

New Techniques and Fine-Grained Hardness for Dynamic Near-Additive Spanners

Thiago Bergamaschi
MIT

Monika Henzinger
University of Vienna

Maximilian Probst Gutenberg
ETH Zurich

Virginia Vassilevska Williams
MIT

Nicole Wein
MIT

Abstract

Maintaining and updating shortest paths information in a graph is a fundamental problem with many applications. As computations on dense graphs can be prohibitively expensive, and it is preferable to perform the computations on a sparse skeleton of the given graph that roughly preserves the shortest paths information. Spanners and emulators serve this purpose. Unfortunately, very little is known about dynamically maintaining sparse spanners and emulators as the graph is modified by a sequence of edge insertions and deletions. This paper develops fast dynamic algorithms for spanner and emulator maintenance and provides evidence from fine-grained complexity that these algorithms are tight. For unweighted undirected m -edge n -node graphs we obtain the following results.

Under the popular OMv conjecture, there can be no decremental or incremental algorithm that maintains an $n^{1+o(1)}$ edge (purely additive) $+n^\delta$ -emulator for any $\delta < 1/2$ with arbitrary polynomial preprocessing time and total update time $m^{1+o(1)}$. Also, under the Combinatorial k -Clique hypothesis, any fully dynamic combinatorial algorithm that maintains an $n^{1+o(1)}$ edge $(1 + \epsilon, n^{o(1)})$ -spanner or emulator for small ϵ must either have preprocessing time $mn^{1-o(1)}$ or amortized update time $m^{1-o(1)}$. Both of our conditional lower bounds are tight.

As the above fully dynamic lower bound only applies to combinatorial algorithms, we also develop an algebraic spanner algorithm that improves over the $m^{1-o(1)}$ update time for dense graphs. For any constant $\epsilon \in (0, 1]$, there is a fully dynamic algorithm with worst-case update time $O(n^{1.529})$ that whp maintains an $n^{1+o(1)}$ edge $(1 + \epsilon, n^{o(1)})$ -spanner.

Our new algebraic techniques allow us to also obtain a new fully dynamic algorithm for All-Pairs Shortest Paths (APSP) that can perform both edge updates and can report shortest paths in worst-case time $O(n^{1.9})$, which are correct whp. This is the first *path-reporting* fully dynamic APSP algorithm with a truly subquadratic query time that beats $O(n^{2.5})$ update time. It works against an oblivious adversary.

Finally, we give two applications of our new dynamic spanner algorithms: (1) a fully dynamic $(1 + \epsilon)$ -approximate APSP algorithm with update time $O(n^{1.529})$ that can report approximate shortest paths in $n^{1+o(1)}$ time per query; previous subquadratic update/query algorithms could only report the distance, but not obtain the paths; (2) a fully dynamic algorithm for near-2-approximate Steiner tree maintenance with both terminal and edge updates.

1 Introduction

Computing shortest paths in a graph is a fundamental problem with many applications. However, as on dense graphs the running time can be prohibitively expensive, it is preferable to perform the computation on a sparser representation of the given graph that approximately preserves the shortest paths distances. Such representations are called *spanners* and *emulators*. Given an undirected, unweighted graph $G = (V, E)$, a subgraph H of G is defined to be an (α, β) -*spanner* if for every pair of vertices $x, y \in V$, we have that

$$\text{dist}_G(x, y) \leq \text{dist}_H(x, y) \leq \alpha \cdot \text{dist}_G(x, y) + \beta.$$

A graph $H = (V, E')$ is defined to be an (α, β) -*emulator* if it fulfills the above constraint, is possibly edge-weighted, but not necessarily a subgraph of G . Thus every spanner is also an emulator but not vice versa. When evaluating the quality of a spanner or emulator H three parameters are of interest: the *multiplicative approximation* α , the *additive approximation* β and the *sparsity* of H that is the number of edges in H .

Spanners and emulators have a variety of applications, ranging from efficient routing to parallel and distributed algorithms to efficient distance oracles, i.e. data structures that answer shortest-path queries. Thus, there exists a large body of work on computing spanners (see below). As real-world graphs are often dynamic, it raises the question whether spanners and emulators can be maintained efficiently when the graph is modified by edge updates. Unfortunately, very little is known about this question. In this article, we are concerned with the design of efficient algorithms to *dynamically* maintain an $(1 + \epsilon, \beta)$ -spanner on an undirected unweighted graph that is undergoing edge insertions and deletions.

If the update sequence is restricted to consist exclusively of insertions, we say that the graph is *incremental* and if it only consists of edge deletions we say that it is *decremental*. If the graph is either incremental or decremental, we also say it is *partially dynamic* and otherwise we say it is *fully dynamic*.

Apart from giving new decremental and fully dynamic deterministic and randomized algorithms that maintain spanners and emulators we also provide evidence from fine-grained complexity that these algorithms are tight. We then use our new algorithms and techniques to give novel fully dynamic approximate and exact all-pairs shortest paths (APSP) algorithms *that can report the corresponding shortest path*, addressing an open question raised by [82]. We further provide applications for other problems such as the maintenance of an approximate Steiner tree of a graph.

Prior Work. In this section, we discuss work that is directly related to the results in our paper. We use \tilde{O} -notation to suppress logarithmic factors, let n and m be the maximum number of vertices and edges respectively in any version of the graph under consideration. Unless otherwise specified, all graphs are undirected and unweighted. To ease the discussion, we assume for the rest of the section that ϵ is a constant.

Spanners and Emulators. Spanners and static algorithms to construct them have been studied in great detail for *multiplicative* approximation [12, 70, 10, 36, 13, 18, 72] culminating in near-optimal algorithms to construct $(2k - 1, 0)$ -spanners of sparsity $\tilde{O}(n^{1+1/k})$. There has also been an extensive line of research on *purely additive* spanners [9, 8, 38, 28, 16, 30] where $(1, 2)$, $(1, 4)$ and $(1, 6)$ -spanners are known of sparsity $\tilde{O}(n^{3/2})$, $\tilde{O}(n^{7/5})$ and $\tilde{O}(n^{4/3})$. While algorithms for fast constructions have been studied (e.g. [86, 63, 64]), no near-optimal algorithm for the construction of any of the above additive spanners is known. For example, the fastest algorithm for constructing a $O(1)$ -additive spanner with $O(n^{4/3})$ edges is $O(n^2)$. Further, Abboud and Bodwin [2] proved that any purely additive spanner of sparsity $\tilde{O}(n^{4/3-\epsilon})$, for any constant $\epsilon > 0$, has at least polynomial in n additive error. Constructions by Bodwin and Vassilevska Williams [27, 26] are known giving sparsity $\tilde{O}(n)$ and additive error $\tilde{O}(n^{3/7+\epsilon})$. Following [2], Huang and Pettie [58] constructed a family of graphs such that any $\tilde{O}(n)$ -sized spanner for an n -node graph in the family must have $\Omega(n^{1/13})$ additive stretch.

Mixed-error (α, β) -spanners were studied in [44, 16, 81, 71, 40, 19]. Most of these results focus on the setting of near-additive spanners, that is (α, β) -spanners where $\alpha = 1 + \epsilon$ for some arbitrarily small constant $\epsilon > 0$. The goal of this setting is to obtain extremely sparse spanners $\tilde{O}(n)$. The best results obtain $(1 + \epsilon, n^{o(1)})$ -spanners with $\tilde{O}(n)$ edges. Abboud et al. [3] further developed a fine-grained hierarchy to give lower bounds for trade-offs between ϵ , additive error and the sparsity of emulators. These lower bounds also apply to the setting of (α, β) -emulators.

Finally, we point out that a related notion to emulators are hopsets: given a graph G , we say that a graph H is a (α, β, h) -hopset if for every two vertices $x, y \in V$, there is a path $\pi_{x,y}$ from x to y in the graph $G \cup H$ consisting of at most h edges, such that $\mathbf{dist}_{G \cup H}(x, y) \leq w(\pi_{x,y}) \leq (1 + \epsilon)\mathbf{dist}_G(x, y) + \beta$ where $w(\pi_{x,y})$ denotes the weight of the path $\pi_{x,y}$. There is a lot of recent work on hopsets, especially $(1 + \epsilon, \beta, h)$ -hopsets, for which there are efficient algorithms [42, 59, 41] with $\beta = 0$ and $h = n^{o(1)}$. The hopset literature builds heavily on previous clustering techniques from near-additive spanners. An excellent survey that highlights this connection was recently given by Elkin and Neiman [43].

Spanners and Emulators in Dynamic Graphs. Spanners have also been extensively studied in the dynamic graph setting, where near-optimal algorithms for multiplicative spanners in fully dynamic graphs exist [11, 39, 17, 14, 25, 23, 48, 20]. For hopsets, the dynamic graph literature has been mainly concerned with maintaining $(1 + \epsilon, n^{o(1)}, n^{o(1)})$ -hopsets in partially dynamic graphs [55, 21, 50] where they were used to derive fast algorithms for the partially dynamic Single-Source Shortest Paths problem. As was observed in [52] (Lemma 4.2) combining [74] with [81] leads to a $(1 + \epsilon, 2(1 + 2/\epsilon)^{k-2})$ -approximate decremental emulator with total time $O((1 + 2/\epsilon)^{k-2}mn^{1/k})$. To our knowledge, additive and near-additive spanners have not been studied in the dynamic graph literature. Also, there are no known conditional lower bounds for dynamic algorithms for maintaining a spanner.

Fully Dynamic Shortest Paths with Worst-Case Update Time. Closely related to maintaining a spanner/emulator is the problem of maintaining shortest paths. There are three problems of focal interest:

- (1) The s - t Shortest Path (st -SP) problem asks for the shortest path between two fixed vertices $s, t \in V$.
- (2) The Single-Source Shortest Paths (SSSP) problem asks for the shortest path tree from a fixed vertex.
- (3) The All-Pairs Shortest Paths (APSP) problem asks for the shortest path between every vertex pair.

For each of these three problems, there is the *distance reporting* and the *path reporting* variant, where the former requires to only return the length of the shortest path, while the latter needs to return the actual shortest path. There is an enormous line of research on these three problems in various settings. Since our fully dynamic algorithms have worst-case guarantees on update time, we focus this discussion on prior work on fully dynamic algorithms with worst-case update time.

For the st -SP problem and the SSSP problem the lower bounds in [4, 56] suggest that the essentially best solution to these problems is to rerun Dijkstra’s algorithm after every update (even when the updates are not required to be worst-case). However, these conditional lower bounds are based on the BMM conjecture and therefore hold only for “combinatorial” algorithms. Indeed, Sankowski [76] has shown that a worst-case update time of $O(n^{1.932})$ and query time of $O(n^{1.288})$ to obtain the *distance* between a pair of vertices is possible and therefore has given the first subquadratic algorithm for the distance-reporting version of the st -SP problem. Recently, this result was further improved to worst-case update time $\tilde{O}(n^{1.863})$ and query time $\tilde{O}(n^{0.45})$ in [82] where distance reporting queries are only required to return a $(1 + \epsilon)$ -approximate distance estimate. Rebalancing their trade-off terms, the authors also obtain an algorithm that maintains $(1 + \epsilon)$ -approximate SSSP with worst-case update time $\tilde{O}(n^{1.823})$ and $(1 + \epsilon)$ -approximate APSP in worst-case update time $\tilde{O}(n^2)$.

A *major drawback* of both approaches is that they cannot answer path reporting queries. The algorithm with fastest worst-case update time that can return *actual* (approximate) shortest-paths for st -SP and SSSP

remains to rerun Dijkstra’s algorithm and for the APSP problem to use a combinatorial data structure where the currently best worst-case update is $\tilde{O}(n^{2+2/3})$ for weighted graphs and $\tilde{O}(n^{2.5})$ time for unweighted graphs (see [79, 6, 51]).

For approximate distance oracles with amortized running time, Abraham et al. [7] achieved an $2^{O(k)}$ -approximation with $O(\sqrt{mn}n^{1/k})$ amortized update time for constant k , and Forster et al. [49] gave an $(O(\log n))^{3k-2}$ -approximation with $O(k \log^2 n)$ query time and $m^{1/k+o(1)}(O(\log n))^{4i-3}$ update time for any $k \geq 2$, being the first to break the \sqrt{m} update time barrier.

Partially Dynamic Shortest Paths. The classic ES-tree data structure [47] initiated the field with a deterministic total time $O(mn)$ algorithm for partially dynamic exact SSSP. In the setting where a $(1 + \epsilon)$ -multiplicative approximation is allowed, Bernstein and Roditty [24] gave the first improvement over the ES-tree for an approximation algorithm with an algorithm for decremental $(1 + \epsilon)$ -approximate SSSP with total time $n^2 2^{O(\sqrt{\log n})}$. Subsequently, Henzinger et al. gave an algorithm [55] with total update time $m^{1+o(1)}$. These algorithms are all randomized and against an oblivious adversary.

Bernstein and Chechik gave the first deterministic partially dynamic $(1 + \epsilon)$ -approximate algorithm that improves upon the ES-tree data structure; it runs in total time $\tilde{O}(n^2)$ [21] and does not report paths, only distances. Chuzhoy and Khanna [34] gave an algorithm with total time $n^{2+o(1)}$ that works against an adaptive adversary and returns paths with $n^{1+o(1)}$ query time. Chuzhoy and Saranurak [35] recently further improved the running time of the path query to $|P|n^{o(1)}$ for an approximate shortest path P . Further, Bernstein and Chechik recently gave an algorithm with total update time $\tilde{O}(mn^{3/4})$ [22], which was then improved to $O(mn^{0.5+o(1)})$ by Probst Gutenberg and Wulff-Nilsen [50]. Neither of these data structures can answer path queries which was recently addressed in [20].

For decremental APSP, Henzinger et al. [55] presented an approximation algorithm with stretch $((2 + \epsilon)^k - 1)$ and total update time $m^{1+1/k+o(1)}$ for any positive integer k . They also gave an algorithm with stretch $(2 + \epsilon)$ or $(1 + \epsilon, 2)$ with total update time $\tilde{O}(n^{2.5})$ in [53] which was recently derandomized by Chuzhoy and Saranurak [35]. Finally Henzinger et al. [53] presented a $(1 + \epsilon)$ -approximate deterministic algorithm with $\tilde{O}(mn/\epsilon)$ update time which derandomized the construction by Roditty and Zwick [74] with matching running time. Later on, Chechik [31] presented a $(2 + \epsilon)k - 1$ -approximate algorithm with update time $mn^{1/k+o(1)}$ for any positive integer k and constant ϵ , whose total update time matches the preprocessing time of static distance oracles [80] with corresponding stretch. Recently, Chen et al. [32] gave an incremental $(2k - 1)$ -approximate algorithm with $O(m^{1/2}n^{1/k})$ worst-case time per operation. We point out that there is an extensive line of work on the decremental APSP problem [62, 15, 37, 73, 79, 24, 74, 5, 55, 53, 54, 6, 31, 46] that is beyond the scope of this overview.

From the lower bounds side, Roditty and Zwick [73] showed that any incremental or decremental algorithm for SSSP in weighted graphs with preprocessing time $p(n)$, query time $q(n)$ and update time $u(n)$ must satisfy $p(n) + n \cdot u(n) + n^2 \cdot q(n) \geq n^{3-o(1)}$ unless APSP has a truly subcubic time algorithm. Similarly, for unweighted graphs, they showed that any *combinatorial* incremental or decremental algorithm must satisfy that equation unless Boolean matrix multiplication (BMM) has a truly subcubic time combinatorial algorithm. Abboud and Vassilevska Williams [4] extended these lower bounds to also hold for *st*-SP, where now the algorithms must satisfy $p(n) + n \cdot (u(n) + q(n)) \geq n^{3-o(1)}$, for weighted graphs under the APSP conjecture, and for unweighted graphs under the combinatorial BMM conjecture.

Hypotheses for Fine-Grained Complexity Our conditional lower bounds rely on two popular hypotheses: the Online Boolean Matrix-Vector Multiplication (OMv) conjecture and the Combinatorial k -Clique hypothesis. In the OMv problem we are given an $n \times n$ matrix M that can be preprocessed. Then, an online sequence of vectors v^1, \dots, v^n is presented and the goal is to compute each Mv^i before seeing the next

vector v^{i+1} . The OMv conjecture was first defined in [56], and has been used many times since.

Conjecture 1.1 (OMv). *For any constant $\epsilon > 0$, there is no $O(n^{3-\epsilon})$ -time algorithm that solves OMv with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The Combinatorial k -Clique hypothesis is defined as follows and has been used a number of times (e.g. [68, 1, 29]).

Hypothesis 1.2 (Combinatorial k -Clique). *For any constant $\epsilon > 0$, for an n -node graph there is no $O(n^{k-\epsilon})$ time combinatorial algorithm for k -clique detection with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

For the special case of k -Clique detection where $k = 3$ (i.e. triangle detection), we also consider non-combinatorial algorithms. Triangle detection can easily be solved using matrix multiplication but it is a big open question whether triangle detection admits a $O(m^{\omega-\epsilon})$ time algorithm, where $\omega < 2.373$ is the matrix multiplication exponent (e.g. [84], [85] Open Problem 4.3(c), and [78] Open Problem 8.1). It is generally believed that such an algorithm does not exist, and our reductions from k -Clique also imply hardness under this hypothesis.

Our results. We present novel algorithms and conditional lower bounds for (α, β) -spanners and emulators as well as faster fully dynamic APSP algorithms. We prove the following for undirected unweighted graphs.

1. **Conditional lower bounds for partially dynamic spanners/emulators.** Under the OMv conjecture, there can be no decremental or incremental algorithm that maintains a $(1, n^{o(1)})$ -emulator (and thus spanner) with $O(m^{1-\epsilon})$ edges for any constant $\epsilon > 0$ with arbitrary polynomial preprocessing time and total update time $O(mn^{1-\epsilon})$. The same result also holds for *all* sparsities m for combinatorial algorithms under the Combinatorial k -Clique hypothesis.

For completeness, we also present algorithms that are tight with our conditional lower bounds. Note that mixed additive/multiplicative error is necessary for these algorithms since there can be no $(O(1), 0)$ -spanner or emulator with $n^{1+o(1)}$ edges (e.g. [69]). Our algorithms rely heavily on prior work. Using techniques similar to [31], for any constant $\epsilon \in (0, 1]$, we maintain whp against an oblivious adversary a partially dynamic $n^{1+o(1)}$ edge $(1 + \epsilon, n^{o(1)})$ -spanner in total update time $m^{1+o(1)}$ time. Using techniques from [50], we also give a deterministic partially dynamic algorithm that maintains a $(1 + \epsilon, n^{\alpha+o(1)})$ -emulator of a graph in total time $O(mn^{1-\alpha+o(1)})$ for any $\alpha > 0$. Using a result in [20], we can further turn the above algorithm into a randomized algorithm that maintains a $(1 + \epsilon, n^{\alpha+o(1)})$ -spanner in expected total time $O(mn^{1-\alpha+o(1)})$ for any $\alpha > 0$ and that works against an adaptive adversary.

2. **Conditional lower bounds for combinatorial fully dynamic spanners.** Under the Combinatorial k -Clique hypothesis, for a graph of any sparsity m , for any constant $\epsilon > 0$, there can be no fully dynamic combinatorial algorithm that maintains an $O(m^{1-\epsilon})$ -edge $(1 + \alpha, n^{o(1)})$ -emulator for small α with preprocessing time $mn^{1-\epsilon}$ and amortized update time $m^{1-\epsilon}$. This conditional lower bound also extends to incremental and decremental algorithms but only for worst-case update times.

For completeness, we also present an algorithm that is tight with our conditional lower bound. This algorithm follows from rerunning a known static algorithm after every update. For any constant $\epsilon \in (0, 1]$, we give a deterministic fully dynamic algorithm with preprocessing time $m^{1+o(1)}$ and worst-case update time $m^{1+o(1)}$ time that maintains an $n^{1+o(1)}$ edge $(1 + \epsilon, n^{o(1)})$ -spanner.

3. **Algebraic fully dynamic spanner algorithms.** The above fully dynamic lower bound only applies to combinatorial algorithms, and we show that this is inherent; we develop an algebraic spanner algorithm

that beats our combinatorial lower bound. For any constant $\epsilon \in (0, 1]$, there is a fully dynamic algorithm with preprocessing time $\tilde{O}(n^2)$ in an initially empty graph and $O(n^{2.373})$ in an initially non-empty graph and worst-case update time $O(n^{1.529})$ that whp maintains an $n^{1+o(1)}$ -edge $(1 + \epsilon, n^{o(1)})$ -spanner and works against an oblivious adversary.¹

The construction from our above lower bound from the Combinatorial k -Clique hypothesis with $k = 3$ (i.e. triangle detection) also gives a conditional lower bound for non-combinatorial algorithms. Unless there is a breakthrough in non-combinatorial algorithms for triangle detection algorithms, there can be no fully dynamic algorithm that maintains an $O(n^{\omega-1-\epsilon})$ -edge $(1 + \alpha, n^{o(1)})$ -emulator for constant $\alpha < 2/3$ with preprocessing time $O(n^{\omega-\epsilon})$ and amortized update time $O(n^{\omega-1-\epsilon})$, where $\omega < 2.373$ is the matrix multiplication exponent. Thus $O(n^{1.372})$ update time and $O(n^{2.372})$ preprocessing time is not possible with current techniques.

We also give a conditional lower bound from the OMv conjecture that precludes algorithms for emulators with more edges and higher preprocessing time than the above lower bound from triangle detection, but at the cost of a lower update time. Under the OMv conjecture, for any constant $\epsilon > 0$, there can be no fully dynamic algorithm that maintains an $O(m^{1-\epsilon})$ -edge $(1 + \alpha, n^{o(1)})$ -emulator for constant $\alpha < 2/3$ with arbitrary polynomial preprocessing time and amortized update time $O(n^{1-\epsilon})$.

Both of these conditional lower bound also extend to incremental and decremental algorithms but only for worst-case update times.

4. Fully dynamic exact path-reporting APSP. To achieve the above results we develop the first fully dynamic APSP data structure that supports distance queries, path reporting queries, and edge updates in *sub-quadratic* time per operation. It uses algebraic techniques, is randomized, and works against an oblivious adversary. Specifically we show the following result, where $\omega(a, b, c)$ is the exponent for multiplying an $n^a \times n^b$ matrix by an $n^b \times n^c$, and κ_* is the solution to $\omega(1, 1, \kappa) = 1 + 2\kappa$. With the current bounds for rectangular matrix multiplication, $\kappa_* \approx 0.529$.

Theorem 1.3. *Let κ be such that $0 < \kappa \leq \kappa_*$, and let D be a distance parameter between 1 and n . There is a randomized fully dynamic data structure that can maintain an unweighted directed graph $G = (V, E)$ supporting the following operations with preprocessing time $\tilde{O}(n^2)$ in an empty initial graph and $\tilde{O}(Dn^\omega)$ in a non-empty initial graph: (a) **edge updates** in worst-case time $\tilde{O}(Dn^{\omega(1,1,\kappa)-\kappa})$ time; (b) **[distance reporting]:** on query $i, j \in V$, return $\mathbf{dist}(i, j)$ if $\mathbf{dist}(i, j) \leq D$, or answer that $\mathbf{dist}(i, j) > D$ otherwise, in worst-case $\tilde{O}(Dn^\kappa)$ time, where the answer is correct whp; (c) **[path reporting]:** on query $i, j \in V$, if $\mathbf{dist}(i, j) \leq D$, return a shortest path from i to j , in $\tilde{O}(D^2n^\kappa)$ time, where the answer is correct whp.*

We believe that this result is of independent interest.

Based on it we build a fully dynamic exact APSP data structure that with preprocessing time $\tilde{O}(n^2)$ on an initially empty graph achieves worst-case time $O(n^{1.9})$ per edge update, $O(n^{1.529})$ per distance query and $O(n^{1.9})$ per path reporting query. This is a significant improvement over the $O(n^{2.5})$ worst-case update time of [6, 51] and closer to the $O(n^2)$ time bound which is achieved by the data structure of [37] which can only support distance reporting queries, but no path reporting queries.

The algorithms in Theorem 1.3 are all Monte Carlo– they are correct with high probability and always run in the desired running time. If they could be made into Las Vegas algorithms (ones that are always correct but have expected running time), our applications of Theorem 1.3, such as our algebraic spanners, would also be Las Vegas, which is a more desirable guarantee. However, there are significant hurdles to overcome in order to make Theorem 1.3 Las Vegas. Like Sankowski’s original data structure [75], Theorem 1.3 heavily

¹Both the bound on the time per update as well as the correctness hold with high probability. If we rebuild the data structure from scratch every polynomially many updates, we can instead achieve an *expected amortized* time bound of $O(n^{1.529})$ per update.

relies on the use of polynomial identity testing (PIT), namely on the fact that PIT is in co-RP and hence has a fast Monte Carlo algorithm. To obtain a Las Vegas algorithm using a similar approach, one would need a ZPP algorithm for PIT. However, obtaining such an algorithm seems extremely difficult, and in fact Impagliazzo and Kabanets [60, 61] showed that such an algorithm would imply strong circuit lower bounds. Thus although the rest of our techniques can be made Las Vegas, making Theorem 1.3 Las Vegas as well is far from possible with current techniques.

5. Applications. We present two applications of our above results: fully dynamic approximate path-reporting APSP, and fully dynamic Steiner tree. Using the above theorem and the above algebraic spanner, we give the first subquadratic fully dynamic $(1 + \epsilon)$ -approximate APSP algorithm. It needs $\tilde{O}(n^2)$ preprocessing time on an empty graph and achieves worst-case time $n^{1+\kappa^*+o(1)} = O(n^{1.529})$ for updates, $n^{1+o(1)}$ for approximate distance reporting and approximate shortest path reporting, whp against an oblivious adversary. Note that all previous subquadratic update/query algorithms could only report distances, not paths.

Our second application of our above results is a fully dynamic algorithm for $(2 + \epsilon)$ -approximate Steiner tree, which can be used, for example, for routing in dynamic networks. Specifically we give the first subquadratic algorithm that maintains a $(2 + \epsilon)$ -approximate Steiner tree for a set S of terminals with both terminal and edge updates. Specifically, it has preprocessing time $\tilde{O}(n^2)$ on an empty initial graph and $n^{\omega+o(1)}$ on a non-empty initial graph, and worst-case time $n^{1+\kappa^*+o(1)} + |S|^2 \cdot n^{1+o(1)}$ per edge update, $|S|n^{1+o(1)}$ per node addition to S , and $|S|n^{o(1)}$ per node removal from S , giving subquadratic update time when $|S| \leq n^{1/2-o(1)}$ whp against an oblivious adversary. By increasing the processing time to $O(n^{2.621})$ using the data-structure of [82], the time for edge updates can be made $O(n^{1.843} + |S|^2 \cdot n^{0.45} + |S| \cdot n^{1+o(1)})$, allowing for more leverage over the size of the terminal set S . The only prior work in general graphs maintains a $(6 + \epsilon)$ -approximate Steiner Tree under changes to S only (no edge updates) and has preprocessing time $\tilde{O}(m\sqrt{n})$ and update time $\tilde{O}(\sqrt{n})$ [65].

Organization In Section 2 we give a technical overview of a selection of our results. Section 3 is the preliminaries. In Section 4, we present our conditional lower bounds. In Section 5, we present our data structure for dynamic APSP with path reporting. In Section 6, we present our algebraic spanner algorithm, which uses the data structure from Section 5. In Section 7 we present two additional applications of the data structure from Section 5: our dynamic algorithm for approximate APSP with path reporting and our dynamic Steiner tree algorithm. Finally, in Section 8, we present our combinatorial dynamic spanner and emulator algorithms.

2 Technical overview

Conditional lower bounds. We first outline our OMv-based constructions. Instead of reducing from the OMv problem, we reduce from the related OuMv problem, which is defined as follows. We are given an $n \times n$ matrix M that can be preprocessed. Then, an online sequence of vector pairs $(u^1, v^1), \dots, (u^n, v^n)$ is presented and the goal is to compute each $(u^i)^\top M v^i$ before seeing the next pair. A reduction from OMv to OuMv is known [56].

For both our fully dynamic and incremental/decremental lower bounds from OMv we begin with the same basic gadget. Given the matrix M from the OuMv instance, we construct a bipartite graph A, B where $A = \{a_1, \dots, a_n\}$, $B = \{b_1, \dots, b_n\}$, and the edge (a_i, b_j) is present if and only if $M_{i,j} = 1$.

The fully dynamic construction is shown in Figure 1. We begin by taking a number c of disjoint copies G_1, \dots, G_c of the basic gadget and an additional set of $c + 1$ isolated vertices w_0, \dots, w_c . Each basic

gadget will introduce error to the approximation, so larger c means that we are showing a lower bound for algorithms with higher approximation factors but faster running times.

After constructing this initial graph, we start n dynamic phases. In phase i , we are given the vectors u^i and v^i of the OuMv instance. For each $1 \leq j \leq c$ and each k with $u_k^i = 1$, insert an edge between w_{j-1} and $a_k \in G_j$. Similarly, for each $1 \leq j \leq c$ and each k with $v_k^i = 1$, insert an edge between w_j and $b_k \in G_j$. We remove these edges after the phase is over.

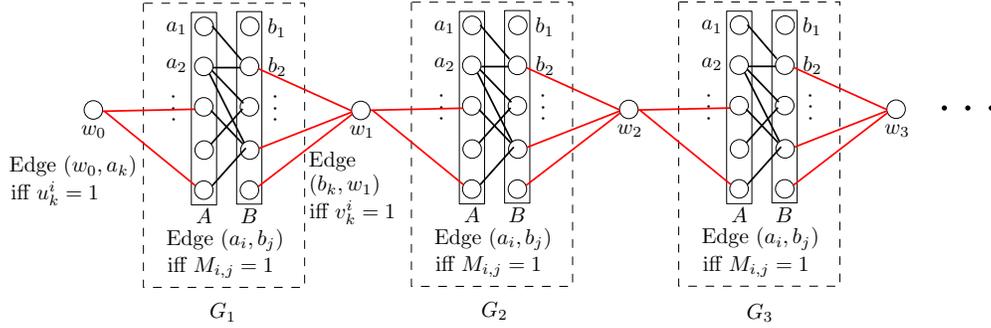


Figure 1: The construction for fully dynamic algorithms. The red edges are dynamically added in phase i .

At the end of each phase, we run Breadth-First Search (BFS) on the dynamic emulator to estimate the distance between w_0 and w_c , which we claim provides the answer to this phase of the OuMv instance. In particular, note that for any i the distance between w_i and w_{i+1} is 3 if and only if $(u^i)^\top M v^i = 1$, and otherwise this distance is at least 5. Also, since the emulator has $O(n^{2-\epsilon})$ edges, the resulting algorithm would solve OuMv in $O(n^{3-\epsilon})$ time.

The incremental construction is similar, however we cannot remove edges at the end of each phase. To get around this, we replace each w_i with a path and insert edges incident to a different vertex in the path at each phase. The resulting construction is shown in Figure 2.

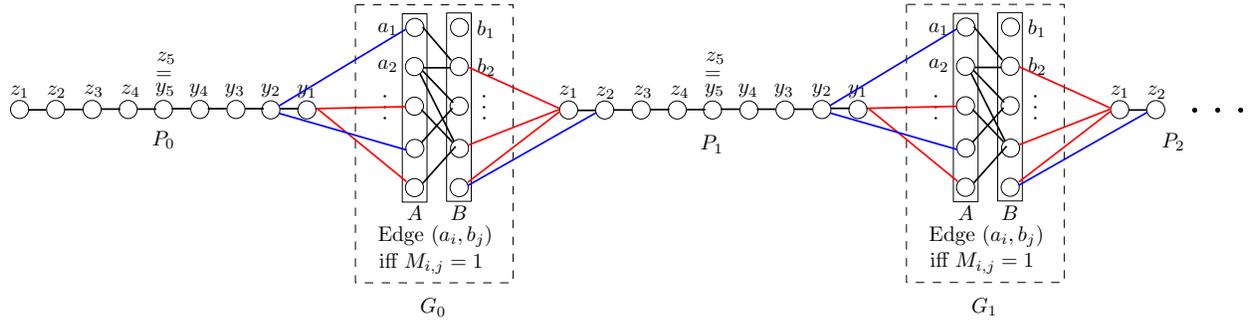


Figure 2: The construction for incremental algorithms. The red edges are dynamically added in phase 1 and the blue edges are dynamically added in phase 2.

The decremental construction is roughly the reverse of the incremental construction.

For the k -clique-based constructions, we instead reduce from the k -cycle problem (a reduction from k -clique to k -cycle is known [68]). The k -cycle constructions for both the fully and partially dynamic settings follow a similar structure to the OuMv constructions but use a different basic gadget. The basic gadget is built by using color coding and taking a layered version of the graph where each color is a layer and only edges between vertices of adjacent colors are present.

Algebraic Fully Dynamic Spanner Algorithm. Let us next present an algorithm to maintain a $(1 + \epsilon, n^{o(1)})$ -spanner on a fully dynamic graph with worst-case update time $O(n^{1.529})$.

Let $k = \sqrt{\log n}$. We sample sets $V = A_0 \supseteq A_1, \dots \supseteq A_k \supseteq A_{k+1} = \emptyset$ where A_i for $i \in [1, k+1]$ is obtained by sampling each vertex in V with probability $n^{-i/k} \log n$ (and to make the sets nesting add it to all A_j where $j \leq i$, we assume that A_{k+1} is empty which can be achieved by resampling a constant number of times in expectation). Given these sets, we say that each $a \in A_\ell \setminus A_{\ell+1}$ is *active* if for no $j > \ell$ there exists a vertex $a' \in A_j \setminus A_{j+1}$ with $\text{dist}_{\tilde{G}}(a, a') \lesssim (\frac{1}{\epsilon})^j$. Using this definition of activeness, we can show by a simple hitting set argument that each active vertex $a \in A_\ell \setminus A_{\ell+1}$ has in its ball to radius $\sim (\frac{1}{\epsilon})^{\ell+1}$ at most $\tilde{O}(n^{(\ell+1)/k})$ vertices w.h.p..

Given this set-up, a natural way to construct a spanner H , would be to find for each level $\ell \in [0, k]$, the active vertices in $A_\ell \setminus A_{\ell+1}$ and to include their shortest path trees truncated at radius $\sim (\frac{1}{\epsilon})^{\ell+1}$. For the number of edges of H , it is not hard to see that each for each $\ell \in [0, k]$, there are at most $\tilde{O}(n^{1-\ell/k})$ vertices that are active in $A_\ell \setminus A_{\ell+1}$. Each of these vertices contributes a single edge for each vertex in its truncated ball (except for itself), and as discussed above we have that each ball is of size at most $\tilde{O}(n^{(\ell+1)/k})$. Thus, we have that H has at most $\tilde{O}(n^{1+1/k}) = n^{1+o(1)}$ edges.

For the stretch factor, observe that for any vertices $s, t \in V$, with shortest path $\pi_{s,t}$, we have that for $s \in A_\ell \setminus A_{\ell+1}$ for some level ℓ , if s is active, we can simply travel along the $\pi_{s,t}$ to some vertex s' that is closer to t (since the truncated shortest path tree of s is included in H) and then expose the shortest path $\pi_{s',t}$ inductively. Or, we have that there is a vertex $a' \in A_j \setminus A_{j+1}$ at distance $\lesssim (\frac{1}{\epsilon})^j$ to vertex s (with $j > \ell$). Choosing a' to be the vertex that is at this distance to a with the largest possible j , we will be able to argue that a' is active. Thus, we can travel from s to a' to a vertex s' on $\pi_{s,t}$ at distance roughly $(\frac{1}{\epsilon})^{j+1}$ along the shortest path tree at a' truncated at depth $\sim (\frac{1}{\epsilon})^{j+1}$ that was included in H . It is not hard to see that the error induced for visiting a' can be subsumed in a multiplicative $(1 + O(\epsilon))$ -approximation. However, this only works if s and t are at distance $\gtrsim (\frac{1}{\epsilon})^{j+1}$, otherwise it induces an additive error of $n^{o(1)}$. This explains why we obtain a $(1 + \epsilon, n^{o(1)})$ -approximation.

Unfortunately, while this set-up of H is sensible, consider the example of the complete graph. Then, there would be a vertex $a \in A_k \setminus A_{k+1}$ (which is active since A_{k+1} is empty) where visiting the truncated ball at a would take time $\tilde{O}(n^2)$, which is by far too expensive for our algorithm.

To overcome this issue, instead of inserting truncated shortest path trees to H , we only insert for any active vertex $a \in A_\ell \setminus A_{\ell+1}$, the shortest paths to other vertices in $A_\ell \setminus A_{\ell+1}$ in the truncated shortest path tree of a . We then fix a threshold $\gamma \approx \lceil 0.529 \cdot k \rceil$, and can use the algebraic data structure from Theorem 1.3 to maintain the distances of vertices in A_γ (and thereby $A_{\gamma+1}, A_{\gamma+2}, \dots, A_k, A_{k+1}$) without explicitly maintaining the balls of the active vertices. For active vertices in some set $A_\ell \setminus A_{\ell+1}$ for $\ell < \gamma$, we can compute the balls explicitly after every update. This is because each such ball only contains $\tilde{O}(n^{(\ell+1)/k})$ vertices, and therefore the induced graph can contain at most $\tilde{O}(n^{2(\ell+1)/k})$ edges, which implies that we overall, spend at most time $\tilde{O}(n^{1-\ell/k}) \cdot n^{2(\ell+1)/k} = \tilde{O}(n^{1+(\ell+2)/k})$ time for computing all such balls. We refer the reader to section Section 6 for a proper analysis of the running time.

Finally, we point out that so far H only contains shortest paths between active vertices in the same set $A_\ell \setminus A_{\ell+1}$ (if they are reasonably close). However, to have a path between vertices on different levels, we also add a $O(\log n)$ spanner \tilde{G} of G to H . Such a spanner is simple to maintain, for example [48] shows how to maintain such a spanner with $\tilde{O}(n)$ edges and $\tilde{O}(1)$ amortized update time.

The idea of the approximation proof then becomes the following for some path $\pi_{s,t}$: Let i be the largest index such that an active vertex $a \in A_i$ is at distance at most $\sim \epsilon^{-i}$ to s . Let w be the farthest vertex from s on $\pi_{s,t}$ such that (1) the distance from s to w is at most $\sim \epsilon^{-(i+1)}$, and (2) w has distance at most $\sim \epsilon^{-i}$ to

an active vertex $a' \in A_i$. Such a vertex w exists since we could have $w = s$ and $a' = a$. It is then apparent that the distance between a and a' is $\lesssim \epsilon^{-(i+1)}$ and since a' is active, we can ensure that the shortest path from a to a' is in H . Further, we can use the paths in the spanner \tilde{G} (which belongs to H) to get from s to a and from a' to w ; since these two distances are small, it suffices to have an $O(\log n)$ multiplicative error for them. Now, observe that this induces additive error along the path segment from s to w of $\tilde{O}(\epsilon^{-i})$ (by the triangle inequality). We either have that a and a' are roughly at distance $\sim \epsilon^{-(i+1)}$ (which suffices to subsume the additive error in the multiplicative error), or we have that for the next path segment of length $\sim \epsilon^{-(i+1)}$, no vertex is close to any active vertex in A_i . Thus, when we repeat the whole argument for the next path segment, we get that vertices on lower levels are active, which means that they induce less additive error. This allows us to subsume the additive error from higher levels into multiplicative error for a series of segments of lower levels.

We refer the reader to Section 6 for the full details of the algorithm.

Fully dynamic APSP with path reporting. Our data-structure of theorem 1.3 is an augmentation of Sankowski's [75] data structure to support fast successor queries. Essentially, Sankowski showed how to reduce the problem of maintaining the short distances in a dynamic unweighted graph to the dynamic matrix inverse problem, by representing the path lengths as degrees of the adjoint of a polynomial matrix. He then showed how to efficiently maintain the inverse of a matrix subject to entry updates, allowing for fast distance queries. We extend his techniques to maintain products of matrices and the inverse, and show how to use these products to extract successor information similarly to Seidel's path reporting algorithm for static APSP [77]. Let us begin here by sketching Sankowski's data-structure [75], formally reviewed in Section 5.2, and then present the high level of our augmentation.

Short Distances to Dynamic Matrix Inverse. [75] showed how to encode path lengths of an unweighed graph in the adjoint of a matrix, that is, given a adjacency matrix A_{ij} with a random integer entry if $(i, j) \in E$, then the lowest degree non-zero term of the polynomial $\text{adj}(\mathbb{1} - uA)_{ij}$ over the variable u is the distance d_{ij} whp (Lemma 5.6). In this manner maintaining $\text{adj}(\mathbb{1} - uA)_{ij} \bmod u^{D+1}$, for some distance parameter D , allows us to query a distance $i \rightarrow j$ correctly whp if $d_{ij} \leq D$. Note that $\text{adj}M = \det M \times M^{-1}$, s.t. it suffices just to maintain $\det M$ and $M^{-1} \bmod u^{D+1}$.

Dynamic Matrix Inverse. We detail the algebraic tools developed by Sankowski [75] to maintain the inverse of a matrix M dynamically and over a ring in Section 5.2. The main idea is to maintain explicitly (i.e. all n^2 entries) two matrices T, N , s.t. we maintain the invariant

$$M^{-1} = T(\mathbb{1} + N) \tag{1}$$

where, initially, $T = M^{-1}$ and $N = 0$, and as later shown each single entry update to M corresponds to a single row update to N (and no modifications to T !). After m updates, N has at most m non-zero rows, and we can exploit this *sparsity* of N to quickly compute its row-updates in $O(mn)$, and every $m = n^\kappa$ updates we reset $T \leftarrow T + TN$, $N \leftarrow 0$, in $O(n^{\omega(1,1,\kappa)-\kappa})$ time on average. In this manner, we guarantee that N is always sparse, and updates take amortized time

$$O(n^{\omega(1,1,\kappa)-\kappa} + n^{1+\kappa}) \tag{2}$$

for some parameter $\kappa \in (0, 1)$ which we can later optimize over. Entry queries (i, j) are now straightforward, as it suffices to compute the dot product

$$M_{ij}^{-1} = e_i^T T(\mathbb{1} + N)e_j = T_{ij} + (e_i^T T) \cdot (Ne_j) \tag{3}$$

which can be done in $O(n^\kappa)$ time since a given column of N has at most $O(n^\kappa)$ non-zero entries. In Corollary 5.8 [75] showed that we can maintain these matrices over a ring mod u^{D+1} by introducing a multiplicative factor of $\tilde{O}(D)$ to the runtimes described above, s.t. now if $M = \mathbb{1} - uA$ we can query any distance $d \leq D$ in the graph in time $\tilde{O}(Dn^\kappa)$.

Successor Queries to Product Maintenance. There are two main ingredients to our augmentation the data-structure of [75]. The first is to reduce the successor query of a pair (i, j) of an unknown number of distinct successors to that of a single successor, using a known sparsification trick used first in Seidel’s algorithm for static, undirected and unweighted APSP [77]. This only introduces a $\log^2 n$ multiplicative factor to the runtime and we defer the formal argument to Lemmas 5.12 and 5.13. The second ingredient is to show how to find a single successor by finding a witness of the product $(A \cdot \text{adj}(\mathbb{1} - uA))_{ij}$. The key new insight is that if the distance $1 < d_{ij} \leq D$, then $\text{adj}(\mathbb{1} - uA)_{ij}$ has minimum degree d_{ij} , and thereby the product $(A \cdot \text{adj}(\mathbb{1} - uA))_{ij}$ must have minimum degree $d_{ij} - 1$. This is since there must exist a *unique* witness s (the single successor!) s.t. A_{is} is non-zero, corresponding to an edge, and $\text{adj}(\mathbb{1} - uA)_{sj}$ has minimum degree $d_{sj} = d_{ij} - 1$, corresponding to the length of the shortest path from s to j .

We can find this single witness by computing its bitwise description, that is, defining $O(\log n)$ versions of the adjacency matrix A , $A^{(l)}$ for $l \in [O(\log n)]$, where we null the p th column of $A^{(l)}$ if the l th bit of p is 0. As there is only a single witness s , the minimum degree of the product $(A^{(l)} \cdot \text{adj}(\mathbb{1} - uA))_{ij}$ is $d_{sj} = d_{ij} - 1$ only if the column of s is selected, that is, the l th bit $s_l = 1$. In this manner, if we query the $O(\log n)$ products $(A^{(l)} \cdot \text{adj}(\mathbb{1} - uA))_{ij}$, the 1-bits $s_l = 1$ are exactly the products l s.t. the minimum degree is correct. This allows us to extract the successor description in a polylog number of queries to products $(E \cdot \text{adj}(\mathbb{1} - uA))_{ij}$ for given matrices E .

Product Maintenance. The last detail in our successor query augmentation is to show how to maintain products $(E \cdot \text{adj}(\mathbb{1} - uA))$, where we can modify entries of E and A , and query entries i, j of the result. Note again that it suffices to maintain $(E \cdot (\mathbb{1} - uA)^{-1})$, as opposed to the adjoint, just by multiplying by the determinant. We do so by following the inverse maintenance algorithm and explicitly maintaining the matrices T, N and $V \equiv ET$, s.t. we maintain the invariant

$$EM^{-1} = ET(\mathbb{1} + N) = V(\mathbb{1} + N) \quad (4)$$

We address updates to E and to A completely differently. Updates to A follow the original lazy construction, where we simply perform a row-update to N , and every n^κ updates we “reset” $V \leftarrow V(\mathbb{1} + N), T \leftarrow T(\mathbb{1} + N), N \leftarrow 0$. We note that correctness follows by associativity, s.t. although matrices V and T are dense we can still exploit the sparsity of N to compute their updates independently in time $O(Dn^{1+\kappa} + Dn^{\omega(1,1,\kappa)-\kappa})$ on average. Entry-Updates to E , $E \leftarrow E + e_{ij}$, are much simpler. We once again use associativity to compute the row update $V \leftarrow V + e_{ij}T$ in $\tilde{O}(Dn)$ time. Finally, to query an entry of the product $(EM^{-1})_{ij}$, we follow analogously to [75] and compute the dot product

$$(EM^{-1})_{ij} = e_i^T V(\mathbb{1} + N)e_j = V_{ij} + (e_i^T V) \cdot (Ne_j) \quad (5)$$

and since V and N are maintained explicitly, this takes time $\tilde{O}(Dn^\kappa)$, the exact same as the distance queries. Overall, this allows for $\tilde{O}(Dn^\kappa)$ time successor queries, and by iterating, short path queries of length $\leq D$ in time $\tilde{O}(D^2n^\kappa)$.

3 Preliminaries

We let $G = (V, E)$ denote an undirected unweighted dynamic input graph, where $n = |V|$ and $m = |E|$. For any graph H , and two vertices $a, b \in V(H)$, we denote by $\mathbf{dist}_H(a, b)$ the distance between the two vertices in G and let $\pi_{a,b,H}$ denote a corresponding shortest path between a and b . If the graph H , especially when we use the input graph G , is clear from the context, we simply use $\pi_{a,b}$. We define $B_H(s, r)$ in the graph H , to be the ball rooted at s with radius r , i.e. the set of vertices $B_H(s, r) = \{w \in V(H) \mid \mathbf{dist}_H(s, w) \leq r\}$. Throughout the article, we often use the data structure stated below that is sometimes referred to as the Even-Shiloach (ES) tree.

Lemma 3.1 (c.f. [47]). *For any vertex $s \in V$, radius r , there is a deterministic data structure on a partially dynamic graph G that reports for every $w \in B(s, r)$, the distance $\mathbf{dist}(s, w)$. In fact, the data structure maintains explicitly the shortest path tree in $G[B(s, r)]$. The total update time of the data structure is $O(mr)$ time where m is the maximum number of edges ever in $G[B(s, r)]$.*

Let ω be the infimum over all reals such that $n \times n$ matrices can be multiplied in $O(n^{\omega+\epsilon})$ time for all $\epsilon > 0$. It is known that $\omega \in [2, 2.373)$ [83, 66]. More generally, let $\omega(a, b, c)$ be the infimum over all reals such that an $n^a \times n^b$ matrix can be multiplied by an $n^b \times n^c$ matrix in $O(n^{\omega(a,b,c)+\epsilon})$ time for all $\epsilon > 0$. A notable result is that for $b \leq 0.313$, $\omega(1, b, 1) = 2$ [67].

4 Conditional lower bounds

We present conditional lower bounds for amortized algorithms in the fully dynamic, incremental, and decremental settings. Our constructions for amortized algorithms in the fully dynamic setting also imply lower bounds for the incremental and decremental settings, but only for worst-case update times. We present separate constructions for the *amortized* incremental and decremental settings.

We first describe why our fully dynamic conditional lower bounds also apply to the incremental and decremental settings for worst-case update times. This is due to the nature of our reductions: all of our reductions produce an initial graph on which we perform update stages that only insert or only delete (we can choose which) a batch of edges, ask a query and undo the changes just made, returning to the initial graph. An incremental (resp. decremental) algorithm can be used for this type of dynamic graph by performing the insertions (resp. deletions) and then rolling back the data structure to the initial graph and repeating.

4.1 Conditional lower bounds from the OMv conjecture

4.1.1 Statement of results

We prove conditional lower bounds from the OMv conjecture for dynamic emulator maintenance in the fully dynamic, incremental, and decremental settings. We prove the following theorem for the fully dynamic setting, which also extends to the incremental and decremental settings but only for worst-case update times.

Theorem 4.1. *Under the OMv conjecture, for an n -vertex fully dynamic graph with at most m edges at all times, for every constant $\epsilon > 0$, there is no algorithm for maintaining a $(1 + \alpha, \beta)$ -emulator for any $\alpha \in [0, 2/3)$ and integer $\beta \geq 0$, with $O(m^{1-\epsilon}(\frac{2-3\alpha}{\beta}))$ edges, polynomial preprocessing time, and amortized update time $O(n^{1-\epsilon}(\frac{2-3\alpha}{\beta})^2)$ such that over a polynomial number of edge updates the error probability is at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

In particular, for the natural setting where α is constant and $\beta = n^{o(1)}$ we have the following corollary:

Corollary 4.2. *Under the OMv conjecture, for an n -vertex fully dynamic graph with at most m edges at all times, for every constant $\epsilon > 0$, there is no algorithm for maintaining a $(1 + \alpha, n^{o(1)})$ -emulator for any constant $\alpha \in [0, 2/3)$ with $O(m^{1-\epsilon})$ edges, polynomial preprocessing time, and amortized update time $O(n^{1-\epsilon})$ such that over a polynomial number of edge updates the error probability is at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

For the incremental and decremental settings, we prove the following theorem.

Theorem 4.3. *Under the OMv conjecture, for an n -vertex incremental or decremental graph with m edge insertions or deletions, for every constant $\epsilon > 0$, there is no algorithm for maintaining a β -additive emulator for any integer $\beta \geq 0$, with $O(m^{1-\epsilon}/\beta)$ edges, polynomial preprocessing time, and total update time $O(mn^{1-\epsilon}/\beta^2)$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

In particular, for the natural setting where $\beta = n^{o(1)}$ we have the following corollary:

Corollary 4.4. *Under the OMv conjecture, for an n -vertex incremental or decremental graph with m edge insertions or deletions, for every constant $\epsilon > 0$, there is no algorithm for maintaining a $n^{o(1)}$ -additive emulator with $O(m^{1-\epsilon})$ edges, polynomial preprocessing time, and total update time $O(mn^{1-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

4.1.2 Preliminaries

Our reductions are from the Online Vector-Matrix-Vector Multiplication problem (OuMv). A reduction from OMv to OuMv is known:

Theorem 4.5 (OuMv: Theorem 2.7 from [56]). *The OMv conjecture implies that for any constant $\epsilon > 0$, there is no algorithm for OuMv with polynomial preprocessing time and computation time $O(n^{3-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

Given an instance of OuMv, we introduce a basic gadget that we will use in all of our constructions.

The basic gadget Let M be the $n \times n$ input matrix for the OuMv instance. We construct a gadget as shown in Figure 3. The gadget consists of a bipartite graph A, B where $A = \{a_1, \dots, a_n\}$, $B = \{b_1, \dots, b_n\}$, and the edge (a_i, b_j) is present if and only if $M_{i,j} = 1$.

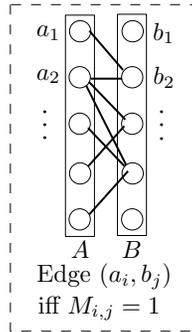


Figure 3: The basic gadget for reductions from OMv.

4.1.3 Reduction for fully dynamic algorithms

In this section we prove Theorem 4.1.

Construction Let $c = \lceil \frac{\beta}{2-3\alpha} \rceil + 1$ and take c disjoint copies G_1, \dots, G_c of the basic gadget (from Figure 3). Let w_0, \dots, w_c be an additional set of $c + 1$ isolated vertices.

Now, we start n dynamic phases. In phase i , we are given the vectors u^i and v^i of the OuMv instance. For each $1 \leq j \leq c$ and each k with $u_k^i = 1$, insert an edge between w_{j-1} and $a_k \in G_j$. Similarly, for each $1 \leq j \leq c$ and each k with $v_k^i = 1$, insert an edge between w_j and $b_k \in G_j$. See Figure 4.

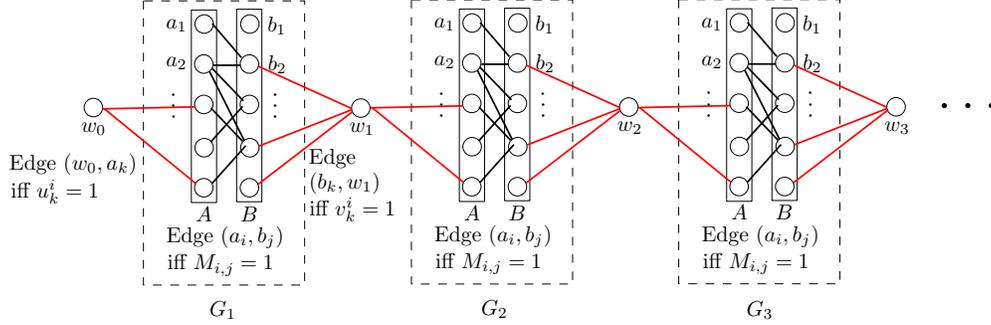


Figure 4: The construction for fully dynamic algorithms. The red edges are dynamically added in phase i .

Throughout all of the edge updates, we maintain our dynamic emulator. At the end of each phase, we run a single source shortest paths computation via Breadth-First Search (BFS) on the emulator to estimate the distance between w_0 and w_c .

If the estimated distance between w_0 and w_c is less than $5c$, we return 1 for this phase of the OuMv instance, and otherwise we return 0.

Following the end of each phase, we remove all of the edges added during that phase.

Correctness First we will show that if $u^i M v^i = 1$ then our algorithm returns 1 for the i^{th} phase. If $u^i M v^i = 1$ then there exists j, k such that $u_j^i = M_{j,k} = v_k^i = 1$. Thus, the basic gadget contains the edge (a_j, b_k) . Also, for all $1 \leq \ell \leq c$, in the i^{th} phase we add an edge between $w_{\ell-1}$ and $a_j \in G_\ell$ and an edge between w_ℓ and $b_k \in G_\ell$. Thus, there is a path of length 3 from $w_{\ell-1}$ to w_ℓ through a_j and b_k . Therefore, $\mathbf{dist}(w_0, w_c) \leq 3c$. Thus, the estimate of $\mathbf{dist}(w_0, w_c)$ returned by our $(\alpha + 1, \beta)$ -emulator is at most $3c(\alpha + 1) + \beta$, which is less than $5c$ since $c > \frac{\beta}{2-3\alpha}$. Thus, our algorithm returns 1 for the i^{th} phase.

Now we will show that if our algorithm returns 1 in the i^{th} phase then $u^i M v^i = 1$. If our algorithm returns 1, then the estimate of $\mathbf{dist}(w_0, w_c)$ returned by our $(1 + \alpha, \beta)$ -emulator is less than $5c$, so the true distance $\mathbf{dist}(w_0, w_c)$ is also less than $5c$.

First, we observe that the layered structure of the graph ensures that every path between w_0 and w_c must contain every w_ℓ in chronological order. That is, every shortest path between w_0 and w_c must contain as a subpath a shortest path from $w_{\ell-1}$ to w_ℓ for all $1 \leq \ell \leq c$. Then since the graph is c identical copies of a gadget, we have that $\mathbf{dist}(w_0, w_c) = c \cdot \mathbf{dist}(w_0, w_1)$.

Since $\mathbf{dist}(w_0, w_c) < 5c$, we know that $\mathbf{dist}(w_0, w_1) < 5$. Furthermore, since the graph is bipartite and w_0 and w_1 are on opposite sides of the bipartition, $\mathbf{dist}(w_0, w_1)$ must be odd so $\mathbf{dist}(w_0, w_1) \leq 3$. Observe that the only possible paths of length 3 between $w_{\ell-1}$ and w_ℓ contain a vertex $a_j \in G_\ell$ followed by a vertex $b_k \in G_\ell$. If such a path exists in the i^{th} phase, then the basic construction ensures that $M_{j,k} = 1$ and the dynamic phase ensures that $u_j^i = 1$ and $v_k^i = 1$. Thus, $u^i M v^i = 1$.

Running time Let n' be the number of vertices in the dynamic graph and let m' be the maximum number of edges ever in the dynamic graph. We first calculate n' and m' . Each basic gadget contains $2n$ vertices and at most n^2 edges. Thus, c copies of the basic gadget contain $2cn$ vertices and cn^2 edges. There are

also an additional $c + 1$ vertices w_ℓ . During each of the n phases we add at most $2nc$ edges. Thus, the total number of vertices is $n' = O(nc)$ and the total number of edge updates over the entire sequence is $O(cn^2)$, so $m' = O(cn^2)$.

Suppose that the emulator has $O(m^{1-\epsilon}(2-3\alpha)/\beta)$ edges for $\epsilon > 0$ and has polynomial preprocessing time and amortized update time $O(n^{1-\epsilon}((2-3\alpha)/\beta)^2)$.

Then, our dynamic emulator algorithm has amortized update time $O(n^{1-\epsilon}(\frac{2-3\alpha}{\beta})^2) = O((nc)^{1-\epsilon}(\frac{2-3\alpha}{\beta})^2)$. Since there are $O(cn^2)$ edge updates, the total update time of the dynamic emulator algorithm is $O(n^{3-\epsilon})$ since $c = \lceil \frac{\beta}{2-3\alpha} \rceil + 1$.

Additionally, n times during the algorithm, we run a single call of BFS on the emulator. The number of edges in the emulator is $O(m^{1-\epsilon}(\frac{2-3\alpha}{\beta})) = O((cn^2)^{1-\epsilon}(\frac{2-3\alpha}{\beta})) = O(n^{2-\epsilon})$ since $c = \lceil \frac{\beta}{2-3\alpha} \rceil + 1$. Thus, running BFS takes total time $O(n^{3-\epsilon})$.

Putting everything together, our dynamic emulator algorithm implies an algorithm for OuMv with polynomial preprocessing time and computation time $O(n^{3-\epsilon})$ for $\epsilon > 0$, contradicting the OMv conjecture.

4.1.4 Reduction for incremental and decremental algorithms.

In this section we prove Theorem 4.3.

Construction The construction will be similar to the fully dynamic construction, but with different interactions between consecutive copies of the basic gadget. We first describe the incremental construction.

Starting with an empty graph, we perform edge insertions to construct the following graph. Take $\beta + 1$ disjoint copies $G_1, \dots, G_{\beta+1}$ of the basic gadget (from Figure 3). Then, add $\beta + 2$ paths $P_0, \dots, P_{\beta+1}$ each on $2n - 1$ new vertices. Call the vertices of each path $z_1, z_2, \dots, z_n = y_n, y_{n-1}, \dots, y_1$. In other words, the middle node of each path has two names, z_n and y_n .

Now, we start n phases. In phase i , we are given the vectors u^i and v^i of the OuMv instance. For each $1 \leq j \leq \beta + 1$ and each k with $u_k^i = 1$, insert an edge between $y_i \in P_{j-1}$ and $a_k \in G_j$. Similarly, for each $1 \leq j \leq \beta + 1$ and each k with $v_k^i = 1$, insert an edge between $z_i \in P_j$ and $b_k \in G_j$. See Figure 5.

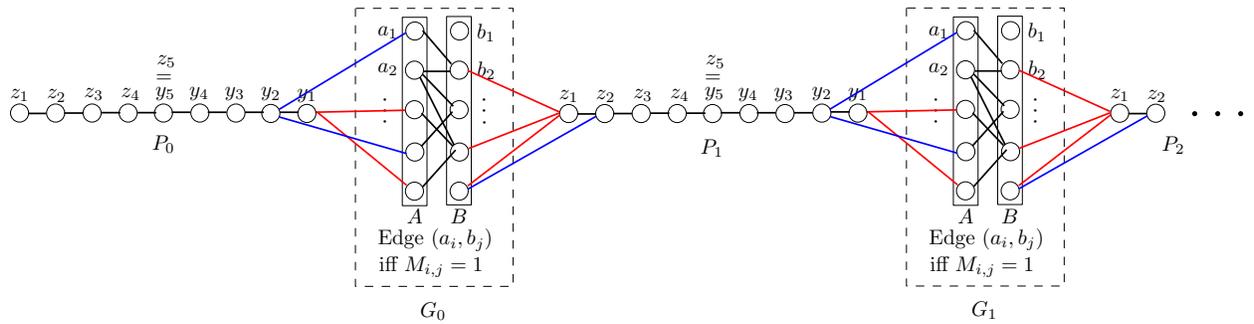


Figure 5: The construction for incremental algorithms. The red edges are dynamically added in phase 1 and the blue edges are dynamically added in phase 2.

Throughout all of the edge updates, we maintain our incremental emulator. At the end of each phase, we run a single call of BFS on the emulator to estimate the distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$.

If the estimated distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ at the end of phase i is less than $4(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$, we return 1 for phase i of the OuMv instance, and otherwise we return 0.

Now, we describe the decremental construction, which is similar to the incremental construction. The initial graph consists of $G_1, \dots, G_{\beta+1}$ and $P_0, \dots, P_{\beta+1}$ from the incremental construction, as well as an

edge for all $1 \leq \ell \leq \beta + 1$ from $z_i \in P_\ell$ to every vertex in $B \subseteq G_\ell$, and an edge for all $1 \leq \ell \leq \beta + 1$ from $y_i \in P_{\ell-1}$ to every vertex in $A \subseteq G_\ell$.

Now, we start n dynamic phases. In phase i , we are given the vectors u^i and v^i of the OuMv instance. For each $1 \leq j \leq \beta + 1$, and each k with $u_k^i = 0$, delete the edge between $y_{n-i+1} \in P_{j-1}$ and $a_k \in G_j$. Similarly, for each $1 \leq j \leq \beta + 1$ and each k with $v_k^i = 0$, delete the edge between $z_{n-i+1} \in P_j$ and $b_k \in G_j$.

Throughout all of the edge updates, we maintain our decremental emulator. At the end of each phase, we run a single call to BFS on the emulator to estimate the distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$.

Following the end of each phase i , for each $1 \leq j \leq \beta + 1$, we delete all edges between $y_{n-i+1} \in P_{j-1}$ and $A \subseteq G_j$ and all edges between $z_{n-i+1} \in P_j$ and $B \subseteq G_j$.

If the estimated distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ at the end of phase i is less than $4(\beta + 1) + (\beta + 2)(2n - 2(n - i + 1)) + 2((n - i + 1) - 1)$, we return 1 for phase i of the OuMv instance, and otherwise we return 0. Note that this threshold is exactly the threshold from the incremental algorithm but with i replaced with $n - i + 1$.

Correctness The following argument is written for the incremental setting but the same argument applies for the decremental setting.

First we will show that if $u^i M v^i = 1$ then our algorithm returns 1 for the i^{th} phase. If $u^i M v^i = 1$ then there exists j, k such that $u_j^i = M_{j,k} = v_k^i = 1$. Thus, the basic gadget contains the edge (a_j, b_k) . Also, for all $1 \leq \ell \leq \beta + 1$, in the i^{th} phase we add an edge between $y_i \in P_{\ell-1}$ and $a_j \in G_\ell$ and an edge between $z_i \in P_\ell$ and $b_k \in G_\ell$.

Thus, for all $1 \leq \ell \leq \beta + 1$, there is a path of length 3 from $y_i \in P_{\ell-1}$ to $z_i \in P_\ell$ through $a_j \in G_\ell$ and $b_k \in G_\ell$. Also, for each $0 \leq \ell \leq \beta + 2$, there is a path along P_ℓ from $z_i \in P_\ell$ to $y_i \in P_\ell$ of length $(2n - 2) - 2(i - 1) = 2n - 2i$. Finally, there is a path of length $i - 1$ from $z_1 \in P_0$ to $z_i \in P_0$ and a path of length $i - 1$ from $y_i \in P_{\beta+2}$ to $y_1 \in P_{\beta+2}$. Concatenating all of these paths, we have that $\mathbf{dist}(z_1 \in P_0, y_1 \in P_{\beta+1}) \leq 3(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$. Thus, the estimate of $\mathbf{dist}(z_1 \in P_0, y_1 \in P_{\beta+1})$ returned by our β -additive emulator is at most $4(\beta + 1) - 1 + (\beta + 2)(2n - 2i) + 2(i - 1)$. Thus, our algorithm returns 1 for the i^{th} phase.

Now we will show that if our algorithm returns 1 in the i^{th} phase then $u^i M v^i = 1$. If our algorithm returns 1, then the estimate of $\mathbf{dist}(z_1 \in P_0, y_1 \in P_{\beta+1})$ returned by our β -additive emulator is less than $4(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$, so the true distance $\mathbf{dist}(z_1 \in P_0, y_1 \in P_{\beta+1})$ is also less than $4(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$.

First, we observe that the layered structure of the graph ensures that every path between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ contains each z_i and y_i in order from P_0 to $P_{\beta+1}$. That is, every shortest path between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ is composed of precisely following subpaths:

- A shortest path from $z_1 \in P_0$ to $z_i \in P_0$. The only simple path connecting these vertices is of length $i - 1$.
- A shortest path from $y_i \in P_{\beta+1}$ to $y_1 \in P_{\beta+1}$. The only simple path connecting these vertices is of length $i - 1$.
- A shortest path from $z_i \in P_\ell$ to $y_i \in P_\ell$ for all $1 \leq \ell \leq \beta + 2$. The only simple path connecting these vertices is of length $(2n - 2) - 2(i - 1) = 2n - 2i$.
- A shortest path from $y_i \in P_{\ell-1}$ to $z_i \in P_\ell$ for all $1 \leq \ell \leq \beta + 1$. Since the graph is a series of identical copies of a gadget, we know that $\mathbf{dist}(y_i \in P_{\ell-1}, z_i \in P_\ell)$ is the same for all ℓ . Furthermore, we

know the length of each of the previous three types of subpaths and we know that $\mathbf{dist}(z_1 \in P_0, y_1 \in P_{\beta+1}) < 4(\beta+1) + (\beta+2)(2n-2i) + 2(i-1)$, so we conclude that each $\mathbf{dist}(y_i \in P_{\ell-1}, z_i \in P_\ell) < 4$.

Due to the layering of the graph, for all $1 \leq \ell \leq \beta + 1$ the shortest path between $y_i \in P_{\ell-1}$ and $z_i \in P_\ell$ must contain vertices $a_j \in G_\ell$ and $b_k \in G_\ell$ for some j, k . Since $\mathbf{dist}(y_i \in P_{\ell-1}, z_i \in P_\ell) < 4$, there are no other vertices on this shortest path. The basic construction ensures that since the edge (a_j, b_k) exists, we have $M_{j,k} = 1$, and the dynamic phase ensures that since the edge $(y_i \in P_{\ell-1}, a_j \in G_\ell)$ exists, we have $u_j^i = 1$ and since the edge $(b_k \in G_\ell, z_i \in P_\ell)$ exists, we have $v_k^i = 1$. Thus, $u^i M v^i = 1$.

Running time Let n' be the number of vertices in the dynamic graph and let m' be number of edge insertions or deletions. We first calculate n' and m' . Each basic gadget contains $2n$ vertices and at most n^2 edges. Thus, $\beta + 1$ copies of the basic gadget contain $2(\beta + 1)n$ vertices and $(\beta + 1)n^2$ edges. Additionally, we have $(\beta + 2)$ paths on $(2n - 1)$ vertices each, for a total of $(2n - 1)(\beta + 2)$ additional vertices. During each of the n phases there are at most $2n(\beta + 1)$ edge updates. Thus, the total number of vertices is $n' = O(\beta n)$ and the total number of edges updates is $m' = O(\beta n^2)$.

Now assume that our incremental or decremental emulator has $O(m'^{1-\epsilon}/\beta)$ edges, polynomial preprocessing time and total update time $O(m' n'^{1-\epsilon}/\beta^2)$.

Then the total update time of the emulator is $O(m' n'^{1-\epsilon}/\beta^2) = O(\beta n^2 (\beta n)^{1-\epsilon}/\beta^2) = O(n^{3-\epsilon})$.

Additionally, n times during the algorithm, we run BFS on the emulator. The number of edges in the emulator is $O(m'^{1-\epsilon}/\beta) = O((\beta n^2)^{1-\epsilon}/\beta) = O(n^{2-\epsilon})$. Thus, running the BFS calls takes total time $O(n^{3-\epsilon})$.

Putting everything together, our incremental or decremental emulator algorithm implies an algorithm for OuMv with polynomial preprocessing time and computation time $O(n^{3-\epsilon})$ for $\epsilon > 0$, contradicting the OMv conjecture.

4.2 Conditional lower bounds from the k -Clique hypothesis

4.2.1 Statement of results

We prove conditional lower bounds from the k -Clique hypothesis for dynamic emulator maintenance in the fully dynamic, incremental, and decremental settings. We prove the following theorem for the fully dynamic setting, which also extends to the incremental and decremental settings but only for worst-case update times.

Theorem 4.6. *Under the Combinatorial k -Clique hypothesis, for every constant $\epsilon > 0$ and every constant integer $\ell \geq 1$, for an n -vertex fully dynamic graph with at most $m = \Theta(n^{1+1/\ell})$ edges at all times, there is no combinatorial algorithm for maintaining a $(1 + \alpha, \beta)$ -emulator for any $\alpha \in [0, \frac{2}{2\ell+1})$ and integer $\beta \geq 0$ with $O(m^{1-\epsilon}(\frac{2-(2\ell+1)\alpha}{\beta}))$ edges, preprocessing time $O(mn^{1-\epsilon}(\frac{2-(2\ell+1)\alpha}{\beta})^2)$, and amortized update time $O(m^{1-\epsilon}(\frac{2-(2\ell+1)\alpha}{\beta})^2)$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

In particular, for the natural setting where α is constant and $\beta = n^{o(1)}$ we have the following corollary:

Corollary 4.7. *Under the Combinatorial k -Clique hypothesis, for every constant $\epsilon > 0$ and every constant integer $\ell \geq 1$, for an n -vertex fully dynamic graph with at most $m = \Theta(n^{1+1/\ell})$ edges at all times, there is no combinatorial algorithm for maintaining a $(1 + \alpha, n^{o(1)})$ -emulator for any constant $\alpha \in [0, \frac{2}{2\ell+1})$ with $O(m^{1-\epsilon})$ edges, preprocessing time $O(mn^{1-\epsilon})$, and amortized update time $O(m^{1-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

When $k = 3$ (i.e. triangle detection), our construction also implies a conditional lower bound under the hypothesis that triangle detection cannot be done in time $O(n^{\omega-\epsilon})$ for any constant ϵ even for non-combinatorial algorithms:

Theorem 4.8. *Under the hypothesis that there is no algorithm for triangle detection in $O(n^{\omega-\delta})$ for any constant δ , for every constant $\epsilon > 0$ for an n -vertex fully dynamic graph with at most m edges at all times, there is no algorithm for maintaining a $(1 + \alpha, \beta)$ -emulator for any $\alpha \in [0, 2/3)$ and integer $\beta \geq 0$ with $O(n^{\omega-1-\epsilon}(\frac{2-3\alpha}{\beta})^{\omega-1})$ edges, preprocessing time $O(n^{\omega-\epsilon}(\frac{2-3\alpha}{\beta})^\omega)$, and amortized update time $O(n^{\omega-1-\epsilon}(\frac{2-3\alpha}{\beta})^\omega)$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

In particular, for the natural setting where α is constant and $\beta = n^{o(1)}$ we have the following corollary:

Corollary 4.9. *Under the hypothesis that there is no algorithm for triangle detection in $O(n^{\omega-\delta})$ for any constant δ , for every constant $\epsilon > 0$ for an n -vertex fully dynamic graph with at most m edges at all times, there is no algorithm for maintaining a $(1 + \alpha, n^{o(1)})$ -emulator for any constant $\alpha \in [0, 2/3)$ with $O(n^{\omega-1-\epsilon})$ edges, preprocessing time $O(n^{\omega-\epsilon})$, and amortized update time $O(n^{\omega-1-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

The same result holds for incremental and decremental algorithms, but for worst-case update time.

For the incremental and decremental settings, we prove the following theorem.

Theorem 4.10. *Under the Combinatorial k -Clique hypothesis, for every constant $\epsilon > 0$ and every constant integer $\ell \geq 1$, for an n -vertex incremental or decremental graph with $m = \Theta(n^{1+1/\ell})$ edge insertions or deletions, there is no combinatorial algorithm for maintaining a β -additive emulator for any integer $\beta \geq 0$ with $O(m^{1-\epsilon}/\beta)$ edges and total time $O(mn^{1-\epsilon}/\beta^2)$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

In particular, for the natural setting where $\beta = n^{o(1)}$ we have the following corollary:

Corollary 4.11. *Under the Combinatorial k -Clique hypothesis, for every constant $\epsilon > 0$ and every constant integer $\ell \geq 1$, for an n -vertex incremental or decremental graph with $m = \Theta(n^{1+1/\ell})$ edge insertions or deletions, there is no combinatorial algorithm for maintaining a $n^{o(1)}$ -additive emulator with $O(m^{1-\epsilon})$ edges and total time $O(mn^{1-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

4.2.2 Preliminaries

Our reductions are from the k -Cycle problem. A reduction from k -clique to k -cycle in graphs of all sparsities is known:

Theorem 4.12 (Combinatorial k -Cycle [68]). *Under the Combinatorial k -Clique hypothesis, for every constant $\epsilon > 0$ and every integer $\ell \geq 1$, there is no combinatorial algorithm for detecting a $k = 2\ell + 1$ -cycle in a directed graph with $m = \Theta(n^{1+1/\ell})$ edges in time $O(mn^{1-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

Given an instance of k -Cycle, we introduce a basic gadget that we will use in all of our constructions.

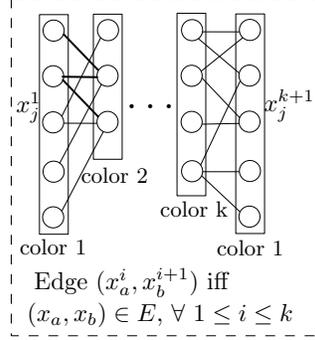


Figure 6: The basic gadget for reductions from k -clique.

The basic gadget Let $G = (V, E)$ with $n = |V|$ and $m = |E|$ be the graph on which we wish to find a directed k -cycle. We use color coding: we color the vertices with colors in $1, 2, \dots, k$ uniformly at random. For all i , let V_i be the set of vertices of color i . We construct a gadget as shown in Figure 6.

The gadget consists of $k + 1$ layers of vertices. For all $1 \leq i \leq k$, layer i contains a copy x_j^i of every $x_j \in V_j$. Layer $k + 1$ contains a copy x_j^{k+1} of every $x_j \in V_1$. For two consecutive layers $i, i + 1$, we include the undirected edge (x_a^i, x_b^{i+1}) if and only if the directed edge (x_a, x_b) is in E .

Our constructions will use several copies of the basic gadget. Each copy uses the same coloring of G so each copy is identical.

4.2.3 Reduction for fully dynamic algorithms

In this section we prove Theorems 4.6 and 4.8.

Construction We say that a k cycle in G is *colorful* if according to the coloring from the basic gadget the k -cycle has exactly one vertex of each color and the vertices are in color order $1, 2, \dots, k$ around the cycle (i.e. the vertex of color 1 is adjacent to the vertex of color k). We will present an algorithm that detects a colorful k -cycle in G if one exists. Any given k -cycle is colorful with probability $1/k^{k-1}$. We repeat the entire algorithm, including construction of the basic gadget, $\Theta(k^{k-1})$ times so that if G contains a k -cycle, then with probability at least $2/3$, for at least one of the repetitions G contains a colorful k -cycle.

We will construct a dynamic graph G' . Let $c = \lceil \frac{\beta}{2-k\alpha} \rceil + 1$ and take c disjoint copies G'_1, \dots, G'_c of the basic gadget. This completes the preprocessing phase.

Now, we start the dynamic phases. There is one dynamic phase for each vertex in G of color 1. In each phase i , we insert an undirected edge between every pair of consecutive gadgets G'_j and G'_{j+1} from $x_i^{k+1} \in G'_j$ $x_i^1 \in G'_{j+1}$. See Figure 7.

Throughout all of the edge updates, we maintain our dynamic emulator. At the end of each phase i , we run a single call to BFS on the emulator to estimate the distance between $x_i^1 \in G'_1$ and $x_i^{k+1} \in G'_c$. If the estimated distance is less than $(k + 2)c - 1$, we return that we have detected a k -cycle.

Following the end of each phase, we remove all of the edges added during that phase.

If after all phases of all $\Theta(k^{k-1})$ repetitions of the algorithm, we have not detected a k -cycle, we return that the graph has no k -cycles.

Correctness First we will show that if the graph G contains a k -cycle then our algorithm detects one. Suppose we are in a repetition of the algorithm where this k -cycle is colorful. Without loss of generality, let x_1, \dots, x_k be the vertices in a k -cycle in G . such that each x_i has color i .

We claim that our algorithm detects this k -cycle at the end of the first phase. The basic construction

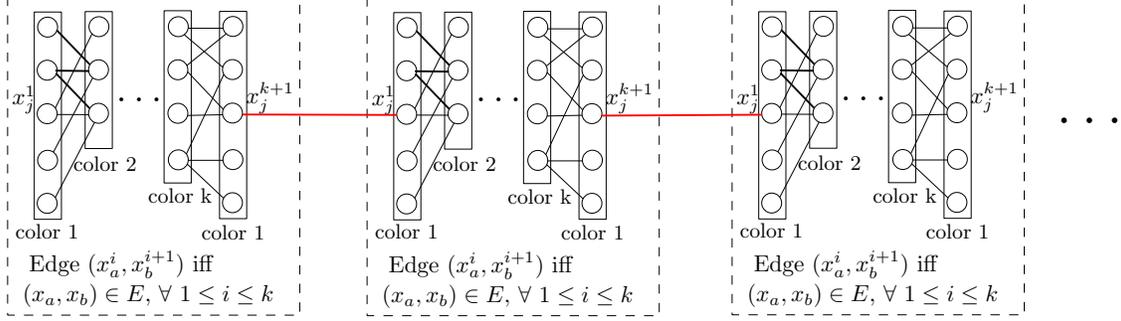


Figure 7: The construction for fully dynamic algorithms. The red edges are dynamically added in phase j .

ensures that for all $1 \leq i \leq k - 1$, in each gadget G'_j , the edge (x_i^i, x_{i+1}^{i+1}) exists and the edge (x_k^k, x_1^{k+1}) exists. Thus, there is a path of length k from x_1^1 to x_1^{k+1} in each basic gadget.

Furthermore, in the first dynamic phase we insert an edge between every two consecutive gadgets G'_j and G'_{j+1} from $x_1^{k+1} \in G'_j$ to $x_1^1 \in G'_{j+1}$. Concatenating these paths, we have a path of length $kc - 1$ from $x_1^1 \in G'_1$ and $x_1^{k+1} \in G'_c$. Thus, the estimate of $\text{dist}(x_1^1 \in G'_1, x_1^{k+1} \in G'_c)$ returned by our $(1 + \alpha, \beta)$ -emulator is at most $(kc - 1)(1 + \alpha) + \beta < (k + 2)c - 1$ by choice of c . Therefore, our algorithm returns that we have detected a k -cycle.

Now we will show that if our algorithm returns that we have detected a k -cycle, then G contains a k -cycle. Let i be the phase that our algorithm detects a k -cycle. Consider G' at the end of phase i . Because our algorithm detected a k -cycle, we know that the distance between $x_i^1 \in G'_1$ and $x_i^{k+1} \in G'_c$ estimated by our emulator is less than $(k + 2)c - 1$. Therefore, the true distance between $x_i^1 \in G'_1$ and $x_i^{k+1} \in G'_c$ is also less than $(k + 2)c - 1$.

We note that since there is only one edge between every pair of adjacent gadgets G'_j , any shortest path between $x_i^1 \in G'_1$ and $x_i^{k+1} \in G'_c$ contains for every edge whose endpoints are in different copies of the basic gadget. That is, this path contains for every $1 \leq j \leq c - 1$ the edge $(x_i^{k+1} \in G'_j, x_i^1 \in G'_{j+1})$. Therefore, for every $1 \leq j \leq c - 1$ this path contains as a subpath a shortest path between $x_i^1 \in G'_j$ and $x_i^{k+1} \in G'_j$. Since G' contains c identical copies of the basic gadget, $\text{dist}(x_i^1 \in G'_j, x_i^{k+1} \in G'_j)$ is the same for all j .

Since $\text{dist}(x_i^1 \in G'_1, x_i^{k+1} \in G'_c) < (k + 2)c - 1$ the edges between gadgets contribute $c - 1$ to this quantity, we know that for all $1 \leq j \leq c$, $\text{dist}(x_i^1 \in G'_j, x_i^{k+1} \in G'_j) < k + 1$. Fix j . Since each basic gadget contains $k + 1$ layers, we know that $\text{dist}(x_i^1 \in G'_j, x_i^{k+1} \in G'_j) \geq k$. Therefore, there is a path from $x_i^1 \in G'_j$ to $x_i^{k+1} \in G'_j$ that contains exactly one vertex from each layer. The construction of the basic gadget ensures that there is an edge $(x_a^i \in G'_j, x_b^{i+1} \in G'_j)$ if and only if the directed edge (x_a, x_b) is in E . Thus, this path from $x_i^1 \in G'_j$ to $x_i^{k+1} \in G'_j$ corresponds to a directed walk of length k in G . In particular the first and last vertex on this walk are both x_i so this is a closed walk. Furthermore, every internal layer of G'_j corresponds to a different color, so every vertex of the closed walk in G has a different color. Thus, every vertex of the closed walk is distinct so the closed walk is indeed a directed k -cycle.

Running time Let n' be the number of vertices in the dynamic graph and let m' be the maximum number of edges ever in the dynamic graph. We first calculate n' and m' . Each basic gadget contains $O(n)$ vertices and $O(m)$ edges. Thus, c copies of the basic gadget contain $n' = O(cn)$ vertices and $O(cm)$ edges. During each of the at most n phases we add $c - 1$ edges so the number of edge updates after preprocessing is $O(cm)$. Thus, $m' = O(cm) = O(cn^{1+1/\ell}) = O(c(n'/c)^{1+1/\ell}) = O(n'^{1+1/\ell})$. We repeat the entire

algorithm $O(k^{k-1}) = O(1)$ times. Thus, the total number of edge updates over the entire sequence is $O(cn)$.

We will now split our running time analysis into two – one for refuting the combinatorial k -Clique conjecture, and one for refuting the hypothesis that triangles cannot be solved faster than matrix multiplication.

Combinatorial k -Clique Let us assume that the dynamic emulator has, for some $\epsilon > 0$, preprocessing time $O(m'n^{1-\epsilon} \left(\frac{2-(2\ell+1)\alpha}{\beta}\right)^2)$, and amortized update time $O(m'^{1-\epsilon} \left(\frac{2-(2\ell+1)\alpha}{\beta}\right)^2)$.

Our dynamic emulator algorithm has preprocessing time $O(m'n^{1-\epsilon} \left(\frac{2-(2\ell+1)\alpha}{\beta}\right)^2) = O((cm)(cn)^{1-\epsilon}/c^2) = O(mn^{1-\epsilon})$ and amortized update time $O(m'^{1-\epsilon} \left(\frac{2-(2\ell+1)\alpha}{\beta}\right)^2) = O((cm)^{1-\epsilon}/c^2) = O(m^{1-\epsilon}/c)$. Since there are $O(cn)$ edge updates, the total running time due to the dynamic emulator algorithm is $O(mn^{1-\epsilon})$.

Additionally, at most n times during each repetition the algorithm, we run BFS on the emulator. The number of edges in the emulator is $O(m'^{1-\epsilon} \left(\frac{2-(2\ell+1)\alpha}{\beta}\right)) = O((cm)^{1-\epsilon}/c) = O(m^{1-\epsilon})$. Thus, the BFS calls take total time $O(nm^{1-\epsilon}) = O(mn^{1-\epsilon})$.

Putting everything together, our combinatorial dynamic emulator algorithm implies a combinatorial algorithm for directed k -cycle detection in time $O(mn^{1-\epsilon})$, thus refuting the combinatorial k -Clique hypothesis.

Non-combinatorial triangle Here we assume that the dynamic emulator has, for some $\epsilon > 0$, preprocessing time $O(n'^{\omega-\epsilon} \left(\frac{2-3\alpha}{\beta}\right)^\omega)$ and amortized update time $O(n'^{\omega-1-\epsilon} \left(\frac{2-3\alpha}{\beta}\right)^\omega)$.

Thus, our dynamic emulator algorithm has preprocessing time $O(n'^{\omega-\epsilon} \left(\frac{2-3\alpha}{\beta}\right)^\omega) = O((cn)^{\omega-\epsilon}/c^\omega) = O(n^{\omega-\epsilon})$ and amortized update time $O(n'^{\omega-1-\epsilon} \left(\frac{2-3\alpha}{\beta}\right)^\omega) = O((cn)^{\omega-1-\epsilon}/c^\omega) = O(n^{\omega-1-\epsilon}/c^{1+\epsilon})$. Since there are $O(cn)$ edge updates, the total running time of the dynamic emulator algorithm is $O(n^{\omega-\epsilon})$.

Additionally, at most n times during each repetition the algorithm, we run BFS on the emulator. The number of edges in the emulator is $O(n'^{\omega-1-\epsilon} \left(\frac{2-3\alpha}{\beta}\right)^{\omega-1}) = O((cn)^{\omega-1-\epsilon}/c^{\omega-1}) = O(n^{\omega-1-\epsilon})$. Thus, the BFS calls take total time $O(n^{\omega-\epsilon})$.

Putting everything together, our dynamic emulator algorithm implies an algorithm for triangle detection in time $O(n^{\omega-\epsilon})$, refuting the hypothesis that triangle detection needs $n^{\omega-o(1)}$ time.

4.2.4 Reduction for incremental and decremental algorithms.

In this section we prove Theorem 4.10.

Construction The construction is similar to the fully dynamic construction, but with different interactions between consecutive gadgets. We first describe the incremental construction.

As in the fully dynamic construction, we say that a k cycle in G is *colorful* if according to the coloring from the basic gadget the k -cycle has exactly one vertex of each color and the vertices are in color order $1, 2, \dots, k$ around the cycle (i.e. the vertex of color 1 is adjacent to the vertex of color k). We will present an algorithm that detects a colorful k -cycle in G if one exists. Any given k -cycle is colorful with probability $1/k^{k-1}$. We repeat the entire algorithm, including construction of the basic gadget, $\Theta(k^{k-1})$ times so that if G contains a k -cycle, then with probability at least $2/3$, for at least one of the repetitions G contains a colorful k -cycle.

We will construct a dynamic graph G' . Starting with an empty graph, we perform edge insertions to construct the following graph. Take $\beta + 1$ disjoint copies $G'_1, \dots, G'_{\beta+1}$ of the basic gadget. Then, we add $\beta + 2$ paths $P_0, \dots, P_{\beta+1}$ each on $2n - 1$ new vertices. Call the vertices of each path $z_1, z_2, \dots, z_n = y_n, y_{n-1}, \dots, y_1$. In other words, the middle node of each path has two names, z_n and y_n .

Now, we start the phases. There is one phase for each vertex in G of color 1. In phase i , for each $1 \leq j \leq \beta + 1$ we insert an edge between $y_i \in P_{j-1}$ and $x_j^1 \in G'_j$. Similarly, for each $1 \leq j \leq \beta + 1$ we insert an edge between $z_i \in P_j$ and $x_j^{k+1} \in G'_j$. See Figure 8.

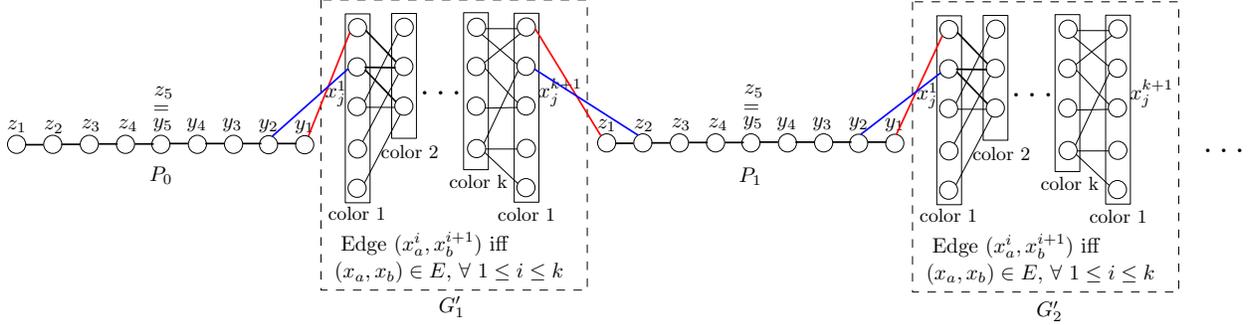


Figure 8: The construction for incremental algorithms. The red edges are dynamically added in phase 1 and the blue edges are dynamically added in phase 2.

Throughout all of the edge updates, we maintain our incremental emulator. At the end of each phase, we run BFS on the emulator to estimate the distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$.

If the estimated distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ at the end of phase i is less than $(k + 3)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$, we return that we have detected a k -cycle.

If after all phases of all $\Theta(k^{k-1})$ repetitions of the algorithm, we have not detected a k -cycle, we return that the graph has no k -cycles.

Now, we describe the decremental construction. The edge updates are exactly the reverse of the incremental construction. That is, the initial graph in the decremental construction is the final graph in the incremental construction. Then, in phase i of the decremental construction, for each $1 \leq j \leq \beta + 1$ we delete the edge between $y_{n-i+1} \in P_{j-1}$ and $x_{n-i+1}^1 \in G'_j$. and delete the edge between $z_{n-i+1} \in P_j$ and $x_{k+1-n-i+1} \in G'_j$.

Throughout all of the edge updates, we maintain our decremental emulator. At the end of each phase, we run BFS on the emulator to estimate the distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$.

If the estimated distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ at the end of phase i is less than $(k + 3)(\beta + 1) + (\beta + 2)(2n - 2(n - i + 1)) + 2((n - i + 1) - 1)$, we return that we have detected a k -cycle. Note that this threshold is exactly the threshold from the incremental algorithm but with i replaced with $n - i + 1$.

If after all phases of all $\Theta(k^{k-1})$ repetitions of the algorithm, we have not detected a k -cycle, we return that the graph has no k -cycles.

Correctness The following argument is written for the incremental setting but the same argument applies for the decremental setting.

First we will show that if the graph G contains a k -cycle then our algorithm detects one. Suppose we are in a repetition of the algorithm where this k -cycle is colorful. Without loss of generality, let x_1, \dots, x_k be the vertices in a k -cycle in G where each x_i is of color i .

We claim that our algorithm detects this k -cycle at the end of the first phase. The basic construction ensures that for all $1 \leq i \leq k - 1$, in each gadget G'_j , the edge (x_i^i, x_{i+1}^{i+1}) exists and the edge (x_k^k, x_1^{k+1}) exists. Thus, there is a path of length k from x_1^1 to x_1^{k+1} in each basic gadget.

Also due to the dynamic edge updates, for each $0 \leq j \leq \beta + 1$, there is an edge between $y_i \in P_{j-1}$ and

x_1^1 and an edge between $z_i \in P_j$ and x_1^{k+1} . Additionally, for each $0 \leq j \leq \beta + 2$, there is a path along P_j from $z_i \in P_j$ to $y_i \in P_j$ of length $(2n - 2) - 2(i - 1) = 2n - 2i$. Finally, there is a path of length $i - 1$ from $z_1 \in P_0$ to $z_i \in P_0$ and a path of length $i - 1$ from $y_i \in P_{\beta+2}$ to $y_1 \in P_{\beta+2}$. Concatenating all of these paths, we have that $\text{dist}(z_1 \in P_0, y_1 \in P_{\beta+1}) \leq (k + 2)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$. Thus, the estimate of $\text{dist}(z_1 \in P_0, y_1 \in P_{\beta+1})$ returned by our β -additive emulator is at most $(k + 2)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1) + \beta < (k + 3)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$.

Now we will show that if our algorithm returns that we have detected a k -cycle, then G contains a k -cycle. Let i be the phase that our algorithm detects a k -cycle. Consider G' at the end of phase i . Because our algorithm detected a k -cycle, we know that the distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ estimated by our emulator is less than $(k + 3)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$. Therefore, the true distance between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ is also less than $(k + 3)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$.

We observe that the layered structure of the graph ensures that every path between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ contains each z_i and y_i in order from P_0 to $P_{\beta+1}$. That is, every shortest path between $z_1 \in P_0$ and $y_1 \in P_{\beta+1}$ is composed of precisely following subpaths:

- A shortest path from $z_1 \in P_0$ to $z_i \in P_0$. The only simple path connecting these vertices is of length $i - 1$.
- A shortest path from $y_i \in P_{\beta+1}$ to $y_1 \in P_{\beta+1}$. The only simple path connecting these vertices is of length $i - 1$.
- A shortest path from $z_i \in P_j$ to $y_i \in P_j$ for all $1 \leq j \leq \beta + 2$. The only simple path connecting these vertices is of length $(2n - 2) - 2(i - 1) = 2n - 2i$.
- A shortest path from $y_i \in P_{j-1}$ to $z_i \in P_j$ for all $1 \leq j \leq \beta + 1$. Since the graph is a series of identical copies of a gadget, we know that $\text{dist}(y_i \in P_{j-1}, z_i \in P_j)$ is the same for all j . Furthermore, we know the length of each of the previous three types of subpaths and we know that $\text{dist}(z_1 \in P_0, y_1 \in P_{\beta+1}) < (k + 3)(\beta + 1) + (\beta + 2)(2n - 2i) + 2(i - 1)$, so we conclude that each $\text{dist}(y_i \in P_{j-1}, z_i \in P_j) < k + 3$.

Due to the layering of the graph, for all $1 \leq p \leq \beta + 1$ the shortest path between $y_i \in P_{p-1}$ and $z_i \in P_p$ must contain vertices $x_j^1 \in G'_p$ and $x_{j'}^{k+1} \in G'_p$ for some j, j' . Fix j, j' . Since $\text{dist}(y_i \in P_{p-1}, z_i \in P_p) < k + 3$, we have that $\text{dist}(x_j^1 \in G'_p, x_{j'}^{k+1} \in G'_p) < k + 1$. Then since each basic gadget contains $k + 1$ layers, we know that there is a path from $x_j^1 \in G'_p$ to $x_{j'}^{k+1} \in G'_p$ that contains exactly one vertex from each layer. The construction of the basic gadget ensures that there is an edge $(x_a^i \in G'_j, x_b^{i+1} \in G'_j)$ if and only if the directed edge (x_a, x_b) is in E . Thus, this path from $x_j^1 \in G'_p$ to $x_{j'}^{k+1} \in G'_p$ corresponds to a directed walk of length k in G . In particular the first and last vertex on this walk are both x_i so this is a closed walk. Furthermore, every internal layer of G'_j corresponds to a different color, so every vertex of the closed walk in G has a different color. Thus, every vertex of the closed walk is distinct so the closed walk is indeed a directed k -cycle.

Running time Let n' be the number of vertices in G' and let m' be number of edge insertions or deletions. We first calculate n' and m' . Each basic gadget contains $O(n)$ vertices and $O(m)$ edges. Thus, $\beta + 1$ copies of the basic gadget contain $O(\beta n)$ vertices and $O(\beta m)$ edges. Additionally, we have $(\beta + 2)$ paths on $(2n - 1)$ vertices each, for a total of $(2n - 1)(\beta + 2)$ additional vertices. Thus, $n' = O(\beta n)$. During each of the at most n phases there are at most $2(\beta + 1)$ edge updates. Thus, $m' = O(\beta m)$. We repeat the entire algorithm $O(k^{k-1}) = O(1)$ times. Thus, the total number of edge updates over all repetitions of the algorithm is $O(\beta m)$.

Let's assume that the incremental or decremental emulator algorithm has total time $O(m'n^{1-\epsilon}/\beta^2) = O(\beta m(\beta n)^{1-\epsilon}/\beta^2) = O(mn^{1-\epsilon})$.

Additionally, at most n times during the algorithm, we run BFS on the emulator. The number of edges in the emulator is $O(m^{1-\epsilon}/\beta) = O((\beta m)^{1-\epsilon}/\beta) = O(m^{1-\epsilon})$. Thus, the BFS calls take total time $O(nm^{1-\epsilon}) = O(mn^{1-\epsilon})$.

Putting everything together, our incremental or decremental emulator algorithm implies an algorithm for directed k -cycle detection in time $O(mn^{1-\epsilon})$, thus refuting the combinatorial k -Clique hypothesis.

5 Algebraic All Pairs Shortest Paths with Path Reporting

The main result of this section is a randomized, fully dynamic algorithm that can maintain and query successors for all pairs shortest paths of up to D edges. Our algorithm is an augmentation of the algebraic all pairs shortest distances algorithm of Sankowski [75], who originally posed as an open problem to use his techniques and the construction of 5.6 to actually report paths. We state our new result formally in the following theorem

Theorem 5.1 (Successor Queries and Short Paths). *For any parameters $\kappa \in (0, 1)$, $D < n$, and an unweighted graph $G = (V, E)$ subject to edge insertions and deletions, there is a dynamic, randomized data-structure \mathcal{P}_κ^D that supports the following operations:*

- *Ins/Delete(e) Inserts/Deletes edge $e \in E$ in worst case time $O(Dn^{\omega(1,1,\kappa)-\kappa} + Dn^{1+\kappa})$.*
- *Short Distance/Successor Query(i, j) Returns the distance $d \leq D$ and a successor on any short, shortest $i \rightarrow j$ path in worst case time $O(Dn^\kappa)$ and is correct whp.*
- *Short Path Queries(i, j) Returns a shortest $i \rightarrow j$ path of length $d \leq D$ by repeatedly finding successors in worst case $O(dDn^\kappa)$ time, and is correct whp.*

with pre-processing time $O(Dn^2)$ on empty graphs and $O(Dn^\omega)$ otherwise.

as also presented in 5.14. We believe this result is of independent interest. We can minimize the update time by the choice of $\kappa = \kappa^*$ that balances the exponents in the runtime, that is, κ^* is the solution to

$$\omega(1, 1, \kappa) - \kappa = 1 + \kappa \tag{6}$$

which numerically can be evaluated to $\kappa^* \approx .529$. The corresponding update time is $O(Dn^{1+\kappa^*}) = O(Dn^{1.529})$

An overview of this section is as follows. We begin by listing a small toolkit of sparse matrix facts, to be used in the following subsections. In subsection 5.2, we detail [75]'s original all pairs shortest distances construction, his path encoding lemma 5.6 and the dynamic matrix inverse algorithm 5.7, all central to our augmentations and key to our proof of 5.1. In subsection 5.3, we build on the matrix inverse algorithm to show how to maintain the product of the adjacency matrix and the inverse, and then use these products to prove the theorem above. Finally, in subsection 5.4, we then use this result to construct the first subquadratic time update and path reporting algorithm for fully dynamic, unweighted APSP.

5.1 Preliminaries

Fact 5.2 (Rectangular Matrix Multiplication). *We denote as $O(n^{\omega(a,b,c)})$ the cost of multiplying two rectangular matrices, the first of size $n^a \times n^b$, the second $n^b \times n^c$. The current exponent of square matrix multiplication $\omega(1, 1, 1) \equiv \omega \leq 2.3729$.*

Fact 5.3 (Row Sparse Matrix Multiplication). *The product of two matrices A, B , the first of $\leq n^\delta$ non-zero rows, and the second of $\leq n^\alpha$ non-zero rows, has at most $n^\delta + n^\alpha$ non-zero rows and takes time $O(n^{\omega(1,\alpha,\delta)})$ to compute.*

Fact 5.4 (Row Sparse Matrix Inverse). *The inverse of a matrix A that differs from identity in at most n^δ rows, differs from identity in at most n^δ rows. Moreover, it can be computed in $n^{\omega(1,\delta,\delta)}$ time.*

Fact 5.5 (The Hitting Set Lemma). *Given a set S of n elements, a random sample $H \subset V$ of size $\geq cn \log n/k$ for a given constant $c > 1$ hits every subset of size k of S with high probability.*

5.2 Algebraic All Pairs Distances

In his PhD thesis, Sankowski [75] showed the following lemma on how to encode path lengths in a matrix.

Lemma 5.6 ([75]). *Let \tilde{A} be the symbolic adjacency matrix of the graph $G = (V, E)$, where each edge $(i, j) \in E$ defines a variable $\tilde{A}_{ij} = x_{ij}$, and $\tilde{A}_{ij} = 0$ if $(i, j) \notin E$. Consider the adjoint $\text{adj}(I - u\tilde{A})_{ij}$ as a polynomial over an additional variable u . The length of the shortest path in G from i to j is equal to the degree of u of the smallest degree non-zero term in $\text{adj}(I - u\tilde{A})_{ij}$.*

Sankowski [75] used this result to construct algorithms for the All Pairs Shortest Distances problem by sampling a uniformly random integer in a field \mathbb{F} of size $n^{O(1)}$ for each symbolic edge-variable and using the Schwartz-Zippel Lemma to guarantee that with high probability, for all $i, j \in V$ the degree d_{ij} (the distance) term in $\text{adj}(I - u\tilde{A})_{ij}$ is non-zero - and thus whp it suffices to read the polynomial entry at i, j to obtain the distance d_{ij} . Sankowski [75] then showed how to maintain and query this adjoint dynamically, and over a ring mod u^{D+1} in worst case update time $O(D(n^{\omega(1,1,\kappa)-\kappa} + n^{1+\kappa}))$ for a given parameter $\kappa \in (0, 1)$. This effectively allowed the *short distance queries*, that is, to return the distance between any pair of vertices i, j correctly if the distance d_{ij} is less than D , however introducing a tradeoff in runtime for large D . We state and re-prove these theorems here, for concreteness.

Theorem 5.7 (Dynamic Integer Matrix Inverse [75]). *Given a constant $\kappa \in (0, 1)$, there is a deterministic, dynamic data-structure \mathcal{D}_κ that maintains the inverse and the determinant of an integer matrix M subject to non-singular entry-wise updates, and supports the following operations:*

- *Update(i, j, v)* Updates entry $M_{ij} = v$ and the data-structure \mathcal{D}_κ in time $O(n^{\omega(1,1,\kappa)-\kappa} + n^{1+\kappa})$
- *Query(i, j)* Returns the value of the inverse at entry (i, j) , M_{ij}^{-1} , in $O(n^\kappa)$ time.

For current $\omega < 2.3729$, the value of κ that balances the update time is $\kappa^* \approx .529$.

The key idea is to write the inverse as $M^{-1} = T(\mathbb{1} + N)$ as the product of two matrices, where one of them, N , is *sparse* and initially null. As we will show, each entry update to M corresponds to a row update to N (without updating T), and every $O(n^\kappa)$ updates we exploit fast matrix multiplication over this sparse N to "reset" $T \leftarrow T + NT, N \leftarrow 0$, guaranteeing the sparsity of N .

Proof. Let us denote $v_{ij}e_i e_j^T$ as the matrix corresponding to the *additive update* to entry ij of M during some update. That is, $M' \leftarrow M + v_{ij}e_i e_j^T$. We construct \mathcal{D}_κ as follows. We maintain explicitly the matrices T, N during the execution, where initially $N = 0$ and we precompute $T = M^{-1}$ in $O(n^\omega)$ time. At every non-singular entry update, we first compute the update to N through the following algorithm:

- Compute the row vector $b = v_{ij}e_j^T M^{-1} = v_{ij}(e_j^T T)(\mathbb{1} + N)$
- Compute the inverse $B = (\mathbb{1} + e_i b)^{-1}$
- Finally, update

$$N' \leftarrow NB + B - \mathbb{1} \quad (7)$$

Correctness of this update follows from plugging in the result into the definition of M^{-1} to guarantee the invariant

$$T(\mathbb{1} + N') = T(NB + B) = M^{-1}B = (M + e_i b M)^{-1} = (M + e_i v_{ij} e_j^T)^{-1} = (M')^{-1} \quad (8)$$

as intended. Note that by Fact 5.4, as $(\mathbb{1} + e_i b)$ has exactly one non-identity row, so does its inverse $B = (\mathbb{1} + e_i b)^{-1}$. By Fact 5.3 this implies that at each update N gains at most one additional non-zero row, and thereby *after k updates N has at most k non-zero rows*. Finally, after $O(n^\kappa)$ updates, we reset the inverse by computing

- After every n^κ updates, $T \leftarrow T(1 + N)$, $N \leftarrow 0$

Since N has at most $O(n^\kappa)$ non-zero rows, computing this product takes time $O(n^{\omega(1,1,\kappa)})$. Note that it takes time $O(n^{1+\kappa})$ to compute the row vector b given the product of the row vector $e_j^T T$ and the row-sparse N , s.t. the average update time is

$$O(n^{1+\kappa} + n^{\omega(1,1,\kappa)-\kappa}) \quad (9)$$

which for current $\omega \approx 2.3729$ is minimized at $\kappa \approx .529$, for a runtime of $O(n^{1.529})$. Finally, to query any entry ij of M^{-1} , we compute the dot product of the row vector $(e_i^T T)$ and the sparse column vector $(N e_j)$ in time $O(n^\kappa)$:

$$e_i^T M^{-1} e_j = e_i^T T(\mathbb{1} + N) e_j = e_i^T T e_j + (e_i^T T)(N e_j) \quad (10)$$

To conclude this proof, we note that the determinant $\det M$ is easily maintained in $O(1)$ by definition of the matrix B above. At an update to entry (i, j) of M , $M' = B^{-1}M \Rightarrow \det M' = \det M \cdot (1 + b_i)$. This allows us to support queries to the adjoint of M . \square

Sankowski then showed how to support the above dynamic matrix inverse algorithm over polynomial matrices. We present the result in the corollary below

Corollary 5.8 (Dynamic Polynomial Matrix Inverse [75]). *Given $\kappa \in (0, 1)$, $D \in [n]$ there is a dynamic data-structure \mathcal{D}_κ^D that can maintain the inverse M^{-1} of polynomial matrix $M = \mathbb{1} - A$ where $M, A \in (\mathbb{F}[X]/\langle X^{D+1} \rangle)^{n \times n}$ subject to entry-wise polynomial updates to M , incurring a multiplicative cost of $\tilde{O}(D)$ to the runtimes of the theorem above.*

The proof of the corollary above arises from Sankowski's extension of Strassen's idea of computing over the formal power series to the dynamic case. We note that over the formal power series mod X^D , $M = \mathbb{1} - A \in (\mathbb{F}[X]/\langle X^{k+1} \rangle)^{n \times n}$ is always invertible as

$$M^{-1} = \frac{1}{\mathbb{1} - A} = \sum_{i=0}^k A^i = \prod_{i=1}^{\lceil \log D \rceil} (\mathbb{1} + A^{2^i}) \quad (11)$$

and moreover can be computed/preprocessed in $\tilde{O}(Dn^\omega)$. The details are in [75]. Corollary 5.8 above, together with the path encoding lemma 5.6 define the fully dynamic data-structure \mathcal{D}_κ^D that can support edge updates and short distance queries, which we base this section off of and later augment. For concreteness, we state this result in the following theorem:

Theorem 5.9 (Short Distance Queries [75]). *Given a unweighted, dynamic graph G subject to edge insertions and deletions, and parameters $\kappa \in (0, 1)$ and $1 < D < n$, there exists a dynamic data-structure \mathcal{D}_κ^D that supports edge updates in worst case $O(D(n^{\omega(1,1,\kappa)-\kappa} + n^{1+\kappa}))$ time and can query any distance $d \leq D$ correctly whp in worst case time $O(Dn^\kappa)$. If otherwise $d > D$, the distance query outputs a failure.*

Which follows simply by defining $M = \mathbb{I} - uA$ in Corollary 5.8, where A is the symbolic adjacency matrix with random integer values whenever non-zero. Additionally, \mathcal{D}_κ^D can also maintain the $H \times H$ submatrix of the inverse, corresponding to the up to D distances between all pairs of vertices in a subset of size $|H| = n^q$ vertices in $O(Dn^{2q})$ time, as follows:

Corollary 5.10 (Submatrix Maintenance [75]). *Given a subset $H \subset [n]$ of size n^q of the column/row indices, \mathcal{D}_κ^D can maintain the $H \times H$ submatrix $M_{H,H}^{-1}$ of the inverse in additional $O(Dn^{2q})$ update time, allowing queries in $O(1)$ to the submatrix.*

Proof. To explicitly maintain $M_{H,H}^{-1}$, it suffices to show how to efficiently perform updates. In the notation of the proof of the theorem above,

$$M_{H,H}^{-1} \leftarrow (M^{-1}B)_{H,H} = M_{H,H}^{-1} + M_{H,\{i\}}^{-1}(B - \mathbb{I})_{\{i\},H} \quad (12)$$

where the update to the inverse is expressed in terms of the matrix B , which only differs from identity at a single row i , and thereby can be quickly computed in $O(Dn^{2q})$. \square

5.3 Successor and Short Path Queries

We dedicate this section to the proof of Theorem 5.1. The key idea in maintaining and querying the successors in short, shortest paths is inspired by Seidel's static algorithm for undirected, unweighted APSP [77]. In order to augment the data-structure of Theorem 5.9 [75], we additionally maintain the product $A \cdot \text{adj}M$ of the boolean adjacency matrix and the adjoint of the polynomial matrix $M = \mathbb{I} - uA$ as defined in Lemma 5.6. The (i, j) entry of the adjoint $(\text{adj}M)_{ij}$ is a polynomial where the degree of the lowest degree non-zero monomial is the length of the shortest path, and if the distance from i to $j \in V$ is $d \leq D$, then we can query it in time $O(Dn^\kappa)$ under Theorem 5.9. Moreover, inspection of the product tells us that $(A \cdot \text{adj}M)_{ij}$ must have smallest degree $d - 1$ as there must exist some witness s (the successor!) s.t. $A_{is} = 1$ and M_{sj}^{-1} has minimum degree $d - 1$. This idea of maintaining successors as witnesses of a polynomial matrix product in addition to a sparsification argument that allows us to reduce the case of a multiple witnesses (successors) to that of a single witness (successor), enables a bitwise selection trick that follows closely to Seidel's successor finding algorithm [77] in the static case.

The outline of this section is as follows. First, we describe how to maintain the product $A \cdot \text{adj}M$. Note again that $A \cdot \text{adj}M = \det M \times A \cdot M^{-1}$, and thus we can instead maintain $A \cdot M^{-1}$, and $\det M$ following 5.7. Then, we show how to find a successor on a given path $i \rightarrow j$ of length $\leq D$, if the successor is unique, using $\log n$ queries to the products described. Finally, we review Seidel's sparsification trick [77] to reduce the case of multiple witnesses to a polylog number of single witness queries, and we conclude with our main theorem on the short path finding algorithm \mathcal{P}_κ^D . In the next subsection, we use this path finding black box as a subroutine to construct novel algorithms for exact dynamic APSP with path reporting. We begin by presenting our theorem on maintaining the products.

Theorem 5.11 (Dynamic Product Maintenance). *Let $\kappa \in (0, 1)$ be a constant, and define the dynamic polynomial matrices $E, A, M \in (\mathbb{F}[X]/\langle X^{D+1} \rangle)^{n \times n}$, where $M \equiv \mathbb{I} - A$, and both E and A are subject to entry-wise updates. There is a data-structure that supports the following operations over the product $E \cdot M^{-1}$:*

- *UpdateE*(i, j, v) Updates the entry $E_{ij} \leftarrow v$ for $i, j \in [n]$ and $v \in (\mathbb{F}[X]/\langle X^{D+1} \rangle)$ in worst case time $O(Dn)$.
- *UpdateA*(i, j, v) Updates the entry $A_{ij} \leftarrow v$ for $i, j \in [n]$ and $v \in (\mathbb{F}[X]/\langle X^{D+1} \rangle)$ in worst case time $O(D(n^{\omega(1,1,\kappa)-\kappa} + n^{1+\kappa}))$
- *Query*(i, j) Returns the entry $(E \cdot M^{-1})_{ij} \in (\mathbb{F}[X]/\langle X^{D+1} \rangle)$ in time $O(Dn^\kappa)$.

Proof. The key idea in this proof is to use Theorem 5.7 and maintain matrices T, N s.t. $M^{-1} = T(\mathbb{1} + N)$, as well as explicitly maintaining a matrix $V \equiv ET$. Using associativity this avoids an intermediary matrix product during queries, and as we will show can be efficiently maintained subject to updates to E and A under slight modifications to the original scheme.

During the execution we maintain explicitly the polynomial matrices E, T, N and the product $V = ET$, such that we maintain the correctness invariant $EM^{-1} = (ET)(\mathbb{1} + N) = V(\mathbb{1} + N)$ by associativity. At every entry update to $E \leftarrow E + v_{ij}e_i e_j^T$, we simply compute the row vector $v_{ij}e_i e_j^T T$, corresponding to the update to the i th row $V \leftarrow V + e_i(v_{ij}e_j^T T)$ in $O(Dn)$ time. Correctness of this update follows from associativity in the expansion

$$(E + v_{ij}e_i e_j^T)M^{-1} = (ET + v_{ij}e_i e_j^T T)(\mathbb{1} + N) = (V + e_i(v_{ij}e_j^T T))(\mathbb{1} + N) \quad (13)$$

showing that it suffices to update the product V . Next, whenever we update $A \leftarrow A + v_{ij}e_i e_j^T$, we simply follow Theorem 5.7 and update the sparse matrix N , and every n^κ updates to A we *reset* $V \leftarrow V \cdot (1 + N)$, then $T \leftarrow T \cdot (1 + N)$, finally $N \leftarrow 0$, in $O(Dn^{\omega(1,1,\kappa)-\kappa} + Dn^{1+\kappa})$ amortized time. Note that this maintains the invariant that $EM^{-1} = V(\mathbb{1} + N)$ as well as $M^{-1} = T(\mathbb{1} + N)$. To conclude this proof, we support queries to $(EM^{-1})_{ij}$ by computing

$$e_i^T EM^{-1} e_j = e_i^T V e_j + (e_i^T V)(N e_j) \quad (14)$$

in $\tilde{O}(Dn^\kappa)$ time due to the dot product of the vectors $(e_i^T V)$ and $(N e_j)$, where $(N e_j)$ is a column vector with a sparse number of non-zero rows, as detailed in the proof of 5.7 in the previous subsection. Note that maintaining V explicitly was key to these fast queries. \square

Effectively, this implies that we can efficiently maintain and query the product above $E \cdot M^{-1}$, for any matrix E , in the same time bounds as Sankowski's original dynamic polynomial matrix inverse construction \mathcal{D}_κ^D of Corollary 5.8. We now show how to use this construction to find single witnesses/single successors.

Lemma 5.12. *Let $i, j \in V$ be vertices s.t. their distance is $1 < d_{ij} \leq D$. If there is a single successor $s \in V = [n]$ on any shortest path from i to j , then we can find it using $O(\log n)$ queries to the data-structure of the theorem above, in $\tilde{O}(Dn^\kappa)$ time.*

Proof. For each $l \in [O(\log n)]$ we define a subset of the column indices $S_l \subseteq [n]$ where each index $s \in [n]$ is in S_l if the l th bit of s 's bitwise description is $s_l = 1$. Define $O(\log n)$ copies of the adjacency matrix A , where $A^{(l)}$ for $l \in [O(\log n)]$ is defined by only picking the subset of columns S_l , and all other columns of $A^{(l)}$ are completely null. We initialize and maintain $O(\log n)$ copies of the data-structure of Theorem 5.11, where $E^{(l)} = A^{(l)}$. As described, for a given pair of vertices $i, j \in V$ if $1 < d_{ij} \leq D$ then there must exist some successor $s \in [n]$ of i s.t. $A_{is} = 1$, $\text{adj}M_{sj}$ has minimum degree $d_{sj} = d_{ij} - 1$, and thereby the product $(A \cdot \text{adj}M)_{ij}$ has minimum degree $d_{sj} = d_{ij} - 1$. However, as by assumption there

is a single witness s , all other $s' \in [n] \setminus \{s\}$ have minimum degree of $\text{adj}M_{s'j} > d_{sj}$, and thus any copy of the data-structure $(A^{(l)} \cdot \text{adj}M)_{ij}$ where s is not one of the selected columns will have minimum degree $> d_{ij} - 1$. This is simply since $A_{is}^{(l)} = 0$ if the l th bit $s_l = 0$. It follows that the minimum degree of the l th product at entry (i, j) is $d_{ij} - 1$ if the l th bit of s is 1, and so if we query the $O(\log n)$ data-structures $(A^{(l)} \cdot \text{adj}M)_{ij}$ for each l , and define a sequence of bits $b_l = 1$ if $(A^{(l)} \cdot \text{adj}M)_{ij}$ has minimum degree $d_{ij} - 1$, then $(b_1, b_2 \cdots b_{O(\log n)}) \equiv s$ is exactly the bitwise description of the single witness s . \square

A straightforward sparsification trick allows us to reduce the case of multiple witnesses to that of a single witness.

Lemma 5.13. *If there is an arbitrary number of distinct successors of $i \in V$ on any shortest path from i to a given $j \in V$ of length $\leq D$, then we can find an arbitrary one of them in $\tilde{O}(Dn^\kappa)$ time.*

Proof. The key construction is to maintain $O(\log^2 n)$ versions of the single-witness data-structure of Lemma 5.12. We first address the case of a known number of witnesses $c \geq 1$ for a given product ij via a sparsification trick, and then show how to efficiently "guess" the number of witnesses for each product in only an additional factor of $\log n$.

If the number of witnesses c of the product $(AM^{-1})_{ij}$ is known to be bounded in a range by $n/2^{w+1} \leq c \leq n/2^w$ for some $w = O(\log n)$, then sampling 2^w columns $C \subset [n]$ uniformly at random gives us a constant probability that we sample only a single witness $s \in C$. This is since

$$\mathbb{P}[|W \cap C| = 1] = 2^w \frac{c}{n} \left(1 - \frac{c}{n}\right)^{2^w - 1} \geq \frac{1}{2} \left(1 - \frac{1}{2^w}\right)^{2^w - 1} \geq 1/2e \quad (15)$$

where W is the set of witnesses of the product ij . It follows that if we maintain $O(\log n)$ versions of the data-structure of lemma 5.12, where in each copy we sample 2^w columns uniformly at random and 0-out the remaining columns, then with high probability for every pair i, j with $n/2^{w+1} \leq c \leq n/2^w$ successors there exists a copy of the data-structure with a single witness of i on a shortest path to j .

Finally, to address the fact that the number of witnesses is unknown and can vary for each pair, we maintain $O(\log n)$ editions of the sparsification above for every $w \in \{1 \cdots O(\log n)\}$, that is, for every possible bound over the number of witnesses. In this manner, whenever we query the successor for a given pair i, j , we query each of the $O(\log^3 n)$ data-structures of theorem 5.11 for the entry i, j , and obtain a set P of $O(\log^2 n)$ potential successors ($\log n$ column samplings per witness-exponent guess). Then, we query the original distance matrix $\text{adj}M_{sj}$ for each $s \in P$, and output any one of them with distance $d_{sj} = d_{ij} - 1$ (a successor) and with an edge $A_{is} = 1$. By construction, we find a successor whp as we check every possible range of the number of successors (parametrized by w), and we always verify the output solution. This takes overall $\tilde{O}(Dn^\kappa)$ to find a successor on any shortest path i, j of length $\leq D$. \square

Now that we have shown how to support successor queries, we formalize our statement on short path finding.

Theorem 5.14 (Successor Queries and Short Paths). *For any parameters $\kappa \in (0, 1), D < n$, and an unweighted graph $G = (V, E)$ subject to edge insertions and deletions, there is a dynamic, randomized data-structure \mathcal{P}_κ^D that supports the following operations:*

- *Ins/Delete(e) Inserts/Deletes edge $e \in E$ in worst case time $O(Dn^{\omega(1,1,\kappa)-\kappa} + Dn^{1+\kappa})$.*
- *Short Distance/Successor Query(i, j) Returns the distance $d \leq D$ and a successor on any short, shortest $i \rightarrow j$ path in worst case time $O(Dn^\kappa)$ and is correct whp.*

- *Short Path Queries*(i, j) Returns a shortest $i \rightarrow j$ path of length $d \leq D$ by repeatedly finding successors in worst case $O(dDn^\kappa)$ time, and is correct whp.

with pre-processing time $O(Dn^2)$ on empty graphs and $O(Dn^\omega)$ otherwise.

Proof. We maintain the polynomial matrix inverse data-structure \mathcal{D}_κ^D of Theorem 5.9, in addition to the data-structure of Lemma 5.13. At each edge update, we update each of the polylog copies of the data-structure of Theorem 5.11, including at most a single entry update to E and to A per copy, per update. To support Short Path Queries of a given pair of vertices i, j of distance $d \leq D$, we simply sequentially query the successor s of i , and recurse on $i \leftarrow s$ while maintaining the path. This takes time $\tilde{O}(d \times Dn^\kappa)$. \square

In the next sections, we apply \mathcal{P}_κ^D as a black box to the problems of Exact and Approximate Dynamic APSP.

5.4 Exact Fully Dynamic Unweighted APSP

In this section we explain how to use our short path finding black box \mathcal{P}_κ^D of Theorem 5.1 to construct the first subquadratic time update and path query algorithm for Exact Dynamic APSP. \mathcal{P}_κ^D allows us to query short length $d \leq D$ paths in $O(dDn^\kappa)$ time, such that it only remains to show how to construct long, up to linear length paths. The construction follows that of the all pairs shortest distances algorithm of [75]. We can use the *path decomposition technique* to sample a hitting set $H \subset V$ of size $\tilde{O}(n/D)$ that whp hits every length $D/2$ path in G by Fact 5.5. In this manner, any i to j shortest path can be decomposed into consecutive subpaths $l_1 \cdots l_k \cdots l_{O(n/D)}$ of $D/2$ vertices, s.t. whp there is a vertex $u_k \in H$ in the hitting set in each one of those subpaths $u_k \in l_k \forall k$. It follows that there exists a subpath decomposition $i \rightarrow u_1, u_1 \rightarrow u_2 \cdots \rightarrow j$ of the shortest path from i to j where whp the distance between each adjacent pair (u_k, u_{k+1}) is "short", that is, less than D , and can be queried efficiently.

This effectively reduces the question to finding the hitting set vertices in a path from arbitrary i to j . At every edge update, we use Corollary 5.10 to explicitly maintain the $H \times H$ submatrix of the inverse in $O(n^2/D)$ time, s.t. we can access the $H \times H$ matrix of distances $D_{H,H}^D$ between any pair of vertices in H if they are $\leq D$ apart in $O(1)$ time. We can use these current distances to construct the induced, weighted graph $G' = (H, E', D_{H,H}^D)$, and compute the APSP via Floyd-Warshall's algorithm on the induced G' in $O(|H|^3) = O(n^3/D^3)$ time. The key idea in constructing long paths is that the hitting set lemma guarantees that on any long shortest path, each pair of hitting set vertices on said path are at most D nodes away whp and thus between vertices in H the shortest distance in G' is also the shortest distance in G . Floyd-Warshall's algorithm therefore not only with high probability gives us the correct distances $D_{H,H}$ between any $p, q \in H$, as it also maintains the paths $p, u_1 \cdots, u_k, \cdots q$ of hitting set vertices in G' between any $p, q \in H$.

At each path query between arbitrary $i, j \in V$, we first check if the distance is short $d_{ij} \leq D$ using \mathcal{P}_κ^D of theorem 5.1, and if not, we compute the distance between i and j by performing the following minimization over the last and first hitting set vertices p, q in the path $i \rightarrow p \cdots q \rightarrow j$:

$$D_{ij} = \min\{D_{ij}^D, \min_{p,q \in H} (D_{ip}^D + D_{H,H}(p, q) + D_{qj}^D)\} \quad (16)$$

where the terms in the inner minimization correspond to the distances $i \rightarrow p \in G$, $p \rightsquigarrow q \in G'$ and $q \rightarrow j \in G$ respectively. We note that it takes time $O(|H|^2) = \tilde{O}(n^2/D^2)$ to perform the minimization, which requires both a row and a column query to \mathcal{D}_κ^D , in particular to the distances (i, H) and (H, j) , s.t. the total runtime of the distance query is $O(|H|^2 + |H|Dn^\kappa) = \tilde{O}(n^2/D^2 + n^{1+\kappa})$.

Note additionally that finding the first and last hitting set vertices p, q in a i, j path also gives us *all* the hitting set vertices $p, u_1, u_2 \dots q \in H$ on the path, due to the successor information in G' . It follows now that whp each pair of adjacent u_k, u_{k+1} are less than $2 \times D/2 = D$ nodes apart, and thus we can perform short path queries between all the adjacent pairs $u_k, u_{k+1} \in H$ using \mathcal{P}_κ^D in $\tilde{O}(Dn^\kappa \sum_k d_{u_k, u_{k+1}}) = \tilde{O}(Dn^{1+\kappa})$ time; this is equivalent to the runtime of $O(n)$ successor queries. Overall, the long path query takes time

$$\tilde{O}\left(Dn^{1+\kappa} + n^2/D^2\right) \quad (17)$$

Putting all the ingredients together gives us a worst case update time of

$$O(Dn^{\omega(1,1,\kappa)-\kappa} + Dn^{1+\kappa} + n^3/D^3 + n^2/D) \quad (18)$$

due to the update time of theorem 5.1, the Floyd-Warshall step, and corollary 5.10 respectively. Finally, the initialization time is dominated by that of initializing Theorem 5.1. We conclude this description with a formal statement on the main result of this subsection:

Theorem 5.15 (Subquadratic, Unweighted APSP). *For any $\kappa \in (0, 1)$, $D \in [n]$, and an unweighted graph $G = (V, E)$ subject to edge insertions and deletions, there is a dynamic, randomized data-structure that supports the following operations:*

- *Ins/Del(e) Inserts/Deletes edge $e \in E$ in worst case time*

$$O(Dn^{\omega(1,1,\kappa)-\kappa} + Dn^{1+\kappa} + n^3/D^3 + n^2/D) \quad (19)$$

- *Distance Query(i, j) Returns the distance from i to j whp, in worst case time*

$$O(n^2/D^2 + n^{1+\kappa}) \quad (20)$$

- *Path Query(i, j) Returns a shortest path from i to j whp, in worst case time*

$$O(n^2/D^2 + Dn^{1+\kappa}) \quad (21)$$

The values of D and κ above establish a trade-off between update and path query runtimes, and we allot the remainder of this section for a discussion on the sub and super quadratic algorithms we can construct with this trade-off. Three cases worth mentioning are that of minimal update time, minimal path query time, and minimal path query subject to subquadratic update time.

Minimal Update Time If we optimize first over κ in the update time, picking the optimum $\kappa = \kappa^* \approx .529$, and then optimizing over D s.t. $D = n^\mu$ and $\mu = \frac{2-\kappa^*}{4} \approx .368$, then we obtain the first fully dynamic exact APSP algorithm with subquadratic time update and path query, in $O(n^{1.897})$ time for both operations.

Minimal Path Query Time we can establish super quadratic update time but efficient path query algorithms by first optimizing over D in the path query expression. The path query runtime is optimized if we pick $D = n^{\frac{1-\kappa}{3}}$, s.t. the runtime is $O(n^{\frac{4+2\kappa}{3}})$. The update time for any $0 < \kappa \leq .313$ is now $O(n^2(n^\kappa + n^{\frac{1-4\kappa}{3}}))$, which itself establishes a trade-off over κ , and allows for $O(n^{2.332})$ update time and $O(n^{1.334})$ path query.

Minimal Path Query subject to Subquadratic Updates We can also establish *slightly* subquadratic update time algorithms, but with efficient path query times. It suffices to first pick the minimum D that allows for a subquadratic update time, $D = n^{1/3+\delta/3 \forall \delta > 0}$ s.t. updates are performed in $O(n^{1/3+\delta/3+\omega(1,1,\kappa)-\kappa} + n^{2-\delta})$, and then finding the minimum κ s.t. the updates are still subquadratic. By numerical optimization we find $\kappa' \approx .334$, s.t. the runtime of path query is $O(n^{1.667})$, while the update time is $O(n^{1.999})$.

6 Algebraic fully dynamic spanner algorithm

The goal of this section is to prove the following theorem.

Theorem 6.1. *For any constant $0 < \epsilon \leq 1$, given an undirected, unweighted fully dynamic graph, there is an algorithm to maintain a $(1 + \epsilon, n^{o(1)})$ -spanner of size $n^{1+o(1)}$ with preprocessing time $\tilde{O}(n^\omega)$ (or $\tilde{O}(n^2)$ if the input graph is empty) and worst update time $n^{1+\kappa^*+o(1)} = O(n^{1.529})$ with high probability against an oblivious adversary.*

Internal Data Structures. In order to obtain this result, we use two internal data structures. We first use a data structure \mathcal{A} that maintains the distance matrix and corresponding shortest paths in a graph G , restricted to a set of sources S . The algorithm is randomized and uses fast matrix multiplication internally. It is a corollary of Theorem 5.1, and follows simply by querying the data structure of Theorem 5.1 for every pairwise distance in S .

Corollary 6.2. *Let κ be such that $0 < \kappa \leq \kappa_* < 0.529$ and let D be a distance parameter in $[1, n^\kappa]$. Suppose we are given a set of vertices S . Then the data structure from Theorem 5.1 can maintain for any arbitrarily small $\delta > 0$, for all pairs of nodes $s, t \in S$, a shortest path between s and t as long as $\text{dist}(s, t) \leq D$ (and can check if $\text{dist}(s, t) > D$), with high probability against an oblivious adversary, with initialization time $\tilde{O}(Dn^\omega)$ (for a nonempty initial graph) or $\tilde{O}(n^2)$ (for an empty initial graph) and worst-case update time $\tilde{O}(Dn^{\omega(1,1,\kappa)-\kappa} + |S|^2 D^2 n^\kappa)$.*

If the depth threshold $D = n^{o(1)}$, the size of the subset $|S| = O(\sqrt{n})$ and we pick the parameter $\kappa = \kappa^ \approx .529$ to minimize the update time as in 5.1, the update time becomes $n^{1+\kappa^*+o(1)} = O(n^{1.529})$.*

Further, we use an algorithm \mathcal{B} to maintain a spanner of G with high multiplicative error efficiently.

Theorem 6.3 (see [48], Theorem 1.4). *Given an unweighted, undirected fully dynamic graph, there exists an algorithm \mathcal{B} that maintains a spanner with multiplicative stretch $\log n$ and expected size $O(n \log n)$ that has expected update time $O(\log^3 n)$ against an oblivious adversary.*

The Algorithm. Equipped with these two powerful data structures, let us state the algorithm that gives Theorem 6.1. Let $k = \sqrt{\log n}$, let $\epsilon' = \frac{\epsilon}{20(k+1)}$, and let $b = \left(\frac{\log n}{\epsilon'}\right)$. We sample sets $V = A_0 \supseteq A_1, \dots \supseteq A_k \supseteq A_{k+1} = \emptyset$ where A_i for $i \in [1, k]$ is obtained by sampling each vertex in V with probability $n^{-i/k} \log n$ (and to make the sets nesting add it to all A_j where $j \leq i$). We maintain two data structures dynamically during the sequence of edge insertions and deletions:

1. For $\gamma = \lfloor \kappa \cdot k \rfloor$, we run \mathcal{A} on graph G with fixed source set A_γ and depth threshold $\frac{1}{8 \log n} b^{k+1}$.
2. Further, we run \mathcal{B} on G and let \tilde{G} be the $\log n$ -approximate spanner.

Let $c_{i,j} = \sum_{y=i+1}^j b^y$. Initially, and after an edge update we construct the spanner H from scratch as follows: We say that $a \in A_\ell \setminus A_{\ell+1}$ is *active* if for no $j > \ell$ there exists a vertex $a' \in A_j \setminus A_{j+1}$ with $\text{dist}_{\tilde{G}}(a, a') \leq c_{\ell,j}/4$. Note that we are using distances in \tilde{G} and *not* in G for this definition. Note that all vertices in A_k are active at all times. In order to determine which vertices are active, we run the following process. We begin by labeling all vertices as active. We then run for each i from k down to 0, a BFS algorithm on the spanner \tilde{G} to depth b^i from every vertex in $A_i \setminus A_{i+1}$ that is still labelled active. We then deactivate each vertex v in $V \setminus A_i$ for which we found a vertex $a \in A_i \setminus A_{i+1}$ that is close enough to establish that v cannot be active. We then construct our spanner H which is initially empty by adding edges from two sources:

1. We add all edges from \tilde{G} to H .
2. For any two vertices $a, a' \in A_i \setminus A_{i+1}$ that are active, we add the shortest path $\pi_{a,a',G}$ to H if $\mathbf{dist}_G(a, a') \leq \frac{1}{8 \log n} b^{i+1}$, that is, if their distance in G is small. To compute these paths $\pi_{a,a',G}$, we distinguish two cases. If $i \geq \gamma$, then we pose a path reporting query to \mathcal{A} . For $i < \gamma$, we run from every such active vertex a BFS on G to depth $\frac{1}{8 \log n} b^{i+1}$.

As we will show, the approximation factor of the spanner always holds and the sparsity holds with high probability. If the spanner becomes too dense, we reinitialize the algorithm.

Spanner Approximation. In the following, we prove that H indeed forms a $(1 + \epsilon, n^{o(1)})$ spanner. The basic idea behind the proof is the following: let s and t be two vertices and we want to analyze $\mathbf{dist}_H(s, t)$ in comparison to $\mathbf{dist}_G(s, t)$. There are basically two cases: If s is “close” to an active vertex a in $A_i \setminus A_{i+1}$, i.e. at distance at most d_i for some d_i , and there is a vertex v that is at distance $\sim 4d_i \log n / \epsilon$ from s on the path from s to t , such that v is “close” to an active vertex a' in $A_i \setminus A_{i+1}$, then we can use the spanner \tilde{G} to get from s to a and from a' to t and have that the shortest path between a and a' is in H by part (2) of its construction. It is not hard to see that this detour only implies a $(1 + \epsilon)$ -multiplicative error. Otherwise, we do not have a vertex within distance $\sim 4d_i \log n / \epsilon$ that is close to any active vertex in $A_i \setminus A_{i+1}$. We can then repeat the same argument for level $i - 1$ along the path segment from s of length $\sim 4d_i \log n / \epsilon$ and we are ensured that we eventually reach a level where vertices are active and where we get a good approximation. This allows us to subsume the additive error from higher levels into multiplicative error for a series of segments of lower levels.

To formalize this concept, fix a value $i \in [1, k]$ and let us say a vertex a is $\geq i$ -far if (1) $a \in A_\ell \setminus A_{\ell+1}$ for some $\ell \in [0, i - 1]$ and (2) for no $j \geq i$ there is a vertex $a' \in A_j \setminus A_{j+1}$ with $\mathbf{dist}_{\tilde{G}}(a, a') \leq b^j - \frac{1}{2}c_{\ell, i}$. Observe that the distance requirement is formulated with regard to the multiplicative spanner \tilde{G} . Note that as $A_{k+1} = \emptyset$, every vertex is trivially $\geq (k + 1)$ -far.

Lemma 6.4. *For any $\ell \in [0, k - 1]$ every vertex $a \in A_\ell \setminus A_{\ell+1}$ that is $\geq (\ell + 1)$ -far is active.*

Proof. Let a be $\geq (\ell + 1)$ -far. It follows that for no $j \geq \ell + 1$ there exists a vertex $a' \in A_j \setminus A_{j+1}$ with $\mathbf{dist}_{\tilde{G}}(a, a') \leq b^j - \frac{1}{2}c_{\ell, \ell+1}$. Note that $b^j - \frac{1}{2}c_{\ell, \ell+1} = b^j - \frac{1}{2}b^{\ell+1} \geq \frac{1}{2}b^j \geq \frac{1}{4}(b^{j+1} - 1)/(b - 1) \geq \frac{1}{4}c_{\ell, j}$, where the second inequality holds since $b - 1 \geq b/2$. Thus, for no $j > i$, there exists an vertex $a' \in A_j \setminus A_{j+1}$ with $\mathbf{dist}_{\tilde{G}}(a, a') \leq \frac{1}{4}c_{\ell, j}$, so a is active. \square

Equipped with this notion, we can prove the following lemma. It immediately implies that H is a $(1 + \epsilon, n^{o(1)})$ -spanner since every vertex is $\geq (k + 1)$ -far.

Lemma 6.5. *For any shortest s - t path $\pi_{s,t}$ in G where every vertex $v \in \pi_{s,t} \setminus \{s, t\}$ is $\geq i$ -far for some $0 < i \leq k + 1$, we have*

$$\mathbf{dist}_H(s, t) \leq (1 + 20i\epsilon')\mathbf{dist}_G(s, t) + b^i \leq (1 + \epsilon)\mathbf{dist}_G(s, t) + n^{o(1)}.$$

Proof. As $\epsilon \leq 1$ it follows that $\epsilon' \leq 1/20$. Let us prove the claim by induction on i .

Base case $i = 1$: Let v_1 be the vertex right after s and let v_2 be the vertex right before t on $\pi_{s,t,G}$. As every vertex v on $\pi_{s,t,G} \setminus \{s, t\}$ is ≥ 1 -far, it follows by the definition of ≥ 1 -far that v belongs to $A_0 \setminus A_1$. That is, all vertices on $\pi_{v_1, v_2, G}$ belong to $A_0 \setminus A_1$. Furthermore, by Lemma 6.4 it follows that each such vertex v is active. Thus the shortest path from v to all vertices in $A_0 \setminus A_1$ at distance at most $\frac{1}{8 \log n} b \geq 1$ are included in H and, in particular, all edges of $\pi_G(v_1, v_2)$ belong to H . As H contains \tilde{G} , $\mathbf{dist}_H(s, v_1) \leq \log n$ and $\mathbf{dist}_H(v_2, t) \leq \log n$. Finally, $\mathbf{dist}_G(v_1, v_2) \leq \mathbf{dist}_G(s, t) + 2$. Thus it follows

that $\mathbf{dist}_H(s, t) \leq \mathbf{dist}_G(v_1, v_2) + 2 \log n \leq \mathbf{dist}_G(s, t) + 2 \log n + 2 \leq \mathbf{dist}_G(s, t) + (\log n)/\epsilon'$, where the last inequality holds since $\epsilon' \leq 1/20$.

Inductive step $i \mapsto i + 1$, for $i \geq 1$: As every vertex v in $\pi_{s,t,G} \setminus \{s, t\}$ is $\geq (i + 1)$ -far, every such vertex must belong to $A_\ell \setminus A_{\ell+1}$ for some $\ell \leq i$.

Let the set \mathcal{C} consist of the vertices on $\pi_{s,t}$ that are $\geq (i + 1)$ -far but not $\geq i$ -far. In other words, a vertex $v \in \pi_{s,t}$ is in \mathcal{C} if and only if there is a vertex $a \in A_i \setminus A_{i+1}$, with $\mathbf{dist}_{\tilde{G}}(v, a) \leq b^i - \frac{1}{2}c_{\ell,i}$ (possibly $a = v$). (Note that \mathcal{C} might be empty.) We prove the following claim for vertices in \mathcal{C} .

Claim 6.6. *For any vertex $v \in \mathcal{C}$, there exists a vertex $a \in A_i \setminus A_{i+1}$ at distance at most $\frac{1}{4}b^i$ in \tilde{G} , such that a is active.*

Proof. Let ℓ be such that $v \in A_\ell \setminus A_{\ell+1}$. If $\ell = i$, then v is active by Lemma 6.4 and, thus, the claim follows. If $\ell < i$, then let $a \in A_i \setminus A_{i+1}$ be the vertex closest to v among all vertices in $A_i \setminus A_{i+1}$. As v is $\geq (i + 1)$ -far, we know that

1. $\mathbf{dist}_{\tilde{G}}(v, a) \leq b^i - \frac{1}{2}c_{\ell,i}$, and
2. for no $j \geq i + 1$ there exists a vertex $a' \in A_j \setminus A_{j+1}$ with $\mathbf{dist}_{\tilde{G}}(v, a') \leq b^j - \frac{1}{2}c_{\ell,i+1}$.

Recall that $b \geq 5$, which implies that $b - 1 \geq 2b/3$. Suppose for contradiction that a is not active. If a is not active there must exist a vertex $a' \in A_p \setminus A_{p+1}$ with $p > i$ such that $\mathbf{dist}_{\tilde{G}}(a, a') \leq c_{i,p}/4$. But this implies that

$$\begin{aligned} \mathbf{dist}_{\tilde{G}}(v, a') &\leq \mathbf{dist}_{\tilde{G}}(v, a) + \mathbf{dist}_{\tilde{G}}(a, a') \leq b^i - \frac{1}{2}c_{\ell,i} + \frac{1}{4}c_{i,p} \\ &= b^i + \frac{1}{2}b^{i+1} + \frac{1}{4} \frac{b^{p+1} - b^{i+1}}{b - 1} - \frac{1}{2}c_{\ell,i+1} \leq b^i + \frac{1}{2}b^{i+1} + \frac{3}{8} \frac{b^{p+1} - b^{i+1}}{b} - \frac{1}{2}c_{\ell,i+1} \\ &\leq \frac{5}{8}b^i + \frac{7}{8}b^p - \frac{1}{2}c_{\ell,i+1} \leq b^p - \frac{1}{2}c_{\ell,i+1}. \end{aligned}$$

This gives a contradiction to the assumption that v is $\geq (i + 1)$ -far. \square

To prove the lemma we partition $\pi_{s,t}$ by constructing a sequence of vertices

$$t_0 = s, s_1, t_1, s_2, t_2, \dots, s_{h-1}, t_{h-1}, s_h = t$$

iteratively as follows: for each s_g with $g \geq 1$, let s_g be the first vertex on the path $\pi_{s,t}(t_{g-1}, t]$ that is in \mathcal{C} or if there is no such vertex, we set $s_g = t$ and $h = g$ and end the sequence. If the sequence does not end with s_g , let t_g with $g \geq 1$ be the farthest vertex from s_g on $\pi[s_g, t] \cap B_G(s_g, \frac{1}{10 \log n} b^{i+1})$ that is in \mathcal{C} . Note that such a vertex always exists since $s_g \in \pi[s_g, t] \cap B_G(s_g, \frac{1}{10 \log n} b^{i+1}) \cap \mathcal{C}$. (If s_g is the only such vertex then $t_g = s_g$.)

Clearly, the path $\pi_{s,t}$ is partitioned by the path segments $\pi_{s,t}[t_{g-1}, s_g]$ for $1 \leq g \leq h$ and the segments $\pi_{s,t}[s_g, t_g]$ for $1 \leq g < h$. Observe that for the former kind of segments, we have that the internal path vertices of a piece $\pi_{s,t}[t_{g-1}, s_g]$, i.e. the vertices in $\pi_{s,t}(t_{g-1}, s_g)$, are not in \mathcal{C} by definition. Thus, we have that all such vertices are $\geq i$ -far (since by assumption every vertex on $\pi_{s,t}$ is $\geq i + 1$ -far and every vertex in \mathcal{C} is not $\geq i$ -far but none of the vertices on $\pi_{s,t}(t_{g-1}, s_g)$ are in \mathcal{C}). By the sub-path property of shortest paths it holds that $\pi_{s,t}(t_{g-1}, s_g)$ is a shortest path between t_{g-1} and s_g . Hence, we can invoke the induction hypothesis on $\pi_{s,t}(t_{g-1}, s_g)$, with leads to the following statement

$$\mathbf{dist}_H(t_{g-1}, s_g) \leq (1 + 20i\epsilon') \mathbf{dist}_G(t_{g-1}, s_g) + b^i. \quad (22)$$

For the path segments of the form $\pi_{s,t}[s_g, t_g]$, claim 6.6 shows that there exist vertices a and a' at distance at most $\frac{1}{4}b^i$ in \tilde{G} from s_g and t_g respectively that are in $A_i \setminus A_{i+1}$ and are active. Then since

$$\begin{aligned} \mathbf{dist}_G(a, a') &\leq \mathbf{dist}_{\tilde{G}}(a, s_g) + \mathbf{dist}_G(s_g, t_g) + \mathbf{dist}_{\tilde{G}}(t_g, a') \\ &\leq \frac{1}{4}b^i + \frac{1}{10 \log n} b^{i+1} + \frac{1}{4}b^i \\ &\leq \frac{1}{8 \log n} \cdot b^{i+1}, \end{aligned}$$

we have that the shortest path $\pi_{a,a'}$ from a to a' is in H , i.e., that $\mathbf{dist}_H(a, a') = \mathbf{dist}_G(a, a')$. By applying the triangle inequality to G and the fact that $\mathbf{dist}_G(x, y) \leq \mathbf{dist}_{\tilde{G}}(x, y)$ (which follows from the fact that \tilde{G} is a spanner of G) it follows that

$$\mathbf{dist}_H(a, a') = \mathbf{dist}_G(a, a') \leq \mathbf{dist}_G(a, s_g) + \mathbf{dist}_G(s_g, t_g) + \mathbf{dist}_G(t_g, a') \leq \frac{1}{2}b^i + \mathbf{dist}_G(s_g, t_g).$$

Since $\tilde{G} \subseteq H$, it holds that $\mathbf{dist}_H(s_g, a) \leq \mathbf{dist}_{\tilde{G}}(s_g, a) \leq \frac{1}{4}b^i$. The same holds for $\mathbf{dist}_H(a', t_g)$. We, thus, have that

$$\mathbf{dist}_H(s_g, t_g) \leq \mathbf{dist}_H(s_g, a) + \mathbf{dist}_H(a, a') + \mathbf{dist}_H(a', t_g) \leq b^i + \mathbf{dist}_G(s_g, t_g). \quad (23)$$

Finally, let us prove that we did not partition the path into many segments: we claim that

$$h \leq \left\lceil \frac{10 \log n \cdot \mathbf{dist}_G(s, t)}{b^{i+1}} \right\rceil$$

which can be seen from carefully studying the requirements to pick given s_g the next t_g and s_{g+1} , which stipulate that each s_g and s_{g+1} are at distance at least $\frac{1}{10 \log n} b^{i+1}$.

We now combine our insights and the fact that $\epsilon' \leq 1/20$, and take the sum over the path segments. This gives

$$\begin{aligned} \mathbf{dist}_H(s, t) &= \sum_{g=1}^h \mathbf{dist}_H(t_{g-1}, s_g) + \sum_{g=1}^{h-1} \mathbf{dist}_H(s_g, t_g) \\ &\leq \sum_{g=1}^h ((1 + 20i\epsilon') \mathbf{dist}_G(t_{g-1}, s_g) + b^i) + \sum_{g=1}^{h-1} (b^i + \mathbf{dist}_G(s_g, t_g)) \\ &\leq \sum_{g=1}^h (1 + 20i\epsilon') \mathbf{dist}_G(t_{g-1}, s_g) + \sum_{g=1}^{h-1} \mathbf{dist}_G(s_g, t_g) + h(b^i + b^i) \\ &\leq (1 + 20i\epsilon') \mathbf{dist}_G(s, t) + \left\lceil \frac{\mathbf{dist}_G(s, t)}{b^{i+1}} \right\rceil \cdot 20 \log n \cdot b^i \\ &\leq (1 + 20i\epsilon') \mathbf{dist}_G(s, t) + 20 \log n \cdot \epsilon' \cdot \frac{\mathbf{dist}_G(s, t)}{\log n} + 20 \log n b^i \\ &\leq (1 + 20(i+1)\epsilon') \mathbf{dist}_G(s, t) + b^{i+1} \\ &\leq (1 + \epsilon) \mathbf{dist}_G(s, t) + n^{o(1)} \end{aligned}$$

as required. □

Sparsity. In order to prove sparsity, let us establish the following key claim. The claim is in fact slightly more powerful than necessary, however, this power will be exploited when bounding the running time.

Claim 6.7. *For any vertex v , and for any $i \leq k$ there are at most $O(n^{1/k} \log n)$ active vertices in $A_i \setminus A_{i+1}$ in*

$$B_{\tilde{G}}\left(v, \frac{1}{8}b^{i+1}\right)$$

with high probability. Further, if a vertex $v \in A_i \setminus A_{i+1}$ is active, we have $|B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})| \leq n^{(i+1)/k}$ with high probability.

Proof. Let us observe that for any vertex $v \in V$, if $|B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})| \geq n^{(i+1)/k}$ there exists at least one vertex a in A_{i+1} that hits $B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})$ with high probability.

We now claim that if there is such a vertex a , no vertex a' in $B_{\tilde{G}}(v, \frac{1}{8}b^{j+1}) \cap A_j$ is active for any $j \leq i$. To see this let us assume for the sake of contradiction that there exists such a vertex a' . Now, let us first consider the case that the vertex a is active. Then, the distance between a and vertex a' is $\mathbf{dist}_{\tilde{G}}(a, v) + \mathbf{dist}_{\tilde{G}}(v, a') \leq \frac{1}{8}b^{i+1} + \frac{1}{8}b^{j+1} = \frac{1}{4}b^{i+1} \leq \frac{1}{4}c_{j,i+1}$ which implies that a must have deactivated a' and therefore leads to a contradiction.

Now, let us consider the case, where a is not active. Then, there exists a vertex $a'' \in A_\ell \setminus A_{\ell+1}$ that is active and with $\ell > i + 1$ such that $\mathbf{dist}_{\tilde{G}}(a, a'') \leq \frac{1}{4}c_{i+1,\ell}$. Such a vertex exists by definition if a is not active so we have that a'' is well-defined.

But then we have by the triangle inequality that

$$\begin{aligned} \mathbf{dist}_{\tilde{G}}(a', a'') &\leq \mathbf{dist}_{\tilde{G}}(a', v) + \mathbf{dist}_{\tilde{G}}(v, a) + \mathbf{dist}_{\tilde{G}}(a, a'') \\ &\leq \frac{1}{8}b^{j+1} + \frac{1}{8}b^{i+1} + \frac{1}{4}c_{i+1,\ell} \\ &\leq \frac{1}{4}c_{j,\ell} \end{aligned}$$

where we use in the last inequality that $c_{j,\ell} = \sum_{y=j+1}^{\ell} b^y = \sum_{y=j+1}^{i+1} b^y + \sum_{y=i+2}^{\ell} b^y = \sum_{y=j+1}^{i+1} b^y + c_{i+1,\ell}$ and since $i \geq j$ we have that $\sum_{y=j+1}^{i+1} b^y \geq b^{i+1}$. This implies that a' cannot be active since a'' is close enough to a' to deactivate it during the algorithm. This again leads to a contradiction and thereby completes the proof of our claim that a' is not active.

Finally, we conclude that we have for any vertex v and level i , that either $|B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})| < n^{(i+1)/k}$ in which case we have by a straightforward application of a Chernoff bound that with high probability at most $O(n^{1/k} \log n)$ of these vertices are sampled into A_i . Thus, there can also be only $O(n^{1/k} \log n)$ active vertices from A_i in the ball.

Otherwise $|B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})| \geq n^{(i+1)/k}$, so by the above claim we have that some vertex a in A_{i+1} hits the ball $B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})$ which results in all vertices in the ball that are in $A_j \setminus A_{j+1}$ to be deactivated for any $j \leq i$, and in particular for i . This proves the first part of the claim.

For the second part of the claim, observe that the contrapositive of the above claim with $a' = v$ implies that there exists no active vertex a'' in A_{i+1} in the ball $B_{\tilde{G}}(v, \frac{1}{4}c_{i,i+1}) \subseteq B_{\tilde{G}}(v, \frac{1}{8}b^{i+1})$, since $c_{i,i+1} = b^{i+1}$. Further, by the contrapositive of our initial observation, we have that then $|B(v, \frac{1}{8}b^{i+1})| < n^{(i+1)/k}$. \square

A special case of the claim is that each vertex $a \in A_i \setminus A_{i+1}$ has $\tilde{O}(n^{1/k})$ active vertices in its ball $B(a, \frac{1}{8}b^{i+1})$ that are in $A_i \setminus A_{i+1}$. Thus, if a is active itself this gives an upper bound on the number of paths that are included in H due to a (if a is not active no paths are added). Since the maximum such path length at any level is $\frac{1}{8}b^{k+1}$, we have that there are only $\tilde{O}(n^{1+1/k}b^{k+1}) = n^{1+o(1)}$ edges in H due to paths.

The spanner \tilde{G} that is additionally added to H contains only $\tilde{O}(n)$ edges which is subsumed in the previous bound.

Running Time. To bound the running time, observe that in order to determine which vertices are active, we run for each $i \leq k$, a BFS to depth $\frac{1}{8}b^{i+1}$ from every active vertex in $A_i \setminus A_{i+1}$ on \tilde{G} . By Claim 6.7, each vertex $v \in V$ is only explored by $\tilde{O}(n^{1/k})$ active vertices on each level. There are only $\sqrt{\log n}$ levels, so each vertex $v \in V$ is only explored by $\tilde{O}(n^{1/k})$ active vertices in total. Thus, every edge incident to each vertex v in \tilde{G} is only explored $\tilde{O}(n^{1/k})$ times. Since \tilde{G} has only $\tilde{O}(n)$ edges, and since each BFS runs linearly in the number of edges explored, we can bound the total running time by $\tilde{O}(n^{1+1/k})$, for each i , and thus also for all values i .

After determining which vertices are active, it remains to bound the time to compute the paths between any two active vertices $a, a' \in A_i \setminus A_{i+1}$ for some i . For $i \geq \gamma$, we can simply check the distances between any pair of such vertices (even the ones that are non-active) by looking up their shortest paths in \mathcal{A} and inserting them into H if the criterion is satisfied. We will calculate the running time of \mathcal{A} at the end. For any active vertex $a \in A_i \setminus A_{i+1}$, where $i < \gamma$, we further run a BFS on G to depth $\frac{1}{8 \log n} b^{i+1}$. We observe that since \tilde{G} is a $\log n$ -spanner of G , we have for every $v \in V$ that

$$B_{\tilde{G}}(v, \frac{1}{8 \log n} b^{i+1}) \subseteq B_G(v, \frac{1}{8} b^{i+1}).$$

By Claim 6.7 we have again that each vertex v is only explored $\tilde{O}(n^{1/k})$ times during these executions of BFS. Using the second fact of Claim 6.7, we further have that each BFS from an active vertex $a \in A_i \setminus A_{i+1}$ only explores at most $n^{(i+1)/k}$ vertices. But this in turn implies that each vertex v that is strictly contained in such a ball $B_G(a, \frac{1}{8 \log n} b^{i+1})$ has degree bounded by $n^{(i+1)/k}$. Since we explore only for $i < \gamma$, we therefore have that each explored vertex has degree at most $n^{\gamma/k} \leq n^\kappa$. Putting everything together, during all BFS explorations, we scan every vertex at most $\tilde{O}(n^{1/k})$ times while the number of edges present at the vertex is at most n^κ and since we have n vertices, the total time required for all explorations can be bounded by $O(n \cdot n^{1/k} \cdot n^\kappa) = n^{1+\kappa+o(1)}$.

Finally, we have to account for the data structures used. While the running time of \mathcal{B} is subsumed in our previous bound, we have that \mathcal{A} has worst-case update time $\tilde{O}(n^\delta (n^{\omega(1,1,\kappa)-\kappa} + (n^{1-\gamma/k})^2 n^\kappa))$ for any constant $\kappa \leq \kappa^*$ and $\delta > 0$ by Corollary 6.2 since $|S| = n^{1-\gamma/k} = n^{1-\kappa+o(1)}$. Setting $\kappa = \kappa^*$ and $\delta = 0.01$, the size $|S| = o(\sqrt{n})$ and thus we obtain worst-case update time $O(n^{1+\kappa^*+o(1)}) = O(n^{1.529})$.

7 Applications

Let $\epsilon > 0$ be an arbitrarily small constant. In this section we give two applications of the data structures developed in the previous sections, namely fully dynamic $(1 + \epsilon)$ -approximate all-pairs shortest paths with path reporting and $(2 + \epsilon)$ -approximate fully dynamic Steiner tree. Both algorithms are the first algorithms that take sub-quadratic worst-case time for these problems.

7.1 Fully dynamic approximate all-pairs shortest paths with path reporting

In this subsection we give a fully dynamic data structure that reports $(1 + \epsilon)$ -approximate all-pairs shortest paths in sub-quadratic worst-case time. Note that all previous such data structure were only able to report the *distance*, i.e. the *length* of the paths, but not the actual paths.

Theorem 7.1. *Let $\epsilon > 0$ be an arbitrarily small constant. There exists an algorithm that maintains $(1 + \epsilon)$ -approximate all-pairs shortest-path in worst-case time $n^{1+\kappa^*+o(1)} = O(n^{1.529})$ per edge update, in worst-*

case time $n^{1+o(1)}$ per path reporting query and per distance reporting query with high probability against an oblivious adversary. The preprocessing time is $n^{\omega+o(1)}$ if the initial graph is non-empty and $\tilde{O}(n^2)$ if the initial graph is empty.

Proof. We maintain the following data structure:

(1) We maintain the fully dynamic path reporting data structure given in the statement of Theorem 5.1 on G with $D = 2n^{o(1)}/\epsilon$ and $\kappa = \kappa^* < 0.529$.

(2) We maintain for G the fully dynamic $(1 + \epsilon/2, n^{o(1)})$ -spanner from Theorem 6.1.

To answer a distance reporting query we first ask a distance query with parameter D in (1). This gives us the exact answer if the distance is less than D . Otherwise we run a static shortest path algorithm on the spanner that we maintain in (2). As the spanner has at most $n^{1+o(1)}$ edges this takes time $n^{1+o(1)}$. Note that, by the choice of D , the shortest path on the spanner gives a $(1 + \epsilon)$ -approximation of the shortest path in G as $(1 + \epsilon/2)\mathbf{dist}(s, t) + n^{o(1)} \leq (1 + \epsilon)\mathbf{dist}(s, t)$ for $\mathbf{dist}(s, t) \geq D = 2n^{o(1)}/\epsilon$. Thus a distance query returns a $(1 + \epsilon)$ -approximate answer and takes time $n^{\kappa+o(1)}/\epsilon + n^{1+o(1)} = n^{1+o(1)}$.

To answer a path reporting query between two nodes s and t we first ask a distance query with parameter D in (1). If the distance is less than D , we ask a path reporting query in (1) in time $O(D^2 n^{\kappa^*}) = o(n)$. Otherwise, we execute a static shortest path algorithm on the spanner that we maintain in (2). Thus a path reporting query returns a $(1 + \epsilon)$ -approximate shortest path and takes time $n^{1+o(1)}$ time. \square

If we increase the preprocessing time to $O(n^{2.621})$, and the update time to $O(n^{1.843})$, we can reduce the cost of a distance reporting query even further.

Corollary 7.2. *Let $\epsilon > 0$ be an arbitrarily small constant. There exists an algorithm that maintains $(1 + \epsilon)$ -approximate all-pairs shortest-path in worst-case time $O(n^{1.843+o(1)})$ per edge update, in worst-case time $O(n^{1+o(1)})$ per path reporting query, and in worst-case time $O(n^{.45})$ per distance reporting query with high probability against an oblivious adversary. The preprocessing time is $n^{\omega+o(1)}$ if the initial graph is non-empty and $\tilde{O}(n^2)$ if the initial graph is empty.*

Proof. Additionally maintain a fully dynamic $(1 + \epsilon)$ -approximate APSP data structure from [82], which takes worst-case time $O(n^{1.843})$ per edge update and worst-case time $O(n^{0.45})$ per distance query. It speeds up the distance reporting queries and, in combination with theorem 7.1, leads to the result. \square

7.2 Steiner trees with terminal vertex and edge updates

In this section we give a further application of the data structure of the previous section. Assume we are given an (unweighted) graph $G = (V, E)$ with a dynamically changing edge set and a dynamically changing terminal set $S \subseteq V$. A *Steiner tree* T_S for a vertex set S is a tree that (1) is a subgraph of G , (2) spans S , and (3) has the minimum number of edges. A β -approximate *Steiner tree* is a tree which fulfills conditions (1) and (2), and the weight of the tree is at most a factor β larger than the weight of the edges of a minimum Steiner tree.

A 2-approximate Steiner tree can be found as follows: Construct a weighted graph \tilde{G} that consists of a clique on the vertices of S such that each edge (u, v) has length $\mathbf{dist}_G(u, v)$, i.e., the length of the shortest path between u and v in G . Find an MST \tilde{T}_S in \tilde{G} . Its weight $w(\tilde{T}_S)$ gives a 2-approximation of the value OPT of the optimal Steiner tree, i.e. $OPT \leq w(\tilde{T}_S) \leq 2OPT$. The reason is as follows: Consider an Eulerian tour $\mathcal{E}(T_S)$ of the optimal Steiner tree T_S in G . It traverses every edge twice and, thus, has length $2OPT$. Now replace the subpath between two consecutive terminal vertices u and v of $\mathcal{E}(T_S)$ by an edge. Note that the length of this subpath is at least $\mathbf{dist}_G(u, v)$ and that there is an edge (u, v) in \tilde{G} with length $\mathbf{dist}_G(u, v)$. Thus $\mathcal{E}(T_S)$ induces a cycle in \tilde{G} whose length is at most $2OPT$. As the minimum spanning

tree \tilde{T}_S in \tilde{G} has a weight that is at most the length of this cycle, its weight is at most $2OPT$. Note that if the weights in \tilde{G} are between $\mathbf{dist}_G(u, v)$ and $(1 + \epsilon/2)\mathbf{dist}_G(u, v)$ for some arbitrarily small $\epsilon > 0$, then the weight of \tilde{T}_S is at most $(2 + \epsilon)OPT$. To construct the corresponding approximate Steiner tree replace each edge in \tilde{T}_S by a shortest path in G between its endpoints and compute a tree \tilde{T}'_S in the resulting graph. The weight of \tilde{T}'_S is at most the weight of \tilde{T}_S .

We consider the following dynamic changes to the input: (i) edge insertions and deletions and (ii) additions and removals from S . In the *fully dynamic unweighted Steiner tree problem* the goal is to maintain a Steiner tree after each modification to the input. In the β -*approximate fully dynamic unweighted Steiner tree problem* the goal is to maintain an β -approximate Steiner tree after each modification of the input. Note that we want to maintain an actual tree, *not just the value* of the β -approximate Steiner tree. We show the following result.

Theorem 7.3. *Let $\epsilon > 0$ be an arbitrarily small constant. There exists an algorithm that solves the $(2 + \epsilon)$ -approximate fully dynamic unweighted Steiner tree problem with high probability against an oblivious adversary, in worst-case time $O(n^{1.529} + s^2 \cdot n^{1+o(1)})$ per edge update and in worst-case time $sn^{1+o(1)}$ per vertex addition to or removal from S , where s is the current size of S . The preprocessing time is $n^{\omega+o(1)}$ if the initial graph is non-empty and $\tilde{O}(n^2)$ if the initial graph is empty.*

Proof. By our discussion before the theorem it suffices to maintain the graph \tilde{G} whose edge weights are a $(1 + \epsilon/2)$ -approximation of the length of their endpoints in G and a minimum spanning tree in \tilde{G} . To build the actual Steiner tree, we need to replace then each edge in this spanning tree into the corresponding shortest path in G .

For convenience, we denote by $A(\epsilon, n)$, the additive error of a $(1 + \epsilon, n^{o(1)})$ -spanner as maintained in Theorem 6.1, i.e. Theorem 6.1 maintains a $(1 + \epsilon, A(\epsilon, n))$ -spanner. From this, it is also easy to see that $A(\epsilon, n) = n^{o(1)}$ for any constant $\epsilon > 0$.

Now, we describe the data structures to maintain the necessary information for updating our Steiner tree approximation:

(1) We maintain the fully dynamic path reporting data structure given in the statement of Theorem 5.1 on G with $D = 4A(\epsilon/4, n)/\epsilon$ and $\kappa = \kappa^* < 0.529$.

(2) We maintain the fully dynamic $(1 + \epsilon/4, A(\epsilon/4, n))$ -spanner from theorem 6.1.

(3) We maintain a $|S| \times |S|$ array that contain the $(1 + \epsilon/4)$ - approximate all-pairs shortest path lengths between any pair of terminals.

(4) We maintain the graph \tilde{G} and a fully dynamic minimum spanning tree data structure from \tilde{G} that takes $n^{o(1)}$ deterministic worst-case update time [33].

To answer a distance query we first ask a distance query with parameter D in (1). This gives us the exact answer if the distance is less than D . Otherwise we run a static shortest path algorithm on the spanner that we maintain in (2). As the spanner has at most $n^{1+o(1)}$ edges this takes time $n^{1+o(1)}$. Note that, by the choice of D , the shortest path on the spanner gives a $(1 + \epsilon/2)$ -approximation of the shortest path in G as $(1 + \epsilon/4)\mathbf{dist}(s, t) + A(\epsilon/4, n) \leq (1 + \epsilon/2)\mathbf{dist}(s, t)$ for $\mathbf{dist}(s, t) \geq D = 4A(\epsilon/4, n)/\epsilon$. Thus a distance query returns a $(1 + \epsilon/2)$ -approximate answer and takes time $n^{\kappa^*+o(1)}/\epsilon + n^{1+o(1)} = n^{1+o(1)}$.

To answer a path reporting query between two nodes s and t we first ask a distance query with parameter D in (1). If the distance is less than D , we ask a path reporting query in (1). Otherwise, we execute a static shortest path algorithm on the spanner that we maintain in (2). Thus a path reporting query returns a $(1 + \epsilon/2)$ -approximate shortest path and takes time $n^{1+o(1)} + n^{\kappa^*+o(1)}/\epsilon = n^{1+o(1)}$ time.

After each update to either E or S we update data structures (1) - (4) as described below. Then we build the approximate Steiner Tree \tilde{T}'_S (in G) from \tilde{T}_S as described above, executing $s - 1$ path reporting queries. This takes time $s \cdot n^{1+o(1)}$, which as we will see is subsumed by the runtime of updating (4).

We are left with describing how we update (1) - (4). Each edge update leads to the corresponding update in (1) and (2). Then we recompute \tilde{G} from scratch using $O(s^2)$ distance queries and compute its minimum spanning tree as well as the dynamic MST data structure from scratch. Afterwards we build \tilde{T}'_S as described above. This takes total time $n^{1+\kappa^*+o(1)} + s^2 \cdot n^{1+o(1)}$.

If a vertex v is added to S , we compute the distance in the spanner from v to all nodes in S , add an edge from v to every vertex in S with the corresponding length to \tilde{G} , and update its dynamic MST data structure. This takes time $sn^{1+o(1)}$. Afterwards we build \tilde{T}'_S as described above.

If a vertex is removed from S , we remove its incident edges from \tilde{G} and its dynamic MST data structure. This takes time $sn^{1+o(1)}$. Afterwards we build \tilde{T}'_S as described above. \square

At the cost of increasing the preprocessing time to $O(n^{2.621})$ we can additionally maintain a fully dynamic $(1 + \epsilon/2)$ -approximate APSP data structure from [82], which takes worst-case time $O(n^{1.843})$ per edge update and worst-case time $O(n^{0.45})$ per distance query. It speeds up the distance query and leads to the following result.

Theorem 7.4. *Let $\epsilon > 0$ be an arbitrarily small constant. There exists an algorithm that solves the $(2 + \epsilon)$ -approximate fully dynamic unweighted Steiner tree problem with high probability against an oblivious adversary in worst-case time $O(n^{1.843+o(1)} + s^2 \cdot n^{0.45} + s \cdot n^{1+o(1)})$ per edge update, in worst-case time $sn^{1+o(1)}$ per vertex addition to or removal from S , where s is the current size of S . The preprocessing time is $O(n^{2.621})$.*

8 Combinatorial dynamic spanner algorithms

In this section, we present our combinatorial algorithms to maintain $(1 + \epsilon, \beta)$ -spanners in dynamic graphs. We first present a simple algorithm to maintain a spanner in a partially dynamic graph, then show how to extend the construction to obtain near-linear update time in fully dynamic graphs. Finally, we present a deterministic algorithm to maintain an emulator with large additive error in partially dynamic graphs.

8.1 A Simple Partially Dynamic Spanner

In this section, we prove the following theorem.

Theorem 8.1. *For any constant $0 < \epsilon \leq 1$, given an undirected, unweighted partially dynamic graph, we can maintain a $(1 + \epsilon, n^{o(1)})$ -spanner of size $n^{1+o(1)}$ in total time $m^{1+o(1)}$ with high probability against an oblivious adversary.*

The Algorithm. For $k = \sqrt{\log n}$, we sample sets $V = A_0 \supseteq A_1, \dots \supseteq A_k \supseteq A_{k+1} = \emptyset$ where A_i for $i \in [1, k]$ is obtained by sampling each vertex in V with probability $n^{-i/k} \log n$ (and to make the sets nesting add it to all A_j where $j \leq i$).

Let $\epsilon' = \epsilon/8$. Throughout the algorithm, we say for each $a \in A_i \setminus A_{i+1}$ that a is *active* if and only if there is no vertex $a' \in A_j$ for $j > i$, with $\text{dist}(a, a') \leq \epsilon'^{-(j+1)} - \epsilon'^{-(i+1)}$. Thus, all vertices in A_k are always active. In the decremental algorithm, once a vertex becomes active, it remains active for the rest of the algorithm, and in the incremental algorithm, once a vertex becomes inactive it remains inactive for the rest of the algorithm. Further, for each active vertex $a \in A_i \setminus A_{i+1}$, we maintain the ball $\mathcal{B}_a = B(a, \epsilon'^{-(i+1)})$. Otherwise, we let $\mathcal{B}_a = \emptyset$. A similar definition of active nodes was introduced by Chechik [31] for an emulator construction.

In order to obtain the spanner, it suffices to take the union of the shortest-path trees maintained from each vertex $a \in V$, truncated at depth r_a , i.e. the spanner is defined at all stages by

$$H = \bigcup_{a \in V, b \in \mathcal{B}_a} \pi_{a,b}$$

We maintain these shortest-path trees using the ES-tree data structure.

Spanner Approximation. At any stage, for an arbitrary shortest path $\pi_{s,t}$ from s to t , we can expose a path in H that approximates $\pi_{s,t}$ as follows.

Let i be chosen to be the smallest index such that $B(s, 2\epsilon'^{-(i+1)}) \cap A_{i+1} = \emptyset$. First we show that s is within distance at most $2\epsilon'^{-i}$ of an active vertex of A_i . Then, we have $i \leq k$ since $A_{k+1} = \emptyset$. By a straightforward application of the Chernoff bound it follows that $|B(s, 2\epsilon'^{-(i+1)})| \leq n^{(i+1)/k}$ with high probability, since otherwise A_{i+1} would likely hit $B(s, 2\epsilon'^{-(i+1)})$. Further, observe that by minimality of i , there exists some vertex a in $B(s, 2\epsilon'^{-i}) \cap A_i \setminus A_{i+1}$, i.e. there exists some vertex $a \in A_i$ at distance at most $2\epsilon'^{-i}$ from s . Now, observe that $B(a, \epsilon'^{-(i+1)}) \cap A_{i+1} \subseteq B(s, 2\epsilon'^{-(i+1)}) \cap A_{i+1} = \emptyset$ by the triangle inequality. Thus a is active.

Now, we distinguish two scenarios:

1. if $t \in \mathcal{B}_a$: then we can use the paths $\pi_{s,a} \circ \pi_{a,t}$ with combined weight less than $2\epsilon'^{-(i+1)} \leq 2\epsilon'^{-(k+1)}$,
2. otherwise, let s' be the vertex in $\mathcal{B}_a \cap \pi_{s,t}$ that is closest to t . As $t \notin \mathcal{B}_a$ it holds that $\text{dist}_G(a, s') = \epsilon'^{-(i+1)}$. As $\text{dist}_G(s, a) \leq 2\epsilon'^{-i}$ and $\epsilon' \leq 1/4$, it follows by the triangle inequality that $\text{dist}_G(s, s') \geq \epsilon'^{-(i+1)} - 2\epsilon'^{-i} \geq \epsilon'^{-(i+1)}/2$. By the triangle inequality it holds that $\text{dist}_G(s', a) \leq \text{dist}_G(s', s) + \text{dist}_G(s, a) \leq \text{dist}_G(s, s') + 2\epsilon'^{-i}$. Thus, we have that the path $\pi_{s,a} \circ \pi_{a,s'}$ has weight at most $\text{dist}_G(s, s') + 4\epsilon'^{-i} = \text{dist}_G(s, s') + 8\epsilon' \cdot (\epsilon'^{-(i+1)}/2) \leq (1 + 8\epsilon')\text{dist}_G(s, s') = (1 + \epsilon)\text{dist}_G(s, s')$.

If we are in the first scenario, we are done. Otherwise, we can recurse on $\pi_{s',t}$. Aggregating all of the path segments from the recursion, we derive

$$\text{dist}_H(s, t) \leq (1 + \epsilon)\text{dist}_G(s, t) + 2\epsilon'^{-(k+1)}.$$

Since this worked with high probability, and we only have a polynomial number of shortest paths to consider (namely n^2 shortest paths after each of the at most m edge deletions), we can take a union bound over all these shortest paths and obtain overall correctness with high probability.

Sparseness of the Spanner. Pick any $i \in [0, k-1]$. Let us observe that for any vertex $v \in V$, as long as

$$|B(v, \epsilon'^{-(i+1)})| \geq n^{(i+1)/k}$$

we have that at least one vertex $a'' \in A_{i+1}$ hits $B(v, \epsilon'^{-(i+1)})$ with high probability. Further, while a'' is in $B(v, \epsilon'^{-(i+1)})$, we claim that no vertex in $B(v, \epsilon'^{-(i+1)}) \cap A_i$ is active. For the sake of contradiction assume that there exists such a vertex in a' , then by the triangle inequality, we have

$$\text{dist}(a', a'') \leq \text{dist}(v, a'') + \text{dist}(v, a') \leq 2\epsilon'^{-(i+1)} < \epsilon'^{-(i+2)} - \epsilon'^{-(i+1)}$$

But that implies that a' should not be active. This implies that vertices in $B(v, \epsilon'^{-(i+1)}) \cap A_i$ can only become active after $|B(v, \epsilon'^{-(i+1)})| < n^{(i+1)/k}$. But then with high probability $|B(v, \epsilon'^{-(i+1)}) \cap A_i| = O(n^{1/k} \log^2 n)$ (again by using a Chernoff bound). Note that vertices in $B(v, \epsilon'^{-(i+1)}) \cap A_i$ are the only nodes of A_i in whose balls v can lie. Thus, each vertex v , only ever participates in $O(n^{1/k} \log^2 n)$ many balls rooted at vertices in A_i . Thus, the total size of the spanner can be upper bounded by $\tilde{O}(n^{1+1/k}) = n^{1+o(1)}$.

Correctness of the Algorithm. While the rest of the algorithm is rather straightforward, we show that by maintaining the balls \mathcal{B} for all active nodes we can maintain for every vertex whether it should become active (in the case of a decremental algorithm) or inactive (in the case of an incremental algorithm).

Consider a vertex a and let i be such that $a \in A_i \setminus A_{i+1}$. Consider the maximal layer j such that there exists a vertex $a' \in A_j \setminus A_{j+1}$, $i < j$, at distance $\mathbf{dist}(a, a') \leq \epsilon^{-(j+1)} - \epsilon^{-(i+1)}$. We will show below that a' is active. Then clearly we have $a \in \mathcal{B}_{a'}$ and a' can therefore inform a that it cannot become active yet. This suffices for an incremental algorithm since vertices only go from active to inactive. We argue next that this also suffices for a decremental algorithm (where vertices only go from inactive to active).

When a' informs a that a cannot be active, we say that a' invokes a *block* on a . This block stays in effect until a' revokes it, which happens when $\mathbf{dist}(a, a')$ grows above $\epsilon^{-(j+1)} - \epsilon^{-(i+1)}$. By maintaining $\mathcal{B}_{a'}$, we can detect when $\mathbf{dist}(a, a')$ grows above $\epsilon^{-(j+1)} - \epsilon^{-(i+1)}$ because if $a \in \mathcal{B}_{a'}$, then $\mathbf{dist}(a, a')$ is explicitly maintained, and if a leaves $\mathcal{B}_{a'}$ then we can detect this and conclude that $\mathbf{dist}(a, a')$ has grown above $\epsilon^{-(j+1)} - \epsilon^{-(i+1)}$. In summary, we determine the status of a by setting a to be active if and only if no other vertex has invoked a block on a (and not revoked it).

To show that a' is active, assume by contradiction that a' is inactive. It follows that there has to be a vertex $a'' \in A_\ell \setminus A_{\ell+1}$ for $j < \ell$ at distance at most $\epsilon^{-(\ell+1)} - \epsilon^{-(j+1)}$ to a' . But by the triangle inequality, we would have

$$\mathbf{dist}(a, a'') \leq \mathbf{dist}(a, a') + \mathbf{dist}(a', a'') \leq \epsilon^{-(j+1)} - \epsilon^{-(i+1)} + \epsilon^{-(\ell+1)} - \epsilon^{-(j+1)} = \epsilon^{-(\ell+1)} - \epsilon^{-(i+1)}$$

contradicting that a' was a vertex satisfying our constraint with maximal index j .

Total Time. Finally, let us bound the total time. Constructing the set A_i takes time $O(n^{1+o(1)})$. We have already shown that each vertex is in at most $O(n^{1/k} \log^2 n)$ balls over the course of the algorithm. Since each edge only participates in balls containing both of its endpoints, we have that each edge is in at most $O(n^{1/k} \log^2 n)$ balls.

Since ES-trees run in time $O(mr)$, each edge is scanned at most $O(mr)$ times for each of the at most $O(n^{1/k} \log^2 n)$ balls it is in. Each ball has radius at most $\epsilon^{-(i+1)}$, so an analysis over all levels and balls shows that the total update time can be bound by

$$\sum_{i \leq k} m \cdot O(n^{1/k} \log^2 n) \cdot \epsilon^{-(i+1)} = \tilde{O}(m^{1+1/\sqrt{\log n}} \epsilon^{-(\sqrt{\log n}+1)}) = m^{1+o(1)}.$$

Note that this also bounds the total number of changes that can occur to the spanner during all deletions as we always explicitly construct the spanner.

8.2 A Simple Deterministic Fully Dynamic Spanner

Although the efficiency of the previous algorithm crucially depends on randomness and the non-adaptiveness of the adversary, Roditty, Thorup and Zwick showed in [72] that the construction given above can be efficiently derandomized in the static setting. Thus, there exist a static algorithm, that runs in time $m^{1+o(1)}$ that computes a $(1 + \epsilon, n^{o(1)})$ -spanner deterministically. Rerunning this algorithm after every update yields the following theorem.

Theorem 8.2. *For any constant $0 < \epsilon \leq 1$, given an undirected, unweighted graph $G = (V, E)$ undergoing vertex updates, we can deterministically maintain a $(1 + \epsilon, n^{o(1)})$ -spanner of size $n^{1+o(1)}$ with preprocessing time $m^{1+o(1)} + O(n)$ and worst update time $m^{1+o(1)}$.*

We point out that a similar result can likely be obtained by using the randomized construction by Elkin and Zhang [45] which can be derandomized using the deterministic sparse neighborhood covers [13].

8.3 Partially Dynamic Algorithms to Maintain Emulators/Spanners with High Additive Error

In this section, we prove the following theorems.

Theorem 8.3. *Given a partially dynamic unweighted graph $G = (V, E)$, a multiplicative approximation parameter $\epsilon > 0$, and an additive approximation parameter $\alpha \geq 0$, there exists a deterministic algorithm that can maintain an $n^{1+o(1)}$ -edge $(1 + \epsilon, n^{\alpha+o(1)})$ -emulator H of G in total time $O(mn^{1-\alpha+o(1)})$.*

Theorem 8.4. *Given a partially dynamic unweighted graph $G = (V, E)$, a multiplicative approximation parameter $\epsilon > 0$, and an additive approximation parameter $\alpha \geq 0$, there exists a randomized algorithm that can maintain an $n^{1+o(1)}$ -edge $(1 + \epsilon, n^{\alpha+o(1)})$ -emulator H of G in total time $O(mn^{1-\alpha+o(1)})$. The algorithm works against an adaptive adversary.*

In order to implement a data structure as required by the theorem above, we heavily rely on the following result from [50] which we extend to incremental graphs.

Definition 8.5 (Refining and Coarsening). *Given a universe U and a collection \mathcal{C} of disjoint subsets of U undergoing changes, we say that \mathcal{C} is refining if every set $C \in \mathcal{C}$ is a subset of a set C' contained in an earlier version of \mathcal{C} . We say that \mathcal{C} is coarsening if every set $C \in \mathcal{C}$ is a superset of a set C' contained in an earlier version of \mathcal{C} .*

Theorem 8.6 (see [50], Definition 4.3, Lemma 4.1). *Given a decremental/incremental unweighted graph $G = (V, E)$, a parameter μ and $0 < \epsilon \leq 1$, then there exists a deterministic algorithm that maintains*

- a refining/coarsening collection of pairwise-disjoint vertex sets $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ where any set C_i has $\text{diam}(G[C_i]) \leq 4(1/\epsilon)^{\sqrt{\log n}} \left\lceil \frac{|E(C_i)|}{\mu} \right\rceil$ and $k \leq m/\mu$, and
- a graph H that forms a subgraph of G such that for every shortest path $\pi_{s,t}$ in G with $\pi_{s,t} \cap (\bigcup_i C_i) = \emptyset$, we have

$$\text{dist}_G(s, t) \leq \text{dist}_H(s, t) \leq (1 + \epsilon)\text{dist}_G(s, t) + O((1/\epsilon)^{\sqrt{\log n}})$$

and we have that $|H| = n^{1+o(1)}$.

The algorithm runs in total time $O(n\mu(1/\epsilon)^{1/\sqrt{\log n}} + m \log^2 n)$.

Proof. Let $\epsilon' = \epsilon/4$. We say a vertex $v \in V$, is μ -heavy if the graph G induced by vertices in the ball $B(v, (1/\epsilon')^{\sqrt{\log n}})$ contains more than μ edges. Otherwise, we say v is μ -light. We further maintain a graph G^{heavy} that is the graph G induced by the vertices $w \in V$ that are in the ball of some μ -heavy vertex $v \in V$. That is, a vertex w is in G^{heavy} if it is either μ -heavy or has some vertex v in its ball to depth $(1/\epsilon')^{\sqrt{\log n}}$ that is μ -heavy. We let \mathcal{C} be the set of connected components in G^{heavy} . Here, we point out that a slightly modified ES-tree can be used to maintain when a vertex transitions from being μ -heavy to μ -light (or vice versa) and while a vertex is μ -light, we can maintain its shortest path tree truncated at depth $(1/\epsilon')^{\sqrt{\log n}}$ explicitly. For partially dynamic graphs this modified ES-tree can be implemented in time $O(\mu(1/\epsilon')^{\sqrt{\log n}})$ per vertex, thus we can maintain the ES-trees in total update time $n^{1+o(1)}\mu$.

Now, consider the diameter d of some component $C_i \in \mathcal{C}$ at any stage in $G[C_i]$. Let u and v be two vertices in C_i at distance d in G . Observe that we then have at least $\lceil d/(4(1/\epsilon')^{\sqrt{\log n}}) \rceil$ vertices in C_i that are at distance at least $4(1/\epsilon')^{\sqrt{\log n}}$ from each other. Since each vertex in C_i is at distance at most $(1/\epsilon')^{\sqrt{\log n}}$ to a μ -heavy vertex, we can thus find $\lceil d/(4(1/\epsilon')^{\sqrt{\log n}}) \rceil$ vertices that have disjoint balls and are μ -heavy. Thus, we derive the upper bound on the diameter $d \leq 4(1/\epsilon')^{\sqrt{\log n}} \lceil \frac{E(C_i)}{\mu} \rceil$, by the pigeonhole principle. It

is straightforward to see that μ -heavy vertices in different components have mutually disjoint balls, so we derive the bound on the number of components by another application of the pigeonhole principle, i.e. we derive that there are at most m/μ connected components in \mathcal{C} at any stage since each component contains at least one μ -heavy vertex.

It is further not hard to see that when G is a decremental graph, then G^{heavy} is a decremental graph since the number of vertices in a ball to fixed radius can only decrease over time, as does the degree of each vertex. Thus, it is straightforward to maintain in $O(m \log^2 n)$ time the components in G^{heavy} using a connectivity data structure (see [57, 87]). For incremental graphs, G^{heavy} is incremental and the same argument applies. This also implies that \mathcal{C} is refining in decremental graphs and coarsening in incremental graphs.

To maintain H , we maintain greedily for each $i \leq \sqrt{\log n}$, a maximal set A_i of *active* vertices at level i such that

1. every vertex in A_i is μ -light, and
2. every vertex $a \in A_i$ has $|B(a, (1/\epsilon')^{i+1})| \leq n^{(i+1)/\sqrt{\log n}}$ and $|B(a, (1/\epsilon')^i)| \geq n^{i/\sqrt{\log n}}$, and
3. any two vertices $a, a' \in A_i$ are at distance more than $(1/\epsilon')^i$ from each other.

We then maintain H as the union of the shortest path trees of vertices in A_i , for every i , truncated at depth $(1/\epsilon')^{i+1}$.

Since each vertex joins A_i at most once and leaves it again once, it is straightforward to see that each A_i can be maintained in time $O(n)$ using the information from the ES-trees from μ -light vertices. Also, the spanner can be maintained in total time $n^{1+o(1)}\mu$ since μ -light vertices have at most μ edges in their truncated shortest-path tree.

To see that for every shortest path $\pi_{s,t}$ in G with $\pi_{s,t} \cap (\bigcup_i C_i) = \emptyset$, we have

$$\mathbf{dist}_G(s, t) \leq \mathbf{dist}_H(s, t) \leq (1 + \epsilon') \mathbf{dist}_G(s, t) + O((1/\epsilon')^{\sqrt{\log n}})$$

we can construct $\pi_{s,t}$ by observing that s is a μ -light vertex and by the pigeonhole principle, there exists an index i such that $|B(s, (1/\epsilon')^{i+1})| \leq n^{1/\sqrt{\log n}} |B(s, (1/\epsilon')^i)|$. But then either s is in A_i (in which case we set $a' = s$), or there is another vertex $a' \in A_i$ at distance at most $(1/\epsilon')^i$ from s . And, from this vertex a' , there is a shortest path tree up to depth $(1/\epsilon')^{i+1}$ in H . If t is contained in the shortest path tree, then we are done since the path from s to a' to t is only of length $n^{o(1)}$. Otherwise, let s' be the vertex on $\pi_{s,t}$ that is farthest from s and still in the truncated shortest path tree of a' . Then, $\mathbf{dist}_G(a', s) = (1/\epsilon')^{i+1}$. Thus, by the triangle inequality, $\mathbf{dist}_G(s, s') \geq (1/\epsilon')^{i+1} - (1/\epsilon')^i$. Thus, the path from s to a' to s' in H is a $\frac{(1/\epsilon')^{i+1} + (1/\epsilon')^i}{(1/\epsilon')^{i+1} - (1/\epsilon')^i} \leq (1 + 4\epsilon')$ -multiplicative approximation of the path in G . We can then repeat for the path $\pi_{s',t}$. We thus derive multiplicative error of a $(1 + 4\epsilon') = (1 + \epsilon)$ on the entire path and an additive error of $n^{o(1)}$ induced by the last segment. To see that H never contains more than $n^{1+o(1)}$ edges, observe that active vertices a' in A_i , for any level i , have disjoint balls $B(a', (1/\epsilon')^{i+1})$ of size at least $n^{i/\sqrt{\log n}}$ and adds at most $n^{(i+1)/\sqrt{\log n}}$ edges to H . Thus, we can amortize the number of edges added over the number of vertices in its ball and obtain straight-forwardly that the total number of edges in H is at most $\sum_i n^{1+1/\sqrt{\log n}} = n^{1+o(1)}$. \square

Let us now give the description of our algorithm.

Algorithm. We maintain the data structure from Theorem 8.6 with $\mu = mn^{-\alpha}$. Our emulator consists of the graph H from Theorem 8.6 augmented with additional edges as described next.

Let us say a vertex $v \in V$ is \mathcal{C} -incident if there is a connected component $C_i \in \mathcal{C}$ with a vertex $w \in C_i$ such that there exists an edge $(v, w) \in E$. In particular, vertices in some C_i are \mathcal{C} -incident. Let \mathcal{C}_{inc} be the set of all vertices that are \mathcal{C} -incident. We maintain a connectivity data structure as described in [57, 87] on the graph $G[\mathcal{C}_{inc}]$ and obtain connected components $\mathcal{D} = \{D_1, D_2, \dots\}$. It is not hard to see that \mathcal{D} is refining in decremental graphs and coarsening in incremental graphs. Further, each connected component D_i in $G[\mathcal{C}_{inc}]$ consists of at least one connected component $C_i \in \mathcal{C}$ and its adjacent vertices from \mathcal{C}_{inc} , and each vertex of \mathcal{C}_{inc} belongs to exactly one connected component of \mathcal{D} .

Now consider some connected component D_i which contains components $C_{i_1}, C_{i_2}, \dots, C_{i_k}$ for some $k \geq 1$, and let $w_i = \sum_{j \leq k} 2(1/\epsilon)^{\sqrt{\log n}} \left\lceil \frac{|E(C_{i_j})|}{\mu} \right\rceil$. We maintain for each such component D_i a balanced binary spanning tree T_i over the vertices (this tree is not a subgraph of G but rather an arbitrary tree over the vertex set of a component), add the tree T_i to the spanner H , and give each edge in T_i weight $2^{\lceil \lg w_i \rceil}$ (that is the w_i rounded up to the nearest power of 2). This completes the description of the algorithm.

Emulator Approximation and Sparsity. For the approximation ratio, consider any shortest path $\pi_{s,t}$ in G . Since $H \setminus \bigcup_i T_i$ forms a subgraph of G and since each edge in T_i for some i has weight larger than the distance of its endpoints by Theorem 8.6, we have that the path is not underestimated in H .

Let G/\mathcal{D} be the graph G after contracting each D_i into a supernode. Observe that since the vertex sets of each D_i are mutually disjoint, there is a one-to-one correspondence between sets D_i and supernodes in G/\mathcal{D} . We also use D_i to denote the supernode corresponding to the connected component $D_i \in \mathcal{D}$.

Consider the shortest path $\pi'_{s,t}$ from s to t in G/\mathcal{D} . The path $\pi'_{s,t}$ clearly has smaller weight than $\pi_{s,t}$ in G . Let d_1, d_2, \dots, d_j be the supernodes on $\pi'_{s,t}$ (in the order that they appear on $\pi'_{s,t}$) that correspond to a connected component D_1, D_2, \dots, D_j in \mathcal{D} respectively.

Letting $d_0 = s$ and $d_{j+1} = t$, we have that every path segment $\pi'_{s,t}[d_i, d_{i+1}]$ for $0 \leq i \leq j$ is a $(1 + \epsilon, n^{o(1)})$ -approximate path segment even after remapping the first and last edge on the path segments again to a vertex in G (i.e. to the endpoint that the edges have in G instead of G/\mathcal{D}).

For each supernode d_i corresponding to a connected component D_i in \mathcal{D} , we have that the two endpoints in G that intersect between $\pi'_{s,t}$ and D_i are connected by a path in T_i consisting of $O(\log n)$ edges since T_i is balanced. If the supernode D_i contains connected components $C_{i_1}, C_{i_2}, \dots, C_{i_k}$, then each edge on this path has weight $w_i \leq 2 \cdot \sum_{j \leq k} 2(1/\epsilon)^{\sqrt{\log n}} \left\lceil \frac{|E(C_{i_j})|}{\mu} \right\rceil$.

We can now obtain an approximate path in H for $\pi_{s,t}$ by taking the union of the path segments and by adding the exposed paths in the in the trees spanning the connected components \mathcal{D} . Since there are at most m/μ supernodes by Theorem 8.6 and since each component $C_i \in \mathcal{C}$ is fully contained in exactly one supernode, we have that the total weight of the path from s to t in H is at most $(1 + \epsilon) \text{dist}_G(s, t) + \frac{m}{\mu} n^{o(1)} \leq (1 + \epsilon) \text{dist}_G(s, t) + n^{\alpha+o(1)}$. The sparsity of $n^{1+o(1)}$ follows by the bound on H in Theorem 8.6 and the sparsity of the trees T_i .

Running time. The running time of the connectivity data structure can be bound by $\tilde{O}(m)$. In decremental graphs, the time to maintain the balanced binary trees over the connected components \mathcal{D} takes time at most $\tilde{O}(m)$ as whenever a connected component splits into two connected components, we can remove the vertices of the smaller connected component from its tree T and repair T and then build a separate tree T' for the smaller component. Because each vertex that leaves T has degree at most 3, both of these operations can be performed in time $\tilde{O}(k)$, where k is the number of vertices of the smaller connected component. For incremental graphs, when we merge two connected components, we merge the corresponding trees by

destroying the smaller tree and adding each of its vertices to the larger tree.

The edge weights w_i can be maintained using the binary tree and since each edge weight in a component can only decrease (increase), and then decreases (increase) by a factor 2 the running time of these operations can be again bound by $\tilde{O}(m)$. Thus, the total update data is dominated by the data structure in Theorem 8.6 which runs in time $mn^{1-\alpha+o(1)}$, as desired.

A Randomized Algorithm for Spanners with High Additive Error. Recently, Bernstein et al. [20] presented a fully-dynamic algorithm, with amortized update time $\tilde{O}(1)$ that maintains an $(\tilde{O}(1), 0)$ -spanner of G of sparsity $\tilde{O}(n)$ that is randomized but works against an adaptive adversary. It is straight-forward to see that using Theorem 8.6 in conjunction with the above spanner, we can reuse the arguments from the approximation analysis of the emulator to derive Theorem 8.4.

9 Acknowledgments

The authors would like to acknowledge Ivan Mikhailin for the ideas he contributed to an initial variant of the conditional lower bounds. The authors would also like to thank Ofir Geri for reminding them of these initial lower bounds.

Monika Henzinger is supported by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement No. 340506 and by the Austrian Science Fund (FWF): Project No. 33775-N.

Maximilian Probst Gutenberg is supported by a start-up grant of Rasmus Kyng at ETH Zurich. This work was partially done while Maximilian Probst Gutenberg was at the University of Copenhagen where he was supported by Basic Algorithms Research Copenhagen (BARC), Thorup’s Investigator Grant from the Villum Foundation under Grant No. 16582, and partially done while visiting MIT, supported by STIBOFONDEN’s IT Travel Grant for PhD Students.

Virginia Vassilevska Williams is supported by an NSF CAREER Award, NSF Grants CCF-1528078, CCF-1514339 and CCF-1909429, a BSF Grant BSF:2012338, a Google Research Fellowship and a Sloan Research Fellowship.

Nicole Wein is supported by NSF Grant CCF-1514339.

References

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is valiant’s parser. *SIAM Journal on Computing*, 47(6):2527–2555, 2018.
- [2] Amir Abboud and Greg Bodwin. The $4/3$ additive spanner exponent is tight. *Journal of the ACM (JACM)*, 64(4):1–20, 2017.
- [3] Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. *SIAM Journal on Computing*, 47(6):2203–2236, 2018.
- [4] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 434–443. IEEE, 2014.
- [5] Ittai Abraham and Shiri Chechik. Dynamic decremental approximate distance oracles with $(1 + \epsilon, 2)$ stretch. *arXiv preprint arXiv:1307.1516*, 2013.

- [6] Ittai Abraham, Shiri Chechik, and Sebastian Krininger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 440–452. SIAM, 2017.
- [7] Ittai Abraham, Shiri Chechik, and Kunal Talwar. Fully dynamic all-pairs shortest paths: Breaking the $o(n)$ barrier. In Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain*, volume 28 of *LIPICs*, pages 1–16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014.
- [8] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.
- [9] Donald Aingworth, Chandra Chekuri, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). In Éva Tardos, editor, *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, 28-30 January 1996, Atlanta, Georgia, USA*, pages 547–553. ACM/SIAM, 1996.
- [10] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- [11] Giorgio Ausiello, Paolo G Franciosa, and Giuseppe F Italiano. Small stretch spanners on dynamic graphs. In *European Symposium on Algorithms*, pages 532–543. Springer, 2005.
- [12] Baruch Awerbuch. Complexity of network synchronization. *Journal of the ACM (JACM)*, 32(4):804–823, 1985.
- [13] Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Near-linear time construction of sparse neighborhood covers. *SIAM Journal on Computing*, 28(1):263–277, 1998.
- [14] Surender Baswana. Dynamic algorithms for graph spanners. In *European Symposium on Algorithms*, pages 76–87. Springer, 2006.
- [15] Surender Baswana, Ramesh Hariharan, and Sandeep Sen. Improved decremental algorithms for maintaining transitive closure and all-pairs shortest paths. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 117–123. ACM, 2002.
- [16] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. New constructions of (α, β) -spanners and purely additive spanners. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 672–681. Society for Industrial and Applied Mathematics, 2005.
- [17] Surender Baswana and Soumojit Sarkar. Fully dynamic algorithm for graph spanners with polylogarithmic update time. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1125–1134. Society for Industrial and Applied Mathematics, 2008.
- [18] Surender Baswana and Sandeep Sen. A simple linear time algorithm for computing a $(2k-1)$ -spanner of $o(n + 1/k)$ size in weighted graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 384–396. Springer, 2003.
- [19] Uri Ben-Levy and Merav Parter. New (α, β) spanners and hopsets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1695–1714. SIAM, 2020.

- [20] Aaron Bernstein, Jan van den Brand, Maximilian Probst Gutenberg, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford, and He Sun. Fully-dynamic graph sparsifiers against an adaptive adversary. *arXiv preprint arXiv:2004.08432*, 2020.
- [21] Aaron Bernstein and Shiri Chechik. Deterministic decremental single source shortest paths: beyond the $o(mn)$ bound. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 389–397. ACM, 2016.
- [22] Aaron Bernstein and Shiri Chechik. Deterministic partially dynamic single source shortest paths for sparse graphs. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 453–469. SIAM, 2017.
- [23] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1899–1918. SIAM, 2019.
- [24] Aaron Bernstein and Liam Roditty. Improved dynamic algorithms for maintaining approximate shortest paths under deletions. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1355–1365. Society for Industrial and Applied Mathematics, 2011.
- [25] Greg Bodwin and Sebastian Krinninger. Fully dynamic spanners with worst-case update time. In Piotr Sankowski and Christos D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPICs*, pages 17:1–17:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [26] Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 855–872. SIAM, 2016.
- [27] Gregory Bodwin and Virginia Vassilevska Williams. Very sparse additive spanners and emulators. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 377–382, 2015.
- [28] Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. *SIAM Journal on Discrete Mathematics*, 19(4):1029–1055, 2005.
- [29] Karl Bringmann, Nick Fischer, and Marvin Künnemann. A fine-grained analogue of schaefer’s theorem in $\exists k\forall$ -quantified first-order graph properties. In *34th Computational Complexity Conference*, pages 1–27. Schloss Dagstuhl, 2019.
- [30] Shiri Chechik. New additive spanners. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 498–512. SIAM, 2013.
- [31] Shiri Chechik. Near-optimal approximate decremental all pairs shortest paths. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 170–181. IEEE, 2018.
- [32] Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. Fast dynamic cuts, distances and effective resistances via vertex sparsifiers. *CoRR*, abs/2005.02368, 2020.

- [33] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. *CoRR*, abs/1910.08025, 2019.
- [34] Julia Chuzhoy and Sanjeev Khanna. A new algorithm for decremental single-source shortest paths with applications to vertex-capacitated flow and cut problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 389–400, New York, NY, USA, 2019. ACM.
- [35] Julia Chuzhoy and Thatchaphol Saranurak. Deterministic algorithms for decremental shortest paths via layered core decomposition. *To appear at SODA'21*, 2020.
- [36] Edith Cohen. Fast algorithms for constructing t -spanners and paths with stretch t . In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 648–658. IEEE, 1993.
- [37] Camil Demetrescu and Giuseppe F Italiano. A new approach to dynamic all pairs shortest paths. *Journal of the ACM (JACM)*, 51(6):968–992, 2004.
- [38] Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. *SIAM Journal on Computing*, 29(5):1740–1759, 2000.
- [39] Michael Elkin. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. *ACM Transactions on Algorithms (TALG)*, 7(2):1–17, 2011.
- [40] Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. *ACM Transactions on Algorithms (TALG)*, 15(1):1–29, 2018.
- [41] Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. *SIAM Journal on Computing*, 48(4):1436–1480, 2019.
- [42] Michael Elkin and Ofer Neiman. Linear-size hopsets with small hopbound, and constant-hopbound hopsets in RNC. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019*, pages 333–341. ACM, 2019.
- [43] Michael Elkin and Ofer Neiman. Near-additive spanners and near-exact hopsets, A unified view. *Bull. EATCS*, 130, 2020.
- [44] Michael Elkin and David Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM Journal on Computing*, 33(3):608–631, 2004.
- [45] Michael Elkin and Jian Zhang. Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models. *Distributed Computing*, 18(5):375–385, 2006.
- [46] Jacob Evald, Viktor Fredslund-Hansen, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Decremental asps in directed graphs versus an adaptive adversary. *arXiv preprint arXiv:2010.00937*, 2020.
- [47] Shimon Even and Yossi Shiloach. An on-line edge-deletion problem. *Journal of the ACM (JACM)*, 28(1):1–4, 1981.

- [48] Sebastian Forster and Gramoz Goranci. Dynamic low-stretch trees via dynamic low-diameter decompositions. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 377–388, 2019.
- [49] Sebastian Forster, Gramoz Goranci, and Monika Henzinger. Dynamic maintenance of low-stretch probabilistic tree embeddings with applications. *CoRR*, abs/2004.10319, 2020.
- [50] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Deterministic algorithms for decremental approximate shortest paths: Faster and simpler. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2522–2541. SIAM, 2020.
- [51] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Fully-dynamic all-pairs shortest paths: Improved worst-case time and space bounds. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2562–2574. SIAM, 2020.
- [52] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A subquadratic-time algorithm for decremental single-source shortest paths. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1053–1072. SIAM, 2014.
- [53] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $o(mn)$ barrier and derandomization. *SIAM Journal on Computing*, 45(3):947–1006, 2016.
- [54] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time maintenance of breadth-first spanning trees in partially dynamic networks. *ACM Trans. Algorithms*, 13(4):51:1–51:24, 2017.
- [55] 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):36:1–36:40, 2018.
- [56] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 21–30, 2015.
- [57] Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *Journal of the ACM (JACM)*, 48(4):723–760, 2001.
- [58] Shang-En Huang and Seth Pettie. Lower bounds on sparse spanners, emulators, and diameter-reducing shortcuts. In *16th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2018, June 18-20, 2018, Malmö, Sweden*, pages 26:1–26:12, 2018.
- [59] Shang-En Huang and Seth Pettie. Thorup–zwick emulators are universally optimal hopsets. *Information Processing Letters*, 142:9–13, 2019.
- [60] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 355–364. ACM, 2003.

- [61] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complex.*, 13(1-2):1–46, 2004.
- [62] Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 81–89. IEEE, 1999.
- [63] Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. In *Scandinavian Workshop on Algorithm Theory*, pages 277–281. Springer, 2014.
- [64] Mathias Bæk Tejs Knudsen. Additive spanners and distance oracles in quadratic time. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [65] Jakub Lacki, Jakub Ocwieja, Marcin Pilipczuk, Piotr Sankowski, and Anna Zych. The power of dynamic distance oracles: Efficient dynamic algorithms for the steiner tree. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 11–20. ACM, 2015.
- [66] François Le Gall. Powers of tensors and fast matrix multiplication. In *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014*, pages 296–303, 2014.
- [67] Francois Le Gall and Florent Urrutia. Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1029–1046, 2018.
- [68] Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1236–1252. SIAM, 2018.
- [69] Grigori A Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1):71–78, 1982.
- [70] David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989.
- [71] Seth Pettie. Low distortion spanners. *ACM Transactions on Algorithms (TALG)*, 6(1):1–22, 2009.
- [72] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 261–272. Springer, 2005.
- [73] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011.
- [74] Liam Roditty and Uri Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. *SIAM J. Comput.*, 41(3):670–683, 2012.
- [75] Piotr Sankowski. *Algebraic graph algorithms*. PhD thesis, Warsaw University, 2005.

- [76] Piotr Sankowski. Subquadratic algorithm for dynamic shortest distances. In *International Computing and Combinatorics Conference*, pages 461–470. Springer, 2005.
- [77] R. Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *Journal of Computer and System Sciences*, 51(3):400 – 403, 1995.
- [78] Jeremy P Spinrad. *Efficient graph representations*. American Mathematical Society, 2003.
- [79] Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 112–119, 2005.
- [80] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.
- [81] Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 802–809. ACM Press, 2006.
- [82] Jan van den Brand and Danupon Nanongkai. Dynamic approximate shortest paths and beyond: Subquadratic and worst-case update time. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 436–455. IEEE Computer Society, 2019.
- [83] Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith Winograd. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 887–898, 2012.
- [84] Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 645–654. IEEE, 2010.
- [85] Gerhard J Woeginger. Open problems around exact algorithms. *Discrete Applied Mathematics*, 156(3):397–405, 2008.
- [86] David P Woodruff. Additive spanners in nearly quadratic time. In *International Colloquium on Automata, Languages, and Programming*, pages 463–474. Springer, 2010.
- [87] Christian Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1757–1769. Society for Industrial and Applied Mathematics, 2013.