

Universal Optimality of Dijkstra via Beyond-Worst-Case Heaps*

Bernhard Haeupler
 INSAIT, Sofia University
 “St. Kliment Ohridski”
 & ETH Zurich

Richard Hladík
 INSAIT, Sofia University
 “St. Kliment Ohridski”
 & ETH Zurich

Václav Rozhoň
 INSAIT, Sofia University
 “St. Kliment Ohridski”

Robert E. Tarjan
 Princeton University

Jakub Tětek
 INSAIT, Sofia University
 “St. Kliment Ohridski”

Abstract

This paper proves that Dijkstra’s shortest-path algorithm is universally optimal in both its running time and number of comparisons when combined with a sufficiently efficient heap data structure.

Universal optimality is a powerful beyond-worst-case performance guarantee for graph algorithms that informally states that a single algorithm performs as well as possible for every single graph topology. We give the first application of this notion to any sequential algorithm.

We design a new heap data structure with a working-set property guaranteeing that the heap takes advantage of locality in heap operations. Our heap matches the optimal (worst-case) bounds of Fibonacci heaps but also provides the beyond-worst-case guarantee that the cost of extracting the minimum element is merely logarithmic in the number of elements inserted after it instead of logarithmic in the number of all elements in the heap. This makes the extraction of recently added elements cheaper.

We prove that our working-set property guarantees universal optimality for the problem of ordering vertices by their distance from the source vertex: The sequence of heap operations generated by any run of Dijkstra’s algorithm on a fixed graph possesses enough locality that one can couple the number of comparisons performed by any heap with our working-set bound to the minimum number of comparisons required to solve the distance ordering problem on this graph for a worst-case choice of arc lengths.

*Partially funded by the Ministry of Education and Science of Bulgaria’s support for INSAIT as part of the Bulgarian National Roadmap for Research Infrastructure. BH and RH were supported in part by the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (grant agreement No. 949272). VR was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (grant agreement No. 853109). RT was partially supported by a gift from Microsoft. JT and RH were supported by the VILLUM Foundation grants 54451 and 16582. JT worked on this paper while affiliated with BARC, University of Copenhagen. RH partially worked on this paper while visiting BARC, University of Copenhagen. VR worked on this paper while visiting MIT and thanks Mohsen Ghaffari and Christoph Grunau for many helpful discussions. RT worked on this paper while visiting the Simons Institute for the Theory of Computing and INSAIT.

Contents

1	Introduction	1
1.1	Beyond the Worst-Case: Universal Optimality	1
1.2	Our Results: Universally Optimal Dijkstra’s Algorithm via Heaps with the Working Set Property	2
1.3	Intuition & Techniques	4
1.4	Related Work	7
2	Preliminaries, Definitions and Our Model	9
2.1	Definitions	9
2.2	Instance & Universal Optimality	10
2.3	Our Model	11
3	Universally Optimal Dijkstra	12
3.1	Framework for Lower Bounding the Optimal Query Complexity	13
3.2	Constructing Barriers in the Exploration Tree	14
3.3	Finding a Good Interval Coloring	17
4	Constructing Heaps with the Working Set Property	19
5	Universally Optimal Dijkstra for Comparisons	25
5.1	Contractions Guided by the Dominator Tree	25
5.2	Algorithm Analysis	26
5.3	Lower Bound in the Number of Vertices	29
5.4	Lower Bound in the Number of Forward Edges	30
5.5	Finishing the Proof of Theorem 5.6	31
5.6	Constructing a Linearization from an SSSP Tree	32
6	Conclusion and Open Problems	34
A	Technical Aspects of Our Heap	38
A.1	Constant-Time Interval Maintenance	38
A.2	Maintaining the Minimum Subheap	38
B	Popular Implementations of Dijkstra Are Not Universally Optimal	40
C	Dijkstra’s Algorithm	42
D	Lower Bound on the Optimal Query Complexity in the Number of Forward Edges	43
E	Deferred Proofs	45

1 Introduction

Universal optimality is a powerful beyond-worst-case performance guarantee for graph algorithms that informally states that a single algorithm runs as fast as possible on every single graph topology.

This paper gives the first application of this notion to the standard sequential model of computation. We prove that Dijkstra’s algorithm, when combined with a sufficiently efficient heap data structure, is universally optimal for the natural problem of ordering nodes by their distance from the source.

Second, we design a sufficiently efficient heap necessary for this result. Our new heap is a strict improvement over Fibonacci heaps [FT87], guaranteeing a stronger beyond-worst-case property for DELETETMIN while keeping the complexities of INSERT and DECREASEKEY a constant.

Specifically, we identify that a natural *working-set property* for heaps previously proposed by Iacono [Iac00] is sufficient for guaranteeing the universal optimality of Dijkstra’s algorithm. Very roughly speaking, the working set property guarantees that the cost of extracting the minimum element from a heap is logarithmic in the number of elements inserted after that element, instead of logarithmic in the number of all elements in the heap (the precise definition is a bit stronger).

Our universal optimality result reveals a surprisingly clean interplay between this property and Dijkstra’s algorithm: Any heap with the working set property enables the algorithm to efficiently leverage every structural attribute of the graph it operates on, to the fullest extent that any comparison-based algorithm possibly can.

We also present a variant of Dijkstra’s algorithm which is universally optimal with respect to both the time complexity and the number of comparisons that are made. We note that while the time complexity of Dijkstra’s algorithm for any graph lies in the relatively narrow range of $\Omega(m+n)$ and $\mathcal{O}(m+n \log n)$, the number of comparisons needed to order the nodes by their distance can be as small as zero.

Beyond our application for Dijkstra’s algorithm, we hope that this paper opens doors for future research in the direction of applying (variants of) universal optimality for problems in the standard sequential model of computation.

1.1 Beyond the Worst-Case: Universal Optimality

The notion of asymptotic worst-case complexity is a key concept, achievement, and cornerstone of theoretical computer science and algorithm design. As the field matures, and our understanding refines, there are more and more problems for which the state-of-the-art algorithms essentially match the best performance achievable in the worst case. However, these guarantees may not be satisfactory – just because there exist some (family of) instances in which one cannot perform well does not mean one should be satisfied with performing equally badly on any easy instance as well.

The field of parameterized complexity sets the goal of finding some parameter $\theta(x)$ of the input x and gives algorithms that have good complexity with respect to both the input size $|x|$ and $\theta(x)$. In the best case, one may be able to get an algorithm that is optimal w.r.t. $|x|$ and $\theta(x)$ – that is, no correct algorithm would have better complexity under this parameterization.

Taking parameterized algorithms to the extreme, one may set $\theta(x) = x$, that is, parameterize by the instance itself. This results in a notion called *instance optimality* [FLN01] (see Sections 1.4 and 2.2 for more discussion). Any algorithm optimal under this parameterization is then at least as good as any correct algorithm on every single input.

Sadly, instance-optimal algorithms rarely exist. However, one commonly used alternative used in graph algorithms on weighted graphs is that of *universal optimality* [HWZ21]. In universal optimality, we want an algorithm A to be optimal if we parameterize by the graph G , but not the

weights w . That is, A is universally optimal if, for any graph G and any other correct algorithm A' , the worst-case running time of A on G across all possible weights is asymptotically no worse than the worst-case running time of A' across all possible weights.

On an intuitive level, a universally optimal algorithm is as fast as possible, on any given graph topology. For example, if there are fast algorithms for planar graphs, then the universally optimal algorithm must be fast when run on any planar graph. Planar graphs here can be replaced by any other subclass of graphs and even concrete graphs. We discuss this notion of optimality formally in the preliminaries in [Section 2.2](#).

1.2 Our Results: Universally Optimal Dijkstra’s Algorithm via Heaps with the Working Set Property

We will now discuss how we can make Dijkstra’s algorithm universally optimal when it uses a heap with the working set property – a certain beyond-worst-case property.

Universal optimality of Dijkstra’s algorithm Dijkstra’s algorithm solves a number of problems related to shortest paths: it computes the distances to all nodes from a source, it finds the shortest path tree, and it even returns an ordering of the nodes by their distance. It is folklore knowledge that the time complexity $\mathcal{O}(m + n \log n)$ of Dijkstra’s algorithm with Fibonacci’s heaps is the “best possible”. But to make this precise, we need to make two commitments: First, we need to work in an analog of the comparison model (see [Section 2.3](#) for the definition of the comparison-addition model). Second, we need to consider a problem that includes the task of ordering the nodes by their distance; only then we can argue that the term $\mathcal{O}(n \log n)$ in the time complexity cannot be improved because the distance ordering problem on a star graph is equivalent to sorting.

We take the same approach and throughout this paper, we consider the following problem.

Definition 1.1 (Distance Ordering problem (DO)). *Given a directed graph G with n nodes, a starting node s , and m edges with lengths, output all vertices of G in the increasing order of their distance from s .*

We prove in [Section 3](#) that Dijkstra’s algorithm is in fact universally optimal for the distance ordering problem if it uses our priority queue from [Definition 1.4](#) that has a certain beyond-worst-case property.

Theorem 1.2. *Dijkstra’s algorithm implemented with any Fibonacci-like priority queue with the working set property (as defined in [Definition 1.4](#)) is a universally optimal algorithm for the distance ordering problem in the comparison-addition model, in terms of running time. This holds both for directed and undirected graphs, and when compared against both deterministic and randomized algorithms.*

Consequently, Dijkstra’s algorithm is also universally optimal for the problems that require us to compute not just the node ordering, but the shortest path tree or the distances as well.

In [Section 5](#), we present a technically more involved algorithm that achieves universal optimality also in terms of the number of comparisons, not just with respect to the time complexity. We discuss this algorithm more in [Section 1.3](#).

Finally, in [Appendix B](#), we show that several natural implementations of a priority queue (e.g. standard Fibonacci heaps or splay trees) do not result in universally optimal Dijkstra’s algorithm.




Figure 1: An illustration of the working set property. Every interval represents an element: the white circle represents an insertion and the black circle a deletion. The set $W_{x,t}$ consists of the items whose associated intervals are highlighted in green and that are intersected by the dashed line.

Priority queues with the working set property One of our technical contributions is identifying the working set property for priority queues as the right property¹ that implies the universal optimality of Dijkstra’s algorithm. The definition of a working set is well-known in the literature of binary search trees [ST85] and for heaps, it first appeared in the work of Iacono [Iac00] that we compare ourselves with later in Section 1.4. The definition is presented next and illustrated in Figure 1.

Definition 1.3 (Working set [ST85; Iac00]). *Consider any priority queue supporting operations INSERT and DELETEMIN. For an element x present in the data structure, we define its working set as follows.*

For each time step t between the insertion and extraction of x , we define the working set of x at time t as the set of elements $W_{x,t}$ inserted after x and still present at time t . We include x itself in $W_{x,t}$.

We fix any time t_0 that maximizes the value of $|W_{x,t}|$; we call the set W_{x,t_0} simply the working set of x and denote it by W_x .

To give an example: For the first element x_1 inserted in the heap, the size of its working set is equal to the largest size of the heap between the insertion and extraction of x_1 . For any element x_2 inserted later, its working set size is defined similarly, but elements inserted before x_2 “do not count”.

We design a data structure that improves upon the Fibonacci heap by having the following working set property.

Definition 1.4 (Fibonacci-like priority queue with the working set property). *We say that a data structure is a Fibonacci-like priority queue with the working set property if it supports the following operations such that their amortized time complexity for any sequence of operations is as follows: INSERT, FINDMIN and DECREASEKEY in $\mathcal{O}(1)$, and DELETEMIN in $\mathcal{O}(1 + \log |W_x|)$, where W_x is the working set of the element being extracted.*

Theorem 1.5. *There is a Fibonacci-like priority queue with the working set property.*

We believe that Theorem 1.5 is of separate interest. We also show in Theorem 4.9 how one can further strengthen it by making its guarantees worst-case instead of amortized, with a $\log \log n$ loss in the complexity of some operations.

¹It seems natural that the “right” property of the priority queue should again be an instance optimality guarantee. However, the work of Munro, Peng, Wild, and Zhang [Mun+19] has shown that it is not possible to achieve the instance optimality guarantee for a natural class of priority queue implementations known as tournament heaps.




Figure 2: An example graph. Note that the basic Dijkstra’s algorithm needs $\Omega(n \log t)$ time to finish. However, for $t \ll n$, Dijkstra’s algorithm with the working set property finishes in time $\mathcal{O}(n + t \log t) = o(n \log t)$, which is optimal.

1.3 Intuition & Techniques

In this section, we first discuss an example graph G_1 from Figure 2 to facilitate intuition. Then, we summarize the proofs of our results.

Example Consider the graph G_1 from Figure 2 with number of leaves $t < n/2$. First, let us compute the time complexity of Dijkstra’s algorithm with the working set on that graph: Up to an additive $\mathcal{O}(n)$ term, the time complexity is dominated by the DELETEMIN operations. The working sets of nodes u_1, \dots, u_t, v_1 are of size $\mathcal{O}(t)$, while the working sets of nodes v_2, \dots, v_r for $r = n - t - 1$ have all size 1. Hence, the algorithm has complexity $\mathcal{O}(n + t \log t)$.

To see that this is the best possible for G_1 , notice that ordering nodes u_1, \dots, u_t by their distance from s is equivalent to the problem of sorting t numbers which requires $\Omega(t \log t)$ time in the comparison model. Moreover, any algorithm needs time $\Omega(m + n) = \Omega(n)$ to read the input.

Our main theorem Theorem 1.2 simply shows that for *all* graphs, the time complexity of Dijkstra’s algorithm with the working set property matches the best lower bound for that graph.

To appreciate this result, consider the following, very different, algorithm tailored to G_1 . The algorithm starts with the linear order $s < v_1 < \dots < v_{n-t-1}$ and repeatedly uses binary search to insert a node u_i into this order. Such an algorithm has complexity $\mathcal{O}(n + t \log n) = \mathcal{O}(n + t \log t)$, i.e., it is also optimal for the graph G_1 . Theorem 1.2 shows that Dijkstra’s algorithm with the working set property is competitive on every graph with every such instance-specific algorithm.

Universal Optimality: intuition behind Theorem 1.2 Up to an additive $\mathcal{O}(m + n)$ term, the time complexity of Dijkstra’s algorithm is dominated by the complexity of its DELETEMIN operations. Complexities of those operations are in turn governed by the sizes of appropriate working sets.

To show the optimality of Dijkstra’s algorithm with the working set on a given graph G , we carefully construct a lower bound distribution over weights on G that “hides” many instances of the sorting problem. More precisely, for an arbitrary graph G and its weights w , consider running Dijkstra’s algorithm on (G, w) . Recall that Dijkstra’s algorithm maintains a priority queue and any node $v \in V(G)$ enters and leaves this queue exactly once; at each point in time, the priority queue contains nodes that are on the “exploration boundary” of the algorithm and this exploration boundary separates the graph into two parts – the already seen and the yet unseen vertices.




Figure 3: Left: Finding the barrier W_u and then the barrier $W_{u'}$. The second barrier is either fully “before” or “after” W_u . In the first case, not considering W_u in the definition of a working set may decrease its size, i.e., we may have $|W_{u'}'| < |W_{u'}|$.

Right: To understand the decrease in working set sizes that results from removing u from consideration, we consider the nodes v_1, \dots, v_k whose working set size decreased (by one). We observe that first all vertices v_k, \dots, v_1, u were inserted in that order and only then they were extracted. This implies a lower bound $|W_{v_i}| \geq i + 1$ on working set sizes.

Consider the node u with the largest working set W_u over the whole run of the algorithm. Observe that for the time t for which $W_u = W_{u,t}$ we have that W_u is exactly the set of nodes in the exploration boundary of Dijkstra’s algorithm at time t (if there was some v on the boundary not in W_u , then $W_u \cup \{v\} \subseteq W_v$, contradicting the maximality of W_u).

The crucial observation is that we can locally change the weights w at edges ingoing to W_u to get a distribution over weights \mathbf{w}_u with the following property: If we order nodes in W_u based on their distance from s under weights \mathbf{w}_u , their order is a uniformly random permutation. That is, the instance (G, \mathbf{w}_u) “embeds” a sorting problem that requires time $\Omega(|W_u| \log |W_u|)$ to solve. By the maximality of $|W_u|$ we know, on the other hand, that our algorithm spends overall time of at most $\mathcal{O}(|W_u| \log |W_u|)$ on DELETETEMIN operations for the nodes in W_u .

That is, with \mathbf{w}_u we managed to “pay for” the nodes in W_u : the time spent by Dijkstra’s algorithm on them in (G, w) is up to a constant the same as the time *any* correct algorithm needs to spend on them on average in an instance sampled from (G, \mathbf{w}_u) .

The rest of the proof of [Theorem 1.2](#) reduces to arguing that we can continue this process until the whole graph is separated into what we call a *barrier sequence*. This way, we ultimately construct a lower bound distribution \mathbf{w} that pays for all the DELETETEMIN operations.

Let us explain how this can be done by showing how the second barrier is selected (see [Figure 3](#)). After the barrier W_u was selected, we continue by finding the node $u' \in V(G) \setminus W_u$ with the largest working set $W_{u'}' = W_{u',t'}$ where we define W_v' as the working set of v if we do not count the nodes of W_u in the definition of a working set.

There are now two cases based on whether $t' > t$ or $t' < t$. Importantly, in the first case, all nodes of $W_{u'}'$ come “after” the nodes of W_u , while in the other case, all nodes come “before”. This property in the end allows us to construct a lower-bound distribution that embeds the sorting lower bound for both W_u and $W_{u'}'$.

In the case $t' < t$ we encounter an additional issue: The set $W_{u'}'$ can potentially be much smaller than the set $W_{u'}$, thus the lower bound of $\Omega(|W_{u'}'| \log |W_{u'}'|)$ may not be able to pay for the corresponding DELETETEMIN cost of $\mathcal{O}(\sum_{v \in W_{u'}'} \log |W_v|)$. To solve this issue, we upper bound

the overall decrease in working set sizes caused by taking W_u out of the picture. In particular, we analyze the value

$$\sum_{v \in V(G) \setminus W_u} \log |W_v| - \log |W'_v|. \quad (1)$$

How much can working set sizes change if we remove just any one node of W_u instead of removing the whole set at once? The size of each working set can decrease at most by one, but it can decrease by one for many nodes inserted before u was. Let us call the nodes whose working set decreased as v_1, \dots, v_k and let us sort them according to their insertion time from the largest to the smallest (Figure 3). We note that each v_i was extracted only after the insertion of u , otherwise, the size of its working set would not decrease. This implies that the working set of each v_i has a size of at least $|W_{v_i}| \geq i + 1$ since this is its size at the time of insertion of u . In the worst case, the removal of u causes that $\log |W_{v_1}|$ changes from $\log 2$ to $\log 1$, $\log |W_{v_2}|$ changes from $\log 3$ to $\log 2$ and so on. Observing a telescoping sum, we note that the overall decrease of working set sizes in the sense of Equation (1) is at most $\log |W_{v_k}|$ which is at most $\log |W_u|$ by our assumption on the maximality of u .

Recall that the above analysis was done in the case when an arbitrary node of W_u was removed. Hence, after removing *all* nodes of W_u one by one, we can upper bound the expression of Equation (1) by $\mathcal{O}(|W_u| \log |W_u|)$. This is very fortunate: We can make the first barrier W_u pay for all the future gaps between lower bounds and upper bounds caused by the mismatch between W and W' .

We can now continue our process until the very end when the whole graph G is decomposed into a sequence of barriers. We can then straightforwardly define a lower bound distribution \mathbf{w} that embeds a sorting lower bound into each barrier.

The working set property: intuition behind Theorem 1.5 Next, we discuss our construction of a heap with the working set property. Our construction is a general black-box construction that uses an arbitrary priority queue (think Fibonacci heap) to create a new heap with the same guarantees that additionally satisfies the working set property of Theorem 1.5.

The main idea behind the construction is straightforward and present in the literature [Iac01; Bäd+07; EFI12]: we keep a list of heaps where the i -th heap has size at most 2^i . Whenever we insert a new element, we insert it into the smallest heap; we will keep it as an invariant that larger heaps contain older elements than smaller heaps. Whenever a heap gets too large, we recursively merge it with its larger neighboring heap as follows.

In general, consider the i -th heap H_i . During merging, a new carry heap C arrives from the merge at position $i - 1$ and we check whether the total number of elements in H_i and C together is at most 2^i . If so, we merge the two heaps together and finish the merging procedure. Else, C replaces H_i as the new i -th heap, while H_i is the new carry heap for the next, $(i + 1)$ -th, level.

Notice that whenever the heap H_i gets “upgraded” to the next position $i + 1$, we know there are at least 2^{i-1} elements in the original carry heap C that just replaced H_i and the smaller heap H_{i-1} together. Otherwise, merging at the i -th level would not have been triggered. This implies that the working set of every element newly upgraded to position $i + 1$ has a size at least 2^{i-1} .

This observation inductively implies that every element in each H_i has the working set of size $\Omega(2^i)$. On the other hand, since H_i has at most 2^i elements, the cost of extraction from H_i is $\mathcal{O}(1 + \log 2^i)$. We can thus write the extraction cost as $\mathcal{O}(1 + \log |W_x|)$, as needed.

Notice that it is not necessary for the i -th heap to be of size roughly 2^i , its size can be up to

2^{2^i} to make our argument work, and we also choose this doubly exponential approach in the formal construction in [Section 4](#).

There are a few more technicalities that need to be addressed. For example, to extract the minimum element, we first need to find out which one of our heaps contains it. As we have at most $\mathcal{O}(\log \log n)$ heaps if we choose their sizes to grow doubly exponentially, we can maintain their minima in a separate heap which results in an additive $\mathcal{O}(\log \log \log n)$ loss. While this type of loss is already very small, we show that with careful data structures, we can make the overhead in our construction negligible and recover the same asymptotic guarantees as the guarantees of the original heap.

Extension to comparisons The most technically challenging part of this paper is the extension of our main result, [Theorem 1.2](#), to comparison complexity which measures how many answers to comparison questions our algorithm needs (see [Section 2.3](#) for details). For example, if the graph G is a path and s is its leaf, the comparison complexity of the distance ordering problem is, in fact, 0, since we do not need to access the weights of the graph at all to figure out how the nodes have to be ordered!

We show in [Section 5](#) that there is an algorithm that is universally optimal for both measures of complexity at once. Dijkstra’s algorithm with the working set property itself is *not* such an algorithm: its comparison complexity on [Figure 2](#) is $\Omega(t \log t + n)$, while the optimal comparison complexity is only $\mathcal{O}(t \log n)$. However, looking at our analysis of Dijkstra’s algorithm, we can notice that it “almost” is universally optimal with respect to comparison complexity – the main problem is that it can happen that in our lower bound construction, we may embed a sorting problem on *exactly* one node. The comparison complexity of sorting one number is then 0, while the algorithm uses $\Theta(1)$ operations, which creates a mismatch between the upper and lower bounds.

This motivates our final algorithm that uses Dijkstra with the working set property as a sub-routine. Concretely, in our final algorithm, we first carefully contract the input graph using 0 comparison queries. Then, we solve the SSSP problem on that contracted graph by Dijkstra with the working set property and uncontract the found shortest path tree to the original graph.

This way, we solve the SSSP problem, but not the distance ordering problem. Fortunately, for trees, there is also a very different universally optimal algorithm for distance ordering based on dynamic programming. We finish by applying this other algorithm to the found SSSP tree.

How exactly should we contract our graph? If it is undirected, it would be enough to contract bridges, but when the graph is directed, we need to find a more complex structure known as the *dominator tree*.

1.4 Related Work

Shortest path Dijkstra’s algorithm [[Dij59](#)] is a foundational algorithm solving variants of the shortest path problems; using a standard implementation of a heap, its worst-case running time is $\mathcal{O}((m + n) \log n)$. A classical result of Fredman and Tarjan [[FT87](#)] shows that shortest paths can be found in $\mathcal{O}(m + n \log n)$ by Dijkstra’s algorithm using Fibonacci heaps. This running time is optimal if we require the algorithm to output the nodes in the sorted order. If we require only to find the distances and not the ordering, a recent breakthrough by Duan, Mao, Shu, and Yin [[Dua+23](#)] solves the shortest path problem in expected time complexity $\mathcal{O}(m\sqrt{\log n \log \log n})$ for undirected graphs.

If we assume that the weights are integers, or the ratio between maximum and minimum weight is bounded, there are many more algorithms for the shortest path problem with time complexity

getting very close to linear [FW93; FW94; Tho00a; Ram96; Ram97; Tho00b; Hag00; Tho99; Tho04].

Universal optimality The term “universal optimality” was coined by researchers in the area of distributed algorithms [GKP98; HRG22; GZ22]. More precisely, the researchers developed techniques for designing distributed algorithms that are close to being universally optimal for many problems. In particular, there are such distributed algorithms for the approximate shortest path problem [HWZ21; Zuz+22; Roz+22]. However, as far as we know, our use of universal optimality outside of distributed algorithms is new.

Instance optimality Our notion of universal optimality is closely related to *instance optimality* [FLN01; VV17; ABC17] (see Section 2.2 for a precise definition) which is a golden standard for the beyond-worst-case analysis. Indeed, in the influential book by Roughgarden [Rou21] containing 30 chapters discussing various approaches to the beyond-worst-case analysis, instance optimality is covered in the third chapter, only after Introduction and Parameterized Algorithms.

The major problem of instance optimality is that it is an extremely strong notion and thus it is very rarely possible to satisfy it. For example, instance optimality is not achievable for our distance ordering problem because an algorithm can guess the linear ordering and the shortest path tree and check whether the guess is correct in linear time (see Section 2 for more discussion). Consequently, it is usual that results from the area of instance optimality both restrict the computational model (as otherwise, it is hard to get lower bounds) and also relax the notion of instance optimality itself (to disallow silly obstructions to it).

Let us give a concrete example. Afshani, Barbay, and Chan [ABC17] develop algorithms for a number of geometrical problems, including the 2D convex hull problem. The authors work in a computational model which is a suitable generalization of the standard comparison model. Moreover, they achieve the guarantee of order-oblivious instance optimality – a relaxed notion of instance optimality. Notice that we can view the convex hull problem as an extension of the standard sorting problem (convex hull sorts via the map $x \rightarrow (x, x^2)$). This is analogous to how our distance ordering problem can be viewed as an extension of sorting (distance ordering sorts via the star graph input). This suggests that problems extending sorting may form an interesting class of problems amenable to instance optimality analysis.

Beyond-worst-case analysis of data structures Perhaps the most famous open question from the area of data structure design – the dynamic optimality conjecture for splay trees [ST85] – is another example of a question from the area of instance optimality. It was shown by Munro, Peng, Wild, and Zhang [Mun+19] that an analog of the dynamic optimality conjecture fails for a large class of heaps known as tournament heaps.

Working set property is among a host of beyond-worst-case properties that are famously attained by splay trees [ST85]. There are analogs of the working set property for heaps and priority queue implementations that achieve them or that achieve different beyond worst-case guarantees [BHM13; EFI13; KS18]. However, none of those implementations seem to suffice for our application.

inline]Add more log log n structures by Iacono etc., even though they are search trees

Iacono [Iac00] showed that our working set property from Definition 1.3 is attained by pairing heaps, a data structure known for excellent theoretical and practical guarantees [Fre+86; Pet05; IÖ14; KS18; ST23]. However, his analysis does not include the DECREASEKEY operation. Elmasry [Elm06] presented an implementation of a heap that achieves our working set property (in fact, a stronger one) but does not implement the DECREASEKEY operation. Elmasry, Farzan, and Iacono

[EFI12] presented an implementation of a heap that again achieves a stronger version of the working set property, but their implementation of DECREASEKEY does not have amortized constant time.

We note that for our purposes, we need a somewhat special variant of the working set property: we need that DECREASEKEY operations are not considered in the definition of working set, but the data structure has to support these operations.

Sorting under partial information After publicizing a pre-print version of this paper, the techniques developed here have been applied in the work of Haeupler, Hladík, Iacono, Rozhoň, Tarjan, and Tětek [Hae+24] to solve the problem of sorting under partial information. The problem, dating back to Fredman [Fre76], is to sort a set of items having an unknown total order by doing binary comparisons of the items, given a list of some already-performed comparisons. This list can be represented as a directed acyclic graph G on the items, and the problem reformulated as that of finding a fixed unknown topological order of G by performing vertex comparisons. The algorithm given in [Hae+24] performs $\mathcal{O}(\log T)$ comparisons and has a running time of $\mathcal{O}(n + m + \log T)$ where n , m and T is the number of items, the number of pre-existing comparisons and the number of topological orders of G , respectively. Both of these bounds are optimal up to constant factors. The algorithm, called *topological heapsort*, combines topological sort with heapsort, using a heap with the working set property to achieve optimality.

There are many parallels between this work and [Hae+24]. Not only can topological heapsort be viewed as a variation on Dijkstra’s algorithm, but the algorithm is also universally optimal, in the sense that for any fixed G , it minimizes the running time and number of comparisons (up to constants) on the worst-case total order consistent with G . Most importantly, the authors prove the universal optimality of topological heapsort using the framework that we develop in this work.

2 Preliminaries, Definitions and Our Model

We now formally state some definitions that we need. For the sake of completeness, we also restate Dijkstra’s algorithm in [Appendix C](#).

2.1 Definitions

All logarithms are binary. We use the convention $x \log x = 0$ for $x = 0$. All undirected graphs in this paper are connected and in directed graphs, all nodes are reachable from the source. We assume all trees are rooted at s . In directed graphs, we assume edges are oriented away from the root. For a rooted tree T and a vertex v , we use the notation $T(v)$ to refer to the subtree of T rooted at v . For two lists A, B , we use the notation $A + B$ to denote their concatenation. We use \emptyset to denote a null value or a null pointer.

Given a (directed or undirected) graph G , the weights w are a function $E(G) \rightarrow \mathbb{R}_{>0}$; in particular, we disallow zero weights. This function then induces the *distances* $d : V(G) \times V(G) \rightarrow \mathbb{R}_{>0} \cup \{+\infty\}$. In shortest path problems, we always assume that the input graph comes with a special node that we call the *source* and denote s . We slightly misuse notation by usually not including the parameterization by s in the definition of our objects, i.e., we write $T_{SSSP}^{G,w}$ instead of $T_{SSSP}^{G,w,s}$ to simplify notation.

Definition 2.1 (SSSP tree, spanning tree, and the exploration tree). *Let us have a fixed (directed or undirected) graph G and a source vertex s .*

- In undirected graphs, a spanning tree is any tree $T \subseteq G$ with $V(T) = V(G)$. In directed graphs, a spanning tree is any tree $T \subseteq G$ with $V(T) = V(G)$ rooted at s where edges go in the direction away from s .
- An SSSP tree $T_{SSSP}^{G,w} \subseteq G$ is a spanning tree that satisfies that for any vertex v , the sv -path in $T_{SSSP}^{G,w}$ is shortest w.r.t. the weights w .
- Given an execution of Dijkstra's algorithm [Algorithm 9](#) on G with weights w , an exploration tree $T_{explore}^{G,w}$ is the tree where the parent of each vertex v is the vertex during whose exploration v was added to the heap. That is, the parent of each $v \neq s$ is the vertex u that minimizes $d(s, u)$ among vertices with $uv \in E(G)$.

The following definition of linearizations is crucial since given any graph G , the set of its linearization is the set of possible outputs to our distance ordering problem [Definition 1.1](#).

Definition 2.2 (Linearizations). *For any tree T , a linearization is a complete linear order of $V(T)$ such that for every edge $uv \in E(T)$, we have $u < v$ in the linear order.*

For any graph G , we say that L is a linearization of G if there is a spanning tree of G such that L is a linearization of T .

For a graph G , we use $\text{Linearizations}(G)$ for the number of different linearizations of G .

For a linearization L and a vertex v , we denote by $L(v) \in \{1, \dots, n\}$ the position of v in the linearization. We write $u \prec_L v$ if and only if $L(u) < L(v)$.

The following claim shows that each linearization is indeed a possible solution to the distance ordering problem for some weights.

Claim 2.3 (Equivalent definition of a linearization). *For any graph G , L is a linearization of G if and only if there exist nonnegative weights w such that*

1. *For every two nodes $u \neq v$ we have $d_w(s, u) \neq d_w(s, v)$,*
2. *$u \prec_L v$ if and only if $d_w(s, u) < d_w(s, v)$.*

Proof. See [Appendix E](#). □

We will also need the following definition of forward edges.

Definition 2.4 (Distance-Forward edges). *Given a directed graph G with source s and weights w such that for every e , $w(e) > 0$, we define $F_{G,w}$ to be the set of edges uv of G such that $d(s, u) < d(s, v)$.*

Analogously, and to make some claims about directed graphs also hold for undirected graphs, for undirected G , we define $F_{G,w}$ as the set of all edges uv of G such that $d(s, u) \neq d(s, v)$.

2.2 Instance & Universal Optimality

Universal optimality is a powerful beyond-worst-case notion that can be applied to any problem on weighted graphs. It is best thought of as a meaningful intermediate notion of optimality between the standard notion of existential worst-case optimality and instance optimality, which in most cases is too much to ask for and provably impossible.

To introduce the formal definition of universal optimality, we first need some notation. An instance of a graph problem consists of a graph topology G and edge or node weights, lengths, capacities, etc. that we denote by w . We say a (possibly randomized) algorithm A is correct for a problem if it outputs a correct solution for any instance (G, w) with probability 1. Denote the

(expected) running time of an algorithm A on instance (G, w) as $A(G, w)$. We also use $\mathcal{G}_{n,m}$ to denote the set of topologies with n nodes and m edges, \mathcal{W}_G to denote the set of all possible weights for a given topology G (or lengths, capacities, \dots). Finally, \mathcal{A} stands for the set of correct algorithms (defined above). We can also restrict \mathcal{A} to e.g. algorithms that work in the comparison-addition model, i.e., that only compare and add weights in w .

A correct algorithm $A \in \mathcal{A}$ is *existentially optimal* (up to a factor of $\alpha(n, m)$) if

$$\forall n, m : \sup_{\substack{G \in \mathcal{G}_{n,m} \\ w \in \mathcal{W}_G}} A(G, w) \leq \alpha(n, m) \inf_{A_{n,m}^* \in \mathcal{A}} \sup_{\substack{G \in \mathcal{G}_{n,m} \\ w \in \mathcal{W}_G}} A_{n,m}^*(G, w).$$

That is, this equation corresponds to being optimal with respect to the classical worst-case complexity, parameterized by n and m .

A correct algorithm $A \in \mathcal{A}$ is *instance optimal* (up to a factor of $\alpha(n, m)$) if

$$\forall n, m, \forall G \in \mathcal{G}_{n,m}, \forall w \in \mathcal{W}_G \quad A(G, w) \leq \alpha(n, m) \inf_{A_{G,w}^* \in \mathcal{A}} A_{G,w}^*(G, w).$$

That is, an instance optimal algorithm is competitive with any correct algorithm on any possible input. Note that this is an extremely strong guarantee: A is supposed to be competitive on any input (G_0, w_0) with algorithms of the type “check that $(G, w) = (G_0, w_0)$; if so, return the (hardcoded) answer for (G_0, w_0) , otherwise run any slow algorithm”. This is the reason why instance optimality is an extremely rare property for an algorithm to have; one can check that this type of “instance-tailored” algorithm implies that instance optimality is impossible for SSSP-like problems in classical computation models.

Finally, a correct algorithm $A \in \mathcal{A}$ is *universally optimal* (up to a factor of $\alpha(n, m)$) if

$$\forall n, m, \forall G \in \mathcal{G}_{n,m} : \sup_{w \in \mathcal{W}_G} A(G, w) \leq \alpha(n, m) \inf_{A_G^* \in \mathcal{A}} \sup_{w \in \mathcal{W}_G} A_G^*(G, w).$$

Note that the universal optimality lies between the worst-case notion of existential optimality and the instance optimality. This expands the class of problems for which we can apply the framework, while we still retain very strong guarantees on the performance of the algorithm.

In this work, we focus solely on (existential, instance-, universal) optimality up to a constant factor, i.e., $\alpha(n, m) = \mathcal{O}(1)$.

2.3 Our Model

In this paper, we work in the standard Word-RAM model. Moreover, we treat edge weights as black boxes that our algorithms can only add and compare. This model is also known as the comparison-addition model [PR02].

Definition 2.5 (Comparison-addition model). *The comparison-addition model implements all operations possible in Word-RAM and additionally, implements operations on an associative array of protected cells storing the edge weights.*

A cell is protected if it is stored in the associative array that contains all edge weights at the beginning of the algorithm, or if it has been created by an ADD operation defined below. Only the following operations can be executed on protected cells.

1. $\text{ADD}(x, y)$ for two protected memory cells x, y . It returns a protected cell containing $x + y$.
2. $\text{COMPARE}(x, y)$ for two protected memory cells x, y . It returns $-1, 0$ or 1 , if $x < y, x = y$ and $x > y$ respectively.

Definition 2.6 (Time and query complexity). For an algorithm A on an input x , we define the time complexity to be the random variable equal to the total number of operations performed. We define the query complexity to be the random variable equal to the total number of COMPARE operations performed.

That is, query complexity measures how many bits of information we are getting by comparing various linear combinations of edge weights.

Optimality For a given graph G , let $\text{OPT}_Q(G)$ be the number of comparison queries of an optimal (correct) algorithm for this graph. Formally, let \mathcal{A} denote the set of all correct algorithms for the distance ordering problem. Then we define $\text{OPT}_Q(G) = \inf_{A \in \mathcal{A}} \sup_{w > 0} A_Q(G, w)$, where $A_Q(G, w)$ denotes the (expected) number of comparison queries issued by A on input (G, w) . Analogously, let $\text{OPT}_T(G) = \inf_{A \in \mathcal{A}} \sup_{w > 0} A_T(G, w)$, where $A_T(G, w)$ is the (expected) number of operations performed by A on input (G, w) . Since a comparison query takes unit time in our model, we always have $\text{OPT}_Q(G) \leq \text{OPT}_T(G)$.

3 Universally Optimal Dijkstra

In this section, we prove [Theorem 1.2](#), i.e., that Dijkstra’s algorithm, equipped with a Fibonacci-like priority queue with the working set property ([Definition 1.4](#)), is universally optimal with respect to the time complexity. In fact, we prove a slightly stronger statement, [Theorem 3.1](#): the algorithm is universally optimal up to an additive $\mathcal{O}(n + \max_{w \in \mathcal{W}_G} |F_{G,w}|)$ factor also with respect to the number of comparisons also. This is later needed in [Section 5](#) to construct the universally optimal algorithm with respect to both the time complexity and the number of comparisons, a construction proving the more refined result of [Theorem 5.1](#).

Throughout this section, we use the notation $\text{OPT}_Q(G)$ and $\text{OPT}_T(G)$ to denote the number of comparison queries and time performed by a universally optimal algorithm. See [Section 2.3](#) for details.

Theorem 3.1. *In the comparison-addition model of [Section 2](#), Dijkstra’s algorithm ([Algorithm 9](#)) implemented using any Fibonacci-like priority queue with the working set property has time complexity $\mathcal{O}(\text{OPT}_Q(G) + m + n)$.*

Moreover, the number of comparison queries made by the algorithm is $\max_{w \in \mathcal{W}_G} |F_{G,w}|$.

We shall prove the theorem in the rest of this section but first, we note that it implies universal optimality with respect to time complexity, since $\text{OPT}_Q(G) + m + n = \mathcal{O}(\text{OPT}_T(G))$. The proof of this fact and of [Theorem 1.2](#) is given in [Appendix E](#).

We proceed with the proof of [Theorem 3.1](#). Let us fix any (directed or undirected) graph G with n vertices and m edges. The bulk of the proof is to understand the cost of DELETETEMIN operations. Formally, we will prove the following proposition.

Proposition 3.2. *Consider all calls of DELETETEMIN performed by Dijkstra’s algorithm ([Algorithm 9](#)) with any Fibonacci-like priority queue with the working set property. The total amortized time complexity of these operations is $\mathcal{O}(\text{OPT}_Q(G) + n)$.*




Figure 4: An example: tree T with a barrier sequence B_1, B_2, B_3, B_4 . Note that the sequence B_2, B_1, B_3, B_4 is also a barrier sequence, as is B_2, B_3, B_4, B_1 , but B_1, B_3, B_2, B_4 is not.

Before we prove [Proposition 3.2](#), let us see why this is sufficient to prove [Theorem 3.1](#). The full proof is deferred to [Appendix E](#).

Proof sketch of [Theorem 3.1](#). We can focus on comparison complexity only, as it dominates the time complexity. Since INSERT and DECREASEKEY cost $\mathcal{O}(1)$ comparisons per operation, there are n INSERT calls, and all DELETETEMIN calls cost $\mathcal{O}(\text{OPT}_Q(G) + n)$ time in total (and thus also at most this many comparisons). Moreover, there are at most $\max_w |F_{G,w}|$ INSERT and DECREASEKEY calls as each call corresponds to a forward edge. \square

3.1 Framework for Lower Bounding the Optimal Query Complexity

The rest of the section proves that the overall time complexity of DELETETEMIN operations is bounded by $\mathcal{O}(\text{OPT}_Q(G) + n)$, i.e., we prove [Proposition 3.2](#).

We will start by developing a simple framework for lower bounding the value of $\text{OPT}_Q(G)$. First, we use the standard entropy argument to show that since each comparison query recovers only $\log 3$ bits of information, the number of comparisons any algorithm needs to make is at least the logarithm of the number of possible outputs of the algorithm.

Lemma 3.3. *For any directed or undirected graph G , any (even randomized) algorithm for the distance ordering problem needs $\Omega(\log(\text{Linearizations}(G)))$ comparison queries in expectation for some weights, i.e., $\text{OPT}_Q(G) = \Omega(\log(\text{Linearizations}(G)))$.*

Proof sketch. See [Appendix E](#) for the full proof. For deterministic algorithms, this is the classical argument that with q comparison, the algorithm can end up in at most 3^q different states and thus answer only 3^q different answers. In the full proof, we 1) show we can ensure that every input has only one valid output (i.e., $d(s, u) \neq d(s, v)$ for all $u \neq v$), which makes the previous argument valid, and 2) use Shannon's source coding theorem instead, which allows us to analyze randomized algorithms as well. \square

To get a more concrete lower bound that can be compared with the actual cost of Dijkstra's algorithm, we will now show how we can efficiently lower bound the number of linearizations using the structure that we call a barrier sequence. See [Figure 4](#) for the intuition.

Definition 3.4 (Barriers and Barrier sequences). *Given a tree T rooted at s , we define:*

1. a barrier $B \subseteq V(T)$ is an incomparable set of nodes of T ; that is, there are no two $u, v \in B$ such that u is the ancestor of v in T ,

2. for two barriers B_1, B_2 with $B_1 \cap B_2 = \emptyset$, we say that $B_1 \prec B_2$ if no node of B_2 is a predecessor of a node of B_1 ,
3. We say that (B_1, B_2, \dots, B_k) is a barrier sequence if whenever $i < j$, we have $B_i \prec B_j$.

Claim 3.5. A sequence (B_1, \dots, B_k) of pairwise-disjoint vertex sets is a barrier sequence if and only if for all $1 \leq i \leq j \leq k$ and all $u \in B_i$ and $v \in B_j$, v is not an ancestor of u in T .

Proof. This is just a condensed version of the above definition.

Given such a sequence, plugging in $i = j$ immediately gives us that all B_i are barriers, and plugging in $i < j$ and using the definition of \prec gives us that $B_i \prec B_j$ and thus it is a barrier sequence.

On the other hand, given a barrier sequence, consider any $1 \leq i \leq j \leq k$. If $i = j$, the ancestry claim holds because B_i is a barrier, and if $i < j$, it holds because $B_i \prec B_j$. \square

Intuitively, whenever we find a sequence of barriers (B_1, B_2, \dots, B_k) , it means that we can set up a distribution over weights w such that any distance ordering algorithm is forced to solve a sequence of sorting problems of sizes $|B_1|, |B_2|, \dots, |B_k|$ where the i -th sorting problem requires $\Omega(|B_i| \log |B_i|)$ comparisons, by the standard sorting lower bound in the comparison model. This is formalized next.

Lemma 3.6 (Barriers give lower bounds). *Let T be any spanning tree of G and (B_1, \dots, B_k) be a barrier sequence for T . Then,*

$$\log(\text{Linearizations}(G)) \geq \log(\text{Linearizations}(T)) = \Omega\left(\sum_{i=1}^k |B_i| \log |B_i|\right).$$

Proof. The first inequality simply follows from the definition of linearization ([Definition 2.2](#)).

Denote by \mathcal{L} the set of linearizations of T . We will prove by induction on (n, k) that for any T with n vertices and a barrier sequence (B_1, \dots, B_k) , we have $|\mathcal{L}| \geq |B_1|! \cdots |B_k|!$. This proves the lemma, as then $\log |\mathcal{L}| \geq \sum_i \log(|B_i|!) = \Omega(\sum_i |B_i| \log |B_i|)$.

If $n = 0$ or $k = 0$, the claim holds trivially. Otherwise, consider a tree T and a barrier sequence (B_1, \dots, B_k) . Delete from T all vertices in B_k and all their descendants, obtaining a tree T' . The sequence (B_1, \dots, B_{k-1}) is now a valid barrier sequence of T' . By induction, the set \mathcal{L}' of all linearizations of T' satisfies $|\mathcal{L}'| \geq \prod_{i=1}^{k-1} |B_i|!$.

For each $v \in B_k$, let L_v denote an arbitrary fixed linearization of $T(v)$, e.g., the DFS pre-order. For any of the $|B_k|!$ possible orderings σ of B_k , let L_σ denote the sequence obtained by concatenating all L_v (for $v \in B_k$) together, taken by the order given by σ . Let \mathcal{L}_σ be the set of all $|B_k|!$ such sequences. We claim that for any $L' \in \mathcal{L}'$ and $L_\sigma \in \mathcal{L}_\sigma$, the concatenation $L := L' + L_\sigma$ is a unique linearization of T . If this is true, then clearly $|\mathcal{L}| \geq |\mathcal{L}'| \cdot |\mathcal{L}_\sigma| = \prod_{i=1}^k |B_i|!$.

L is a linearization thanks to all L_v being linearizations, B_k being a barrier (so all relative orders of L_v in L_σ are valid) and (B_1, \dots, B_k) being a barrier sequence (so no vertex in L' can be a successor of a vertex in L_σ). Uniqueness follows from the fact that changing L' or L_σ changes L as well. \square

3.2 Constructing Barriers in the Exploration Tree

In the rest of this section, we show that we can find suitable barriers in the exploration tree that “pay for” the cost of the DELETETEMIN operations. We now sketch the big picture idea behind our proof of how we relate to each other the working sets and the barriers.




Figure 5: Example of a graph with radically different $T_{SSSP}^{G,w}$ and $T_{explore}^{G,w}$. $T_{SSSP}^{G,w}$ only has one linearization, while $T_{explore}^{G,w}$ has $n!$ linearizations.

Consider running Dijkstra’s algorithm until some time t . Let us say that S is the set of nodes that are in the priority queue of the algorithm at time t , and choose t such that S is the largest possible. On one hand, the total cost the algorithm will pay for the DELETEMIN of all elements in S is $\mathcal{O}(|S| \log |S|)$, since by definition every working set has size at most $|S|$.

On the other hand, if we imagine building the exploration tree $T_{explore}^{G,w}$ in parallel to running Dijkstra’s algorithm, then S is a suitable barrier in $T_{explore}^{G,w}$, and per the previous section, it contributes $\Omega(|S| \log |S|)$ to the lower bound. To see that S is a barrier, note that S are the leaves of the partial exploration tree built so far, which is itself a subgraph of the final exploration tree $T_{explore}^{G,w}$, and therefore, S is an incomparable set in the final exploration tree.

If we can continue this process recursively and decompose the whole vertex set into a barrier sequence, then we can relate our barrier-sequence-based lower bound with the actual cost of DELETEMINS, and thus prove optimality.

In order to make the reasoning about possible barriers easier, we will work in the following framework: imagine drawing intervals, one for each vertex, denoting the timespan between inserting this vertex and extracting it. Every set of vertices whose intervals all overlap at some time t is a barrier, precisely because of the argument about the exploration tree above. Therefore, if we *color* the intervals so that all intervals of one color have a nonempty intersection, it holds that each color induces one barrier, and – a fact that we shall prove – the barriers form a barrier sequence, and can therefore be used to prove a lower bound. The crucial claim of this section is that if we proceed analogously to our choice of S above – that is, if we greedily color the largest possible intersecting interval set and solve the rest recursively, the resulting barrier set gives us a good enough lower bound.

We note that it is important to work with the exploration tree here, not the shortest path tree. As an example, consider the graph G being equal to the fan in [Figure 5](#). If weights are $w(s, v_i) = i$ and $w(v_i, v_{i+1}) = 0.5$, $T_{SSSP}^{G,w}$ is a path. Yet, even in this case, running Dijkstra’s algorithm with a Fibonacci-like priority queue with the working set property results in an $\Omega(n \log n)$ time complexity, since the vertices may be inserted to the heap in the order v_1, \dots, v_{n-1} and then if we run the algorithm, half of the nodes have working set of size at least $n/2$, so their DELETEMIN cost is $\Omega(\log n)$. However, in this example, the exploration tree $T_{explore}^{G,w}$ is equal to the star, so if we pick a single barrier $B = \{v_1, \dots, v_{n-1}\}$, this barrier can “pay” $\Omega(|B| \log |B|) = \Omega(n \log n)$, i.e., it can pay for the operations done by Dijkstra’s algorithm, exactly as we need.

We now need to introduce additional notation. Whenever we run Dijkstra’s algorithm on G , we can generate a transcript of INSERT and DELETEMIN operations. Such a transcript allows us to talk about insertions, deletions, working sets, etc. without the need of talking about the graph G itself. More formally, every run of Dijkstra’s algorithm on some graph G corresponds to a set I

of n time intervals defined as follows.

Definition 3.7 (Dijkstra-induced interval set). *An interval set I of size n is a set of n nonempty closed intervals such that all the $2n$ start and end times of all intervals are distinct.*

A run of Dijkstra's algorithm induces an interval set, namely each vertex v induces an interval $\iota(v) = [\ell, r]$, where ℓ is the time v was first discovered and the value associated with vertex v was inserted into the heap, and r is the time v was extracted from the heap. All times are measured in the amount of INSERT and DELETETEMIN operations performed on the heap so far. We say such an interval set is Dijkstra-induced.

We now define the working set of an interval. This definition is analogous to [Definition 1.3](#).

Definition 3.8 (Working set of an interval). *For every $x \in I$, working set at time $t \in x$ is defined as the set of intervals $y \in W_{x,t}$ starting after x starts and for which $t \in y$. Formally, for $x = [\ell_x, r_x]$,*

$$W_{x,t} = \{ [\ell_y, r_y] \in I \mid \ell_x \leq \ell_y \leq t \leq r_y \}.$$

The working set W_x of x is then defined as W_{x,t^} for arbitrary $t^* = \arg \max_{t \in x} |W_{x,t}|$.*

Definition 3.9. *We say that the cost of $x \in I$ is equal to $\text{cost}(x) = \log |W_x|$. We also define for an interval set I that $\text{cost}(I) = \sum_{x \in I} \log |W_x|$.*

We will use the fact that $\text{cost}(x)$ corresponds to the amortized cost of the DELETETEMIN operation performed on x . Similarly, $\text{cost}(I)$ is the overall amortized cost of all DELETETEMIN operations performed on the set I .

Definition 3.10 (Intersecting coloring, energy). *An intersecting coloring of I with k colors is a function $C : I \rightarrow [k]$ that assigns a color to every interval, and, additionally, for every color $i \in [k]$, the intersection of all intervals having color i is nonempty.*

Given a coloring C , let $c_i = |\{a \mid C(a) = i\}|$. The energy of C is defined as $\mathcal{E}(C) = 2 \sum_{i=1}^k c_i \log c_i$.

For Dijkstra-induced interval sets, if some intervals are colored with one color, the definition above states that there exists a time t such that all their corresponding vertices were present in the heap at time t . In fact, every intersecting coloring induces a barrier sequence in the exploration tree $T_{\text{explore}}^{G,w}$, and thus, indirectly, a lower bound on the number of linearizations:

Lemma 3.11 (Intersecting coloring gives a lower bound). *Let I be the interval set induced by running Dijkstra's algorithm ([Algorithm 9](#)) on a graph G , and let C be its arbitrary intersecting coloring of I . Then*

$$\log(\text{Linearizations}(G)) \geq \log \left(\text{Linearizations} \left(T_{\text{explore}}^{G,w} \right) \right) = \Omega(\mathcal{E}(C)).$$

Proof. For each color c , define $B_c = \{v \in V(G) \mid C(\iota(v)) = c\}$ and $t_c = \min\{t \mid \forall v \in B_c : t \in \iota(v)\}$, i.e., the minimum time when all vertices whose corresponding intervals have color c were in the heap. Without loss of generality, assume $t_1 < \dots < t_k$, where to simplify notation, we are using that the colors are the first k integers for some value k ; if these inequalities do not hold, we permute the colors so as to make them true (this can always be done as no two intervals start at the same time). We claim that (B_1, \dots, B_k) is a barrier sequence in $T_{\text{explore}}^{G,w}$.

Per [Claim 3.5](#), all we have to show is that for any $i \leq j$, $u \in B_i$ and $v \in B_j$, v is not a predecessor of u in $T_{\text{explore}}^{G,w}$. Take such i, j, u, v , and let t_i and t_j be the times defined above. We know that $t_i \leq t_j$.

What would it mean for v to be a predecessor of u in the exploration tree? By definition, a parent p of each vertex w in the exploration tree is the vertex during whose exploration w was added to the heap. But due to the nature of Dijkstra's algorithm, this means the insertion of w into the heap only happened after p was removed from the heap, and the same is true for any predecessor of w and not just a parent. Specially, the interval $\iota(p)$ ends before $\iota(w)$ starts. But this means that if v were a predecessor of u , then necessarily $\iota(u) \ni t_i > t_j \in \iota(v)$, which is a contradiction.

Hence, (B_1, \dots, B_k) is a barrier sequence in $T_{explore}^{G,w}$. The rest follows from [Lemma 3.6](#) and the fact that $|B_i| = c_i$ and $\mathcal{E}(C) = 2 \sum_{i=1}^k |B_i| \log |B_i|$. \square

The interval formalism allows us to transform a problem about proving time complexity lower bounds into a problem about coloring intervals. As we argue, $\text{cost}(I)$ gives an asymptotic upper bound on the total cost of the DELETETMIN operations. At the same time, $\mathcal{E}(C)$ gives an asymptotic lower bound on $\text{OPT}_Q(G)$. We can, therefore, prove [Proposition 3.2](#) by proving there always exists C satisfying $\mathcal{E}(C) \geq \Omega(\text{cost}(I))$, which we will do in the rest of this section. We now formalize this argument:

Lemma 3.12 (Finding a good coloring proves optimality). *Run Dijkstra's algorithm ([Algorithm 9](#)) with any Fibonacci-like priority queue with the working set property on graph G . Let I be the interval set induced by this run. If there exists a coloring C of I such that $\mathcal{E}(C) \geq \text{cost}(I)$, then the total time complexity of all DELETETMIN calls performed by the algorithm is $\mathcal{O}(\text{OPT}_Q(G) + n)$.*

Proof. On one hand, the overall time complexity of all DELETETMIN operations can be upper bounded by

$$\begin{aligned} \sum_{v \in V(G)} \mathcal{O}(1 + \log |W_v|) &= \mathcal{O}(n + \sum_{v \in V(G)} \log |W_v|) \\ &= \mathcal{O}(n + \sum_{x \in I} \log |W_x|) \\ &= \mathcal{O}(n + \text{cost}(I)) = \mathcal{O}(n + \mathcal{E}(C)), \end{aligned}$$

where the first equality holds because $\sum_{i=1}^n \mathcal{O}(f_i(k)) = \mathcal{O}(n + \sum_{i=1}^n f_i(k))$, the second because the heap-based and interval-based definitions of working sets coincide, and the third follows from the definition of $\text{cost}(I)$.

On the other hand, per [Lemmas 3.3](#) and [3.11](#), $\text{OPT}_Q(G) = \Omega(\mathcal{E}(C))$. Put together, the overall amortized complexity of DELETETMIN operations is upper bounded by $\mathcal{O}(n + \text{OPT}_Q(G))$, as needed. \square

3.3 Finding a Good Interval Coloring

Ultimately, we want to find the intersecting coloring recursively. This means that we need to understand how the cost of I changes if we delete an element from it.

Lemma 3.13 (Deleting intervals from I). *Consider any interval set I and let $x \in I$. Let $k = \max_t |\{y \in I \mid t \in y\}|$; that is, k is the largest number of intervals overlapping at any given time. Then,*

$$\text{cost}(I) \leq \text{cost}(I \setminus \{x\}) + \log |W_x| + \log k. \quad (2)$$

Moreover, for any element x with $|W_x| = k$ we have

$$\text{cost}(I) \leq \text{cost}(I \setminus W_x) + 2|W_x| \log |W_x|. \quad (3)$$

Proof. We start by proving [Equation \(2\)](#). Consider removing x from I and then putting it back: the overall cost increase $\text{cost}(I) - \text{cost}(I')$ can be bounded first by $\text{cost}(x) = \log |W_x|$ for the newly added interval x . Then, some intervals of I' may increase the size of their working set by one, due to x being in their working set in I . To prove [Equation \(2\)](#), we need to bound the overall increase by $\log k$.

Let t be the starting time of x . Let $e_1, \dots, e_{k'} \in I'$ be all the intervals containing t , ordered in the decreasing order of their start time. Notice that $k' \leq k - 1$ since $x, e_1, \dots, e_{k'}$ all overlap at time t .

We note that $e_1, \dots, e_{k'}$ are the only intervals (besides x) whose contributions to $\text{cost}(I')$ and $\text{cost}(I)$ differ from each other: every other interval y has $t \notin y$, which either means that y ends before t , or y starts after t . In both cases, x is by definition not included in $W_{y,t'}$ for any t' , and therefore it also is not included in W_y .

Moreover, notice that the working set of e_i is of size at least $i + 1$, since one candidate for W_{e_i} is the set $W_{e_i,t} = \{x, e_1, \dots, e_i\}$, i.e., the working set of e_i at the time t . Since the size of the working set W_{e_i} is either $|W'_{e_i}|$ or $|W'_{e_i}| + 1$ where W'_{e_i} is the working set of e_i in I' , we conclude that

$$\begin{aligned} & \text{cost}(I) - \text{cost}(I \setminus \{x\}) - \log |W_x| \\ &= (\log |W_{e_1}| - \log |W'_{e_1}|) + \dots + (\log |W_{e_{k'}}| - \log |W'_{e_{k'}}|) \\ &\leq (\log 2 - \log 1) + \dots + (\log(k' + 1) - \log k') \\ &= \log(k' + 1) - \log 1 \\ &\leq \log k \end{aligned}$$

as needed. On the third line, we used the fact that \log is concave and therefore $f(n) = \log(n + 1) - \log(n)$ is increasing.

To prove [Equation \(3\)](#), we simply remove the elements of W_x from I one by one and each time we apply [Equation \(2\)](#). We note that the maximum number of intervals intersecting at any given time cannot increase by deleting intervals from I , thus the overall difference between $\text{cost}(I)$ and $\text{cost}(I \setminus W_x)$ can be upper bounded by $2k \log k = 2|W_x| \log |W_x|$. \square

Finally, we will now use [Lemma 3.13](#) recursively to find an intersecting coloring with high energy.

Lemma 3.14 (Greedy analysis). *For any interval set I , there exists an intersecting coloring C such that $\mathcal{E}(C) \geq \text{cost}(I)$.*

Proof. We construct C by induction on $|I|$. For $I = \emptyset$, set $C = \emptyset$. Otherwise, find the interval $x \in I$ with the largest W_x , and use induction on $I' = I \setminus W_x$ to find a coloring C' which satisfies $\mathcal{E}(C') \geq \text{cost}(I')$. Now we define $C(x) = C'(x)$ if $x \in I'$ and $C(x) = c^*$ otherwise, where c^* is a new color not used in C' .

On one hand, $\mathcal{E}(C) = \mathcal{E}(C') + 2|W_x| \log |W_x|$, by the definition of \mathcal{E} . On the other hand, per [Lemma 3.13](#), $\text{cost}(I) \leq \text{cost}(I') + 2|W_x| \log |W_x|$.

Putting this all together, we have $\mathcal{E}(C) = \mathcal{E}(C') + 2|W_x| \log |W_x| \geq \text{cost}(I') + 2|W_x| \log |W_x| \geq \text{cost}(I)$, as needed. \square

Finally, we may now prove [Proposition 3.2](#):

Proof of Proposition 3.2. Per [Lemma 3.14](#), there always exists C such that $\mathcal{E}(C) \geq \text{cost}(I)$. Per [Lemma 3.12](#), if such a C exists, [Proposition 3.2](#) holds. \square

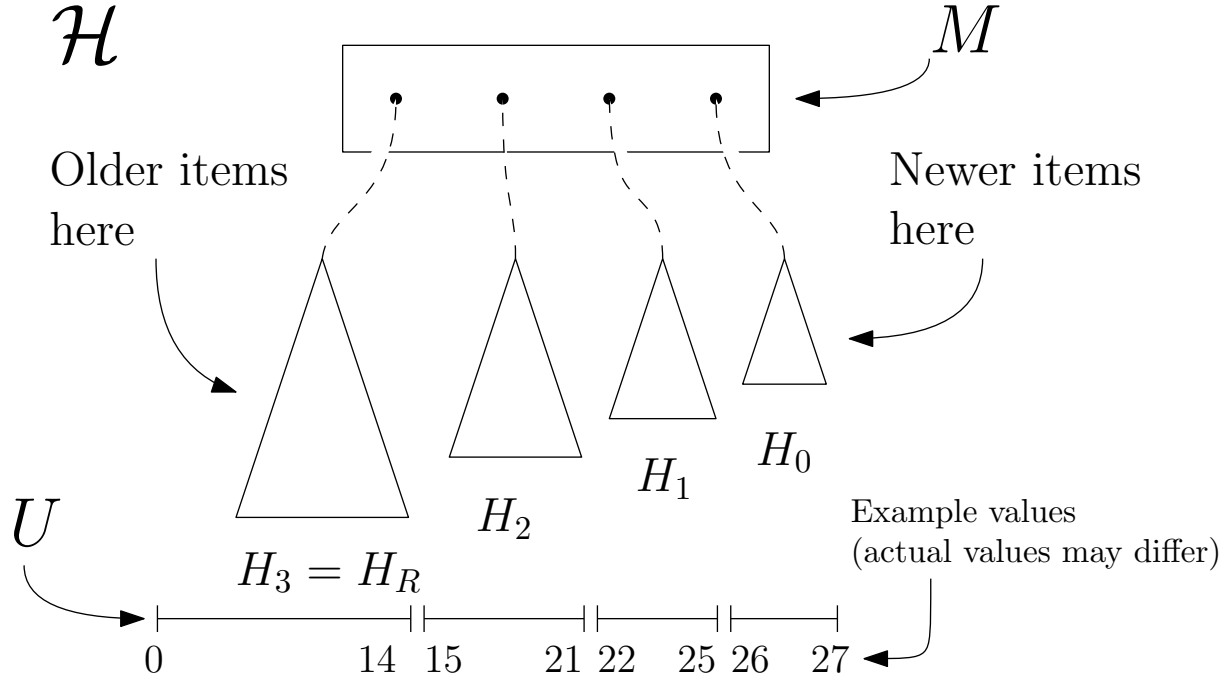


Figure 6: Illustration of the data structure \mathcal{H} . Note that the intervals in U depend on the exact sequence of INSERTS and DELETEMINS.

4 Constructing Heaps with the Working Set Property

In this section, we design a Fibonacci-like priority queue with the working set property, and by doing that, prove [Theorem 1.5](#):

Theorem 1.5. *There is a Fibonacci-like priority queue with the working set property.*

We now give a high-level description of the data structure, then we state the invariants that we require it satisfies, and then we give a high-level description of some of the operations. The formal description of the operations then follows.

High-level description of the data structure The general idea is to have $\mathcal{O}(\log \log n)$ Fibonacci heaps² of sizes roughly $2^{2^0}, 2^{2^1}, 2^{2^2}, \dots$ such that recently inserted elements are in smaller heaps, and their DELETEMINS are therefore cheap.

The data structure consists of a collection \mathcal{H} of Fibonacci heaps containing the elements, together with auxiliary data structures that allow us to operate efficiently on \mathcal{H} . Each of the heaps gets assigned a *rank* $r \in \mathbb{N}_0$, and we have one (possibly implicitly stored, empty) heap for each rank. Denote by H_r the heap of rank r , and by $|H_r|$ its size (zero if $H_r = \emptyset$). For an element $x \in \mathcal{H}$, denote $r(x)$ the rank of the heap x is stored in. Finally, each $x \in \mathcal{H}$ has assigned a unique *insertion time* $t(x) \in \mathbb{N}$. Formally, instead of storing raw elements, each H_r stores elements as pairs $(x, t(x))$. When we insert a new element y into \mathcal{H} , its $t(y)$ will be the last used $t(x)$ incremented by 1.

Invariants We will be maintaining the following invariants:

1. For all r , $|H_r| \leq 2^{2^r}$.

²Our reduction is black-box and holds for any data structure with the functionality of Fibonacci heaps.

2. Let R be the maximum R such that $H_R \neq \emptyset$. If $R \geq 2$, then $|H_R| + |H_{R-1}| \geq 2^{2^{R-1}}$.
3. Define a partial order \prec on heaps as follows: $H_a \prec H_b$ if for all $x \in H_a$ and $y \in H_b$, we have $t(x) < t(y)$. Then \prec is actually a linear order and $H_0 \succ H_1 \succ \dots \succ H_R$. In other words, elements inserted earlier are in higher-rank heaps.

Note that **Invariant 2** implies that the maximum rank is $\mathcal{O}(\log \log n)$ where n is the number of elements currently in \mathcal{H} , as the existence of rank $> \log \log n + 1$ would imply there are more than $2^{2^{\log \log n}} = n$ elements in the last two heaps. Also observe that \prec is indeed a partial order and that $H_a \prec H_b$ and $H_a \prec H_c$ implies $H_a \prec H_b \cup H_c$.

High-level description of the operations With the invariants in mind, we describe the idea behind the operations. A heap with rank r can grow by merging smaller heaps into it, and once it grows too much (it would violate **Invariant 1**), we *promote* it to the next rank by merging it into the heap with rank $r + 1$. INSERT creates a heap of size 1 and then triggers a series of merges and promotions. DECREASE and DELETEMIN first find the correct heap using auxiliary data structures (described below), and then perform the operation there. If the two top-rank heaps shrink too much after an DELETEMIN, we merge them together into H_{R-1} to preserve **Invariant 2**.

Auxiliary data structures Before we can show how to implement the operations, we describe the auxiliary data structures associated with \mathcal{H} . We now specify their interfaces and prove their existence later in **Sections A.1** and **A.2**. We note that the data structure from **Theorem A.1** is used both as a standalone data structure and as a building block in **Theorem A.2**.

Theorem A.1 (Interval maintenance). *Assume the Word-RAM model with word size w . There exists an ordered-dictionary-like data structure U for storing right-open, pairwise non-overlapping intervals. Interval boundaries are integers from $\{0, \dots, 2^{\mathcal{O}(w)}\}$. There may be at most ℓ intervals in the set at once, for some $\ell = \mathcal{O}(\log w)$. The data structure supports the following operations, each with $\mathcal{O}(1)$ worst-case time complexity:*

- SET(a, b, x): Set $U[[a, b]] = x$. Caller guarantees $[a, b]$ does not overlap with any other interval present in U .
- DELETE(a, b): Delete $[a, b]$ from U if it exists.
- GET(a, b): Return $U[[a, b]]$ or \emptyset if such interval is not in U .
- FIND(t): Return (unique) a, b such that $t \in [a, b) \in U$, or \emptyset if no such interval exists.
- PREV(t): Return largest c, d such that $[c, d) \in U$ and $d \leq t$.
- NEXT(t): Return smallest c, d such that $[c, d) \in U$ and $t < c$.

Proof sketch. See **Section A.1** for the full proof. It is folklore knowledge that the interval maintenance problem can be straightforwardly reduced to ordered set maintenance. With fusion trees, this leads to $\mathcal{O}(\log_w w) = \mathcal{O}(1)$ time complexity. Incidentally, as $\ell = \mathcal{O}(\log \log n)$, even with AVL trees we obtain the same guarantees up to a $\mathcal{O}(\log \log \log n)$ factor. \square

Theorem A.2 (Minimum-keeping). *Assume the Word-RAM model with word size w . Then there exists a data structure that maintains an array M of length $\ell = \mathcal{O}(\log w)$ that supports the following operations:*

- GET(i): return $M[i]$.

- **DECREASE**(i, x): given $x \leq M[i]$, set $M[i] := x$.
- **CHANGEPREFIX**(P): given an array P of length i , set $M[0 : i] = P$. P is allowed to be larger than the current length of the array.
- **FINDMIN**(i): return $\arg \min_i M[i]$.
- **POP**(i): decrease the length of M by one. It is promised that the last two elements of M contain the same value beforehand.

The elements stored may only be loaded, stored and compared, all in $\mathcal{O}(1)$ time per operation. The amortized time complexity of those operations is as follows: $\mathcal{O}(1)$ for **GET**, **DECREASE**, **FINDMIN** and **POP**, and $\mathcal{O}(i)$ for **CHANGEPREFIX**, where i is the length of the prefix to be changed. The worst-case time complexity is the same, except for **DECREASE**, which takes $\mathcal{O}(i)$ time worst-case where i is the index being decreased.

Proof sketch. See [Section A.2](#) for the full proof. The problem can be solved by efficiently maintaining a compressed array S of suffix minima of M . Then, **FINDMIN** returns $S[0]$, **CHANGEPREFIX** only has to recompute first $|P|$ suffix minima, and the complexity of **DECREASE** can be argued to be $\mathcal{O}(1)$ amortized, since a **DECREASE** taking i units of time decreases the number of unique values of S by roughly i . Storing S efficiently can be achieved using the interval maintenance structure from [Theorem A.1](#). In particular, as an aside, we can again use an AVL tree instead of a fusion tree with only a $\mathcal{O}(\log \log \log n)$ multiplicative slowdown to all operations. \square

We now state how we use the auxiliary data structures to store information on \mathcal{H} . In the minimum-keeping data structure, which we will call M , we store in $M[r]$ the current minimum of H_r . Clearly, using **FINDMIN** on M then returns the globally minimal element, which will be important when implementing **DELETEMIN**.

In the interval data structure, which we will call U , we maintain an interval for each non-empty heap corresponding to the range of insertion times such that items with these insertion times are stored in the given heap; note that this is indeed an interval by [Invariant 3](#). Specially, when we delete an item from H_r , we do not update U . Thus, each interval stored in U is actually a *superset* of the insertion time range of the given heap. We need U for the following subtle reason: The implementation of **DECREASEKEY** in an inner Fibonacci heap needs not only a pointer to an element that is being decreased, but also a pointer to the heap this element is in. U provides those pointers.

If the values stored in M, U are as described, we say that the data structures are valid.

Formal description of the operations We start by specifying an important subroutine, [Algorithm 1](#), that describes a merge of inner heaps. First, we need the following definition. In the process of promoting heaps to larger ranks, we will temporarily create additional *carry heaps*:

Definition 4.1. Given a heap collection \mathcal{H} , a carry heap of rank r is a heap C_r such that $|C_r| \leq 2^{2^r}$ and $H_r \succ C_r \succ H_{r+1}$.

Lemma 4.2. [Algorithm 1](#), i.e., **PROMOTIONSTEP**, has $\mathcal{O}(1)$ amortized time complexity. It maintains [Invariants 1](#) and [3](#) (if they held before executing **PROMOTIONSTEP**, they hold after). It also maintains validity of U . Furthermore, the overall set of elements in the data structure (including the carry heaps) is preserved.

Algorithm 1 PROMOTIONSTEP

Input: $r \in \mathbb{N}$, \mathcal{H} satisfying **Invariants 1** and **3**, carry heap C_{r-1} of rank $r - 1$

Output: $\mathcal{H}' := \mathcal{H} - \{H_r\} \cup \{H'_r\}$ satisfying **Invariants 1** and **3**; carry heap C_r of rank r (with respect to \mathcal{H}')

```
1: if  $|H_r| + |C_{r-1}| \leq 2^{2^r}$  then
2:    $[a, b] \leftarrow U.\text{FIND}(t(H_r.\text{FINDMIN}()))$ ,  $[c, d] \leftarrow U.\text{FIND}(t(C_{r-1}.\text{FINDMIN}()))$ 
    $\triangleright$  Get the insertion time ranges associated with  $H_r$  and  $C_{r-1}$  by querying  $U$  with an arbitrary
   representant.
3:    $C_r \leftarrow \emptyset$ ,  $H'_r \leftarrow H_r \cup C_{r-1}$   $\triangleright$  Here  $\cup$  denotes the merge operation on two Fibonacci heaps.
4:    $U.\text{REMOVE}(a, b)$ ,  $U.\text{REMOVE}(c, d)$ ,  $U.\text{SET}(a, d, H'_r)$   $\triangleright$  Merge the two intervals in  $U$ .
5: else
6:    $C_r \leftarrow H_r$ ,  $H'_r \leftarrow C_{r-1}$ 
```

Proof. Recalling that Fibonacci heaps can be merged in $\mathcal{O}(1)$ amortized time suffices to prove the first part of the lemma.

For correctness, first recall that the input is guaranteed to satisfy $H_{r-1} \succ C_{r-1} \succ H_r \succ H_{r+1}$. Also, note that for **Invariants 1** and **3**, we only need to look at H'_r , since no other heap changed. We consider two cases:

If the if-statement evaluated as true, then $C_r = \emptyset$ is trivially a carry heap of rank r . Also $|H'_r| = |H_r| + |C_{r-1}| \leq 2^{2^r}$, so **Invariant 1** holds. Lastly, by properties of \succ and \cup , we have $H_{r-1} \succ H'_r = H_r \cup C_{r-1} \succ H_{r+1}$, which proves **Invariant 3** for \mathcal{H}' .

Otherwise, we have $H'_r = C_{r-1}$ and $C_r = H_r$. From C_{r-1} being of rank $r - 1$, we get $|H'_r| \leq 2^{2^{r-1}} \leq 2^{2^r}$, so **Invariant 1** is satisfied. Similarly, from H_r being of rank r , we get $|C_r| \leq 2^{2^r}$. Finally, by rewriting the heap ordering from before, we get $H_{r-1} \succ C_{r-1} = H'_r \succ H_r = C_r \succ H_{r+1}$, which proves both **Invariant 3** and the fact that C_r is a carry heap of rank r .

We now argue the validity of U is preserved. If $|H_r| + |C_{r-1}| \leq 2^{2^r}$, we perform a merge operation on heaps and perform a corresponding merging of intervals in U , retaining validity. In the other case, the set of heaps does not change, meaning we do not perform any operations on U , so validity is also clearly preserved.

Finally, note that the algorithm only redistributes the elements between heaps, so the overall set of elements is indeed preserved. \square

Now we finally describe the operations. We start with INSERT, given as **Algorithm 2**.

Algorithm 2 INSERT(x)

$t(x) \leftarrow 1 + \text{last used } t(y)$

$C_{-1} \leftarrow \{(x, t(x))\}$

$U.\text{SET}(t(x), t(x) + 1, C_{-1})$

$P \leftarrow []$

for $r = 0, 1, \dots$ **do**

$C_r = \mathcal{H}.\text{PROMOTIONSTEP}(r, C_{r-1})$

$P.\text{APPEND}(H'_r.\text{FINDMIN}())$

if $C_r = \emptyset$ **then**

\square **break**

$M.\text{CHANGEPREFIX}(P)$

\triangleright Mutate \mathcal{H} in place.

We now prove that this algorithm is correct and efficient.

Lemma 4.3. INSERT preserves *Invariants 1 to 3* and the validity of M and U .

Proof. INSERT consists of iteratively applying PROMOTIONSTEP, which already preserves *Invariants 1* and *3* and the validity of U . The very first call with C_{-1} is also valid as $|C_{-1}| = 1 \leq 2^{2^{-1}}$ and the interval stored in U is correct. Let us show *Invariant 2* is preserved. Consider two cases: either R increased, or it stayed the same.

If R stayed the same, then look at H_{R-1} . If $H'_{R-1} = H_{R-1}$, then also $H'_R = H_R$ and we are done. Otherwise, $|H'_{R-1}| + |H'_R| \geq |H_{R-1}| + |H_R|$, because the elements only move from smaller to larger ranks and at the same time H_R is still the largest non-empty heap. Either way, *Invariant 2* holds.

Finally, if R increased to R' , this means that the algorithm did not stop for $r = R' - 1$, but it stopped for $r = R'$, i.e. that $H'_{R'} = C_{R'-1} = H_{R'-1}$ and $H'_{R'-1} = C_{R'-2} = H_{R'-2}$. Additionally, $|H_{R'-1}| + |C_{R'-2}| > 2^{2^{R'-1}}$, otherwise the procedure would have stopped earlier. But given the first set of equalities, this is precisely what *Invariant 2* says.

Finally, the validity of M is preserved because the minima of all heaps that changed get recomputed. \square

We prove the constant amortized complexity of INSERT later in [Lemma 4.8](#). Next, the description of DELETEMIN is given in [Algorithm 3](#).

Algorithm 3 DELETEMIN

```

1:  $r \leftarrow M.\text{FINDMIN}()$ 
2:  $m \leftarrow H_r.\text{DELETEMIN}()$ 
3:  $m' \leftarrow H_r.\text{FINDMIN}()$   $\triangleright$  The new minimum after deleting  $m$ .
4: set  $M[r] \leftarrow m'$  by calling  $M.\text{CHANGEPREFIX}(M[:r] + [m'])$ 
5: if  $r = R$  and  $|H_R| + |H_{R-1}| < 2^{2^{R-1}}$  then
6:   merge  $H_R$  into  $H_{R-1}$  and merge the respective intervals in  $U$  (analogously to Line 4 in Algorithm 1)
7:    $H_R \leftarrow \emptyset$ ;  $R \leftarrow R - 1$ ,  $M.\text{POP}()$ 
return  $m$ 

```

Lemma 4.4. DELETEMIN preserves *Invariants 1 to 3* and the validity of M and U .

Proof. Neither *Invariants 1* and *3* can be violated by deleting elements and *Invariant 2* is preserved explicitly. The only problem that could arise is that even after merging H_R into H_{R-1} , *Invariant 2* is still not satisfied. But this cannot happen since the fact that all invariants held before the operation means that $|H_R| + |H_{R-1}| = 2^{2^{R-1}} - 1 \geq 2^{2^{R-2}}$ (the last inequality holds for $R \geq 2$).

No maintenance is needed for U as no element moves from one heap to another, except when R decreases, and then we maintain U correctly. For M , only the heaps H_r , H_{R-1} and H_R change, and we update M accordingly. \square

Lemma 4.5. Cost of DELETEMIN is $\mathcal{O}(2^r)$ amortized, where H_r is the heap containing the minimum.

Proof. DELETEMIN consists of calling DELETEMIN on a Fibonacci heap of size $\leq 2^{2^r}$, which costs $\mathcal{O}(2^r)$. Additionally, we spend $\mathcal{O}(r)$ on FINDMIN to find H_r and the same amount of time in CHANGEPREFIX. Finally, we may spend $\mathcal{O}(1)$ additional time to maintain *Invariant 2*. \square

Lemma 4.6. DELETETMIN has amortized time complexity $\mathcal{O}(1 + \log |W_x|)$.

Proof. Per Lemma 4.5, if the element x being extracted is in H_r , the cost of DELETETMIN is $\mathcal{O}(2^r)$. H_r must at some time t have been promoted from rank $r - 1$. In order for INSERT to reach $r - 1$ and not stop earlier, we must have had $|H_{r-2}| + |C_{r-3}| \geq 2^{2^{r-2}}$ when calling PROMOTIONSTEP on rank $r - 2$. Furthermore, $C_{r-3} \succ H_{r-2} \succ H_{r-1}$, which means that at time t , there were at least $2^{2^{r-2}}$ elements newer than any element in H_{r-1} , and in particular, newer than x . Hence, $|W_x| \geq |W_{x,t}| \geq 2^{2^{r-2}}$, and thus $\log |W_x| \geq 2^{r-2} = \Theta(2^r)$, which gives us that the cost is $\mathcal{O}(2^r) = \mathcal{O}(1 + \log |W_x|)$. \square

Now we give the description of DECREASEKEY in Algorithm 4.

Algorithm 4 DECREASEKEY(pointer to a key x, v)

```

1:  $r \leftarrow U.FIND(t(x))$   $\triangleright$  Recall that we store  $t(x)$  together with the key  $x$ 
2:  $H_r.DECREASEKEY(x, v)$ 
3: if  $v < M[r]$  then
4:    $M.DECREASE(r, v)$ 

```

Lemma 4.7. DECREASEKEY preserves Invariants 1 to 3 and the validity of M and U .

Proof. Invariants 1 to 3 are preserved trivially since the operation does not change the heap sizes or insertion times at all.

No maintenance is needed for U . For M , only the heap whose element is decreased changes, and we update M accordingly. \square

Lemma 4.8. The amortized time complexity of both INSERT and DECREASEKEY is $\mathcal{O}(1)$.

Proof. DECREASEKEY consists of $\mathcal{O}(1)$ calls to methods on U , H_r and M , all of which take $\mathcal{O}(1)$ time. In the rest of the proof, we focus on INSERT.

We define a potential Ψ_t equal to the number of ranks r such that $|H_r| > 2^{2^{r-1}}$ at time t . Note that DELETETMIN only changes this potential by a constant and DECREASEKEY not at all, meaning that using this analysis does not affect their amortized complexity subject to this analysis. We thus only focus on INSERT.

We claim that if INSERT stopped after reaching rank \check{r} , then $\Psi_{t+1} \leq \Psi_t - \check{r} + 3$. Changes to H_0 and $H_{\check{r}}$ increased the potential by at most 2, and we claim that all of $H_1, \dots, H_{\check{r}-1}$ previously contributed to the potential, but their counterparts $H'_1, \dots, H'_{\check{r}-1}$ no longer do. Take any r for $0 < r < \check{r}$. We know INSERT has not stopped after reaching r , which means that $H'_r = C_{r-1}$ and $|H'_r| = |C_{r-1}| \leq 2^{2^{r-1}}$. This means H'_r no longer contributes to Ψ_{t+1} . Also, since $|H_r| + |C_{r-1}| \geq 2^{2^r}$, this means that $|H_r| \geq 2^{2^r} - |C_{r-1}| \geq 2^{2^r} - 2^{2^{r-1}} > 2^{2^{r-1}}$ (holds for $r > 0$), so H_r contributes to Ψ_t .

The (non-amortized) time complexity is $\mathcal{O}(\check{r})$ as this is the amount of time we spend in the algorithm itself as well as the amount of time we spend in the call of CHANGEPREFIX. Since the potential decreases by $\check{r} - 2$, we can use that to pay for the cost, leading to an amortized cost of $\mathcal{O}(1)$. \square

Finally, we note that FINDMIN can be implemented by just calling $M.FINDMIN$ and returning the result. This is clearly correct and it has $\mathcal{O}(1)$ time complexity.

Putting this all together, we are ready to prove Theorems 1.5 and 4.9:

Proof of Theorem 1.5. Correctness follows from Lemmas 4.3, 4.4 and 4.7. Amortized costs of the operations follow from Lemmas 4.5 and 4.8. \square

In fact, we can get worst-case guarantees with only a $\mathcal{O}(\log \log n)$ slowdown of INSERT and DECREASEKEY:

Theorem 4.9. *There is a priority queue satisfying the guarantees of [Theorem 1.5](#) such that, additionally, FINDMIN has worst-case cost $\mathcal{O}(1)$, DELETEMIN has worst-case cost $\mathcal{O}(1 + \log |W_x|)$, and INSERT and DECREASEKEY have worst-case cost $\mathcal{O}(\log \log n)$.*

Proof sketch. We use the heap from [Theorem 1.5](#) and replace Fibonacci with strict Fibonacci heaps [[BLT12](#)] that strengthen the amortized guarantees of Fibonacci heaps to worst-case ones. Guarantees of the auxiliary structures of [Theorems A.1](#) and [A.2](#) hold also in the worst case, except for $M.\text{DECREASE}(i, x)$ whose worst-case complexity is $\mathcal{O}(\log \log n)$, not $\mathcal{O}(i)$.

Then we need to reanalyse [Lemmas 4.5](#) and [4.8](#), but this time with worst-case lens. INSERT and DECREASE run in $\mathcal{O}(\log \log n)$ time in the worst case, as they perform $\mathcal{O}(1)$ work on each of $R \leq \log \log n$ heaps in the worst case, and then update the auxiliary structures in $\mathcal{O}(\log \log n)$ time. The analysis of DELETEMIN holds also in the worst case, provided that we use worst-case priority queues for all H_r . Finally, FINDMIN trivially works in $\mathcal{O}(1)$ worst-case time. \square

5 Universally Optimal Dijkstra for Comparisons

In [Section 3](#), we showed an algorithm that is universally optimal w.r.t. time complexity and universally optimal up to an additive term w.r.t. query complexity. The goal of this section is to get rid of this additive term and thus prove the following theorem.

Theorem 5.1. *There exists an algorithm for the distance ordering problem that is universally optimal both with respect to time complexity and query complexity, on directed graphs and when compared against correct deterministic algorithms.*

We remark that [Theorem 5.1](#) can be extended to undirected graphs by slightly modifying the analysis, which we will not do here. We also believe that it can be extended to randomized algorithms, although this seems to require a different analysis.

We prove [Theorem 5.1](#) in a series of claims that analyze [Algorithm 5](#), given below. [Algorithm 5](#) is based on Dijkstra’s algorithm ([Algorithm 9](#)) with a Fibonacci-like priority queue with the working set property from [Theorem 3.1](#). It has two parts. In the first part ([Algorithm 6](#)), we only attempt to find the SSSP tree in the input graph and not the linearization. In the second part ([Algorithm 8](#)), we use a straightforward dynamic programming algorithm to find the linearization given the SSSP tree.

To find the SSSP tree, we need to be more careful than in [Section 3](#) – to give an example, if the input graph is an undirected path and the source is its second vertex, the optimal number of comparisons for distance ordering is $\Theta(\log n)$, so we cannot use Dijkstra’s algorithm directly as it will perform up to $\Theta(n)$ comparisons. It turns out that we can solve this problem in the simpler case of undirected graphs by contracting all bridges in the graph and running Dijkstra only on the contracted graph. In our more complicated, directed, case, we perform an analogous operation – we use the *dominator tree* of G to find a directed equivalent of bridges that we can contract.

5.1 Contractions Guided by the Dominator Tree

We next present the definitions and results related to the dominance relation and the dominator trees in directed graphs.

Definition 5.2 (Dominator [LM69; LT79]). *Given a directed graph G , a source vertex s , and $u, v \in V(G)$, $u \neq v$, we say that*

- u dominates v if every path from s to v also visits u ,
- u immediately dominates v if u dominates v and every other dominator of v dominates u as well.

It is known [LM69] that every vertex other than s has exactly one immediate dominator, possibly equal to s . This justifies the following definition:

Definition 5.3 (Dominator tree [LT79]). *The dominator tree T_D^G is a tree on the vertex set of G . It is rooted at s and the parent of each vertex is its immediate dominator. Note that T_D^G is not necessarily a subgraph of G , i.e., u may immediately dominate v even though $uv \notin E(G)$.*

Fact 5.4 (Common properties of dominators [LT79]). *Let G be a directed graph and T_D^G its dominator tree. The dominance relation is a partial order on vertices, i.e. it is transitive, reflexive and antisymmetric. For two vertices $u, v \in V(G)$, u is an ancestor of v in T_D^G if and only if u dominates v in G .*

Theorem 5.5. *The dominator tree of G can be constructed in linear time. It can be used to answer queries of the form “Does u dominate v in G ?” in $\mathcal{O}(1)$ time.*

Proof. There are many linear-time algorithms for constructing the dominator tree [Har85; Als+99; GT04; Fra+13]. Per **Fact 5.4**, dominance in G is equivalent to ancestry in T_D^G , and that can be solved in $\mathcal{O}(n)$ precomputation and $\mathcal{O}(1)$ time per query by running DFS on T_D^G and remembering for every node v the times $in(v)$ and $out(v)$ DFS entered and left this node respectively. Then u dominates $v \neq u$ if and only if $in(u) < in(v) \leq out(v) < out(u)$. \square

We now give **Algorithm 5** and its first part, **Algorithm 6**, which computes an SSSP tree. Recall that $OPT_Q(G)$ is defined as the query complexity of an optimal algorithm computing the linearization, and thus the existence of **Algorithm 6** serves as a proof that, in some sense, constructing an SSSP tree is easier or equally hard as the distance ordering.

Algorithm 5 Query-Universally optimal algorithm

Input: G, w, s , we assume every node of G is reachable from s

Output: linearization L of nodes of G with respect to w

- 1: Get $T_{SSSP}^{G,w}$ by running **Algorithm 6**
 - 2: Calculate L from $T_{SSSP}^{G,w}$ by running **Algorithm 8**
-

5.2 Algorithm Analysis

We next state the theorem that we are going to prove in **Sections 5.2 to 5.5**. Recall that **Algorithm 6** covers the first part of the overall optimal algorithm **Algorithm 5**.

Theorem 5.6. ***Algorithm 6** correctly finds $T_{SSSP}^{G,w}$. It has $\mathcal{O}(OPT_Q(G))$ query complexity and $\mathcal{O}(OPT_T(G))$ time complexity.*

Algorithm 6 SSSP tree construction in $\text{OPT}_Q(G)$ queries

Input: G, w, s , we assume every node of G is reachable from s

Output: SSSP tree $T_{SSSP}^{G,w}$ of G with respect to w

- 1: $T_D^G \leftarrow$ the dominator tree of G
 - 2: Drop from G all edges uv where v dominates u .
 - 3: **while** T_D^G has a node u with outdegree one **do**
 - 4: $v \leftarrow$ the (only) child of u
 - 5: Contract uv in both G and T_G and recompute weights w (see [Claim 5.7](#)).
 - 6: Let $T_D^{G''}, G'', w''$ be the result.
 - 7: Compute G' and w' from G'' and w'' by deduplicating parallel edges using [Algorithm 7](#). Let $T_D^{G'} \leftarrow T_D^{G''}$.
 - 8: Compute the SSSP tree $T_{SSSP}^{G',w'}$ by Dijkstra's algorithm ([Algorithm 9](#) modified to return a SSSP tree as well), using any Fibonacci-like priority queue with the working set property.
 - 9: Get $T_{SSSP}^{G,w}$ by uncontracting $T_{SSSP}^{G',w'}$; in particular, all edges that got contracted on [Line 5](#) are included in $T_{SSSP}^{G',w'}$.
-

Algorithm 7 Multiedge deduplication

Input: multigraph G , edge weights w

Output: simple graph G' , weights w'

- 1: Let G' and w' be empty.
 - 2: **for** all $u, v \in V(G)$ such that an edge from u to v exists in G **do**
 - 3: Add edge $e = uv$ to G' .
 - 4: Set $w'(e)$ to compute $\min_{\substack{e_i \in E(G) \\ \text{start}(e_i)=u \\ \text{end}(e_i)=v}} w(e_i)$ the first time it is accessed and lazily return it.
- return** G', w'
-

Roadmap The rough roadmap is as follows: The rest of this subsection proves the correctness of the algorithm.

Then we prove that since no contractions are possible in G' , the execution of [Algorithm 9](#) (which, per [Theorem 3.1](#), performs $\mathcal{O}(\text{OPT}_Q(G') + |V(G')| + \max_w |F_{G',w}|)$ comparison queries) on the remaining graph G' actually has $\mathcal{O}(\text{OPT}_Q(G'))$ query complexity, that is, the additive terms are all paid for by $\mathcal{O}(\text{OPT}_Q(G'))$. Specifically, in [Section 5.3](#), we show that $\text{OPT}_Q(G') = \Omega(|V(G')|)$, and in [Section 5.4](#), we show that $\text{OPT}_Q(G') = \Omega(\max_w |F_{G',w}| - |V(G')|)$. The latter lower bound also analogously applies to G and we use it to show that $\max_w |F_{G,w}| - n = \mathcal{O}(\text{OPT}_Q(G))$ and hence we can afford to do the deduplication in [Algorithm 7](#).

Finally, in [Section 5.5](#), we assemble all parts to finish the proof of [Theorem 5.6](#).

Structure of contractions and correctness The following technical claim collects a few technical properties of how dominators and contractions interact. This allows us to argue e.g. that we can contract our graph, find the SSSP tree in the contracted graph, and then uncontract; properties like these ensure that our algorithm is correct.

While [Claim 5.7](#) talks about contracting a single edge in the dominator tree, it makes sense to think about it in the context of contracting a whole path u_1, u_2, \dots, u_k in the dominator tree T_D^G (by path we mean a sequence of nodes whose outdegree is exactly 1 in the dominator tree). The claim justifies that this path also “looks like a path” in the original graph G . More precisely, on one hand, G has to contain edges $u_1u_2, u_2u_3, \dots, u_{k-1}u_k$ (As u_i is the only node directly dominated by u_{i-1} , how else would it be dominated if not for this direct edge?).

On the other hand, consider any u_i for $i > 1$. It turns out that $u_{i-1}u_i$ is the only edge incoming to u_i ! This is because on one hand, vertices not dominated by u_{i-1} cannot have an outgoing edge to u_i lest u_{i-1} is not dominating u_i . On the other hand, edges outgoing from vertices dominated by u_i were deleted in [Line 2](#). What remains are vertices dominated by u_{i-1} but not by u_i , but those do not exist by our assumption on u_{i-1} having outdegree one.

Note that in the following claim, we denote both nodes in a graph G and its dominator tree T_D^G with the same letters u, v, \dots . Similarly, an edge $e = uv$ may be both an edge in G and in T_D^G . For the graph G/e that we get by contracting an edge $e = uv \in E(G)$, we use the notation $[uv]$ to denote the newly created vertex in $V(G')$ while we use $[x], x \in V(G)$ to denote all other nodes in $V(G')$. We sometimes slightly abuse notation and do not carefully distinguish between objects in G and G/e and between parallel edges.

Claim 5.7 (Structure of contractions). *Let T_D^G be the dominator tree of a multigraph G rooted at s , let w be weights on G , and $e = uv \in E(T_D^G)$ an edge such that u is a node with outdegree one in T_D^G and v its unique child in T_D^G . Assume that every node in G is reachable from s . Also, assume that G may have duplicate edges from one vertex to another, but there are no duplicate edges from a dominator vertex to a dominated vertex. Furthermore, assume that there are no edges $xy \in E(G)$ where y dominates x .*

We claim:

1. *The edge e also exists in G and it is the only incoming edge to v in G .*
2. *There is no edge $uz \in E(G)$ other than e for which $z \in T_D^G(v)$.*

Next, define G', T', w' as follows. We set $G' = G/e$ and $T' = T_D^G/e$. Let $\phi(x) : V(G) \rightarrow V(G')$ be the mapping performed by the contraction, i.e., $\phi(u) = \phi(v) = [uv]$ and $\phi(x) = [x]$ otherwise. We set $w' = w$ with the exception that for every edge $f = vx \in E(G)$ and its counterpart $f' = [uv][x] \in E(G')$, we define $w'(f') = w(u, v) + w(f)$. In particular, for any edge $f = ux$ and its counterpart $f' = [uv][x]$, we set $w'(f') = w(f)$. Then, the following holds:

3. Every node in G' is reachable from $\phi(s)$. G' has no two edges e_1, e_2 such that they both go from x to y and x dominates y . Also, there are no edges xy where y dominates x .
4. $T' = T_D(G')$.
5. $T_{SSSP}^{G',w'} = T_{SSSP}^{G,w}/e$.
6. $F_{G',w'} = F_{G,w}/e$, $|F_{G',w'}| = |F_{G,w}| - 1$ and $|V(G')| = |V(G)| - 1$.

Proof. See [Appendix E](#). □

Corollary 5.8. *Algorithm 6 correctly finds $T_{SSSP}^{G,w}$.*

Proof. [Claim 5.7](#) guarantees that each contraction is valid and all necessary invariants are preserved. Specifically, part 4 proves that we do not have to recompute the dominator tree after each contraction and part 5 proves that e is always used in $T_{SSSP}^{G,w}$ and that we can indeed calculate $T_{SSSP}^{G',w'}$ and then uncontract it to find $T_{SSSP}^{G,w}$. The correctness of [Algorithm 6](#) then follows from applying induction on [Claim 5.7](#). □

5.3 Lower Bound in the Number of Vertices

In this section, we prove that the $+n$ term in [Theorem 3.1](#) is negligible for G' that arises in [Algorithm 6](#), i.e., that $|V(G')| = \mathcal{O}(\text{OPT}_Q(G'))$.

We first show that after contracting the nodes of outdegree one in the dominator tree, the new, contracted, tree cannot be too deep.

Lemma 5.9. *Let G be a directed graph and T_D^G its dominator tree. Assume that no node in T_D^G has outdegree 1. Let v_1, \dots, v_k be any sequence of distinct vertices in G such that v_i dominates v_{i+1} for all i . Then $k \leq (n+1)/2$.*

Proof. Since every non-leaf vertex of T_D^G has ≥ 2 children, there are ≥ 2 vertices on each but the very first level of T_D^G . Thus, its height is at most $(n-1)/2$.

By the definition of T_D^G , if v_i dominates v_{i+1} , then $v_{i+1} \in T_D^G(v_i)$. We thus have a sequence $T_D^G(v_1) \supset T_D^G(v_2) \supset \dots \supset T_D^G(v_k)$ of trees of gradually decreasing height. Since $T_D^G(v_1)$ has height at most $(n-1)/2$, this means that $k \leq (n+1)/2$ as needed. □

This implies that the number of linearizations is at least exponential in the number of vertices:

Lemma 5.10. *Let G be a directed graph and T_D^G its dominator tree. Assume that no node in T_D^G has outdegree 1. Then,*

$$\log(\text{Linearizations}(G)) = \Omega(n).$$

Proof. Define $\ell(v) = d_{G,\mathbf{1}}(s, v)$, where $\mathbf{1}$ is a vector of unit weights and $d_{G,\mathbf{1}}$ therefore measures BFS distance. We partition vertices of G by BFS layers, i.e., define $B_i = \{v \in V(G) \mid \ell(v) = i\}$. In the first part of the proof, we want to upper bound the number of layers for which $|B_i| = 1$. Denote this number k .

Let v_1, \dots, v_k be all the vertices for which $B_{\ell(v_i)} = \{v_i\}$, ordered by increasing $\ell(v_i)$. Fix some i and look at a path from s to v_{i+1} . Due to the properties of BFS, before a path visits a vertex in layer L , it must visit a vertex in each previous layer. In particular, a path to v_{i+1} must visit layer $\ell(v_i)$ and its single vertex v_i . Thus v_i dominates v_{i+1} , and this is true for all i . Hence we may apply [Lemma 5.9](#) on v_1, \dots, v_k to conclude that $k \leq (n+1)/2$.

Now let $T_{BFS}^G := T_{SSSP}^{G,1}$ be an arbitrary BFS tree of G . Observe that B_0, \dots, B_R (with $R = \max_v \ell(v)$) is a barrier sequence with respect to T_{BFS}^G . By [Lemma 3.6](#), $\log(\text{Linearizations}(G)) = \Omega(\sum_i |B_i| \log |B_i|)$. Now use that $a \log a \geq a - 1$ for $a \geq 1$ and $a \log a \geq a$ for $a \geq 2$ to write:

$$\sum_i |B_i| \log |B_i| \geq \left(\sum_i |B_i| \right) - k = n - k \quad \geq n - (n + 1)/2 \geq n/2 - 1 = \Omega(n).$$

In the first inequality, we used the fact that at most k of the $|B_i|$ are equal to 1. \square

Corollary 5.11. *Let G' be the result of [Algorithm 6](#) performing contractions on G . Then $|V(G')| = \mathcal{O}(\text{OPT}_Q(G'))$.*

Proof. G' and $T_D^{G'}$ satisfy the conditions of [Lemma 5.10](#), so it can be applied and $|V(G')| = \mathcal{O}(\log(\text{Linearizations}(G')))$. By [Lemma 3.3](#), $\log(\text{Linearizations}(G')) = \mathcal{O}(\text{OPT}_Q(G'))$, and thus $|V(G')| = \mathcal{O}(\text{OPT}_Q(G'))$ as needed. \square

5.4 Lower Bound in the Number of Forward Edges

Next, we prove that every deterministic algorithm needs to make at least $\Omega(|F_{G,w}| - n + 1)$ comparison queries. (In contrast to [Section 5.3](#), we make no assumptions on G .) The idea of our lower bound is as follows: the algorithm makes linear queries, so if it does not do enough of them, we have enough degrees of freedom to find a linear subspace such that the algorithm cannot tell the inputs from that subspace apart. On the other hand, for at least one other input in that subspace, the answer to the distance ordering problem is different.

The full proof is somewhat technical and can be found in [Appendix D](#). Here we present its main argument, [Lemma 5.12](#), and briefly mention the details needed to finish the lower bound. Note that in [Lemma 5.12](#) below, we talk about running the algorithm on instances with possibly zero-weight edges, which we otherwise forbid (see [Section 2.1](#)). This is fine as long as we do not use the instance to claim incorrectness. Making the weights strictly positive is one detail that we sort out in [Appendix D](#).

Lemma 5.12. *Run a deterministic correct algorithm A on input (G, w) (with $w > 0$), and assume it makes at most $|F_{G,w}| - n$ queries, and let L be the linearization it outputs. Then there exist weights $w' \geq 0$ such that A run on (G, w') performs the same queries and gets the same answers, and furthermore, there is an edge $uv \in F_{G,w}$ such that $w'(uv) = 0$ and $L(v) > L(u) + 1$. That is, v comes after u in L , but they are not neighboring in L .*

Proof. Each query can be represented as $a_i^\top w = b_i$, where the algorithm supplies a_i and gets back $\text{sgn}(b_i)$. If the algorithm makes at most $|F_{G,w}| - n$ queries, there is an affine space $W = \{x \in \mathbb{R}^m \mid Ax = b\}$ having $\dim W \geq m - |F_{G,w}| + n$ such that the (deterministic) algorithm cannot distinguish between w and any $w' \in W$. If we add at most $m - |F_{G,w}| + n - 1$ additional constraints requiring that all edges but those in $F_{G,w} \setminus \{uv \mid L(v) = L(u) + 1\}$ keep their weight, we are still left with a nontrivial affine subspace $K \subset W$ having $\dim K \geq 1$.

We can choose an arbitrary $0 \neq d \in \mathbb{R}^m$ such that $w + d \in K$. Assume d has at least one negative component, otherwise take $-d$. If we now move from w in the direction of d , there exists $\alpha \in \mathbb{R}^+$ such that $w' := w + \alpha d \geq 0$, but for some edge e , $w'(e) = 0$.

Since $w' \in K \subseteq W$, the queries posed by A cannot distinguish between w and w' . Additionally, the edge uv having $w'(uv) = 0$ must also have $uv \in F_{G,w}$ and $L(v) > L(u) + 1$, because any f not satisfying one of these conditions has $w'(f) = w(f) > 0$ by the definition of K . \square

One may be tempted to conclude that we have shown that we can turn (G, w) with a linearization L into (G, w') with a linearization $L' \neq L$ such that A behaves the same on both inputs, and that this is therefore a contradiction with the correctness of A . However, there are two problems with this argument: first, (G, w') can have multiple valid linearizations if $d_w(s, x) = d_w(s, y)$ for some $x \neq y$; specially, L may still be valid for (G, w') . And second, the lemma gives us $w' \geq 0$, not $w' > 0$, and so the input (G, w') is not a valid input to A and we cannot use it to derive a contradiction.

Both of these problems can be solved by applying appropriate perturbations to w and w' , see [Appendix D](#). The rough idea is that we can perturb w so that the answers to the comparison queries are never 0. Since the algorithm only gets $\text{sgn}(b_i)$ and not b_i as a response to a query, this then gives us enough flexibility to be able to perturb w' without changing the answers to queries in such a way that $w' > 0$ and L is inconsistent with w' . This gives us the following result:

Lemma 5.13. *Let (G, w) be a weighted (directed or undirected) graph. Then, any valid deterministic algorithm for the distance ordering problem in the comparison-addition model from [Section 2](#) needs to make at least $|F_{G,w}| - n$ comparison queries when run on some graph (G, w') .*

Proof. See [Appendix D](#). □

We get the following corollary.

Corollary 5.14. *For any G , $\max_w |F_{G,w}| - n + 1 \leq \text{OPT}_Q(G)$.*

5.5 Finishing the Proof of [Theorem 5.6](#)

We are ready to finish the proof of [Theorem 5.6](#).

First, we observe that the optimal comparison complexity only decreases after we contract G to G' .

Proposition 5.15. *Let G' be the contraction of G produced by [Algorithm 6](#). Then $\text{OPT}_Q(G') \leq \text{OPT}_Q(G)$.*

Proof. One way of computing a linearization of G' is to first “reduplicate” it to G'' and then uncontract G'' to G , duplicating weights in the first step and putting zero weights on the newly uncontracted edges in the second step, and then run an optimal algorithm for G . This approach gives an upper bound on $\text{OPT}_Q(G')$ as needed. □

We also check that the number of comparisons used for parallel edge removal is bounded by the number of forward edges.

Proposition 5.16. *Run [Algorithm 5](#) on (G, w) . The total number of comparison queries caused by [Algorithm 7](#) (that is, the total number of lazy queries evaluated later in the algorithm) is at most $\max_w |F_{G,w}| - n + 1$.*

Proof. See [Appendix E](#). □

Before proving [Theorem 5.6](#), we focus on time complexity separately:

Theorem 5.17. *[Algorithm 6](#) can be implemented in $\mathcal{O}(m + n + \text{OPT}_Q(G))$ time.*

Proof. By [Theorem 5.5](#), the construction of T_D^G and the subsequent dropping of edges in G going against the direction of dominance can be done in $\mathcal{O}(n + m)$ time. To implement the edge contraction on [Line 5](#), we first find all vertices of degree one in T_D^G , these vertices form a graph whose connected components are oriented paths. Then, we contract each such path at once. With

a careful implementation, both the contraction and deduplication can be performed in $\mathcal{O}(n + m)$ time. By [Theorem 3.1](#), the [Algorithm 9](#) invocation runs in $\mathcal{O}(\text{OPT}_Q(G') + |V(G')| + |E(G')|)$ time. As $|V(G')| \leq n$, $|E(G')| \leq m$, and $\text{OPT}_Q(G') \leq \text{OPT}_Q(G)$ (per [Proposition 5.15](#)), this is $\mathcal{O}(\text{OPT}_Q(G) + n + m)$. Finally, the uncontractions can be done in $\mathcal{O}(n + m)$ time. Hence, the total time complexity is $\mathcal{O}(n + m + \text{OPT}_Q(G))$ as claimed. \square

Now we are ready to prove [Theorem 5.6](#):

Proof of Theorem 5.6. By [Corollary 5.8](#), [Algorithm 6](#) is correct. By [Theorem 5.17](#), its time complexity is $\mathcal{O}(m + n + \text{OPT}_Q(G))$ and by [Lemma E.1](#), this is $\mathcal{O}(\text{OPT}_T(G))$.

Finally, by [Theorem 3.1](#) and [Proposition 5.16](#), the query complexity is

$$\mathcal{O}\left(\text{OPT}_Q(G') + |V(G')| + \max_w |F_{G',w}| + (\max_w |F_{G,w}| - |V(G)| + 1)\right).$$

Using [Corollaries 5.11](#) and [5.14](#) (the latter one applied twice), we conclude that this quantity is at most $\mathcal{O}(\text{OPT}_Q(G') + \text{OPT}_Q(G))$, and by [Proposition 5.15](#), this is $\mathcal{O}(\text{OPT}_Q(G))$ as needed. \square

5.6 Constructing a Linearization from an SSSP Tree

It remains to describe the algorithm $DP(G, w, T_{SSSP}^{G,w})$ that constructs the linearization of (G, w) in universally optimal time, given the shortest path tree of (G, w) as an additional input. The algorithm does not take into account edges in $E(G) \setminus T_{SSSP}^{G,w}$ as those do not, by definition, have any effect on the linearization.

Thus, we only need to solve the easier version of the original problem: Given a tree (T, w) , we need to find its linearization. [Algorithm 8](#) straightforwardly solves this problem using dynamic programming. For convenience, it also passes around $d_{G,w}(s, \cdot)$, the computed distances from s . Those can be easily precomputed from (T, w) in linear time and 0 comparison queries.

[Algorithm 8](#) internally uses a query-optimal procedure for merging two sorted lists:

Lemma 5.18. *There is an algorithm MERGESORTEDLISTS for merging two sorted lists L_1 and L_2 of arbitrary lengths $|L_1| \leq |L_2|$ that uses at most $2 \log \binom{|L_1|+|L_2|}{|L_1|}$ comparisons.*

Proof. We use the algorithm of Hwang and Lin [[HL72](#)]. In their [Theorem 2](#), they prove that their algorithm uses at most $\log \binom{|L_1|+|L_2|}{|L_1|} + |L_1| - 1$ comparisons, and we will prove that this is at most $2 \log \binom{|L_1|+|L_2|}{|L_1|}$.

We show the following inequality: for $0 \leq k \leq n/2$, it holds that $\log \binom{n}{k} \geq k$. For $k = 0$, we have $\log 1 = 0 \geq 0$ as needed. For $k > 0$, we start with proving the known inequality that $\binom{n}{k} \geq \frac{n^k}{k^k}$. First note that for $a > b > x > 0$, $\frac{a}{b} \geq \frac{a-x}{b-x}$, as we can write

$$\frac{a}{b} \cdot \frac{b-x}{a-x} = \frac{ab-ax}{ab-bx} > 1,$$

since $a > b$ and thus $ab - ax > ab - bx$. Hence,

$$\binom{n}{k} = \frac{n}{k} \cdot \frac{n-1}{k-1} \cdots \frac{n-k+1}{1} \geq \frac{n}{k} \cdot \frac{n}{k} \cdots \frac{n}{k} = \frac{n^k}{k^k}.$$

Now we can write:

$$\log \binom{n}{k} \geq \log \frac{n^k}{k^k} = k \log \frac{n}{k} \geq k \log \frac{n}{n/2} = k \log 2 = k.$$

Thus, the algorithm of Hwang and Lin [HL72] uses at most $\log \binom{|L_1|+|L_2|}{|L_1|} + |L_1| \leq 2 \log \binom{|L_1|+|L_2|}{|L_1|}$ comparison queries, as needed. \square

Algorithm 8 DP – Universally optimal distance ordering on a tree

Input: a rooted tree T , its root r , array d of distances from s

Output: a linearization of T

```

1: function DP( $T, r, d$ )
2:    $L \leftarrow []$ 
3:   Let  $u_1, \dots, u_t$  be children of  $r$ .
4:   for  $i = 1, \dots, t$  do
5:      $L \leftarrow \text{MERGESORTEDLISTS}(L, \text{DP}(T(u_i), u_i, d))$   $\triangleright$  (use  $d$  to perform the comparisons)
6:    $L \leftarrow [r] + L$ 
7:   return  $L$ 

```

Lemma 5.19. *There exists a time- and query- universally optimal algorithm for the distance ordering problem on trees.*

Proof. Consider Algorithm 8. Observe that at every point in the algorithm, the list $[r] + L$ is a linearization of some tree T_{part} , where T_{part} is a subgraph of T consisting of vertex r and the subtrees of the first t' sons of r for some $0 \leq t' \leq t$.

We prove by induction that at all times, the number of comparison queries performed so far (including all recursive calls) is at most $2 \log(\text{Linearizations}(T_{part}))$.

Initially, $L = []$, T_{part} is a single vertex r and $\log(\text{Linearizations}(T_{part})) = 0$, which is also the number of comparisons performed. In the inductive case, let $T_i := T(u_i)$, denote L_i the linearization of T_i , and let T_{part}, L and T'_{part}, L' be the partial tree and linearization before and after processing the i -th son of r . We can assume from induction that we have already proven the claim for all smaller T and also for all $i' < i$ on this particular T .

By induction, the algorithm has so far used at most $2 \log(\text{Linearizations}(T_{part}))$ comparison queries on T_{part} and at most $2 \log(\text{Linearizations}(T_i))$ comparison queries on T_i . Merging L and L_i together takes at most $2 \log \binom{|L|+|L_i|}{|L|}$ queries. So, on one hand, the total query complexity after constructing L' is at most

$$2 \log \left(\text{Linearizations}(T_{part}) \cdot \text{Linearizations}(T_i) \cdot \binom{|L| + |L_i|}{|L|} \right).$$

On the other hand, T'_{part} has exactly $\text{Linearizations}(T_{part}) \cdot \text{Linearizations}(T_i) \cdot \binom{|L|+|L_i|}{|L|}$ linearizations, since every linearization of T'_{part} is uniquely determined by the choice of L and L_i and the way in which to interleave them, and vice versa.

Thus, we used at most $2 \log \text{Linearizations}(T'_{part})$ comparison queries to compute L' . By induction the total number of comparison queries to compute the linearization of T is at most $2 \log(\text{Linearizations}(T))$, as needed.

Using Lemma 3.3, we conclude that the algorithm is universally optimal with respect to query complexity. Its time complexity is proportional to query complexity (up to an additive term linear in the input size), and so it is universally optimal with respect to time complexity as well. \square

We can now finish the proof of Theorem 5.1.

Proof of Theorem 5.1. The correctness of Algorithm 5 follows from Corollary 5.8 and the correctness of Algorithm 8.

Algorithm 5 spends $\mathcal{O}(\text{OPT}_T(G))$ time in Algorithm 6 and $\mathcal{O}(\text{OPT}_T(T_{SSSP}^{G,w}))$ time in Algorithm 8. The latter is bounded by $\mathcal{O}(\text{OPT}_T(G))$ as well, since every problem on $T_{SSSP}^{G,w}$ can be solved by running the optimal algorithm on G and giving the edges in $E(G) \setminus E(T_{SSSP}^{G,w})$ infinite weights. Hence, the overall time complexity of Algorithm 5 is $\mathcal{O}(\text{OPT}_T(G))$.

Finally, by the same argument, its query complexity is $\mathcal{O}(\text{OPT}_Q(G)) + \mathcal{O}(\text{OPT}_Q(T_{SSSP}^{G,w})) = \mathcal{O}(\text{OPT}_Q(G))$. \square

6 Conclusion and Open Problems

In general, it would be very interesting to see more applications of the concepts of universal and instance optimality. Here we state a concrete open problem that we find interesting.

Universally optimal MST Is there a deterministic universally-optimal algorithm for the minimum spanning tree (MST) problem in the comparison model? We note that by the result of Chazelle [Cha00], the deterministic complexity of MST is already extremely close to linear (up to the inverse Ackermann function). Moreover, the result of Pettie and Ramachandran [PR02] provides an explicit worst-case asymptotically optimal MST algorithm. It is unclear how to achieve universal optimality using their approach: their algorithm uses the idea that if we reduce the problem of computing MST in a graph with m edges and n vertices to computing MST in a graph with $m/2$ edges and $n/2$ vertices in $\mathcal{O}(\text{OPT}(m, n))$ time, we can recurse and solve the MST problem fully in $\mathcal{O}(\text{OPT}(m, n))$ time. This line of reasoning does not seem to be applicable if we want to achieve universal optimality.

References

- [ABC17] Peyman Afshani, J r my Barbay, and Timothy M. Chan. “Instance-Optimal Geometric Algorithms”. In: *J. ACM* 64.1 (Mar. 2017). ISSN: 0004-5411. DOI: [10.1145/3046673](https://doi.org/10.1145/3046673). URL: <https://doi.org/10.1145/3046673>.
- [Als+99] Stephen Alstrup, Dov Harel, Peter W Lauridsen, and Mikkel Thorup. “Dominators in linear time”. In: *SIAM Journal on Computing* 28.6 (1999), pp. 2117–2132.
- [B d+07] Mihai B doiu, Richard Cole, Erik D Demaine, and John Iacono. “A unified access bound on comparison-based dynamic dictionaries”. In: *Theoretical Computer Science* 382.2 (2007), pp. 86–96.
- [BHM13] Prosenjit Bose, John Howat, and Pat Morin. “A history of distribution-sensitive data structures”. In: *Space-Efficient Data Structures, Streams, and Algorithms: Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday* (2013), pp. 133–149.
- [BLT12] Gerth St lting Brodal, George Lagogiannis, and Robert E Tarjan. “Strict fibonacci heaps”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. 2012, pp. 1177–1184.
- [Cha00] Bernard Chazelle. “A minimum spanning tree algorithm with inverse-Ackermann type complexity”. In: *Journal of the ACM (JACM)* 47.6 (2000), pp. 1028–1047.
- [Dij59] E. W. Dijkstra. “A note on two problems in connexion with graphs”. In: *Numerische Mathematik* 1 (1959), pp. 269–271.

- [Dua+23] Ran Duan, Jiayi Mao, Xinkai Shu, and Longhui Yin. “A Randomized Algorithm for Single-Source Shortest Path on Undirected Real-Weighted Graphs”. In: *arXiv preprint arXiv:2307.04139* (2023).
- [Elm06] Amr Elmasry. “A priority queue with the working-set property”. In: *International Journal of Foundations of Computer Science* 17.06 (2006), pp. 1455–1465.
- [EFI12] Amr Elmasry, Arash Farzan, and John Iacono. “A priority queue with the time-finger property”. In: *Journal of Discrete Algorithms* 16 (2012), pp. 206–212.
- [EFI13] Amr Elmasry, Arash Farzan, and John Iacono. “On the hierarchy of distribution-sensitive properties for data structures”. In: *Acta informatica* 50.4 (2013), pp. 289–295.
- [FLN01] Ronald Fagin, Amnon Lotem, and Moni Naor. “Optimal Aggregation Algorithms for Middleware”. In: *Proceedings of the Twentieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. PODS ’01. Santa Barbara, California, USA: Association for Computing Machinery, 2001, pp. 102–113. ISBN: 1581133618. DOI: [10.1145/375551.375567](https://doi.org/10.1145/375551.375567). URL: <https://doi.org/10.1145/375551.375567>.
- [Fra+13] Wojciech Fraczak, Loukas Georgiadis, Andrew Miller, and Robert E Tarjan. “Finding dominators via disjoint set union”. In: *Journal of Discrete Algorithms* 23 (2013), pp. 2–20.
- [Fre76] Michael L Fredman. “How good is the information theory bound in sorting?” In: *Theoretical Computer Science* 1.4 (1976), pp. 355–361.
- [Fre+86] Michael L Fredman, Robert Sedgwick, Daniel D Sleator, and Robert E Tarjan. “The pairing heap: A new form of self-adjusting heap”. In: *Algorithmica* 1.1-4 (1986), pp. 111–129.
- [FT87] Michael L Fredman and Robert Endre Tarjan. “Fibonacci heaps and their uses in improved network optimization algorithms”. In: *Journal of the ACM (JACM)* 34.3 (1987), pp. 596–615.
- [FW93] Michael L. Fredman and Dan E. Willard. “Surpassing the information theoretic bound with fusion trees”. In: *Journal of Computer and System Sciences* 47.3 (1993), pp. 424–436. ISSN: 0022-0000. DOI: [https://doi.org/10.1016/0022-0000\(93\)90040-4](https://doi.org/10.1016/0022-0000(93)90040-4). URL: <https://www.sciencedirect.com/science/article/pii/0022000093900404>.
- [FW94] Michael L. Fredman and Dan E. Willard. “Trans-dichotomous algorithms for minimum spanning trees and shortest paths”. In: *Journal of Computer and System Sciences* 48.3 (1994), pp. 533–551. ISSN: 0022-0000. DOI: [https://doi.org/10.1016/S0022-0000\(05\)80064-9](https://doi.org/10.1016/S0022-0000(05)80064-9). URL: <https://www.sciencedirect.com/science/article/pii/S0022000005800649>.
- [GKP98] Juan A Garay, Shay Kutten, and David Peleg. “A sublinear time distributed algorithm for minimum-weight spanning trees”. In: *SIAM Journal on Computing* 27.1 (1998), pp. 302–316.
- [GT04] Loukas Georgiadis and Robert Endre Tarjan. “Finding dominators revisited”. In: *SODA*. 2004, pp. 869–878.
- [GZ22] Mohsen Ghaffari and Goran Zuzic. “Universally-Optimal Distributed Exact Min-Cut”. In: *arXiv preprint* (2022).

- [Hae+24] Bernhard Haeupler, Richard Hladík, John Iacono, Václav Rozhoň, Robert Tarjan, and Jakub Tětek. *Fast and Simple Sorting Using Partial Information*. 2024. arXiv: [2404.04552](https://arxiv.org/abs/2404.04552) [cs.DS].
- [HRG22] Bernhard Haeupler, Harald Räcke, and Mohsen Ghaffari. “Hop-constrained expander decompositions, oblivious routing, and distributed universal optimality”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 1325–1338.
- [HWZ21] Bernhard Haeupler, David Wajc, and Goran Zuzic. “Universally-optimal distributed algorithms for known topologies”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1166–1179.
- [Hag00] Torben Hagerup. “Improved Shortest Paths on the Word RAM”. In: *Automata, Languages and Programming*. Ed. by Ugo Montanari, José D. P. Rolim, and Emo Welzl. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 61–72. ISBN: 978-3-540-45022-1.
- [Har85] Dov Harel. “A linear algorithm for finding dominators in flow graphs and related problems”. In: *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*. 1985, pp. 185–194.
- [HL72] Frank K. Hwang and Shen Lin. “A simple algorithm for merging two disjoint linearly ordered sets”. In: *SIAM Journal on Computing* 1.1 (1972), pp. 31–39.
- [Iac01] John Iacono. “Alternatives to splay trees with $O(\log n)$ worst-case access times”. In: *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*. 2001, pp. 516–522.
- [Iac00] John Iacono. “Improved upper bounds for pairing heaps”. In: *Scandinavian Workshop on Algorithm Theory*. Springer. 2000, pp. 32–45.
- [IÖ14] John Iacono and Özgür Özkan. “A tight lower bound for decrease-key in the pure heap model”. In: *arXiv preprint arXiv:1407.6665* (2014).
- [KS18] László Kozma and Thatchaphol Saranurak. “Smooth heaps and a dual view of self-adjusting data structures”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 801–814.
- [LT79] Thomas Lengauer and Robert Endre Tarjan. “A fast algorithm for finding dominators in a flowgraph”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 1.1 (1979), pp. 121–141.
- [LM69] Edward S Lowry and Cleburne W Medlock. “Object code optimization”. In: *Communications of the ACM* 12.1 (1969), pp. 13–22.
- [Mun+19] J Ian Munro, Richard Peng, Sebastian Wild, and Lingyi Zhang. “Dynamic Optimality Refuted—For Tournament Heaps”. In: *arXiv preprint arXiv:1908.00563* (2019).
- [PT14] Mihai Pătraşcu and Mikkel Thorup. “Dynamic Integer Sets with Optimal Rank, Select, and Predecessor Search”. In: *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18–21, 2014*. IEEE Computer Society, 2014, pp. 166–175. DOI: [10.1109/FOCS.2014.26](https://doi.org/10.1109/FOCS.2014.26). URL: <https://doi.org/10.1109/FOCS.2014.26>.
- [Pet05] Seth Pettie. “Towards a final analysis of pairing heaps”. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*. IEEE. 2005, pp. 174–183.

- [PR02] Seth Pettie and Vijaya Ramachandran. “An optimal minimum spanning tree algorithm”. In: *Journal of the ACM (JACM)* 49.1 (2002), pp. 16–34.
- [Ram96] Rajeev Raman. “Priority Queues: Small, Monotone and Trans-Dichotomous”. In: *Proceedings of the Fourth Annual European Symposium on Algorithms*. ESA ’96. Berlin, Heidelberg: Springer-Verlag, 1996, pp. 121–137. ISBN: 3540616802.
- [Ram97] Rajeev Raman. “Recent Results on the Single-Source Shortest Paths Problem”. In: *SIGACT News* 28.2 (June 1997), pp. 81–87. ISSN: 0163-5700. DOI: [10.1145/261342.261352](https://doi.org/10.1145/261342.261352). URL: <https://doi.org/10.1145/261342.261352>.
- [Rou21] Tim Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2021. DOI: [10.1017/9781108637435](https://doi.org/10.1017/9781108637435).
- [Roz+22] Václav Rozhoň, Christoph Grunau, Bernhard Haeupler, Goran Zuzic, and Jason Li. “Undirected (1+eps)-Shortest Paths via Minor-Aggregates: Near-Optimal Deterministic Parallel and Distributed Algorithms”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, pp. 478–487. ISBN: 9781450392648. DOI: [10.1145/3519935.3520074](https://doi.org/10.1145/3519935.3520074). URL: <https://doi.org/10.1145/3519935.3520074>.
- [Sha48] Claude Elwood Shannon. “A mathematical theory of communication”. In: *The Bell system technical journal* 27.3 (1948), pp. 379–423.
- [ST23] Corwin Sinnamon and Robert E Tarjan. “Efficiency of Self-Adjusting Heaps”. In: *arXiv preprint arXiv:2307.02772* (2023).
- [ST85] Daniel Dominic Sleator and Robert Endre Tarjan. “Self-adjusting binary search trees”. In: *Journal of the ACM (JACM)* 32.3 (1985), pp. 652–686.
- [Tho00a] Mikkel Thorup. “Floats, Integers, and Single Source Shortest Paths”. In: *J. Algorithms* 35.2 (May 2000), pp. 189–201. ISSN: 0196-6774. DOI: [10.1006/jagm.2000.1080](https://doi.org/10.1006/jagm.2000.1080). URL: <https://doi.org/10.1006/jagm.2000.1080>.
- [Tho04] Mikkel Thorup. “Integer priority queues with decrease key in constant time and the single source shortest paths problem”. In: *Journal of Computer and System Sciences* 69.3 (2004). Special Issue on STOC 2003, pp. 330–353. ISSN: 0022-0000. DOI: <https://doi.org/10.1016/j.jcss.2004.04.003>. URL: <https://www.sciencedirect.com/science/article/pii/S002200000400042X>.
- [Tho00b] Mikkel Thorup. “On RAM Priority Queues”. In: *SIAM Journal on Computing* 30.1 (2000), pp. 86–109. DOI: [10.1137/S0097539795288246](https://doi.org/10.1137/S0097539795288246). URL: <https://doi.org/10.1137/S0097539795288246>.
- [Tho99] Mikkel Thorup. “Undirected Single-Source Shortest Paths with Positive Integer Weights in Linear Time”. In: *J. ACM* 46.3 (May 1999), pp. 362–394. ISSN: 0004-5411. DOI: [10.1145/316542.316548](https://doi.org/10.1145/316542.316548). URL: <https://doi.org/10.1145/316542.316548>.
- [VV17] Gregory Valiant and Paul Valiant. “An automatic inequality prover and instance optimal identity testing”. In: *SIAM Journal on Computing* 46.1 (2017), pp. 429–455.
- [Zuz+22] Goran Zuzic, Goramoz Goranci, Mingquan Ye, Bernhard Haeupler, and Xiaorui Sun. “Universally-Optimal Distributed Shortest Paths and Transshipment via Graph-Based L1-Oblivious Routing”. In: *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022.

A Technical Aspects of Our Heap

In this section, we present two data structures needed to maintain auxiliary data about our heap from [Section 4](#).

A.1 Constant-Time Interval Maintenance

First, we describe a standard data structure for maintaining a set of disjoint intervals in a search tree. It is used in [Section 4](#) to keep track of which elements are in which heap, and also in [Section A.2](#) to efficiently maintain a run-length encoded array. In terms of the interface that we specified back in [Section 4](#), in this section we are proving [Theorem A.1](#). Recall:

Theorem A.1 (Interval maintenance). *Assume the Word-RAM model with word size w . There exists an ordered-dictionary-like data structure U for storing right-open, pairwise non-overlapping intervals. Interval boundaries are integers from $\{0, \dots, 2^{\mathcal{O}(w)}\}$. There may be at most ℓ intervals in the set at once, for some $\ell = \mathcal{O}(\log w)$. The data structure supports the following operations, each with $\mathcal{O}(1)$ worst-case time complexity:*

- **SET**(a, b, x): Set $U[[a, b]] = x$. Caller guarantees $[a, b)$ does not overlap with any other interval present in U .
- **DELETE**(a, b): Delete $[a, b)$ from U if it exists.
- **GET**(a, b): Return $U[[a, b]]$ or \emptyset if such interval is not in U .
- **FIND**(t): Return (unique) a, b such that $t \in [a, b) \in U$, or \emptyset if no such interval exists.
- **PREV**(t): Return largest c, d such that $[c, d) \in U$ and $d \leq t$.
- **NEXT**(t): Return smallest c, d such that $[c, d) \in U$ and $t < c$.

Proof. We may use any ordered set data structure D that allows attaching data to stored elements. **SET**(a, b, x) stores $D[a] = (b, x)$, **DELETE**(a, b) checks if $D[a][0] = b$ and if yes, deletes $D[a]$. **GET**(a, b) checks if $D[a][0] = b$ and if yes, returns $D[a][1]$. **FIND**, **PREV**, **NEXT** can all be implemented by calling the **PREV** and **NEXT** methods of the ordered set D .

If we use e.g. an AVL tree to represent D , we get a $\mathcal{O}(\log \ell)$ time complexity of all operations. On the other hand, using a fusion tree [\[FW93\]](#) (or, more precisely the dynamic variant of Pătraşcu and Thorup [\[PT14\]](#)), the time complexity we get is $\mathcal{O}(\log_w \ell) = \mathcal{O}(\log_w \log w) = \mathcal{O}(1)$, as promised. \square

A.2 Maintaining the Minimum Subheap

This section explains how to efficiently find the heap containing the minimum in our collection of $\mathcal{O}(\log \log n)$ Fibonacci heaps. In terms of the interface that we specified back in [Section 4](#), in this section we are proving [Theorem A.2](#). Recall:

Theorem A.2 (Minimum-keeping). *Assume the Word-RAM model with word size w . Then there exists a data structure that maintains an array M of length $\ell = \mathcal{O}(\log w)$ that supports the following operations:*

- **GET**(i): return $M[i]$.
- **DECREASE**(i, x): given $x \leq M[i]$, set $M[i] := x$.
- **CHANGEPREFIX**(P): given an array P of length i , set $M[0 : i] = P$. P is allowed to be larger than the current length of the array.

- **FINDMIN()**: return $\arg \min_i M[i]$.
- **POP()**: decrease the length of M by one. It is promised that the last two elements of M contain the same value beforehand.

The elements stored may only be loaded, stored and compared, all in $\mathcal{O}(1)$ time per operation. The amortized time complexity of those operations is as follows: $\mathcal{O}(1)$ for **GET**, **DECREASE**, **FINDMIN** and **POP**, and $\mathcal{O}(i)$ for **CHANGEPREFIX**, where i is the length of the prefix to be changed. The worst-case time complexity is the same, except for **DECREASE**, which takes $\mathcal{O}(i)$ time worst-case where i is the index being decreased.

In the rest of this section, we prove this theorem. We start with a correct but slow idea and show how to speed it up. We ignore the **POP** operation until the very end, as we require it in [Section 4](#) only because of a minor technicality.

Suffix minima If there were no **CHANGEPREFIX** calls, we could just remember the global minimum and update it on every **DECREASE** in $\mathcal{O}(1)$ per operation. However, this approach is too costly if we reintroduce **CHANGEPREFIX** calls, since then we have to iterate over the whole M every time.

Instead, we will maintain an array S of suffix minima: $S[i]$ will contain the minimum over $M[i :]$. This means $S[0]$ contains the global minimum, so **FINDMIN** is trivial. **CHANGEPREFIX** can now be processed in $\Theta(i)$ time, since the only values of S that are affected by a change of $M[: i]$ are $S[: i]$, and we can use $S[i]$ to get the minimum of $M[i :]$ instead of iterating over the whole array. Unfortunately, the cost of **DECREASE** is now also $\Theta(i)$, since decreasing $M[i]$ may possibly change all of $S[: i]$. To satisfy [Theorem A.2](#), we need to make it $\mathcal{O}(1)$ amortized.

Compressing values A helpful observation is that the changes to S made by **DECREASE** are very special. Since S is nondecreasing, we can implement **DECREASE** by taking $j = i, i - 1, \dots, 0$, and performing $S[j] \leftarrow M[i]$ until we find j^* with $M[j^*] \leq M[i]$, at which point we stop. This takes $\Theta(i - j^*)$ time, which does not help the worst-case time complexity. However, observe that we set all values in $S[j^* : i]$ to the same value. Hence, if we had a way to store S in a compressed form, compressing each interval where S is constant into a single value, we could use a simple amortization argument, taking the number of distinct values as a potential function, to prove that the time complexity of **DECREASE** is actually $\mathcal{O}(1)$ amortized. Fortunately, the following theorem holds:

Theorem A.3 (Skippable array). *Assume the Word-RAM model with word size w . There exists an array-like data structure that can represent (resizable) arrays of size $2^{\mathcal{O}(w)}$ and additionally supports the following two operations. Initially, the array is empty. All operations take $\mathcal{O}(1)$ time worst-case.*

- **SKIP**(i): return largest $j < i$ such that the array value at j and i differ.
- **CHANGESKIP**(i, x): set $A[j] = A[j + 1] = \dots = A[i] = x$ for $j = \text{SKIP}(i)$. Return j .

Proof. Let A be the array we want to represent. It will be represented using D from [Theorem A.1](#) as follows: we partition A into maximal same-valued intervals $[a, b)$ where $A[a] = \dots = A[b - 1] = x$, and store $D[[a, b)] = x$. Both accessing $A[i]$ and the **SKIP** operation can be trivially implemented in $\mathcal{O}(1)$ time using **FIND** and **GET**. Changing $A[i]$ can be implemented in $\mathcal{O}(1)$ time by finding $[a, b) \ni i$ and splitting it into at most 3 new intervals, updating $D[[i, i + 1)]$, looking at the neighboring intervals and possibly merging some of them if their values are now the same. **CHANGESKIP** can

be implemented by splitting the interval $[a, b) \ni i$ into two and changing the value of the left one. Finally, pushing can be implemented similarly to storing a value, and popping can be implemented by finding the appropriate interval in D and shrinking it by 1. \square

With this data structure, we can prove [Theorem A.2](#):

Proof of [Theorem A.2](#). We store M as a regular array. Additionally, we store its suffix minima S as a skippable array from [Theorem A.3](#). GET and CHANGEPREFIX are implemented trivially in $\mathcal{O}(1)$ and $\mathcal{O}(i)$ time respectively, FINDMIN returns $S[0]$ in $\mathcal{O}(1)$ time. DECREASE(i, x) is implemented by starting from $i = j$ and alternating between CHANGESKIP and SKIP on $S[j]$ until we reach $j < i$ where $M[j] < x$. Define a potential function that counts the number of distinct elements in S . Then if DECREASE performs s SKIPS, this takes $\mathcal{O}(s)$ time and the potential decreases by $s - 1$. Therefore, the potential decrease can pay for the cost of DECREASE and the amortized cost is thus $\mathcal{O}(1)$. Finally, POP can be implemented by just shortening M and S . Since the last two elements of M are promised to be equal, this means that S does not need to be recomputed. \square

B Popular Implementations of Dijkstra Are Not Universally Optimal

In this section, we examine various priority queue data structures and, prove that if we use them in Dijkstra’s algorithm ([Algorithm 9](#)), the resulting algorithm will not be universally optimal. Note that any such priority queue also cannot have the working set property. In [Lemma B.3](#) and [Corollary B.4](#), we show this for some common heaps (notably, Fibonacci heaps), and in [Lemma B.5](#) and [Corollary B.8](#), we show this for some common search trees (notably, splay trees).

Both lower bound claims use specific graph constructions. The first lower bound uses the construction from [Figure 2](#). The second lower bound uses a similar, denser construction, which we will now describe.

Example B.1 (Dense counterexample). *For any $k \in \mathbb{N}$ and $n = k^2$, define the following directed graph $G = (V, E)$ with $n + k$ vertices:*

$$\begin{aligned} V &= V_n \cup V_k \text{ for } V_n = [n], V_k = \{n + 1, \dots, n + k\}, \\ E &= P_n \cup V_n \times V_k \text{ for } P_n = \{(i, i + 1) \mid 1 \leq i < n\} \end{aligned}$$

In other words, the graph consist of a directed path and then a set V_k of k additional vertices such that every vertex of the path has an edge to all vertices in V_k . Set $s = 1$.

Lemma B.2. *Let $G = (V, E)$ be the graph from [Example B.1](#). An universally optimal algorithm can solve the distance ordering problem on this graph in $\mathcal{O}(n^{3/2})$ time.*

Proof. Take Dijkstra’s algorithm with a Fibonacci-like priority queue with the working set property. Note that the working set size of all vertices is at most \sqrt{n} , and for all but \sqrt{n} vertices, it is actually 1. Therefore, the cost of all DELETETMIN calls is $\mathcal{O}(n + \sqrt{n} \log \sqrt{n}) = \mathcal{O}(n)$. The cost of all INSERT calls is $\mathcal{O}(|V|) = \mathcal{O}(n)$, and the cost of all DECREASEKEY calls is $\mathcal{O}(|E|) = \mathcal{O}(n^{3/2})$, and besides that, the algorithm performs only $\mathcal{O}(|V| + |E|) = \mathcal{O}(n^{3/2})$ additional work. \square

Now we are ready to give both lower bounds. First, we prove a simple lower bound for priority queues taking $\Omega(\log n)$ time per DELETETMIN.

Lemma B.3. *Let Q be a priority queue such that DELETETMIN on Q takes $\Omega(\log |Q|)$ time. Then the instantiation of Dijkstra’s algorithm (Algorithm 9) using Q as the priority queue is not universally optimal.*

Proof. Let G be the graph from Figure 2 with $t = \sqrt{n}$ and $r = n - \sqrt{n}$. Define weights w such that $w(v_i v_{i+1}) = 1$ and $w(su_i) = n + 42$ for all i . We prove the statement by showing that the algorithm runs on (G, w) in $\Omega(n \log n)$ time. Indeed: since the vertices u_i get inserted into the heap immediately at the start and are only extracted at the end, the algorithm performs $\Omega(n)$ DELETETMIN calls on a heap of size $\Omega(\sqrt{n})$, so in total, it spends $\Omega(n \log n)$ time on DELETETMIN calls. On the other hand, as we have already seen in Section 1.3, an universally optimal algorithm solves the distance ordering on G in linear time. \square

Corollary B.4. *Dijkstra’s algorithm that uses binary heaps, binomial heaps or Fibonacci heaps as the priority queue data structure is not universally optimal.*

Proof. For binomial and Fibonacci heaps, this follows directly from Lemma B.3, as in both cases, DELETETMIN internally performs bucket sorting into $\Theta(\log |Q|)$ bins.

For binary heaps, we have no lower bound on the complexity of DELETETMIN, but we can note that INSERT of a vertex $1 < v \leq n$ takes $\Theta(\log n)$ time, since the value is inserted to the last level of a heap of size $\Theta(\sqrt{n})$ and then it “bubbles up” all the way to the top. Therefore, we can use the argument from Lemma B.3, but on INSERTS instead. \square

Next, we turn our attentions to a large class of search trees that includes, among others, AVL trees, (a, b) -trees, red-black trees and splay trees.

Lemma B.5. *Let Q be a priority queue such that it is possible, via read-only white-box access to the data structure, to enumerate all elements of Q in sorted order in time $o(|Q| \log |Q|)$. Then the instantiation of Dijkstra’s algorithm (Algorithm 9) using Q as the priority queue is not universally optimal.*

We will prove Lemma B.5 by showing that a universally optimal Dijkstra with such a Q would imply a faster-than-possible algorithm for the following sorting problem:

Definition B.6 (Successive sorting). *Given n arrays a_1, \dots, a_n , each of them being a permutation of the set $\{1, \dots, k\}$, output n arrays S_1, \dots, S_n , each of length k , such that for each i , $a_i[S_i[1]] < a_i[S_i[2]] < \dots < a_i[S_i[k]]$. We are working in the comparison (Word-)RAM model defined in Definition 2.5, i.e., the elements $a_i[j]$ cannot be accessed directly, but only summed and compared.*

Lemma B.7. *For any n and k , any (deterministic or randomized) algorithm solving the successive sorting problem needs to perform $\Omega(nk \log k)$ comparisons in the worst case.*

Proof sketch. The proof is analogous both to the classic proof for (regular) sorting and to the proof of Lemma 3.3. Here we prove the lemma only for deterministic algorithms. There are $(k!)^n$ possible inputs and every one of them has a unique valid output. After q comparisons, the algorithm can only be in 3^q possible states. Therefore, if the algorithm never uses more than q comparisons, it can produce at most 3^q possible outputs. Therefore, it needs $\Omega(\log((k!)^n)) = \Omega(nk \log k)$ queries in the worst case in order to sort all inputs correctly. \square

Now we are ready to prove Lemma B.5:

Proof of Lemma B.5. Let A be Dijkstra’s algorithm instantiated with the Q from the statement as the priority queue. Assume A is universally optimal. Let G be the graph from [Example B.1](#). We show that we can solve the successive sorting problem for $k = \sqrt{n}$ by running an algorithm A' that encodes its input into suitable weights w for the graph G and then runs a slightly modified version of A on (G, w) . We show that we can do this in $o(n^{3/2} \log n)$ time, which contradicts the successive sorting lower bound, and therefore proves that A cannot in fact be universally optimal.

Define w as follows: Set $\varepsilon = 1/100n^2$. For $e \in P_n$, define $w(e) := \varepsilon$. For $e = (i, n + j)$ (recall the definition of V_k in [Example B.1](#)), define $w(e) := n - i + a_i[j]/n$.

Now define $d_i(v)$ as the distance from s to v if we are only allowed to use vertices $V_k \cup \{1, \dots, i\}$. By definition, $d_i(v) \geq d_{i+1}(v)$ for all i . Then we can prove by induction that for all $v \in V_k$ and all $1 \leq i < n$, we have $n - i \leq d_i(v) = n - i + a_i[j]/n + (i - 1)\varepsilon < n - i + \sqrt{n}/n + 1/100n < n - i + 1$. Therefore, it holds that $a_i[x] < a_i[y]$ if and only if $d_i(n + x) < d_i(n + y)$.

We also know how A behaves: it starts in s , insert all vertices from V_k into Q and then repeats the following: insert the next $i \in V_n$ into Q , immediately pop it, iterate over all edges $e \in \{i\} \times V_k$ and decrease all values in the queue.

We let algorithm A' run a modification of A that behaves as follows: first, instead of storing just the weights in Q , we also, for each weight, store the identifier of the vertex this weight refers to. (Internally, we may achieve this by storing ordered pairs (x, v) in Q instead of just x .) Then, every time after the algorithm pops the vertex $i \in V_n$ from Q and performs all the DECREASES, we perform a read-only traversal of Q in sorted order and store the associated vertices into an array S_i . Now S_i stores the vertex numbers $\{n + 1, \dots, n + k\}$ sorted by $d_i(v)$, and by the above observation, they are also sorted by $a_i[j - n]$. After the modified A ends, we output S_1, \dots, S_n , with all numbers decreased by n .

We see that A' , together with the initial preprocessing, solves successive sorting in $o(n^{3/2} \log n)$: The original algorithm A takes $\Theta(n^{3/2})$ time in total, and we additionally traverse Q of size \sqrt{n} , n times in total. By our assumption, each traversal takes $o(\sqrt{n} \log n)$ time, so $o(n^{3/2} \log n)$ in total, and this is also the total complexity of A' . But by [Lemma B.7](#), every algorithm solving the successive sorting problem needs $\Theta(n^{3/2} \log n)$ time, which is a contradiction. \square

Corollary B.8. *Dijkstra’s algorithm that uses AVL trees, red-black trees, (a, b)-trees or splay trees as the priority queue data structure is not universally optimal.*

Proof. Follows immediately, as all those search trees can be traversed by DFS in $\Theta(|Q|)$ time. \square

Initially, this result may seem at odds with the dynamic optimality conjecture for splay trees [[ST85](#)], as we have found a data structure – namely the heap from [Section 4](#) – that is asymptotically more efficient for a sequence of operations than the splay tree. However, the conjecture only applies to binary search trees, which the heap from [Section 4](#) clearly is not. Therefore, our result says nothing about (and is not contradicted by) the dynamic optimality conjecture.

C Dijkstra’s Algorithm

For the sake of completeness, we now state the Dijkstra’s algorithm in [Algorithm 9](#). Especially note that edges uv such that $d_s(u) > d_s(v)$ are always ignored and do not contribute to the query complexity.

Algorithm 9 Dijkstra’s Algorithm

Input: A graph $G = (V, E)$, weights w , a source vertex $s \in V$ Output: A linearization L of G , i.e., vertices of G sorted according to distance from s

```
1:  $S \leftarrow \{\}$   $\triangleright$  A set of explored vertices, initially empty.
2:  $L \leftarrow [s]$   $\triangleright$  A list representing the linearization.
3:  $Q \leftarrow \{\}$   $\triangleright$  Priority queue storing (distance, vertex) pairs, sorted by distance. Initially empty.
4:  $P \leftarrow [\emptyset, \dots, \emptyset]$   $\triangleright$  For each vertex  $v$ , a pointer into  $Q$ , or  $\emptyset$ , if  $v$  has not been inserted yet.
5:  $D \leftarrow [+\infty, \dots, +\infty]$   $\triangleright$  For each vertex, the current best distance from  $s$ .
6:  $P[s] \leftarrow Q.\text{INSERT}((0, s)), D[s] \leftarrow 0$ 
7: while  $|Q| \neq \emptyset$  do
8:    $d_u, u \leftarrow Q.\text{DELETETEMIN}()$ 
9:    $L \leftarrow L + [u], S \leftarrow S \cup \{u\}$ 
10:  for all  $uv \in E$  where  $v \notin S$  do
11:    if  $P[v] = \emptyset$  then
12:       $P[v] \leftarrow Q.\text{INSERT}((+\infty, v))$   $\triangleright$  Insert a dummy value first.
13:       $D[v] \leftarrow \min(D[v], d_u + w(uv))$   $\triangleright$  Use  $d_u$  to update the current best distance to  $v$ .
14:       $Q.\text{DECREASEKEY}(P[v], (D[v], v))$   $\triangleright$  For simplicity, call even if  $D[v]$  did not change.
15: return  $L$ 
```

D Lower Bound on the Optimal Query Complexity in the Number of Forward Edges

Here we present the technical proof of [Lemma 5.13](#) presented in [Section 5.4](#). The goal is to prove that for all (G, w) , $\text{OPT}_Q(G) \geq |F_{G,w}| - n$.

Given a deterministic algorithm A that uses at most $|F_{G,w}| - n$ queries on (G, w) and outputs L , [Lemma 5.12](#) guarantees the existence of weights w' such that A cannot tell w and w' apart and, furthermore, $d_{w'}(s, v) = d_{w'}(s, u)$ for some u, v non-adjacent in L . However, this is not sufficient to get a contradiction, as L may still be consistent with (G, w') , for example if all $d_{w'}(s, x)$ are the same. Also $w'(uv) = 0$, while the distance ordering problem requires $w' > 0$, so (G, w') is not even a valid input to A .

The core idea is to perform appropriate perturbations to w and w' . Concretely, we use that the comparison-addition model is actually strictly weaker than we assumed in [Section 5.4](#), since the algorithm does not get the result $a_i^\top w$ for each query, but only its sign. In [Lemma D.1](#), we show one can without loss of generality assume that A never gets an “is equal” answer to any of its queries, and if not, we can perturb w to make it so without weakening the bound on queries. With such a guarantee, we are free to perturb w' , since for a sufficiently small ε , any ε -perturbation will not change the answers of the queries. In [Lemma D.2](#), we show that with this freedom, we are able to find a ε -perturbation w'' of w' that is strictly positive and for which L is not a valid linearization. The proof of [Lemma 5.13](#) then ties all these steps together.

Lemma D.1 (Preprocessing weights). *Let G be a (directed or undirected) graph and $w > 0$ arbitrary weights. Let A be any correct deterministic algorithm for the distance ordering problem that always terminates. Then, there are weights $w' > 0$ such that:*

1. $F_{G,w} \subseteq F_{G,w'}$.
2. If we run A on (G, w') , the answer to every query is ± 1 .

Proof. We start by defining $\varepsilon_0 := \min_{uv \in F_{G,w}} |d(s, v) - d(s, u)|/10n$ and taking a ℓ_∞ -ball $B_0 =$

$w + (0, \varepsilon_0)^m$. Note that $\varepsilon_0 > 0$ by the definition of $F_{G,w}$, and also that for all $\hat{w} \in B_0$, we have $F_{G,w} \subseteq F_{G,\hat{w}}$. To prove this, note that $d_w(s, v) < d_{\hat{w}}(s, v) < d_w(s, v) + n\varepsilon_0$, because the length of every path increases by some $\delta \in (0, n\varepsilon_0)$ when transitioning from w to \hat{w} . Therefore, for any $uv \in F_{G,w}$, we have $d_{\hat{w}}(s, v) - d_{\hat{w}}(s, u) > d_w(s, v) - d_w(s, u) - n\varepsilon_0 \geq 10n\varepsilon_0 - n\varepsilon_0 > 0$, and thus $uv \in F_{G,\hat{w}}$.

Now we construct a sequence of open, nonempty ℓ_∞ -norm balls $B_0 \subseteq B_1 \subseteq \dots \subseteq B_k$ by induction as follows:

We run algorithm A on G , with the weight vector changing adversarially (but so that it is at all times consistent with previous queries). Once the algorithm asks the i -th query $a_i^\top w$, we let $H_i := \{x \in \mathbb{R}^m \mid a_i^\top x = 0\}$ be the hyperplane of all the weight vectors that make the i -th query evaluate to 0. Then we find an arbitrary nonempty open ℓ_∞ -norm ball B_i satisfying

$$B_i \subseteq B_{i-1} \setminus H_i$$

Such B_i always exists, since the set on the right-hand side is open and nonempty. Then we pick arbitrary $w_i \in B_i$ and report $\text{sgn}(a_i^\top w_i)$ as the answer to the query. Note that $\text{sgn}(a_i^\top x) = \text{sgn}(a_i^\top y)$ for any two $x, y \in B_i$, since the whole B_i is on one side of H_i . Since $w_i \in B_i \subseteq B_j$ for all $j < i$, this means that w_i gives retroactively valid answers to all previous queries as well.

After the algorithm finishes with B_k as the final ball, we take an arbitrary $w' \in B_k$ as the final weight vector. By the above paragraph, $w' \in B_k$ is consistent with all queries the algorithm made, and thus running A on (G, w') will produce the same set of queries that we observed, and the answer to each of them will be ± 1 .

Additionally, $w' \in B_0$, and thus $F_{G,w} \subseteq F_{G,w'}$, so w' satisfies the first property. Lastly, $w' > 0$, since $w' > w > 0$. \square

Lemma D.2 (Postprocessing weights). *Given (directed or undirected) $G, w \geq 0, \varepsilon > 0$ and an edge uv such that $w(uv) = 0$, it is possible to find weights $w' \in w + (0, \varepsilon)^m$ such that any linearization L consistent with w' has $L(v) \leq L(u) + 1$.*

Proof. Let $B = w + (0, \varepsilon)^m$ be an ℓ_∞ -ball. For every $\emptyset \neq E_1, E_2 \subseteq E(G), E_1 \neq E_2$, consider the hyperplane

$$H_{E_1, E_2} = \left\{ w \in \mathbb{R}^m \mid \sum_{e \in E_1} w(e) = \sum_{e \in E_2} w(e) \right\}.$$

Now define

$$B' = B \setminus \bigcup_{\substack{\emptyset \neq E_1, E_2 \subseteq E(G) \\ E_1 \neq E_2}} H_{E_1, E_2}.$$

Clearly, B' is nonempty, since we are subtracting a finite amount of $(m-1)$ -dimensional sets from a m -dimensional set. Now take any $w' \in B'$. We claim that $d_{w'}(s, x) \neq d_{w'}(s, y)$ for $x \neq y$. This is true because every $d_{w'}(s, x)$ can be expressed as $\sum_{e \in E_x} w'(e)$ where E_x are the edges on a shortest path from s to x , and equality would imply the existence of two nonempty $E_x \neq E_y$ giving the same sum.

Finally, set $\varepsilon_F := \min_{x \in V(G)} |d_{w'}(s, u) - d_{w'}(s, x)|/10$ and $\varepsilon' := \min(\varepsilon, \varepsilon_F)$. Define w'' by taking $w'' := w'$ and setting $w''(uv) := \varepsilon'/2$.

Now we claim that w'' satisfies the conditions. Clearly $w'' \in w + (0, \varepsilon)^m$ (remember that $w(uv) = 0$). Take any L consistent with w'' . If $L(v) \leq L(u) + 1$, we are done. Assume $L(v) > L(u) + 1$, i.e., assume $d_{w''}(s, u) \leq d_{w''}(s, z) \leq d_{w''}(s, v)$ for some z . We claim this cannot happen: clearly, $d_{w''}(s, u) = d_{w'}(s, u)$ and $d_{w''}(s, v) \leq d_{w'}(s, v) + \varepsilon'/2$, as uv is the only edge whose weight

decreased. On the other hand, we claim that $d_{w''}(s, z) > d_{w''}(s, u) + \varepsilon'/2$. That is because either $d_{w''}(s, z) = d_{w'}(s, z)$, in which case this holds trivially, or $d_{w''}(s, z) < d_{w'}(s, z)$. But then the shortest path to z in (G, w'') uses the edge uv , in which case it consists of the shortest path to v , the edge uv and at least one more edge from v (since $x \neq v$).

Hence, necessarily $L(v) \leq L(u) + 1$, which concludes the proof. \square

Now we are ready to prove [Lemma 5.13](#) that we restate for convenience.

Lemma 5.13. *Let (G, w) be a weighted (directed or undirected) graph. Then, any valid deterministic algorithm for the distance ordering problem in the comparison-addition model from [Section 2](#) needs to make at least $|F_{G,w}| - n$ comparison queries when run on some graph (G, w') .*

Proof. Fix the algorithm A . First, we use [Lemma D.1](#) to pass from w to some w' . Let $L_{w'}$ be the linearization returned by A for (G, w') . Assume A makes at most $|F_{G,w}| - n \leq |F_{G,w'}| - n$ queries. By [Lemma 5.12](#), there exist weights $w_0 \geq 0$ such that the algorithm cannot distinguish between (G, w') and (G, w_0) and such that there is a forward edge $uv \in F_{G,w'}$ with $L_{w'}(v) > L_{w'}(u) + 1$ and $w'(uv) = 0$. Our goal is to perturb w_0 in a way that leads to a contradiction.

First, there exists $\varepsilon > 0$ such that every $\hat{w} \in w_0 + (0, \varepsilon)^m$ is indistinguishable from w_0 by A . That is because for every query A makes, the set of weight vectors indistinguishable from w_0 by that query is an open half-space, and the finite intersection of all such half-spaces is thus an open set $S \ni w_0$. By openness, $w_0 + (-\varepsilon, \varepsilon)^m \subseteq S$ for some $\varepsilon > 0$, and thus also $w_0 + (0, \varepsilon)^m \subseteq S$.

Now we call upon [Lemma D.2](#) to find a $w^* \in w_0 + (0, \varepsilon)^m$ for which all linearizations L^* have $L^*(v) \leq L^*(u) + 1$. But this is a contradiction: $w^* > 0$ and by the previous paragraph, running A on (G, w') and (G, w^*) leads to the same queries and thus the algorithm answers with the same linearization $L_{w'}$. On one hand, we know that $L_{w'}(v) > L_{w'}(u) + 1$, but on the other hand, we know that if $L_{w'}$ were a valid linearization for (G, w^*) , it would need to have $L_{w'}(v) \leq L_{w'}(u) + 1$. Therefore, A is not correct. \square

inline]Add a remark to section 5 that it's $m - n$ for undirected graphs (without proof)

E Deferred Proofs

In this section, we include some short-but-technical proofs from the main part of the paper. We always restate the claim and then prove it.

Claim 2.3 (Equivalent definition of a linearization). *For any graph G , L is a linearization of G if and only if there exist nonnegative weights w such that*

1. For every two nodes $u \neq v$ we have $d_w(s, u) \neq d_w(s, v)$,
2. $u \prec_L v$ if and only if $d_w(s, u) < d_w(s, v)$.

Proof of Claim 2.3. We start by proving the forward implication. For every edge $uv \in E$, we define $w(uv) = |\{x \mid u \prec_L x \preceq_L v\}|$ if $u \prec_L v$ and $w(uv) = \infty$, otherwise. Let s be the minimum w.r.t. L . We argue that for any vertex v , it then holds $d_w(s, v) = |\{x \mid x \prec_L v\}|$. This will imply both properties.

Since L is a linearization, there is a spanning tree T such that L is a linearization of this T . Consider the s - v path in T . Since L is a linearization of T , this path visits vertices in increasing order under \prec_L , meaning that there is a path with finite length.

Let $s = v_1, v_2, \dots, v_k = v$ be the shortest s - v path. All the edges have to be such that $v_i \prec v_{i+1}$ as otherwise the weight would be infinite by the definition of w . This allows us to write

$$|\{x \mid x \prec v\}| = \sum_{i=1}^{k-1} |\{x \mid v_i \prec_L x \preceq_L v_{i+1}\}| = \sum_{i=1}^{k-1} d(v_i, v_{i+1})$$

which is the total length of the path, showing that $d(s, v) = |\{x \mid x \prec v\}|$ since the path in question is the shortest s - v path.

We now prove the reverse implication. Suppose there exist nonnegative weights w on the edges of G such that the two properties are satisfied. We now prove that L is a linearization. Consider the tree $T_{SSSP}^{G,w}$. It is easy to verify that L is a linearization of this tree, since for any $u \prec_L v$, it holds $d(s, u) < d(s, v)$, and u thus cannot be a successor of v . \square

Theorem 3.1. *In the comparison-addition model of Section 2, Dijkstra's algorithm (Algorithm 9) implemented using any Fibonacci-like priority queue with the working set property has time complexity $\mathcal{O}(\text{OPT}_Q(G) + m + n)$.*

Moreover, the number of comparison queries made by the algorithm is $\max_{w \in \mathcal{W}_G} |F_{G,w}|$.

Proof of Theorem 3.1. First, let us bound the query complexity of the algorithm. By properties of Fibonacci heap with working set property from Definition 1.4, we know that the amortized complexity of INSERT and DECREASEKEY is $\mathcal{O}(1)$. Thus, the total amortized query complexity of INSERT operations is $\mathcal{O}(n)$.

Next, we focus on DECREASEKEY operations. We notice that due to our implementation of Dijkstra's algorithm in Appendix C, the DECREASEKEY calls correspond 1-to-1 to edges $e = uv$ such that u comes before v in the output linearization. In fact, by Claim 2.3 there exist weights w' that define the same linearization and where additionally $d_{w'}(s, u) \neq d_{w'}(s, v)$ for $u \neq v$. If we run Dijkstra on G with weights w' , it will perform the exact same DECREASEKEY calls, and additionally, the edges $e = uv$ that those calls correspond to have $d_{w'}(s, u) < d_{w'}(s, v)$ (as opposed to $d_w(s, u) \leq d_w(s, v)$ with the original weights w). That is, they are forward edges. Thus, the number of all DECREASEKEY calls is at most $|F_{G,w'}|$ and we can bound the overall cost of DECREASEKEY operations by $\mathcal{O}(\max_{\tilde{w}} |F_{G,\tilde{w}}|)$.

Finally, we use Proposition 3.2 to bound the cost of DELETETMIN operations by $\mathcal{O}(\text{OPT}_Q(G) + n)$. Together, we get that the number of queries is bounded by $\mathcal{O}(\text{OPT}_Q(G) + n + \max_w |F_{G,w}|)$, as needed.

The time complexity is dominated by priority queue operations and additional linear time for bookkeeping, thus we can similarly bound the time complexity of the algorithm by $\mathcal{O}(\text{OPT}_Q(G) + m + n)$. \square

Lemma E.1. $\text{OPT}_Q(G) + m + n = \mathcal{O}(\text{OPT}_T(G))$.

Proof. Since every comparison query takes unit time, we have $\text{OPT}_Q(G) \leq \text{OPT}_T(G)$. We note that the output of the distance ordering problem has size $\Omega(n)$, thus $\text{OPT}_T(G) = \Omega(n)$.

It remains to argue that $\text{OPT}_T(G) = \Omega(m)$. Consider any fixed weights w such that all pairwise distances by the weights are distinct. Let L be the output ordering of the vertices under w . Assuming that the input graph is given as a list of edges, we prove that any correct algorithm needs to read non-zero amount of data about at least $\Omega(m)$ edges.

Suppose this is not the case. This means that there is a pair of vertices $u \prec_L v$ for which we have seen neither of the two possible directed arcs uv, vu (or, the undirected edge uv in the case when G is undirected), and at the same time they are not adjacent in L . It may be, however, the

case that the edge that we have not read is the edge uv with arbitrarily small weight. This makes the returned ordering L invalid since with a sufficiently small weight, v has to immediately follow u in L , proving that the algorithm is not correct. \square

Theorem 1.2. *Dijkstra’s algorithm implemented with any Fibonacci-like priority queue with the working set property (as defined in Definition 1.4) is a universally optimal algorithm for the distance ordering problem in the comparison-addition model, in terms of running time. This holds both for directed and undirected graphs, and when compared against both deterministic and randomized algorithms.*

Proof of Theorem 1.2. The theorem follows immediately from the combination of Theorem 3.1 and Lemma E.1. \square

Lemma 3.3. *For any directed or undirected graph G , any (even randomized) algorithm for the distance ordering problem needs $\Omega(\log(\text{Linearizations}(G)))$ comparison queries in expectation for some weights, i.e., $\text{OPT}_Q(G) = \Omega(\log(\text{Linearizations}(G)))$.*

Proof of Lemma 3.3. Let A be any algorithm that solves the distance ordering problem. Let us first assume that A is deterministic.

By Claim 2.3, for any linearization L of T , there exists a setting of weights w_L such that L is the unique solution to the distance ordering problem on (G, w_L) . For any given L , we fix one such w_L arbitrarily and denote \mathcal{W} the set of all w_L . Consider a uniform distribution over \mathcal{W} . Its entropy is $H = \log(|\mathcal{W}|) = \log(\text{Linearizations}(G))$. Now sample $w_L \in \mathcal{W}$ from this distribution, run A on (G, w_L) , and denote by $C_L \in \{-1, 0, 1\}^*$ the sequence of answers to the comparison queries made by A .

By the Shannon’s source coding theorem for symbol codes [Sha48], any ternary prefix-free code $C : \mathcal{W} \rightarrow \{-1, 0, 1\}^*$ has to have an expected length of at least $H/\log 3$. But as A is correct and deterministic, the mapping that maps each w_L to C_L is such a code. Namely, if $C_L = C_{L'}$, then due to determinism, the algorithm asked the same queries and got the same answers, and hence it output $L = L'$ and we necessarily have $w_L = w_{L'}$. For similar reasons, C_L cannot be a prefix of $C_{L'}$. Therefore, A has to perform at least $\log_3(\text{Linearizations}(G))$ queries in expectation.

Now let A be a randomized algorithm and let X be the random variable denoting the number of comparisons performed by A . By the previous statement, for any choice of random bits r , we have $\mathbb{E}[X \mid R = r] \geq \log_3(\text{Linearizations}(G))$. By the law of total expectation, we also have $\mathbb{E}[X] \geq \log_3(\text{Linearizations}(G))$. Finally, we can write $\mathbb{E}[X] = \sum_{w_L \in \mathcal{W}} \mathbb{E}[X \mid w = w_L]/|\mathcal{W}|$, and as X is nonnegative, there exists an input w_L such that $\mathbb{E}[X \mid w = w_L] \geq \log_3(\text{Linearizations}(G)) = \Omega(\log(\text{Linearizations}(G)))$. In other words, for some weights w_L , algorithm A needs to make $\Omega(\log(\text{Linearizations}(G)))$ queries in expectation, just as we wanted.

Note that this lower bound holds also in Word-RAM, since every query only supplies $\log_2 3 = \Theta(1)$ bits of information. \square

Proposition 5.16. *Run Algorithm 5 on (G, w) . The total number of comparison queries caused by Algorithm 7 (that is, the total number of lazy queries evaluated later in the algorithm) is at most $\max_w |F_{G,w}| - n + 1$.*

Proof. First, use the same argument as in the proof of Theorem 3.1 to claim that we can in fact transition to some other weights \hat{w} such that $F_{G,w} \subseteq F_{G,\hat{w}}$, Dijkstra’s algorithm on (G, \hat{w}) asks the same queries as Dijkstra’s algorithm run in (G, w) , and furthermore it only asks queries about edges in $F_{G,\hat{w}}$. Although not immediately clear, one can observe that the same holds also for (G', w') and

(G', \hat{w}') , and for $F_{G,w'}$ and $F_{G,\hat{w}'}$. As $|F_{G,w}| \leq |F_{G,\hat{w}}|$, we can carry on with the proof by setting $w = \hat{w}$ (and the conclusion will apply also to the original w as well).

Let $f = |F_{G,w}|$, $n = |V(G)|$ and $f'' = |F_{G'',w''}|$, $n'' = |V(G'')|$. By [Claim 5.7](#) and by induction, $f - n = f'' - n''$. Therefore, it is enough to show that [Algorithm 7](#) causes at most $f'' - n'' + 1$ queries.

[Algorithm 7](#) causes at most one query per each $e \in F_{G'',w''}$. Partition all edges in G'' by their endpoints, namely, for $u, v \in V(G'')$, let $C_{u,v}$ contain all edges from u to v in G'' . Note that for each choice of u and v , either all edges of $C_{u,v}$ are forward, or none is. For each pair $u, v \in V(G'')$, define $D_{u,v}$ as $C_{u,v}$ if the edges in $C_{u,v}$ are forward, and as an empty set otherwise. Note that $|F_{G'',w''}| = \sum_{u,v \in V(G'')} |D_{u,v}|$.

We can see that for each $D_{u,v}$, we actually need to pay only $|D_{u,v}| - 1$ comparisons, not $|D_{u,v}|$, since finding a minimum among k elements can be done using $k - 1$ comparisons. Therefore, the total number of comparisons we need to perform is $\sum_{u,v \in V(G'')} \max(0, |D_{u,v}| - 1) = |F_{G'',w''}| - h$, where h is the number of pairs (u, v) such that $D_{u,v}$ is nonempty. If we recall that G' is defined as the deduplication of G'' , we immediately get that $h = |E(G')|$. Since G' is connected, we have $h \geq |V(G')| - 1 = |V(G'')| - 1 = n'' - 1$.

Thus, the total number of comparisons is $f'' - n'' + 1$, as needed. \square

Claim 5.7 (Structure of contractions). *Let T_D^G be the dominator tree of a multigraph G rooted at s , let w be weights on G , and $e = uv \in E(T_D^G)$ an edge such that u is a node with outdegree one in T_D^G and v its unique child in T_D^G . Assume that every node in G is reachable from s . Also, assume that G may have duplicate edges from one vertex to another, but there are no duplicate edges from a dominator vertex to a dominated vertex. Furthermore, assume that there are no edges $xy \in E(G)$ where y dominates x .*

We claim:

1. *The edge e also exists in G and it is the only incoming edge to v in G .*
2. *There is no edge $uz \in E(G)$ other than e for which $z \in T_D^G(v)$.*

Next, define G', T', w' as follows. We set $G' = G/e$ and $T' = T_D^G/e$. Let $\phi(x) : V(G) \rightarrow V(G')$ be the mapping performed by the contraction, i.e., $\phi(u) = \phi(v) = [uv]$ and $\phi(x) = [x]$ otherwise. We set $w' = w$ with the exception that for every edge $f = vx \in E(G)$ and its counterpart $f' = [uv][x] \in E(G')$, we define $w'(f') = w(u, v) + w(f)$. In particular, for any edge $f = ux$ and its counterpart $f' = [uv][x]$, we set $w'(f') = w(f)$. Then, the following holds:

3. *Every node in G' is reachable from $\phi(s)$. G' has no two edges e_1, e_2 such that they both go from x to y and x dominates y . Also, there are no edges xy where y dominates x .*
4. *$T' = T_D(G')$.*
5. *$T_{SSSP}^{G',w'} = T_{SSSP}^{G,w}/e$.*
6. *$F_{G',w'} = F_{G,w}/e$, $|F_{G',w'}| = |F_{G,w}| - 1$ and $|V(G')| = |V(G)| - 1$.*

Proof. In the following, we abuse multigraph notation, so that, for example, $xy \in E(G)$ means “there exists an edge from x to y in G ”. Also,

1. Let z be any vertex such that $f = zv \in E(G)$. Such a z exists, otherwise v would not be reachable from s . Take an arbitrary path P in G from s to v through z . As u dominates v , P must go through u , and since this holds for all P , necessarily $z \in T_D^G(u)$. But $z \notin T_D^G(v)$, since by our assumption, there cannot exist an edge $zv \in E(G)$ where v dominates z . Then, necessarily, $z = v$, as v is the only child of u in T_D^G . Finally, f has to be a unique multiedge as u dominates v .

2. Take any such $z \neq v$ and $f = uz$. On one hand, since $z \neq u$ and $z \in T_D^G(u)$, necessarily $z \in T_D^G(v)$. On the other hand, take any simple path P from s to z using the edge f . P cannot go through v , as it would have to visit it before u , and that would imply u does not dominate v . But then v does not dominate z , as we have found a path from s to z not going through v . Thus, $z \in T_D^G(v)$ and $z \notin T_D^G(v)$ at the same time, which is a contradiction.
- 3+4. For arbitrary $z \in V(G)$, take any path P from s to z . Necessarily, the path $P' = P/e$ is a path in G' from $\phi(s)$ to $\phi(z)$, and thus, $\phi(z)$ is reachable from $\phi(s)$ in G' . As ϕ is surjective, all vertices in $V(G')$ are reachable from $\phi(s)$.

By a similar argument, whenever a path P contains a vertex $x \in V(G)$, then a path P/e contains $\phi(x)$. Hence, if z dominates x , then $\phi(z)$ dominates $\phi(x)$, and vice versa (with the only exception that v does dominate u), and the dominance relation is thus preserved. By the way the contraction happened, we may conclude that $T' = T_D(G')$.

Suppose we have two edges e_1, e_2 , both going from x to y such that x dominates y . Since G has no such parallel edges, this necessarily means either that $x = [uv]$ and $u\phi^{-1}(y), v\phi^{-1}(y) \in E(G)$, or $y = [uv]$ and $\phi^{-1}(x)u, \phi^{-1}(x)v \in E(G)$. In the first case, since $T' = T_D(G')$, this means that u dominates $\phi^{-1}(y)$. By point 2, $u\phi^{-1}(y)$ cannot exist as $\phi^{-1}(y) \neq v$, which is a contradiction. In the second case, we similarly have $\phi^{-1}(x)v \in E(G)$, but by point 1, such an edge cannot exist for $\phi^{-1}(x) \neq u$.

Finally, assume x dominates y in G' and there is an edge from y to x . Since this does not happen in G , then necessarily either $x = [uv]$ or $y = [uv]$. Consider the first case. Necessarily u and v both dominate $\phi^{-1}(y)$, since $T' = T_D(G')$. Then either $\phi^{-1}(y)u \in E(G)$ or $\phi^{-1}(y)v \in E(G)$, and both contradict the assumptions on G . Similarly, in the second case, $\phi^{-1}(x)$ dominates both u and v and the necessary existence of either $u\phi^{-1}(x)$ or $v\phi^{-1}(x)$ again contradicts the assumptions on G . Therefore, no such x and y exist.

5. By the fact that u dominates v and there is only one edge from u to v , necessarily any path to v , including the shortest one, must use edge e , so $e \in T_{SSSP}^{G,w}$. On the other hand, one can verify that w' is constructed precisely in such a way that, for all P , $w(P) = w'(P/e)$. Therefore, if P is a shortest sz -path in G , then P/e is a shortest $\phi(s)\phi(z)$ -path in G .
6. Since u and v get merged into a single vertex, we have $|V(G')| = |V(G)| - 1$. By checking the definition of w' , we can validate that if an edge is forward, it will be forward also after the contraction, except for e , which disappears. Hence also $|F_{G',w'}| = |F_{G,w}| - 1$, as needed.

□