices of the items stored in the arrays as their new weights. As the orders of indices of the items preserve the order of the items, this is without loss of generality for bottleneck paths. Now for every edge update in the graph, we might need to perform up to $\tilde{O}(1)$ weight updates in amortization, which is affordable. From now on, we assume all the edge weights are integers from $[1, n^2]$.

Next, we aim to use Demetrescu and Italiano's dynamic APSP algorithm [DI04] in a white-box way. In particular, if we simply store the path weights as described in the intuition, it would be too costly. Instead, we use some data structure to maintain the path weights.

The weight of any path in our graph (under the new weights described in the intuition) can be represented as a length $2n^2$ array, where the value on the i -th entry of the array equals the number of times the path use an edge of weight i . Therefore, say this array is $(a_i)_{i=1}^{2n^2}$, then the weight of the path is $\sum_{i=1}^{2n^2} a_i M^{2n^2-i}$, for some large integer M . We use a persistent segment tree [DSST89] to maintain such arrays efficiently.

Next, we run Demetrescu and Italiano's dynamic APSP algorithm [DI04] that has $\tilde{O}(n^2)$ amortized running time per operation, but we use the data structure to support all operations their algorithm performs on path weights. Initially, we create a persistent segment tree maintaining the all-zero array, representing the empty path. Then, our data structure needs to support the followings efficiently (i.e., in $\tilde{O}(1)$ time per operation).

1. **Concatenate one edge to an existing path.** Say the new edge has weight w , and the array corresponding to the existing path is a , then the array corresponding to the new path is a , with a_w incremented by 1. Therefore, we can simply copy the segment tree for the existing path (as we maintain persistent segment trees, this takes $\tilde{O}(1)$ time), and then update the w -th entry of the new segment tree in $\tilde{O}(1)$ time.
2. **Compare the weight of two paths.** For two paths whose weight correspond to arrays a and b respectively, their weights are $\sum_{i=1}^{2n^2} a_i M^{2n^2-i}$ and $\sum_{i=1}^{2n^2} b_i M^{2n^2-i}$, for some large integer M . To compare these two weights, it suffices to find the longest common prefix of the arrays a and b , and then compare the first differing position of a and b .

Then similar to Aho, Hopcroft and Ullman's algorithm for Tree Isomorphism [AH74], we assign all subtrees of all persistent segment tree an identifying number, so that isomorphic trees will have the same identifying number while distinct trees will have a different number. We maintain a perfect hash function H that maps integers or pairs of identifying numbers to identifying numbers, and we can maintain this hash function using a binary search tree. Leaf nodes in the segment trees represent an entry of an array, so it is associated with an integer. If this integer is c , then its identifying number would be $H(c)$. For internal nodes whose left subtree and right subtree have identifying numbers L and R respectively, its identifying number would be $H((L, R))$.

Then, we can perform binary search on two persistent segment trees augmented with identifying numbers to find the longest common prefix of the underlying arrays: To find the longest common prefix between two nodes N_1 and N_2 , we first compare the identifying numbers of their left subtrees. If these two numbers are different, we know that there must be a different array entry corresponding to the left subtree, so we recursively compare the left subtrees; otherwise, we recursively compare the right subtrees. As the height of the segment trees is $\tilde{O}(1)$, this takes $\tilde{O}(1)$ time.

□

4.5 Conditional Lower Bounds for Bottleneck Paths

We now present our conditional lower bound for partially dynamic bottleneck paths. We observe that due to Remark 4.12, our lower bounds hold only for directed graphs. In contrast, our lower bounds for shortest paths hold for undirected graphs as well.

Theorem 4.16. *Under the Combinatorial 4-Clique hypothesis, any combinatorial algorithm computing incremental/decremental stBP on directed graphs requires $n^{4-o(1)}$ total update and query time.*

Proof. We describe the reduction for incremental stBP, noting that in the decremental case we can simply run the reduction in reverse.

Suppose for contradiction there is an algorithm \mathcal{A} for incremental stBP with total time $O(n^{4-c})$ for some $c > 0$. Consider a 4-Clique Detection instance with vertex sets A, B, C, D of size n . Throughout the reduction, assume that the vertex sets A, B, C, D are indexed from 0 to $n - 1$. We design an algorithm with total time $O(n^{4-c})$ for some positive $c > 0$, contradicting the Combinatorial 4-Clique hypothesis.

We construct a graph with vertices

$$\{s\} \cup (A_1 \cup A_2 \cup A_3) \cup (B_1 \cup B_2 \cup B_3) \cup (C_1 \cup C_2 \cup C_3) \cup (\hat{A}_1 \cup \hat{A}_2 \cup \hat{A}_3) \cup \{t\}.$$

Each vertex set $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, \hat{A}_1, \hat{A}_2$ and \hat{A}_3 has n vertices. For each $a \in A$, we create a copy $a^{(i)} \in A_i$ and a copy $\hat{a}^{(i)} \in \hat{A}_i$ for $i \in \{1, 2, 3\}$. Similarly, for each $b \in B$ and $c \in C$, we create a copy $b^{(i)} \in B_i$ and $c^{(i)} \in C_i$ for $i \in \{1, 2, 3\}$. Initially, insert edges $(s, a^{(1)})$ and $(\hat{a}^{(3)}, t)$ with ∞ weight for all $a \in A$. Furthermore, insert edges with ∞ weight between $(a^{(3)}, b^{(1)})$ (resp. $(b^{(3)}, c^{(1)})$, $(c^{(3)}, \hat{a}^{(1)})$) if and only if $(a, b) \in E$ (resp. if $(b, c) \in E$, $(c, a) \in E$). This creates a graph with $O(n)$ vertices

We now proceed to more edge insertions of the reduction. In the outer loop we iterate over $d_i \in D$ in increasing order.

1. For each $b_j \in N(d_i)$, insert edges $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$ and $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ with weight $(i+1) \cdot n$.

2. For each $c_j \in N(d_i)$, insert edges $(c_j^{(1)}, c_{j+i \bmod n}^{(2)})$ and $(c_{j+i \bmod n}^{(2)}, c_j^{(3)})$ with weight $(i+1) \cdot n$.
3. In the inner loop, we iterate over $a_k \in A$ in increasing order.

- (a) If $a_k \in N(d_i)$, insert the following edges of weight $i \cdot n + (k+1)$:

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}).$$

- (b) Query \mathcal{A} to get $\mathbf{b}(s, t)$. If $\mathbf{b}(s, t) \geq i \cdot n + (k+1)$, return TRUE.

Then, return FALSE after all iterations.

Clearly, the bottleneck of this algorithm is to run \mathcal{A} , so the running time is $O(n^{4-c})$.

We show the following lemma in order to show correctness:

Lemma 4.17. *Consider the query when iterating over d_i and a_k . The result of the query is $\mathbf{b}(s, t) \geq i \cdot n + (k+1)$ if and only if (d_i, a_k) are in a 4-clique.*

Proof. Suppose d_i, a_k are in a 4-clique with vertices b_j, c_ℓ . Then, the following (s, t) -bottleneck path exists in the graph and has capacity $i \cdot n + (k+1)$:

$$(s, a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)}, b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)}, c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)}, \hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}, t).$$

We verify that this path exists and has the required capacity. $(s, a_k^{(1)})$ has capacity ∞ . Since $a_k \in N(d_i)$, each edge in the sub-path $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)})$ exists and has capacity $i \cdot n + (k+1)$. Next, $(a_k, b_j) \in E$ implies that the edge $(a_k^{(3)}, b_j^{(1)})$ exists with ∞ capacity. Since $b_j \in N(d_i)$, each edge in the sub-path $(b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ exists and has capacity $(i+1) \cdot n$. Then, $(b_j, c_\ell) \in E$ implies that the edge $(b_j^{(3)}, c_\ell^{(1)})$ exists with ∞ capacity. Since $c_\ell \in N(d_i)$, each edge in the sub-path $(c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)})$ exists and has capacity $(i+1) \cdot n$. Then, $(c_\ell, a) \in E$ implies that the edge $(c_\ell^{(3)}, \hat{a}_k^{(1)})$ exists with ∞ capacity. Finally, as $a_k \in N(d_i)$, we have the sub-path $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$ with capacity $i \cdot n + (k+1)$. We conclude by observing that the edge $(\hat{a}_k^{(3)}, t)$ has ∞ capacity.

Conversely, suppose $\mathbf{b}(s, t) \geq i \cdot n + (k+1)$. Then, since edges are directed left to right, there is a path of the form

$$(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, b_{i_4}^{(1)}, b_{i_5}^{(2)}, b_{i_6}^{(3)}, c_{i_7}^{(1)}, c_{i_8}^{(2)}, c_{i_9}^{(3)}, \hat{a}_{i_{10}}^{(1)}, \hat{a}_{i_{11}}^{(2)}, \hat{a}_{i_{12}}^{(3)}, t)$$

achieving this capacity. At the time of the query in the i -th iteration over D and k -th iteration over A in the inner-loop, the maximum weight of any edge between $A_1 \times A_2$, $A_2 \times A_3$, $\hat{A}_1 \times \hat{A}_2$, and $\hat{A}_2 \times \hat{A}_3$ is $i \cdot n + (k+1)$. Furthermore, the only edges attaining this weight are $(a_k^{(1)}, a_{k+i \bmod n}^{(2)})$, $(a_{k+i \bmod n}^{(2)}, a_k^{(3)})$, $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)})$, and $(\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$. If these edges exist, then $(a_k, d_i) \in E$ and $k = i_1 = i_3 = i_{10} = i_{12}$. The maximum weight of any edge between $B_1 \times B_2$ and $B_2 \times B_3$ is $(i+1) \cdot n$. Furthermore, the only edges attaining this weight are $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$, $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ for $b_j \in N(d_i)$. In particular, we have $i_4 = i_6 = j$ for some $b_j \in N(d_i)$. By a similar argument, $i_7 = i_9 = \ell$ for some $c_\ell \in N(d_i)$.

Finally, it remains to show a_k, b_j, c_ℓ form a triangle. Since $\text{wt}(a_k^{(3)}, b_j^{(1)}) = \infty$ we have $(a_k, b_j) \in E$. A similar argument shows that $(a_k, c_\ell), (b_j, c_\ell) \in E$. \square

From Lemma 4.17, the above procedure solves the 4-Clique Detection instance, contradicting the Combinatorial 4-Clique hypothesis. \square

Using an essentially identical reduction, we show that under the OMv3 hypothesis, any algorithm computing stBP with polynomial preprocessing requires $n^{\omega+1-o(1)}$ total update time.

Theorem 4.18. *Under the OMv3 hypothesis, any algorithm computing incremental/decremental stBP on directed graphs with polynomial preprocessing time requires $n^{\omega+1-o(1)}$ total update time.*

Proof. Suppose for contradiction there is an incremental stBP algorithm \mathcal{A} with polynomial preprocessing time and total update time $O(n^{\omega+1-c})$ for some $c > 0$. The decremental case can be handled using similar modifications as Theorem 3.8. We design an efficient algorithm for OMv3.

In the preprocessing phase we receive a Boolean matrix A . We construct a graph with vertices

$$\{s\} \cup (A_1 \cup A_2 \cup A_3) \cup (B_1 \cup B_2 \cup B_3) \cup (C_1 \cup C_2 \cup C_3) \cup (\hat{A}_1 \cup \hat{A}_2 \cup \hat{A}_3) \cup \{t\}.$$

Each vertex set $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, \hat{A}_1, \hat{A}_2, \hat{A}_3$ has n vertices, each indexed by $[n]$. Insert edges $(s, a^{(1)})$ and $(\hat{a}^{(3)}, t)$ with ∞ weight for all $a \in A$. Furthermore, insert edges with ∞ weight between $(a^{(3)}, b^{(1)})$ (resp. $(b^{(3)}, c^{(1)}), (c^{(3)}, \hat{a}^{(1)})$) if and only if $A[a, b] = \text{TRUE}$ (resp. if $A[b, c] = \text{TRUE}, A[a, c] = \text{TRUE}$). We then run the preprocessing step of \mathcal{A} on this $O(n)$ -vertex graph.

We now proceed to the dynamic phase of the reduction. Throughout, we assume that both queries and coordinates are indexed starting at 0 (just as vertices as indexed from 0 in Theorem 4.16). Suppose we have received the i -th query, $\vec{u}_i, \vec{v}_i, \vec{w}_i$.

1. For all j with $\vec{v}_i[j] = \text{TRUE}$, insert edges $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$ and $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ with weight $(i+1) \cdot n$.
2. For all j with $\vec{w}_i[j] = \text{TRUE}$, insert edges $(c_j^{(1)}, c_{j+i \bmod n}^{(2)})$ and $(c_{j+i \bmod n}^{(2)}, c_j^{(3)})$ with weight $(i+1) \cdot n$.
3. Now, we iterate over the coordinates k of vector \vec{v}_i in increasing order.

- (a) If $\vec{u}_i[k] = \text{TRUE}$, insert edges of weight $i \cdot n + (k+1)$ between

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}).$$

- (b) Compute $b(s, t)$ by querying \mathcal{A} . If $\mathbf{b}(s, t) \geq i \cdot n + (k+1)$, return TRUE .

Otherwise, return FALSE if none of the queries return TRUE .

During the dynamic phase, the bottleneck of the algorithm is running \mathcal{A} , so the total update and query time is $O(n^{\omega+1-c})$.

The following lemma shows that our algorithm correctly computes each OMv3 query.

Lemma 4.19. *Consider the query $(\vec{u}_i, \vec{v}_i, \vec{w}_i)$ and the query after updating the k -th coordinate of \vec{u} . The result of the query is $\mathbf{b}(s, t) \geq i \cdot n + (k+1)$ if and only if*

$$\bigvee_{j, \ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k, j] \wedge A[j, \ell] \wedge A[\ell, k]) = \text{TRUE}.$$

Since this lemma is essentially identical to Lemma 4.17, we defer the proof to Appendix A.1.

Note that at the query after the i -th query and the k -th iteration of the inner loop, we have constructed the same graph as in Theorem 4.16. In particular, Lemma 4.19 shows that this query returns true if and only if

$$\bigvee_{j,\ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k,j] \wedge A[j,\ell] \wedge A[\ell,k]) = \text{TRUE}.$$

Thus, if the above is satisfied for any k , we return TRUE. Otherwise, we return FALSE, answering the query correctly in either case. Therefore, using \mathcal{A} we have obtained an algorithm for the OMv3 instance with polynomial preprocessing and total update and query time $O(n^{\omega+1-c})$, contradicting the OMv3 hypothesis. \square

Furthermore, we show that any partially dynamic SSBP algorithm computes the min-witness matrix product with n updates. In particular, the following reduction shows that any algorithm with amortized update time $O(n^{1.5-c})$ for $c > 0$ gives an improved algorithm for computing the min-witness product when $\omega = 2$. We leave as an interesting open question how to extend the following lower bound to n^2 updates.

Theorem 4.20. *Under the Minimum-Witness Product hypothesis, any algorithm computing incremental/decremental SSBP on directed graphs requires $n^{2.5-o(1)}$ total update time over n updates. Equivalently, any algorithm computing SSBP on directed graphs requires $n^{1.5-o(1)}$ amortized update time over n updates.*

Proof. Suppose for contradiction there is an incremental SSBP algorithm with pre-processing time $O(n^{2.5-c})$ and amortized update time $O(n^{1.5-c})$ for some $c > 0$, describing modifications for the decremental version as necessary. We design an algorithm computing a min-witness product of $n \times n$ matrices A, B in $O(n^{2.5-c})$ time.

Construct a 4 layered graph

$$\{s\} \cup U \cup V \cup W,$$

where $|U| = |V| = |W| = n$ and index each vertex set from 1 to n . We initialize the graph as follows. For every entry $A[i,k] = \text{TRUE}$, add an edge (u_i, v_k) with weight $i \cdot n - (k-1)$. For every entry $B[k,j] = \text{TRUE}$, add an edge (v_k, w_j) with weight ∞ . In the decremental reduction, we additionally add edges (s, v_i) of weight ∞ for all i . This creates a graph with n vertices, so we can preprocess in $O(n^{2.5-c})$ time.

We now proceed to the dynamic phase. We iterate over $i \in [n]$ in increasing order. For each i , we add an edge (s, u_i) of weight ∞ and query SSBP. Suppose $\mathbf{b}(s, w_j) > (i-1) \cdot n$. Then, we note the minimum witness of i, j as $i \cdot n + 1 - \mathbf{b}(s, w_j)$. Otherwise, we note that there is no witness for i, j . In the decremental setting, we iterate over $[n]$ in decreasing order and instead delete the edge (s, u_i) before each query.

We argue that this correctly computes the minimum-witness product. Consider the query after inserting edge (s, u_i) . We claim $\mathbf{b}(s, w_j) > (i-1) \cdot n$ if and only if there is some k for which $A[i,k] = B[k,j] = \text{TRUE}$, and furthermore this path has bottleneck capacity $i \cdot n - (k-1)$. In particular, this shows that $i \cdot n + 1 - \mathbf{b}(s, w_j)$ is precisely the minimum such k (if such a path exists) and otherwise there is no witness.

Suppose there is some k for which $A[i,k] = B[k,j] = \text{TRUE}$. Then, there is the path (s, u_i, v_k, w_j) with bottleneck capacity $i \cdot n + k$. Conversely, if there is a path with bottleneck capacity at least $(i-1) \cdot n$, then it must have the form (s, u_i, v_k, w_j) for some k since the graph is directed and there are no edges from $s_{i'}$ for $i' > i$. Thus, for this k it must be the case that $A[i,k] = B[k,j] = \text{TRUE}$, as desired.

The overall reduction makes n updates and computes the minimum-witness product in $O(n^{2.5-c})$, contradicting the Minimum-Witness Product hypothesis. \square

5 Earliest Arrivals (or Non-decreasing Paths)

For the earliest arrivals problem, we obtain a simple, efficient algorithm for the incremental single-source problem via a reduction to reachability (Theorem 5.4). In contrast, under the weight dynamic model, we show that any partially dynamic algorithm can be reduced from 4-Clique Detection. The earliest arrivals problem (sometimes also called non-decreasing paths problem) is defined below.

Definition 5.1 (Earliest Arrivals). Let $G = (V, E, \text{wt})$ be a directed, weighted graph. For any path P is a **valid itinerary** if the edges on the path are non-decreasing in weight. For a valid itinerary, let $\mathbf{a}(P) = \max_{e \in P} \text{wt}(e)$ be the weight of the final edge. For any pair of nodes u, v , let $\mathbf{ap}(u, v)$ denote the **earliest arrival**, the valid itinerary between u, v minimizing $\mathbf{a}(P)$ (breaking ties arbitrarily). Let $\mathbf{a}(u, v) = \mathbf{a}(\mathbf{ap}(u, v))$.

The (s, t) -**Earliest Arrival** problem (stEA) asks to compute $\mathbf{a}(s, t)$ for fixed nodes s, t . The **Single Source Earliest Arrivals** problem (SSEA) asks to compute $\mathbf{a}(s, v)$ for a single source s and all $v \in V$. The **All Pairs Earliest Arrivals** problem (APEA) asks to compute $\mathbf{a}(u, v)$ for all $u, v \in V$.

5.1 Decremental Algorithm for SSEA

We give a simple near-linear total time algorithm for partially dynamic SSEA using a reduction from SSEA to Single-Source Reachability (SSR) in sparse graphs:

Lemma 5.2 (Reduction from SSEA to SSR). *We can reduce an instance of SSEA in n -vertex m -edge graph G to an instance of SSR in an $O(m)$ -vertex $O(m)$ -edge graph. The reduction runs in near-linear time.*

Proof. Let $G = (V, E, \text{wt})$ be the graph for SSEA. We create the following graph for SSR.

1. For every $e \in E$, we create two vertices p_e and q_e . The vertex set of the SSR instance is $\{s\} \cup \{p_e, q_e\}_{e \in E}$.
2. For every $e \in E$, we add an edge from p_e to q_e .
3. For every $v \in E$, let i_1, i_2, \dots, i_x be the set of incoming edges, and let o_1, o_2, \dots, o_y be the set of outgoing edges. We sort the vertices $q_{i_1}, q_{i_2}, \dots, q_{i_x}, p_{o_1}, p_{o_2}, \dots, p_{o_y}$ based on the weights of their corresponding edges (if there is a tie, we rank incoming edges before outgoing edges since we require the path to be non-decreasing), and add a chain from the smallest to the largest. This way, q_{i_a} can reach p_{o_b} if and only if $\text{wt}(i_a) \leq \text{wt}(o_b)$.
4. For every $(s, v) \in E$, we add an edge from s to $p_{(s, v)}$.

We observe that the reduction satisfies the following property.

Claim 5.3. *Let e_1, e_2, \dots, e_x be the set of incoming edges of v . Then*

$$\mathbf{a}(s, v) = \min\{\text{wt}(e_i) \text{ s.t. } p_{e_i} \text{ is reachable from } s\}.$$

Then, to answer a query it suffices to keep track of the minimum e_i such that p_{e_i} is reachable from s . To bound the run-time, note that the set of reachable p_{e_i} can only increase/decrease, so maintaining all the relevant heaps requires only $\tilde{O}(m)$ total time. \square

Proof of Claim 5.3. Suppose p_{e_i} is reachable from s . Then, there is a valid itinerary (non-decreasing path) from s to the origin of e_i in the SSEA instance, since every edge (p_e, q_e) taken is an edge in the graph and they must be taken in non-decreasing order by the chain construction. In particular, q_{e_i} is reachable from s so that v is reachable from s via a non-decreasing path ending at e_i , or $\mathbf{a}(s, v) \leq \text{wt}(e_i)$.

Otherwise, let $P = \mathbf{ap}(s, v)$ to be the earliest arrival path with final edge e_j . We claim p_{e_j} is reachable from s . This follows easily as the path is a valid itinerary (non-decreasing path) so that there is a path in the constructed graph that follows the appropriate edges (p_e, q_e) as well as edges within the chain. \square

By maintaining the reduction dynamically, we can obtain a faster algorithm for partially dynamic SSEA.

Theorem 5.4. *There is an algorithm for incremental/decremental SSEA in $\tilde{O}(m)$ total time. This algorithm holds against an oblivious, adaptive adversary.*

Proof. We can maintain the reduction in Lemma 5.2 dynamically: For inserting e , we need to insert the edge from p_e to q_e , and appropriately insert p_e and q_e to the chains (note that we do not need to delete edges that are already on the chain, i.e., if we have a chain $1 \rightarrow 2 \rightarrow 4 \rightarrow 5$, and need to insert 3, we just add (2, 3) and (3, 4) and do not need to delete (2, 4)). For deleting e , we only need to delete the edge from p_e to q_e , and do not need to update the chain.

Incremental SSR can be solved (easily) in $O(m)$ total time. Decremental SSR can be solved in $\tilde{O}(m)$ total time [BPW19]. \square

5.2 Conditional Lower Bound for Weight Dynamic Earliest Arrivals

Theorem 5.5. *Under the Combinatorial 4-Clique hypothesis, any combinatorial algorithm computing stEA with incremental/decremental weights requires $n^{2-o(1)}$ amortized update time. Furthermore, each edge weight is modified only once.*

Proof. We describe the reduction for incremental stEA (weight non-decreasing), noting that the decremental reduction can be obtained by reversing the reduction.

Suppose for contradiction there is an algorithm \mathcal{A} with $O(n^{2-c})$ amortized update time for some $c > 0$. Consider a 4-Clique Detection instance with vertex sets A, B, C, D of size n . Throughout the reduction, assume that the vertex sets A, B, C, D are indexed from 1 to n . We design an algorithm with total time $O(n^{4-c})$ for some positive $c > 0$, contradicting the Combinatorial 4-Clique hypothesis.

We construct a graph with vertices

$$\{s\} \cup (A_1 \cup A_2 \cup A_3) \cup (B_1 \cup B_2 \cup B_3) \cup (C_1 \cup C_2 \cup C_3) \cup (\hat{A}_1 \cup \hat{A}_2 \cup \hat{A}_3) \cup \{t\}.$$

Each vertex set $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, \hat{A}_1, \hat{A}_2$ and \hat{A}_3 has n vertices. For each $a \in A$, we create a copy $a^{(i)} \in A_i$ and a copy $\hat{a}^{(i)} \in \hat{A}_i$ for $i \in \{1, 2, 3\}$. Similarly, for each $b \in B$ and $c \in C$, we create a copy $b^{(i)} \in B_i$ and $c^{(i)} \in C_i$ for $i \in \{1, 2, 3\}$. Initially, insert edges $(s, a^{(1)})$ with weight 1 and edges $(\hat{a}^{(3)}, t)$ with weight 9 for all $a \in A$. Furthermore, insert edges with weight 3 between $(a^{(3)}, b^{(1)})$ (resp. weight 5 between $(b^{(3)}, c^{(1)})$, weight 7 between $(c^{(3)}, \hat{a}^{(1)})$) if and only if $(a, b) \in E$ (resp. $(b, c) \in E$, $(c, a) \in E$). We also insert all edges between $A_1 \times A_2, A_2 \times A_3, B_1 \times B_2, B_2 \times B_3, C_1 \times C_2, C_2 \times C_3, \hat{A}_1 \times \hat{A}_2, \hat{A}_2 \times \hat{A}_3$ with weight 0. This creates a graph with $O(n)$ vertices.

We now proceed to the more updates of the reduction. In the outer loop we iterate over $d_i \in D$ in increasing order.

1. For each $b_j \in N(d_i)$, increase the weights of $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$ and $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ to 4.
2. For each $c_j \in N(d_i)$, increase the weights of $(c_j^{(1)}, c_{j+i \bmod n}^{(2)})$ and $(c_{j+i \bmod n}^{(2)}, c_j^{(3)})$ to 6.
3. In the inner loop, we iterate over $a_k \in A$ in increasing order.

(a) If $a_k \in N(d_i)$, increase the weights of

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)})$$

to 2 and increase the weights of

$$(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$$

to 8.

(b) Computing $\mathbf{a}(s, t)$ by querying \mathcal{A} . If $\mathbf{a}(s, t) = 9$, return TRUE.
(c) After the query, increase the weights of

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$$

to 10.

4. After all queries of this round, increase the weights of

$$(b_j^{(1)}, b_{j+i \bmod n}^{(2)}), (b_{j+i \bmod n}^{(2)}, b_j^{(3)}), (c_j^{(1)}, c_{j+i \bmod n}^{(2)}), (c_{j+i \bmod n}^{(2)}, c_j^{(3)})$$

for all $b_j, c_j \in N(d_i)$ to 10.

Then, return FALSE after all iterations.

For each d , we increase/decrease the weights of at most $16n$ edges, so the total number of updates is $O(n^2)$. The bottleneck of this algorithm is to run \mathcal{A} , so the running time is $O(n^{4-c})$.

Lemma 5.6. *Consider the query when iterating over d_i and a_k . The result of the query is $\mathbf{a}(s, t) = 9$ if and only if (d_i, a_k) are in a 4-clique.*

Proof. Note that all edges into t have weight 9. Thus, the arrival time $\mathbf{a}(s, t) \in \{9, \infty\}$, and $\mathbf{a}(s, t) = 9$ if and only if there is a valid itinerary from s to t .

Suppose d_i, a_k are in a 4-clique with vertices b_j, c_ℓ . Then, the following valid itinerary exists in the graph

$$(s, a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)}, b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)}, c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)}, \hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}, t).$$

We verify that this path exists and is valid (i.e. has non-decreasing weights). $(s, a_k^{(1)})$ has weight 1. Since $a_k \in N(d_i)$, each edge in the sub-path $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)})$ has weight 2. Next, $(a_k, b_j) \in E$ implies that the edge $(a_k^{(3)}, b_j^{(1)})$ exists with weight 3. Since $b_j \in N(d_i)$, each edge in the sub-path $(b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ has weight 4. Then, $(b_j, c_\ell) \in E$ implies that the edge $(b_j^{(3)}, c_\ell^{(1)})$ exists with weight 5. Since $c_\ell \in N(d_i)$, each edge in the sub-path $(c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)})$ has weight 6. Then, $(c_\ell, a) \in E$ implies that the edge $(c_\ell^{(3)}, \hat{a}_k^{(1)})$ exists with weight 7. Finally, as $a_k \in N(d_i)$, we have the sub-path $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$ with weight 8. We conclude by observing that edge $(\hat{a}_k^{(3)}, t)$ has weight 9.

Conversely, suppose $\mathbf{a}(s, t) = 9$. Then, since edges are directed left to right, there is a path of the form

$$(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, b_{i_4}^{(1)}, b_{i_5}^{(2)}, b_{i_6}^{(3)}, c_{i_7}^{(1)}, c_{i_8}^{(2)}, c_{i_9}^{(3)}, \hat{a}_{i_{10}}^{(1)}, \hat{a}_{i_{11}}^{(2)}, \hat{a}_{i_{12}}^{(3)}, t)$$

and is a valid itinerary.

At the time of the query in the i -th iteration over D and k -th iteration over A in the inner-loop, the only edges in $A_1 \times A_2, A_2 \times A_3, \hat{A}_1 \times \hat{A}_2, \hat{A}_2 \times \hat{A}_3$ with weight 2 are $(a_k^{(1)}, a_{k+i \bmod n}^{(2)})$, $(a_{k+i \bmod n}^{(2)}, a_k^{(3)})$, $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)})$, and $(\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$. Note that all other weights have weight either 0 or 10, and thus cannot participate in any valid itinerary (since edges into A_1 have weight 1 and edges out of A_3 have weight 3). If these edges exist with weight 2, then $(a_k, d_i) \in E$ and $k = i_1 = i_3 = i_{10} = i_{12}$. The only valid edges between $B_1 \times B_2$ and $B_2 \times B_3$ with weight 4 are $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$, $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ for $b_j \in N(d_i)$ (and no other edges can participate in a valid itinerary since edges into B_1 have weight 3 and edges out of B_3 have weight 5). Since these paths are vertex disjoint, we have $i_4 = i_6 = j$ for some $b_j \in N(d_i)$. By a similar argument, $i_7 = i_9 = \ell$ for some $c_\ell \in N(d_i)$.

Finally, it remains to show a_k, b_j, c_ℓ form a triangle. Since $(a_k^{(3)}, b_j^{(1)})$ form an edge, we have $(a_k, b_j) \in E$. A similar argument shows that $(a_k, c_\ell), (b_j, c_\ell) \in E$. \square

From Lemma 5.6, the above procedure solves the 4-Clique Detection instance, contradicting the Combinatorial 4-Clique hypothesis. \square

Again, using an essentially identical reduction we obtain a lower bound for arbitrary algorithms under the OMv3 hypothesis.

Theorem 5.7. *Under the OMv3 hypothesis, any algorithm computing incremental/decremental stEA on directed graphs with polynomial preprocessing time requires $n^{\omega-1-o(1)}$ amortized update and query time. Furthermore, each edge weight is incremented only once.*

We defer the proof to Appendix A.1. In particular, for stEA, we have shown that modifying edge weights and inserting/removing edges lead to different complexities in the partially dynamic model. Specifically, while edge insertions/deletions can be handled efficiently, it is hard to maintain earliest arrivals under weight modifications. Since the weight-dynamic is more general than the insertion/deletion model, our lower bounds for both shortest paths and bottleneck paths continue to hold in the weight-dynamic model.

Remark 5.8. In both Theorems 5.5 and 5.7, we have restricted the lower bound to only modify each edge weight once. We can extend our lower bound to handle many increases per edge as follows. After modifying each edge weight once, increase all edge weights by some large enough constant (e.g. 10) so that all edges that started at weight 0 are now weight 10, all weights that started at weight 1 are weight 11, and so on, and run the reduction again. Using this modification, we can compute Unbalanced 4-Clique, or OMv3 with arbitrarily (polynomially) many queries. Note that this maintains the fact that all weight modifications are increase only.

References

[ABDN18] Amir Abboud, Karl Bringmann, Holger Dell, and Jesper Nederlof. More consequences of falsifying SETH and the orthogonal vectors conjecture. In *50th Annual ACM SIGACT Symposium on Theory of Computing, STOC, 2018*. 3

[ACK17] Ittai Abraham, Shiri Chechik, and Sebastian Krinninger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2017. [1](#)

[AGI⁺19] Amir Abboud, Loukas Georgiadis, Giuseppe F. Italiano, Robert Krauthgamer, Nikos Parotsidis, Ohad Trabelsi, Przemyslaw Uznanski, and Daniel Wolleb-Graf. Faster algorithms for all-pairs bounded min-cuts. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, 2019. [4](#)

[AH74] Alfred V. Aho and John E. Hopcroft. *The design and analysis of computer algorithms*. Pearson Education India, 1974. [35](#)

[AHR⁺19] Bertie Ancona, Monika Henzinger, Liam Roditty, Virginia Vassilevska Williams, and Nicole Wein. Algorithms and hardness for diameter in dynamic graphs. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, 2019. [2](#)

[AMO93] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. Network flows: Theory, algorithms and applications. *New Jersey: Rentice-Hall*, 1993. [3](#)

[AMV20] Kyriakos Axiotis, Aleksander Madry, and Adrian Vladu. Circulation control for faster minimum cost flow in unit-capacity graphs. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2020. [1](#)

[AV14] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *55th Annual Symposium on Foundations of Computer Science, FOCS*, 2014. [2](#), [6](#), [7](#), [19](#)

[AVW14] Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In *41st International Colloquium on Automata, Languages, and Programming, ICALP*, 2014. [3](#)

[BC16] Aaron Bernstein and Shiri Chechik. Deterministic decremental single source shortest paths: beyond the $o(mn)$ bound. In *48th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2016. [1](#)

[BC17] Aaron Bernstein and Shiri Chechik. Deterministic partially dynamic single source shortest paths for sparse graphs. In *28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2017. [1](#)

[BDT16] Arturs Backurs, Nishanth Dikkala, and Christos Tzamos. Tight hardness results for maximum weight rectangles. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP*, 2016. [3](#)

[Bel58] Richard Bellman. On a routing problem. *Quarterly of applied mathematics*, 16(1):87–90, 1958. [1](#)

[Ber16] Aaron Bernstein. Maintaining shortest paths under deletions in weighted directed graphs. *SIAM J. Comput.*, 45(2), 2016. [1](#)

[Ber17] Aaron Bernstein. Deterministic partially dynamic single source shortest paths in weighted graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP*, 2017. [1](#)

[BFvdH⁺24] Karl Bringmann, Nick Fischer, Ivor van der Hoog, Evangelos Kipouridis, Tomasz Kociumaka, and Eva Rotenberg. Dynamic dynamic time warping. In *2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 208–242, 2024. [3](#)

[BGL17] Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular expression membership testing. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2017. [4](#)

[BGMW20] Karl Bringmann, Paweł Gawrychowski, Shay Mozes, and Oren Weimann. Tree edit distance cannot be computed in strongly subcubic time (unless APSP can). *ACM Trans. Algorithms*, 16(4):48:1–48:22, 2020. [3](#)

[BGS20] Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental reachability, scc, and shortest paths via directed expanders and congestion balancing. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2020. [1](#)

[BGS21] Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental SSSP and approximate min-cost flow in almost-linear time. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2021. [1](#)

[BNW22] 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. [1](#)

[BPW19] Aaron Bernstein, Maximilian Probst, and Christian Wulff-Nilsen. Decremental strongly-connected components and single-source reachability in near-linear time. In *51st Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2019. [33](#), [40](#)

[BR11] Aaron Bernstein and Liam Roditty. Improved dynamic algorithms for maintaining approximate shortest paths under deletions. In *22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2011. [1](#)

[Bra] Jan van den Brand. Complexity term balancer. www.ocf.berkeley.edu/~vdbrand/complexity/. Tool to balance complexity terms depending on fast matrix multiplication. [6](#)

[BT17] Arturs Backurs and Christos Tzamos. Improving viterbi is hard: Better runtimes imply faster clique algorithms. In *34th International Conference on Machine Learning, ICML*, 2017. [3](#)

[Cha05] Timothy M. Chan. All-pairs shortest paths with real weights in $O(n^3 / \log n)$ time. In *9th International Workshop on Algorithms and Data Structures, WADS*, 2005. [1](#)

[Cha10a] Timothy M. Chan. More algorithms for all-pairs shortest paths in weighted graphs. *SIAM J. Comput.*, 39(5), 2010. [1](#), [3](#)

[Cha10b] Timothy M. Chan. A (slightly) faster algorithm for klee’s measure problem. *Comput. Geom.*, 43(3):243–250, 2010. [4](#)

[CK19] Julia Chuzhoy and Sanjeev Khanna. A new algorithm for decremental single-source shortest paths with applications to vertex-capacitated flow and cut problems. In *51st Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2019. [1](#)

[CKL07] Artur Czumaj, Miroslaw Kowaluk, and Andrzej Lingas. Faster algorithms for finding lowest common ancestors in directed acyclic graphs. *Theor. Comput. Sci.*, 380(1-2), 2007. [5](#), [24](#), [25](#)

[CKL⁺22] 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. [1](#)

[CKT⁺16] Shiri Chechik, Haim Kaplan, Mikkel Thorup, Or Zamir, and Uri Zwick. Bottleneck paths and trees and deterministic graphical games. In *33rd Symposium on Theoretical Aspects of Computer Science (STACS 2016)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2016. [3](#)

[CL09] Artur Czumaj and Andrzej Lingas. Finding a heaviest vertex-weighted triangle is not harder than matrix multiplication. *SIAM J. Comput.*, 39(2), 2009. [3](#)

[CZ23] Shiri Chechik and Tianyi Zhang. Faster deterministic worst-case fully dynamic all-pairs shortest paths via decremental hop-restricted shortest paths. In *ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2023. [1](#)

[DI04] Camil Demetrescu and Giuseppe F. Italiano. A new approach to dynamic all pairs shortest paths. *J. ACM*, 51(6), 2004. [1](#), [34](#)

[Dij59] Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 50:269–271, 1959. [1](#)

[DJW19] Ran Duan, Ce Jin, and Hongxun Wu. Faster algorithms for all pairs non-decreasing paths problem. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, 2019. [3](#)

[DLX18] Ran Duan, Kaifeng Lyu, and Yuanhang Xie. Single-source bottleneck path algorithm faster than sorting for sparse graphs. In *45th International Colloquium on Automata, Languages, and Programming, ICALP*, 2018. [3](#)

[Dob90] Włodzimierz Dobosiewicz. A more efficient algorithm for the min-plus multiplication. *Int. J. Computer Math.*, 32, 1990. [1](#)

[DP09] Ran Duan and Seth Pettie. Fast algorithms for (max, min)-matrix multiplication and bottleneck shortest paths. In *20th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2009. [3](#), [28](#), [29](#), [30](#)

[DSST89] James R. Driscoll, Neil Sarnak, Daniel Dominic Sleator, and Robert Endre Tarjan. Making data structures persistent. *J. Comput. Syst. Sci.*, 38(1):86–124, 1989. [34](#)

[DWZ23] Ran Duan, Hongxun Wu, and Renfei Zhou. Faster matrix multiplication via asymmetric hashing. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2023. [2](#)

[EG04] Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theor. Comput. Sci.*, 326(1-3), 2004. [4](#)

[EK72] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, 1972. [3](#)

[ES81] Shimon Even and Yossi Shiloach. An on-line edge-deletion problem. *J. ACM*, 28(1):1–4, 1981. [1](#), [33](#)

[Fin24] Jeremy T. Fineman. Single-source shortest paths with negative real weights in $\tilde{O}(mn^{8/9})$ time. In *the 56th Annual ACM Symposium on Theory of Computing, STOC*, page to appear, 2024. [1](#)

[For56] Jr Ford. Lr network flow theory. *Paper P-923, The RAND Corporation, Santa Monica, California*, 1956. [1](#)

[Fre76] Michael L. Fredman. New bounds on the complexity of the shortest path problem. *SIAM J. Comput.*, 5(1):83–89, 1976. [1](#)

[FT87] Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987. [1](#)

[FW93] Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *J. Comput. Syst. Sci.*, 47(3):424–436, 1993. [1](#)

[FW94] Michael L. Fredman and Dan E. Willard. Trans-dichotomous algorithms for minimum spanning trees and shortest paths. *J. Comput. Syst. Sci.*, 48(3):533–551, 1994. [1](#)

[Gab83] Harold N. Gabow. Scaling algorithms for network problems. In *24th Annual Symposium on Foundations of Computer Science, FOCS*, 1983. [1](#)

[GIL⁺21] Fabrizio Grandoni, Giuseppe F. Italiano, Aleksander Lukasiewicz, Nikos Parotsidis, and Przemyslaw Uznanowski. All-pairs LCA in dags: Breaking through the $O(n^{2.5})$ barrier. In *2021 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2021. [28](#)

[Gol95] Andrew V. Goldberg. Scaling algorithms for the shortest paths problem. *SIAM J. Comput.*, 24(3):494–504, 1995. [1](#)

[GT88] Harold N Gabow and Robert E Tarjan. Algorithms for two bottleneck optimization problems. *Journal of Algorithms*, 9(3):411–417, 1988. [3](#)

[GT89] Harold N. Gabow and Robert Endre Tarjan. Faster scaling algorithms for network problems. *SIAM J. Comput.*, 18(5):1013–1036, 1989. [1](#)

[GV21] Aparna Gupta and Vinod Vaikuntanathan. The fine-grained hardness of sparse linear regression. *CoRR*, abs/2106.03131, 2021. [3](#)

[GVW20] Maximilian Probst Gutenberg, Virginia Vassilevska Williams, and Nicole Wein. New algorithms and hardness for incremental single-source shortest paths in directed graphs. In *52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2020. [1](#), [2](#), [4](#), [5](#), [6](#), [10](#), [19](#), [20](#)

[GW20a] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Decremental SSSP in weighted digraphs: Faster and against an adaptive adversary. In *31st ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2020. [1](#)

[GW20b] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Deterministic algorithms for decremental approximate shortest paths: Faster and simpler. In *31st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2020. [1](#)

[GW20c] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Fully-dynamic all-pairs shortest paths: Improved worst-case time and space bounds. In *2020 ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2020. [1](#)

[Hag00] Torben Hagerup. Improved shortest paths on the word RAM. In *27th International Colloquium on Automata, Languages and Programming, ICALP*, pages 61–72, 2000. [1](#)

[Han04] Yijie Han. Improved algorithm for all pairs shortest paths. *Inf. Process. Lett.*, 91(5), 2004. [1](#)

[Han06] Yijie Han. Achieving $O(n^3 / \log n)$ time for all pairs shortest paths by using a smaller table. In *21st International Conference on Computers and Their Applications, CATA*, 2006. [1](#)

[HKN14a] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs. In *46th Annual Symposium on the Theory of Computing, STOC*, 2014. [1](#)

[HKN14b] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A subquadratic-time algorithm for decremental single-source shortest paths. In *25th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2014. [1](#)

[HKN18] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. *J. ACM*, 65(6), 2018. [1](#)

[HKNS15] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *47th Annual ACM on Symposium on Theory of Computing, STOC*, 2015. [4](#), [19](#)

[HLSW23] Zhiyi Huang, Yaowei Long, Thatchaphol Saranurak, and Benyu Wang. Tight conditional lower bounds for vertex connectivity problems. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1384–1395, 2023. [4](#)

[HPS21] Monika Henzinger, Ami Paz, and Stefan Schmid. On the complexity of weight-dynamic network algorithms. In *IFIP Networking Conference, IFIP Networking*, 2021. [5](#)

[HRWN15] Jacob Holm, Eva Rotenberg, and Christian Wulff-Nilsen. Faster fully-dynamic minimum spanning forest. In *Algorithms-ESA 2015: 23rd Annual European Symposium*. Springer, 2015. [33](#)

[HT12] Yijie Han and Tadao Takaoka. An $O(n^3 \log \log n / \log^2 n)$ time algorithm for all pairs shortest paths. In *13th Scandinavian Workshop on Algorithm Theory, SWAT*, 2012. [1](#)

[Hu61] TC Hu. The maximum capacity route problem. *Oper. Res.*, 9(6), 1961. [3](#), [33](#)

[IKR81] Alon Itai, Alan G Konheim, and Michael Rodeh. A sparse table implementation of priority queues. In *8th Colloquium on Automata, Languages and Programming, ICALP*. Springer, 1981. [34](#)

[IR78] Alon Itai and Michael Rodeh. Finding a minimum circuit in a graph. *SIAM J. Comput.*, 7(4), 1978. [2](#)

[JX22] Ce Jin and Yinzhan Xu. Tight dynamic problem lower bounds from generalized BMM and OMv. In *54th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2022. [2](#), [4](#), [7](#), [24](#)

[KL22] Miroslaw Kowaluk and Andrzej Lingas. A multi-dimensional matrix product - A natural tool for parameterized graph algorithms. *Algorithms*, 15(12):448, 2022. [4](#), [5](#), [24](#)

[Li19] Jason Li. Faster minimum k-cut of a simple graph. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2019. [4](#)

[LPV20] 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. [5](#), [24](#)

[LVW18] Andrea Lincoln, Virginia Vassilevska Williams, and R. Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2018. [2](#), [3](#)

[Mao24] Xiao Mao. Fully-dynamic all-pairs shortest paths: Likely optimal worst-case update time. In *56th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, page to appear, 2024. [1](#)

[Min58] George J. Minty. A variant on the shortest-route problem. *Operations Research*, 6(6):882–883, 1958. [3](#)

[Moo59] Edward F. Moore. The shortest path through a maze. In *International Symposium on the Theory of Switching*, pages 285–292. Harvard University Press, 1959. [1](#)

[OV22] James B Orlin and László Végh. Directed shortest paths via approximate cost balancing. *Journal of the ACM*, 70(1):1–33, 2022. [1](#)

[Pet04] Seth Pettie. A new approach to all-pairs shortest paths on real-weighted graphs. *Theoretical Computer Science*, 312(1):47–74, 2004. [1](#)

[Pol60] Maurice Pollack. The maximum capacity through a network. *Oper. Res.*, 8, 1960. [3](#)

[Pol24] Adam Polak. Personal communication, 2024. [7](#)

[PR05] Seth Pettie and Vijaya Ramachandran. A shortest path algorithm for real-weighted undirected graphs. *SIAM J. Comput.*, 34(6):1398–1431, 2005. [1](#)

[Ram96] Rajeev Raman. Priority queues: Small, monotone and trans-dichotomous. In *4th Annual European Symposium on Algorithms, ESA*, 1996. [1](#)

[Ram97] Rajeev Raman. Recent results on the single-source shortest paths problem. *SIGACT News*, 28(2):81–87, 1997. [1](#)

[RZ11] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2), 2011. [2](#), [6](#), [7](#)

[Sch11] Markus Schulze. A New Monotonic, Clone-Independent, Reversal Symmetric, and Condorcet-Consistent Single-Winner Election Method. *Soc. Choice Welf.*, 36(2):267–303, 2011. [3](#)

[Shi53] Alfonso Shimbel. Structural parameters of communication networks. *Bull. Math. Biol.*, 15:501–507, 1953. [1](#)

[ST83] Daniel D. Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3), 1983. [33](#)

[SYZ07] Asaf Shapira, Raphael Yuster, and Uri Zwick. All-pairs bottleneck paths in vertex weighted graphs. In *18th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2007. [3](#)

[Tak92] Tadao Takaoka. A new upper bound on the complexity of the all pairs shortest path problem. *Inf. Process. Lett.*, 43(4), 1992. [1](#)

[Tak98] Tadao Takaoka. Subcubic cost algorithms for the all pairs shortest path problem. *Algorithmica*, 20(3), 1998. [1](#)

[Tho96] Mikkel Thorup. On RAM priority queues. In *7th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 1996. [1](#)

[Tho99] Mikkel Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM*, 46(3):362–394, 1999. [1](#)

[Tho00] Mikkel Thorup. Floats, integers, and single source shortest paths. *J. Algorithms*, 35(2):189–201, 2000. [1](#)

[Tho04] Mikkel Thorup. Integer priority queues with decrease key in constant time and the single source shortest paths problem. *J. Comput. Syst. Sci.*, 69(3), 2004. [1, 3](#)

[Tho05] Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *37th Annual ACM Symposium on Theory of Computing, STOC*, 2005. [1](#)

[ULH09] Ehsan Ullah, Kyongbum Lee, and Soha Hassoun. An algorithm for identifying dominant-edge metabolic pathways. In *International Conference on Computer-Aided Design, ICCAD*, 2009. [3](#)

[Vas10] Virginia Vassilevska Williams. Nondecreasing paths in a weighted graph or: How to optimally read a train schedule. *ACM Trans. Algorithms*, 6(4), 2010. [3](#)

[vdBLN⁺20] Jan van den Brand, Yin Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Bipartite matching in nearly-linear time on moderately dense graphs. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2020. [1](#)

[vdBNS19] Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. Dynamic matrix inverse: Improved algorithms and matching conditional lower bounds. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2019. [28](#)

[VW06] Virginia Vassilevska and Ryan Williams. Finding a maximum weight triangle in $n^{3-\delta}$ time, with applications. In *38th Annual ACM Symposium on Theory of Computing, STOC*, 2006. [3](#)

[VW13] Virginia Vassilevska Williams and Ryan Williams. Finding, minimizing, and counting weighted subgraphs. *SIAM J. Comput.*, 42(3), 2013. [3](#)

[VW18] Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5), 2018. [3](#), [4](#)

[VWY06] Virginia Vassilevska, Ryan Williams, and Raphael Yuster. Finding the smallest H -subgraph in real weighted graphs and related problems. In *33rd International Colloquium on Automata, Languages and Programming, ICALP*, 2006. [3](#)

[VWY07] Virginia Vassilevska Williams, Ryan Williams, and Raphael Yuster. All-pairs bottleneck paths for general graphs in truly sub-cubic time. In *39th Annual ACM Symposium on Theory of Computing, STOC*, 2007. [28](#)

[VWY09] Virginia Vassilevska Williams, Ryan Williams, and Raphael Yuster. All pairs bottleneck paths and max-min matrix products in truly subcubic time. *Theory Comput.*, 5(1):173–189, 2009. [3](#)

[VXXZ24] Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega. In *2024 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2024. [2](#), [4](#), [6](#)

[Wil18] R. Ryan Williams. Faster all-pairs shortest paths via circuit complexity. *SIAM J. Comput.*, 47(5), 2018. [1](#)

[Yus09] Raphael Yuster. Efficient algorithms on sets of permutations, dominance, and real-weighted APSP. In *20th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2009. [3](#), [13](#)

[Zwi04] Uri Zwick. A slightly improved sub-cubic algorithm for the all pairs shortest paths problem with real edge lengths. In *15th International Symposium on Algorithms and Computation, ISAAC*, 2004. [1](#)

A Omitted Proofs for Partially Dynamic Lower Bounds

In this appendix, we provide several omitted lower bounds for partially dynamic path problems.

A.1 Lower Bounds from the OMv3 Hypothesis

First, we show that bottleneck queries have the desired equivalence with OMv3.

Lemma 4.19. *Consider the query $(\vec{u}_i, \vec{v}_i, \vec{w}_i)$ and the query after updating the k -th coordinate of \vec{u} . The result of the query is $\mathbf{b}(s, t) \geq i \cdot n + (k + 1)$ if and only if*

$$\bigvee_{j, \ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k, j] \wedge A[j, \ell] \wedge A[\ell, k]) = \text{TRUE}.$$

Proof of Lemma 4.19. First, consider the converse direction. Let j, ℓ denote the indices such that the statement is true. We claim the following (s, t) -bottleneck path exists in the graph and has capacity $i \cdot n + (k + 1)$,

$$(s, a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)}, b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)}, c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)}, \hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}, t)$$

We verify that this path exists and has the required capacity. $(s, a_k^{(1)})$ has capacity ∞ . Since $\vec{u}_i[k] = \text{TRUE}$, each edge in the sub-path $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)})$ exists and has capacity $i \cdot n + (k + 1)$. Next, $A[k, j] = \text{TRUE}$ implies that the edge $(a_k^{(3)}, b_j^{(1)})$ exists with ∞ capacity. Since $\vec{v}_i[j] = \text{TRUE}$, each edge in the sub-path $(b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ exists and has capacity $(i + 1) \cdot n$. Then, $A[j, \ell] = \text{TRUE}$ implies that the edge $(b_j^{(3)}, c_\ell^{(1)})$ exists with ∞ capacity. Since $\vec{w}_i[\ell] = \text{TRUE}$, each edge in the sub-path $(c_\ell^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_\ell^{(3)})$ exists and has capacity $(i + 1) \cdot n$. Then, $A[\ell, k] = \text{TRUE}$ implies that the edge $(c_\ell^{(3)}, \hat{a}_k^{(1)})$ exists with ∞ capacity. Finally, as $\vec{u}_i[k] = \text{TRUE}$, we have the sub-path $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$ with capacity $i \cdot n + (k + 1)$. We conclude by observing that edge $(\hat{a}_k^{(3)}, t)$ has ∞ capacity.

Conversely, suppose $\mathbf{b}(s, t) \geq i \cdot n + (k + 1)$. Then, since edges are directed left to right, there is a path of the form,

$$(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, b_{i_4}^{(1)}, b_{i_5}^{(2)}, b_{i_6}^{(3)}, c_{i_7}^{(1)}, c_{i_8}^{(2)}, c_{i_9}^{(3)}, \hat{a}_{i_{10}}^{(1)}, \hat{a}_{i_{11}}^{(2)}, \hat{a}_{i_{12}}^{(3)}, t)$$

achieving this capacity. In both incremental and decremental cases, at the time of the query after updating k -th coordinate of \vec{u}_i during the OMv3 query $\vec{u}_i, \vec{v}_i, \vec{w}_i$, the maximum weight of any edge between $A_1 \times A_2$, $A_2 \times A_3$, $\hat{A}_1 \times \hat{A}_2$, and $\hat{A}_2 \times \hat{A}_3$ is $i \cdot n + (k + 1)$. Furthermore, the only edges attaining this weight are $(a_k^{(1)}, a_{k+i \bmod n}^{(2)})$, $(a_{k+i \bmod n}^{(2)}, a_k^{(3)})$, $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)})$, and $(\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$. If these edges exist, then $\vec{u}_i[k] = \text{TRUE}$ and $k = i_1 = i_3 = i_{10} = i_{12}$. Note also that $i_2 = i_{11} = k + i \bmod n$.

During this query, in both incremental and decremental cases, the maximum weight of any edge between $B_1 \times B_2$ and $B_2 \times B_3$ is $(i + 1) \cdot n$. Furthermore, the only edges attaining this weight are $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$, $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ for $b_j \in N(d_i)$. In particular, we have $i_4 = i_6 = j$ and $i_5 = j + i \bmod n$ for some $\vec{v}_i[j] = \text{TRUE}$. By a similar argument, $i_7 = i_9 = \ell$ and $i_8 = \ell + i \bmod n$ for some $\vec{w}_i[\ell] = \text{TRUE}$.

Finally, it remains to show $A[k, j] = A[j, \ell] = A[\ell, k] = \text{TRUE}$. Since $\text{wt}(a_k^{(3)}, b_j^{(1)}) = \infty$ we have $A[k, j] = \text{TRUE}$. A similar argument shows that $A[j, \ell] = A[\ell, k] = \text{TRUE}$. \square

Finally, we prove the lower bound for partially weight dynamic stEA under the OMv3 hypothesis.

Theorem 5.7. *Under the OMv3 hypothesis, any algorithm computing incremental/decremental stEA on directed graphs with polynomial preprocessing time requires $n^{\omega-1-o(1)}$ amortized update and query time. Furthermore, each edge weight is incremented only once.*

Proof of Theorem 5.7. We describe the reduction for incremental stEA (weight increasing), noting that the decremental reduction can be obtained by modifying the reduction as in Theorem 3.8.

Suppose for contradiction there is an algorithm \mathcal{A} with polynomial preprocessing time and total update and query time $O(n^{\omega+1-c})$ for some $c > 0$. Consider an OMv3 instance with matrix A and queries $(\vec{u}_i, \vec{v}_i, \vec{w}_i)_{i=1}^n$. Throughout the lower bound, assume that the matrices and vectors are indexed 1 to n . We design an algorithm with polynomial preprocessing time and total update and query time $O(n^{\omega+1-c})$, contradicting the OMv3 hypothesis.

We construct a graph with vertices,

$$\{s\} \cup (A_1 \cup A_2 \cup A_3) \cup (B_1 \cup B_2 \cup B_3) \cup (C_1 \cup C_2 \cup C_3) \cup (\hat{A}_1 \cup \hat{A}_2 \cup \hat{A}_3) \cup \{t\}$$

Each vertex set $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, \hat{A}_1, \hat{A}_2$ and \hat{A}_3 has n vertices. For each $a \in A$, we create a copy $a^{(i)} \in A_i$ and a copy $\hat{a}^{(i)} \in \hat{A}_i$ for $i \in \{1, 2, 3\}$. Similarly, for each $b \in B$ and $c \in C$, we create a copy $b^{(i)} \in B_i$ and $c^{(i)} \in C_i$ for $i \in \{1, 2, 3\}$. Initially, insert edges $(s, a^{(1)})$ with weight 1 and edges $(\hat{a}^{(3)}, t)$ with weight 9 for all $a \in A$. Furthermore, insert edges with 3 weight between $(a^{(3)}, b^{(1)})$ (respectively weight 5 between $(b^{(3)}, c^{(1)})$ and weight 7 between $(c^{(3)}, \hat{a}^{(1)})$) if and only if $A[a, b] = \text{TRUE}$ (respectively if $A[b, c] = \text{TRUE}$ and $A[c, a] = \text{TRUE}$). We also insert all edges between $A_1 \times A_2, A_2 \times A_3, B_1 \times B_2, B_2 \times B_3, C_1 \times C_2, C_2 \times C_3, \hat{A}_1 \times \hat{A}_2, \hat{A}_2 \times \hat{A}_3$ with weight 0.

This creates a graph with $O(n)$ vertices, so that by running \mathcal{A} we can preprocess the graph in polynomial time. We now proceed to the dynamic phase of the reduction.

Consider an OMv3 query $\vec{u}_i, \vec{v}_i, \vec{w}_i$.

1. For each $\vec{v}_i[j] = \text{TRUE}$, increase the weights of $(b_j^{(1)}, b_{j+i \bmod n}^{(2)})$ and $(b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ to 4.
2. For each $\vec{w}_i[j] = \text{TRUE}$, increase the weights of $(c_j^{(1)}, c_{j+i \bmod n}^{(2)})$ and $(c_{j+i \bmod n}^{(2)}, c_j^{(3)})$ to 6.
3. In the inner loop, we iterate over coordinates k of \vec{u} in increasing order.

- (a) If $\vec{u}_i[k] = \text{TRUE}$, increase the weights of

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)})$$

to 2 and increase the weights of,

$$(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$$

to 8.

- (b) Compute $\mathbf{a}(s, t)$ by querying \mathcal{A} . If $\mathbf{a}(s, t) = 9$, return TRUE .
- (c) After the query, increase the weights of

$$(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}), (\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$$

to 10.

4. After all queries of this round, increase all weights of

$$(b_j^{(1)}, b_{j+i \bmod n}^{(2)}), (b_{j+i \bmod n}^{(2)}, b_j^{(3)}), (c_j^{(1)}, c_{j+i \bmod n}^{(2)}), (c_{j+i \bmod n}^{(2)}, c_j^{(3)})$$

for all $\vec{u}_i[k] = \vec{w}_i[k] = \text{TRUE}$ to 10. Then, return FALSE .

The bottleneck of this algorithm is to run \mathcal{A} , so the total query time is $O(n^2 \cdot n^{\omega-1-c}) = O(n^{\omega+1-c})$.

Lemma A.1. *Consider the query when iterating over d_i and a_k . The result of the query is $\mathbf{a}(s, t) = 9$ if and only if,*

$$\bigvee_{j, \ell} (\vec{u}_i[k] \wedge \vec{v}_i[j] \wedge \vec{w}_i[\ell] \wedge A[k, j] \wedge A[j, \ell] \wedge A[\ell, k]) = \text{TRUE}$$

Proof. Note that all edges into t have weight 9. Thus, the arrival time $\mathbf{a}(s, t) \in \{9, \infty\}$, and $\mathbf{a}(s, t) = 9$ if and only if there is a valid arrival path from s to t .

Suppose there is clause indexed by j, ℓ that is TRUE. Then, the following arrival path exists in the graph,

$$(s, a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)}, b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)}, c_{\ell+i \bmod n}^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_{\ell}^{(3)}, \hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)}, t)$$

We verify that this path exists and is valid. $(s, a_k^{(1)})$ has weight 1. Since $\vec{u}_i[k] = \text{TRUE}$, each edge in the sub-path $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}, a_k^{(3)})$ has weight 2. Next, $A[k, j] = \text{TRUE}$ implies that the edge $(a_k^{(3)}, b_j^{(1)})$ exists with weight 3. Since $\vec{v}_i[j] = \text{TRUE}$, each edge in the sub-path $(b_j^{(1)}, b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ has weight 4. Then, $A[j, \ell] = \text{TRUE}$ implies that the edge $(b_j^{(3)}, c_{\ell}^{(1)})$ exists with weight 5. Since $\vec{w}_i[\ell] = \text{TRUE}$, each edge in the sub-path $(c_{\ell}^{(1)}, c_{\ell+i \bmod n}^{(2)}, c_{\ell}^{(3)})$ has weight 6. Then, $A[\ell, k] = \text{TRUE}$ implies that the edge $(c_{\ell}^{(3)}, \hat{a}_k^{(1)})$ exists with weight 7. Finally, as $\vec{u}_i[k] = \text{TRUE}$, we have the sub-path $(\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$ with weight 8. We conclude by observing that edge $(\hat{a}_k^{(3)}, t)$ has weight 9.

Conversely, suppose $\mathbf{a}(s, t) = 9$. Then, since edges are directed left to right, there is a path of the form,

$$(s, a_{i_1}^{(1)}, a_{i_2}^{(2)}, a_{i_3}^{(3)}, b_{i_4}^{(1)}, b_{i_5}^{(2)}, b_{i_6}^{(3)}, c_{i_7}^{(1)}, c_{i_8}^{(2)}, c_{i_9}^{(3)}, \hat{a}_{i_{10}}^{(1)}, \hat{a}_{i_{11}}^{(2)}, \hat{a}_{i_{12}}^{(3)}, t)$$

and is a valid arrival path.

In both incremental and decremental cases, at the time of the query after updating the k -th coordinate of the i -th OMv3 query $\vec{u}_i, \vec{v}_i, \vec{w}_i$, the only edges in $A_1 \times A_2, A_2 \times A_3, \hat{A}_1 \times \hat{A}_2, \hat{A}_2 \times \hat{A}_3$ with weight 2 are $(a_k^{(1)}, a_{k+i \bmod n}^{(2)}), (a_{k+i \bmod n}^{(2)}, a_k^{(3)}), (\hat{a}_k^{(1)}, \hat{a}_{k+i \bmod n}^{(2)}),$ and $(\hat{a}_{k+i \bmod n}^{(2)}, \hat{a}_k^{(3)})$. Note that all other weights have weight either 0 or 10, and thus cannot participate in any valid arrival path (since edges into A_1 have weight 1 and edges out of A_3 have weight 3). If these edges exist with weight 2, then $\vec{u}_i[k] = \text{TRUE}$ and $k = i_1 = i_3 = i_{10} = i_{12}$. The only valid edges between $B_1 \times B_2$ and $B_2 \times B_3$ with weight 4 are $(b_j^{(1)}, b_{j+i \bmod n}^{(2)}), (b_{j+i \bmod n}^{(2)}, b_j^{(3)})$ for $\vec{v}_i[j] = \text{TRUE}$ (and no other edges can participate in an arrival path since edges into B_1 have weight 3 and edges out of B_3 have weight 5). Since these paths are vertex disjoint, we have $i_4 = i_6 = j$ for some $\vec{v}_i[j] = \text{TRUE}$. By a similar argument, $i_7 = i_9 = \ell$ for some $\vec{w}_i[\ell] = \text{TRUE}$.

Finally, it remains to show $A[k, j] = A[j, \ell] = A[\ell, k] = \text{TRUE}$. Since $(a_k^{(3)}, b_j^{(1)})$ form an edge, we have $A[k, j] = \text{TRUE}$. A similar argument shows that $A[j, \ell] = A[\ell, k] = \text{TRUE}$. \square

From Lemma A.1, the above procedure solves the OMv3 instance, contradicting the OMv3 hypothesis. \square

A.2 Lower Bounds for Sparse Graphs

We have given strong conditional lower bounds for a variety of partially dynamic problems on dense graphs, where $m = \Theta(n^2)$. The following lemma shows that our lower bounds similarly apply to graphs with all sparsities. In fact, we obtain a simple reduction for any graph problem where the solution does not change when adding isolated vertices.

Proposition A.2 (Graph Sparsification). *For each of the following problems:*

1. stSP, SSSP *on undirected, weighted graphs*
2. nw-stSP, nw-SSSP *on undirected, node-weighted graphs*
3. stBP, SSBP *on directed, weighted graphs*

Suppose any incremental/decremental algorithm for the given problem requires $n^{t-o(1)}$ total update time on n -vertex graphs.

Then, for any $m \leq \binom{n}{2}$, any incremental/decremental algorithm for \mathcal{P} requires $m^{t/2-o(1)}$ total update time on n -vertex graphs with m edges.

Proof. Let $m \leq \binom{n}{2}$ and set $n_0 = \lfloor \sqrt{m} \rfloor \leq n$. By assumption, any incremental/decremental algorithm requires $n_0^{t-o(1)}$ total time on n_0 -vertex graphs. Given any instance of a partially dynamic n_0 -vertex graph we can construct a partially dynamic n -vertex graph by adding $n - n_0$ isolated vertices, which we will never add edges to. Suppose for contradiction there is an algorithm with $O(m^{t/2-c})$ total time on the augmented graph for some $c > 0$. Then, we obtain an algorithm for the n_0 -vertex graph simply by outputting the queries on the n_0 -vertex sub-graph of the augmented n -vertex graph, obtaining an $O(m^{t/2-c}) = O(n_0^{t-2c})$ total time algorithm, a contradiction. \square